

Evaluation of the Recognition Distances of Safety Signs in VR Considering Vision Impairments

DIPLOMARBEIT

zur Erlangung des akademischen Grades

Diplom-Ingenieur

im Rahmen des Studiums

Visual Computing

eingereicht von

Dominik Bauer, BSc

Matrikelnummer 1227005

an der Fakultät für Informatik

der Technischen Universität Wien

Betreuung: Assoc. Prof. Dipl.-Ing. Dipl.-Ing. Dr.techn. Michael Wimmer Mitwirkung: Dipl.-Ing. Katharina Krösl, BSc

Wien, 21. Jänner 2018

Dominik Bauer

Michael Wimmer



Evaluation of the Recognition Distances of Safety Signs in VR Considering Vision Impairments

DIPLOMA THESIS

submitted in partial fulfillment of the requirements for the degree of

Diplom-Ingenieur

in

Visual Computing

by

Dominik Bauer, BSc

Registration Number 1227005

to the Faculty of Informatics

at the TU Wien

Advisor: Assoc. Prof. Dipl.-Ing. Dipl.-Ing. Dr.techn. Michael Wimmer Assistance: Dipl.-Ing. Katharina Krösl, BSc

Vienna, 21st January, 2018

Dominik Bauer

Michael Wimmer

Erklärung zur Verfassung der Arbeit

Dominik Bauer, BSc Sonnwendgasse 31/22, 1100 Wien

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst habe, dass ich die verwendeten Quellen und Hilfsmittel vollständig angegeben habe und dass ich die Stellen der Arbeit – einschließlich Tabellen, Karten und Abbildungen –, die anderen Werken oder dem Internet im Wortlaut oder dem Sinn nach entnommen sind, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht habe.

Wien, 21. Jänner 2018

Dominik Bauer

Danksagung

Ich möchte zunächst Katharina Krösl danken, die mir während der Entstehung dieser Arbeit jederzeit für Fragen zur Verfügung stand und mir während des Schreibens dieser Arbeit hilfreiches Feedback gab. Gleichermaßen gilt mein Dank Michael Wimmer für das konstruktive Feedback, das mir half, diese Arbeit stets zu verbessern. Ich möchte außerdem allen am VRVis danken, die mir mit Rat und Tat zu Hilfe standen.

Bei meiner Freundin möchte ich mich für das ständige Korrekturlesen und die moralische Unterstützung danken, vom Proposal bis zur Enstehung dieser Arbeit. Besonderer Dank gilt natürlich meinen Eltern, die mir dieses Studium ermöglicht haben und mir dabei immer unterstützend zur Seite standen.

Kurzfassung

Die sichere Evakuierung eines Gebäudes wird durch Sicherheitszeichen ermöglicht, welche so entlang des Rettungswegs platziert werden müssen, dass sie für die Nutzer des Gebäudes lesbar sind. Während Normen und gesetzliche Vorgaben Empfehlungen zur Auswahl und Platzierung von Sicherheitszeichen enthalten, bieten sie keine ausreichenden Überlegungen in Bezug auf Personen, die an einer Sehbehinderung leiden. Eine der Hauptursachen von Sehbehinderungen sind altersbedingte Augenkrankheiten. Das häufigste Symptom ist dabei der Verlust der Sehschärfe.

Wir untersuchen mittels einer neuen Methodik den Einfluss der Sehschärfe auf die Fähigkeit, Sicherheitszeichen zu erkennen, und evaluieren dadurch bestehende Normen in Hinblick auf Sehbehinderungen: Wir kalibrieren die Sehschärfe aller Testpersonen auf das selbe Niveau unter Verwendung eines standardisierten medizinischen Tests in VR. Dies wird ermöglicht, indem wir die Auswirkungen einer Sehbehinderung für Testpersonen mit normaler oder korrigierter Sehkraft in VR simulieren. Weiters stellen wir ein Tool für Lichtdesigner vor, welches ihnen ermöglicht, ihre Designs in Bezug auf die maximale Erkennungsweite von Sicherheitszeichen zu überprüfen, um problematische Bereiche entlang von Fluchtwegen zu untersuchen.

Unter Verwendung unserer neuen User-Study Methodik bestimmen wir die Erkennungsweiten von Sicherheitszeichen für zwei verschiedene Sehschärfeniveaus und unter variierendem Betrachtungswinkel. Zusätzlich bestimmen wir den Einfluss der HTC Vive auf die in VR erreichbare Sehschärfe. Wir schließen aus unseren Ergebnissen, dass die bestehenden Normen die maximale Erkennungsweiten von Sicherheitszeichen für Betrachter mit reduzierter Sehschärfe nicht korrekt abschätzen können.

Abstract

To facilitate the safe evacuation of buildings, escape-route safety signs need to be placed along the whole escape route such that they are legible for building occupants. While standards and legal requirements provide suggestions on how to select and place safety signs to achieve this, they do not provide sufficient considerations concerning people suffering from vision impairments. A main cause of vision impairment are age-related eye diseases, with the most common symptom being the loss of visual acuity.

We investigate the influence of visual acuity on the ability to recognize safety signs using a novel methodology, evaluating existing standards concerning vision impairments: We calibrate the visual acuity of the test subjects to the same level via a standardized medical test in VR. This is achieved by using test subjects with normal or corrected vision and simulating the impairment in VR. Furthermore, we present a tool for lighting designers which enables them to check their designs considering maximum recognition distances to investigate problematic areas along an escape route.

Using our novel user-study methodology, we determined the recognition distances for safety signs, observed under two different levels of visual acuity and varying observation angles. In addition, we determined the impact of the HTC Vive's HMD on the visual acuity achievable in VR. We conclude that the existing standards fail to correctly estimate the maximum recognition distances of safety signs for observers suffering from reduced visual acuity.

Contents

\mathbf{K}	Kurzfassung i Abstract >						
A							
Contents							
1	Intr	oduction	1				
	1.1	Motivation	1				
	1.2	Problem Statement	2				
	1.3	Contributions	3				
	1.4	Structure of the Work	4				
2	Background						
	2.1	Emergency Lighting	5				
	2.2	Visual Perception	7				
	2.3	Eye Diseases	13				
3	State of the Art						
	3.1	Design and Placement of Safety Signs	17				
	3.2	Visibility Computation	23				
	3.3	Simulation of Vision Impairments	28				
4	Impairment Simulation Tool						
	4.1	Motivation for a Simulation in VR	33				
	4.2	Setup	35				
	4.3	Simulating a Loss of Visual Acuity	35				
	4.4	Simulating Age-related Eye Diseases	36				
	4.5	Plausible Scene Lighting	38				
	4.6	Wheelchair Movement	40				
5	Use	User Study Design					
	5.1	Overview and Study Design	44				
	5.2	Calibration of Visual Acuity	45				
	5.3	Experiment: MRD Determination	49				

	5.4	Experiment: Escape Scenario	51					
6	Design Evaluation Tool							
	6.1	Environment	56					
	6.2	Visibility Computation	56					
	6.3	User Interface	60					
	6.4	Visualizing the User-Study Data	61					
7	Implementation Details							
	7.1^{-1}	Simulation Tool	63					
	7.2	User Study	73					
	7.3	Evaluation Tool	76					
8	Results and Discussion							
	8.1	Simulation Tool	79					
	8.2	User Study	80					
	8.3	Evaluation Tool	87					
	8.4	Limitations	89					
	8.5	Discussion	91					
9	Future Work and Conclusion							
	9.1	Conclusion	95					
	9.2	Future Work	96					
\mathbf{Li}	st of	Figures	99					
\mathbf{Li}	List of Tables							
\mathbf{Li}	List of Algorithms							
Bi	Bibliography							

CHAPTER

Introduction

1.1 Motivation

In emergency situations, the general lighting in a building may fail. But sufficient lighting is key to the building occupants' orientation inside the building, the prevention of the risk of accidents and the occupants' perception of safety. Emergency lighting is designed and placed in buildings to step-in in these situations. Especially along the escape routes of a building, emergency lighting must be adequately placed and should provide sufficient levels of illumination, enabling the safe exit from the building. As an integral part of emergency lighting, escape-route safety signs support the building occupants by providing directional information and illumination. Therefore, their placement is of special concern to facilitate the safe evacuation of a building.

The safety signs need to be visible and legible along the whole escape route. It is preferable to determine a placement that fulfills these demands before the construction of a building, during the design phase, to avoid additional costs due to late changes after construction. There are standards and legal requirements in place that specify how to select the properties and position of safety signs to achieve this. These specifications include viewing angles, important areas that need to be considered and maximum recognition distances. The latter are defined in terms of a safety sign's height and illumination. However, these standards and legal requirements mainly focus on individuals with nonimpaired vision and thus do not provide normative considerations concerning people suffering from vision impairments.

In particular, buildings whose occupants are at high risk of suffering from vision impairments, such as retirement homes, should be equipped with a guidance system that takes these factors into account. Age-related eye diseases are a main cause of vision impairment for these elderly people, with the most common symptom being the loss of visual acuity. The standards and legal requirements currently in place might not be sufficient to design a placement of safety signs that considers the vision abilities of elderly people suffering from a loss of visual acuity. To our knowledge, there are no studies that evaluate and quantify the influence of visual acuity on a person's ability to see and follow safety signs.

1.2 Problem Statement

The standards and legal requirements provide only informal suggestions on how to account for vision impairments: For example, with regards to elderly people, EN 1838 states that they require higher light intensities and more time to adapt to the lighting conditions, but does not consider the age-related eye diseases common for elderly people. Concerning reduced visual acuity, ISO 3864-1 gives an informal recommendation to scale the maximum recognition distances by a person's decimal acuity value. For example, a person with a visual acuity of 0.1 decimal acuity can see the same detail from 1m as an unimpaired person can see from 10m. In this case, the standard recommends to scale the maximum recognition distance (MRD) by 0.1. While visual acuity is the most common vision impairment resulting from age-related eye diseases, no data is available that quantifies the actual influence of a reduced visual acuity on a person's ability to recognize a sign. Therefore, a user study to determine this influence is required.

Such a user study is difficult to carry out in the real world as it is hard to find an appropriate group of participants: This would require a subgroup per visual acuity level to be tested, where each participant suffers from a loss of visual acuity that results in the same level of visual acuity. Furthermore, visual acuity is seldom experienced as an isolated symptom but rather as a part of a more complex set of symptoms that result from an eye disease. Hence, to enable control over the level of visual acuity experienced by the participants, the effect of reduced visual acuity has to be simulated. This simulation has to offer varying degrees of severeness to account for different visual acuity levels.

There exist several approaches to the simulation of the effects of age-related eye diseases that can be broadly categorized into simulator goggles, renderings on 2D screens as well as immersive AR and VR simulations. Simulator goggles use welding or ski goggles with customized lenses to simulate the effects of age-related eye diseases in the real world. Thus, the users of such goggles are restricted to their real-world environments, which limits the possible scenarios that can be evaluated using such a simulation. Renderings on 2D screens either alter static 2D images or use 3D games as basis for the simulation. They allow for arbitrary environments and can incorporate simulations of evacuation scenarios. Yet, only the contents of the 2D screen are altered, not the full vision field of the user. Therefore, the immersiveness of the simulation and thus the realism of the data gathered using such a simulation are limited.

Simulations using AR or VR HMDs allow this realism: The users are fully immersed into an emergency scenario and their whole vision can be effectively altered. However, existing approaches are mainly based on artistic impressions of the effects of age-related eye diseases. Persons actually suffering from age-related eye diseases experience a set of symptoms, each of which impair their vision to a certain degree. These symptoms' strength has to be calibrated to realistic levels to provide a simulation that lets users experience the actual effects experienced by persons suffering from these age-related eye diseases.

1.3 Contributions

The limitations of current approaches in related work prohibit the conduction of a user study on the impact of vision impairments on recognition distances that yields comparable data. Neither the used simulations of vision impairments nor the employed methodologies are suited to this task.

To tackle these limitations, we created a tool to simulate the effects of reduced visual acuity in VR. This thesis provides a more realistic simulation of this symptom than previously achieved by calibrating the severity of the simulated reduction of visual acuity in VR using a standardized medical test. Inside the simulation, the experienced visual acuity can be calibrated to any acuity value and accounts for the user's real-world visual acuity. In contrast to previous work, this allows calibrating all users of this simulation to the same level of perceived visual acuity, enabling a user study on the impact of visual acuity on the MRDs of safety signs. Furthermore, we provide a simulation of two of the most common age-related eye diseases, cataract and dry age-related macular degeneration. These simulations are also based on the calibrated simulation of the loss of visual acuity.

The simulation tool presented in this thesis further provides a novel VR locomotion technique suited to our case study: Elderly occupants of retirement homes. Our wheelchair movement method is based on the motion controllers of the HTC Vive and a common office chair, making it easy to recreate. The feedback from our user study suggests that this locomotion technique is an effective and intuitive alternative to more common VR locomotion techniques. In addition to this, the tool chain we used to create the assets for our user study allows creating detailed building models that are illuminated using the physically plausible light-planning software HILITE. This results in a configurable, yet realistic virtual test environment that, combined with the realistic simulation of the loss of visual acuity, serves as the basis for a quantitative user study on the influence of visual acuity on the MRDs of safety signs.

We employ a new methodology that, for the first time, allows conducting a user study with participants calibrated to the same level of experienced visual acuity. Compared to existing approaches, which are either limited by the availability of an appropriate group of test subjects or the insufficient simulation of symptoms, our new methodology allows us to gather comparable data in VR, validated through a standardized medical test. This is achieved by simulating this vision impairment in VR and using test subjects with normal or corrected-to-normal vision (using glasses or contact lenses). Using this methodology, we determined the impact of the HTC Vive's HMD on the visual acuity achievable in VR as well as the recognition distances for safety signs, observed under three different levels of visual acuity. Furthermore, we determined the recognition distances for safety

1. INTRODUCTION

signs during an evacuation scenario, allowing us to compare the MRDs achieved in these two experiments. Using this evaluation of the influence of visual acuity on the MRDs of safety signs, we evaluate the MRDs suggested by existing standards. We determined that their informal suggestions and formulas for the computation of MRDs are unsuited for the case of people suffering from reduced visual acuity. We provide recommendations for future extensions to these standards as well as suggestions for further experiments to determine a more detailed relationship between visual acuity and the MRDs of safety signs.

Furthermore, we provide lighting designers with an additional evaluation tool to check their designs considering MRDs to investigate problematic areas along an escape route during the design phase of a building. This is achieved by adapting the shadow-mapping algorithm to the computation of the visible area inside a safety sign's MRD. In the evaluation tool, this area is visualized parallel to the building's floor-map and along the escape route. Additionally, this tool allows the visualization of the data that can be gathered using our simulation tool and our new user-study methodology. The evaluation tool thus enables lighting designers to dynamically evaluate their designs considering the influence of reduced visual acuity.

1.4 Structure of the Work

Following this introduction, Chapter 2 provides an overview of the context of this thesis, covering emergency lighting, visual perception and eye diseases. Chapter 3 discusses the previous work in the field, dealing with existing standards and legislation in emergency lighting, existing approaches to visibility computation in this field as well as with previous approaches to the simulation of vision impairments.

The methodology part of this thesis is structured into three chapters: First, the impairment simulation tool is presented in Chapter 4, providing a simulation of age-related eye diseases and a calibrated loss of visual acuity in VR. The user study is presented in Chapter 5 and aims to answer whether the assumptions on the influence of a loss of visual acuity found in the relevant standards can be confirmed or whether these standards need to be extended. Finally, in Chapter 6, a design evaluation tool for lighting designers is presented. This tool allows identifying problematic sections along an escape route through a visualization of the visibility of safety signs.

Details on the implementation of the simulation tool and the evaluation tool are presented in Chapter 7. The chapter also discusses how the user study and its experimental design were embedded into the simulation tool. In addition to this, an overview of the technical environment of this thesis is given. Chapter 8 then presents the results of this thesis, and existing limitations, also comparing this work with previous approaches. The emphasis of this chapter lies on the evaluation of the results of the user study. In addition, the simulation tool and the design evaluation tool presented in this thesis are discussed in terms of their applicability. Finally, Chapter 9 discusses possibilities for future extensions to this work and gives a conclusion based on the presented results.

CHAPTER 2

Background

The following chapter provides definitions for important terms and concepts used in this thesis. The structure of this chapter therefore reflects the general aspects covered by this thesis: First, a definition of the different parts of emergency lighting is provided. This is followed by the discussion of three concepts from visual perception central to this thesis, namely brightness, visual angle and visual acuity. Finally, the problem of vision impairment is discussed and common age-related eye diseases are presented, with a focus on the two eye diseases simulated in Chapter 4.

2.1 Emergency Lighting

In situations where the general lighting in a building fails – such as fire or power shortages–, the so-called emergency lighting must step in. This is essential to risk prevention, as light is fundamental to our perception of safety, and darkness induces fear which results in an increased risk of accidents [WF17]. However, emergency lighting covers more than the replacement for the general lighting. This becomes clear from the classification by the standard EN 1838 [DIN13], shown in Figure 2.1 and described in the following.

Complementary to the replacement of the general lighting through standby lighting, emergency escape lighting should guarantee that work processes can be safely terminated and that the safe exit from a building is possible [WF17]. In accordance with EN 1838, this can be further categorized into:

• Escape route lighting allows the safe exit from a room or building. It guarantees sufficient visual conditions and orientation along the escape routes, especially near firefighting and safety facilities, so that they can be easily found and used. An integral part of escape route lighting are escape-route safety signs (safety signs), which provide the orientation needed to easily find and use escape routes. They may



Figure 2.1: Classification of different types of emergency lighting. Adapted from EN 1838 [DIN13].

also provide additional lighting themselves, helping to create appropriate visual conditions.

- Anti-panic lighting, or open-area lighting, prevents the development of panic situations and allows safely reaching the escape routes by guaranteeing sufficient visual conditions and orientation. This also includes appropriate illumination of obstacles.
- High-risk task area lighting enables the safe termination of work processes by sustaining an appropriate amount of lighting in case of emergency. This ensures the safety of persons in potentially dangerous work processes and endangered persons in the affected area.

It should be noted that escape route lighting is only considered to the extent necessary to guarantee the validity of the depicted brightness of escape-route safety signs used in the user study carried out in this thesis. Other than that, this thesis merely deals with escape-route safety signs and their visibility in the context of vision impairments.

2.2 Visual Perception

Based on measures of a sign's brightness, requirements for the properties of escape-route safety signs are specified in standards and regulative texts. Since these measures are different from the common notion of "brightness", it is important to understand which phenomenon they intend to describe and they are therefore described in the first part of this section.

In this thesis, the influence of the reduction of a person's visual acuity on a person's ability to recognize escape-route safety signs is studied. Therefore, this section further includes definitions of visual acuity and visual angle, illustrated by examples. In addition, a standardized medical test to measure a subject's visual acuity is presented.

2.2.1 Brightness

What is commonly termed "brightness" can be looked at from two different perspectives [Tho11]: The radiometric view defines brightness in terms of the properties of the electromagnetic radiation. The photometric view weights these properties by the spectral sensitivity response function of the human visual system.



Figure 2.2: Different measures of brightness, from the light source to to the eye of the observer. Based on an illustration from [WF17].

Furthermore, different measures of brightness exist that characterize the light emitted from a light source or surface and received by an observer. As such, these measures are based on a photometric view on brightness. For reference, example values for different light sources and lighting settings are given in Table 2.1 and an illustration of the different phonomena described by these measures is given by Figure 2.2. According to Witting [Wit14], the brightness measures are defined as follows: Luminous Flux Φ, measured in Lumens (lm): The luminous flux of a light source is its total emitted radiant power in all directions, weighted by the spectral sensitivity response function of the human visual system. Thus, the more radiant power is emitted in the spectrum between 380-780nm, i.e., as visible light, the higher the luminous flux. Radiant power emitted outside this spectrum, i.e., UV radiation and heat as IR radiation, does not contribute to either vision or lighting and consequently does not contribute to the luminous flux.

Luminous Flux Φ [lm]

• Luminous Intensity I, measured in Candela (cd): The luminous intensity describes the part of the total luminous flux Φ_P that is emitted per solid angle Ω . It is therefore a direction dependent measure. To characterize the luminous intensity of a light source in all directions, a luminous intensity distribution curve can be determined.

Luminous Intensity $I = \Phi_P / \Omega \left[lm \cdot sr^{-1} = cd \right]$

• Illuminance and Luminous Emittance E, measured in Lux (lx): The illuminance is a measure for the luminous flux Φ incident on an area A. Conversely, the luminous emittance measures the luminous flux emitted from an area. The effect of this is not directly visible to the human eye, as the optical impression depends on the reflection factor of the observed surface.

Illuminance/Luminous Emittance $E = \Phi/A [lm/m^2 = lx]$

• Luminance L, measured in Candela per m^2 (cd/m^2) : The luminance of a light source is the part of the total luminous flux Φ_P that is emitted per solid angle Ω and illuminates the area A^* , the projected area as seen from an observer with viewing direction α . Therefore, the luminance describes how a light source or an illuminated surface are perceived by a human observer.

Luminance $L = I/A^* [cd/m^2]$ with $A^* = A \cdot cos\alpha$

2.2.2 Visual Angle

To define the recognizable size of an object independent of the distance between the object and the observer, the angular extent on the retina, i.e., the visual angle the object subtends, and not the actual physical size of an object is considered [Tho11]. This is motivated by the fact that objects subtending the same visual angle will be similarly detectable, independent of their actual size. As reference, when held at arm's length, the thumbnail subtends an angle of 1.5° and the first subtends an angle of $8 - 10^{\circ}$ [Tho11]. Figure 2.3 illustrates the latter example.

Measure	Source/Setting	Value
Luminous Flux Φ [lm]	Candle Light Bulb (75W) Sun	$\begin{array}{c} 12 \\ 935 \\ 3.63 \cdot 10^{28} \end{array}$
Luminous Intensity I [cd]	Candle Light Bulb (60W) Sun	$1 \\ 60 \\ 1.04 \cdot 10^{28}$
Luminous Emittance E [lx]	Corridors, Stairways, Entrance Halls Office Rooms, Meeting Rooms Operating Room in a Hospital	$ \begin{array}{r} 100 \\ 500 \\ 10^4 \end{array} $
Luminance L $[cd/m^2]$	Perception Threshold Average Cloudy Sky Candle Average Clear Sky Light Bulb (60W)	$ \begin{array}{r} 3 \cdot 10^{-6} \\ 2000 \\ 7500 \\ 8000 \\ 5 - 35 \cdot 10^{6} \end{array} $

Table 2.1: Reference values for examples of light sources and lighting settings, given in the presented measures. The values can be found in Witting [Wit14].



Figure 2.3: A visualization of the concept of visual angle: Held at arm's length, the fist subtends a visual angle of 8-10°. Reprinted from [Tho11].

2.2.3 Acuity

Directly related to the visual angle an object subtends is the question of whether the human visual system is able to detect this object. This ability depends on the spatial resolution of the visual system, called the visual acuity. According to Thompson [Tho11], a subject's visual acuity can be tested by using a pattern of high-contrast, parallel light and dark bars. One such pair is called a cycle. When tested in bright light, a normally sighted person can distinguish the bars in patterns of 50-60 cycles/degree. This is equal to the ability of distinguishing two objects separated by a visual angle of 1 arcminute. This ability is dependent on lighting conditions, contrast and "the nature of the pattern under view" [Tho11]. For example, under moonlight the distinguishable pattern is reduced to 20-30 cycles/degree and to only 10 cycles/degree when viewed under starlight.

2. Background

It is also of interest to consider the finest detail a digital display can convey. For example, as the display of the HMD used in this thesis, effectively reduces the visual acuity of the user. As Thompson demonstrates, the finest detail displayable on a 147cm diagonal 1080p HDTV when viewed from 3m distance would be equal to 40 cycles/degree, i.e., below a normally sighted person's visual acuity. The considerations for the display of the HMD used in this thesis are provided in Sections 4.2 and 5.2.

The human visual system itself is limited in its visual acuity by two factors: Optical blur and the effects of the distribution of the photodetector cells in the retina. The optical blur is mainly due to focusing errors and aperture diffraction at the pupil [Tho11]. Focusing errors are for example a consequence of refraction errors of the lens or of a too long or too short eyeball [Pal99].



Figure 2.4: Density of the photoreceptive cells as a function of their position on the retina. Reprinted from [Pal99].

Considering the retina, the most important factor influencing visual acuity is the spatial configuration of the cones which are responsible for vision in most normal lighting conditions. As depicted in Figure 2.4, the cones are densely packed around the fovea, peaking around $1.5 \cdot 10^5$ cones per mm^2 . The densely-packed area spans about 1° to 5.2° of visual angle in the central visual field. This results in a high spatial sampling frequency in this area and a lower sampling frequency in the less densely populated areas elsewhere on the retina, influencing the visual acuity. Since the discrete spatial configuration of photodetector cells allows applying digital sampling theory [Tho11], frequencies for which aliasing effects should be visible can be deduced. However, the visual system is able to avoid or mask these effects under normal conditions.

Further, the signals of the photoreceptive cells in the retina are combined several times [Tho11]. As a result, cells in the optic nerve and in the visual cortex have different receptive fields, varying in location and size. Light that hits the retina outside a cell's receptive field has no influence on its behavior. The signals resulting from light hitting the retina inside a cell's receptive field are effectively filtered. As the photoreceptive cells are more densely packed near the fovea and more thinly distributed farther away from the fovea, the farther away a cell's receptive field is from the fovea, the larger the size of this filtering is. This results in a sharp drop of visual acuity when moving away from the fovea, as the area over which the brightness is averaged through the filtering increases [Tho11].

In addition to these two factors, eye diseases can impair the visual acuity of a person. For example, consider the blood vessels that supply the cells in the retina: The light must pass through these blood vessels to reach the photoreceptive cells of the retina, casting shadows. The brain normally fills out the parts shadowed by the vessels [Pal99]. However, a thickening or leaking of these vessels can impair visual acuity. Causes for this and other impairments are described in Section 2.3

It is important to note that there exist other measures of acuity, for example the so-called Vernier Acuity. Thompson [Tho11] defines this measure as follows: Vernier Acuity measures the ability to detect small differences in position, correlating to the visual angle that the distance between two objects subtends. A normal sighted person can detect such offsets of about 5 arcseconds. The visual angle that corresponds to this value is about 1/10th of the finest detail that is detectable in terms of visual acuity and has a projected size of about 1/5th of the width of a single cone cell. Note that the objects themselves still must be of detectable size, i.e., must subtend a visual angle of at least 1 arcminute. For this "hyperacuity", the visual system "compares the outputs from neighboring photoreceptors" [Tho11].

2.2.4 Measuring Visual Acuity

The standard ISO 8596 [Int09] defines test symbols and procedures for testing a subject's distance visual acuity under daytime conditions. The test symbol suggested is the so-called Landolt ring: a black ring on a white background with a gap on one side, as shown in Figure 2.5. Its variations provide more consistent readability than letters or numbers used in other test procedures. For example, to test for a visual acuity of 1.0 decimal value, the diameter of the ring d should be chosen such that the ring spans 5 arc minutes and its thickness and the gap g span 1 arc minute when viewed from the selected test distance. The angular extent of the gap g is the reciprocal of the decimal acuity value dA. Therefore, to test for other visual acuity levels, the diameter of a ring (in arc minutes) can be computed as:

$$d = 5 \cdot g = \frac{5}{dA}.$$

11

The position of the gap should be horizontally left or right, vertically up or down or diagonally in-between for a total of eight possible positions. ISO 8596 suggests to use at least 5 rings per acuity level to be tested, with the position being as random and different as possible. A test subject has reached the limit of visual acuity when less than 60% of rings in this line can be correctly identified. The actual visual acuity of the test subject is that corresponding to the last correct line.



Figure 2.5: The Landolt ring as defined by ISO 8596: The diameter of the ring is d, the thickness of the ring and the gap are sized d/5. To use this symbol to measure a subject's visual acuity, the size of the gap should be such that it subtends a visual angle whose angular extent is reciprocal to the decimal acuity value to be tested. Reprinted from [Int09].

Additionally, the standard provides definitions and relations between common measures for visual acuity:

- Decimal Acuity Value: As mentioned above, the decimal acuity value is defined in context of the Landolt ring as the reciprocal of the gap size in arc minutes.
- Snellen Fraction: Another common notation is the Snellen fraction, given by the test distance as numerator and the distance at which a normal sighted person could still correctly identify the same symbol as denominator. The decimal value of this fraction is equal to the decimal acuity value.
- LogMAR acuity: The LogMAR acuity is defined by ISO 8596 as the "logarithm (base 10) of the minimum angle of resolution in minutes of arc" [Int09]. Consequently, the decimal acuity value dA can be computed from the LogMAR value lM as

$$dA = 10^{-lM}$$

The diameter or the angular extent of the gap in the Landolt ring used to test for a certain visual acuity level can therefore be directly converted to one of these three common measures for visual acuity and vice versa.

2.3 Eye Diseases

The primary cause of vision impairment and blindness in the United States are age-related eye diseases, as Shoemaker and Friedman [SF02] note in their study on the prevalence of eye diseases in the U.S. population over the age of 40. This problem will become even more pronounced as the number of affected people is "expected to double within the next three decades" [SF02]. This will especially concern people aged 75 and above because the prevalence in this group is particularly high. It is important to note that the number of completely blind in the U.S. is low in comparison with the number of people that suffer from vision impairment. The latter suffer from a permanent partial vision loss with largely varying severity.

In the U.S., vision impairment and legal blindness are defined in terms of visual acuity, according to Shoemaker and Friedman [SF02]:

- Vision Impairment is defined as a visual acuity of 20/40 or worse in the better eye with corrected eyesight. In most U.S. states, this would prevent the affected person from obtaining an (unrestricted) driver's license.
- Legal Blindness is defined as a visual acuity of 20/200 or worse in the better eye with corrected eyesight. Alternatively, a person with a visual field of less than 20° in diameter is also considered "legally blind".



Number of cases of age-related eye diseases in the U.S. population aged 40 years and above [in millions]

Figure 2.6: Most common eye disease in the U.S. population aged over 40. Adapted from the National Eye Institute (NEI) [NEI].

While the most common cause for this is a refractive error, resulting in an unfocused image on the retina, elderly people are especially prone to suffer from a set of diseases classified as "age-related eye diseases". The most frequent age-related eye diseases in the U.S. population aged over 40 are shown in Figure 2.6. The definitions provided by Shoemaker and Friedman [SF02] characterize them as follows:

2. Background

- Cataract: The affected person suffers from a clouding of the lens. The size of this clouding can vary widely, as well as its position and opacity, resulting in several distinct types and classifications. While the disease can be treated, for example by replacing the clouded lens with an artificial one, Cataract remains the major cause of vision impairment in the U.S. It is strongly correlated with age and more than 50% of the U.S. population suffers from cataract by the age of 80. While the exact triggers for Cataract remain unclear, further influences may be the constant exposure to UV radiation in sunlight, cigarette smoking, diet or alcohol consumption. Cataract can also be caused by eye damage resulting from an injury, toxic substances or radiation. This disease can also occur in babies or children, where it is normally caused by genetic defects or developmental disorders.
- Diabetic Retinopathy: The disease affects the blood vessels that supply the cells in the retina. They get blocked, can break and leak and, as a result, impair the affected person's vision. Diabetic Retinopathy is common in people suffering from diabetes and the risk increases with the time they have diabetes. Also, the risk of suffering from further eye diseases is increased for people with diabetes.
- Glaucoma: The cells of the optic nerve of an affected person gradually die, resulting in a slow loss of vision in the damaged area. Usually, peripheral vision is affected first. The disease and its progression are often not noticed until the effects are already severe, resulting in 50% of the affected persons being unaware of their impaired vision. While elevated fluid pressure in the eye is found to be correlated to this disease, the exact cause of most types of Glaucoma remains undetermined.
- Age-related Macular Degeneration (AMD): The affected person suffers from a loss of central vision. AMD is the primary cause of vision impairment and legal blindness in the U.S, as central vision is essential to most everyday vision tasks. Two main forms of AMD can be differentiated in terms of severity and cause:

The more common form is the so-called dry AMD. Its early stage is characterized by the depositing of fat under the photoreceptive cells in the retina. This results in limited vision loss, mainly in the central visual field. In the later stage, the supportive layer under these cells may fade. A more severe loss of central vision is the consequence. The other form of AMD is called wet AMD, as it results in tiny blood vessels breaking open and leaking blood into the retina. Not only does the leaked blood impair vision but also the scar tissue that forms as a result of the breaking of the vessels can distort the affected person's vision. In contrast to the dry AMD, this form of AMD progresses fast. The exact cause for this disease remains undetermined and there is currently no cure for AMD, although laser therapies can help decelerating the progression of wet AMD by destroying the affected blood vessels. Nevertheless, age and smoking cigarettes seem to be correlated with the emergence of the disease.



Figure 2.7: Artistic impressions of (a) Cataract, (b) Diabetic Retinopathy, (c) Glaucoma and (d) AMD. Adapted from NEI [NEI].

It is important to note that the causes as well as the effects of one and the same disease can vary widely between the affected, especially over time. However, to give an impression of how a person suffering from one of these diseases might see the world, artistic impressions provided by NEI [NEI] are shown in Figure 2.7.

The simulation of the effects of eye diseases, as described in Chapter 4, focuses on Cataract and dry AMD. Simulating more severe types of eye diseases prohibits the recognition of safety signs and thus makes the measurement of recognition distances in the user study, presented in Chapter 5, impossible. However, there exist less severe types of these diseases that still allow recognizing signs at a distance. These are investigated in the user study. Furthermore, experiments to simulate all eye diseases discussed above were carried out during the implementation of the simulation that is presented in this work. They suggested that Cataract and dry AMD would result in the most realistic impression without requiring the use of an eye-tracker. By simulating the effects of these diseases on a large area of the visual field, the same vision impairment as produced by the actual eye disease can be achieved. Otherwise, test subjects would be able to look to either side of the affected area, evading the effect of vision impairment and producing measurements unrelated to the simulated eye diseases.

CHAPTER 3

State of the Art

The following chapter discusses related work in the area of lighting design, visibility computation and the simulation of vision impairments. Section 3.1 explains how the placement of safety signs is addressed in the legal regulations and the standards in force. It presents best-practice guidelines by corporations and professional associations as well as methods to evaluate the resulting lighting designs. Two approaches for the visibility computation, needed for the evaluation of a specific placement of safety signs after the initial design phase, are presented in Section 3.2. Finally, in Section 3.3, related work on the simulation of the symptoms of eye diseases is discussed, ranging from modified goggles for educational purposes to sophisticated AR and VR simulations for architectural design evaluation.

3.1 Design and Placement of Safety Signs

A series of international and national laws and regulations are in force that prescribe a general concept for emergency lighting in buildings and the placement of safety signs in particular. Standards further specify how the targets defined in the legislation can be fulfilled. Additionally, corporations and professional organizations provide condensed best-practice approaches to the problem of placing safety signs, even exceeding the targets prescribed by law. Together, these regulations and recommendations point out important spots in a building which must be signposted, the installation height, the maximum recognition distances (MRDs) depending on the size of a sign and its brightness (luminous emittance or luminance, depending on sign type), as well as viewing angle dependent considerations.

Based on this information, a lighting designer must place each safety sign appropriately during the planning phase of a building. Since this is a manual procedure, the placement also depends on the knowledge as well as the experience of the designer and must be evaluated after this initial design phase. An overview of the main sources of information for lighting designers is given in this chapter while a more detailed discussion with a focus on the situation in Europe and Germany can be found in [WF17].

3.1.1 Legal Regulations

There are laws and regulations in place that justify the use of safety signs and give general considerations on the placement of signs while staying relatively abstract. For Austria, relevant legislation is laid down in the ArbeitnehmerInnenschutzgesetz (ASchG), the Kennzeichnungsverordnung (KennV), the Arbeitsstättenverordnung (AStV) as well as the building regulations of the individual federal states. The ASchG for example simply notes that emergency routes must be clearly visible and permanently signposted. The KennV extends on this by requiring safety signs to be easily recognizable and mounted near "relevant" locations.

With regard to employees with disabilities, the AStV states that employers must ensure that the "appropriate" technical or organizational manners are in place that enable these employees to timely recognize a danger, allowing a quick and safe exit from the building. For more concrete suggestions, one has to consult the relevant standards such as EN 1838 [DIN13] and ISO 3864-1 [Int11], the former being concerned with emergency lighting in general and the latter dealing particularly with the design of safety signs.

3.1.2 Standards

EN 1838 [DIN13] deals with all aspects of emergency lighting described in Section 2.1: emergency escape lighting, standby lighting, escape route lighting, anti-panic lighting, high risk task area lighting and escape-route safety signs. The aim of escape-route safety signs (safety signs) is, according to the standard, to make escape routes easy to find and follow. Additionally, safety signs are thought to reduce fear and irritation by making exits and the path that leads to them clearly visible. For this, safety signs should be clearly visible through their internal or external illumination and offer orientation. In terms of vision conditions, EN 1838 notes that since eye sight differs between individuals, the required illumination level and the time needed for adaptation to changes of illumination may vary. In general, elderly people are regarded as requiring a higher level of illumination and a longer time to adapt to the conditions present in emergency routes.



Figure 3.1: Visualization of the MRD proposed by EN 1838. Reprinted from ISO 3864-1 [Int11].

As guidance for lighting designers to make escape routes easy to find and follow, EN 1838 provides suggestions on the placement of individual signs: They should not be mounted higher than 2m above floor level; where possible, also not higher than 20° above the horizontal viewing direction at the MRD of this sign. The MRD – the maximum distance from which a sign should still be recognized by a normal-sighted observer – is defined by the standard as the sign's height times a distance factor, as illustrated in Figure 3.1. The distance factor is assumed to be 100 for externally illuminated signs and 200 for internally illuminated signs. Let z_0 be the distance factor and h be the height of a sign, then the MRD l of a sign is calculated as follows:

$$l = z_0 \cdot h$$

As Xie et al. [XFG⁺07] note, while the distance factors provided by EN 1838 are in general determined through standardized tests, they only assume an observation angle α of 0° (hence z_0). However, people are expected to view the sign from different angles in real-world situations. This of course has an effect on the MRD and thus the distance factors listed in EN 1838. The authors therefore introduce an angle-dependent distance factor that is deduced from considering the visual angle between elements of a sign. If this angle becomes too small, the elements will no longer be separable and the sign will as a consequence no longer be legible. These theoretical considerations can be approximated in practice using a maximum recognition sphere that touches the sign centrally, illustrated in Figure 3.2.

For validation of this concept, Xie et al. $[XFG^+07]$ carried out an experiment with 48 participants: They were asked to walk down a 39m long corridor and walk up to a sign at the end. As soon as they could recognize the information on the sign, the recognition distance was recorded. This procedure was repeated for different safety signs, varying in the symbols used and their height, with subjects approaching from different observation angles. The assumption was that the average of the recorded recognition distances would approximate the MRD for this sign and observation angle. The authors note that the luminance of the signs was not considered, since the illumination along the corridor was assumed to be strong and consistent enough. Also, 55% of the participants had corrected eye sight. The resulting MRDs were consistent with the theoretical model although the data showed slightly longer recognition distances than expected for observation angles between about 30 and 60°.

Similarly, the informal appendix of ISO 3864-1 [Int11] suggests the use of distance factors that are computed based on the angle under which a sign is observed and depending on a sign's brightness. The viewing angle α influences the recognition distance via $\cos\alpha$. This is illustrated in Figure 3.2 and is in accordance with the concept developed by Xie et al. [XFG⁺07]. For the brightness, a table of scaling factors is provided for a few discrete values of luminous emittance and luminance. The illuminance values are for use with externally illuminated signs, the luminance values for internally illuminated signs. Further factors identified by the standard, that influence the ability to recognize a



Figure 3.2: Visualization of the MRD proposed by ISO 3864-1. This also approximates the theoretical model presented by Xie [XFG⁺07]. (1) vertical ring, diameter z_0 , (2) horizontal ring, diameter z_0 , (3) MRD for $\alpha = 0^{\circ}$, (4) MRD when viewed under angle $\alpha \neq 0^{\circ}$, $z_{\alpha} = z_0 \cdot \cos\alpha$ (5) safety sign. Reprinted from ISO 3864-1 [Int11].

safety sign, are the size of the sign's elements, the contrast of the sign's elements with the background, the illumination conditions, the familiarity of the observer with the sign's elements and the observer's visual acuity.

In relation to the latter, ISO 3864-1 suggests to scale the MRD by an observer's visual acuity as follows: If a person has a visual acuity of 20/20, it is scaled by 1.0 – hence, the scaling factor is the person's decimal visual acuity. As is pointed out in ISO 3864-1, with increasing recognition distance, the visual angle spanned by a sign element decreases. Thus, an ever-smaller portion of observers can recognize this sign element. For practical use, the standard suggests to define the MRD such that 85% of observers can still recognize the sign elements from that distance. Additionally, the same should hold true for 50% of observers from 1.5 times that distance and for 15% of observers for 2.5 times the MRD.

3.1.3 Design Guidelines

In addition to the laws, regulations and standards, professional associations publish planning guidelines such as the TRVB E 102 [OB05], issued by the Austrian firefighter's

association. Besides safety signs, this guideline deals with guiding systems mounted on or near the floor and even suggests using electro-acoustic alert systems in regions where persons with impaired visions might be endangered. Considering the case where emergency routes lead through smoke filled corridors, it is advised by the guideline to mount safety signs no higher than 2.6m above floor level. If the safety sign is mounted lower than this, the distance from smoke possibly forming under the ceiling is increased, avoiding the obstruction of the visibility of the sign through the smoke. For the placement of safety signs, TRVB E 102 05 also suggests using MRDs, although using the simpler $l = z_0 \cdot h$ model used in EN 1838.

Finally, corporations publish guidelines that condense all the regulations and suggestions described above to best-practice advices for lighting designers. An example are guidelines given by the Zumtobel Lighting GmbH [ZLG16]. They stress that the use of emergency lighting is essential to fire protection although scenarios including fire are only seen as worst case. As illustrated in the guidelines, any event leading to a loss of orientation, like a power shortage, might induce panic and thus can result in personal damage. To prevent such a loss of orientation, safety signs are demanded to be clearly recognizable in accordance with the environmental conditions and independent of a person's familiarity with a building.

3.1.4 Design Evaluation and Evacuation Scenarios

After the design phase, evacuation models are used to evaluate the evacuation capacities of a building. Such tools allow simulating the behavior of pedestrians, including their interaction with signage. According to Xie [Xie11], the number of publicly available evacuation models is above 60. While these tools are able to simulate an evacuation, the selection of parameters that control the pedestrian's behavior is still up to the user. Furthermore, some tools do not take the visibility of the signs into account.

The evacuation model buildingEXODUS [Xie11] simulates the interactions of occupants of a building between each other, their interactions with the building itself and its guidance systems as well as their interactions with fire. For this simulation, several heuristics that characterize typical behaviors and movements in an evacuation scenario are used. These include for example the probability to detect a sign that is theoretically visible to the observer and the probability that this observer then complies with the information conveyed by the sign. However, the movement of simulated occupants as well as the visibility computation are restricted to a 2D grid of square cells.

In addition to simulations using evacuation models, evaluations of way-finding behavior in emergency situations have been carried out in user studies using serious games and VR simulations. Almeida et al. [AJF⁺14] developed a serious first-person game in Unity [Uni] in which users should find an exit during emergency scenarios. These scenarios vary in terms of safety signage, depend on whether a crowd of people is escaping in the opposite direction or whether smoke and fire are present. A group of fire safety-, emergency- and building evacuation planning experts validated the scenarios in terms of

3. STATE OF THE ART

realism. Test subjects had to start in an office room and decide whether they followed a safety sign by turning either left or right when exiting the room. Almeida et al. evaluated how the users reacted to these different scenarios in terms of their way-finding behavior. While they found that there is a general tendency for subjects to turn right when exiting a room with no safety signage present, all but one followed the escape direction conveyed by the safety sign. Even if a group of people is running away from that direction, 63% of subjects still obeyed the escape direction. This scenario is illustrated in Figure 3.3 on the right. Yet, when there was smoke in that direction, 82% disobeyed the escape direction. In contrast to the simulation tool presented in this thesis, the simulation by Almeida et al. does not consider impaired vision and no information on the recognition distances of safety signs was determined. Thus, an evaluation of an actual placement of safety signs is not possible using this approach.





Figure 3.3: (left) Top view of the building model used by Almeida et al. (right) A crowd running against the direction indicated by the safety sign. Adapted from [AJF⁺14].

Similarly, Cosma et al. [CRN16] use a rail-tunnel evacuation simulation in VR based on Unity to evaluate way-finding performance of users. Cosma et al. compare the use of two different types of bright, green LED stripes as guidance system. In the experiment, test subjects could move freely in a railway tunnel filled with smoke and were asked to escape to an exit. The subjects' position was recorded for discrete time steps. The positions and times were used to compare evacuation times, though the variation between subjects was found to be insignificant. The authors explain this with the fixed walking speed in the simulation. Also, Cosma et al. note that the subjects will always know that they are not actually in a dangerous situation which is expected to influence their behavior. Nevertheless, Cosma et al. conclude that VR simulations are a cost-effective way for the evaluation of guidance systems with the benefit of a high control over the experimental conditions that would not be possible in the real world. As Cosma et al. use a non-standard signage system, the results cannot be transferred to the evaluation of buildings that use safety signs as described in EN 1838.
3.2 Visibility Computation

During the placement of safety signs, it is of interest whether they will be visible from all relevant areas in the building, such as escape routes or ways leading up to them. This can either be the basis for an evacuation simulation using an evacuation model, or a static test, for example in a lighting-planning software. In the following section, visibility computation methods used in evacuation models are presented. In addition, we discuss methods from computer graphics which allow a more precise computation of visibility as compared to those currently used in evacuation models.

3.2.1 Visibility Computation in Evacuation Models

The aim of evacuation models is similar to that of the evaluation tool presented in Chapter 6 of this thesis: They are used to evaluate the placement of safety signs after the initial design phase of a building.



Figure 3.4: The view volume defined in MASSEgress. Point P is considered visible as it is inside the view volume and the ray tracing test's distance is equal to the distance between the observer and P. Reprinted from [CPLL15].

In MASSEgress [Pan06], visibility is determined from the observer's point of view: Objects inside an observer's perception range and viewing angle are considered to be theoretically visible to the observer. This is illustrated in Figure 3.4. The perception range is not dependent on the observed object's size or other object properties, limiting the realism of this approach. For example, differently sized safety signs of different brightness would be assumed to be visible from the same distance. Yet, the benefit of this method is that it would allow for a straight forward extension to account for vision impairments. In addition to the consideration of the observer's view volume, MASSEgress determines whether an object is occluded by another object. This is carried out by a ray tracing test between the observer's eye position and the object's location is larger than the traced distance, it is occluded by another object and thus not visible to the observer.

For example, the occluded agent's distance in Figure 3.4 is larger than the distance of the ray tracing test which terminates at the obstacle.

In SAFEgress [CPLL15], instead of determining the visibility from the observer's pointof-view during simulation, the regions from which a safety sign will be visible are pre-computed. The so-called visibility map used in this method is constructed by first discretizing the obstacle free building space into square cells. For each cell of the resulting 2D grid, the area that would be visible to an observer from this position is computed.

Similarly, in building EXODUS, the computation of this so-called VCA [Xie11] is implemented by first discretizing the building layout into 2D cells. The size of these cells is chosen such that they represent the space occupied by one observer. For each cell, a line-of-sight search is carried out from the height of an average observer to the center of the lower edge of a sign. The lower edge is assumed to be visible if the rest of the sign is also visible. Also, the distance to a safety sign is considered in the visibility computation, based on the $l = z_0 \cdot h$ model provided by standards and guidelines.



(b)

Figure 3.5: (a) The model of a hypothetical supermarket in buildingEXODUS. (b) The VCA resulting from the eight safety signs in the supermarket model. Adapted from [Xie11].

Xie et al. [XFG⁺07] extend the visibility computations originally used in buildingEXO-DUS. The authors include considerations on how the viewing angle influences the MRD of a sign, resulting in circular VCA. In [Xie11], Xie further extends the visibility model by including occlusions caused by smoke. The resulting VCA for a supermarket model is shown in Figure 3.5. In terms of vision impairments, Xie assumes a safety factor of 2 for all individuals with reduced visual acuity, i.e., halving the MRD. Considering the possible range of visual acuity values, this assumption can only be used as an approximation to the actual factors of visual acuity influencing MRDs.

3.2.2 Visibility Computation in Computer Graphics

The visibility computation methods used in the evacuation models of the state-of-the-art methods presented in the previous section are based on 2D grid representations of a building. None of the methods – besides buildingEXODUS – incorporate the safety signs' MRDs into the visibility computation. Furthermore, these methods do not incorporate considerations on the signs' extent into the visibility computation. They only consider visibility for the center point of a safety sign.

In computer graphics [BW03, COCSD03], methods that consider visibility from a specific position are called "from-point visibility". In contrast to this, "from-region visibility" algorithms consider whole regions of space at a time. They allow determining whether the whole region is visible or occluded, i.e., the classification into "visible" or "occluded" is valid for any point inside this region. While some form of preprocessing is commonly required by such methods, this additional computational burden is made up for as the determined region visibility remains valid for static geometry, e.g., the building geometry in architectural walkthrough applications.

A way of exploiting the typical layout of buildings is the usage of "cell-and-portal" methods. They assume that an observer in one room of the building (the cell) can only see an object in another room (another cell) through a door, a window or similar (the portal). Therefore, only objects visible through portals of connected cells have to be considered for the visibility computation.

Further criteria that need to be considered for visibility-computation algorithms, according to Cohen-Or et al. [COCSD03], include the conservativeness of a method, its support for generic occluders and a possible need for precomputation. The conservativeness describes whether the method guarantees to include all visible objects, possibly overestimating the set of visible objects (conservative), or whether the method omits this guarantee to either speed up computation (approximate sampling) or to avoid overestimation (aggressive approximation). Approximate sampling techniques (structured or random sampling) accept to miss some actually visible objects. Aggressive approximative methods weaken the guarantee of including all visible objects to minimize the chance of overestimating the set of visible objects. Bittner et al. [BW03] further identify exact methods, whose result contains exactly the set of visible objects. Concerning the supported types of occluder geometry, some methods are limited to computing visibility for convex occluders, which can require additional precomputation for non-convex geometry. Other methods use volumetric representations or simplified objects that represent the original occluders.



Figure 3.6: (a) The visibility is computed from 5 sampling points with shrunk occluders. (b) The individual umbrae of the actual, unshrunk occluders (dark) as compared to the fused umbra resulting from the conservative visibility computation using sampling and occluder shrinking. Reprinted from [COCSD03].

Wonka et al. [WWS00] obtain a conservative from-region visibility estimation by utilizing a sampling approach in combination with a shrinking of the individual occluders. The amount ϵ by which the occluders have to be shrunk depends on the sampling density. The idea behind this is that an object that is determined to be occluded by the smaller, shrunk occluder will still be occluded by the actual occluder, if the observers viewpoint remains in the ϵ neighborhood of the original viewpoint. One sample approximates an ϵ sized viewing region and the joint umbra of all samples are a conservative approximation for the whole viewing region that is sampled. As Figure 3.6 illustrates, the resulting umbra, i.e., the fully occluded area, is larger than just the joint umbrae of the individual occluders. The authors use 2.5D geometry to enable sensible occluder shrinking. Their presented system is also able to directly extract building facades as occluders. However, in the context of safety-sign visibility, geometry is expected to be created in the course of an architectural process. This may yield arbitrary geometry for which occluder shrinking is non-trivial, e.g., building interiors with complex furnishings. To apply this method to general geometry created in common CAD tools, the development of additional preprocessing procedures would be required.

The extended projections method by Durand et al. [DDTP00] uses projections of the occludees and occluders onto a projection plane to determine a conservative visibility estimation. The projection of an occludee overestimates its size; the projection of the occluder underestimates its size. This guarantees the conservativeness of the process. As Figure 3.7 illustrates, in the most simple case where the projection plane is inbetween the occluder and occludee, an occludee is invisible from a viewing cell if its projection is completely covered by the occluder's projection. Analogous, for more than one occluder,



Figure 3.7: Extended Projections. (a) The occludee (green) is projected to different positions of the projection plane, depending on the viewing position in the viewing cell. The union of all its projections is its extended projection. Similarly, the occluder's projection (gray) onto the projection plane changes with the viewing position. Its extended projection is the intersection of these projections. (b) The projections effectively fuse occluders A and B. Reprinted from [DDTP00].

the occludee's projection has to be completely covered by the joint projection of the occluders, effectively fusing occluders. However, the projection plane might be hard to position in more complex cases. The authors suggest a reprojection operator to handle these situations. The positioning of the projection plane itself is carried out in a brute-force manner, placing a projection plane behind each occluder and evaluating the projection size in a preprocessing step. Furthermore, only convex occluders are handled out of the box. To use concave occluders, they need to be sliced, requiring additional preprocessing. In the context of safety-sign visibility, a scene can be complex, for instance, due to the amount of furnishing or structural building elements contained in the scene. This inflates the cost of positioning the projection planes in the preprocessing procedure. As the concave geometry of many furnishing objects needs to be sliced in order to be used with this method, the preprocessing burden is further increased. For complex building models, this might become a limiting factor as the safety signs' positions and sizes will change throughout the planning process with each change requiring the recomputation of the preprocessing step.

In this thesis, we decided to use an approximate sampling technique based on omnidirectional shadow mapping to compute the visibility of safety signs. A safety sign is sampled at each corner and in the center. For each sample, we determine the umbra using the shadow-mapping algorithm. We consider the union of these umbrae as the complete umbra of the safety sign. This approach is motivated by its ease of implementation and integration with the recognizability computation via MRDs as well as the lack of any required preprocessing. As a result, we are able to directly use general mesh data as input for the visibility computation, without any additional geometry processing. A detailed motivation and discussion of the method used in this thesis is provided in Section 6.2.

3.3 Simulation of Vision Impairments

There are different approaches to simulate age-related eye diseases: The simplest method are modified goggles with painted lenses. Such goggles try to emulate the symptoms of an age-related eye disease and can be used to simulate how a person suffering from a particular eye diseases might see. They are commonly used in medical education and medical studies.

3.3.1 Simulator Goggles

Zagar et al. [ZB10] developed a set of simulator goggles for educational purposes that reproduce the typical advanced-stage symptoms of diseases such as glaucoma or cataract. The goggles are reported to simulate moderate to severe levels of impairment, which are defined by the authors as a visual field of less than 40° and visual acuity of 0.29 decimal acuity or less. To produce this effect, the lens of a pair of goggles was painted with clear fingernail polish until a visual acuity test for a level of 0.29 decimal acuity could no longer be successfully mastered.

Using similar goggles, Wood et al. [WCCC10] study the influence of cataract and refractive blur on night-time driving performance and pedestrian recognition. The visual acuity of the test subjects was reduced using goggles until it reached a level of 0.6 decimal acuity. For this, the so-called "Bailey Lovie Chart" was used. Subjects had to read the lines of this chart from 3m distance and were instructed to guess when unsure. In another test phase, subjects had to recognize road signs and avoid low-contrast hazards on the road. The results show that visual acuity and contrast sensitivity are strong predictors on driving performance and the number of recognized signs.

Simulator googles are limited in the range of symptoms and severetiy levels thereof they can capture. The goggles themselves decrease the users' vision field and the handcrafting of the lenses limits the reproducibility of the experiments.

3.3.2 2D Renderings and 3D Games

The application of computer graphics for the task of simulating eye diseases opens up the opportunity to simulate a more diverse set of symptoms. The resulting methods either manipulate 2D images or apply post-processing filters on 3D first-person games.

Lewis et al. [LBCM11] use a first-person game to simulate common eye diseases using overlays and post processing effects in Epic's Unreal Engine 3 (UE3). Opticians and a consultant who runs impairment awareness workshops evaluated the simulation. They concluded that, while the simulation seems useful and accurate, the severity of symptoms should be controllable and that characters representing other people should be included, as these are problematic when navigating an environment with impaired vision. The authors find UE3 suitable for this task but the need to install Epic's Unreal Development Kit was found unsuitable for use in medical education. To overcome this problem, Lewis et al. [LSB12] developed another simulator in a XNA-based framework. The effects used are based on the descriptions of symptoms and were evaluated by an optician. The simulation was found to not accurately represent the impairments but to give a good-enough impression to raise awareness of the implications on mobility of people with impaired vision.



Figure 3.8: The top row shows the results of the UE3 based simulation and the bottom row shows the XNA based simulation, presented by Lewis et al. in [LSB12] and [LBCM11]. (a) Glaucoma. (b) Cataract. Adapted from [LBCM11] and [LSB12].

There are various methods ([HVD06, HvD08, JAR05]) based on measuring the loss of acuity and contrast sensitivity using medical tests. These methods alter images such that they convey the reduced vision resulting from these impairments. Hogervorst et al. [HVD06] simulate "fixation independent", "fixation dependent" and glare effects of eye diseases. Fixation independent effects are related to the task of object recognition, while fixation dependent effects are related to visual search tasks. The authors found a linear correlation between visual acuity as measured using the Landolt C Chart and a just-recognizable threshold for blurring an image, depending on the local contrast in an image region. The just-recognizable blur threshold also relates to the eccentricity of a pixel in an image. The experiments were carried out with subjects with unimpaired as well as impaired vision. Based on these findings, the authors provide a tool to alter 2D images such that they convey the respective impairing effects. This is intended to, on the one hand, educate people with unimpaired vision on the constraints faced by people suffering from vision impairment. On the other hand, the authors find the tool to be suitable for the evaluation of designs with vision impairment in mind, e.g., in architecture. As Hogervorst et al. [HVD06] note, it may not be possible to simulate how a person suffering from vision impairment sees the world. However, the visual limitations faced by these persons can be visualized by removing information, that would not be visible to them, from an image.

While 2D renderings and 3D games are able to produce plausible impressions of the symptoms of age-related eye diseases, these impressions are limited to 2D screens. Carrying out real-world tasks is not possible with such simulations.

3.3.3 Simulation in AR/VR

To allow for an immersive experience of the effects of age-related eye diseases, AR and VR simulators effectively alter the users' vision. They are used to simulate the symptoms of eye diseases in medical education, accessibility inspection and to raise awareness for eye diseases. No HMD available at the time of writing provides a resolution on par with the human visual system, neither in terms of resolution nor field-of-view. Thus, a certain impairment due to the HMD cannot be prevented using currently available hardware for simulations in AR or VR.

Velázquez et al. [VVR16] propose an AR setup using a HMD and a monocular camera to simulate symptoms of common eye diseases by applying filters or masks to the rendered images. Using a perimetry exam, the effective visual field of the setup is determined and is found to span 56°. The simulation of eye diseases is based on a description of common symptoms and artistic impressions provided by the National Eye Institute [NEI], shown in Figure 2.7, and was evaluated in a user study. The subjects of the study had to carry out every-day reading and mobility tasks, e.g., navigating in a wheelchair without colliding with obstacles. The time to solve a task was used as a measure to compare different impairments. However, Velázquez et al. found that task completion time varied heavily across subjects and less so between the different impairments that were simulated. Figure 3.9 (a) presents the resulting simulations of age-related eye diseases.

A similar AR setup is used by Ates et al. [AFF15] to increase awareness and allow accessibility inspection for interface designers. The Oculus Rift DK1 used as HMD offers a diagonal field-of-view of 110° and the wide-angle camera used has a field-of-view of 85°. Thus, the setup can only simulate a narrow vision field compared to the full human vision field of 210° horizontally. The simulation is also based on photos provided by the National Eye Institute [NEI] and implemented through a VR media player. The output of this media player can be altered using HLSL shaders to apply filters to the video stream of the camera. Ates et al. [AFF15] evaluated this setup through a user study: Subjects had to control an email program while using the AR setup. A problem encountered was the low resolution of the Oculus Rift DK1 of only 640x800 pixels per eye. A selection of simulations presented in the work by Ates et al. is shown in Figure 3.9 (b).



Figure 3.9: Simulation of different age-related eye diseases as presented in (a) Velázquez et al. [VVR16], (b) Ates et al. [AFF15], (c) Väyrynen et al. [VCH16] and (d) Jin et al. [JAR05]. Reprinted from the respective papers.

The camera used for the AR setups limits the resolution and field-of-view in which the surroundings can be observed by users. Additionally, a monocular camera setup prevents binocular 3D vision, increasing the risk of users hurting themselves while using the simulator. The use of VR simulations reduces this risk as compared to the use of simulator goggles or AR setups in a real-world environment.

Väyrynen et al. [VCH16] base their VR simulation of eye diseases on online simulators that produce static images and hardware-based simulations. The authors use Unity to allow users to navigate through a city model while carrying out way-finding tasks. The aim is to allow designers to understand problems faced by impaired persons early in architectural design. The evaluation of the presented method was carried out once without impairment as baseline and with simulated eye diseases as comparison. The way-finding and target-location tasks were evaluated based on a time basis. The times did not differ across diseases, only between subjects. The authors thus conclude that a time-based evaluation is not suitable for these tasks. Still, the subjects found the simulation to be a useful tool to consider vision impairment as early as possible during the design phase of buildings. Additionally, the VR simulation was found to allow to better "empathize with vision impaired" [VCH16] as compared to simulations using static images. Figure 3.9 (c) shows the resulting simulations of the work by Väyrynen et al.

In contrast to the methods based on artistic impressions, Jin et al. [JAR05] base their VR simulation on data from real patients: The perimetry exam data from patients suffering from AMD, Glaucoma and Diabetic Retinopathy serves as example to create a scotoma texture. In the simulation, this texture is used to select regions where vision is deteriorated. By letting users navigate through a home, they are educated what difficulties people suffering from vision impairment must cope with on a daily basis. Jin et al. [JAR05] thus find such a simulation to be suitable for training physicians and students. However, the scotoma textures are the same for all users and the simulated symptom strength does not accommodate for the users' vision capabilities. The resulting simulations are shown in Figure 3.9 (d).

CHAPTER 4

Impairment Simulation Tool

One objective of this thesis is to provide a VR simulation of the effects of vision impairment caused by age-related eye diseases. Such a simulation can be used to dynamically evaluate a concrete emergency lighting design during the design phase of a building. This includes the evaluation of the placement of safety signs by allowing interactive walkthroughs with impaired vision. Furthermore, we use the plausible simulation of an individual symptom to empirically determine its influence on participants' performance in recognizing safety signs. This allows us to extend the computation of MRDs, suggested in existing standards, by factors that depend on an observer's vision capabilities. The experiments to determine these factors are part of the user study carried out during this work and are based on the simulation presented in this section.

4.1 Motivation for a Simulation in VR

The experiments carried out in the user study presented in Chapter 5 could be carried out in the real world with people affected by eye diseases. Yet, the symptoms of a certain eye disease vary strongly between affected people in terms of which symptoms are experienced, as well as the strength of their effects on vision.

While it is tedious to find a large enough number of people affected by the same set of symptoms of the same magnitude, it is possible to simulate the effects caused by eye diseases using a VR HMD and post-processing effects. In VR, symptoms can be simulated in isolation and with a controlled strength, which heavily reduces the variability between individual participants. Additionally, this allows the group tested with vision impairment and the control group to be the same set of people: A participant can carry out the evaluation tasks once without the simulated impairments and again with simulated impairments presented at different levels of severity. We present previous work in simulating the effects of age-related eye diseases in Section 3.3. The results of the discussed work suggest that VR-based simulations, based on artistic impressions of the diseases' symptoms provided by the NEI [NEI], are a useful tool for raising awareness and for educational purposes. However, since these simulations are neither calibrated to the users' visual capability nor validated, for example through comparison with existing medical tests, they are unsuited for carrying out quantitative experiments. Furthermore, as illustrated by Figure 3.9, the existing simulations provide no consistent simulation of the strongly varying symptoms of eye diseases, also often not allowing for specific control of individual parts of the simulation. Two major limitations can be identified for most approaches:

- 1. Dependence on Viewing Direction: Many symptoms depend on the viewing direction, for example the loss of vision in the central field of vision. In this example, when a user moves the eyes, the simulated loss of vision should follow the gaze. The effect should remain in the central field of vision. Therefore, an accurate simulation of such symptoms requires an eye-tracker. None of the discussed methods ([VVR16, AFF15, VCH16, JAR05]) incorporates this in their simulation, but some avoid this issue by increasing the area a simulated symptom affects. While this prevents users from looking past the affected area by turning their heads and looking at the unaffected border regions, the simulation becomes inaccurate as it also deteriorates vision in areas that would be unaffected in reality.
- 2. Selection and Strength of Symptoms: Since the symptoms experienced by an individual who is affected by a certain eye disease can vary strongly, it is difficult to select a representative subset of symptoms for a specific eye disease. This is further complicated by the varying severity of the symptoms. Two of the discussed methods ([VVR16, AFF15]) simplify the problem by selecting a combination of symptoms of a certain strength that best resembled the artistic impressions of the effects of eye diseases provided by the NEI [NEI]. The authors of a third method ([VCH16]) fit the symptom strength of their simulation to the static images produced by online eye disease simulations and hardware-based simulations. Besides ignoring the variability of the simulated eye diseases, these approaches have the downside of being independent of the user's visual capabilities. For example, slightly blurring the displayed image may result in a slight reduction of visual acuity for a normal-sighted person. However, this does not affect the vision of a person that already suffers from a loss of visual acuity to the same extent as that being simulated. As a result, two persons exposed to the same simulated symptoms would experience different levels of impairment depending on preexisting conditions. The discrepancy between simulated impairment and perceived impairment would make the comparison of data gathered from a study using such a simulation difficult, as it would only be representative for the respective perceived impairment. Furthermore, the implementations of most methods discussed ([VVR16, AFF15, VCH16]) lack direct control over the strength of the simulated symptoms.

The first limitation could not be overcome in this thesis due to the lack of an appropriate, affordable eye-tracker compatible with the VR HMD used. Therefore, we selected those age-related eye diseases that could be most realistically simulated without the use of an eye-tracker, namely cataract and dry AMD. The symptoms of these two diseases are simulated such that they affect a large area, preventing users from looking at unaffected border areas.

The second limitation is tackled by focusing on a specific symptom, namely the loss of visual acuity, and basing the simulation on a calibrated representation of its effect. The calibration is achieved through the use of a standardized medical exam, as outlined in Section 5.2. The loss of visual acuity is one of the most common symptoms of age-related eye diseases. Thus, we base simulations of more complex age-related eye diseases on the calibrated simulation of the loss of visual acuity, increasing the degree of realism of the resulting simulations. In addition to this, we assume a loss of visual acuity to strongly influence the achievable recognition distances for safety signs, which is evaluated as part of the user study described in Chapter 5.

4.2 Setup

The simulation is implemented in a state-of-the-art game engine, Epic's Unreal Engine 4 (UE4) [Epi], and offers an easy-to-use control over the simulated symptoms. Additionally, the game engine contains a set of standard effects and tools, as well as content-editing tools, making the simulation easily extendable for future use. Furthermore, UE4 provides rendering modes particularly tailored to VR. A detailed description of the implementation is provided in Section 7.1.

The VR system used in this thesis is the HTC Vive. Its display has a resolution of 2160x1200 pixels and a refresh rate of 90Hz. We assume that these properties enable a reasonably realistic simulation of the symptoms of eye diseases in VR. In addition, the tracking system of the HTC Vive supports room-scale tracking. The tracking information provided by the system can be directly accessed in UE4 through an interface for the SteamVR API, utilized by the HTC Vive. Furthermore, two motion controllers are part of this VR system, which allow various inputs and are also tracked by the system. These controllers serve as the basis of the wheelchair locomotion, presented in Section 4.6, as well as for the interactions in the user study.

4.3 Simulating a Loss of Visual Acuity

In this thesis, the primary symptom simulated is the loss of visual acuity. A loss of visual acuity is one of the main symptoms experienced by people suffering from the most common age-related eye diseases outlined in Section 2.3. Especially diseases affecting the retina often result in a loss of visual acuity in the impaired area, which can eventually lead to a complete loss of vision in these areas. It is of interest to simulate this symptom as plausibly as possible because it can serve as a basis for the simulation of more complex

conditions. To achieve this plausibility, the parameters controlling this effect must be chosen in a medically sensible way, as outlined in Section 5.2.

To simulate a loss of visual acuity, post-processing blurs are an effective choice. This has been shown in prior work, for example by Hogervorst et al. [HVD06] who determine a relation between the σ parameter of a just recognizable Gaussian blur and the visual acuity of a person. Since the simulation we present in this thesis should be rendered at 90 frames-per-second on the VR HMD, an efficient implementation of the Gaussian blur as a post-processing effect is needed. Using a naive approach to implement a Gaussian blur, h^2 texture lookups and multiplications are required per pixel in the filtered image, with h being the size of the Gaussian filter kernel. A more efficient implementation considers the separability of the Gaussian filter kernel. This only takes $2 \cdot h$ texture lookups and multiplications, one order of magnitude less than the naive implementation.

A further speed-up can be achieved by utilizing filtered texture lookups, sampling the texture at intermediary points using linear interpolation. The filtering effectively covers a larger area in the image using the same filter size and thus requires fewer texture lookups. A detailed explanation of this method is provided in Section 7.1. This approach takes $2 \cdot \lceil h/2 \rceil$ texture lookups and multiplications, thus half the effort required for the separable Gaussian blur kernel. Using this Gaussian blur implementation, we can efficiently use large filter sizes, making the simulation of large amounts of blur possible. The exact amount of blur, i.e., the σ of the Gaussian distribution to simulate a certain visual acuity, is determined in a calibration step per participant. This process is described in detail in Chapter 5. The resulting simulation of a strong loss of visual acuity is demonstrated in Figure 4.1.

4.4 Simulating Age-related Eye Diseases

Based on the plausible simulation of the loss of visual acuity, we simulate two of the most common age-related eye diseases outlined in Section 2.3 by additionally utilizing texture overlays and contrast effects provided by the game engine: cataract and dry age-related macular degeneration (AMD). While the symptoms of these diseases affect the central field of vision, simulating the effects on a larger area of the visual field with low strength allows us to represent the effect experienced in an early phase of the diseases' progression.

To recognize a sign, the user must look at it directly, and the peripheral vision plays a minor role. By simulating the effects of a disease on a larger area, the user is prevented from looking past the affected area while still experiencing a plausible vision impairment. Other age-related eye diseases affecting the central field of vision, such as Diabetic Retinopathy, have a more spotted, localized effect that cannot be reasonably simulated without using an eye-tracker.

The effect of an impaired peripheral vision might be interesting in the context of evacuation performance, in terms of the impact on speed and number of errors. An impaired peripheral vision might increase the time needed to localize the safety signs or may even



Figure 4.1: The effect of the simulated loss of visual acuity.

make their detection impossible. This would also introduce errors in wayfinding. Yet, this thesis is only concerned with the ability to recognize a sign. The models for the determination of the MRDs of a sign assume that an observer directly looks at a sign, i.e., keeps the sign in the central field of vision. Thus, we do not consider eye diseases affecting the peripheral vision.

The overall goal when simulating cataract is to achieve the effects of a fogged vision that reduces visual acuity, reduces perceivable contrast, and introduces a yellowish tint. In this thesis, the reduction of visual acuity is achieved by using the calibrated Gaussian blur presented in Section 4.3. We simulate a reduced contrast by manipulating the contrast controls provided by the engine. Finally, the tint is introduced by lerping between a yellowish tint color and the image. The impact of these three effects is bounded by an alpha-map that falls off to the border of the image. The resulting cataract simulation can be seen in Figure 4.2a. It is important to note that the yellowish tinting also affects the perceived color contrast of the safety signs.

To achieve the effects of dry AMD, the central field of vision is impaired in terms of perceivable intensity and visual acuity. We achieve this by darkening the image using an overlay texture and blurring it using the calibrated Gaussian blur. Again, an alpha-map is utilized to control the affected area. The resulting dry AMD simulation is illustrated in Figure 4.2b.



Figure 4.2: Simulation of age-related eye diseases. (a) The effect of the simulated cataract: The vision is fogged by a yellowish clouding of the lens, reducing visual acuity, perceivable contrast and tinting the perceived colors. (b) The effect of the simulated dry AMD: The central field of vision is darkened and reduced in terms of visual acuity.

4.5 Plausible Scene Lighting

The realism of the simulation of how a person suffering from vision impairment might experience an emergency scenario also depends on the geometry of the building and on the lighting present in that building.

The geometry should be as complex as necessary for the evaluation tasks and as simple as possible to allow fast rendering. Thus, we constructed simple corridors and a grid-like evaluation building in the 3D interior design software pCon.planner [Eas]. The software provides tools for the quick creation of architectural models and provides a variety of interior models, spanning from furniture and luminaries to plants and other decoration. The use of such interior models allows the creation of a reasonably complex building geometry.

The resulting building model can then be imported into the light-planning software HILITE [VRVb]. In HILITE, luminaries that feature the luminous intensity distribution curve (see Section 2.1) of their real-world counterparts are inserted into the building model. Using these luminous intensity distribution curves and a realistic material model, HILITE enables the rendering of physically plausible lighting. This also allowed us to control the illuminance of the safety signs present in the building model, which is important to guarantee the validity of the lighting in terms of the target values prescribed by emergency lighting standards and regulations. Additionally, HILITE offers the possibility to export the building model including the luminaries, the rendered lighting baked into a lightmap and the lightmap UVs as FBX file.

This FBX file is directly imported into UE4 with only minor adjustments to the model needed inside the game engine. For a detailed description of how to import a model from HILITE into UE4, see Section 7.1. Inside UE4, we used the model as environment during the simulation runs, lit using the lighting from HILITE as premultiplied lightmap. Figure 4.3 shows how the final model rendered in UE4 looks like.



Figure 4.3: The building model rendered in UE4 as used in the user study. It was designed in pCon.planner and features physically plausible lightmaps rendered in HILITE.

4.6 Wheelchair Movement

The VR system used in this thesis – the HTC Vive – provides room-scale tracking in a space of up to 4.5x4.5m. However, the experiments of the user study presented in Chapter 5 require the participants to navigate a corridor of 40m length and a building sized 154x29m. Therefore, to navigate through the scenes during the user study, an appropriate locomotion technique is needed. Since the participants of the user study should be able to give feedback via (at least one of) the motion controllers of the HTC Vive, the interaction method for the movement must be simple.

We tested different common VR locomotion techniques in an experiment, resulting in the following conclusions:

- **Teleportation** is unsuitable for the task of determining recognition distances as the jumping to user-selected points in the scene used by this method does not allow to determine the exact distance inside which a sign can be recognized. Users are expected to jump beyond this threshold point, resulting in an underestimation of the actual MRD.
- There are several approaches to "**natural movement**", for example tracking head and hand movements and defining thresholds for when a movement is counted as a "step". Per step, the user then moves into the direction indicated by either one of the hands or the viewing direction. Tracking hand movements was discarded as it was expected to interfere with giving feedback via the controller. Although tracking the head movement seemed like an intuitive mean of control, it was expected to be too exhausting for the participants during the course of a 30min simulation run.
- Strafing by clicking the trackpad of the Vive controller or swiping over the trackpad in a certain direction provides a more familiar locomotion technique for people with 3D game experience. The user moves in the direction of the click or swipe, either by being teleported a "step" in that direction or by strafing the same way as in a classical first-person game. The problem with being teleported a step is the same as with the common VR teleportation technique. Classical strafing is also no viable option as it is known to cause motion sickness in VR, which is especially a concern when considering an extended VR session during the user study and even more so for people without VR experience.

None of the above techniques provided satisfying results, therefore we developed an alternative locomotion technique.

The aim of this thesis is to evaluate the impact of age-related eye diseases on the recognizability of safety signs. Thus, the simulation must convey the impressions of a person suffering from such a disease. Elderly persons that suffer from age-related eye diseases are especially endangered by misplaced safety signs due to their reduced mobility. Therefore, in the simulation, we consider people who have to use a wheelchair. This also

provides another case that is often not considered during the design phase of emergency lighting – building occupants with an eye-height lower than the average population.

Because the feedback for the user study can be carried out using only one of the two motion controllers tracked by the VR system, the other can be used to track the position and orientation of an office chair that serves as the real-world part of a wheelchair simulator. For this purpose, the motion controller is mounted to the back of the chair. We used a simple loop made from tape for the wheelchair simulator as shown in Figure 4.4 on the left. When the chair is rotated or moved, the motion controller is tracked accordingly. Therefore, by attaching the transformation of the virtual wheelchair to the transformation of the motion controller, both the virtual and the physical wheelchair move together. As the virtual wheelchair also moves the tracking space of the HMD and thus the user's view in the simulation, the user can move through the scene using this wheelchair simulator.

The user can see the virtual representation of the wheelchair in the simulation and rotate it in VR by rotating the office chair in the real world as needed. The virtual model is shown in Figure 4.4 on the right. To move in the forward direction of the wheelchair, the user now simply has to press the trigger button of the free motion controller used for inputting feedback. Depending on how far the trigger button is pressed, the speed of the virtual wheelchair varies from 0, equivalent to not pressing the trigger button at all, to a maximum speed when fully pressing the trigger button.



Figure 4.4: The wheelchair simulator: (left) The office chair with the mounted motion controller used for the wheelchair simulator viewed from behind and (right) the virtual model of the wheelchair in the simulation viewed from above.

CHAPTER 5

User Study Design

The subsequent section is structured as follows: First, Section 5.1 presents the general concept of the user study carried out in the course of this thesis and gives an overview of the experiments. Next, Section 5.2 provides details on the calibration process for the simulation of loss of visual acuity. Section 5.3 then discusses the experiments to determine MRDs for different levels of visual acuity. Finally, Section 5.4 presents the escape scenario and discusses the building used as test environment.

The main goal of the user study presented in this section is the determination of the influence of varying visual acuity on a participant's ability to recognize safety signs. We calibrate the simulation of a loss of visual acuity used in the user study via a standardized medical test per participant. The data gathered is expected to be a plausible estimation for the real-world effects of this symptom of vision impairment. This allows carrying out an experiment in VR to determine MRDs for safety signs. Such a VR experiment enables extensive control over the experimental conditions. Not only can the lighting be tuned to accurately fulfill the prescriptions of the relevant standards, but the vision capabilities of the test subjects can also be altered and adjusted to a comparable level. The latter further allows using the data gathered during this experiment to determine whether the considerations of existing standards are sufficient for people suffering from vision impairments. In addition, participants have to escape from a specially designed test building as part of a second experiment. It is used to evaluate whether the reduction of the recognition distances observed in the first experiment is consistent with the recognition distances observed when the participants additionally have to navigate through a complex environment. Compared to the static assumptions of the concept of MRDs, the visibility, the observation angles and the distances to the signs change dynamically in such a situation.

5.1 Overview and Study Design

Our user study consists of three phases: The calibration procedure and two rounds of experiments. First, each participant has to carry out the calibration procedure outlined in Section 5.2. After determining the bias on the visual acuity caused by the HTC Vive, participants are calibrated to two predefined visual acuity levels, 0.2 and 0.125 decimal acuity.

The experiments following the calibration are performed in two rounds in order to test for learning effects. The experiments are presented in Section 5.3 and Section 5.4. The purpose of the first experiment is to determine recognition distances for varying safety signs, observation angles and visual acuity levels. The second experiment allows studying the participants' behavior during a simulated evacuation scenario and it also makes the study less tiring for them. The participants' task in this experiment is to escape from a building using the information conveyed by the safety signs. While the first experiment is restricted to the analysis of the influence of visual acuity on MRDs, the second experiment also includes the simulation of age-related eye diseases.

In the first round of experiments, the first two runs are carried out without any simulated impairments to acquaint the participants with the experimental environment. The experiment protocol for the user study is as follows:

- 1. Calibration of Visual Acuity
 - Determine HMD's influence on visual acuity; visual acuity level "no blur"
 - Determine blur factor for 0.2 decimal acuity; "weak blur"
 - Determine blur factor for 0.125 decimal acuity; "strong blur"
- 2. First round of experiments
 - MRD determination, no blur
 - Escape scenario, no blur
 - MRD determination, weak blur
 - MRD determination, strong blur
 - Escape scenario, weak blur
- 3. Second round of experiments
 - MRD determination, no blur
 - Escape scenario, cataract with weak blur
 - MRD determination, weak blur
 - MRD determination, strong blur
 - Escape scenario, macular degeneration with weak blur

The illuminance of the safety signs in both, the MRD determination experiments and escape scenario experiments, is examined by measurements of a physically plausible lighting simulation done in HILITE. The scene lights' brightness is then adjusted to guarantee that the safety signs presented to the participants are illuminated such that the distance factor z_0 of 100 for externally illuminated safety signs suggested by EN 1838 [DIN13] can be assumed. The standard ISO 3864-1 notes that if a sign features a vertical illuminance of 51x and above, a distance factor of 95 can be assumed. The next larger value specified is 1001x and above, resulting in a distance factor of 170. Therefore, by adjusting the lights' brightness such that the safety signs' average illuminance is 301x, it can be safely assumed that the corresponding distance factor is at least 100 and thus that the illumination fulfills EN 1838. Figure 5.1 shows the measurements in HILITE.



Figure 5.1: Measuring the illuminance of a safety sign in HILITE: A measurement surface is placed on the safety sign, as can be seen in the false color visualization (bottom left). The measured values (top right) show an average illuminance of 30lx on this measurement surface.

5.2 Calibration of Visual Acuity

Though the test subjects should have normal vision or corrected-to-normal vision, their visual acuity is expected to vary to some degree. This becomes especially problematic when their visual acuity is worse than the loss of visual acuity that is being simulated. In this case, the data reflects the impact of their worse real-world visual acuity and not that of the simulated loss of visual acuity. Such data distorts the relationship between visual acuity and MRDs. Therefore, for the evaluation of the data gathered from the user study, we detect these outliers and exclude them from the evaluation data. Also, the HTC Vive's display already introduces a loss of visual acuity caused by its resolution. The properties of the Fresnel lenses used in the optical system of the HTC Vive furthermore result in a blurred image along the edges of the field-of-view of the HMD, reducing the

5. User Study Design

visual acuity in these peripheral regions. It is of interest to determine the influence of the HMD on the level of visual acuity reachable in VR, as it allows interpreting the data gathered from the experiments with non-impaired vision and can provide useful information for future work in this area.



Figure 5.2: The calibration process: (top left) The black color of the Landolt rings implies that no input was recognized yet. The position of the participant's thumb, and thus input to be made, is indicated by a white sphere on the trackpad of the motion controller model. (top right) The participant has determined the orientation of three rings. Their green color does not indicate whether these inputs were correct or not. The next ring that requires an input is the forth from the left, as indicated by its black color. (bottom left) When determining the σ parameter of the Gaussian blur, the participant's vision becomes increasingly blurred for consecutive lines. Note the slightly blurred vision and the smaller size of the rings, compared with the previous examples. (bottom right) A detail view of a line of Landolt rings.

Taking all this into account, we require some sort of calibration process before carrying out the actual experiments. As the visual acuity experienced in VR has to be comparable across participants, this calibration process must be performed in VR too. Further, since the simulations used in the experiments depict the symptoms of eye diseases, a calibration based on medically sensible measures brings certain benefits: Valid test methods to determine a person's visual acuity exist that can be reimplemented for use in VR. The result of such a test method is a visual acuity value that relates to a certain degree of vision impairment. Also, by calibrating participants to a certain visual acuity value through such a standardized procedure, the data gathered in the user study becomes comparable to real-world experiments. Hence, in this thesis, a standardized medical test to determine a participant's visual acuity – namely the test outlined in the norm ISO 8596 [Int09] – is adapted for usage in VR. See Section 2.2.4 for more detailed information on the medical test outlined in ISO 8596.

In the standard procedure outlined by ISO 8596, a test chart consisting of several lines of Landolt rings is used. Each line corresponds to a certain level of visual acuity and consists of (at least) 5 Landolt rings of appropriate size. However, in VR, there is no need to create the whole chart with fixed visual acuity steps beforehand. New lines can be generated as needed, with varying steps and randomized Landolt ring orientations. The latter is especially beneficial when determining the blur required to simulate different levels of visual acuity, as is the case with our calibration procedure.



Figure 5.3: The trackpad of the motion controller as seen by the user in VR. The green area in the middle accepts no inputs as mistakes would be too likely to happen with such small distances between different directions. For the experiments, the number of possible directions is reduced to four.

The calibration procedure presents the participant only one line of 5 Landolt rings at a time. The participant has to determine the orientation of the gaps in the Landolt rings and reports this orientation by clicking the according orientation on the trackpad of the motion controller, shown in Figure 5.3. To support the test subjects in this task and to prevent unwanted inputs, the participants can see the motion controller and an indicator of their input in VR. Also, once an input for a Landolt ring is registered, the corresponding ring turns green as illustrated in Figure 5.2. At any point during the calibration procedure, the participant knows for which Landolt ring the orientation is asked and what input will be registered by the system when clicking the trackpad. In accordance with ISO 8596, a line is rated "correct" when at least 3 out of the 5 Landolt rings' orientations were determined correctly. As long as a participant answers the lines correctly, the test conditions change depending on the measure to be calibrated: Either the size of consecutive lines is decreased, increasing the visual acuity required to detect the gaps in the Landolt rings, or the amount of blur is increased, effectively reducing the experienced visual acuity and making the correct answering of a line harder. Once a line is answered incorrectly, i.e., with less than 3 rings' orientations correctly determined, the limit of the participant's visual acuity for the current vision conditions is reached: The participant could correctly answer the previous line and therefore has a visual acuity at least corresponding to that line's conditions. However, the participant could not correctly answer the current line and thus has a visual acuity less than that corresponding to the current line's conditions. Consequently, the participant's visual acuity is assumed to be the level corresponding to the last correct line's conditions. This procedure is carried out three times under different vision conditions:

- 1. To determine the HTC Vive's influence on the reachable level of visual acuity, the calibration procedure is carried out without any further impairment besides the one introduced by the HMD. After each correctly answered line, the size of the Landolt rings decreases. Thereby, the reachable level of visual acuity is determined analogous to the original test in ISO 8596.
- 2. When determining the amount of blur needed to result in a loss of visual acuity equivalent to 0.2 decimal acuity (0.7 logMAR or 6/30 Snellen), the size of the Landolt rings of the test lines are kept constant, i.e., sized corresponding to the visual acuity level of 0.2 decimal acuity. Instead, starting from no blur, the σ parameter of the Gaussian blur which is used to simulate the loss of visual acuity is increased after each correctly answered line. This increases the resulting blur and decreases the simulated visual acuity. The result is the inverted procedure compared to the previous setting: The visual acuity experienced by the participant is decreased until the line corresponding to the desired visual acuity level can no longer be correctly answered. This blur level is also referred to as "weak blur" in the remainder of this thesis.
- 3. The same procedure is carried out to determine the σ parameter for a visual acuity of 0.125 decimal acuity (0.9 logMAR or 6/54 Snellen), albeit using a

line of larger Landolt rings corresponding to the desired visual acuity level. This blur level is also referred to as "strong blur" in this thesis.

It is important to note that the participants' visual acuity might already be worse in real life than it would be due to the HMD. However, it is assumed that people with normal or corrected-to-normal vision have a high level of visual acuity that does not influence the outcome of the calibration. In part, this limitation applies to the determination of the blur levels too: The amount of blur needed to simulate a visual acuity level better than that of the participant's real-world visual acuity can of course not be determined. Nevertheless, the visual acuity levels corresponding to the weak and strong blur levels are so severe that such a test subject would experience difficulties in everyday tasks.

Thus, by selecting participants with normal or corrected-to-normal vision, this limitation is not expected to apply to the participants of this user study: Participants start from a (close to) unimpaired visual acuity level of about 1.0 decimal acuity that is expected for normal or corrected-to-normal vision. Their effective visual acuity is reduced by the HMD's optics and display itself. As described in Chapter 8, this reduced the visual acuity of participants to about 0.4 decimal acuity. The calibration to the "weak blur" and "strong blur" levels is equal to a calibration to 0.2 and 0.125 decimal acuity, respectively. Therefore, if we assume that participants have a visual acuity higher than 0.4 decimal acuity, neither their original real-world visual acuity nor the HMD's deterioration of visual acuity impact the calibration process – their visual acuity is effectively capped at about 0.4 decimal acuity through the use of a HMD anyhow.

During the calibration, the participants' visual acuity is then further reduced by introducing blur until they can no longer recognize the line of Landolt rings that relates to the respective blur level. If we consider a real-world visual acuity exam, people that fail at the same line of Landolt rings are assumed to have the same level of visual acuity. Thus, the experienced visual acuity for the two blur levels is expected to be the same for all participants, since they also fail at the same line of Landolt rings.

5.3 Experiment: MRD Determination

The experiment which is designed to determine the influence of a participant's visual acuity on the MRD of a safety sign is inspired by the one carried out by Xie et al. in $[XFG^+07]$. The general procedure carried out in this experiment is as follows: A safety sign, featuring the design prescribed by the standard ISO 7010 [Eur12], is placed at the end of a corridor. The arrow on the sign randomly points to the left, to the right, up or down. The participants have to approach the sign and report the direction the arrow points to as soon as they think they can recognize it correctly. If this input is wrong, the participants get feedback by the motion controller which then vibrates twice. This means that the participants should move closer and try again as soon as they are certain. If the input is correct, the motion controller vibrates once, the safety sign at the end of the corridor.

The distance for correct inputs is recorded and serves as basis for the determination of the MRD. This procedure is repeated with different sign heights, different observation angles and – most importantly – different levels of visual acuity. Figure 5.4 illustrates the different stages of the experiment.



Figure 5.4: The determination of MRDs: (left) The participant starts at the end of the corridor and cannot yet recognize the information on the safety sign. (middle) After approaching further, the participant can recognize the direction the arrow on the sign is pointing to and reports this direction using the trackpad. If this input is correct, the next sample is presented. (right) The same situation as in the middle image but for the visual acuity level "strong blur", equivalent to a decimal acuity value of 0.125. It is not possible to recognize the information on the safety sign, hence the participant has to move closer. Overall, this results in a shorter MRD.

The sign height h is either 15 or 30cm and thus, the resulting MRDs should theoretically differ by a factor of 2. The selected observation angles α are 0°, 30° and 60°. Using the factor $\cos\alpha$ suggested by ISO 3864-1 [Int11], this should influence the MRDs by factors of 1.0, 0.87 and 0.5 respectively. The levels of visual acuity simulated are either "no impairment", equal to the loss of visual acuity caused by the HMD and presumably different for each participant, "weak blur" (0.2 decimal acuity) or "strong blur" (0.125 decimal acuity). The goal of this experiment is to determine how these visual acuity levels influence the MRDs. The assumption made by ISO 3864-1 is that the MRDs should scale with the decimal acuity value, thus by a factor of 0.2 for the "weak blur" level and a factor of 0.125 for the "strong blur" level, respectively. In total, 18 measurements are obtained per participant and test run: One measurement for 2 sign heights at 3 angles with 3 acuity levels each.

It should be noted that while the participants are instructed not to simply guess but to establish a certain confidence before reporting the direction they determined, some might still do so and therefore increase their chances to report correctly – after a wrong input, it becomes obvious that only one of the other options will be correct. This could distort the recorded recognition distances and might result in an overestimation for the respective MRDs, if these measurements are not detected as outliers.

5.4 Experiment: Escape Scenario

In the second experiment, participants freely navigate a building and have to try to escape through an emergency exit. To find the emergency exit, they have to follow the directions which are conveyed by the safety signs placed along the escape route of the building. Each time they recognize the direction displayed on such a safety sign, they are asked to report this direction just as the interaction in the MRD experiment before. The motivation for this is to evaluate whether the MRDs represent a meaningful measure for placement of safety signs, not only during the design phase of a building, but also during a simulated evacuation. Figure 5.5 shows the main steps of this experiment. In addition, by carrying out this experiment with a simulated loss of visual acuity, the factors determined in the course of the MRD experiment can be evaluated and compared to the suggestions found in the standards and guidelines. Furthermore, the two age-related eye diseases implemented in the simulation are evaluated for their influence on the participants' performance.



Figure 5.5: The escape scenario. (left) The participants start in a room with several obstacles and have to locate the escape route.(middle) The participants have to navigate through the grid-like part of the building by following the directions provided by the safety signs. The distance to the next sign varies, and the escape route itself changes between individual runs of the escape scenario. (right) After following the escape route to the end, the participants have to navigate through a large hall featuring three emergency exits.

The motivation for the building layout, shown in Figure 5.6, is to cover a range of possible scenarios: In the first room, several obstacles complicate navigation while the decorative elements distract the participants from the safety signs. The task in this room is to find the start of the escape route – a door with a safety sign on top. This door can be seen from most of the first room since this room is separated from the start of the escape route by glass walls. Once the participants have located the start of the escape route, safety signs guide their way through the main section of the test building, constructed as a grid of corridors. The idea behind the use of a grid is that this allows defining different escape routes of same total length as well as same segment lengths.

5. User Study Design

These properties enable the comparison between runs using different escape routes, as the participants have to recognize the same number of signs from the same distances, albeit in a varying order. Furthermore, the grid-like structure is assumed to be confusing enough for participants such that they can neither learn an escape route nor locate themselves, even after consecutive runs through the building. After the grid of corridors, the participants enter a room featuring a large obstacle in the middle. This room is intended to be representative for problematic areas that might be present along an escape route in the real world: Here, the next safety sign is not visible at first. Also, the path along the direction suggested by the last safety sign before this room is blocked by the obstacle. This forces participants to explore the room and look for the continuation of the escape route. Finally, the escape route leads through a spacious hall, reminiscent of an entrance hall found in large buildings. There are three possible emergency exits at the end of this hall. However, only one of them is the actual exit. This means that participants cannot simply rely on finding a door but have to find the one featuring a safety sign on top. Since the distance from one end of the hall to the emergency exits exceeds the MRD of the safety signs, participants are forced to move through most of the hall before being able to detect the information on the sign. In addition, two of the possible emergency exits are occluded by obstacles when viewed from large parts of the hall.



Figure 5.6: The building used for the evacuation experiment as rendered in UE4, featuring the lightmaps imported from HILITE, viewed from above. Areas not accessible to the participants are crossed out. (A) points to the participants' starting point, (B) points to the three possible emergency exits.

The safety signs are placed along the escape route as suggested by the standards and regulations: There is a sign before each direction change, at each crossing, before doors and above (one of) the emergency doors. However, they are only visible when following the escape route along the correct direction, and only one exact escape route is signposted. This means there is only one route leading to the exit. This simplification is introduced to prevent distractions caused by safety signs of other escape routes leading from or to other points inside the building. Two of the escape routes used in the user study are shown in Figure 5.7.



Figure 5.7: Two of the escape routes used in the user study are highlighted in green. The two other paths are mirrored versions of these, albeit with the final exit at the left and right respectively. The safety signs are depicted as green rectangles. As can be seen, the length of the escape route segments and the number of segments of same length are equal for both (therefore all) paths, making the results from the different routes comparable.

CHAPTER 6

Design Evaluation Tool

Especially when dealing with buildings for special use such as retirement homes, a lighting designer has to estimate the influence of eye diseases and other vision impairments on the recognizability of the safety signs and adapt the design accordingly. Existing tools such as building EXODUS [Xie11] enable a lighting designer to import a building design, including safety signs, and to determine the regions from which the signs will be recognizable according to MRDs. However, this information is only available for a coarse grid representing the floor space of the building, evaluated at a fixed height.

Therefore, we developed an extended design evaluation tool for lighting designers, which allows them to determine problematic areas along the escape route of a building during the design phase. In contrast to previous work, our design evaluation tool enables lighting designers to interactively change the vertical position of the horizontal slice through the building for which the recognizability is evaluated. For example, when considering wheelchair users, it is not only important to take into account the average eye height of building occupants, but also that of specific groups of people, particularly if they are especially endangered in case of emergency, e.g., the lowered eye height of wheelchair users. In this case, a lighting designer has to be able to account for average standing eye heights as well as average seated eye heights. It is important to note that similar considerations have to be made for buildings that are mainly frequented by children or adolescents, which also have to be able to recognize safety signs from a lowered eye level compared to adults. Therefore, we enable the lighting designer to easily and interactively check the building design at various eye levels.

Additionally, to enable a convenient way to examine the coverage of the escape route in terms of recognizability, our evaluation tool visualizes vertical slices along the escape route. This make it easier to detect problematic areas along the escape route by covering the escape route at all possible eye levels at once. In combination, these two views cover the areas for which a good recognizability of the safety signs is most important. Finally, the evaluation tool presented in this thesis offers the possibility to visualize the data gathered

from experiments that are carried out using the simulation tool discussed in Chapter 4. On the basis of these visualizations, an interactive comparison of the theoretical MRDs and the actual recognition distances of the participants of the user study can be performed, offering insight into problematic areas considering vision impairment.

6.1 Environment

Our design evaluation tool is implemented using the open-source rendering platform Aardvark [VRVa]. This is motivated by the idea to directly access scene information from the light planning software HILITE – which is also based on Aardvark – in a future version of the tool. In addition to the rendering capabilities of Aardvark, the evaluation tool employs the Aardvark.Media framework for the creation of the GUI. The central part of the design evaluation tool is the visibility computation considering MRDs. We were able to achieve real-time framerates without any effort required for additional optimizations, since the Aardvark platform provided efficient functionality out of the box.

6.2 Visibility Computation

The selection of the appropriate approach for the visibility computation depends on the visualizations that it should enable: In this thesis, a horizontal and a vertical slice through the building are visualized, showing the areas covered by the MRDs of the safety signs and the areas occluded by other geometry. The MRDs are determined by the signs' height, their brightness and their illumination type.

6.2.1 Theoretical Considerations

The accuracy of the visibility computation has to be at a level which allows the resulting visualizations to be used by a lighting designer to decide on the adequacy of the coverage of the escape route. Using a coarse scale, the regions for which the visibility is computed are larger than a human observer. In this case, a region might be classified as visible considering representative points inside this region, but for an actual observer the respective safety signs might still be occluded. The resulting visualizations would indicate sufficient coverage of the escape route while actual observers would experience incomplete visibility when moving along the escape route. The regions for which visibility is determined have to be at least so small that they represent a human observer. The smaller we choose the scale of the regions for which we determine visibility, the more precisely we can visualize the actual visibility. However, the resolution of the rendered visualization remains as limiting factor. Reducing the scale of the regions below the size of a pixel in the visualization therefore does not benefit the informational content of the visualization.

Simple line-of-sight tests on a coarse grid as utilized in building EXODUS are rigid in the regions that can be examined and lack in detail due to the underlying grid structure.

SAFEgress [CPLL15] suffers from the same problem due to the underlying coarse 2D grid structure, although the employed determination of visibility from the observer's point-of-view could be extended to consider the effects of vision impairment. Similarly, the approach of MASSEgress [Pan06], which considers an observer's view volume, could be extended to account for vision impairments. However, this would still require to compute and save such view volumes for all possible positions and viewing directions that are covered by the escape route, both horizontally and vertically.

By determining the visibility as seen from the safety signs, we can reduce the viewing directions to consider to the safety sign's normal. Still, we would need to consider all possible positions inside the building per sign. By further limiting a safety sign's visibility to the regions inside its MRD, i.e., the region where a sign is also recognizable, we can heavily reduce the number of required computations: Positions outside the respective MRD are never considered and only positions for which the MRDs overlap are considered more than once. The computation models for MRDs result in a hemisphere (EN 1838) or a sphere (ISO 3864-1), both positioned in front of the safety sign. Therefore, we can further consider all points that lie behind a safety sign to be occluded, as they would not be recognizable.



Figure 6.1: The region from which a safety sign is recognizable is delimited by its MRD (green). We only consider this region for visibility computation, approximated by a hemicube (solid black). The dashed black lines indicate the view frusta used for rendering.

To carry out the visibility computation as seen from the safety sign, we propose the use of well-known methods from real-time rendering: omni-directional shadow mapping (ODSM) and percentage-closer filtering (PCF). Each safety sign is represented by a point light that casts an omni-directional shadow, delimited by the sign's MRD. By computing which parts of the scene are lit or in shadow, the visibility of a sign is effectively computed as well. Since we consider points that lie behind a safety sign to be occluded, we do neither require the backwards facing side nor the parts of sides of the cubemap used in ODSM that lie behind the safety sign. Instead, we can use a hemicube of a full cube map for the shadow mapping, as indicated in Figure 6.1.

The proposed ODSM approach, however, would only allow us to compute from-point visibility for a selected sample on the safety sign's area, e.g., its center. This would result in an overestimation of the area from which the safety sign is fully visible, i.e., its information can be seen and understood: While the center point might be visible from a certain position, half of the sign might still be occluded, as Figure 6.2 illustrates. In the worst case, only the center is visible and the rest of the sign is occluded. This problem is



Figure 6.2: Sampling visibility at the center of a safety sign using ODSM to compute the umbra: The visible region and the umbra (gray) are separated by the solid black line (left). The safety sign should be fully visible from all points that lie in the visible region (middle). However, using only a single sample at the sign's center, only this point is guaranteed to be visible. Parts of the safety sign might be occluded anyways (right).

reduced by also considering the corners of a safety sign, as shown in Figure 6.3. While the use of 5 samples is sufficient for larger occluders that make up a building's geometry (Figure 6.3a), it may fail for smaller occluders (Figure 6.3b). Further increasing the number of samples would reduce the falsely classified areas but would also increase the computational cost and is therefore not practical.

To correctly solve this problem, we would require a conservative, from-area visibility algorithm, as discussed in Section 3.2. Yet, the discussed methods ([WWS00, DDTP00]) require preprocessing of the geometry for which the visibility should be computed. General geometry, including non-convex objects, cannot be used directly – and providing geometry preprocessing for such general geometry would go beyond the scope of our evaluation tool. Since the sampling-based method shown in Figure 6.3 is assumed to be sufficient to determine visibility for the building geometry itself, the visibility computation in the evaluation tool is carried out accordingly: We extend the proposed ODSM-based approach to an approximative, structured-sampling based method by sampling visibility at each of the corners of a safety sign and its center.

6.2.2 Proposed Approach

The first step in our visibility computation is the computation of ODSM for each sampling point of each safety sign: For standard ODSM, the basic shadow-mapping algorithm has to be carried out for the 6 sides of a cube, each representing one direction in which light is cast by a point light, hence "omni-directional" shadow-mapping. Each side of the cube captures the depth information for 90° horizontally and vertically. In our implementation, we can restrict this to a hemicube, as explained above. The rendering to the hemicube is implemented through the use of a geometry shader. For each input vertex, the geometry shader emits transformed vertices for each of the 5 directions.

The second step is the evaluation of the shadow maps: For each fragment that is rasterized for the visualization slices, a depth value is determined from the hemicube maps. These maps contain the shadow maps for each of the 5 samples per safety sign. To avoid an


(a) The umbra (gray) is determined as the union of the umbrae of each sample (left). For reasonably large occluders, e.g., walls and other building geometry, the safety sign is fully visbile from all points inside the visible area (right).



(b) For smaller occluders, e.g., building interior such as lamps or plants, the individual umbrae may be disjoint (left). While the safety sign is correctly classified as (partly) occluded from points in the umbra, the safety sign may still be partly occluded from points in the visible area (right).

Figure 6.3: Sampling visibility at the safety sign's center and its corners.

aliased shadow border, the determination of the depth values is carried out in a filtered manner. The resulting shadows, or visibility regions, are illustrated in Figure 6.4.

By using ODSM, we do not need to recompute the shadow maps as the slices are moved around. Using Aardvark, the rendering framework we use to implement the evaluation tool, the shadow map is only redrawn if the occluding geometry or its position changes. Since we assume static building geometry, this is only the case once at the start of the tool. Also, as ODSM allows considering viewpoints above safety signs, we can handle cases where the escape-route level is above the safety sign's height, e.g., at staircases.

Since the standards suggest that the MRD of a safety sign is the limit for its recognizability, it is assumed in our evaluation tool that they are also not visible beyond the distance from the sign given by its MRD. The visibility computation considers this by evaluating, for each fragment on a visualization slice, whether it is inside a safety sign's MRD or not. The visibility is evaluated for all safety signs at once. If the fragment is not inside any sign's MRD, the fragment is regarded as shadowed. Otherwise, the shadow map is evaluated for this fragment as described above. The MRD computation itself allows the use of two different models that were discussed in Section 3.1:



Figure 6.4: Visualization of the visibility regions inside the test building as rendered by the evaluation tool, without considering MRDs. Areas rendered blue are either occluded by geometry or outside the MRDs of the safety signs. At least one safety sign can be recognized from areas rendered green. (top) The horizontal slice and (bottom) the vertical slice.

- 1. The standard model suggested by EN 1838: $l = z_0 \cdot h$
- 2. The model presented by Xie et al. [XFG⁺07] and suggested in the appendix of ISO 3864-1. This also considers the observation angle: $l = z_{\alpha} \cdot h$, with $z_{\alpha} = \cos \alpha$ and α being the observation angle.

Computing and evaluating the shadow maps only for two visualization slices allows for a higher resolution as well as for a faster and simpler computation compared to previous work in evacuation models ([XFG⁺07, CPLL15, Pan06]). The two slices can be represented by simple polygons that are "shadowed" by the safety signs. Thereby, the slices' position can be changed continuously and the slices themselves can be directly constructed from building information model (BIM) data specifying the path of the escape route (if available for this building).

6.3 User Interface

Users of the evaluation tool are able to load building geometry, including safety signs, load log files from simulation runs and specify parameters for the visualization. The latter include the eye level for the slices, the height and width of the slices, the model to be used for the computation of the MRDs as well as the height and the distance factor of the safety signs. No information on the illumination of the safety signs is directly available in the evaluation tool, as neither the BIM data for internally illuminated signs nor measurements for externally illuminated signs can be gathered from HILITE at the moment. Therefore, the distance factor cannot be computed and therefore has to be manually selected by the user. The GUI illustrated in Figure 6.5 enables these interactions.



Figure 6.5: The GUI of the evaluation tool: The user can load building models and experiment data. The visibility model to be used can be selected as well as the parameters controlling slices. The visualization itself is darkened while the sidebar is opened.

The appropriate view on the visualization itself depends on the building geometry, the size and height of the visibility slices and the user's intentions. To reproduce the participants' path through a building, for example, it might be of interest to directly follow their movements at the respective eye level. To enable such interactions, the user can freely navigate the visualized scene using a fly-through camera. Furthermore, we visualize glyphs that represent the data from the user-study logs, corresponding to the participants' movements and inputs.

6.4 Visualizing the User-Study Data

The aim of visualizing static MRDs, overlaid with information from the simulation runs, is to allow a comparison of the MRDs suggested by one of the two models available in the evaluation tool to the recognition distances of the participants in the user study. Furthermore, this visualization allows lighting designers to evaluate their designs in VR with different safety-sign configurations as well as with participants with controllable vision capabilities.

The different types of logs recorded during an experiment are visualized using different glyphs. These glyphs are placed at the positions of the HMD, i.e., the participant's perspective, at the time the log was recorded. To also visualize the movement of the participant, a line is drawn between the glyphs, connecting them in chronological order.

CHAPTER

Implementation Details

While Chapters 4, 5 and 6 deal with the theoretical considerations and the methodological approach to the solutions presented in this work, this chapter details how the resulting tools were implemented. In addition, the technical environments in which this implementation is carried out are presented. Section 7.1 discusses the simulation tool by describing the implementation of the interaction and movement methods, the simulation of vision impairments, as well as the import of models from HILITE. Closely related to this, Section 7.2 discusses how the user-study design and the experimental procedure are embedded in the simulation tool as game logic and how the data from the user study is recorded. Finally, Section 7.3 details the implementation of the visibility computation in the rendering framework Aardvark [VRVa]. Furthermore, the visualization of the data gathered from the user study as glyphs and the implementation of the evaluation tool's GUI in the Aardvark.Media framework are discussed.

7.1 Simulation Tool

In this thesis, we want to achieve a realistic simulation of the effects of vision impairment in VR, usable for a quantitative user study. This requires a responsive interaction and rendering. Therefore, the technical environment in which the simulation tool is realized must support the use of a VR HMD and has to allow for an efficient implementation of the simulated symptoms of eye diseases. To enable the calibration procedure and the experiments of the user study, user interactions and the implementation of our wheelchair movement method are required, which are also based on the simulation tool's technical environment. To fulfill these requirements, the simulation tool is implemented in Epic's Unreal Engine 4 (UE4) [Epi] and uses the HTC Vive as VR HMD and tracking system. Figure 7.1 gives an overview of this environment and the simulation tool's structure.



Figure 7.1: Overview of the simulation tool. The input handling and rendering are covered in this section. The game logic is discussed in Section 7.2.

7.1.1 Environment

UE4 is selected as basis for our simulation tool due to its integration of VR APIs and a rendering pipeline adjusted to the needs of VR. Optimizations such as Instanced Stereo Rendering work out of the box, and the deferred rendering pipeline normally used by the engine can be swapped for a forward rendering pipeline which allows for comparably higher frame rates. UE4 offers a direct integration of SteamVR, the API used by the HTC Vive. Thus, the camera settings for VR rendering are automatically adapted to the HTC Vive, and tracking data as well as inputs from the motion controllers can be easily accessed.

UE4 provides ready-to-use implementations of widely used post-processing effects, for example bloom, ambient occlusion or motion blur. The precomputation of lightmaps for the scene geometry, including indirect lighting, as well as particle systems are supported too. These features, combined with a WYSIWYG editor and content-editing tools, enable a fast development of visually appealing 3D games. Additionally, the engine allows defining materials that not only control the rendering of scene objects, but can also be used to control post-processing effects. The latter is exploited for our implementation of the simulation of symptoms of eye diseases. Materials can be altered at runtime through material-parameter collections: The material is defined in dependence of such parameters and their values are altered through code. Further, custom nodes enable the use of HLSL shaders inside a material, although the pipeline itself remains rigid, limiting the usability for custom rendering. Material parameter collections and custom nodes are used in the implementation of the Gaussian blur post-processing effect discussed in Section 7.1.3.

While it would also be possible to define the game logic via a visual scripting language called Blueprints in UE4, the simulation tool is implemented in C++. The main motivation is the ability to better represent mathematical functions and computations in C++ as compared to Blueprints, since the visual representation can get cluttered for such applications. Such functions are needed for example in the computation of the Gaussian blur's weights. Additionally, C++ code is expected to be more maintainable.

The HTC Vive is a VR system developed by HTC and Valve, which is compatible with the requirements defined by SteamVR. Its HMD and the motion controllers are tracked using the IR-laser based Lighthouse tracking system. It allows for 6-degrees-of-freedom tracking in a tracking space of up to 4.5x4.5m. The HMD features a display with a resolution of 1080x1200 pixels per eye and a refresh rate of 90Hz. The field of view (FOV) is specified as 110°. However, it is not clear from the specifications whether the FOV is measured horizontally or diagonally. The motion controllers themselves offer various input methods: A clickable trackpad, two buttons and a trigger button.

7.1.2 Interaction and Movement

Thanks to the direction integration of SteamVR, UE4 allows interacting with the HMD and the motion controllers in C++ via UE4's VR interfaces, namely *IHeadMountedDisplay* and *UMotionControllerComponent*. Through IHeadMountedDisplay, the position and orientation of the HMD relative to the origin of the tracking space can be queried. The transformation of the tracking space is applied to this position and orientation to get the HMD's position in world coordinates, equal to the camera position and orientation used for rendering. Similarly, models representing the motion controllers in VR can be transformed to the correct position and orientation relative to the HMD. The models are simply attached to the transformation of the UMotionControllerComponent.

In addition to these VR-specific interfaces, UE4 allows mapping the inputs from controllers, mouse or keyboard to so-called *UInputComponents*. There are two types of inputs, *actions* and *axis*: The former are related to buttons and similar input mechanics, which have a pressed, down, up and released state. The latter are related to trigger buttons, joysticks and trackpads where the input can be mapped to a continuous interval.

In our simulation tool, an action is used to trigger a feedback function whenever the user clicks the trackpad. Depending on the x- and y-axis values, resulting from the touch position on the trackpad, a certain direction is detected as input. The x- and y-value are converted to polar coordinates. If the resulting radius is beyond a safe zone, the input is considered valid. The angle then determines the desired direction. This is illustrated in Figure 7.2. Using this interaction, the user can report the recognized direction a safety sign (or Landolt ring) is pointing to. Figure 7.3 shows the input mappings used for the user interactions in the simulation tool.



Figure 7.2: From left to right: A touch on the trackpad is input as an x-y position on the trackpad with $x, y \in [-1, 1]$. The cartesian coordinates (x, y) of the input are converted to polar coordinates (r, α) , allowing easier mapping of input to selected direction. The texture visible to the user: Clicks inside the green area (r < 0.5) are ignored. Clicks outside this area are mapped to the related direction using α , e.g., for 45° the direction "top-right" is assumed.



Figure 7.3: Input mappings in UE4.

Both motion controllers provided by the HTC Vive are utilized for the movement: The transformation of the wheelchair, i.e., the real-world office chair, is determined using one of the motion controllers (controller A). Controller A is attached to the back of the office chair in which the user sits. Its tracking position and orientation are queried using the transformation of the UMotionControllerComponent related to this controller. The other controller is given to the user (controller B). It is used for giving feedback during the user study as well as to let the user control the speed of the movement. Whenever the user presses controller B's trigger button, the character in the simulation moves into the direction suggested by the wheelchair. This interaction is illustrated in Figure 7.4.



Figure 7.4: The wheelchair movement. (a) Controller A, attached to the back of an office chair. The tracked yaw angle of this controller changes when the user rotates the office chair. (b) The user controls the movement speed using the trigger button of controller B. (c) The capsule and thus the origin of the tracking space as well as the user's view are moved into the forward direction of the wheelchair by this interaction. (d) The required component attachments in UE4.

To allow for user movement beyond the tracking space, the origin of the tracking space itself is moved by the wheelchair movement. The tracking origin is represented by a capsule for collision detection. When a trigger-button input is detected for controller B, the capsule component is moved into the direction suggested by the yaw angle of controller A. The required component attachments are shown in Figure 7.4. The tracking space can only move to positions the capsule component can move to. Therefore, the users cannot walk through walls or cut corners.

7.1.3 Impairment Simulation

The simulation of the impairments caused by age-related eye diseases is implemented as a series of post-processing effects in UE4. The basic symptom simulated for all impairments is a loss of visual acuity based on a Gaussian blur. As outlined in Section 4.3, an efficient way to implement the Gaussian blur is by using bilinear filtering provided by the texture mapping hardware of modern GPUs. We want to construct this method by first considering the naive approach to the Gaussian blur.

0,000	0,000	0,001	0,002	0,001	0,000	0,000		0,004	0,000	0,000	0,001	0,002	0,001
0,000	0,003	0,013	0,022	0,013	0,003	0,000		0,054	0,000	0,003	0,013	0,022	0,013
0,001	0,013	0,059	0,097	0,059	0,013	0,001		0,242	0,001	0,013	0,059	0,097	0,059
0,002	0,022	0,097	0,159	0,097	0,022	0,002		0,399	0,002	0,022	0,097		0,097
0,001	0,013	0,059	0,097	0,059	0,013	0,001		0,242	0,001	0,013	0,059	0,097	0,059
0,000	0,003	0,013	0,022	0,013	0,003	0,000		0,054	0,000	0,003	0,013	0,022	0,013
0,000	0,000	0,001	0,002	0,001	0,000	0,000		0,004	0,000	0,000	0,001	0,002	0,001
							-	\bigotimes	0.004	0.054	0.242	0 399	0.242

(a)

(b)

0.000

0,003

0.022

0,003

0,000

0,054

0.000

0,000

0,001

0.002

0,001

0,000

0,000

0,004

Figure 7.5: The weights of a Gaussian kernel: (a) The full kernel, (b) the separable kernel. Note that multiplying the separable kernels results in the original weights of the full kernel.

In that case, an image is blurred by (discrete) convolution with the $h \times h$ sized Gaussian kernel, shown in Figure 7.5a. The values for the Gaussian kernel are sampled from the probability density function (PDF) of the Gaussian distribution $\mathcal{N}(\mu = 0, \sigma^2)$. The amount of blur resulting from the convolution with the Gaussian kernel is dependent on the kernel size h and the standard deviation σ of \mathcal{N} . For the convolution, the following procedure is repeated for all pixels of the source image: The kernel is centered over the currently considered pixel. The $h \times h$ pixels in the source image, which are now covered by the kernel, are first multiplied with their respective weight from the Gaussian kernel and then summed up. This sum is the new value for this pixel in the convolved, blurred image. Put into mathematical terms, the pixel at position (x, y) in the source image I is convolved with the $h \times h$ sized kernel k using Equation (7.1).

$$F(x,y) = I(x,y) * k = \sum_{i} \sum_{j} I(x+i,y+j) \cdot k(i+h/2,j+h/2)$$
with $i,j \in [-h/2,h/2]$
(7.1)

where * is the convolution operator and F(x, y) is the pixel value of the blurred image at position (x, y). We are assuming that the kernel sums to 1 and are not considering the handling of the image borders. As can be seen, this results in h^2 multiplications and texture lookups as well as $(h - 1)^2$ additions per pixel of the source image.

A more efficient implementation considers the separability of the Gaussian kernel. Separability means that the Gaussian kernel can be separated into a horizontal part with a $h \times 1$ kernel and a vertical part with a $1 \times h$ kernel. When applying these two kernels one after another, the result is the same as with the full $h \times h$ kernel. This is also illustrated by the multiplication of the two kernels, shown in Figure 7.5b. Note that both kernels feature the same weights, albeit transposed. The pixel at position (x, y) in I is convolved with the $h \times 1$ horizontal kernel k_h using Equation (7.2).

$$F_{h}(x,y) = I(x,y) * k_{h} = \sum_{i} I(x+i,y) \cdot k_{h}(i)$$
with $i \in [-h/2, h/2]$
(7.2)

where F_h is the horizontally blurred image. Analogous, by applying the vertical kernel k_v to the source image I, the vertically blurred image F_v can be computed. As shown in Equation (7.3), if we apply k_v to F_h , we get F, i.e., the image convolved with the original $h \times h$ Gaussian kernel. Applying k_h to F_v also results in F, meaning that the application of the separated kernels is interchangeable.

$$F_h(x,y) * k_v = \sum_j F_h(x,y+j) \cdot k_v(j)$$

=
$$\sum_j [\sum_i I(x+i,y+j) \cdot k_h(i)] \cdot k_v(j)$$

=
$$F(x,y) = F_v(x,y) * k_h$$

with $i,j \in [-h/2,h/2]$ (7.3)

The benefit of using the separable variant instead of the full kernel is, however, that they can be applied one after another: Computing F_h , or F_v , takes h multiplications and texture lookups as well as h - 1 additions. The same amount of operations is needed for applying the second kernel. Therefore, the effective effort for applying the separable Gaussian kernels one after another to compute F results in $h + h = 2 \cdot h$. This equals

7. Implementation Details

an effort in the order of O(h), one order of magnitude less than required for the naive Gaussian kernel which is in the order of $O(h^2)$.

As shown in previous work [Str], even further speedup can be achieved by utilizing the bilinear filtering provided by the texturing hardware of modern GPUs: Instead of sampling the source image at the center of the respective pixels, we chose sampling positions in-between two pixels. Using bilinear filtering, we can gather information on both pixels using a single texture lookup. The weights for the Gaussian blur using bilinear filtering, however, must be adapted according to the new sampling points, such that each pixel value is weighted as with the original weights in the separable Gaussian blur. This is achieved by applying Algorithm 7.1. The derivation from the separable Gaussian blur is illustrated by Figure 7.6.



Figure 7.6: From a separable Gaussian blur kernel to the weights and sampling positions for a Gaussian blur using bilinear filtering: (a) The weights of the separable Gaussian blur kernel. (b) The center pixel is sampled twice for simpler convolution code. Since the kernel would no longer sum to 1 if we would just duplicate the center weight, its weight is halved. (c) Weights and sampling positions for the interpolation via bilinear filtering: The weights are summed up pair-wise and the sampling positions are adapted (see line 14 in Algorithm 7.1) such that, when sampling them in a bilinear-filtered manner, the texture values are weighted as with the original Gaussian blur.

The computation of the weights w_i and sampling positions o_i for the Gaussian blur using bilinear filtering in our simulation tool is given by Algorithm 7.1: The weights g_i for a standard separable Gaussian blur kernel, sized $h \times 1$, are computed in lines 1 to 4. $\mathcal{N}(x|0,\sigma^2)$ computes the value of the PDF of the Gaussian distribution $\mathcal{N}(0,\sigma^2)$ at x. We take advantage of the symmetry of the kernel and only compute the n_1 weights from the center of the kernel at i = 0 to one end of the kernel. We will apply the kernel once in positive and once in negative direction. As a result, the center pixel would be sampled twice in this case. This is accounted for in line 5 by halving its weight g_0 . To normalize the resulting kernel such that it sums to 1.0, all weights are summed up and are then divided by this sum in line 7. Without this normalization the kernel would sum to a value $\neq 1.0$, resulting in the convolved image becoming darker or brighter. The w_i and o_i themselves are computed in lines 12 to 14: The weights w_i are simply the sum of the weights they replace, thus the normalization is preserved. The sampling positions o_i are adapted to these new weights in line 14. When sampling at the positions o_i , using bilinear filtering, the texture values are weighted as with the original Gaussian blur. Finally, if the outermost weight only replaces one weight, no bilinear filtering is required and the sampling position remains unchanged. This is accounted for in the computation in lines 8 to 10 by simply adding a zero weight to g_i .

Algorithm 7.1: Computation of the weights and sampling positions for a Gaussian blur using bilinear filtering.

```
Input: h, \sigma
    Output: w_i and o_i
 1 n_1 = \lceil h/2 \rceil;
 2 for i \leftarrow 0 to n_1 - 1 do
 \mathbf{3} \quad | \quad g_i = \mathcal{N}(i|0,\sigma^2);
 4 end
 5 g_0 = g_0/2;
 6 normalization = \sum g_i;
 7 g_i = g_i / normalization;
 s if n_1 \mod 2 \not\equiv 0 then
 9
         g_{n_1} = 0;
10 end
11 n_2 = \lceil n_1/2 \rceil;
12 for i \leftarrow 0 to n_2 - 1 do
13
         w_i = g_{2i} + g_{2i+1};
         o_i = 2i \cdot \frac{g_{2i+1}}{w_i} ;
\mathbf{14}
15 end
16 return w_i, o_i
```

The kernel generated by Algorithm 7.1 is applied in positive and negative direction per axis, effectively resulting in an $\lceil h/2 \rceil$ sized kernel per axis that is equivalent to an h sized separable Gaussian blur kernel. For example, the 16×1 kernel resulting from Algorithm 7.1 is equivalent to a 31×1 separable Gaussian blur kernel or a 31×31 naive Gaussian blur kernel. Using bilinear filtering, we can thus reduce the number of texture lookups and multiplications needed to 32, down from 62 for the separable or 961 for the naive Gaussian blur kernel, respectively.

The resulting weights w_i and sampling positions o_i are passed to the material controlling the post-processing effect via a material parameter collection. Since material parameter collections only accept vectors of length 4, the weights and sampling positions must be split into multiple vector parameters. In the post-processing material, these vector parameters are used as an input for a custom node, implementing the convolution of the post-processing input texture, i.e., the rendered scene.

Algorithm 7.2: Gaussian blur filter applied in x-direction.

Input: $h, w_i, o_i, uv, texture$ Output: sum 1 stepX = 1.0/texture.width;2 $kernelSize = \lceil h/2 \rceil /2;$ 3 sum = 0;4 for $i \leftarrow 0$ to kernelSize - 1 do 5 $| offset = [o_i \cdot stepX, 0.0];$ 6 | value = texture.sample(uv + offset) + texture.sample(texture, uv - offset);7 $| sum = sum + value \cdot w_i;$ 8 end 9 return sum

Algorithm 7.2 outlines the HLSL shader executed by the custom node for the x-direction. As for the separable Gaussian blur kernel, this has to be applied once for the x-direction and once for the y-direction. For the y-direction, the step size in line 1 depends on the texture height, and the offset in line 5 is applied to the sampling position in y-direction. The resulting post-processing material is directly applied to the HMD's rendering camera via so-called blendables, which allow overlaying several post-processing effects. For the plain Gaussian blur, the post-processing material is overlayed twice – once set to x-direction, once set to y-direction – resulting in the desired Gaussian blur effect.

The simulations of cataract and dry AMD use this Gaussian blur to simulate the loss of visual acuity. Additionally, for the cataract simulation, another post-processing material is overlayed as blendable. It tints the input texture with a yellowish color (RGB 0.5, 0.3, 0.0) by lerping between the tint color (0.05) and the input texture's color (0.95). A fog texture is further transparently overlayed by lerping between the fog texture (0.25) and the input texture (0.75) depending on the fog texture's alpha channel. For the dry AMD simulation, the post-processing material darkens the input texture by lerping between black (0.3) and the input texture's color (0.7). The effect is limited to the area given by the alpha channel of a masking texture.

7.1.4 Importing Models from HILITE

The building model from HILITE [VRVb] can be imported into UE4 using FBX as interchange format. Unfortunately, the FBX file produced by HILITE and the format

expected by UE4 seem to be incompatible and thus some model transformations are lost. These models have to be manually aligned in the graphical editor of UE4.

The imported materials are broken in terms of the lightmaps imported from HILITE. UE4 only supports lightmaps that have been baked in UE4 as static lighting. Yet, as the lightmap's UVs are imported correctly, the lightmap can be applied as pre-multiplied texture. The materials of models affected by a lightmap have to be adapted by adding a texture-sampler node in the Material Editor that samples from the lightmap using the lightmap UVs of the model. The sampled lightmap value is then multiplied with the diffuse color of the model.

Finally, to be able to navigate through a model in UE4, the geometry's bounding volumes for the collision detection are adapted. By default, the imported models only feature a simple bounding box. Thankfully, using the building geometry itself as bounding volume has no noticeable impact on performance in our test scene.

7.2 User Study

As can be seen in Figure 7.1, the experiments and the gathering of data thereof are implemented as a part of the simulation tool in UE4. The calibration procedure and the two experiments are realized as game logic with certain interactions and goals that vary between levels, i.e., experiments. The data gathered and the log files created during the user study serve as the basis for the evaluation in Chapter 8.

7.2.1 Game Logic

The simulation tool's behaviors for the experiments of the user study, as described in Section 5.1, are implemented as variations of an UE4 AGameModeBase in C++ . The GameMode defines behaviors and goals for the calibration and the two experiments. The different building models that are used are represented by levels containing the relevant geometry. Since the GameModeBase object is reinitialized when the level is changed and thus looses all state information, a UGameInstance is used to enable persistence across the different levels used per simulation run. The experiment protocol, the related levels and impairment settings as well as the calibration settings for the specific participant are saved in this GameInstance.

For the calibration procedure, the GameMode provides a method to generate lines of Landolt rings that relate to a certain visual acuity value. The position of the gap, i.e., the rotation of the ring, is generated randomly for each ring. The Landolt ring elements shown in the calibration level are then scaled and rotated accordingly. Furthermore, the visual acuity value as well as the σ parameter of the Gaussian blur are dynamically controlled by the GameMode as described in Section 5.2. The GameInstance defines the test distance, the visual acuity values for the weak and the strong blur as well as the steps for the Landolt ring sizes. In addition, the determined visual acuity and blur level parameters are saved by the GameInstance for use in the experiments.

To enable the first experiment, the determination of MRDs, different sign configurations have to be instantiated. The GameMode allows enabling the relevant sign model, apply the lightmap generated for this configuration, change the simulated impairment and to generate random directions for the safety signs. In contrast to the calibration procedure, which allows 8 directions, the possible directions for the experiments are limited to up, down, left and right. Also, on standardized signs, the arrow symbol for left or right usually points away from the pictogram. It is thus assumed that participants could rule out the directions "left" or "right" by recognizing the pictogram's position on a standardized sign. Therefore, we allow the position of the arrow symbol on the sign to be either to the left or to the right side of the pictogram.

As in the first experiment, the safety sign configurations have to be adapted for the escape scenario experiment. The different escape routes are defined as strings, encoding the sign's positions and directions they point to, in the GameInstance. These strings are parsed in the GameMode and the safety signs and their textures are changed accordingly. As before, the simulated impairment can be changed through the GameMode. In addition to the behavior in the MRD experiment, the participant's position is tracked by the GameMode through trigger volumes placed in the level. This allows us to associate the participant's feedback with the correct safety sign.

7.2.2 Reporting

Various data is recorded during each of these GameModes and serves as the basis for the evaluation presented in Chapter 8. The following is an example of a log file created during the MRD determination experiment. Additional line breaks were added for better readability.

```
Changed impairment to Clear - 00:00:00.000 -

X=-125.741 Y=5.664 Z=113.835 - X=5.901 Y=-11.200 Z=-4.483 - none - -1.0

Evaluation - 00:01:34.500 -

X=854.876 Y=-41.331 Z=108.926 - X=4.255 Y=-11.544 Z=3.215 - right - 1049.898

Evaluation - 00:02:19.338 -

X=924.169 Y=-10.776 Z=108.244 - X=4.873 Y=-14.987 Z=1.212 - right - 980.194

Evaluation - 00:02:58.769 -

X=1280.672 Y=-13.500 Z=109.044 - X=6.770 Y=-7.993 Z=2.225 - down - 626.117

Evaluation - 00:03:16.329 -

X=-402.538 Y=-11.642 Z=108.582 - X=4.924 Y=-13.553 Z=1.456 - right - 2304.382

Evaluation - 00:03:41.473 -

X=82.003 Y=-18.718 Z=108.982 - X=5.502 Y=-9.400 Z=1.876 - up - 1820.370

Evaluation - 00:04:02.898 -

X=206.548 Y=-20.453 Z=108.184 - X=6.154 Y=-16.236 Z=6.619 - left - 1696.062
```

The first line shows that the impairment for this run was set to "Clear", thus no impairment was simulated. The lines starting with "Evaluation" are records logged each time the participant recognized the direction indicated by the safety sign in the tested configuration, resulting in 6 logs. We do not consider incorrect participant feedback as

we want to compute the MRD of a safety sign based on the logged data: The MRD is defined by ISO 3864-1 [Int11] as the distance from a safety sign inside which 85% of observers can correctly recognize the sign. Therefore, only distances from which a safety sign is correctly recognized are required for the computation of the MRD and, thus, we only log the distances for correct input by the participants. Each log features the time it was recorded in *hh:mm:ss.sss*, the participants position (X, Y, Z) and rotation in Euler angles (X, Y, Z) as well as the reported direction and the distance to the safety sign.

The logs for the calibration procedure contain the calibration mode ("HMD calibration", "blur weak calibration" or "blur strong calibration") and the corresponding parameters determined for a participant. For the determination of the HMD's influence on visual acuity, this parameter is the last size of the last line of Landolt rings for which the participant was able to answer correctly, expressed as visual angle. This value is directly related to the visual acuity level the participant experienced while using the HMD. For the determination of the σ parameter for the two blur levels, the logged parameter is the last σ value for which the participant could correctly answer the line of Landolt rings. The size of this line is fixed to a visual angle of 5.0 for the weak blur and 8.0 for the strong blur, which corresponds to a decimal acuity value of 0.2 and 0.125, respectively.

For each run of the MRD determination, the log firstly contains the impairment that was simulated during this run. This is followed by 6 lines, one per sign configuration, containing the participant's position, the participant's orientation, the reported direction and the distance to the safety sign. These values are logged when the participant correctly recognizes the direction the arrow on the safety sign points to.

The log file created for the runs of the escape scenario experiment are the most extensive: Again, the simulated impairment is recorded first. Next, each second, the participant's position and orientation are logged. The participant's position and orientation are also logged each time a safety sign along the escape route is recognized correctly. The logs further include the reported direction, the safety sign's internal ID as well as the distance to the safety sign. The internal ID is the same that was used in the string to generate the safety sign configuration for this run. This generation string is logged as well.

7.2.3 Anti-Aliasing

During tests with the calibration procedure it became apparent that the finely detailed Landolt rings produced strong aliasing effects. Especially in combination with small head movements, the almost pixel sized gap in the Landolt rings for high visual acuity values begins to flicker. Since the rendering should not introduce an additional worsening of the perceived visual acuity, some form of anti-aliasing is needed.

To avoid aliasing effects, a combination of two approaches is applied in the simulation tool: First, the supersampling settings of UE4 are set to 200%. This results in the scene being rendered to a framebuffer of 200% of the display's resolution. The final output is sampled from this texture, heavily reducing the aliasing artifacts. In addition, the Temporal Anit-Aliasing (TempAA) provided by UE4 is activated. This method reduces

the flickering resulting from small movements over time by taking samples from different frames. While TempAA by itself seems to reduce the sharpness of the output, this effect could not be observed in combination with supersampling set to 200%.

7.3 Evaluation Tool

As already discussed in Chapter 6, we implemented our design evaluation tool using the Aardvark rendering platform and the accompanying UI framework Aardvark.Media. We were able to create the evaluation tool almost entirely using the functionality already provided by Aardvark, with slight modifications required to enable geometry-shader based cube-map rendering. In the Aardvark.Media framework, the interface elements are specified as HTML-elements and backed by program logic in F#. The rendered view itself is controlled using these interface elements: The user inputs are generated by the HTML-elements that are presented on the client. The inputs are then sent to a server, which carries out the rendering of the scene. Depending on the inputs, the rendering is adapted. The resulting rendered images are compressed and sent back to the client, where they are displayed as an image stream inside a HTML-container. As a side effect, this allows using the evaluation tool on low-performance hardware by leaving the actual rendering to a server. The client only has to handle the input and the decompression of the image stream.

7.3.1 Rendering to a Cube-Map Texture

To implement omni-directional shadow mapping using a geometry shader, we required the ability to render to a framebuffer object (FBO) that is backed by a depth-only cube-map texture as target. This was not possible using Aardvark, but could be fixed by exposing the corresponding flag in the FBO creation.

Still, since this lower-level functionality was not exposed until then, also the higher-level rendering functions did not support such FBOs. Namely, RenderTask and AdaptiveRenderingResult had to be adapted such that they allow FBOs with a cube-map texture as targets as well as to allow the execution of a RenderTask on such a target.

7.3.2 Visualizing the User-Study Data

To create a visualization of the data gathered during the evacuation scenario experiment of the user study, the evaluation tool offers the possibility to load log files generated by the simulation tool. The visualization is based on the recorded positions and inputs of the participants: For each record, we plot a glyph representing the type of record at the respective position. The type can either be a position record, which is logged each second, or a user-input record, which is logged when a participant reports the recognized direction. Further, a line connecting the records in chronological order is plotted to illustrate the participants' movement path. Since a participant might have paused at a certain position, these glyphs can overlap. For glyphs that represent a position record, this is unproblematic. However, if a participant paused to recognize and report the recognized direction of a sign, the glyph related to this input might become occluded by position glyphs. Users of the evaluation tool would no longer be able to keep track of this participant's input. Therefore, after the parsing of a log file, we cleanup the records to remove position glyphs in close proximity to user-input glyphs.

This procedure assumes that the participants do not cross their own path around a user-input record's position. We want to use all user-input records and only use position records whose positions are farther than 0.1 meters from a user-input record. To achieve this, we first create a list of user-input records and set the first entry as our current user-input record. Now we iterate through the position records in chronological order: We compute the distance of the position record to the user-input record. If the distance is smaller than 0.1 meters, the position record is omitted and thus no position records can overlap the user-input record. If the position record is beyond that distance and comes chronologically after the user-input record, the index for the current user-input record is incremented. The position record again must not lie inside a 0.1 meter radius from this record. If we run out of user-input records, all remaining position records can be kept.

After this cleanup procedure, we have a list of user-input records and position records, represented by their position in world space. In addition, we encode the type and – in case of a user-input record – the reported direction into an info parameter. This controls which texture is applied to the glyphs.

During the rendering of the visualization, we pass the records as vertices with their world-space positions and the info parameters as associated vertex arrays to a geometry shader. This geometry shader emits a quad per record, centered at the record's position and oriented towards the camera. The latter is achieved by offsetting the quad's vertices from the world-space position of the record vertex by an offset in screen space. The quad's texture coordinates are determined by the value of the info parameter. In the fragment shader, we sample from a texture containing textures for all glyphs: Depending on the texture coordinates, the texture shows a flag (position record) or an arrow into one of three directions (user-input record). The circular form of the final glyphs is achieved by only rendering fragments that are inside a radius from the glyph center. Similarly, we plot a line between each record: Another geometry shader emits a quad between each two records' positions.

CHAPTER 8

Results and Discussion

In the following chapter, the results of this thesis, as well as its limitations, are presented and discussed. First, Section 8.1 presents the results achieved using our simulation tool. Next, Section 8.2 discusses the results of the user study. The data gathered from the calibration and the two experiments is evaluated. Finally, Section 8.3 illustrates the usage of the design evaluation tool to examine the escape-route coverage for a specific placement of safety signs on the example of the building used for the user study.

In addition, the limitations of the presented tools and the results from the user study are discussed in Section 8.4, offering insight into possible starting points for future extensions. Section 8.5 then completes this chapter with the comparison of our approach and the results presented in this thesis with the previous work discussed in Chapter 3.

8.1 Simulation Tool

The 90 FPS required for a VR application using the HTC Vive could be constantly maintained on an Intel i7 7700K and a NVIDIA Geforce GTX 1080, even though the simulation tool had to process extensive geometry and various post-processing effects in the course of the user study. However, when tested on hardware featuring a weaker GTX 1070, the main limitation was found to be supersampling. Reducing the supersampling factor (see Section 7.2.3 for a detailed description) from 200% down to 150% resulted in a sufficient speedup to maintain the 90FPS, without visibly introducing aliasing artifacts that would influence the calibration procedure. Even though the Gaussian blur post-processing effect, presented in Chapter 4, is used with a kernel size equivalent to a 31×31 naive Gaussian blur kernel, the frame rate never dropped below 90FPS in our experiments.

The wheelchair movement (see Section 4.6) used as locomotion technique, presented in this thesis, was accepted well by the participants of the user study. We carried out informal interviews during and after the experiments. The participants reported that they were able to intuitively navigate through the scenes used in the user study. While some participants reported a slight feeling of dizziness after the 30min VR sessions, only one participant had to stop and the overall feedback was very positive. In addition, the interaction for reporting the recognized safety-sign directions could be used by all participants right away. Even participants without experience with VR reported no problems in terms of locomotion or input.

8.2 User Study

The user study was conducted with 30 participants aged between between 23 and 42, 10 of them female and 20 male. 50% of the participants had normal vision and the other 50% needed correction, wearing either glasses or contact lenses. However, two of those needing correction did not wear any during the user study. Most of the participants needing correction were shortsighted, with some also having astigmatism. All but one participant had previous experience with computer games, while two thirds also already tried VR before. In the course of the user study, one participant got motion sick and the experiments for this participant had to be stopped. Some other participants reported a slight feeling of dizziness after the 30min VR session.

8.2.1 Outliers

We analyzed the dataset gathered from the user study for outliers. First, the participant who had to stop because of motion sickness was removed from the dataset. While inspecting the calibration data, two additional participants stuck out: One, for whom the calibration to "weak blur" (equal to a mild reduction of visual acuity) as well as "strong blur" (equal to a more severe reduction of visual acuity) resulted in a comparably minor reduction of visual acuity, and another one, for whom "weak blur" and "strong blur" were calibrated to the same reduction of visual acuity. Naturally, the first participant, for whom the visual acuity was hardly reduced by the simulation, achieved comparably high recognition distances in the MRD determination experiment. The equal visual acuity in "weak blur" and "strong blur" for the second participant make the reported recognition distances ambiguous. Therefore, the data related to those two participants was also removed from the dataset. We assume that the first participant might have misunderstood to which Landolt ring his/her following input will be matched and thus failed to correctly carry out the calibration procedure, failing before the actual limit of visual acuity was reached. Similarly, as a result of erroneous input, the calibration procedure at the "strong blur" level for the second participant might have ended before the actual limit of visual acuity was reached, thus resulting in no further reduction in the simulated visual acuity from the "weak blur" level.

In the data gathered from the MRD determination experiment some comparably extreme values could be observed: Very long recognition distances are assumed to result from participants guessing a direction without actually recognizing the sign, although they were explicitly asked not to guess. Very short distances may be a result of participants being inattentive or accidentally pressing another direction than the one recognized. Since pressing the wrong direction results in a short vibration, indicating that the input was false, the participant then might assume to have falsely recognized the direction. Therefore, participants would then approach closer than required, confirming the original guess but shortening the measured distance. To reduce the influence of such outliers, measurements that deviate more than 3 standard deviations from the mean were removed from the dataset.

8.2.2 Calibration



Figure 8.1: The impact of the HTC Vive's display on the visual acuity achievable in VR. None of the participants could achieve a visual acuity better than 0.4 decimal acuity. This corresponds to a considerable vision impairment.

Using the calibration procedure presented in Section 5.2, the maximum level of visual acuity achievable when using the HTC Vive could be determined. Figure 8.1 shows the visual acuity achieved by participants. As can be seen, no participant could correctly detect lines of Landolt rings sized smaller than 2.5 angular minutes, equivalent to 0.4 decimal acuity. In general, the limitation on achievable visual acuity introduced by the HTC Vive seems to be between 0.3 and 0.4 decimal acuity. This variation is expected to stem from whether a participant was able to guess the line sized 3.2 angular minutes correctly or not. Furthermore, Figure 8.1 features the values of three participants who could not successfully complete a line of Landolt rings equivalent to a visual acuity worse than 0.4 decimal acuity. Their visual acuity was either already worse than the limitation introduced by the HTC Vive, or they mistakenly input wrong directions.

8.2.3 MRD Determination

Figure 8.2, Table 8.1 and Table 8.2 show the results of the MRD determination experiment. The participants' average recognition distances, as well as the standard deviation, are determined per run, split up into the different blur levels and safety sign configurations.

It can be observed that the average recognition distance approximately doubles as the size of the safety sign doubles. This is consistent with the MRD calculation models [DIN13, Int11, XFG⁺07] where the sign's height h is a multiplicative factor. Further,



Figure 8.2: The boxplots show the recognition distances in centimeters, determined during the two runs of experiments with no blur (0.4 decimal acuity caused by the HTC Vive), weak blur (0.2 decimal acuity) and strong blur (0.125 decimal acuity). The red curves shows the MRDs computed from the recognition distances of both runs. For comparison, the upper gray curves shows the MRD as suggested in EN 1838. The lower gray curves shows the MRDs as computed according to the suggestions in ISO 3864-1. They depend on the observation angle by a factor of $cos(\alpha)$ and are scaled by the decimal acuity value.

depending on the observation angle α , the recognition distances decrease similar to $\cos(\alpha)$. Such a decrease is suggested by ISO 3864-1 and Xie et al. [XFG⁺07], yet our data suggests that this is additionally influenced by the reduced visual acuity: For example, the reduction factor for the decimal acuity values 0.4, 0.2 and 0.125 were 0.69, 0.52 and 0.42 for the first run of the 15cm safety sign at 0°. ISO 3864-1 suggests that the recognition distances are reduced by a factor equal to the decimal acuity of the observer. As Figure 8.2 shows, this underestimates the recognition distances recorded during the user study. Figure 8.2 also shows how EN 1838 generally overestimates the recognition

		first	run	second run					
	weak	blur	strong	g blur	weak	blur	strong blur		
sign size, rotation	\bar{x}	σ	\bar{x}	σ	\bar{x}	σ	\bar{x}	σ	
$15 \text{cm}, 0^{\circ}$	$783 \mathrm{cm}$	$148 \mathrm{cm}$	$622 \mathrm{cm}$	$174 \mathrm{cm}$	847cm	$197 \mathrm{cm}$	$570 \mathrm{cm}$	141cm	
$15 \text{cm}, 30^{\circ}$	$778 \mathrm{cm}$	$223 \mathrm{cm}$	$521 \mathrm{cm}$	$128 \mathrm{cm}$	747cm	$162 \mathrm{cm}$	$507 \mathrm{cm}$	$109 \mathrm{cm}$	
$15 \mathrm{cm}, 60^{\circ}$	$598 \mathrm{cm}$	$163 \mathrm{cm}$	$371 \mathrm{cm}$	$95 \mathrm{cm}$	590cm	$89 \mathrm{cm}$	$388 \mathrm{cm}$	$101 \mathrm{cm}$	
30 cm, 0°	$1675 \mathrm{cm}$	$340 \mathrm{cm}$	$1028 \mathrm{cm}$	$226 \mathrm{cm}$	1665cm	$381 \mathrm{cm}$	$1069 \mathrm{cm}$	$251 \mathrm{cm}$	
30 cm, 30°	$1550 \mathrm{cm}$	$264 \mathrm{cm}$	$1009 \mathrm{cm}$	$266 \mathrm{cm}$	1570cm	$312 \mathrm{cm}$	$1065 \mathrm{cm}$	$256 \mathrm{cm}$	
$30 \mathrm{cm}, 60^{\circ}$	$1112 \mathrm{cm}$	$242 \mathrm{cm}$	$761 \mathrm{cm}$	$191 \mathrm{cm}$	1189cm	$214 \mathrm{cm}$	$833 \mathrm{cm}$	$258 \mathrm{cm}$	

Table 8.1: The mean \bar{x} and standard deviation σ of the recognition distances over all observations per run and safety sign configuration.

	firs	t run	second run		aver: both	age of 1 runs	$\mathbf{MRD:}\\ cos(\alpha)$	$\mathbf{MRD:}\\ cos(\alpha) \cdot dA$				
valid for 85%												
sign size,	weak	strong	weak	strong	weak	strong		weak	strong			
rotation	\mathbf{blur}	\mathbf{blur}	blur	blur	blur	blur		blur	blur			
$15 \text{cm}, 0^{\circ}$	6.4m	4.4m	6.9m	4.1m	6.7m	$4.3 \mathrm{m}$	$15.0\mathrm{m}$	$3.0\mathrm{m}$	$1.9\mathrm{m}$			
$15 \mathrm{cm}, 30^{\circ}$	6.2m	4.2m	6.1m	4.0m	6.1m	4.1m	$13.0\mathrm{m}$	2.6m	1.6m			
$15 \mathrm{cm}, 60^{\circ}$	4.6m	$2.8\mathrm{m}$	5.0m	$3.0\mathrm{m}$	4.8m	$2.9\mathrm{m}$	$7.5\mathrm{m}$	1.5m	$0.9\mathrm{m}$			
$30 \text{cm}, 0^{\circ}$	13.5m	8.5m	13.0m	$8.9\mathrm{m}$	13.2m	$8.7\mathrm{m}$	$30.0\mathrm{m}$	$6.0\mathrm{m}$	$3.8\mathrm{m}$			
$30 \mathrm{cm}, 30^{\circ}$	$13.1\mathrm{m}$	$7.1\mathrm{m}$	13.2m	8.2m	13.1m	7.6m	$26.0 \mathrm{m}$	5.2m	3.2m			
$30\mathrm{cm},60^\circ$	9.4m	5.8m	10.5m	6.2m	9.9m	$6.0 \mathrm{m}$	$15.0\mathrm{m}$	$3.0\mathrm{m}$	$1.9\mathrm{m}$			

Table 8.2: The MRDs determined in the user study, valid for 85% of the participants. For comparison, the MRDs resulting from the angle-dependent model $z_0 \cdot h \cdot cos(\alpha)$ suggested by ISO 3864-1 as well as the MRDs resulting from a further scaling by the decimal acuity value, $z_0 \cdot h \cdot cos(\alpha) \cdot dA$. The visual acuity levels "weak blur" and "strong blur" relate to decimal acuity values of 0.2 and 0.125, respectively.

distances. To deduce a model that relates the decimal acuity exactly to the expected recognition distances, measurements for more visual acuity levels and more observation angles are needed. The latter should be investigated because our data already gives a good indication that there might be a correlation between these two influence factors.

As ISO 3864-1 suggests, the MRD should be valid for 85% of observers. Therefore, the distance covering 85% of the participants' recorded recognition distances represents the MRD for a visual acuity level and safety sign configuration. The resulting values are shown in Table 8.2. Again, the reduction of the MRDs by the decimal acuity values would result in an underestimation. For example, the visual acuity dependent reduction factors for the first run of the 15cm, 0° configuration and decimal acuity values of 0.4, 0.2 and 0.125 are determined to be 0.55, 0.51 and 0.38.

Judging from this data, neither the MRD computation models suggested by EN 1838 nor ISO 3864-1 seem to be able to capture the dependence of an observer's recognition

distance on the experienced visual acuity correctly: EN 1838 overestimates the MRDs. When used to determine a placement of safety signs, this would result in an incomplete coverage of the escape route as the actual recognition distances would be shorter than suggested by the MRDs computed using the model of EN 1838. Conversely, ISO 3864-1 underestimates the MRDs. This would suggest the use of more safety signs than actually required, thus increasing costs. The results presented in this thesis serve as the basis for the determination of the relationship between an observer's recognition distance and the experienced visual acuity. Yet, it would be of interest to repeat the MRD determination experiment in a more detailed manner in the future to determine the exact nature of this relationship, providing factors which are dependent on visual acuity to be used in the MRD computation models.



Figure 8.3: Recognition distances in centimeters, determined during the two runs of experiments with no blur, weak blur and strong blur. The left column shows the results for the normal-sighted participants. The right column shows the results for the participants requiring vision correction.

8.2.4 Validity Checks

As can be seen from the experiment protocol in Section 5.1, each participant had to carry out the MRD determination experiment with all three visual acuity levels twice, alternating with runs of the escape scenario experiment. Resulting from this, the participants had breaks of 10 to 15 minutes between two MRD determination experiment runs with the same impairment level. These runs are also illustrated by the top and bottom row of Figure 8.2. A t-test was used to check for a learning effect in-between the two runs. Because all p-values are above the standard $\alpha = 0.05$ cutoff value, it can be concluded that there is no evidence for a learning effect.

Figure 8.3 shows the recognition distances for the sub-groups of normal-sighted participants and participants requiring vision correction. Welch's t-test was used to test for a significant difference between these two sub-groups: For the first run of experiments with weak and strong blur, evidence for a significant difference was found. Yet, for the second run of experiments, the t-test shows no evidence for a significant difference between normal-sighted participants and participants requiring vision correction. This suggests that there is no systematic error in the user study design. Yet, the t-tests that show evidence for a significant difference require further experiments and analysis to determine the exact cause for this difference. It may be due to a too small sample size or yet unknown parameters. Moreover, the 4 out of 12 t-tests that showed a significant difference might be false positives.

Analysis showed no evidence for an influence of the participants' previous experience with computer games or with VR. Furthermore, there is no evidence that gender influenced the participants' performance in the user study.

8.2.5 Escape Scenario

As with the MRD determination experiment using the corridor map, we evaluated the recognition distances in the escape scenario experiment. However, since the participants could move freely through the building, each participant could approach the safety signs differently. Some tried to stay away from the walls of the building, centering themselves in the rooms and corridors while moving around. Others stayed close to the walls, trying to cut corners as much as possible and even peeking around corners. Therefore, we defined two measures: The farthest distance and the the nearest distance from which a particular safety sign could theoretically be recognized, meaning those for which it is not occluded by geometry. These distances are shown in Figure 8.4 as dashed, black lines. Because participants occasionally forgot to report the direction for some safety signs, there are no recognition distances for these samples. We omitted these distances from the evaluation, resulting in broken lines in Figure 8.4.

For the remaining recognition distances, we computed the mean recognition distance per run and per sign and compared these means to the recommendations by EN 1838, ISO 3864-1 and our MRD determination experiment for safety signs of 15cm height, observed head-on. The results are shown in Figure 8.4. Note that we did not consider



Figure 8.4: The recognition distances per safety sign in the evacuation scenario experiment. The semi-transparent lines show the recognition distances per user. The dashed lines illustrate the farthest distance (top), within which the sign becomes theoretically visible, and the nearest distance (bottom), within which the sign is guaranteed to be theoretically visible. The gray lines show the MRDs as suggested by EN 1838 (top) and ISO 3864-1 (bottom). Finally, the MRD suggested by our experiments is shown in red and the mean recognition distances of all participants is shown in green.

the observation angle for either ISO 3864-1 or the MRDs determined by our experiments, as the participants could face most signs only (approximately) head-on. The evaluation shows that, as with the MRD determination experiment, EN 1838 overestimates the achievable recognition distances, while ISO 3864-1 underestimates them.

The MRDs that result from our MRD determination experiment seem to be a usable approximation to the achievable recognition distances. We can also see that, compared to the simple setting of moving down a corridor, when moving freely, participants might not recognize a safety sign right away but only after having oriented themselves. Furthermore, the geometry of the building itself becomes an important factor. The results for safety sign number 12 in Figure 8.4 show that one cannot separate recognizability, considered by the MRDs, from visibility: While the sign would be recognizable from 6.7m based on our MRD, it only becomes visible from 5.5m (farthest) or 3.5m (nearest), respectively. A combination of both considerations is provided by the evaluation tool presented in this thesis.

8.3 Evaluation Tool

The design evaluation tool presented in this thesis offers the inspection of building models and the safety signs placed therein in terms of escape-route coverage. Using a responsive UI and rendering of the scene in real time, lighting designers can interactively examine the visibility of the safety signs, evaluated for a horizontal and a vertical slice along the escape route. However, the users cannot move the safety signs directly in the evaluation tool, which, in combination with the potentially time-consuming import of huge building models, is a limiting factor for the applicability of the evaluation tool. Still, experiments using the evaluation tool and the building model used in the user study suggest that it provides a useful extension to the set of existing light-planning tools: The immediate update of the visualizations as parameters of the evaluation tool are changed allow experimenting with different safety sign heights and models for MRD computation. Figures 8.5 and 8.6 illustrate how this facilitates the determination of a gap in the escape-route coverage: The slices representing the escape route indicate recognizability of at least one safety sign for the model suggested by EN 1838. However, when using the model suggested by ISO 3864-1 that additionally considers the observation angle, a segment of the escape route is no longer covered. A person moving along the escape route would not be able to recognize the next safety sign from this region.



Figure 8.5: Visualization of the visibility regions inside the test building as rendered by the evaluation tool. The MRDs are generated by the model which is suggested by EN 1838. The model is not able to represent the reduction of recognizability of signs following from observation under a flat angle. This results in an overestimation of the escape-route coverage.

Our evaluation tool allows lighting designers to detect such problematic areas and it helps them to adapt their designs accordingly. We found the visualizations provided by the evaluation tool to be helpful, as the vertical and horizontal coverage can be examined



Figure 8.6: Visualization of the visibility regions inside the test building as rendered by the evaluation tool. The MRDs are generated by the model which is suggested by ISO 3864-1 and Xie et al. It incorporates the observation angle and the corresponding visualization reveals a gap in the escape-route coverage (blue).

in combination and their position can be easily changed. Adding further viewing modes, e.g., a top-down view, might provide a better overview compared to the fly-through camera used at the moment.

Nevertheless, the fly-through camera was found to be a suitable viewing mode to compare the visibility visualization of the evaluation tool to the data gathered from the user study. Directly following the participants' movements allows better understanding the participants' feedback, letting lighting designers take the view of the participant at problematic areas along the escape route. In general, the visualization of the data gathered from the user study is useful to determine areas where participants had to come closer or could stay farther away than suggested by static visibility and the MRDs. Figure 8.7 illustrates an example with data gathered from the user study presented in this thesis. Especially when extending the simulation by further symptoms, the deviation from participants' recognition positions and static visibility might increase: The safety signs' visibility becomes restricted for participants who are exposed to symptoms such as scotoma in the central vision field, which means they have a partial loss of vision.

In terms of its usability for future work, the evaluation tools can be used with other building geometry and for further user studies using the simulation tool: Reports generated by the simulation tool can be imported and visualized, as shown in Figure 8.7. Also, building geometry compatible with the used Open Asset Import Library (Assimp) [The] can be imported, although only a subset of the features of Assimp is accessible through the evaluation tool. Depending on the intended model, it might either be required to convert the model or to extend the capabilities of the evaluation tool.



Figure 8.7: Visualization of the data from the experiments inside the test building: Green glyps show the direction that was reported for the safety sign, gray glyphs show the position of the participant in 1sec intervals and the gray line illustrates the chronological order of these glyphs.

In addition, the Aardvark.Media framework allows easily extending the provided UI by further parameters, controlling the visualization, safety signs and building geometry. Yet, there is currently no possibility to get information considering illumination per safety sign from HILITE. A closer integration with HILITE would enable to individually adapt the size, illumination and position of each safety sign as required by the determined escape-route coverage. Additionally, this would allow extending the MRD computation model used in the evaluation tool by the illumination data that can be gathered in HILITE.

8.4 Limitations

We presented an impairment simulation tool, experiments with impaired vision in VR and a design evaluation tool for lighting designers. Each of these proved valuable to test our new user-study methodology, resulting in promising results that can guide further extensions. Still, certain limitations of the tools and approaches discussed in this work could be determined.

8.4.1 Simulation Tool

As with previous work in the simulation of vision impairments, the simulation tool is not able to simulate the dependence of some symptoms on the user's viewing direction, such as the vision field losses resulting from diabetic retinopathy. This is due to the simple fact that the user's viewing direction is not tracked. An eye-tracking device that can be fit inside the HMD would be required for this task.

Furthermore, as many age-related eye diseases are accompanied by several symptoms of varying severity, a realistic simulation would incorporate a calibration for each of these symptoms. Since the scope of this work is limited to the calibration and realistic simulation of a loss of visual acuity, more complex disease patterns can only be simulated by hand-crafting effects. While the creation of these effects is based on descriptions of the individual symptoms [VVR16, AFF15, VCH16, JAR05, NEI] as well as the results of related work [LBCM11, LSB12, VVR16, AFF15, VCH16, JAR05], they can only be an educated guess of how the actual impairment might look like and are fixed in terms of the overall severity. Still, by also using the calibrated visual acuity, one of the symptoms of the eye diseases is expected to be plausibly simulated.

Finally, the physically plausible lighting generated by HILITE might differ from the subjective brightness experienced by the users: The safety signs' brightness correctly relates to the overall scene's brightness, but the absolute brightness perceived by the users also depends on the HMD's brightness. We would have to adapt the brightness of a pixel in the rendered scene according to a mapping between the pixel's value and the resulting luminance in the display. For the HTC Vive, this mapping could not be determined.

8.4.2 User Study

Since the HTC Vive was found to introduce a reduction of visual acuity itself, no data for "unimpaired" vision could be gathered. Thus, no direct comparison with the MRDs suggested by the standards and regulations is possible.

Another limitation in the data gathered is that only two visual acuity levels were tested. However, the 30min VR sessions that were required to test only these two levels seem like the maximum in terms of duration, as several participants reported slight dizziness afterwards.

For the evacuation scenario experiment, the participants could freely move through a largely empty building. Yet, in a real-world scenario, one would assume that other people are trying to escape, using the same escape route, and that some sort of danger would be present. For example, some segments of the building should be filled with smoke or fire.

8.4.3 Evaluation Tool

In our current implementation, the placement of safety signs and escape routes cannot be interactively changed in the evaluation tool, which limits the lighting designer's ability to react to problematic areas, determined using the evaluation tool. The paths of the escape routes themselves are normally provided together with the building model through so-called building information model (BIM) data, but this is not yet provided by the HILITE export. Therefore, a closer integration of the evaluation tool with HILITE or vice-versa is desired for a future version.

As mentioned in Section 6.2, the visibility computation carried out in the evaluation tool is only an approximation to the from-region visibility that would be required to determine whether the full safety sign is visible to an observer. The solution presented in this work assumes that it is sufficient to consider only the corners and the center of a safety sign. In cases where this assumption fails, also the visibility computation of the evaluation tool returns invalid results. Further, the visibility computation visualized in the tool does not account for the viewing direction and the visual field of an observer. Thus, some theoretically visible signs might not be visible to an actual observer who is moving along the escape route.

To allow users to freely move around the evaluation slices of the visualization in the evaluation tool without requiring a redraw of the shadow maps, we decided to let the view frusta of the cube mapping span the full hemi cube from which a safety sign is theoretically visible. By tightly fitting the view frusta around these two slices, either the accuracy or the required resolution of the shadow map could be optimized as compared to our current solution. However, this would require to redraw the shadow maps each time the slices are moved or rescaled.

8.5 Discussion

The presented results and the limitations that remain with this work allow setting this thesis in context with its related work. While the tools and approaches presented in this work perform favorably for the studied application of determining recognition distances for safety signs under impaired vision, starting points for further extensions can still be identified.

8.5.1 Impairment Simulation and User Study

We simulated the loss of visual acuity using the simulation tool presented in this thesis. In most of the previous work, the simulation of this symptom is done by blurring the output until it subjectively "looks blurry enough". In contrast to previous work [AFF15, VCH16, LBCM11, LSB12], the simulation tool provides a representation of the effect of a reduced visual acuity, whose severity is calibrated per participant. The simulation of a reduced visual acuity presented in this thesis is calibrated using a standardized medical test and can therefore be used for quantitative experiments. Compared to our method, previous work [JAR05, AFF15, LBCM11, LSB12] mainly deals with qualitative studies on whether the simulated symptoms convey an impression strong enough for the mostly educational purposes of the particular work. In this thesis, further simulations of mild forms of age-related eye diseases are based on the calibrated simulation of a loss of visual acuity. As a result, the two simulated age-related eye diseases presented in this work are also assumed to better represent the actual impairment experienced by affected

8. Results and Discussion

persons compared to previous work. Age-related eye diseases therein [VVR16, AFF15] are simulated based on the artistic impressions provided by the NEI [NEI].

Additionally, the impairment simulation tool features realistic lighting. As the lighting conditions directly influence the perception of safety signs, this is an important factor to consider when studying the influence on MRDs. This also allowed us to carry out quantitative experiments in the course of a user study. Previous work, discussed in Section 3.3, uses common lighting models from real-time rendering, thus not reproducing the lighting scenarios to be experienced in the real world.

However, the simulation tool does not account for other people in evacuation scenarios or smoke and fire along the escape route. Almeida et al. [AJF⁺14] found such distractions to affect the decision making of participants. Cosma et al. [CRN16] include smoke in their evacuation scenario. Still, extending the simulation tools by these features should be straight-forward, as UE4 allows the simulation of other characters and provides an efficient implementation of particle systems that can be used to simulate smoke and fire effects.

8.5.2 Visibility Computation and Design Evaluation

The visualization of a safety sign's visibility in buildingEXODUS [Xie11] as well as the extensions presented by Xie et al. [XFG⁺07, Xie11] are restricted to a coarse 2D grid. The height at which the grid is placed is selected beforehand. No information on the visibility along the escape route is provided. Compared to this, the evaluation tool presented in this thesis allows inspecting the regions in which the safety signs are visible more easily and in more detail: The visibility slices parallel to the building's floor can be moved freely, allowing to quickly check different eye heights. In addition to this, it is possible to inspect the visibility along an escape route through a vertical slice. This enables lighting designers to determine problematic areas along escape routes where the visibility of safety signs might be limited. Furthermore, we created a visualization of the data gathered using the simulation tool, enabling lighting designers to consider building occupants suffering from the simulated eye diseases.

Compared to the from-point visibility methods used for the visibility computation in evacuation models [Xie11, Pan06, CPLL15], our evaluation tool uses an approximate, structured sampling to consider the safety signs' areal extent in the visibility computation. While the sampling of visibility for the corners and the center of the safety sign applied in this thesis can only approximate the from-region visibility of the safety signs, conservative from-region visibility algorithms [WWS00, DDTP00] would determine regions from which it is guaranteed that the safety sign is fully visible.

The methods presented by Pan et al. [Pan06] and Chu et al. [CPLL15] compute the safety signs' visibility from the observer's perspective, thus enabling the consideration of the observer's vision field. This would also allow extending the visibility computation by considerations on symptoms of eye diseases, e.g., partial vision loss in certain regions of the vision field of an affected person. As we compute the visibility from the safety

sign's point of view, such considerations cannot be added easily. For our presented use case, it seems more appropriate to consider the sign's point of view. However, in case the behavior of building occupants should be incorporated in a future version of the evaluation tool, a switch to a computation from the observer's point of view might be advisable.

Existing evacuation models, such as buildingEXODUS, not only allow the computation of visibility areas but also the simulation of evacuation scenarios. In such evacutation scenarios, characters representing the building occupants are programmed to try to reach an emergency exit. Their movements are based on behavior and wayfinding models, including information on the safety signs' visibility and further heuristics. With respect to this, the evaluation tool is limited to visualizing data gathered from the user study through glyphs. By extending the evaluation tool by functionality such as in buildingEXODUS, behavioral aspects of the observers could be included into the considerations of the escape-route coverage.
CHAPTER 9

Future Work and Conclusion

This final chapter summarizes the contributions of this thesis and reflects on their applicability in the domain of lighting design and beyond. In addition, we present several ideas for extensions of the simulation tool and the evaluation tool presented in this work.

9.1 Conclusion

As could be shown in this thesis, the simulation of the symptoms of eye diseases in VR is a powerful tool, providing a controlled testing environment and allowing experiments that would be otherwise hard to carry out in the real world.

We provide a more realistic simulation of the reduction of visual acuity than previously achieved. Our calibration procedure for the simulated symptoms using standardized medical tests allows us to calibrate the experienced visual acuity to any decimal acuity value, and it accounts for the user's actual visual acuity. The medical test used in our calibration procedure could be transferred to VR easily, and further tests can be added similarly in future versions of our simulation tool.

We determined the impact of the HTC Vive's HMD on the visual acuity achievable in VR using a new simulation approach and determined that the visual acuity achievable using the HTC Vive corresponds to a considerable vision impairment. Furthermore, we determined the recognition distances for safety signs, which were observed under two different levels of visual acuity. This allowed us to compute MRDs for specific safety sign sizes and different angles. Also, we could evaluate how these values translate into an actual escape scenario where participants have to follow these safety signs to escape from a building. We conclude that the existing standards fail to correctly estimate the MRDs that result from the recognition distances we determined. While EN 1838 overestimates the MRDs, which results in an incomplete coverage of the escape route, ISO 3864-1 underestimates the MRDs, requiring more safety signs and thus increasing costs. We

could show that in an escape scenario, the MRDs determined using our procedure more closely approximate the recognition distances achievable by participants of our user study. Nevertheless, further experiments are needed to validate our preliminary results as well as to enable the definition of visual acuity dependent factors for MRDs to be used for planning.

In addition, this thesis presents a design evaluation tool which allows lighting designers to inspect their designs in terms of visibility and recognizability of safety signs. Moreover, we provide the possibility to visualize and retrace individual simulation runs from the user study in the evaluation tool. Combining these two features, lighting designers can effectively determine problematic areas along an escape route, where no safety sign can be recognized and thus no information on how to escape from the building is available. Since different building models and data from further studies using our simulation tool can be easily loaded into our evaluation tool, it can serve as basis for further VR experiments to determine factors influencing escape behavior and MRDs.

We found VR to be an effective way to carry out experiments in this domain and especially for simulating dangerous escape scenarios, VR offers a more controllable setting that can be incrementally extended for future use in lighting design and related fields.

9.2 Future Work

Based on the limitations and discussion addressed in Chapter 8, we want to motivate possible future work as well as approaches to further increase the realism of the escape scenario and the calibrated symptoms used in the user study.

9.2.1 Simulation Tool

The limitation of the simulation tool in terms of simulating viewing direction dependent symptoms could be solved by using an eye-tracking solution that can be fit inside the HMD. The rendering would need to be adapted accordingly. However, no affordable eye-tracker was available at the time of implementing the simulation tool. Nevertheless, it is expected that such a solution will be available in the near future and could serve a future version of the impairment simulation.

Moreover, the accurate simulation of a mixture of several symptoms would be of interest, since the combined effect might differ from the implications of the individual symptoms. Additional medical tests should be recreated in VR to calibrate these symptoms in order to provide a valid representation of the effects experienced by affected persons. For example, standardized tests for gray-level contrast exist. A spotted loss of visual acuity or a spotted, total loss of vision could be validated using a perimetry exam in VR.

The luminance perceived by the user of the simulation depends on the HMD's display. Therefore, the physically plausible lighting from HILITE can only correctly represent the relative differences between elements in the scene, not absolute values. The rendered scene's brightness would need to be adapted according to the mapping between a pixel's value and the luminance of the display for this value. Yet, this mapping could not be determined for the HTC Vive. HTC neither provides such values nor is the exact build of the display publicly known. The only information known in this context is that the display is an AMOLED display. Since an AMOLED display's maximum luminance per pixel changes with the overall displayed pixels' values, defining an exact mapping between the illuminance captured in the lightmap and the luminance of the display would be considerably hard. These values could, however, maybe be determined through an appropriate measuring scheme. Alternatively, a future version of the HTC Vive might offer more insight into these values or another HMD could be used.

9.2.2 User Study

The validity of the simulated impairments could not be tested with people actually affected by these symptoms for comparison. It is assumed to be difficult to find participants suffering only from one isolated symptom, for example only a loss of visual acuity. A more complex simulation featuring a more diverse set of symptoms might offer the opportunity for such a comparison.

It would be of interest to compare the visual acuity value determined in VR to a real-world visual acuity exam using the same test and settings. The outliers in the visual acuity values determined for the HTC Vive might be traced back to participants who already had a visual acuity less than that tested. This could be shown by a visual acuity exam in the real world. It can be expected that future generations of VR HMDs will provide higher resolution displays, overcoming the limitations on visual acuity in VR that could be observed for the HTC Vive. Such HMDs would enable tests using the full range of visual acuity levels, from normal vision to severe impairment. Testing for more levels of visual acuity would allow deducing a relationship between visual acuity and the achievable recognition distances. Such a general visual acuity factor for MRD computation could then be used in the evaluation tool.

The distractions resulting from other people escaping from a building, smoke or fire might influence the participants' ability to follow the escape route or increase the average recognition distance. It would therefore be of interest to use the effects and tools provided by UE4 to extend the simulation tool by these factors and incorporate them in a future user study.

9.2.3 Evaluation tool

To allow lighting designers to directly adapt their designs when they determine problematic areas using the evaluation tool, the placement of safety signs and the path describing the escape route should be editable in the tool. It might be necessary to integrate the evaluation tool more closely with HILITE or vice-versa. This would also offer the possibility to check further prescriptions by standards and regulations: For example, a certain illuminance is to be kept along the escape route. This condition could be checked using measurement planes in HILITE. Such an integration would also offer an access to the building information model (BIM) data available in HILITE. BIM data provides a consistent representation of additional data related to a building's 3D model [Ais86], such as the escape routes or exits. In addition, the exact lighting conditions could be considered for the visibility computation. As ISO 3864-1 [Int11] suggests, the MRD of a safety sign should be scaled depending on its brightness. Measurement planes in HILITE would allow determining this brightness and let it influence the visibility computation.

To provide a more accurate visibility computation that considers the from-region visibility for the safety signs, the current approximative solution could be replaced by a conservative, from-region visibility algorithm. If the required preprocessing, including the slicing of non-convex objects, can be incorporated into the evaluation tool, the approach by Durand et al. [DDTP00] presented in Section 3.2 would provide a sensible solution. By handling all required geometry processing inside the evaluation tool, lighting designers would still be free to use their building models directly, making their workflow less cumbersome. Also, using this approach, the evaluation tool could provide a conservative visibility estimation for general geometry objects.

Furthermore, the viewing direction and visual field of a person fleeing along the escape route could be incorporated in the visualizations of the evaluation tool. To achieve this, the visible area of a safety sign in an observer's visual field could be determined, weighted depending on its distance from the central vision field. This value could then be used as an estimation for the likelihood of the sign being observed from this point by an observer moving along the escape route. A visualization of this value would serve as an additional guidance to the lighting designer.

List of Figures

2.1	Classification of different types of emergency lighting. Adapted from EN 1838 [DIN13]	6
<u></u>	Different measures of brightness from the light source to to the ave of the	0
2.2	observer. Based on an illustration from [WF17].	7
2.3	A visualization of the concept of visual angle: Held at arm's length, the fist subtends a visual angle of 8-10°. Reprinted from [Tho11].	9
2.4	Density of the photoreceptive cells as a function of their position on the retina.	10
2.5	The Landolt ring as defined by ISO 8596: The diameter of the ring is d, the thickness of the ring and the gap are sized $d/5$. To use this symbol to measure a subject's visual acuity, the size of the gap should be such that it subtends a visual angle whose angular extent is reciprocal to the decimal acuity value to	10
	be tested. Reprinted from [Int09]	12
2.6	Most common eye disease in the U.S. population aged over 40. Adapted from the National Eye Institute (NEI) [NEI].	13
2.7	Artistic impressions of (a) Cataract, (b) Diabetic Retinopathy, (c) Glaucoma and (d) AMD. Adapted from NEI [NEI]	15
3.1	Visualization of the MRD proposed by EN 1838. Reprinted from ISO 3864-1 [Int11].	18
3.2	Visualization of the MRD proposed by ISO 3864-1. This also approximates the theoretical model presented by Xie [XFG ⁺ 07]. (1) vertical ring, diameter z_0 , (2) horizontal ring, diameter z_0 , (3) MRD for $\alpha = 0^\circ$, (4) MRD when viewed under angle $\alpha \neq 0^\circ$, $z_\alpha = z_0 \cdot \cos\alpha$ (5) safety sign. Reprinted from	
	ISO 3864-1 [Int11]	20
3.3	(left) Top view of the building model used by Almeida et al. (right) A crowd running against the direction indicated by the safety sign. Adapted from	
	$[AJF^+14]$	22
3.4	The view volume defined in MASSEgress. Point P is considered visible as it is inside the view volume and the ray tracing test's distance is equal to the	
	distance between the observer and P. Reprinted from [CPLL15]	23

3.5	(a) The model of a hypothetical supermarket in buildingEXODUS. (b) The VCA resulting from the eight safety signs in the supermarket model. Adapted from [Xie11].	24
3.6	 (a) The visibility is computed from 5 sampling points with shrunk occluders. (b) The individual umbrae of the actual, unshrunk occluders (dark) as compared to the fused umbra resulting from the conservative visibility computation using sampling and occluder shrinking. Reprinted from [COCSD03] 	26
3.7	Extended Projections. (a) The occludee (green) is projected to different positions of the projection plane, depending on the viewing position in the viewing cell. The union of all its projections is its extended projection. Similarly, the occluder's projection (gray) onto the projection plane changes with the viewing position. Its extended projection is the intersection of these projections. (b) The projections effectively fuse occluders A and B. Reprinted from [DDTP00].	27
3.8	The top row shows the results of the UE3 based simulation and the bottom row shows the XNA based simulation, presented by Lewis et al. in [LSB12] and [LBCM11]. (a) Glaucoma. (b) Cataract. Adapted from [LBCM11] and [LSB12].	29
3.9	Simulation of different age-related eye diseases as presented in (a) Velázquez et al. [VVR16], (b) Ates et al. [AFF15], (c) Väyrynen et al. [VCH16] and (d) Jin et al. [JAR05]. Reprinted from the respective papers	31
4.1	The effect of the simulated loss of visual acuity.	37
4.2	Simulation of age-related eye diseases. (a) The effect of the simulated cataract: The vision is fogged by a yellowish clouding of the lens, reducing visual acuity, perceivable contrast and tinting the perceived colors. (b) The effect of the simulated dry AMD: The central field of vision is darkened and reduced in terms of visual acuity	38
4.3	The building model rendered in UE4 as used in the user study. It was designed in pCon.planner and features physically plausible lightmaps rendered in HILITE.	39
4.4	The wheelchair simulator: (left) The office chair with the mounted motion controller used for the wheelchair simulator viewed from behind and (right) the virtual model of the wheelchair in the simulation viewed from above	41
5.1	Measuring the illuminance of a safety sign in HILITE: A measurement surface is placed on the safety sign, as can be seen in the false color visualization (bottom left). The measured values (top right) show an average illuminance of 30lx on this measurement surface	45

5.2	The calibration process: (top left) The black color of the Landolt rings implies
	that no input was recognized yet. The position of the participant's thumb,
	and thus input to be made, is indicated by a white sphere on the trackpad
	of the motion controller model. (top right) The participant has determined
	the orientation of three rings. Their green color does not indicate whether
	these inputs were correct or not. The next ring that requires an input is
	the forth from the left, as indicated by its black color. (bottom left) When
	determining the σ parameter of the Gaussian blur, the participant's vision
	becomes increasingly blurred for consecutive lines. Note the slightly blurred
	vision and the smaller size of the rings, compared with the previous examples.
	(bottom right) A detail view of a line of Landolt rings

5.3 The trackpad of the motion controller as seen by the user in VR. The green area in the middle accepts no inputs as mistakes would be too likely to happen with such small distances between different directions. For the experiments, the number of possible directions is reduced to four.

- 5.4 The determination of MRDs: (left) The participant starts at the end of the corridor and cannot yet recognize the information on the safety sign. (middle) After approaching further, the participant can recognize the direction the arrow on the sign is pointing to and reports this direction using the trackpad. If this input is correct, the next sample is presented. (right) The same situation as in the middle image but for the visual acuity level "strong blur", equivalent to a decimal acuity value of 0.125. It is not possible to recognize the information on the safety sign, hence the participant has to move closer. Overall, this results in a shorter MRD.
- 5.5 The escape scenario. (left) The participants start in a room with several obstacles and have to locate the escape route.(middle) The participants have to navigate through the grid-like part of the building by following the directions provided by the safety signs. The distance to the next sign varies, and the escape route itself changes between individual runs of the escape scenario. (right) After following the escape route to the end, the participants have to navigate through a large hall featuring three emergency exits.

5.6 The building used for the evacuation experiment as rendered in UE4, featuring the lightmaps imported from HILITE, viewed from above. Areas not accessible to the participants are crossed out. (A) points to the participants' starting point, (B) points to the three possible emergency exits.

5.7 Two of the escape routes used in the user study are highlighted in green. The two other paths are mirrored versions of these, albeit with the final exit at the left and right respectively. The safety signs are depicted as green rectangles. As can be seen, the length of the escape route segments and the number of segments of same length are equal for both (therefore all) paths, making the results from the different routes comparable.

47

46

50

51

52

53 101

6.1	The region from which a safety sign is recognizable is delimited by its MRD (green). We only consider this region for visibility computation, approximated by a hemicube (solid black). The dashed black lines indicate the view frusta used for rendering.	57
6.2	Sampling visibility at the center of a safety sign using ODSM to compute the umbra: The visible region and the umbra (gray) are separated by the solid black line (left). The safety sign should be fully visible from all points that lie in the visible region (middle). However, using only a single sample at the sign's center, only this point is guaranteed to be visible. Parts of the safety sign might be occluded anyways (right)	58
6.3	Sampling visibility at the safety sign's center and its corners. \ldots .	59
6.4	Visualization of the visibility regions inside the test building as rendered by the evaluation tool, without considering MRDs. Areas rendered blue are either occluded by geometry or outside the MRDs of the safety signs. At least one safety sign can be recognized from areas rendered green. (top) The horizontal slice and (bottom) the vertical slice	60
6.5	The GUI of the evaluation tool: The user can load building models and experiment data. The visibility model to be used can be selected as well as the parameters controlling slices. The visualization itself is darkened while the sidebar is opened	61
7.1	Overview of the simulation tool. The input handling and rendering are covered in this section. The game logic is discussed in Section 7.2.	64
7.2	From left to right: A touch on the trackpad is input as an x-y position on the trackpad with $x, y \in [-1, 1]$. The cartesian coordinates (x, y) of the input are converted to polar coordinates (r, α) , allowing easier mapping of input to selected direction. The texture visible to the user: Clicks inside the green area $(r < 0.5)$ are ignored. Clicks outside this area are mapped to the related direction using α , e.g., for 45° the direction "top-right" is assumed	66
7.3	Input mappings in UE4	66
7.4	The wheelchair movement. (a) Controller A, attached to the back of an office chair. The tracked yaw angle of this controller changes when the user rotates the office chair. (b) The user controls the movement speed using the trigger button of controller B. (c) The capsule and thus the origin of the tracking space as well as the user's view are moved into the forward direction of the wheelchair by this interaction. (d) The required component attachments in UE4	67
7.5	The weights of a Gaussian kernel: (a) The full kernel, (b) the separable kernel. Note that multiplying the separable kernels results in the original weights of the full kernel.	68

7.6	From a separable Gaussian blur kernel to the weights and sampling positions for a Gaussian blur using bilinear filtering: (a) The weights of the separable Gaussian blur kernel. (b) The center pixel is sampled twice for simpler convolution code. Since the kernel would no longer sum to 1 if we would just duplicate the center weight, its weight is halved. (c) Weights and sampling positions for the interpolation via bilinear filtering: The weights are summed up pair-wise and the sampling positions are adapted (see line 14 in Algorithm 7.1) such that, when sampling them in a bilinear-filtered manner, the texture values are weighted as with the original Gaussian blur	70
8.1	The impact of the HTC Vive's display on the visual acuity achievable in VR. None of the participants could achieve a visual acuity better than 0.4 decimal	01
8.2	acuity. This corresponds to a considerable vision impairment The boxplots show the recognition distances in centimeters, determined during the two runs of experiments with no blur (0.4 decimal acuity caused by the HTC Vive), weak blur (0.2 decimal acuity) and strong blur (0.125 decimal acuity). The red curves shows the MRDs computed from the recognition distances of both runs. For comparison, the upper gray curves shows the MRDs as computed according to the suggestions in ISO 3864-1. They depend on the observation angle by a factor of $cos(\alpha)$ and are scaled by the decimal acuity	81
8.3	value	82
8.4	The recognition distances per safety sign in the evacuation scenario experiment. The semi-transparent lines show the recognition distances per user. The dashed lines illustrate the farthest distance (top), within which the sign becomes theoretically visible, and the nearest distance (bottom), within which the sign is guaranteed to be theoretically visible. The gray lines show the MRDs as suggested by EN 1838 (top) and ISO 3864-1 (bottom). Finally, the MRD suggested by our experiments is shown in red and the mean recognition distances of all participants is shown in green.	86
8.5	Visualization of the visibility regions inside the test building as rendered by the evaluation tool. The MRDs are generated by the model which is suggested by EN 1838. The model is not able to represent the reduction of recognizability of signs following from observation under a flat angle. This	07
8.6	Visualization of the visibility regions inside the test building as rendered by the evaluation tool. The MRDs are generated by the model which is suggested by ISO 3864-1 and Xie et al. It incorporates the observation angle and the corresponding visualization reveals a gap in the escape-route coverage (blue)	88
	corresponding insumination reveals a Sap in the obsape route coverage (blue).	00

8.7	Visualization of the data from the experiments inside the test building: Green	
	glyps show the direction that was reported for the safety sign, gray glyphs show	
	the position of the participant in 1sec intervals and the gray line illustrates	
	the chronological order of these glyphs.	89

List of Tables

2.1	Reference values for examples of light sources and lighting settings, given in the presented measures. The values can be found in Witting [Wit14]	9
8.1	The mean \bar{x} and standard deviation σ of the recognition distances over all	
	observations per run and safety sign configuration.	83
8.2	The MRDs determined in the user study, valid for 85% of the participants. For	
	comparison, the MRDs resulting from the angle-dependent model $z_0 \cdot h \cdot cos(\alpha)$	
	suggested by ISO 3864-1 as well as the MRDs resulting from a further scaling	
	by the decimal acuity value, $z_0 \cdot h \cdot cos(\alpha) \cdot dA$. The visual acuity levels	
	"weak blur" and "strong blur" relate to decimal acuity values of 0.2 and 0.125,	
	respectively.	83

List of Algorithms

7.1	Computation of the weights and sampling positions for a Gaussian blur	
	using bilinear filtering	71
7.2	Gaussian blur filter applied in x-direction	72

Bibliography

- [AFF15] Halim Cagri Ates, Alexander Fiannaca, and Eelke Folmer. Immersive simulation of visual impairments using a wearable see-through display. In Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction, pages 225–228. ACM, 2015.
- [Ais86] Robert Aish. Building modelling: the key to integrated construction cad.
 In CIB 5th International Symposium on the Use of Computers for Environmental Engineering Related to Buildings, volume 5, pages 7–9, 1986.
- [AJF⁺14] João Emílio Almeida, João Tiago Pinheiro Neto Jacob, Brígida Mónica Faria, Rosaldo JF Rossetti, and António Leça Coelho. Serious games for the elicitation of way-finding behaviours in emergency situations. In Information Systems and Technologies (CISTI), 2014 9th Iberian Conference on, pages 1–7. IEEE, 2014.
- [BW03] Jiří Bittner and Peter Wonka. Visibility in computer graphics. *Environment* and Planning B: Planning and Design, 30(5):729–756, September 2003.
- [COCSD03] Daniel Cohen-Or, Yiorgos L Chrysanthou, Claudio T. Silva, and Frédo Durand. A survey of visibility for walkthrough applications. *IEEE Transactions* on Visualization and Computer Graphics, 9(3):412–431, 2003.
- [CPLL15] Mei Ling Chu, Paolo Parigi, Jean-Claude Latombe, and Kincho H Law. Simulating effects of signage, groups, and crowds on emergent evacuation patterns. Ai & Society, 30(4):493–507, 2015.
- [CRN16] Giovanni Cosma, Enrico Ronchi, and Daniel Nilsson. Way-finding lighting systems for rail tunnel evacuation: a virtual reality experiment with oculus rift[®]. Journal of Transportation Safety & Security, 8(sup1):101–117, 2016.
- [DDTP00] Frédo Durand, George Drettakis, Joëlle Thollot, and Claude Puech. Conservative visibility preprocessing using extended projections. In Proceedings of the 27th annual conference on Computer graphics and interactive techniques, pages 239–248. ACM Press/Addison-Wesley Publishing Co., 2000.

- [DIN13] DIN German Institute for Standardization. Lighting applications Emergency lighting; EN 1838:2013. Standard, DIN German Institute for Standardization, Berlin, Germany, October 2013.
- [Eas] EasternGraphics GmbH. pCon.planner. http://pcon-planner.com/ en/. Accessed: 2017-09-04.
- [Epi] Epic Games, Inc. Unreal Engine 4. https://www.unrealengine.com/. Accessed: 2017-09-04.
- [Eur12] European Committee for Standardization. Graphical symbols Safety colours and safety signs - Registered safety signs (ISO 7010:2011). Standard, European Committee for Standardization, Brussels, Belgium, July 2012.
- [HVD06] MA Hogervorst and WJM Van Damme. Visualizing visual impairments. Gerontechnology, 5(4):208–221, 2006.
- [HvD08] Maarten A Hogervorst and Wim JM van Damme. Visualizing the limits of low vision in detecting natural image features. Optometry & Vision Science, 85(10):E951–E962, 2008.
- [Int09] International Organization for Standardization. ISO 8596:2009 Ophthalmic optics – Visual acuity testing – Standard optotype and its presentation. Standard, International Organization for Standardization, Geneva, CH, July 2009.
- [Int11] International Organization for Standardization. ISO 3864-1:2011 Graphical symbols – Safety colours and safety signs – Part 1: Design principles for safety signs and safety markings. Standard, International Organization for Standardization, Geneva, CH, April 2011.
- [JAR05] Bei Jin, Zhuming Ai, and Mary Rasmussen. Simulation of eye disease in virtual reality. In Engineering in Medicine and Biology Society, 2005. IEEE-EMBS 2005. 27th Annual International Conference of the, pages 5128–5131. IEEE, 2005.
- [LBCM11] James Lewis, David Brown, Wayne Cranton, and Robert Mason. Simulating visual impairments using the unreal engine 3 game engine. In Serious Games and Applications for Health (SeGAH), 2011 IEEE 1st International Conference on, pages 1–8. IEEE, 2011.
- [LSB12] James Lewis, L Shires, and DJ Brown. Development of a visual impairment simulator using the microsoft xna framework. In Proc. 9th Intl Conf. Disability, Virtual Reality & Associated Technologies, Laval, France, 2012.
- [NEI] NEI Office of Science Communications, Public Liaison, and Education. NEI Photos and Images. https://nei.nih.gov/photo/. Accessed: 2017-09-15.

- [OB05] Österreichischer Bundesfeuerwehrverband. TRVB E 102 Fluchtweg-Orientierungsbeleuchtung und bodennahe Sicherheitsleitsysteme. Technical Guideline, Österreichischer Bundesfeuerwehrverband, 2005.
- [Pal99] Stephen E. Palmer. Vision science. A Bradford book. MIT Press, Cambridge, Mass. [u.a.], 3. print. edition, 1999.
- [Pan06] Xiaoshan Pan. Computational modeling of human and social behaviors for emergency egress analysis. PhD thesis, Stanford University, 2006.
- [SF02] John A Shoemaker and David S Friedman. Vision problems in the US: prevalence of adult vision impairment and age-related eye disease in America. National Eye Institute, 2002.
- [Str] Filip Strugar. An investigation of fast real-time GPU-based image blur algorithms. https://software.intel.com/enus/blogs/2014/07/15/an-investigation-of-fast-realtime-gpu-based-image-blur-algorithms. Accessed: 2017-09-24.
- [The] The assimp team. Open Asset Import Library. http://assimp. sourceforge.net/. Accessed: 2017-12-08.
- [Tho11] William B. Thompson. Visual perception for computer graphics perspective. An A K Peters book. CRC Press, Boca Raton, Fla. [u.a.], 2011.
- [Uni] Unity Technologies. Unity. https://unity3d.com/. Accessed: 2017-09-04.
- [VCH16] Jani Väyrynen, Ashley Colley, and Jonna Häkkilä. Head mounted display design tool for simulating visual disabilities. In Proceedings of the 15th International Conference on Mobile and Ubiquitous Multimedia, pages 69–73. ACM, 2016.
- [VRVa] VRVis Research Center. Aardvark. https://www.vrvis.at/ research/projects/aardvark/. Accessed: 2017-09-04.
- [VRVb] VRVis Research Center. HILITE. https://www.vrvis.at/research/ projects/hilite/. Accessed: 2017-02-13.
- [VVR16] Ramiro Velázquez, Jorge Varona, and Pedro Rodrigo. Computer-based system for simulating visual impairments. *IETE Journal of Research*, 62(6):833– 841, 2016.
- [WCCC10] Joanne Wood, Alex Chaparro, Trent Carberry, and Byoung Sun Chu. Effect of simulated visual impairment on nighttime driving performance. *Optometry* and vision science, 87(6):379–386, 2010.

- [WF17] Bruno Weis and Hans Finke. *Not- und Sicherheitsbeleuchtung.* de-Fachwissen. Hüthig, München Heidelberg, 2., völlig neu bearbeitete auflage edition, 2017.
- [Wit14] Walter Witting. Licht und physik. lichttechnische grundlagen. In *Licht. Sehen. Gestalten*, chapter 2, pages 49–154. Birkhäuser, Basel, 2014.
- [WWS00] Peter Wonka, Michael Wimmer, and Dieter Schmalstieg. Visibility preprocessing with occluder fusion for urban walkthroughs. In *Rendering Techniques 2000*, pages 71–82. Springer, 2000.
- [XFG⁺07] Hui Xie, Lazaros Filippidis, Steven Gwynne, Edwin R Galea, Darren Blackshields, and Peter J Lawrence. Signage legibility distances as a function of observation angle. *Journal of fire protection engineering*, 17(1):41–64, 2007.
- [Xie11] Hui Xie. Investigation into the interaction of people with signage systems and its implementation within evacuation models. PhD thesis, University of Greenwich, 2011.
- [ZB10] Michelle Zagar and Scott Baggarly. Low vision simulator goggles in pharmacy education. *American journal of pharmaceutical education*, 74(5):83, 2010.
- [ZLG16] Zumtobel Lighting GmbH. Normen und Planungshilfen für die Sicherheitsbeleuchtung. Brochure, 2016.