Real-time statistical simulation of dynamic laser speckle

Mátyás Komáromi

School of Science

Thesis submitted for examination for the degree of Master of Science in Technology.

Espoo 28.02.2022

Supervisor

Prof. Jaakko Lehtinen
Prof. Demi Libertario

Advisor

Dr. Jorma Palmén

Aalto University
School of Science

Throughout the writing of this thesis I have received a great deal of support and assistance.
From Ladimo Oy, I want to thank my ever changing advisors, Sami Ruuskanen, Tuomas Savolainen, and Jorma Palmén.
My friends, for believing that I can finish this thesis, but especially Richárd Bagi, who spared no brain cells helping me, and Gaia Businaro, for the final stretches.

Aalto University
School of Science

Aalto University, P.O. BOX 11000, 00076 AALTO
www.aalto.fi
Abstract of the master's thesis

Author Mátyás Komáromi

Title Real-time statistical simulation of dynamic laser speckle

Degree programme ICT Innovation

Major Visual Computing and Communication **Code of major** SCI3102

Supervisor Prof. Jaakko Lehtinen
Prof. Demi Libertario

Advisor Dr. Jorma Palmén

Date 28.02.2022 **Number of pages** 70 **Language** English

Abstract

Lasers have several applications in the industry, such as cutting, engraving, and drilling. A specific use of lasers is taking distance measurements, by shining beams of light at objects and observing the hit points with an infrared camera. However, the depth measurements are inevitably inaccurate. This is mainly due to the manufacturing and design errors in beam splitters and dynamic speckle. The goal of this thesis is not to reduce the contribution of these effects, but to introduce them into a computer simulation, to make the virtual model as close as possible to reality. Studying the works of J. W. Goodman, Donald D. Duncan, and others in the field of Fourier Optics gave a solid theoretical foundation of the beam splitter and dynamic speckle. To tailor the general theory to this specific case, various physical experiments were carried out. Based on the theory and the experiment results, a way to extend an already existing physically based rendering engine was proposed. In conclusion, this extension produces similar results in the simulation to what is observable in real life. This is achieved with a small computational overhead on a modern graphics processor. Due to these properties, the technique can be used for more robust testing of depth estimation and reconstruction algorithms. Moreover, it also raises the quality of machine learning data that can be collected in large volumes from a computer simulation of this setup, leading to better downstream performance.

Keywords Speckle pattern, Real-time computer graphics, Physically based rendering, Fourier optics

Table of Contents

Abstract **3**

Table of Contents **4**

Symbols and abbreviations **6**

1 Introduction **8**
1.1 Company Case . . . 9
1.2 Physical effects affecting image formation . . . 10
1.3 Related Works . . . 12
1.3.1 Joseph W. Goodman . . . 13
1.3.2 Simulation of speckle . . . 13
1.3.3 3D imaging . . . 14
1.3.4 Other applications . . . 14

2 Background **15**
2.1 From Maxwell's equations to the Huygens-Fresnel principle . . . 15
2.2 Fresnel approximation of the Huygens-Fresnel principle . . . 18
2.3 Speckle in imaging systems . . . 22
2.4 Surface roughness . . . 26
2.5 Speckle statistics . . . 28
2.6 Beam splitter DOE . . . 31
2.7 Approximation constraints . . . 33
2.8 Laser beam profile . . . 34
2.9 Polarization of light . . . 34

3 Methodology **36**
3.1 Tools . . . 36
3.2 Experiments . . . 37
3.2.1 Image processing model . . . 37
3.2.2 General setup . . . 38
3.2.3 Stationary particleboard . . . 39
3.2.4 Displacement along the X-axis . . . 41
3.2.5 Displacement along the Z-axis . . . 42
3.2.6 Rotation about the Y-axis . . . 43
3.2.7 Airflow . . . 43
3.2.8 Human experiments . . . 44
3.2.9 Recovering the beam profile . . . 45
3.2.10 Spatial consistency . . . 47
3.3 Proposed Algorithm . . . 47

4 Results **50**
4.1 Performance . . . 50
4.2 Evaluation . . . 52

5 Conclusion **56**
5.1 Future Work 56

References **58**

A Varying F number in simulation **65**

B Varying F number in measurements **66**

C Measurement Setups **67**

D Dot patches with speckle (simulated) **68**

E Pictures with and without speckle from the simulator **69**

Symbols and abbreviations

Symbols

Symbol	Description
$\vec{A}$	A vector quantity with typically 2 or 3 components.
$\mathbf{i}, \mathbf{j}, \mathbf{k}$	Unit vectors in 3D Cartesian space
∇	The nabla operator. A vector differential operator that can express gradient, divergence, or curl, depending on the usage.
π	The circumference of a circle with radius=0.5.
$\mathbb{E}(X)$	Expected value of the random variable X.
$Im(a)$	The imaginary part of complex number 'a'.
$Re(a)$	The real part of complex number 'a'.
$\|v\|$	Length of vector 'v'.
mod	The fractional modulo operator
λ	Denotes the wavelength of the light field at hand.
$\mathcal{F}_{\alpha,S}$	Scaled Fourier transformation.
$\mathcal{F}$	The traditional Fourier transformation.
$\mathcal{O}$	Big O notation. It characterizes the complexity of a function with an upper bound.
Ω	Similarly to the O notation, it characterizes the complexity of a function with a lower bound.
$\gg$	Significantly larger. Typically one or two orders of magnitude of difference.

Abbreviations

Abbreviation	Expansion
FT	Fourier Transform.
IFT	Inverse Fourier Transform.
FFT	Fast Fourier Transform.
IR	Infrared light.
CLT	Central Limit Theorem.
GS	Gerchberg-Saxton algorithm.
ACF	Autocorrelation Function.
RMS	Root Mean Square.
DOE	Diffractive Optical Element.
EM	Electromagnetic (wave).
SNR	Signal to Noise Ratio.
DBS	Direct Binray Search.
ML	Machine Learning.
CCD	Charge-Coupled Device.
PBR	Physically Based Rendering.
TOF	Time Of Flight.

1 Introduction

Since its invention, the laser has found many applications. These range from precision cutting of components through engraving names into trophies to laser eye surgeries. Among these, an industrial application of lasers is conducting depth measurements.

Several depth measurement devices that use lasers are available on the market. They are based on a range of different technologies. A class of approaches, called laser-based time-of-flight (TOF) [1] consists of measuring the amount of time a laser pulse takes to arrive back at the sending position. See Figure 1 for an illustration. Based on the elapsed time, a depth estimation can be calculated. TOF cameras are used in automotive applications, gaming, machine vision, and even in several modern smartphones.

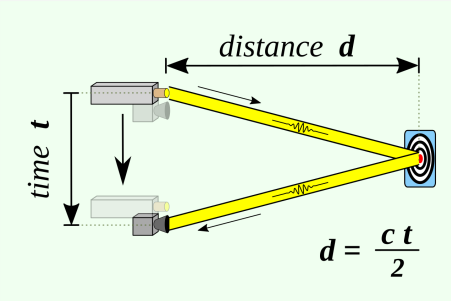

Figure 1: Illustration of the principle of time of flight distance measurement devices. Image created by R. Craig, shared under CC BY-SA 4.0.

Another approach is employed by the famous Microsoft Kinect [2] and its competitors, the ASUS Xtion, and FaceID by Apple. Both work by projecting a fixed laser pattern (see Figure 2) and observing the scene with an infrared (IR) camera. The depth estimation is done based on the deformation of the observed pattern. The now discontinued project Google Tango [3] was aiming at bringing the same technology to Android phones and tablets.

Figure 2: Semi-random laser pattern projected by the Microsoft Kinect. Photograph by J. Flatley, from his Engadget article.

Ladimo Oy is developing a 3D camera, similar in principle to the Kinect. Instead of a semi-random pattern, the company is using a high number of independent laser beams, going in specific directions. Each beam contributes a single depth

measurement, and therefore, this technology enables high precision and repeatability, while maintaining a low computational cost.

In real life, when such a distance measurement is conducted, the captured IR image is always distorted. This deviation is due to several physical effects. In this work we will analyse these effects and isolate the most important ones.

Ladimo Oy is interested in creating a virtual model of the measurement device for rapid testing and machine learning (ML) data collection purposes. The current version of this virtual model, however, is an overly simplified model of reality and does not account for the aforementioned deviations. Therefore, an approach needs to be designed to investigate the sensitivity of the reconstruction algorithms to these errors. It is important to note that a technique might work perfectly on ideal simulated data but fail miserably in the presence of these effects.

Therefore, this thesis hypothesizes that the physical effects that influence the captured IR image can be adapted into an already existing rendering engine. Furthermore, that the simulation can be done in real-time.

To support this hypothesis, an approach to extend the existing simulation with these effects is devised. This approach is based on the theoretical work in the field of Fourier Optics done by J. W. Goodman, Donald D. Duncan, and others. The applied simplifications enable real time computing of the effects. These simplifications are justified through physical experiments.

Finally, to place the simulation results in contrast with theory and practice, the measurements from the three were compared in various setups. We find that the simulation matches reality and theory on useful key aspects. Due to this and the low computational cost this extension enables a more robust testing of depth estimation algorithms and collection of higher quality ML data .

Note that details were omitted in cases of company secrets in this work.

The rest of this work is structured as follows. Sections 1.1 and 1.2 introduce the setup we are examining and the physical limitations that apply to this setup. Section 1.3 dives into the contribution of other authors to the same or similar concepts in a different setting. Sections from 2.1 to 2.5 introduce and derive the theoretical - mathematical and physical - background of speckle, essential factors in our simulation. In Section 2.6, we look into what deviations are presented by the optical element that splits the laser beam into many other beams. Sections under 3 elaborate on the experiments conducted, their conclusions, and the tools used in carrying them out. In Section 3.3, we present the actual extension of the rendering engine in. We show the improvements and costs this modification brings in sections 4.1 and 4.2. We examine how the simulation compares to reality in Section 4.2. In 5.1, we discuss the shortcomings and possible future extensions of our approach.

1.1 Company Case

To better understand how the physical setup works, the depth measurement technology created by Ladimo Oy will be discussed in detail in this section.

The software developed by Ladimo Oy estimates true scale depth values using a monochromatic light source and an IR camera in the setup described below (Fig 3).

The monochromatic light source (laser) has a narrow wavelength span around 940 nm. It emits a single beam of light with a high-quality Gaussian profile, which is then split into thousands of beams by a small Diffractive Optical Element (DOE). These beams then propagate through the air and hit objects in a scene; in our case, we consider the setting to be an ordinary home or office environment, and as such, the typical bodies hit by the light are humans, office desks, keyboards, etc. The distance of these objects from the laser source is between 70 cm and 2 m. The scattered light is then reflected back towards the emitter; it propagates through the air again and gets captured by an IR camera. The image of the beams shows up as spatial Gaussians, which the company terms as dots. The next step is to estimate the center of dots with subpixel accuracy and use these results in further geometrical calculations to find the true-scale 3D position of the centers. The accuracy of the center estimation is key to the company since it directly affects the quality of the depth estimation. Therefore, it is of great interest to understand the physical processes that influence the process of image formation, which is the topic of this thesis.

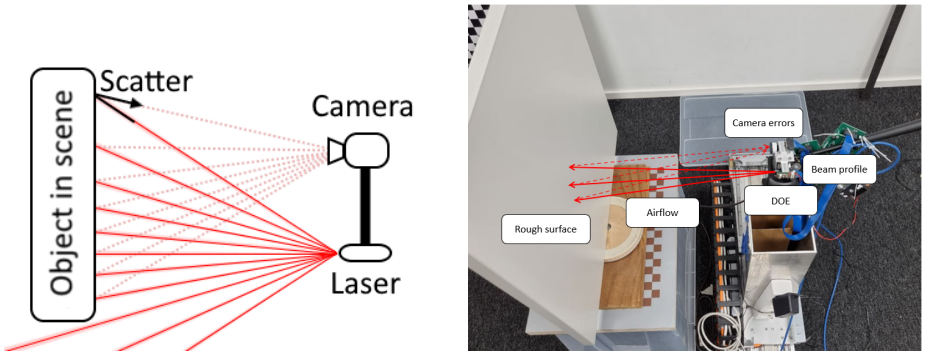

Figure 3: Schematic drawing of the setup and the physical phenomena that influence the final result, illustrated on a real setup. The laser emits beams which are then scattered by the object(s) in the scene and the reflected light is captured by an IR camera.

Ideally, we would see small Gaussian dots on the captured image corresponding to a single laser beam. However, we observe fluctuations and a distorted image in practice, resulting in alternating high and low-intensity regions on the image. This difference is illustrated in Figure 4. The deviation from perfect imaging results from multiple physical properties of the above-described scenario. We discuss these properties deeper in Section 1.2.

1.2 Physical effects affecting image formation

During the journey from the laser emitter to the electron wells of the IR camera, the light in question undergoes several changes that affect the registered values in said camera (see Figure 3). However, not all of these effects contribute significantly to the final results. Below we attempt to give an as complete as possible listing of these effects and their significance. The listing is ordered chronologically.

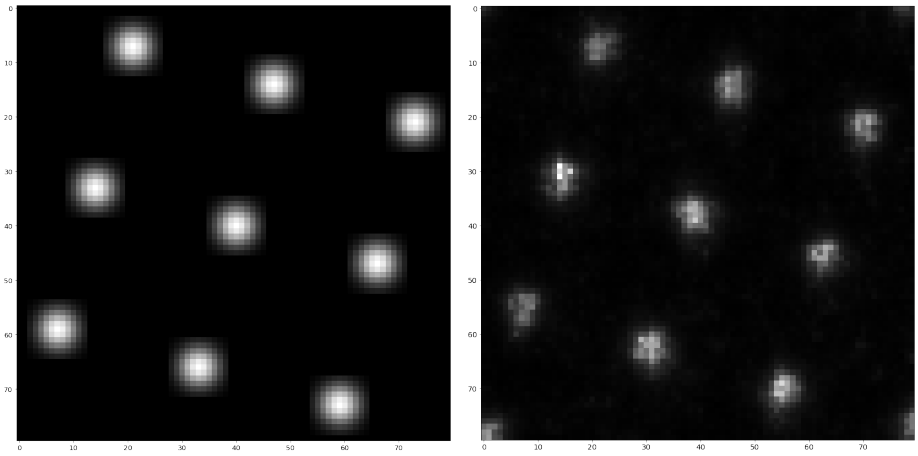

Figure 4: Perfect, spatially Gaussian dots and measured realistic dots. To the left: Ideally shaped dots, no speckle. To the right: Gaussian dots distorted mainly by the dynamic speckle.

1. The beam profile of the laser emitter deviates from the perfect Gaussian. This effect is important since the mean of a skewed distribution will naturally be different from an ideal one and therefore influence the estimation. The image postprocessing that follows the image acquisition can take this defect into account. The company is working with high-quality lasers; therefore, we will not explore this problem in-depth. See Section 2.8 for a more complete treatment of this topic.

2. No practical monochromatic light source can be truly monochromatic. Every laser emitter has a characteristic frequency profile, which describes what wavelengths are emitted. When the beam splitter is hit with a different frequency than it was designed for, it projects the outgoing beams into different directions than intended. See Figure 5 and the first three chapters of Diffraction Physics [4]. In our case, the range of emitted wavelengths is around 1µm wide (see Figure 18), which introduces only an insignificant deviation. Therefore, the simulation of this property is only considered as possible future work.

3. We require the output to be composed of around 5303 beams in specific directions for our beam splitter. Our older design used to employ 2509 beams. Designing such a laser beam splitter is a challenging task, and the geometric calculations that follow the image processing depend on the directions being precise. Imperfections in the design and manufacturing lead to ghost dots appearing between useful dots, some of the light being scattered between dots, a significant amount of energy is undiverted, and the shape of dots being deformed (see Figure 16). For further details, see Section 2.6.

4. While the laser light is propagating through the air, due to the movement and nonuniform temperature of the medium, different parts of the light field propagate with different speeds. This results in phase differences that lead to the

appearance of an interference pattern. Simulating this effect is computationally expensive [5]. It is insignificant effect in our general use case compared to the cost its simulation bears. See Section 3.2.7 for the experiments conducted on this topic.

5. After hitting the first object in the scene, the laser light might bounce into the direction of another object and thus perform multiple bounces before finally heading towards the camera. This effect is modeled without considering the wave nature of light, which speeds up computations significantly. This treatment makes sense in our case since the finite aperture of the camera introduces dynamic speckle.

6. When light bounces back from an object with a sufficiently rough surface, the phase of that light field is replaced by a random phase. It results in an interference pattern after propagation (see Section 2.4). Simulating the exact way this happens would require a complete map of the microstructure of that object, which is computationally too taxing in most cases as well as in ours. Our approach, therefore, is of statistical origins and is concerned with the statistical behavior of such phenomena. Combined with the finite aperture of the camera this phase offset causes what we refer to as *dynamic speckle* (see Figure 4 and [6, 7, 8, 9, 10]). Dynamic speckle is the single largest contributor to the deviation from perfect Gaussian dots on our image and thus the main topic of this thesis.

7. Finally, while capturing the image, the IR camera introduces additional effects, such as noises and distortions, typical to such cameras. We model the effects of the dark current, bayering, geometric distortion, and hot pixels. These effects can also be temperature-dependent, but they can be avoided or compensated for during image acquisition in our experiments. The simulation already takes some of these effects into account. Expanding it with the rest is considered outside of the scope of this work, but it can be considered in the future. For a more detailed description of the used model, we refer to Section 3.2.1.

Other factors of the environment can also affect the performance of the system. The temperature in which the laser resides can cause it to shift frequency. The natural light coming from the Sun also contains frequencies overlapping with that of our laser. The material of objects in the scene might not be perfectly diffuse or might exhibit subsurface scattering.

1.3 Related Works

Our investigation on the phenomena influencing the final IR image in the 3D measurement setup quickly showed that the two most significant factors are dynamic speckle and the beam splitter. We focus our attention on these two effects, especially on speckle. This section covers the materials we have used, built upon, or inspired us.

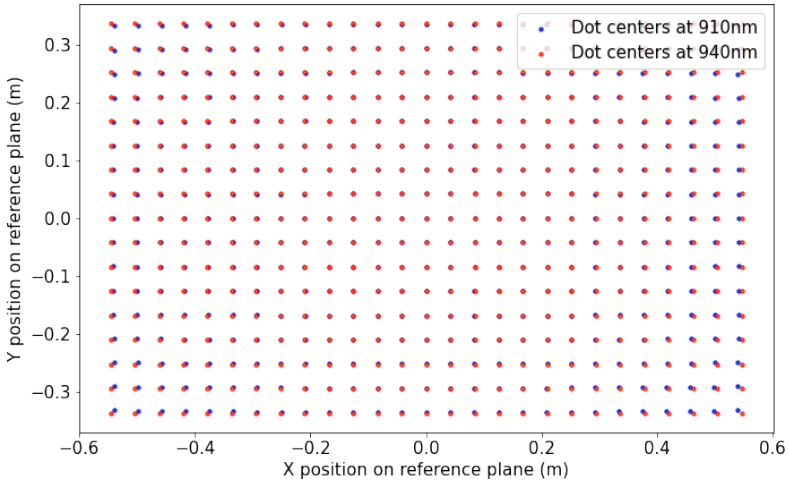

Figure 5: Shift of dot centers due to deviating laser frequency. The DOE was designed for 940 nm (red dots) but used with 910 nm (blue dots). The shift is more apparent for dots farther from the origin.

1.3.1 Joseph W. Goodman

The theoretical foundation of the diffraction-based beam splitter and dynamic speckle is Fourier optics. In his books, *Fourier Optics* [10], *Statistical Optics* [11], and in his 1975 paper [8] and 1976 article [7], Joseph W. Goodman introduced the tools and methodology of analysing the phenomenon of speckle. He qualitatively and quantitatively described why and how speckle develops. The books go way beyond our application of infrared imaging and discuss, among many topics, imaging in general, optical data processing, holography, noise reduction techniques, and imaging with partially coherent light. The statistical model of speckle, established in his articles based on the model of a random walk (also called drunkard's walk) became the standard statistical model of speckle.

1.3.2 Simulation of speckle

Donald D. Duncan, Sean J. Kirkpatrick [12] created a comprehensive survey of algorithms for speckle simulation. They described the generation of objective speckle on static and dynamic targets alike. Dynamic targets exhibiting small scale movements were considered, exhibiting the well-known effect of speckle boiling.

Victor Nascov, Cornel Samoilă, and Doru Ursuţiu [13] created fast computation algorithms for speckle pattern simulation. The focus of their work was on tilted and rotated surfaces.

Lipei Song et al. [14] developed a technique to generate pairs of speckle images with predefined correlation maps. Originally we wanted to use this technique to decorrelate the generated speckle pattern slowly as the camera moves. Although this turned out to be unnecessary in our case, the work gave an insight into how speckle loses correlation.

1.3.3 3D imaging

L. G. Shirley and Gregory R. Hallerman [15] described a technique to conduct noncontact contour measurement of 10 µm accuracy. What is even more impressive is that the range resolution of their approach does not degrade as the distance to the object increases. However, this technique requires a reference point to be added to the scene. This highly reflective small object causes the autocorrelation function of the speckle image to include the depth map. It is only a matter of extracting the right part of the autocorrelation image from then on. This work aided us greatly in understanding how the speckle is formed and what causes it.

1.3.4 Other applications

Speckle shows up in several fields, and countless applications have been discovered already. In the following, we discuss works that influenced our work, inspired us, or connected to what we worked on.

Lennart Bargsten et al. [16] integrated speckle into a neural network. They added a speckle layer into a generative adversarial network in order to generate more realistic training data for medical applications. Their goal was to augment ultrasound, optical coherence tomography, and other data sets. Their work inspired us to consider using ML to generate speckle in our case. The conventional method provides comparable, if not higher performance, and we can be certain that it is performing what we expect it to do.

YiChang Shih et al. [17] raises a completely different application. Instead of medical imagining or 3D imaging, they worked on forensics. Their work proposes an approach to detect if a surface has been touched between two imaging steps and, if so, where. All this without physically touching the object. The key is the change in the micro-surface of the object. Taking IR images, the laser speckle will encode information about the micro-surface. Comparing the changed speckle reveals the tampering. This work shows excellently just how many different applications speckle can have.

As mentioned earlier, speckle appears in medical applications often. While measuring the flow of bodily or other fluids such as blood (See Dynamic Imaging of Cerebral Blood Flow Using Laser Speckle [18]), the moving media blurs the resulting speckle and introduces a pattern with smaller contrast. Imaging parts of the human body is fundamental in tomography. Since highly coherent waves are used, speckle also shows up in these applications (see Speckle in Optical Coherence Tomography [19]). Other applications include continuous deformation tracking [20], keeping track of tomato ripening [21], detecting heat damage on plants [22], measuring soldering quality [23], ultrasonic imaging [24], stellar imaging [25], and plant water status monitoring [26].

2 Background

The theoretical foundation of dynamic speckle and, by the extension of Fourier Optics, the beam splitter are analysed in this section. This will be mainly achieved by following the work of Joseph W. Goodman in his books Introduction to Fourier Optics [10], Statistical Optics [11], and in his articles [8, 7] about speckle. To accommodate modern mathematical standards the notation will be occasionally slightly changed. In each section, the meaning of the acquired results and their relation to this specific industrial application is considered. For additional discussion on Fourier optics that is outside of the scope of this application, the reader is highly advised to refer to the aforementioned books and articles.

In Section 2.1, the kind of materials and scales, on which the following theory is applicable are established. It is done starting from a fundamental theory, the Maxwell equations and developed into at a geometrical expression, the Huygens-Fresnel principle. This expression is then reduced using the Fresnel and Fraunhofer approximation in Section 2.2. This section simplifies the previous geometric expression, granted some assumptions about the scale of sizes involved is large enough. Here, the first expression that could be put into practice is reached. Further discussion about the specifics of the imaging setup in Section 2.3 brings the topic to the relation between the light field at the object plane and the detector plane of a camera. To let speckle come into play another element is taken under consideration in Section 2.4, the importance of the roughness of the observed objects.

Based on the above theoretical foundation, some statistics of speckle is discussed in Section 2.5. The results shown here are later used to validate the results of the rendering algorithm. After the theoretical treatment of speckle, the role of the beam profile follows. The beam profile directly influences the depth estimations, since this is the shape that dots assume on average. In Section 2.8 it is shown that the beam profile can be modelled as close to ideal. The general equations in Section 2.3 only work on polarized light, but the setup of the company is working with fully unpolarized light. To bridge this gap, Section 2.9 discusses how the polarization of the emitted light can be handled.

2.1 From Maxwell's equations to the Huygens-Fresnel principle

In classical optics, light is fundamentally treated as an electromagnetic (EM) field, also referred to as a wave field or propagating wave [27]. EM fields are composed of two intertwined complex vector fields, the electric field and the magnetic field, denoted $\vec{\mathcal{E}} = (\mathcal{E}_x, \mathcal{E}_y, \mathcal{E}_z)$ and $\vec{\mathcal{H}} = (\mathcal{H}_x, \mathcal{H}_y, \mathcal{H}_z)$, respectively. These two fields are a function of both position and time; their value is always a three-dimensional vector. The evolution of these two fields over time can be described by a system of partial

differential equations known as Maxwell's equations [28]:

$$\begin{aligned} \nabla \times \vec{\mathcal{E}} &= -\mu \frac{\partial \vec{\mathcal{H}}}{\partial t} \\ \nabla \times \vec{\mathcal{H}} &= \epsilon \frac{\partial \vec{\mathcal{E}}}{\partial t} \\ \nabla \cdot \epsilon \vec{\mathcal{E}} &= 0 \\ \nabla \cdot \mu \vec{\mathcal{H}} &= 0, \end{aligned} \tag{1}$$

where $\nabla = (\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z})$ is the nabla operator. As for operations with nabla, we recognize $\nabla\times$ as the circulation (curl) and $\nabla\cdot$ as the divergence (div) when applied to a vector field. Note that the above two fields are in fact functions of position $\mathbf{P}$ and time t. μ and ϵ are characteristic values of the medium in which the wave is propagating in, referred to as permeability and permittivity, respectively. μ and ϵ can also be a function of position and time.

When an EM wave is propagating through a simple medium, we can use some attributes of this medium to simplify the above equations. These are

- Linearity: The medium satisfies the linearity constraints. See Chapter 2 in Introduction to Fourier Optics for more on linearity.
- Isotropy: μ and ϵ are independent of the orientation of $\mathcal{E}$ and $\mathcal{H}$.
- Homogeneous: ϵ is constant with respect to both position and time.
- Nondispersivity: ϵ is independent of the wavelength of the propagating wave.
- Nonmagnetic: μ is constant with respect to both position and time.

Assuming these properties means that Maxwell's equations can be expressed in the form of six independent equations. These equations are formed by substituting u in

$$\nabla^2 u - \frac{n^2}{c^2} \frac{\partial^2 u}{\partial t^2} = 0 \tag{2}$$

with $\mathcal{E}_x$,$\mathcal{E}_y$,$\mathcal{E}_z$,$\mathcal{H}_x$,$\mathcal{H}_y$, and $\mathcal{H}_z$, respectively. Here, ∇^2 is the Laplacian operator, u is a complex function of both position and time, n is the medium's refractive index, and c is the propagation speed in vacuum. These assumptions hold well in the case of light propagating in air and also in lenses. However, the rim of an aperture or the edge of a lens creates boundaries. We cannot assume the above properties to hold close to these boundaries. This is because in that case, the different components of $\vec{\mathcal{E}}$ and $\vec{\mathcal{H}}$ get coupled - they behave in an interdependent way. Luckily the effect of the boundary conditions only extends within the order of magnitude of the wavelength.

Thus, studying Equation 2 with respect to a general wave function u is sufficient to characterize EM fields on large scales. The results from this approximation are generally referred to as the field of *Scalar Theory*. It gives precise predictions on the propagation of light on macroscopic levels where apertures, lenses, and other optical elements can be assumed to be significantly larger than the wavelength of light.

In the case of a monochromatic wave, the scalar field can be written explicitly as

$$u(\mathbf{P}, t) = A(\mathbf{P})cos(2\pi\nu t + \phi(\mathbf{P})), \tag{3}$$

where $A(\mathbf{P})$ is the amplitude, $\phi(\mathbf{P})$ is the phase of the field at point $\mathbf{P}$ and ν is the optical frequency of the wave in the propagating medium. The time dependence of u is readily available. Using this fact, substituting Equation 3 into Equation 2 yields the Helmholtz equation:

$$(\nabla^2 + k^2)U = 0, \tag{4}$$

where $U(\mathbf{P}) = A(\mathbf{P})exp(-j\phi(\mathbf{P}))$ is the *complex phase* describing the field spatially and $k = \frac{2\pi}{\lambda}$ is the wave number corresponding to the wavelength λ.

The discussion up until now has been concerned with characterizing the field at individual points using the derivatives of said field and properties of the media of propagation. To arrive at an equation that relates values of the field at different points, we employ Green's theorem [29]:

> **Given U and G complex-valued functions of position and a closed surface, S, bounding volume V, if the first and second partial derivatives of both U and G are continuous in V and on S then**
>
> $$\iiint_V (U\nabla^2 G - G\nabla^2 U)d\mathbf{P} = \oiint_S \left(U\frac{\partial G}{\partial n} - G\frac{\partial U}{\partial n}\right) d\mathbf{P}, \tag{5}$$
>
> **where $\frac{\partial}{\partial n}$ denotes the derivative in the direction away from volume V.**

Green's theorem creates a connection between the values and the derivatives of our complex field U in a volume and on the closed surface of that volume. In exchange, Green's theorem requires an auxiliary function, G, termed Green's function. Different choices of such an auxiliary function give rise to different integral theorems, such as the integral theorem of Helmholtz and Kirchhoff.

Without settling on a specific Green's function, we can employ clever analytic techniques and reasonable assumptions on the wave function to have Green's theorem yield

$$U(\mathbf{P}) = \frac{1}{4\pi}\iint_S \left(\frac{\partial U}{\partial n}G - \frac{\partial G}{\partial n}U\right) d\mathbf{P_1}, \tag{6}$$

where S is an infinite plane, and the integration is over that plane. This equation connects values of U on a plane to the value of U at an arbitrary single point and therefore forms the first step towards expressing how waves propagate through space. For the details on the derivation of Equation 6, the reader is directed to Chapters 3.3 and 3.4 of Introduction to Fourier Optics [10].

Sommerfeld defines a suitable auxiliary function [30] as

$$G_-(\mathbf{P_1}) = \frac{exp(jkr)}{r} - \frac{exp(jk\tilde{r})}{\tilde{r}}, \tag{7}$$

where j denotes the complex unit, r is the distance between $\mathbf{P}$ and $\mathbf{P_1}$ and $\tilde{r}$ is the distance between $\mathbf{P_1}$ and the mirror image of $\mathbf{P}$ over S. Plugging G_- into Equation 6 results in

$$U(\mathbf{P}) = \frac{1}{j\lambda} \iint_S U(\mathbf{P_1}) \frac{exp(jk|\vec{r}|)}{|\vec{r}|} cos(\theta) d\mathbf{P_1} \tag{8}$$

where λ is the wavelength of the wave, S is an infinite plane, $\mathbf{r}$ is the vector pointing from $\mathbf{P}$ and $\mathbf{P_1}$, and $\vec{n}$ is the unit normal vector of S at $\mathbf{P_1}$ facing $\mathbf{P}$. We denote the angle between $\vec{n}$ and $\vec{r}$ as θ.

Finally, assuming Kirchhoff's boundary conditions and that $|\vec{r}| \gg \lambda$, we can restrict the integration from the whole plane of S to a finite area, denoted Σ in the followings. This result is known as the Huygens-Fresnel principle, and it will play a key role in our speckle simulation. For a complete derivation and other details, refer to Chapter 3.5 in Introduction to Fourier Optics [10].

In this section, the foundation of the theory behind the beam splitter and the dynamic speckle has been formulated. To arrive to an equation that can actually be put into use in practice, it is necessary to introduce further approximations in the next section.

2.2 Fresnel approximation of the Huygens-Fresnel principle

The derivation in Section 2.1 concluded with the Huygens-Fresnel principle, which expresses the relation between the values of a complex light field on an infinite plane and a single point outside of that plane. Consider now the setup in Figure 6.

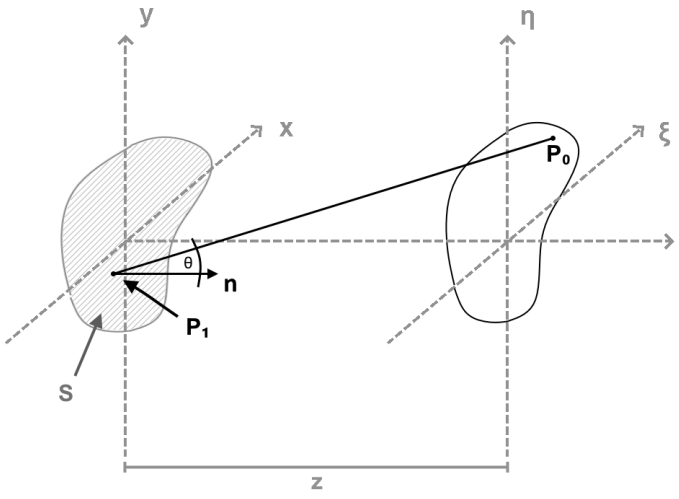

Figure 6: The geometry of the complex light field propagation from one plane to another.

In this setup, light is propagating from the xy-plane towards the $\xi\eta$-plane. The illumination originates from a closed surface S on the former of the two planes. No

light originating from outside of S on the xy plane is taken into consideration. The two coordinate systems are chosen to have their corresponding axes parallel, and the third axis is shared between them. The separation of the planes is denoted by z. The normal vector $\vec{n}$ is perpendicular to both planes and faces the $\xi\eta$-plane from the xy-plane. $\mathbf{P_1}$ is any point on S and $\mathbf{P_0}$ is any point on the $\xi\eta$-plane. Finally, θ is the angle between $\vec{n}$ and the vector pointing from $\mathbf{P_1}$ to $\mathbf{P_0}$.

Figure 6 is an adaptation of Figure 4.1 in [11], with some changes in the notation. The conclusion will be that monochromatic light goes through a Fourier transform while propagating. Therefore, we denoted the plane of origin of the light with xy, signifying the spatial coordinates. On the other hand, $\xi\eta$ will turn out to be coordinates in the frequency domain regarding the original light field.

Let the three-dimensional spatial function U denote the complex amplitude of the light field. In our simulation, we know the values of U on the xy-plane and wish to get a formula for the values on the $\xi\eta$-plane. Let us start developing an approximation to the Huygens-Fresnel principle by restating it in our current setup at hand. Equation 6 then becomes

$$U(\mathbf{P_0}) = \frac{1}{j\lambda} \iint_S U(\mathbf{P_1}) \frac{exp(jk|\vec{r}|)}{|\vec{r}|} cos(\theta) d\mathbf{P_1}, \tag{9}$$

where λ is the wavelength of the light, k is the wavenumber of the propagating light, and $\vec{r}$ is the vector pointing from $\mathbf{P_1}$ to $\mathbf{P_0}$. Note that the region of integration, although still denoted by S, has now been restricted to a finite closed surface.

The cosine can be replaced by its geometrical definition, since we know the cathetus and hypotenuse enclosing θ in the right triangle in which it resides. Therefore, $\theta = \frac{z}{|\vec{r}|}$. Expanding $\mathbf{P_0} = (\xi, \eta)$ and $\mathbf{P_1} = (x, y)$ results in

$$U(\xi, \eta) = \frac{z}{j\lambda} \iint_S U(x, y) \frac{exp(jk|\vec{r}|)}{|\vec{r}|^2} dxdy \tag{10}$$

It is perhaps interesting to take a step back and try to interpret this expression geometrically. Noting that $k|\vec{r}| = \frac{2\pi}{\lambda}|\vec{r}|$, we can see how this value is none other than the phase change, that a beam of light would experience traveling from $\mathbf{P_1}$ to $\mathbf{P_0}$. The multiplication by $exp(jk|\vec{r}|)$ then applies this phase change to $U(x, y)$. Meanwhile, the division by $|\vec{r}|^2$ is comparable to how the power of light is spread out after propagating for a distance of $|\vec{r}|$ in free space. The same amount of energy occupies a spherical surface of radius $|\vec{r}|$ thus at $\mathbf{P_0}$ it is attenuated by a factor of $|\vec{r}|^2$. This quasi-physical reasoning, however, does not explain the scaling factor of $\frac{z}{j\lambda}$. As J. W. Goodman writes, "It is perhaps expecting too much to find such an explanation."

Let us expand the length of the vector $\vec{r}$ as

$$|\vec{r}| = \sqrt{z^2 + (x - \xi)^2 + (y - \eta)^2}. \tag{11}$$

We can see from here that Equation 10 can already be used to actually compute

speckle images since it can be stated as a convolution. The convolution kernel is

$$h(X,Y) = \frac{z}{j\lambda}\frac{exp(jk\sqrt{z^2+X^2+Y^2})}{z^2+X^2+Y^2}, \tag{12}$$

while the convolution would take the form of

$$U(\xi,\eta) = \iint_S U(x,y)h(\xi-x,\eta-y)dxdy. \tag{13}$$

Such a convolution can be calculated using the Convolutional Theorem [31]. The technique is composed of Fourier transforming U and h, taking their point-wise product, and then transforming the result back.

Tempting it might be to use this raw result, we must look at its caveats. The exponential in the numerator changes spatially rapidly thanks to the magnitude of the wavenumber k. What makes this a problem, is that the phase of a single light beam changes completely in the span of a wavelength. Furthermore, this wavelength is typically less than $1\mu m$ in our case. Our targets, such as humans or desks, are of higher order in size, spanning up to meters. This contrast in size means that high precision is required to store these distances. Another computational difficulty arises from the fact that three Fourier transforms are required for each image generation. This fact deems the approach too costly for real-time applications in our case. Later we will see that we can make the simulation computationally much cheaper given certain assumptions about the setup.

Therefore, next, we turn our attention to approximating the Huygens-Fresnel principle. The candidate of approximation is the distance between $\mathbf{P_0}$ and $\mathbf{P_1}$, $|\vec{r}|$. To be able to approximate $|\vec{r}|$, we factor the distance between the two planes (z) out of the square root in Equation 11. We are allowed to do that since z is always positive. The result is then

$$|\vec{r}| = z\sqrt{1+\left(\frac{x-\xi}{z}\right)^2+\left(\frac{y-\eta}{z}\right)^2}. \tag{14}$$

We want to replace the square root with a simpler expression. Since x and y can freely vary in the integration, we can generalize the expression as $\sqrt{1+b}$ where b is any non-negative real number. Let us now look at the Taylor series [32] of $\sqrt{1+b}$ around 0.

$$\sqrt{1+b} = 1+\frac{1}{2}b-\frac{1}{8}b^2+\dots \tag{15}$$

Replacing the two appearing $|\vec{r}|$ expressions in Equation 10 with an approximation introduces errors. These errors depend on how many terms of the Taylor expansion we include. For the $|\vec{r}|^2$ in the denominator, it is generally acceptable to use just one term, z. In the exponential in the numerator, however, $|\vec{r}|$ is multiplied by the wavenumber k. In our case, the wavenumber can be as large as $6.7\times10^7\frac{1}{m}$ in the case of IR waves (940 nm). This means we need a better estimate than merely a single term since even a half π change in the phase can introduce significantly different

results. For this reason, we also keep the second term of the Taylor series in the exponent. The two approximations are then

$$|\vec{r}| \approx z \left[1 + \frac{1}{2} \left(\frac{y-\eta}{z} \right)^2 + \frac{1}{2} \left(\frac{x-\xi}{z} \right)^2 \right] \tag{16}$$

$$|\vec{r}|^2 \approx z^2. \tag{17}$$

Substituting the above two approximations into the Huygens-Fresnel principle and factoring out constants from the integration, we find that

$$U(\xi,\eta) \approx \frac{e^{jkz}}{j\lambda z} \iint_S U(x,y) exp\left(j\frac{k}{2z}((x-\xi)^2 + (y-\eta)^2) \right) dxdy. \tag{18}$$

Geometrically, this approximation replaces the spherical wavefronts with quadratic surfaces. We must ensure that this substitution does not introduce significant errors to have accurate results. For a discussion on the conditions for accuracy and caveats, see Section 2.7.

The next step is to transform the expression on the right-hand side in Equation 18. We expand the square of differences and separate the exponential of sums into a product of exponentials. Finally, we factor constants, with respect to x and y, out of the integration.

$$U(\xi,\eta) \approx \frac{e^{jkz}}{j\lambda z} e^{j\frac{k}{2z}(\xi^2+\eta^2)} \iint_S U(x,y) e^{j\frac{k}{2z}(x^2+y^2)} e^{-j\frac{k}{z}(x\xi+y\eta)}) dxdy. \tag{19}$$

We note that $\frac{1}{j} = -j = e^{-j\frac{\pi}{2}}$ and that the multipliers in front of the integral can be condensed into a single exponential. Introducing the two auxiliary functions

$$\begin{aligned} M(\xi,\eta) &= e^{j(-\frac{\pi}{2} + kz + \frac{k}{2z}(\xi^2+\eta^2))} \\ T(x,y) &= e^{j\frac{k}{2z}(x^2+y^2)} \end{aligned} \tag{20}$$

enables us to look at the equation on a higher level. After expanding $k = \frac{2\pi}{\lambda}$ and moving the multiplication by $\frac{k}{z}$ in the last exponent, we find that in the heart of this formulation, we are dealing with a Fourier transform.

$$U(\xi,\eta) \approx \frac{1}{\lambda z} M(\xi,\eta) \iint_S U(x,y) T(x,y) e^{-j2\pi(x\frac{\xi}{\lambda z} + y\frac{\eta}{\lambda z})} dxdy \tag{21}$$

Notice how the multiplication by both M and T only have an effect on the phase of light. This effect is also an addictive one. We are going to exploit this property later in order to arrive at an even more straightforward methodology. To see what happens with our monochromatic light when it propagates from the xy-plane to the $\xi\eta$-plane we can put Equation 21 into words as follows

- The complex amplitude $U(x,y)$ gets a phase shift in the form of $T(x,y)$. This phase shift is independent of the light field and only varies spatially.
- The phase-shifted signal gets Fourier transformed.

- The transformed signal is scaled by a factor of λz. This scaling accounts for the density change in the light field and therefore preserves energy.
- Lastly, the result is phase-shifted again, this time by $M(\xi, \eta)$.

As a quick sanity check, we question whether the transformation preserves the energy of light. Here, by energy, we mean the total intensity of the light field (see Section 4.1.1 of [10]). The energy of the two light fields are therefore defined as

$$\begin{aligned} E_1 &= \iint |U(\xi, \eta)|^2 d\xi d\eta \\ E_2 &= \iint |U(x, y)|^2 dx dy, \end{aligned} \tag{22}$$

where both integrations are over the entirety of the corresponding plane, since M and T only change the phase of the light, they do not influence the amount of energy. The Fourier transform is also energy preserving; this theorem is known as Parseval's Theorem. The scaling of ξ and η introduces a scaling factor into Equation 22, but this is counteracted by the $\frac{1}{\lambda z}$ factor. That means that even though the Fresnel approximation simplifies the Huygens-Fresnel principle considerably, it keeps it energy preserving.

To make the computation of the beam splitter and the dynamic speckle feasible, further approximations were introduced. Although the computation is now possible for the former, it is still impractical in terms of required numerical precision for the latter. In the following, we look at how to introduce and simplify the effect of the camera aperture.

2.3 Speckle in imaging systems

In Section 2.2 we explored what happens to monochromatic light when it travels between two parallel planes in free space. We arrived at a practical expression in the form of Equation 21 known as the Fresnel approximation of the Huygens-Fresnel principle. This result shows up in the analysis of J. W. Goodman, M. Francon, and others (see [7], [33], [9], and [34]). The setup is more complicated in an imaging system such as a camera. In the following, we explore what happens when an aperture separates the scene and the imaging plane. The analysis follows the work of J. W. Goodman [8].

The model is shown in Figure 7 and described as follows. An object with a macroscopically flat surface is illuminated by a monochromatic light source. For the discussion at hand, it is sufficient to consider the macrostructure of the object. The roughness of the material comes into play in Section 2.4.

The backscattered light propagates through free space from the surface of the object to the plane of the aperture. The light field goes through a transformation while propagating. The aperture then clips the light field, leaving it invariant inside the opening but entirely blocking it off outside of it. The aperture might also be partially transparent and thus attenuate light close to the rim. The clipped field then

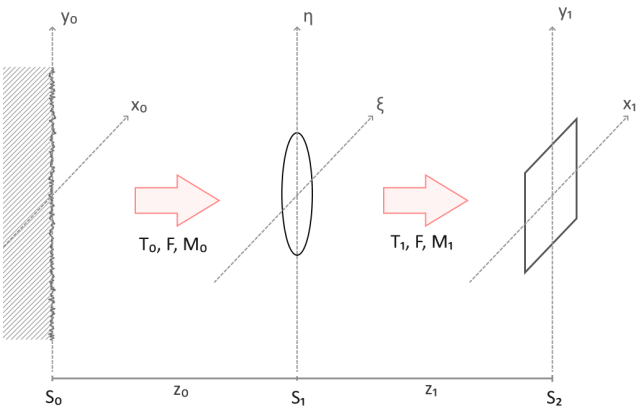

Figure 7: Schematic representation of light propagating from a rough surface (S_0-plane) to an imaging plane (S_2-plane), obstructed by an aperture (S_1-plane). The arrows denote the transformation from Equation 21 with different auxiliary functions.

propagates once more towards the imaging plane, where the intensity is registered. The before mentioned three planes are denoted by S_0, S_1, and S_2 in order and their axes are as shown in Figure 7. The planes in question are parallel; their separations are denoted by z_0 and z_1. We denote the complex amplitude of light on the surface of the object as $U_0(x_0, y_0)$, on the aperture plane as $U_a(\xi, \eta)$ and on the imaging plane as $U_1(x_1, y_1)$.

As seen in Equation 21, the transformation that light goes through during the propagation between two parallel planes is characterized by two phase-shift functions, T and M, the distance between the planes, z and the wavelength of light, λ. Notice that T and M depend on z. Here we are dealing with two consecutive propagation events, and therefore we denote the values that define the two transformations as T_0, M_0, z_0, and T_1, M_1, z_1 respectively. The wavelength is the same in both cases, λ. These variables are also noted in Figure 7. For the sake of completeness, we list the above introduced auxiliary functions.

$$\begin{aligned} M_0(\xi, \eta) &= e^{j(-\frac{\pi}{2} + kz_0 + \frac{k}{2z_0}(\xi^2+\eta^2))} \\ T_0(x_0, y_0) &= e^{j\frac{k}{2z_0}(x_0^2+y_0^2)} \\ M_1(x_1, x_1) &= e^{j(-\frac{\pi}{2} + kz_1 + \frac{k}{2z_1}(x_1^2+y_1^2))} \\ T_1(\xi, \eta) &= e^{j\frac{k}{2z_1}(\xi^2+\eta^2)} \end{aligned} \tag{23}$$

To be able to describe the clipping effect of the aperture, we need to formalize the effect it has. Therefore, we use an aperture function [35], which describes the shape of the aperture. We denote it as $Ap(\xi, \eta)$. It is a real-valued function of two variables representing spatial position. The values it can assume range from zero to one, zero being completely opaque and one being completely transparent. A typical

example of an aperture function is that of a circular aperture of radius r:

$$Ap(\xi,\eta) = \begin{cases} 1 & \text{if } \xi^2+\eta^2 < r^2 \\ 0 & \text{otherwise} \end{cases} \tag{24}$$

To describe the effect an aperture has on a propagating light field, we multiply the aperture function Ap with the field's complex amplitude. Zeros will cancel the light intensity, stopping it from contributing to the final result. Ones will leave the light field unchanged. Figure 8 shows some common aperture functions.

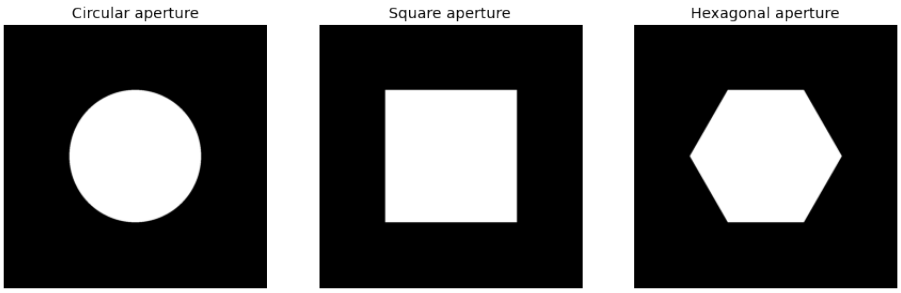

Figure 8: Different aperture functions. The mask is generated on the fly in the simulator.

To make the next step easier to follow, we introduce a scaled version of the Fourier transform as

$$\mathcal{F}_{\alpha,S}\{f\}(x,y) = \frac{1}{\alpha}\iint_S f(\xi,\eta)e^{-j2\pi(\xi\frac{x}{\alpha}+\eta\frac{y}{\alpha})}d\xi d\eta \tag{25}$$

Where f is a Fourier-transformable complex function of two variables and α is the scaling factor. The region of integration is a plane and is denoted by S in the equation. Here $\mathcal{F}_{\alpha,S}$ is, similarly to the regular Fourier transform, $\mathcal{F}$, a function of functions.

Using the scaled Fourier transform, we are able to express the two transformations at hand. We substitute the corresponding value of T,M and z in the Fresnel approximation in operator notation and get

$$\begin{aligned} U_a &\approx M_0\mathcal{F}_{\lambda z_0,S_0}\{U_0T_0\} \\ U_1 &\approx M_1\mathcal{F}_{\lambda z_1,S_1}\{(ApU_a)T_1\} \end{aligned} \tag{26}$$

We can omit the function parameters since we are describing an approximation on the entire domain of the functions, S_1 and S_2, respectively. This concise notation captures the essence of the transformation presented at the end of Section 2.2. Note that the integrations hidden in the Fourier transforms are over two different planes.

In a real physical setup, we have a positive lens instead of a simple aperture. Having S_0 coincide with the front focal plane of the lens and S_2 with the back focal plane, we can make useful simplifications. This focusing acts as if we moved the planes very far away from each other (see Section 10.2 in [6]). To incorporate this

into our model, we assume that z_0 and z_1 are significantly larger than the diameter of the aperture. Denoting the radius of the aperture as r, we have

$$\xi^2 + \eta^2 < r^2 \tag{27}$$

Note that such an upper bound can be given, even if the aperture is not circular in shape. In that case r is the maximal distance of a point inside of the aperture from the origin. The assumption we take into the model is then

$$\begin{aligned} z_0 &\gg \frac{k}{2} r^2 \\ z_1 &\gg \frac{k}{2} r^2. \end{aligned} \tag{28}$$

This approximation of the Huygens-Fresnel diffraction is known as *The Fraunhofer Approximation* [6]. It generally has pretty harsh requirements. However, by adding the focusing lens into the equation, we can utilize it. Using these assumptions, we can simplify M_0 and T_1 as

$$\begin{aligned} M_0(\xi, \eta) &\approx e^{j(-\frac{\pi}{2} + kz_0)} \\ T_1(\xi, \eta) &\approx 1. \end{aligned} \tag{29}$$

Notice how M_0 and T_1 no longer depend on the input coordinates. Putting this result into Equation 26 we have

$$\begin{aligned} U_a &\approx e^{j(-\frac{\pi}{2} + kz_0)} \mathcal{F}_{\lambda z_0, S_0}\{U_0 T_0\} \\ U_1 &\approx M_1 \mathcal{F}_{\lambda z_1, S_1}\{A p U_a\}. \end{aligned} \tag{30}$$

Combining the two together

$$U_1 \approx e^{j(-\frac{\pi}{2} + kz_0)} M_1 \mathcal{F}_{\lambda z_1, S_1}\{A p \mathcal{F}_{\lambda z_0, S_0}\{U_0 T_0\}\}. \tag{31}$$

Our sensor records the power of light hitting it. This power is proportional to the intensity of light (see Section 4.1.1 of Introduction to Fourier Optics [10]) which we will denote with I.

$$I = |U_1|^2 \approx |\mathcal{F}_{\lambda z_1, S_1}\{A p \mathcal{F}_{\lambda z_0, S_0}\{U_0 T_0\}\}|^2. \tag{32}$$

We got rid of the constants in front of the outer Fourier transform since they only change the phase of the resulting field and not the intensity.

Now we use a handy property of the Fourier transform, time reversal [36]. It states that two Fourier transforms, one nested in another, equates to a simple flipping of the domain about the origin. This property carries over to our scaled version of the transform. It holds even with the filter function Ap between the two transforms (this can be proven using the convolutional theorem). This domain flipping manifests in the image on the image plane being flipped about the center. The effect is present in camera systems, it appears even in the most ancient technology. See Figure 9 for a nearly 270 years old example.

Figure 9: Depiction of the *camera obscura* by James Ayscough, 1755. The image inside the house is flipped compared to reality.

Modern digital cameras, such as the one we use, are able to scan the image from any corner of the sensor array and thus flip the image back before any software postprocessing happens. Taking that into account, we can conclude that dynamic speckle in imaging systems manifests as a phase shift followed by a high-frequency filter in a scaled-Fourier domain.

This result enables us to simulate how a known light field gets distorted by a finite aperture. But speckle is more than that, as it also depends on the rough surface of the target object. In the following, we examine how the surface roughness comes into play.

2.4 Surface roughness

Everyday office objects such as chairs, desks, or walls may appear flat to the human observer, but that is just an illusion of our macro-world. We humans also have an intuitive understanding of surface roughness. Touch a smooth plastic chair, a windowpane, a wooden fence, or even sandpaper, and the distinction is imminent. Zooming in on the surface of any of these objects reveals that their microscopic landscape is not flat at all. Take, for example, glass. It seems quite smooth to the human perception, but as seen in Figure 10 it is not flat on tiny scales.

These small variations, spanning only fractions of a millimeter, become important when we look at how laser light bounces off of surfaces. Since different parts of the propagating light field travel different distances before they bounce back from the surface of the object, the light field that is propagating back acquires a wrinkled wavefront (See Figure 11 for illustration). These wrinkles on the wavefront cause interference during the propagation and show up as speckle on the detector. See Section 1 in [8] for more information.

To understand the interaction mentioned above, we need a surface roughness model. Since the surface is macroscopically flat, we say that the average height of the surface is zero. The roughness value assigned to a surface should capture how much the microstructure differs from a perfectly smooth surface. Using the root

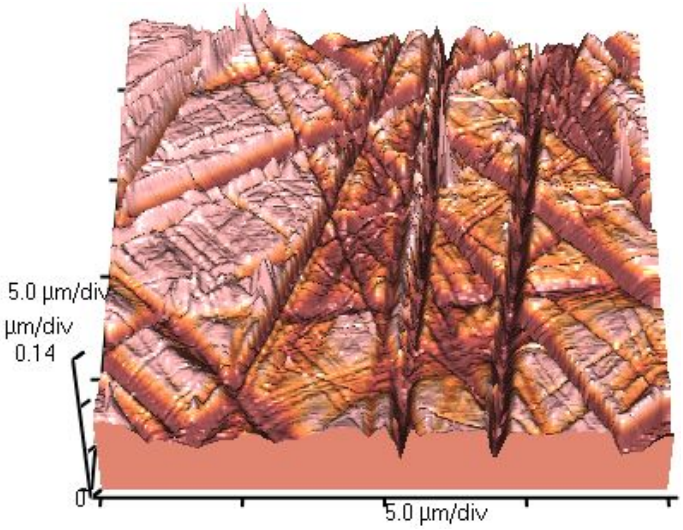

Figure 10: The microscopic landscape of a macroscopically flat glass pane. The capture was created by the Nanorobotics Laboratory at Carnegie Mellon University (see http://nanolab.me.cmu.edu).

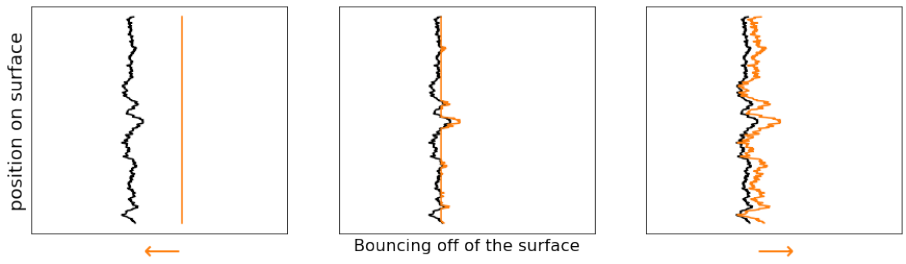

Figure 11: Illustration of how a planar wave (orange) gets wrinkled by bouncing off of a surface (black). The varying surface height introduces phase differences.

mean square (RMS) of the height deviations is common practice [37]. This measure gives a good picture about how the surface varies in height, it is, however, insensitive to the speed of varying, see Figure 12.

The height profile can be thought of as a random process, with an expected value of zero [38]. Since the height at a specific location is ultimately determined by a large number of factors, such as manufacturing errors or micro imperfections in the materials, we can model the height distribution as a normal distribution [39], see Figure 13 for an example. The roughness then corresponds to the standard deviation of the normal distribution. Since the correlation of the height profile is significantly smaller than the height deviation, we can ignore this correlation and look at samples from the surface as if they were independent realizations of the same normal distribution. This approach follows that of J. W. Goodman [8].

Consider a single point on a light wavefront that is incident to the surface of interest. That single point experiences a phase shift, corresponding to twice the height of the surface at the incidence point. We denote the wavelength of our light

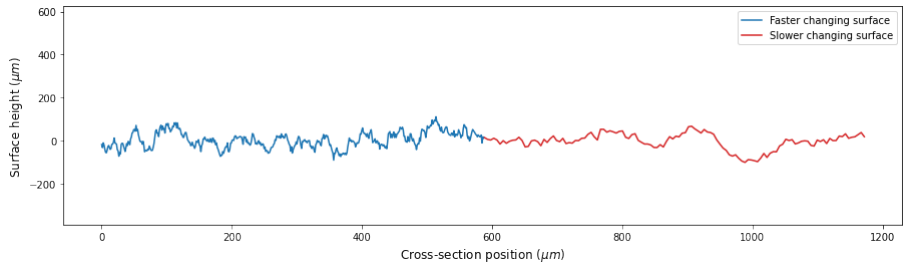

Figure 12: A faster (blue) and a slower (red) varying surface profile. The blue profile has a roughness of 36.84 µm while the red profile has 36.69 µm, nearly identical. The blue surface changes more rapidly, therefore we visually perceive it as rougher.

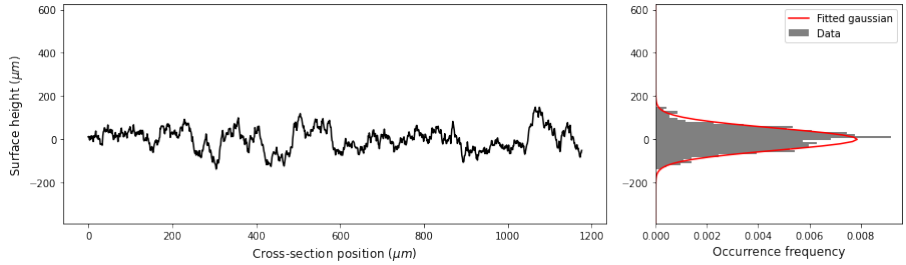

Figure 13: Simulated height profile of a rough surface and its histogram. The standard deviation of the heights is 50.7 µm while the correlation distance is 19.4 µm. The distribution of height values is approximately Gaussian.

field as λ and the random variable describing the surface height as $h \sim N(0, \xi)$. The applied phase offset will be $\frac{2h}{\lambda} \bmod 2\pi$, where mod denotes the fractional modulo operator. Looking at $\frac{2h}{\lambda}$, we see that it is a simple transform of the surface height; it is still a normal distribution with 0 expected value. The modulo operator can be thought of as collecting slices of this spread-out distribution and summing them into a single 2π interval. As ξ grows large, relative to λ, our normal distribution spreads out. The phase offset tends then to approach a uniform distribution, see Figure 14. When dealing with the statistical properties of speckle in Section 2.5, we will build on this observation.

The effect of surface roughness on the phase of an incident coherent light field is established. Therefore, in the following we can reason about the statistical behavior of speckle. This discussion will aid the experimentation and validation later on.

2.5 Speckle statistics

Looking at how the micro-level unevenness introduces a phase-shift in Section 2.4, we concluded that during the reflection of a single beam of light, the phase of that beam gets replaced by a uniformly random phase. Following the analysis of J. W. Goodman [8], we proceed by evaluating the statistical behavior of the complex field on the imaging plane.

To create a statistical model, it is necessary to establish a model of the surface

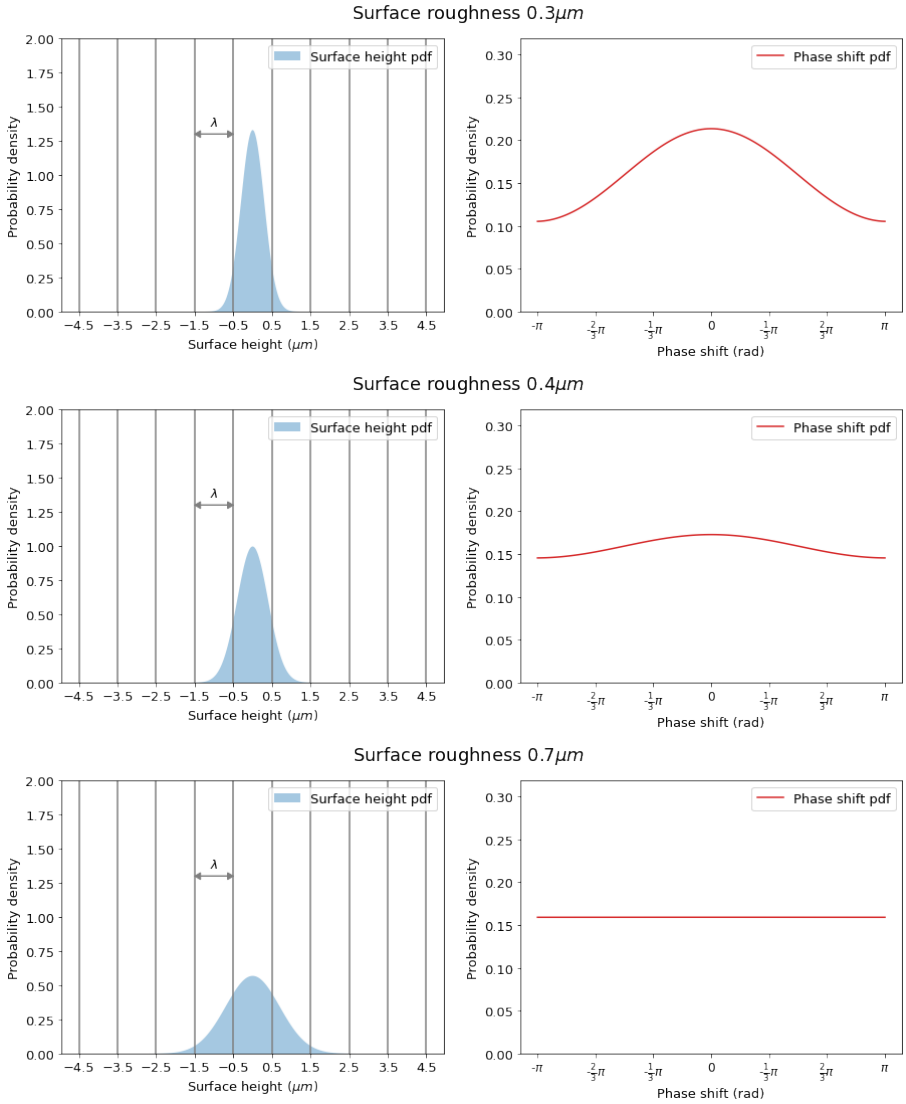

Figure 14: Relation between the probability density function (pdf) of the surface roughness and the introduced phase shift. The plotted surface roughness values are $0.3, 0.4$, and 0.7 µm, respectively, while the wavelength is $\lambda = 1$ µm. Notice how the pdf of the phase shift is already close to uniform when the roughness is comparable to λ.

that takes microscopic features into account. Our model assumes that the area that contributes to the complex amplitude's value at a single point is divided into N parts with equal area. Goodman calls these parts *elementary scattering areas* [8]. We model the contribution of each elementary scattering area using a complex random variable, $\frac{1}{\sqrt{N}} a_k exp\,(i\phi_k)$. Here a_k and ϕ_k are real-valued random variables. Furthermore, we assume that the different elementary scattering areas are independent of one another.

The resulting complex amplitude, denoted as A_N, is therefore

$$A_N = \sum_{k=1}^{N} \frac{1}{\sqrt{N}} a_k exp\left(i\phi_k\right) = \frac{1}{\sqrt{N}} \sum_{k=1}^{N} a_k exp\left(i\phi_k\right) \tag{33}$$

Note that a_k are independent of the number and size of the elementary scattering areas, and the factor of $\frac{1}{\sqrt{N}}$ is included so that the reflected intensity of each element is proportional to the actual area of that element. In the following, we take the same assumptions as Goodman in [7]:

1. (i) The phase and amplitude of all elementary scattering areas are independent. a_i independent of ϕ_j $\forall i, j \in 1..N$.
2. (ii) The phase variables are uniform random variables, meaning that $\phi_i \sim U[-\pi, \pi]$ $\forall i \in 1..N$.

The first assumption arises from our model of the surface, while the second is a direct application of the statistics derived in Section 2.4. Complex numbers can be interpreted as two-dimensional vectors and summing up independent vectors is the classical problem of a random walk. Random walks are fundamental problems of probability [40, 41, 42]. In our case, the most important result follows when we look at the real and imaginary part of A_N:

$$\begin{aligned} Re(A_N) &= \frac{1}{\sqrt{N}} \sum_{k=1}^{N} a_k cos\left(\phi_k\right) \\ Im(A_N) &= \frac{1}{\sqrt{N}} \sum_{k=1}^{N} a_k sin\left(\phi_k\right). \end{aligned} \tag{34}$$

As the number of elementary scattering areas is huge, we can assume that N is large enough to apply the Central Limit Theorem (CLT) (See Chapter 7 in [43]). Individual scattering areas can behave differently, but we can apply the CLT with the Lyapunov criterion using the assumption of finite moments. Applying the CLT yields that as $N \to \infty$, the real and imaginary component of A_N will converge in distribution to a normal random variable, A. The CLT gives a guarantee on the expected value and variance of the limit:

$$\begin{aligned} \mathbb{E}(Re(A)) &= \lim_{N\to\infty} \mathbb{E}\left(\frac{1}{\sqrt{N}} \sum_{k=1}^{N} a_k cos\left(\phi_k\right)\right) = \lim_{N\to\infty} \frac{1}{\sqrt{N}} \sum_{k=1}^{N} \mathbb{E}(a_k)\mathbb{E}(cos\left(\phi_k\right)) = 0 \\ \mathbb{E}(Im(A)) &= \lim_{N\to\infty} \mathbb{E}\left(\frac{1}{\sqrt{N}} \sum_{k=1}^{N} a_k sin\left(\phi_k\right)\right) = \lim_{N\to\infty} \frac{1}{\sqrt{N}} \sum_{k=1}^{N} \mathbb{E}(a_k)\mathbb{E}(sin\left(\phi_k\right)) = 0 \end{aligned} \tag{35}$$

Here have we used the linearity of the expected value operator and the multiplicity over independent variables. The utilized independence stems from assumption (i) above. The expectation of $cos\left(\phi_k\right)$ and $sin\left(\phi_k\right)$ are trivially zero, as per assumption (ii).

The variances can also be derived similarly (see [7]). However, for our purposes, it is only important that they exist and are of finite value. The variances of the real and imaginary components are also equal. Let us call this variance σ.

Consequently, we are dealing with a complex random variable that is composed of a Gaussian random variable for both real and imaginary parts, with the same 0 expected value and σ variance. Such a random variable is commonly known as a circular Gaussian random variable [44].

Imaging systems detect the intensity of the incident light field [10], and therefore the most important statistic about A is the intensity, defined as:

$$I = |A|^2 = |Re(A)|^2 + |Im(A)|^2. \tag{36}$$

One can show that the intensity will have a negative exponential distribution [7].

$$p_I(x) = \begin{cases} \frac{1}{2\sigma^2} exp\left(-\frac{x}{2\sigma^2}\right) & \text{if x} \geq 0 \\ 0 & \text{otherwise.} \end{cases} \tag{37}$$

So far, we have been considering fully polarized light. When dealing with partially polarized or unpolarized light, we can treat the light field as if it was two independent fully polarized fields [11]. These two imaginary fields then get added up on an intensity basis in the detector. It can be shown that a fully polarized speckle pattern, I_2 has Rayleigh distribution [11]:

$$p_{I_2}(x) = \begin{cases} \frac{x}{\sigma^2} exp\left(-\frac{x}{2\sigma^2}\right) & \text{if x} \geq 0 \\ 0 & \text{otherwise.} \end{cases} \tag{38}$$

There are other properties of speckle that are useful in industrial settings, such as the contrast of the pattern. Speckle contrast is widely used in biological imaging, where the flow of some fluid is under observation [45]. Monitoring blood flow in the brain is one such area [46, 47]. Measuring the flow of blood must be done in a non-intrusive way, which Speckle Contrast Imaging is very much suited for. Other biological applications include measuring tissue perfusion [48], measurement of biotissue vibrations [49], and many other clinical applications [50].

On the other hand, high speckle contrast reduces the ability to resolve fine detail [8]. Therefore, in several applications, it is beneficial to decrease speckle contrast. Techniques for reducing it rely on combining multiple independent patterns to exploit the statistical nature of the process. This independence can be realized through a difference in time, space, frequency, or polarization [8]. In this specific case, speckle contrast is also to be reduced, since the loss of fine details leads to uncertainty in the depth estimations.

2.6 Beam splitter DOE

We employ an optical beam splitter to produce the high number of (5303 currently, 2509 in an older design) beams required for the industrial application at hand. A beam splitter is usually designed for a single wavelength and breaks up the incoming laser light into many theoretically identical beams.

The core idea of shaping laser beams [51] into arbitrary distributions originates from the Fourier transform property derived in Section 2.2. Essentially, it is realized using a filter that alters the laser light's wavefront to produce the desired pattern in the far field. One can achieve this alteration in various ways. The simplest and cheapest way is amplitude modulation (see Section 1.2 in [52]), which involves dimming or completely blocking parts of the laser light. On the one hand, this solution is cost-efficient, and it produces good quality output in shape. On the other hand, it operates with low efficiency, losing a large portion of the intensity coming from the laser.

To retain a higher efficiency, phase-only filters can be used. These are small, thin layers of some transparent material that imposes a varying phase shift on the laser light that passes through it. This shift can be achieved by using different materials at different locations or by alternating the thickness of the layer. See Figure 15 for a visualization.

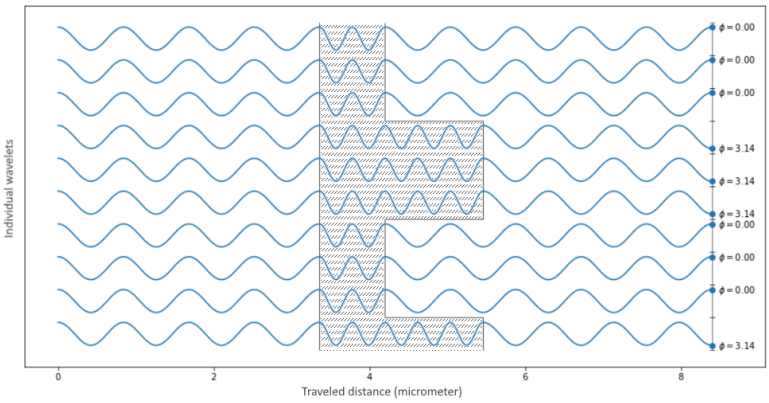

Figure 15: Schematic illustration of a thickness-based DOE. The individual wavelets start with the same phase but get delayed by different amounts due to the thickness of the media.

The interference from the applied phase delay belongs to the phenomena of diffraction. Therefore, the beam-shaping thin layer is called a Diffractive Optical Element (DOE). To make manufacturing easier and the DOE more robust against surface deterioration, binary designs are usually employed. In a binary design, the incoming phase is either shifted by some fixed amount or π radians more. The surface profiles of DOEs are computationally designed.

The two fundamental approaches to DOE design are Direct Binary Search (DBS) and Iterative Fourier Transform (IFT), also known as the Gerchberg-Saxton (GS) algorithm. The basic DBS starts with a random phase map and incrementally changes random pixels but only keeps the change if it improves the result. The IFT, on the other hand, exploits the known Fourier transform relationship between the input and output planes. The algorithm consists of iterations, starting with an initial guess on the phase profile. The guessed phase profile is then used to calculate the output field. This output will be imperfect and as a next step, the intensity of this calculated output field is replaced with the desired output intensity. This corrected

output is then backtransformed to yield a new approximation for the phase profile. To create a new guess for the phase profile, the old and the new ones are combined. The actual way this fusion is happening can follow many different strategies and gives rise to a family of algorithms. The process is repeated until convergence.

These basic algorithms have several limitations that render them insufficient for an advanced industrial application, but they provide excellent ground for improved algorithms. Countless improvements have been proposed to both DBS [53, 54, 55] and GS [56, 57, 58].

Creating DOEs for specific applications gives rise to an entire industrial sector. Companies, such as HoloOr [59] and HoloEye [60] design, manufacture, and sell these optical components. Ladimo Oy works in close collaboration with some of these manufacturers to obtain the DOEs used.

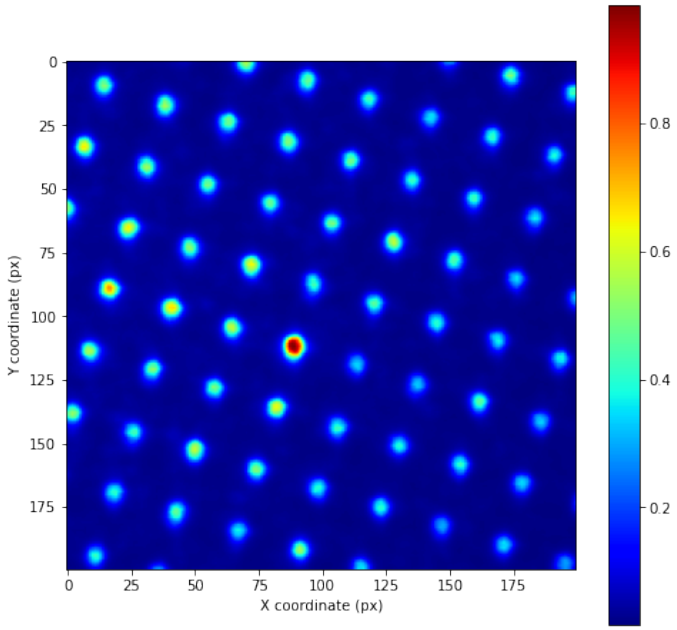

Figure 16: Non-ideal distribution of laser energy introduced by the beam splitter. Notice that the dots do not have a perfectly round shape. The center dot is clearly identifiable, for it has much higher energy than the rest. There is energy scattered between the dots. The image is the mean of 100 images taken in independent configurations in order to cancel out other kinds of random effects.

The design quality of the DOE is a fundamental factor in the pipeline of the company; it directly influences the shape of the dots. Design and manufacturing limitations make the resulting IR image deviate from ideal in multiple ways. These are shown in Figure 16.

2.7 Approximation constraints

Care must be taken when applying approximations to physics problems. This section examines whether the applied simplifications are justified in our case.

The assumptions of Scalar Theory hold well in our case since we are only interested in the microscopical behaviors of the light field. We examine laser light traveling through air and lenses and finally hitting a detector in our setup. In this setup, all components fulfill the assumptions of Scalar Theory.

In the derivation of the Huygens-Fresnel principle and subsequently, in the Fresnel approximation, it is assumed that the light travels a significantly larger distance than the wavelength associated with it. The laser we are using emits light in a narrow band around 940 nm wavelength. Compared to this, the typical object distance of 40cm - 2m is millions of times larger.

The approximation of the Fraunhofer diffraction imposes the most severe conditions. It imposes a minimum requirement on the object distance from the aperture. Looking at Equation 29 and substituting the typical values in our case ($\lambda = 940$ nm and $r \leq 1$ cm) we get

$$z_0 \gg \frac{2\pi \times 1 \text{ cm}^2}{2 \times 940 \text{ nm}} = 3.34m \tag{39}$$

This means that the objects in the scene should be significantly farther away than 3 m. Potentially kilometers. This would not be possible in the slightest in our case. Luckily cameras are composed of more than just a simple pinhole. We have a lens system behind the aperture that focuses the incoming light ray, and it essentially puts the object into virtual infinity. The same can be said about the second condition in Equation 29. This approximation, therefore, is only valid close to the focal plane. Treating out-of-focus objects requires a model that is valid in the near field.

2.8 Laser beam profile

The beam profile of a laser is the spatial distribution of intensity coming from the laser emitter measured along a plane at a fixed distance. Beam profiles can assume several different shapes [61, 62], based on their application, source, and build quality. The fundamental beam profile is the Gaussian profile. Through cleverly designed optical elements, it can be shaped [63] into other profiles used in the industry. Typical beam profiles include square profile, round profile, and Gaussian. Some of these are shown in Figure 17. Modern lasers can produce high-quality beam profiles close to Gaussian.

The beam profile plays a key role in the industrial application at hand. The goal is to estimate the center of the dots appearing on the camera image with as high accuracy as possible. Therefore, any disparity in the beam profile trickles down the pipeline.

2.9 Polarization of light

Light is an electromagnetic wave. This means that at every point in space, a complex vector **E** and **H** characterizes it. Light is also a transverse wave, meaning that **E** and **M** are always perpendicular to one another. Polarization arises because the orientation of oscillation is not predefined by the two fields oscillating. Thanks to

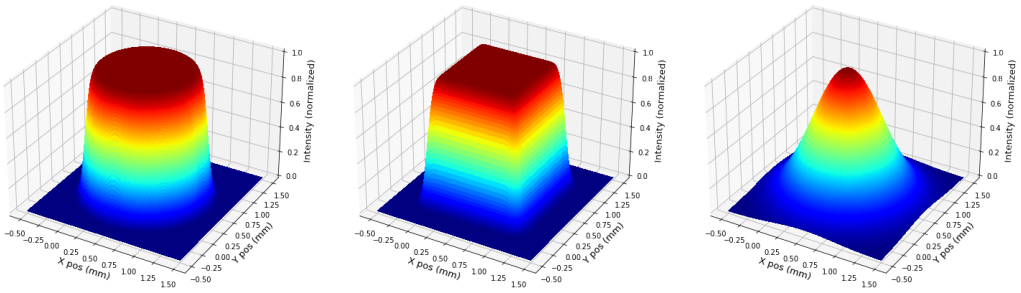

Figure 17: Typical laser beam profiles: flat-top circle, square and Gaussian profile. In practice, these profiles get distorted by manufacturing imperfections.

the perpendicularity property, we can characterize the wave by the orientation of the electric field **E**. A vertically oscillating **E** is referred to as vertically polarized light. Circular polarization comes into play when there is a phase lag between the two fields. Individual waves are always polarized, therefore. But light, in general, can be unpolarized or partially polarized, which means that the polarization of individual photons are uncorrelated.

Throughout the previous sections, we followed the implications of Scalar Theory and implicitly treated light as perfectly polarized. In our setup, the structured laser light is not polarized. Going through the simulation step twice is sufficient to extend our model to describe unpolarized light. As J. W. Goodman concludes on page 200 in Statistical Optics [11]:

> "As a consequence, we can always regard the intensity of a partially polarized wave as being the sum of the intensities I_1, and I_2, of two uncorrelated field components. The average intensities of these components are expressible in terms of the degree of polarization p as $I = \frac{1}{2}(1-p)I_1 + \frac{1}{2}(1+p)I_2$."

Therefore, given the degree of polarization, we can simulate speckle with arbitrary polarization with this equation:

$$I = \frac{1}{2}(1-p)I_1 + \frac{1}{2}(1+p)I_2, \tag{40}$$

where I_1 and I_2 are realizations of fully polarized speckle and the yielded I is of polarization p.

3 Methodology

This section is devoted to the measurements and experiments we took while working towards understanding how speckle behaves. The various experiments provide insights into the nature of this phenomenon. While the theoretical background is laid down in Section 2, here we focus on the practical foundations. Some simplifications in the applied approach were justified only through actually measuring that property.

3.1 Tools

While conducting the experiments, we utilized the following tools. Some of the listed equipment is part of the actual device configuration of the company. Other tools are only used during the setup or evaluation of the measurements conducted in this thesis.

Laser distance meter: To correctly set up the distance between the camera and the plate, we used a Bosch PLR 50 Digital Laser Rangefinder. According to its manual, it is able to measure distances from 5 cm to 50 m with 2 mm of accuracy. It also has a built-in level that we used to ensure that the camera is facing towards the plate without any tilt.

Laser and DOE: Made by Jabil Optics GmbH these lasers have a beam profile close to Gaussian (see Experiment 3.2.9). The emitted wavelength is centered around 940 nm, as evident from the measurement shown in Figure 18. The laser operates on an eye-safe power with a narrow frequency band. The beam splitter design and assembly procedure are proprietary to Jabil. Following the format of Ladimo, the beam splitter splits the laser beam into 5303 (2509 in the older design) beams.

Laser Frequency meter: To determine the energy distribution with respect to wavelength in the emitted laser light, we used a CCS175 Thorlabs Compact CCD Spectrometer. It has sub-nanometer spectral accuracy. See measurement results in Figure 18.

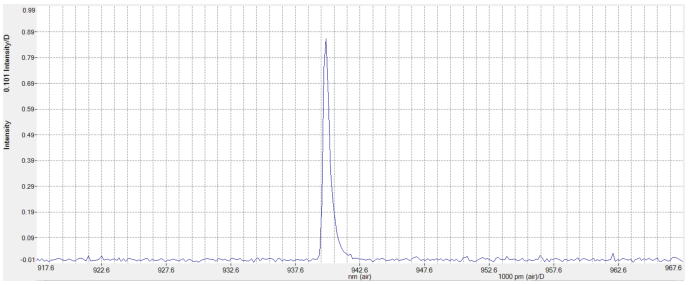

Figure 18: The wavelength profile of our laser emitter. Observe that almost all of the emitted energy is in a narrow band of $1nm$. The measurement was carried out with a Thorlabs Compact CCD Spectrometer CCS175.

Particleboard: The particleboard was the test plate we were using for our experimental measurements. It is composed of rigid medium-density fibreboard with white plastic coating (a.k.a. MDF, common material for office desks), and it exhibits

mostly diffuse scattering properties. The plate is 0.6m x 0.6m, flat on the macroscopic level but rough compared to the wavelength of our laser emitter. The plate is shown in Appendix C, in the experimental setups.

Linear actuator: Also visible in Appendix C, the linear actuator is a device that is capable of precise movement, with a small payload. This payload, in our case, is the measurement device, fixed either parallel or perpendicular to the axis of motion. The actuator is capable of micrometer accuracy movements. It is connected to an Arduino microcontroller which controls the actuator actions. In our experiments, we used this device to capture a large number of images with precise spatial intervals.

IR Camera: The camera in our setup is a custom made one for our purposes. It has a Omnivision OS04A10 light sensor. The attached infrared filter restricts the incoming light to the 850 nm - 940 nm range. In this range the bayering filter on the sensor acts the same in the case of both R, G, and B filters (see Nyxel 2 technology [64]). Thanks to this, the bayering does not influence our results. The aperture can be toggled between f/1.8 and f/11, while the focus can be anywhere between 10 cm and virtual infinity. It is important to note that this camera was used to capture the experiment images, but one with a fixed f-number is used in real devices.

3.2 Experiments

Our simulation is designed to match the application of the developed technology. The setting in which this device is deployed is a typical home or office environment. Usually there is a person sitting in front of a monitor or on the couch. Items such as coffee mugs, a keyboard, or a notebook lay on the table. We set up a series of experiments to understand how the changes in this environment, such as the human moving around or a chair being pushed aside, affect the unfolding speckle pattern. These experiments aim to isolate a single changing variable in the system and analyze the response arising in the speckle pattern.

3.2.1 Image processing model

Previously we considered what happens to a coherent light field traveling through the aperture of the camera. In our imaging systems, however, there is a multitude of further steps happening before the measurement data is stored digitally. Therefore, these effects influence our measurements and must be taken into account.

Even when no light is entering the recording apparatus, the recorded digital image is not constant zero. This is due to electrons spontaneously moving around in photosensitive devices even in the absence of photo-stimuli. This effect is called **dark current** (See [65]). The average value of this noise is a characteristic of the sensor, and it is dependent on the temperature. A map of the average values of this noise under a given constant temperature can be created by averaging a series of measurements in complete darkness. Such a map is termed **dark frame** (See [66]).

A similar challenge is posed by **hot (and cold) pixels**. Due to manufacturing defects, these sensors consistently produce a high (or low) output value. Cold pixels

can be easily filtered out based on their low maximal value. On the other hand, hot pixels can be identified using the captured dark frames.

Gamma correction is commonly employed to more efficiently utilize the intensity resolution of the stored images. This technique exploits the non-linear characteristics of the human vision (See page 260 in [67]) and voids the linear properties of incoming intensity. Therefore, it is troublesome for our measurements and is turned off in our cameras.

Almost all digital color cameras capture colors using a **Bayer filter** (or a modification thereof). This mosaic filter controls which photosensors on a grid will measure red, green, and blue light intensities (See patent). The missing values are interpolated with widely varying approaches. In our case, we are interested in wavelength that is outside of the visible range: typically close to 940 nm. Our cameras operate with Nyxel 2 technology [64] from Omnivision. The efficiency of the color filter around the wavelength we use is closely matched by both the red, green, and blue filters. Therefore, we can turn the demosaicing off and work directly with the captured values (in a RAW format). The efficiency difference between neighboring pixels is negligible.

The measurement images are captured with a 10-bit resolution. The values are stored as integers, and during our handling, we convert them into a 32-bit floating-point representation in the range $[0, 1]$. This conversion is lossless.

Finally, we needed to address the **lens distortion**. This is handled through the existing software solution of the company. The cameras are always calibrated before measurements are taken, and the resulting calibrations are used to undistort the captured images.

3.2.2 General setup

Although the following experiments focus on different changes that the simulation should handle, they share a common starting setup. This setup is composed of the measurement device, as described in Section 1.1, fixed steadily on a linear actuator. The linear actuator is a device that is capable of micrometer precision motion on a single axis. The orientation of the measuring device is fixed compared to the actuator but is different among experiments. The device faces an object with a flat, rough surface, usually a particleboard, similar to the material of our office desks.

The recordings are done at night, lights turned off, minimizing the contribution of ambient light. The camera is set up with the correct focus and exposure time and left for several hours to have the temperature stabilize over the detector plane. After the temperature became stable, we take 1000 images with the laser being turned off. These images, termed dark images, give a measure of the contribution of dark current and point out hot pixels (see CCD Data Reduction Guide [68] Chapter 3). See Figure 19 for an example of dark frames. Finally, the laser is turned on, and the actual measurement images are captured with the same aperture and exposure, ensuring that the taken dark images are applicable to the measurements.

As per the results in Section 2.5 we assume each pixel to sample a Rayleigh distribution with different parameters. This, however, is to be understood as a

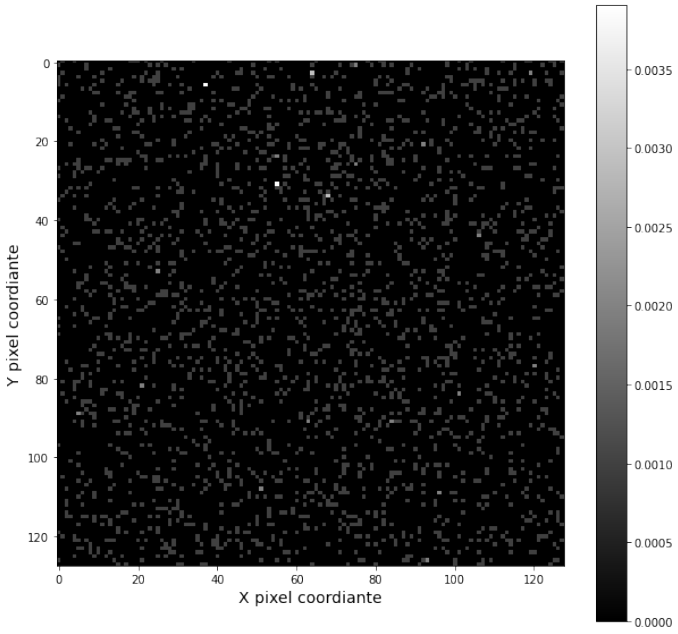

Figure 19: Example of a 128x128 dark image. The image is displayed with normalized colors so that the noise is visible. The scale is relative to the maximum intensity detectable by the camera. The added noise is small, but for a few pixels, it can be significant.

sampling in space but, rather a sampling in time. As time goes on, there is an additive noise that changes, but the underlying Rayleigh distribution gives the same value. It only changes when the microstructure that the camera observes moves.

To be able to measure the correlation of stability in value, we select some pixels from the recorded images. These pixels must satisfy the following criteria:

- The pixel must not be overexposed on any of the measurements.
- The average intensity of the pixel should be high enough (at least 10% of maximal value) so we are not measuring mostly noise.
- The pixel cannot be a hot pixel. That is, it should have a low dark current.
- Depending on the experiment, the selected pixels should be far enough from each other so as to not be correlated.

3.2.3 Stationary particleboard

The particleboard is placed one meter away from the measuring device, facing it. The local Z-axis of the laser is therefore perpendicular to the plane of the board. While capturing the images, the device was standing still in this experiment. We captured 150 images over the course of 5 seconds with equal spacing in time.

We looked at ACF of the captured time series as a first step. Figure 20 shows the ACF and the average of the ACF of 150 pixels. It is clear that the correlation

drops to 0 very quickly. This, however, is due to the underlying expected value being constant, rather than changing. In the case of a constant expected value, we are measuring the correlation of the added noise, which should diminish instantaneously.

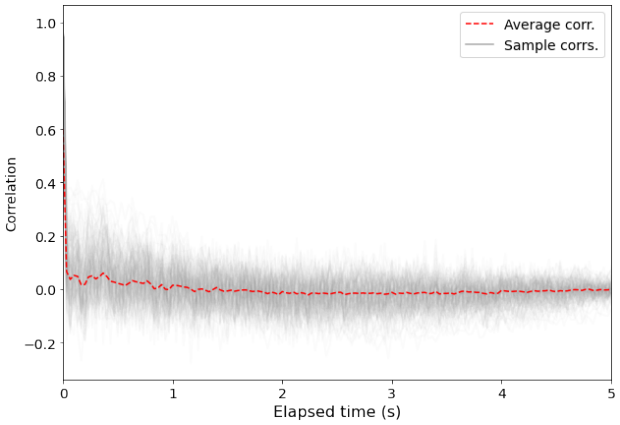

Figure 20: Autocorrelation functions of pixel values in the stationary measurement scenario. The average ACF is also plotted. On average, the correlation vanishes instantaneously.

It would be beneficial to be able to employ a Gaussian time series model in this stationary case. This model would imply that over time the pixel values should fluctuate around a true value by an additive noise. Above we have already shown that this is the case. To see how the variance of this additive noise evolves, we plotted the absolute difference of the selected pixel values from their first measurement. Figure 21 shows an example of a pixel value fluctuating over time and an aggregation of many selected values.

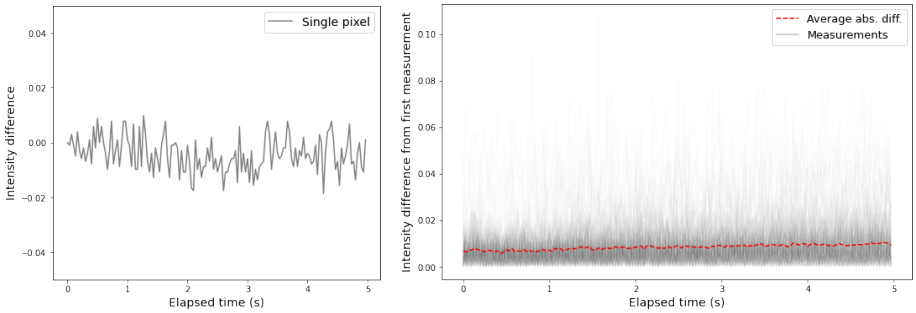

Figure 21: The intensity fluctuation of a single pixel (left) and the aggregated behavior of all 250 selected pixels (right). Observe that the average absolute difference slowly increases over time.

The average absolute differences increase over time but very slowly. This increase is due to subtle effects out of our hands. This could be the dust or other molecules in the air landing on the board surface, the airflow changing, or even the mechanical parts of the camera moving. However, the scale on which this increase happens is really small and, for our purposes, can be neglected.

Conclusion: The Gaussian time series model implies that the fluctuations over time can be countered by taking several images over time and working with their average. The time-dependent additive noise is a Gaussian noise. We can calculate that in these measurements, the standard deviation of the noise is $\sigma = 0.015$. To reduce it below 0.005, it is sufficient to take the average of 9 time samples. In the following experiments, we employ this strategy, usually with 16 images averaged.

3.2.4 Displacement along the X-axis

The board is placed 50 cm away from the device, perpendicularly to the viewing direction of the laser. The axis of the linear actuator is parallel to the plane of the board. See Appendix C for illustration. The actuator moves a total of 7 millimeters by 7 micrometer increments. In each position, we capture 16 images and store the mean.

Displacing the device along the X-axis far enough results in uncorrelated pixel values. This experiment aims to shed light on how far along the X-axis the device should be displaced for that. The question might arise whether the roughness of the surface of the object influences this correlation distance. As long as the surface is rough enough compared to the wavelength of our laser, the captured speckle should exhibit the same statistical behavior [7].

Treating the measurements of each selected pixel as a time series, we can plot their autocorrelation function (ACF). When the setup is stationary, the ACF should stay high over time. However, in the case of this experiment, we expect the ACF to start out at 1.0. As the lag increases, it should then drop to 0 on average.

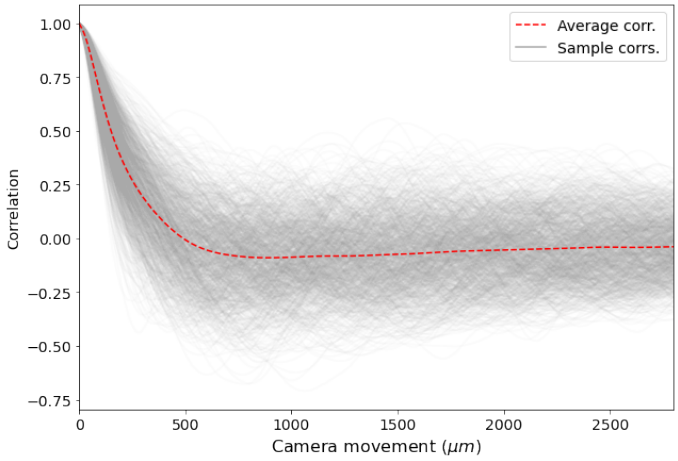

Figure 22: Autocorrelation functions of pixel values upon moving along the X-axis. The average ACF is also plotted. On average, the correlation diminishes after 0.5 mm of horizontal movement.

Individual ACFs fluctuate due to the probabilistic nature of the measurements. However, on average, they drop from 1 to 0 throughout 0.5 mm and after that they stay on average 0. See Figure 22.

Conclusion: The pixel values decorrelate on a distance that is small on the scale of our everyday life. Depressing a key on a keyboard, laying forward, or blinking are all on a larger scale. Displacement along the Y-axis or any axis on the XY plane is assumed to yield the same results.

3.2.5 Displacement along the Z-axis

Similarly to the previous experiment, the particle board is placed perpendicular to the viewing direction of the laser at 1 m of distance this time. The difference is that the movement axis of the actuator is parallel to the viewing direction. See Appendix C for illustration. The actuator moves 7 millimeters away from the board by 7 micrometer increments. In each position, we capture 16 images and store the mean.

Care must be taken to account for the shifting of dots. As the laser moves backward, the location where the dots hit the board slowly shift. This shift is different for different dots and is large enough to influence the ACF. To avoid this problem, we only look at pixels of the center dot. The center dot is produced by the laser beam aligned with the laser's Z-axis. On the one hand, this beam hits the object at the same place regardless of the distance along the Z-axis. On the other hand, the area covered by the dot changes. The dot also moves on the image, the movement, however, is minuscule and can be disregarded.

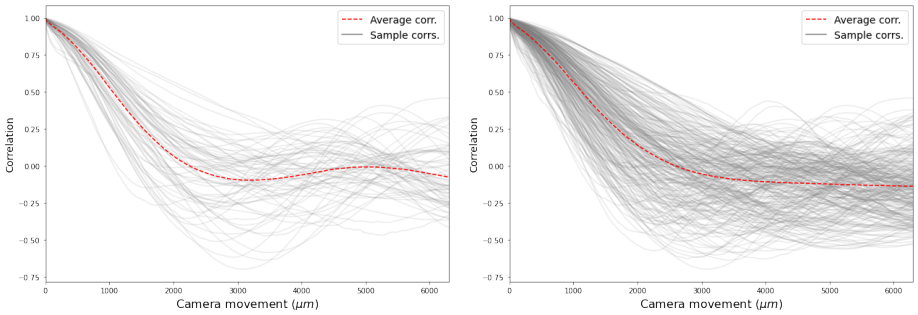

Figure 23: Autocorrelation functions of pixel values upon moving along the Z-axis. The average ACF is also plotted. Pixels selected solely from the center dot (left) and from the central 3x3 dot (right) are plotted.

The center dot is composed of 69 pixels that fit the requirements. If we select pixels from the neighboring dots too (a total of 9 dots in a 3x3 pattern), we can work with 272 pixels. These two ACFs paint a very similar picture, namely that the correlation diminishes after 3mm of movement (Figure 23). The ripples on the average ACF, in the case with fewer pixels, can be attributed to the low number of samples and should flatten out if we could use more data.

Conclusion: Although the correlation spans farther along the Z-axis than on the X-axis, it is still on a small scale compared to everyday movements. Consider a person moving their mug in their hand while having a video meeting. This mug will travel distances on the scale of decimeters.

3.2.6 Rotation about the Y-axis

Moving along different axes causes the laser beams to be reflected from (not necessarily entirely) different areas of the target surface. In this experiment, we set out to discover how hitting the same micro-structure from different angles influences the detected values. The plate is rotated for a total of 2.5 degrees and images are captured in 70 different positions.

Comparably to moving along the Z-axis, we are to consider how the dots shift during the rotation. It is once again sensible to restrict our selection to the pixels of the center dot. Refer to Figure 24 for a comparison between considering only the center dot and also considering the 3x3 neighborhood of it.

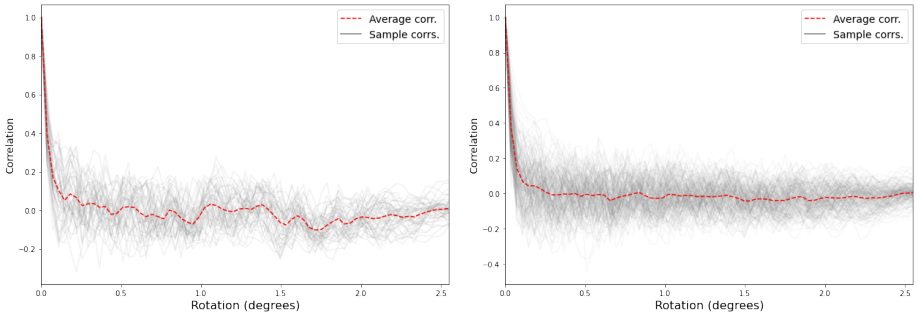

Figure 24: Autocorrelation functions of pixel values upon rotating about the Y-axis. The average ACF is also plotted. Pixels selected solely from the center dot (left) and from the central 3x3 dots (right) are plotted.

Conclusion: It is clear that the pixel values decorrelate extremely fast. A change of merely 0.25 degrees causes the intensities to decorrelate. This is an important aspect since natural movements in 3D are rarely composed of only one directional motion.

3.2.7 Airflow

While the laser light is propagating from the emitter towards the scene (and later back), it travels through non-stationary air. Simulating these effects is far from trivial. A novel approach at simulating these effects by Bar Chen et al. [69] uses Monte Carlo simulation to simulate this effect.

In this experiment, we investigate how strong the effect of airflow is. We replicated the configuration presented in the stationary setup and recorded three different situations. In the first case, a disturbance is generated by blowing in front of the camera. In the second case, we moved around in the room but outside of the field of view of the camera, causing disturbances in the air. Finally, in the third case, a piece of cardboard is used to fan by hand. We moved it up and down next to the field of view as if airing on a hot day. In each case, 300 images were recorded over the course of 10 seconds.

Yet again, we are considering how much the pixel values deviate from their first measured value. The results are plotted in Figure 25. We can clearly see the three

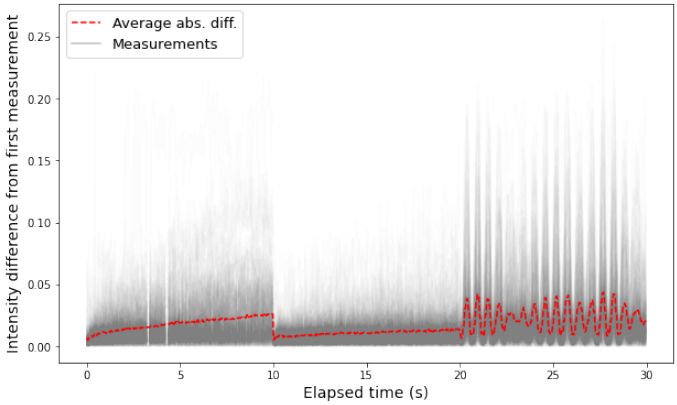

Figure 25: Absolute difference of pixel values in three cases: blowing air, moving around in the room, and fanning with a cardboard. The three segments separate clearly.

segments separating. In the first case, the deviation is greater than in the stationary setup since we blew air directly in front of the laser. The second case matches the base case of no manual airflow closely. Finally, in the third case, we can clearly see the "flops" of the cardboard.

Conclusion: With greater airflow comes greater noise, as we could expect. Accurately simulating the effect of this disturbance is computationally too expensive compared to the magnitude with which it affects the results. If the airflow is insignificant, such as when no one is fanning, and there is no wind in the room, we can still model the effect with a Gaussian time series.

As a possible future expansion of the simulator, it would be possible to conduct deeper research on the subject and come up with simplifications. This could potentially reduce the problem of summing up specially generated independent speckle on an intensity basis.

3.2.8 Human experiments

A key element of the scene observed in our use case is the human user. In a 3D call or a game streaming scenario, the focus is on the actual person in front of the camera. The appearance of people includes several complicating factors in this setup. We constantly move, even when sitting as still as possible. The fabric that we wear and the hair on our heads have a complicated microstructure. Our skin exhibits subsurface scattering. These properties interact with the wave nature of light and influence the resulting speckle.

In this experiment, we sat in front of the laser light and tried to sit as still as possible. Back leaning to the wall and not even breathing. Even though the laser is eye-safe, we used safety goggles. Images were captured over 5 seconds with 30 frames per second. We created a mask by hand to separate two regions on the body: skin and fabric. Pixel values were then extracted from these two regions in a similar fashion to the stationary experiment, and the correlation functions were plotted in

Figure 26. We selected 744 pixels from the skin region and 2446 pixels from the shirt region.

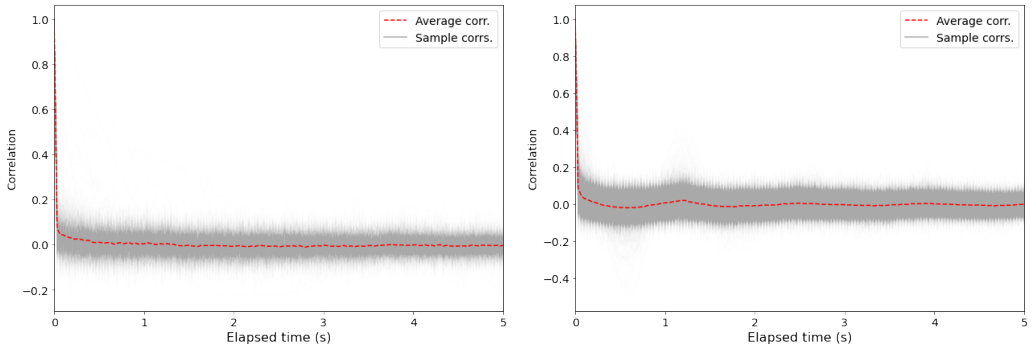

Figure 26: Autocorrelation of pixel values selected from a human sitting very still. Pixels selected from both regions of skin (left) and shirt (right) lose correlation instantly.

Before conducting the experiment, we expected to see similar behavior to the stationary experiment since nothing was moving much in the scene. According to the plots, however, the small movements humans are constantly making and other effects like the subsurface scattering of our skin decorrelate the speckle fast.

Conclusion: Speckle in laser light that falls on human skin or fabric changes rapidly. Initially, we planned to do further tests with humans, such as just sitting without holding our breath or typing on a keyboard. These setups would result in even more movement of the subject, and therefore, it is safe to assume that the results would be identical.

3.2.9 Recovering the beam profile

As established in Section 2.8, the beam profile plays a vital role in the precision of our application. In this experiment, we attempted to recover the beam profile solely from the captured IR images. This profile is supposed to be very close to Gaussian. We captured images in the same setup as in the experiment about moving along the X-axis, with the difference that each move was large enough to decorrelate the pixel values (0.5 mm).

Although several effects contribute to deviation the resulting dot shapes from the ideal Gaussan, we can get rid of all of them and estimate the original beam profile. This can be done with the power of the Central Limit Theorem (CLT). The following conditions capitalize on the CLT:

- To counter the additive noise, we take a large number of measurements, around 100.
- To attenuate speckle, we take only independent measurements, by moving the device between each take.

- To nullify the distortion of the dot shapes by the beam splitter, we average over all 5000 dots, assuming their shape to be independently deformed.

Thanks to these steps, we are expecting to end up with an intensity distribution close to Gaussian after averaging.

We must note that here we are dealing with patches rather than individual pixels, and therefore we need a different over-exposure rejection strategy than in previous experiments. We drop any patch with at least one pixel overexposed in more than 1% of the measurements. The exposure is also kept low enough so that we only lose a few patches. The dark frame values are subtracted as before.

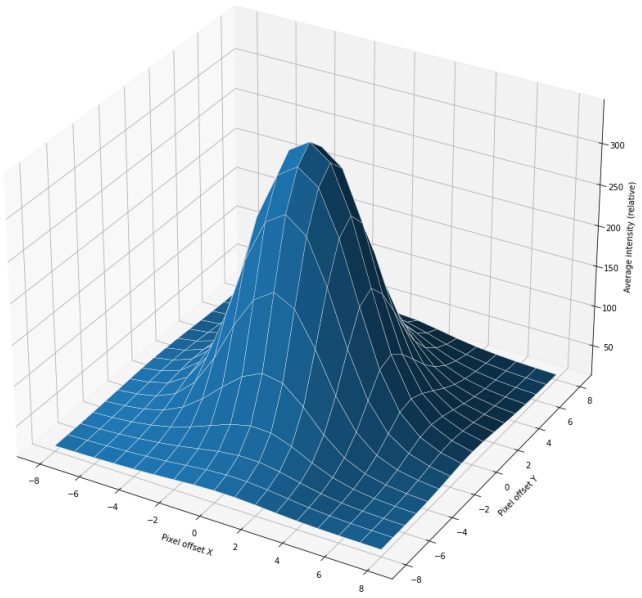

Figure 27: Recovered beam profile. The plotted profile is the mean of 5303 dots over 100 independent measurements.

Finally, the extracted patches might have a subpixel misalignment. An approach to fixing this problem would be to offset the images by that amount before adding up. The techniques for subpixel shifting can distort the spatial distribution unless we explicitly assume the underlying distribution to begin with. Instead of this approach, we rely on a large number of independent dots being captured. The subpixel misalignment is always smaller than half a pixel, and we assume these offsets to be independent of one another. Therefore, averaging over 5000 dot images should turn this misalignment into a slight Gaussian blur.

Figure 27 shows the recovered profile. It clearly exhibits a Gaussian shape.

This experiment confirms that we can model the beam profile as being close to ideal. The distortions we see on the projected image are a byproduct of the DOE rather than the beam profile.

3.2.10 Spatial consistency

In this experiment, we are comparing two different scenarios. The setup is yet again the same as in the stationary experiment. In scenario one, measurements are taken throughout some seconds without changing anything on the configuration. In scenario two, the camera moved along the X-axis with an amount that already decorrelates the speckle pattern. This movement was done between every consecutive measurement.

With this setup, we effectively inspect the results of having a different microstructure while every macro-property stays the same. This can be interpreted as a re-randomization of the underlying phase offset map discussed in Section 2.5. What we expect to see is that, in scenario one, the depth estimations - done by the algorithm of the company - fluctuate significantly less than in scenario two.

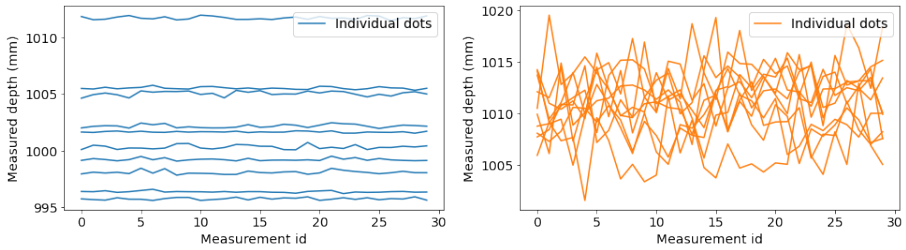

Figure 28: Depth estimations on 10 separate dots over 30 measurements. Stationary setup (left). The camera is displaced along the plane of the target (right). Observe how the fluctuation is significantly higher in the latter case.

We plotted the depth estimates in the two described scenarios in Figure 28. It is straightaway obvious that the two behave differently. In the stationary scenario, each dot has a characteristic bias coming from the unchanging speckle. As soon as this speckle changes due to the movements between each measurement, the depth estimations experience a larger fluctuation.

Conclusion: The depth estimating algorithm is sensitive to the changing nature of the speckle. In practical terms, this means that completely stationary objects will have a depth estimation with a constant bias but much less variance than moving surfaces. Another takeaway from this experiment is that the time-dependent Gaussian noise has an order of magnitude smaller effect on the depth estimation than the changing speckle pattern.

3.3 Proposed Algorithm

To extend the simulator with the new simulated effects, we first look into the mechanisms of the DOE. As shown in Section 2.6 the beam splitter can be modeled as the Fourier transform of the beam profile, phase-modulated by the DOE design. The DOE design is taken as an input by the simulator. The beam profile is either read from an image or generated on the fly, based on parameters that control the shape. These parameters for a Gaussian profile are the total intensity and the covariance

matrix. With these two at hand, we can already generate the Ladimo pattern. The resulting pattern is downscaled to 2048x2048 and cached. Alternatively, we can choose to use the previously employed perfect representation of the pattern. The two approaches have the same runtime cost and can be switched between at any time.

The Ladimo pattern represents the cross-section of the propagating laser light at a unit distance from the emitter. The laser light is already in the far field by then and therefore does not change shape along the Z-axis other than expanding (See Section 2.2). This fact enables us to efficiently find the incident light energy at any point by treating the Ladimo pattern as a texture and sampling it at the right place. The intensity is scaled by the distance to the power of -2.

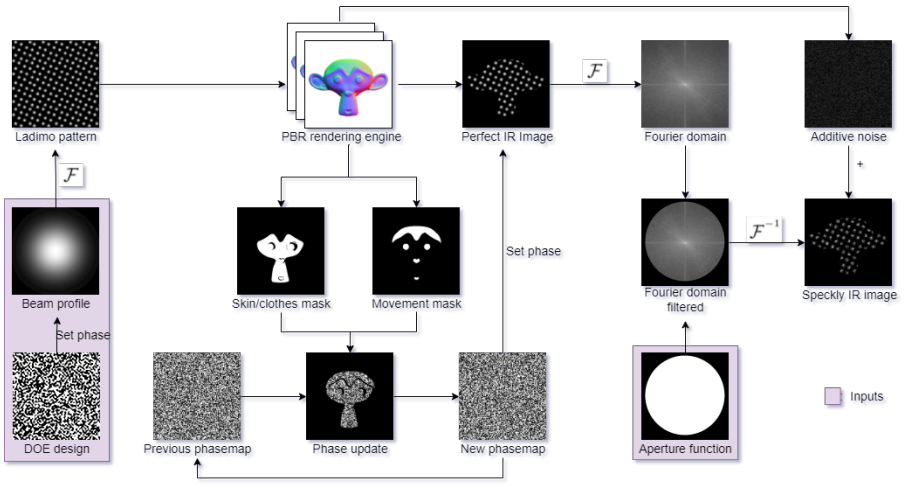

Figure 29: Schematics of the proposed algorithm. The purple background signifies input variables.

The rendering engine takes care of rendering what an IR camera with infinite aperture would see. To introduce the dynamic speckle into the mix, we further filter this ideal IR image. The laser light that bounces back from a rough surface acquires a random phase (see Section 2.5). To implement this, we set the phase of the ideal IR image to a uniform random phase map. The random phase depends on the contributing microstructure of the hit surface (see Experiment 3.2.10). Therefore, this phase map is updated every frame with new values at locations where the seen geometry moved relative to the camera. Additionally to moving surfaces, complicated surfaces such as those exhibiting subsurface scattering (like human skin) are also changing their phase contributions (see Experiment 3.2.8). We also update the phase map at locations corresponding to those surfaces.

The phase-modulated ideal IR image is the Fourier transformed one. The transformed image is what arrives at the aperture of the camera. This field is then clipped by the aperture. We model this clipping as masking by an aperture function. This mask is also an input of the simulator but can be generated on the fly from variables that control the aperture shape. For a circular aperture, the parameter is simply the relative diameter of the opening. Finally, the clipped image travels from the aperture

to the imaging sensor and goes through an inverse Fourier transform to yield the final result. Multiple clever tricks are at play here that make this step so simple:

1. Looking back at Equation 32 in Section 2.3 we notice some differences. Even with the assumption of focus and suitable object distance, we still need to apply a phase modulation, denoted T_0 there. A peek at the definition of T_0 in Equation 23 reveals that it is an additive change to the phase. Since our input has a uniform random phase, we can disregard this offset.

2. Still in Equation 32 we note that two Fourier transforms should be done, not a Fourier transform and an inverse. This is shown later to yield the same results, only flipped. This flip naturally occurs in imaging systems and is compensated for during reading. Hence, our simulation can skip this step.

3. The last issue to address is the question of scaling. The Fourier transforms above are scaled Fourier transforms. The solution here is that the different images we work with are in different coordinate systems. We theoretically deal with an object space, the aperture space, and the space of the imaging plane. In practice, we can represent each of them with just one intensity texture.

As mentioned in Section 2.2, we could directly use the Huygens-Fresnel principle to compute speckle in an asymptotically fast way. All we would need to do is to state that formula as a convolution. Convolutions can be calculated as fast as Fourier transforms. This approach would result in accurate predictions on extremely short distances. The problem with it arises when we consider the magnitude of values at play. Different values in this formula have vastly different orders of magnitude. We are talking about meters where nanometer offsets can make a significant difference. Using the Fraunhofer approximation and some reformulation, we are able to hide these magnitude differences in the implicit coordinate system in which our images are defined. Using such an approximation means that we restrict our simulation to the far field. This is not a problem in our case since it starts after just some millimeters.

This process is suitable to produce polarized speckle. To be able to generate unpolarized or partially polarized speckle, we follow Equation 40. This introduces a new input, the degree of polarization. We keep a separate second randomized phase map that is updated using the same mask as the first one. Then following the same pipeline, a second IR image with speckle is generated. The two images are then combined on an intensity basis before any further processing.

The output image that has the added speckle effect then goes through the image postprocessing. This step consists of adding various noises such as a time-dependent Gaussian noise and the camera noise. The complete extension to the existing pipeline is depicted in Figure 29.

4 Results

Ladimo Oy has already developed a tool to simulate physical setups in which their algorithm can be tested. This simulator also serves as a tool to rapidly accumulate training data for learning models developed by the company. This tool, referred to as the simulator from now on, renders the captured images of the IR camera(s) assuming ideal conditions. In the actual physical setup, several effects contribute to the final result, making it deviate from ideal. For more details, we refer to Section 1.2. In order to increase the utility of the simulator, we set out to devise an extension of the simulator to accommodate the significant effects out of those listed above. In this section, based on our physical experiments and theoretical background (see Section 2), we propose this extension. We also evaluate the resource requirements it has (both theoretically and empirically) and compare the quality of the output with that of the experiments and theoretical predictions. The added effect of speckle in the simulator is shown in Appendix E.

4.1 Performance

In the following, we examine the computational and memory costs of the proposed algorithm theoretically and empirically. As per the company's requirements, our goal was to conceive an approach suitable for real-time applications.

Memory usage: The extra memory requirement is easy to judge. The Ladimo pattern and the final IR image were already taking up space in memory. The update masks and the actual update are calculated on the go, not requiring any memory. We added the phase map, the phase channel to the two IR images, and the Fourier domain image. The aperture function can also be modeled on the go or be an actual texture. The images are of size 2048x2048 and of single-precision floating-point format. The added memory cost is therefore 16MB/image * (2 + 3) = 80MB. Generating partially polarized speckle requires a second phase map and thus adds another 16MB to the memory need. We can avoid using a second speckly IR image by reusing the existing perfect IR image for that purpose. Additional structures like the random number generator do not consume more than 1MB. This amount of GPU memory is considerable but still far from problematic.

Theoretical treatment: We are going to assess the asymptotic running time of our algorithm. We denote the number of pixels on the used images with n. Observe that all operations are simple element-wise operations such as phase-setting, masking, or Fourier transforms. This is true regarding polarized and unpolarized speckle alike. The time complexity of the Fourier transform is lower bounded by $\Omega(nlog(n))$ as far as research is at our time. Since the operations are in sequence, the dominant step will be the Fourier transform. The popular divide-and-conquer approaches, such as the split-radix FFT algorithm, deal well with composite values of n. Pixel counts that are the product of only a few primes can be handled by padding the image that is being transformed. In our case, n always has many prime factors. Typical values include 1024x1024, 1024x1920, and 2048x2048. The total running time shall fall in the class of $\mathcal{O}(nlog(n))$.

Empirical treatment: We tested our implementations on various image sizes to put actual numbers on the theoretical results. We created two separate implementations. A CPU implementation using Python 3.6.9 and NumPy 1.19.5 and a GPU implementation using C++17 and Cuda (NVCC version 11.1.105). The CPU version was created for prototyping and statistical analysis, while the GPU version is the one integrated into the simulator. CPU tests were carried out on an AMD Ryzen 7 3700X 8-Core Processor. NumPy is able to take advantage of advanced features in the CPU such as the vector instructions and instruction-level parallelism but only utilizes one physical core. Our memory access pattern is quite optimal since the required operations are simple in nature. GPU tests were run on a machine with GeForce RTX 2070 SUPER and the same CPU as mentioned before. We randomly selected resolutions with many prime factors in the range of 8x8 to 2048x2048 and performed a single frame of speckle simulation on a uniformly white input image. Note that the running time is independent of the actual content of the images. Partially polarized speckle images should take roughly twice the effort to generate as unpolarized ones. The performance plot is shown in Figure 30.

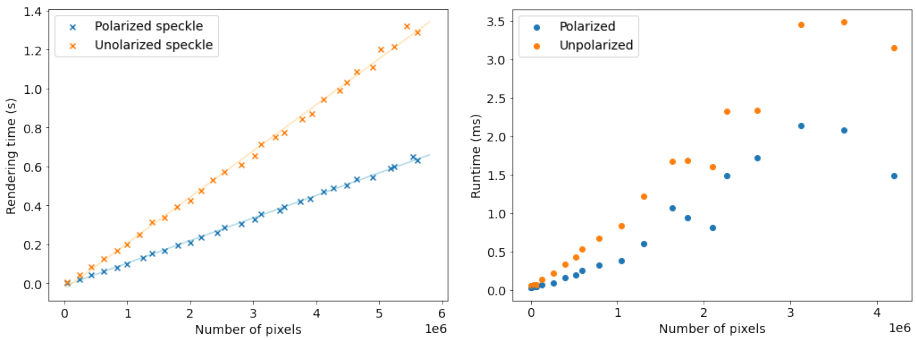

Figure 30: The performance of CPU (left) and GPU (right) implementation across various pixel counts. Note that the CPU implementation is on the order of seconds, while the GPU one runs on the order of milliseconds.

The CPU times show a linear trend. This is probably explained by the memory access. Having a high pixel count means that the L1 and even the L2 caches are too small to fit an entire image, and therefore only little memory reuse can happen. This slow access hides the logarithmic factor, because once the data has arrived, the actual calculations are carried out very fast. We can see that some sizes are more beneficial than others on the GPU side. We notice significant speedups at the resolutions 1024x1024, 1024x2048, and 2048x2048. This is probably due to the implementation of FFT in Cuda that we suspect to pad the input to a power of two size. The overall curve tends more upwards than a simple linear curve. The logarithmic factor is visible here since the massively parallel architecture is able to access data in parallel fast enough to show the computational cost.

Nevertheless, we can conclude that both implementations are satisfactory in speed to their respective applications. This is true for unpolarized and partially polarized speckle alike.

4.2 Evaluation

Validating the quality of the output of our algorithm can only happen through statistical means since setting up the exact same conditions in real life and simulation is infeasible. In the following, we create specific scenarios that are simple enough to be evaluated theoretically or can be statistically compared to physical measurements.

Distribution of a single pixel value: Suppose a setup of uniform lighting of a flat, rough surface facing the laser and camera. We investigate the distribution of the intensity of a single pixel. In Section 2.5 we established that the PDF of the intensities in such a setup should follow an exponential distribution. We created 64 thousand 128x128 images to probe this property, then added speckle to each independently. After this, we plotted the histogram of a single pixel location. Finally, we fit an exponential distribution to the acquired data. We applied the same treatment to the unpolarized case, expecting a Rayleigh distribution this time.

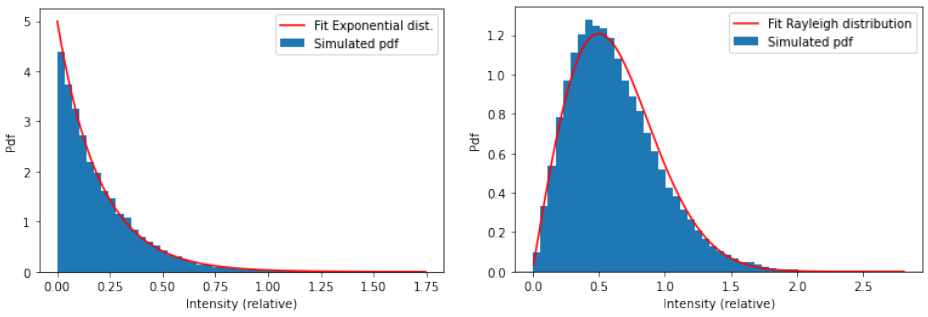

Figure 31: Distribution of polarized (left) and unpolarized (right) speckle intensities. In both cases we generated a large number of independent samples and fitted the corresponding expected distribution.

Examining the plots in Figure 31, we can conclude that the simulation matches this basic property well. The slight fluctuations and offset might come from the small size of the image (128x128).

Depth estimation with changing f-number: We created the same setup as in Experiment 3.2.3 with a different camera. This camera is capable of changing the aperture freely in a range. We took three measurements with three different f-numbers: f/2.8, f/4, and f/5.6. By increasing the f-number, we reduce the size of the aperture that lets in light. This results in a lower SNR and thus increases uncertainty in the depth estimation. More importantly, the effect of speckle increases rapidly as the aperture gets smaller. We manually set a higher exposure time on consecutive measurements to compensate for the smaller aperture. The depth estimations are plotted in Figure 36.

We recreated the very same arrangement inside the simulator. Additionally, f/0 is also plotted, which is physically impossible to achieve. Having the f-number as f/0 showcases the effect of other noises since no speckle is observed here. The generated intensity maps were scaled to match in total intensity. The resulting depth estimations are plotted in Figure 35.

Looking at the two results, we can conclude that the effect is similar. In both cases, the depth estimations acquired an uncertainty, which manifested in a normal distribution of depth values around the actual distance. We also observe a difference in the two cases. The simulated depth values have a lower variance for the same f-number, which is due to other effects also exhibiting an exposure-dependent behavior. This relation poses a possible future addition to the simulator. The current match between the measured and simulated results is satisfactory for present testing purposes.

Recovering beam profile: Speckle removes information from our rendered scene; it is basically a Fourier domain low-pass filter. But treating the resulting image as a random variable reveals that not all information is lost. Each and every pixel of a recorded image is a sample taken from an exponential distribution (or Rayleigh in the case of unpolarized speckle). Therefore, the random variables corresponding to each pixel have an expected value. If we average many generated dot images, we should get the laser beam profile back. This property is analogous to Experiment 3.2.9.

We have generated an image with around 2500 perfectly Gaussian dots to test this property. Then we added unpolarized speckle to the image with an f-number of 2.8. Samples from these patches are shown in Figure 38. We immediately notice that the nice Gaussian shape is gone, and the center is easy to make out.

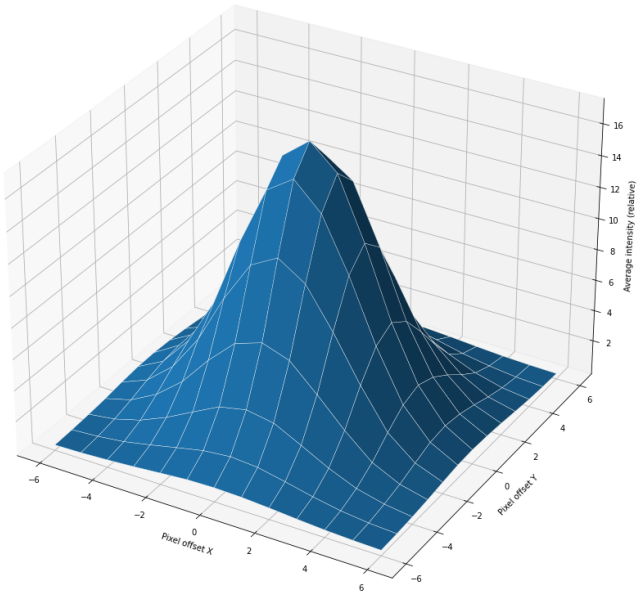

Figure 32: Recovered beam profile from simulated speckly dots. The plotted profile is the mean of around 2500 dots over a single measurement. Observe that the fitting quality is on the level of Figure 27.

Averaging these 2500 measurements yields an estimation of the beam profile, which we know to be Gaussian. The mean of the patches is plotted in Figure 32. A precise match is seen, and the average is close to Gaussian. We only worked with a

single measurement as opposed to the 100 in Experiment 3.2.9. This is due to only adding speckle and no other noises or effects in this test.

Intensity attenuation: Reducing the aperture size results in less light arriving at the imaging sensor. Given a circular aperture, it is easy to see that the total light intensity entering through it is quadratically proportional to the radius. We generated a series of images with perfectly Gaussian dots and added speckle to them with smaller and smaller apertures to test this property.

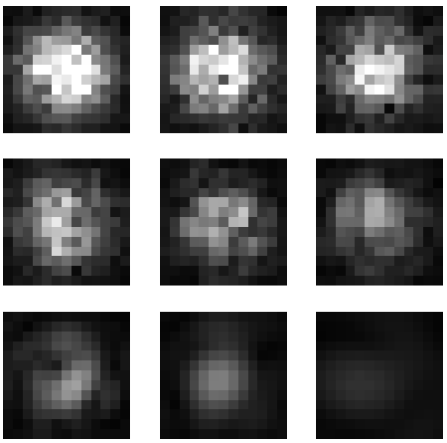

Figure 33: Gaussian dot with added speckle. The aperture is largest on the top-left image and decreases linearly to 20% of it by the bottom-right image. Note that the intensities were Gamma corrected in order to be able to show them in the same scale.

Although the underlying mathematical operations do not guarantee this property, it should naturally emerge in our simulation. Looking at Figure 33, we can see that as the aperture gets smaller and smaller, the details get lost, and the total intensity reduces.

Plotting the total intensity on the speckly images against the relative aperture size reveals a good agreement with a quadratic curve (Figure 34). The simulation, therefore, behaves conformant to reality.

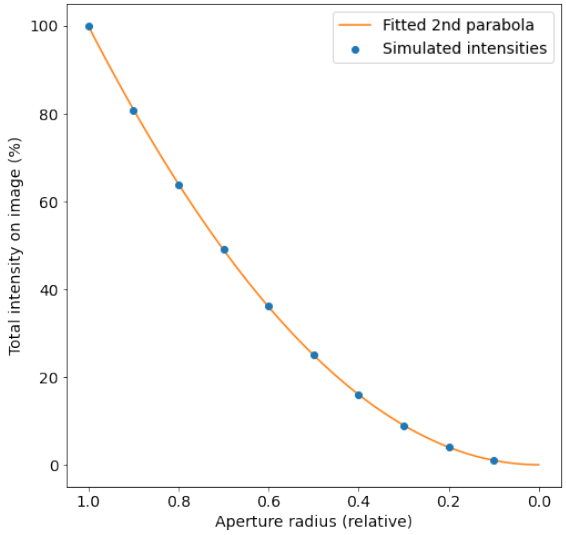

Figure 34: Total intensity remaining when decreasing the aperture size. The intensity from a fully open aperture is taken as 100%. The data points approximate a quadratic curve (orange) well.

5 Conclusion

Our goal was to devise an approach of extending an existing PBR rendering engine with the capability to render realistic IR images. These images should statistically resemble their real counterparts and be generated in real-time. We have shown that the important physical effects that influence the final IR image and make it deviate from ideal can be simulated.

We discovered the different sources that can contribute to the appearance of the final IR image and how they make it deviate from ideal. We identified two significant contributors: the beam splitter DOE and the dynamic speckle.

We thoroughly investigated the theoretical background of both, which turned out to be founded on Fourier optics. To complement the theory, we experimentally evaluated the sensitivity of speckle to changing individual parameters of the setup, and our research led to useful simplifications.

Using the gained insights, we proposed and implemented an algorithm to simulate dynamic speckle and the effect of the DOE in the simulation software of Ladimo Oy. Based on our performance tests, the simulation can be done in real-time on a modern graphics processor.

Finally, we evaluated the quality of the generated images. We compared these images to the ones captured during real-world applications of the technology. Our approach and results satisfy the industrial application it was designed for but can be further extended in the future.

The added realism of IR images is of high value in the industry since it enables the robust testing of depth estimation algorithms, and it makes the collection of high quality ML training data possible.

5.1 Future Work

We have explored the different effects that contribute to the captured final IR image. We found two factors - the dynamic speckle and the beam splitter - significant enough to include in the simulator. As the work at Ladimo Oy continues and we discover alternative applications of the developed 3D sensing technology, we might have to reconsider what effects we pay attention to.

We modeled the laser light as perfectly monochromatic. This assumption might not hold with lasers of lower quality. Currently, there is no plan on using an emitter with a broader wavelength spread. In the future, it might be considered, since that way the emitter can be made to occupy less space. To accommodate such a scenario into the simulation, we could run the simulation steps on multiple wavelengths and combine the results, taking the correlation between them into account. On the one hand, this will result in a smearing of the projected pattern thanks to the way DOEs operate. On the other hand, the combination of partially correlated speckle decreases the speckle contrast.

The current application is strictly indoors, in an office, or home environment. As soon as the technology of Ladimo is taken outdoors, a number of factors come into play. The ambient light might increase significantly. Winds are common in open

spaces. The increased airflow will distort the propagating pattern significantly. If we consider creating a version of the tool that fits commercial phones, we must consider this. Simulating the effect of higher airflow is computationally costly, but it might be worth investigating later.

Additionally, the temperature dependence of the entire system is not handled in the current pipeline. This is partly due to the early stage of the technology. Changing the temperature by a larger amount influences several components. The thermal expansion of the DOE might distort the projected pattern. The laser might emit a different range of wavelengths. Furthermore, the characteristic noise of the camera increases with temperature. Even the lens distortion is a function of temperature. We plan to look into these effects once the company actually needs to address them and the required equipment is available.

Finally, treating out-of-focus measurements poses a significant challenge. Being out of the focal plane results in near field diffraction that cannot be approximated with such a single expression as in the far field. A more costly convolution operation is needed, but the resolution and numerical accuracy must then be considered. Currently, our cameras provide a sufficient depth of field to not suffer much from this phenomenon.

References

[1] Larry Li et al. "Time-of-flight camera–an introduction". In: *Technical white paper* SLOA190B (2014).

[2] Zhengyou Zhang. "Microsoft Kinect Sensor and Its Effect". In: *IEEE Multi-Media* 19.2 (2012), pp. 4–10. DOI: 10.1109/MMUL.2012.24.

[3] Mark Froehlich, Salman Azhar, and Matthew Vanture. "An investigation of Google Tango® tablet for low cost 3D scanning". In: *ISARC. Proceedings of the International Symposium on Automation and Robotics in Construction*. Vol. 34. IAARC Publications. 2017.

[4] JOHN M. COWLEY. "CHAPTER 1, 2, AND 3". In: *Diffraction Physics (Third Edition)*. Ed. by JOHN M. COWLEY. Third Edition. North-Holland Personal Library. Amsterdam: North-Holland, 1995, pp. 3–73. DOI: https://doi.org/10.1016/B978-044482218-5/50004-3. URL: https://www.sciencedirect.com/science/article/pii/B9780444822185500043.

[5] Chen Bar et al. *A Monte Carlo Framework for Rendering Speckle Statistics in Scattering Media*. 2019. arXiv: 1901.06931 [physics.optics].

[6] E. Hecht. "Fraunhofer Diffraction". In: *Optics*. Pearson Education, Incorporated. Chap. 10.2, pp. 465–505. ISBN: 9780133977226. URL: https://books.google.fi/books?id=ZarLoQEACAAJ.

[7] J. W. Goodman. "Some fundamental properties of speckle∗". In: *J. Opt. Soc. Am.* 66.11 (Nov. 1976), pp. 1145–1150. DOI: 10.1364/JOSA.66.001145. URL: http://www.osapublishing.org/abstract.cfm?URI=josa-66-11-1145.

[8] J. W. Goodman. "Statistical Properties of Laser Speckle Patterns". In: *Laser Speckle and Related Phenomena*. Ed. by J. C. Dainty. Berlin, Heidelberg: Springer Berlin Heidelberg, 1975, pp. 9–75. ISBN: 978-3-662-43205-1. DOI: 10.1007/978-3-662-43205-1_2. URL: https://doi.org/10.1007/978-3-662-43205-1_2.

[9] M. FRANÇON. "CHAPTER II - Speckle Produced at a Finite Distance by a Diffusing Object Illuminated by a Laser". In: *Laser Speckle and Applications in Optics*. Ed. by M. FRANÇON. Academic Press, 1979, pp. 21–34. ISBN: 978-0-12-265760-3. DOI: https://doi.org/10.1016/B978-0-12-265760-3.50005-9. URL: https://www.sciencedirect.com/science/article/pii/B9780122657603500059.

[10] Joseph W Goodman. "Introduction to Fourier Optics, Roberts & Co". In: *Publishers, Englewood, Colorado* (2005). http://www.opticsjournal.net/post/ifo.pdf.

[11] Joseph W Goodman. *Statistical optics*. John Wiley & Sons, 2015.

[12] D Duncan and Sean Kirkpatrick. "Algorithms for simulation of speckle (laser and otherwise)". In: Jan. 2008.

[13] Victor Nascov, Cornel Samoilă, and Doru Ursuţiu. "Fast computation algorithms for speckle pattern simulation". In: *AIP Conference Proceedings* 1564.1 (2013), pp. 217–222. DOI: 10.1063/1.4832821. eprint: https://aip.scitation.org/doi/pdf/10.1063/1.4832821. URL: https://aip.scitation.org/doi/abs/10.1063/1.4832821.

[14] Lipei Song et al. "Simulation of speckle patterns with pre-defined correlation distributions". In: *Biomed. Opt. Express* 7.3 (Mar. 2016), pp. 798–809. DOI: 10.1364/BOE.7.000798. URL: http://www.osapublishing.org/boe/abstract.cfm?URI=boe-7-3-798.

[15] L. G. Shirley and Gregory R. Hallerman. "Nonconventional 3 D Imaging Using Wavelength-Dependent Speckle". In: 1999.

[16] Lennart Bargsten and Alexander Schlaefer. "SpeckleGAN: a generative adversarial network with an adaptive speckle layer to augment limited training data for ultrasound image processing". In: *International Journal of Computer Assisted Radiology and Surgery* 15.9 (Sept. 2020), pp. 1427–1436. ISSN: 1861-6429. DOI: 10.1007/s11548-020-02203-1. URL: https://doi.org/10.1007/s11548-020-02203-1.

[17] Yi Chang Shih et al. "Laser speckle photography for surface tampering detection". In: *2012 IEEE Conference on Computer Vision and Pattern Recognition*. 2012, pp. 33–40. DOI: 10.1109/CVPR.2012.6247655.

[18] Andrew K. Dunn et al. "Dynamic Imaging of Cerebral Blood Flow Using Laser Speckle". In: *Journal of Cerebral Blood Flow & Metabolism* 21.3 (2001). PMID: 11295873, pp. 195–201. DOI: 10.1097/00004647-200103000-00002. eprint: https://doi.org/10.1097/00004647-200103000-00002. URL: https://doi.org/10.1097/00004647-200103000-00002.

[19] Joseph M. Schmitt, S. H. Xiang, and Kin Man Yung. "Speckle in optical coherence tomography". In: *Journal of Biomedical Optics* 4.1 (1999), pp. 95–105. DOI: 10.1117/1.429925. URL: https://doi.org/10.1117/1.429925.

[20] Torgny E. Carlsson and An Wei. "Phase evaluation of speckle patterns during continuous deformation by use of phase-shifting speckle interferometry". In: *Appl. Opt.* 39.16 (June 2000), pp. 2628–2637. DOI: 10.1364/AO.39.002628. URL: http://www.osapublishing.org/ao/abstract.cfm?URI=ao-39-16-2628.

[21] GG Romero et al. "Bio-speckle activity applied to the assessment of tomato fruit ripening". In: *Biosystems engineering* 103.1 (2009), pp. 116–119.

[22] Yang Zhao et al. "Point-wise and whole-field laser speckle intensity fluctuation measurements applied to botanical specimens". In: *Optics and Lasers in Engineering* 28.6 (1997), pp. 443–456. ISSN: 0143-8166. DOI: https://doi.org/10.1016/S0143-8166(97)00056-0. URL: https://www.sciencedirect.com/science/article/pii/S0143816697000560.

[23] Wei An and Torgny E. Carlsson. "Measurement of deformation of solder joint during heating with speckle interferometry". In: *Optical Diagnostics for Fluids/Heat/Combustion and Photomechanics for Solids*. Ed. by Soyoung Stephen Cha, Peter John Bryanston-Cross, and Carolyn R. Mercer. Vol. 3783. International Society for Optics and Photonics. SPIE, 1999, pp. 382–388. DOI: 10.1117/12.365759. URL: https://doi.org/10.1117/12.365759.

[24] John G. Abbott and F.L. Thurstone. "Acoustic speckle: Theory and experimental analysis". In: *Ultrasonic Imaging* 1.4 (1979), pp. 303–324. ISSN: 0161-7346. DOI: https://doi.org/10.1016/0161-7346(79)90024-5. URL: https://www.sciencedirect.com/science/article/pii/0161734679900245.

[25] J. C. Dainty. "Stellar Speckle Interferometry". In: *Laser Speckle and Related Phenomena*. Ed. by J. C. Dainty. Berlin, Heidelberg: Springer Berlin Heidelberg, 1975, pp. 255–280. ISBN: 978-3-662-43205-1. DOI: 10.1007/978-3-662-43205-1_7. URL: https://doi.org/10.1007/978-3-662-43205-1_7.

[26] Xu Zhong et al. "Normal vector based dynamic laser speckle analysis for plant water status monitoring". In: *Optics Communications* 313 (2014), pp. 256–262. ISSN: 0030-4018. DOI: https://doi.org/10.1016/j.optcom.2013.10.052. URL: https://www.sciencedirect.com/science/article/pii/S0030401813009802.

[27] Gabriel Laufer. *Introduction to optics and lasers in engineering*. Cambridge University Press, 1996.

[28] Daniel Fleisch. *A student's guide to Maxwell's equations*. Cambridge University Press, 2008.

[29] K. F. Riley, M. P. Hobson, and S. J. Bence. *Mathematical Methods for Physics and Engineering: A Comprehensive Guide*. 3rd ed. Cambridge University Press, 2006. DOI: 10.1017/CBO9780511810763.

[30] Fabrizio Frezza. "Green's FunctionsGreen's Function". In: *A Primer on Electromagnetic Fields*. Cham: Springer International Publishing, 2015, pp. 151–165. ISBN: 978-3-319-16574-5. DOI: 10.1007/978-3-319-16574-5_5. URL: https://doi.org/10.1007/978-3-319-16574-5_5.

[31] Eric W. Weisstein. *Convolution Theorem. From MathWorld—A Wolfram Web Resource*. Last visited on 7/10/2021. URL: https://mathworld.wolfram.com/ConvolutionTheorem.html.

[32] Eric W. Weisstein. *Taylor Series. From MathWorld–A Wolfram Web Resource*. Last visited on 7/10/2021. URL: https://mathworld.wolfram.com/TaylorSeries.html.

[33] G.H. Sendra et al. "Numerical model for simulation of dynamic speckle reference patterns". In: *Optics Communications* 282.18 (2009), pp. 3693–3700. ISSN: 0030-4018. DOI: https://doi.org/10.1016/j.optcom.2009.06.022. URL: https://www.sciencedirect.com/science/article/pii/S0030401809005720.

[34] M. FRANÇON. “CHAPTER I - Speckle in the Image of an Object Illuminated with Laser Light”. In: *Laser Speckle and Applications in Optics*. Ed. by M. FRANÇON. Academic Press, 1979, pp. 1–19. ISBN: 978-0-12-265760-3. DOI: https://doi.org/10.1016/B978-0-12-265760-3.50004-7. URL: https://www.sciencedirect.com/science/article/pii/B9780122657603500047.

[35] M. Fernández-Guasti and M. Heredia. “Diffraction Pattern of a Circle/square Aperture”. In: *Journal of Modern Optics - J MOD OPTIC* 40 (June 1993), pp. 1073–1080. DOI: 10.1080/09500349314551141.

[36] Ronald Newbold Bracewell and Ronald N Bracewell. *The Fourier transform and its applications*. Vol. 31999. McGraw-Hill New York, 1986.

[37] David Whitehouse. “3 - Profile and areal (3D) parameter characterization”. In: *Surfaces and Their Measurement*. Ed. by David Whitehouse. Oxford: Kogan Page Science, 2002, pp. 48–95. ISBN: 978-1-903996-01-0. DOI: https://doi.org/10.1016/B978-190399601-0/50003-7. URL: https://www.sciencedirect.com/science/article/pii/B9781903996010500037.

[38] David Whitehouse. “2 - Identification and separation of surface features”. In: *Surfaces and Their Measurement*. Ed. by David Whitehouse. Oxford: Kogan Page Science, 2002, pp. 16–47. ISBN: 978-1-903996-01-0. DOI: https://doi.org/10.1016/B978-190399601-0/50002-5. URL: https://www.sciencedirect.com/science/article/pii/B9781903996010500025.

[39] David Whitehouse. “1 - Introduction”. In: *Surfaces and Their Measurement*. Ed. by David Whitehouse. Oxford: Kogan Page Science, 2002, pp. 1–15. ISBN: 978-1-903996-01-0. DOI: https://doi.org/10.1016/B978-190399601-0/50001-3. URL: https://www.sciencedirect.com/science/article/pii/B9781903996010500013.

[40] Karl Pearson. *A mathematical theory of random migration*. Dulau and co., 1906. DOI: 10.5962/bhl.title.57440. URL: https://doi.org/10.5962/bhl.title.57440.

[41] Frank Spitzer. *Principles of random walk*. Vol. 34. Springer Science & Business Media, 2013.

[42] Karl Pearson. “The problem of the random walk”. In: *Nature* 72.1867 (1905), pp. 342–342.

[43] Edwin T Jaynes. *Probability theory: The logic of science*. Cambridge university press, 2003.

[44] Ildar Abdulovich Ibragimov and Yurii Antol’evich Rozanov. *Gaussian random processes*. Vol. 9. Springer Science & Business Media, 2012.

[45] David Briers et al. “Laser speckle contrast imaging: theoretical and practical limitations”. In: *Journal of Biomedical Optics* 18.6 (2013), pp. 1–10. DOI: 10.1117/1.JBO.18.6.066018. URL: https://doi.org/10.1117/1.JBO.18.6.066018.

[46] David A. Boas and Andrew K. Dunn. "Laser speckle contrast imaging in biomedical optics". In: *Journal of Biomedical Optics* 15.1 (2010), pp. 1–12. DOI: 10.1117/1.3285504. URL: https://doi.org/10.1117/1.3285504.

[47] J. David Briers and Sian Webster. "Laser speckle contrast analysis (LASCA): a nonscanning, full-field technique for monitoring capillary blood flow". In: *Journal of Biomedical Optics* 1.2 (1996), pp. 174–179. DOI: 10.1117/12.231359. URL: https://doi.org/10.1117/12.231359.

[48] Matthijs Draijer et al. "Review of laser speckle contrast techniques for visualizing tissue perfusion". In: *Lasers in Medical Science* 24.4 (Dec. 2008), p. 639. ISSN: 1435-604X. DOI: 10.1007/s10103-008-0626-3. URL: https://doi.org/10.1007/s10103-008-0626-3.

[49] Valery V. Tuchin, Vladimir P. Ryabukho, and Sergey S. Ulyanov. "Speckle interferometry in the measurements of biotissue vibrations". In: *Holography, Interferometry, and Optical Pattern Recognition in Biomedicine II*. Ed. by Halina Podbielska M.D. Vol. 1647. International Society for Optics and Photonics. SPIE, 1992, pp. 125–136. DOI: 10.1117/12.60198. URL: https://doi.org/10.1117/12.60198.

[50] Wido Heeman et al. "Clinical applications of laser speckle contrast imaging: a review". In: *Journal of Biomedical Optics* 24.8 (2019), pp. 1–11. DOI: 10.1117/1.JBO.24.8.080901. URL: https://doi.org/10.1117/1.JBO.24.8.080901.

[51] David L. Shealy and Fred M. Dickey. "Laser Beam Shaping". In: *Optical Engineering* 42.11 (2003). DOI: 10.1117/1.1624611. URL: https://doi.org/10.1117/1.1624611.

[52] Andrew Bañas et al. "GPC light shaper: static and dynamic experimental demonstrations". In: *Opt. Express* 22.20 (Oct. 2014), pp. 23759–23769. DOI: 10.1364/OE.22.023759. URL: http://www.osapublishing.org/oe/abstract.cfm?URI=oe-22-20-23759.

[53] Brian K. Jennison, Jan P. Allebach, and Donald W. Sweeney. "Direct Binary Search Computer-Generated Holograms: An Accelerated Design Technique And Measurement Of Wavefront Quality". In: *Holographic Optics: Optically and Computer Generated*. Ed. by Ivan Cindrich and Sing H. Lee. Vol. 1052. International Society for Optics and Photonics. SPIE, 1989, pp. 2–9. DOI: 10.1117/12.951480. URL: https://doi.org/10.1117/12.951480.

[54] "Design of binary computer generated holograms by simulated annealing: coding density and reconstruction error". In: *Optics Communications* 94.6 (1992), pp. 491–496. ISSN: 0030-4018. DOI: https://doi.org/10.1016/0030-4018(92)90592-F. URL: https://www.sciencedirect.com/science/article/pii/003040189290592F.

[55] Michael A. Seldowitz, Jan P. Allebach, and Donald W. Sweeney. "Synthesis of digital holograms by direct binary search". In: *Appl. Opt.* 26.14 (July 1987), pp. 2788–2798. URL: http://www.osapublishing.org/ao/abstract.cfm?URI=ao-26-14-2788.

[56] T. Haist, M. Schönleber, and H.J. Tiziani. "Computer-generated holograms from 3D-objects written on twisted-nematic liquid crystal displays". In: *Optics Communications* 140.4 (1997), pp. 299–308. ISSN: 0030-4018. DOI: https://doi.org/10.1016/S0030-4018(97)00192-2. URL: https://www.sciencedirect.com/science/article/pii/S0030401897001922.

[57] Zeev Zalevsky, David Mendlovic, and Rainer G. Dorsch. "Gerchberg–Saxton algorithm applied in the fractional Fourier or the Fresnel domain". In: *Opt. Lett.* 21.12 (June 1996), pp. 842–844. DOI: 10.1364/OL.21.000842. URL: http://www.osapublishing.org/ol/abstract.cfm?URI=ol-21-12-842.

[58] Chunde Huang. *Binarized Gerchberg Saxton Algorithm for Hologram Generation Using a Digital Micromirror Device*. 2020. arXiv: 1910.00658 [eess.IV].

[59] HOLO/OR LTD. *HoloOr website - beam shapers and splitters*. Last visited on 16/12/2021. URL: https://www.holoor.co.il.

[60] HOLOEYE Photonics AG. *HoloEye website - Diffractive Optical Elements*. Last visited on 16/12/2021. URL: https://holoeye.com.

[61] Carlos B. Roundy and Ph. D. Ophir-Spiricon. "Current Technology of Laser Beam Profile Measurements". In: 1999.

[62] Carlos B Roundy and KD Kirkham. "Current technology of laser beam profile measurements". In: *Laser beam shaping* (2014), pp. 463–524.

[63] Fred M Dickey and Scott C Holswade. "Gaussian laser beam profile shaping". In: *Optical Engineering* 35.11 (1996), pp. 3285–3295.

[64] OmniVision Technologies Inc. *Nyxel® Technology Generation 2*. Last visited on 15/01/2022. URL: https://www.ovt.com/technologies/nyxel-technology-generation-2/.

[65] Dr. Rüdiger Paschotta. *Article on Dark Current in the RP Photonics Encyclopedia*. Last visited on 15/01/2022. URL: https://www.rp-photonics.com/dark_current.html.

[66] Ralf Widenhorn et al. "Computation of dark frames in digital imagers". In: *Sensors, Cameras, and Systems for Scientific/Industrial Applications VIII*. Ed. by Morley M. Blouke. Vol. 6501. International Society for Optics and Photonics. SPIE, 2007, pp. 18–25. DOI: 10.1117/12.714784. URL: https://doi.org/10.1117/12.714784.

[67] C. Poynton, Inc Books24x7, and Engineering Information Inc. *Digital Video and HD: Algorithms and Interfaces*. Computer Graphics. Elsevier Science, 2003, p. 260. ISBN: 9781558607927. URL: https://books.google.fi/books?id=ra1lcAwgvq4C.

[68] Craig Matthew. *CCD Data Reduction Guide - Chapter 3, Dark current and hot pixels*. Last visited on 14/12/2021. URL: https://www.astropy.org/ccd-reduction-and-photometry-guide/v/dev/notebooks/03-01-Dark-current-The-ideal-case.html.

[69] Chen Bar et al. "A Monte Carlo framework for rendering speckle statistics in scattering media". In: *ACM Transactions on Graphics* 38 (July 2019), pp. 1–22. DOI: 10.1145/3306346.3322950.

A Varying F number in simulation

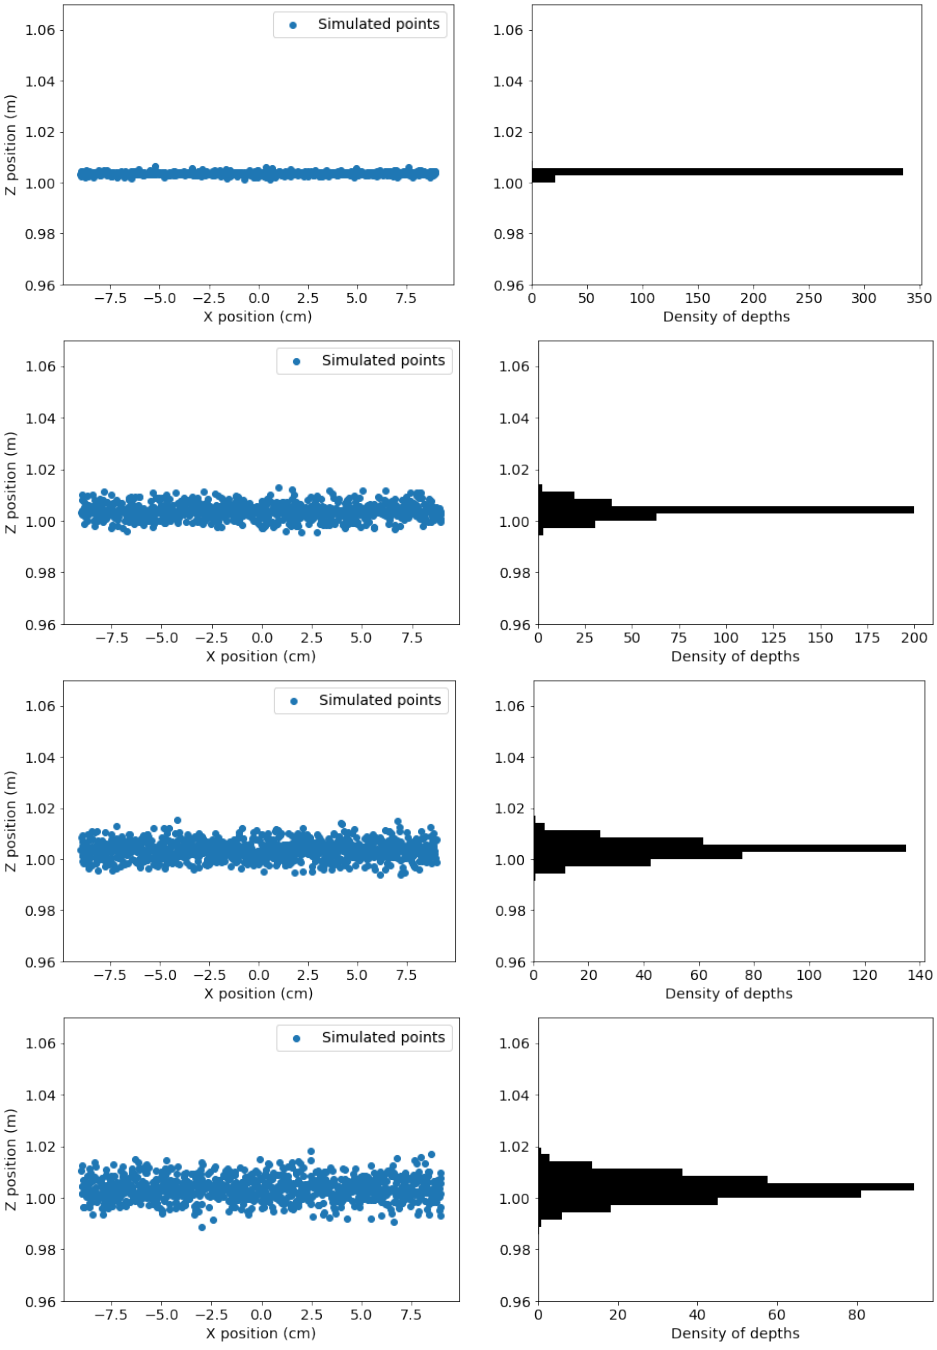

Figure 35: Simulated depth estimations of a plane at about 1 m distance. Points (left) and their depth distribution (right). Four different apertures (top to bottom): f/0, f/2.8, f/4, f/5.6

B Varying F number in measurements

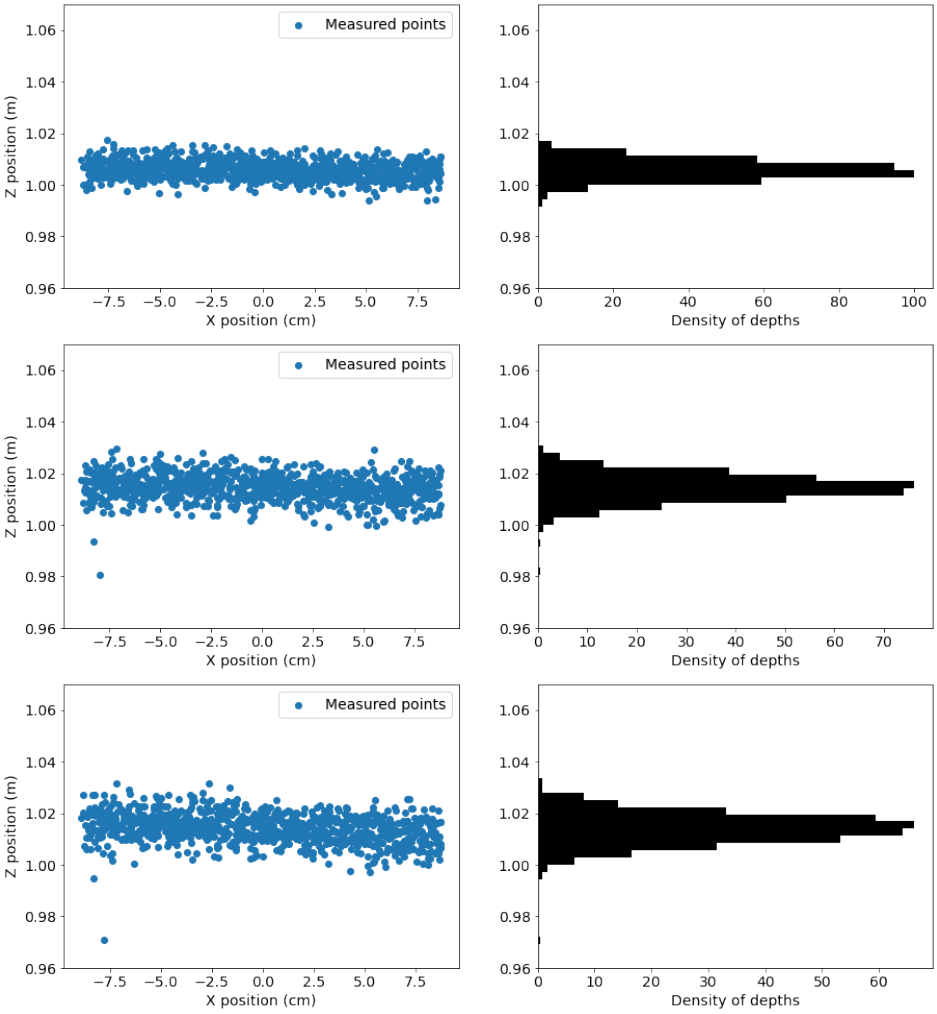

Figure 36: Measured depth estimations of a plane at about 1 m distance. Points (left) and their depth distribution (right). Three different apertures (top to bottom): f/2.8, f/4, f/5.6

C Measurement Setups

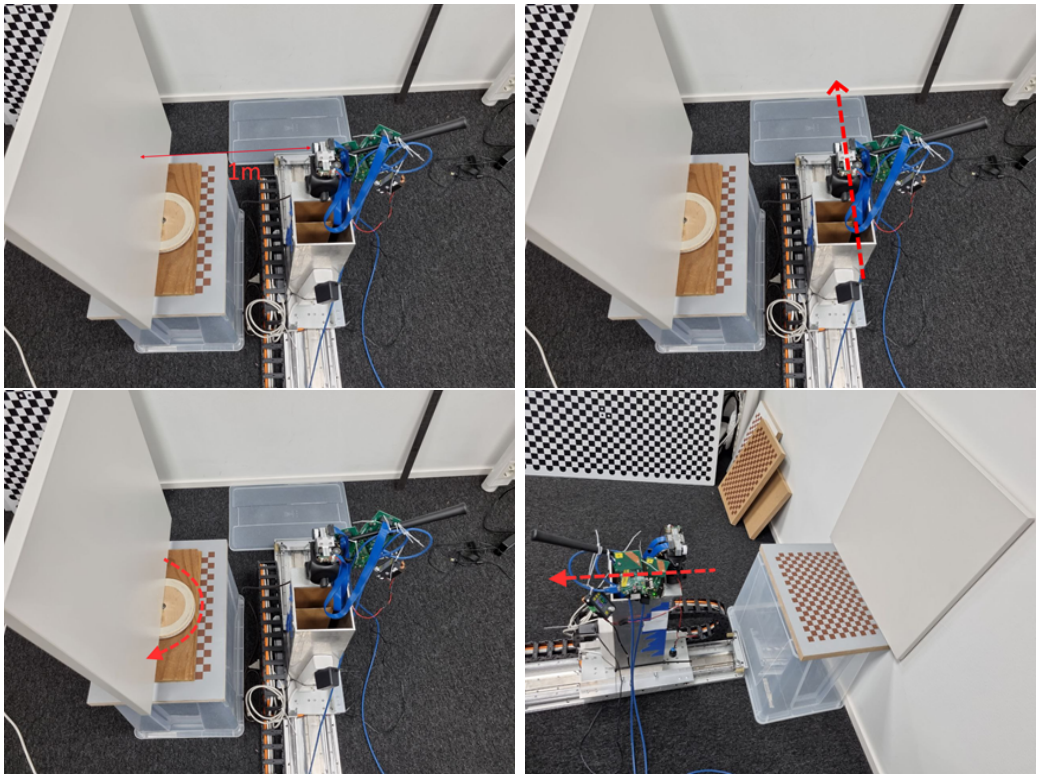

Figure 37: Setups used to take our measurements. The red arrows indicate movement of the object or the camera. Stationary setup (top-left), moving along X-axis (top-right), rotating the object around the Y-axis (bottom-left), and moving the camera along the Z-axis (bottom-right). The lights were on when the pictures were taken but the for the measurements, we turned them off.

D Dot patches with speckle (simulated)

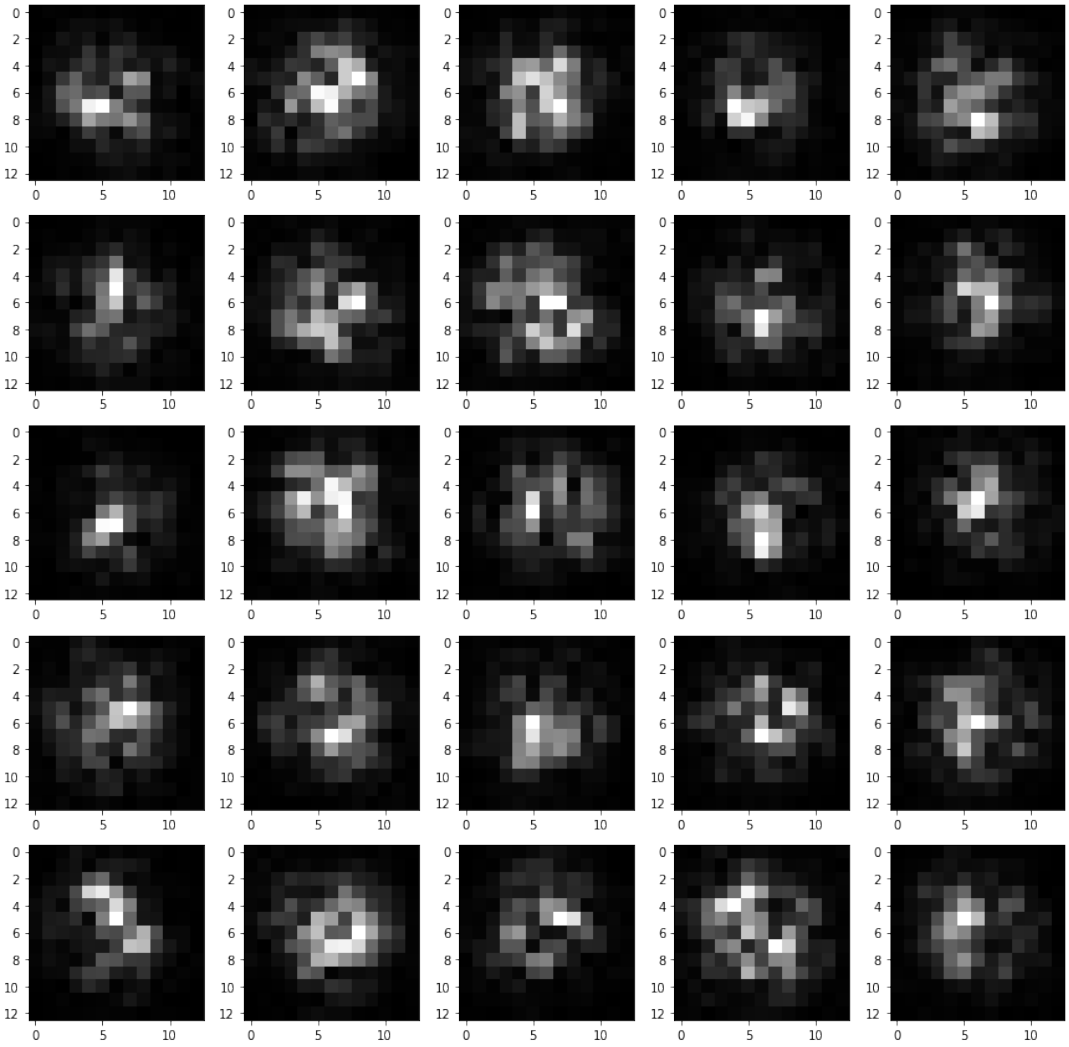

Figure 38: 25 dots from an image with 2500 dots and added speckle. The speckle was generated with an F number of 2.8. The intensities were scaled to [0,1] to show the structure of the speckly dot.

E Pictures with and without speckle from the simulator

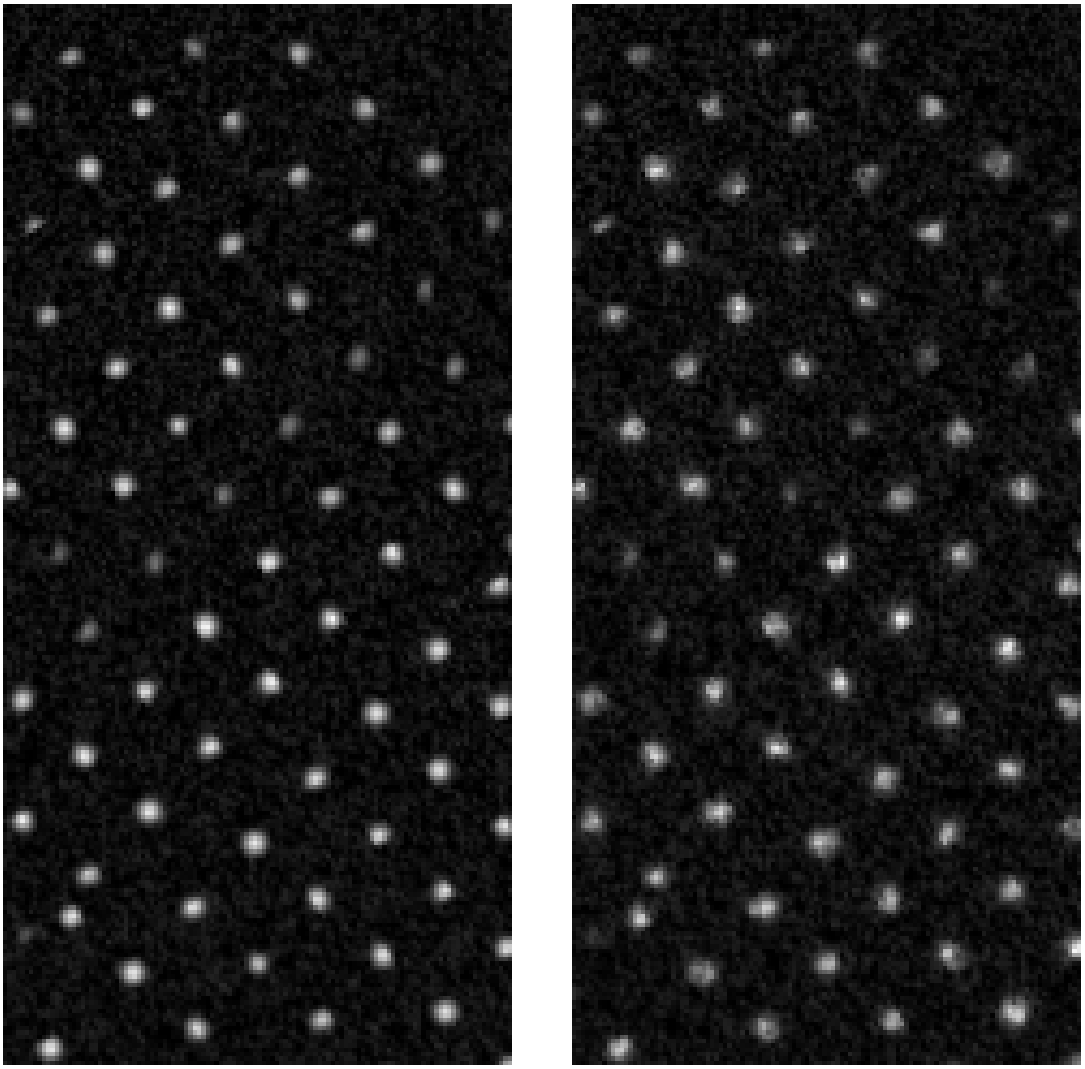

Figure 39: Capture from the simulator. Before adding speckle (left) and after adding a small amount of speckle (right). Note that the additive noises are present in both cases, but the dots are more round on the left image.