

**Digital cameras as radiance
sensors: high dynamic range
imaging and application to
characterize carbon nanotube
depositions**

Master's thesis, 19.6.2024

Author:

SAMI VOUTILAINEN

Supervisors:

MARKUS AHLKOG, ANNA-MARIA RAITA-HAKOLA



UNIVERSITY OF JYVÄSKYLÄ
DEPARTMENT OF PHYSICS

© 2024 Sami Voutilainen

This publication is copyrighted. You may download, display and print it for Your own personal use. Commercial use is prohibited. Julkaisu on tekijänoikeussäännösten alainen. Teosta voi lukea ja tulostaa henkilökohtaista käyttöä varten. Käyttö kaupallisiin tarkoituksiin on kielletty.

Abstract

Voutilainen, Sami

Digital cameras as radiance sensors: high dynamic range imaging and application to characterize carbon nanotube depositions

Master's thesis

Department of Physics, University of Jyväskylä, 2024, 193 pages.

Pairing an optical microscope and a digital camera is a useful tool in observing small samples. However accurate quantitative measurements in terms of relating image brightness to radiance are not possible as is for most cameras due to the nonlinear response of digital cameras. In this thesis we formulate and implement a method to linearize the response of a digital camera and merge multiple images of different exposure times into a high dynamic range image. Additionally imaging system artifacts, such as vignetting and hot pixels, are modeled and corrected for. The performance of the implemented image processing pipeline is quantified, evaluated and finally utilized in the characterizations of carbon nanotube depositions. The depositions are characterized via relative radiance measurements based on the merged high dynamic range images and additionally local thickness and autocorrelation measurements based on both high dynamic range images and scanning electron microscopy images. A theory for the optical properties of the depositions is proposed based on the performed measurements. The results of this thesis indicate that with sufficient calibration and modeling a normal digital camera can be used as a radiance sensor.

Keywords: Carbon nanotube, CNT, high dynamic range, HDR, microscopy, optical microscopy, scanning electron microscopy, SEM, camera response function, CRF, ICRF, linearization, image processing, image analysis

Tiivistelmä

Voutilainen, Sami

Digitaalikamerat radianssisensorina: laajan dynaamisen alueen kuvantaminen ja sovellus hiilinanoputkidepositioiden karakterisointiin

Pro Gradu-tutkielma

Fysiikan laitos, Jyväskylän yliopisto, 2024, 193 sivua

Optisen mikroskoopin ja digitaalikameran yhdistelmä on hyödyllinen työkalu pienikokoisten näytteiden tutkimiseen. Tarkat mittaukset, joissa rinnastetaan kuvien kirkkaus radianssiin, eivät kuitenkaan ole sellaisenaan mahdollisia johtuen digitaalikameroiden epälinearisesta vasteesta. Tässä tutkielmassa muotoilemme ja toteutamme menetelmän linearisoida kameran vaste ja yhdistää useita eri valotusajoilla otettuja kuvia yhdeksi laajan dynaamisen alueen kuvaksi. Lisäksi kuvantamisjärjestelmän virheitä, kuten vinjetointi ja kuumat pikselit, mallinnetaan ja virheitä korjataan. Kuvien prosessointiputken suoriutumista kvantifioidaan, arvioidaan ja lopulta hyödynnetään hiilinanoputkidepositioiden karakterisointiin. Depositioita karakterisoidaan suhteellisten radianssimittausten avulla perustuen laajan dynaamisen alueen kuviin. Lisäksi paikallisen paksuuden ja autokorrelaation mittauksia hyödynnetään karakterisoinnissa perustuen sekä laajan dynaamisen alueen kuviin, että pyyhkäisyelektronimikroskopian kuviin. Mittausten perusteella esitämme teorian depositioiden optisille ominaisuuksille. Tutkielman keskeiset tulokset viittaavat siihen, että tavallista digitaalista kameraa voi käyttää radianssisensorina riittävän kalibroinnin ja mallintamisen pohjalta.

Avainsanat: Hiilinanoptuki, CNT, laaja dynaaminen alue, HDR, mikroskopia, optinen mikroskopia, pyyhkäisyelektronimikroskopia, SEM, kameran vastefunktio, CRF, ICRF, linearisointi, kuvaprosessointi, kuva-analyysi

Contents

Abstract	3
Tiivistelmä	5
I Light and microscopy	12
1 Introduction	12
2 Fundamentals of light and light–matter interactions	15
2.1 Light as an electromagnetic wave	15
2.1.1 Basic properties	15
2.1.2 Coherence	16
2.1.3 Polarization	17
2.2 Light as a particle	19
2.2.1 Photoelectric effect	19
2.2.2 Blackbody radiation	20
2.3 Simple lenses	20
2.4 Interference and diffraction	22
2.4.1 Diffraction gratings	23
2.4.2 Thin film interference	24
2.5 Models of scattering	25
2.6 Reflectance, transmittance and absorbtance	26
2.6.1 Filters	26
3 The optical microscope	29
3.1 General structure	29
3.2 Magnification	31
3.2.1 Compound lenses and infinity-corrected optics	31
3.2.2 Lens aberrations	32

3.2.3	Objectives	34
3.2.4	Oculars	37
3.3	Light delivery	37
3.3.1	Light sources	37
3.3.2	Filters in optical microscopes	40
3.3.3	Collector and condenser lenses	41
3.3.4	Diaphragms	42
3.4	Contrast modes	42
3.4.1	Bright field illumination	42
3.4.2	Darkfield illumination	45
3.5	Olympus BX53M optical microscope	45
3.6	Beyond the diffraction limit	47
4	Samples and basic inspection	49
4.1	Sample fabrication	50
4.2	Sample inspection with different microscopy techniques	51
4.2.1	Optical images	51
4.2.2	AFM images	56
4.2.3	SEM images	60
II	Digital images	65
1	Introduction to digital imaging	65
2	Digital image acquisition	69
2.1	Images as data	70
2.2	Noise	72
2.3	Mean values	73
2.4	Uncertainty values	75
3	The camera response function	77
3.1	Empirical model of the CRF	78
3.2	Measuring linearity	80
3.3	Recovering the ICRF	83
3.3.1	Recovery from exposure series	84

3.4	The ICRF and temporal mean	86
4	High dynamic range images	87
5	The linear effects	91
5.1	Dark current and hot pixels	92
5.2	Non-uniformity	93
5.3	Practical considerations of the flat-field image	95
5.4	Uncertainty	96
6	Other image-based measurements	99
6.1	Autocorrelation	99
6.2	Local thickness	103
7	Full image processing pipeline	105
III	Measurements and results	108
1	Solving the ICRF	109
1.1	The minimization algorithm	109
1.2	Optimal parameters for the linearity measure	110
1.3	Optimal principal components and variability of minimization	114
1.4	Comparison to other data sets	120
2	Olympus SC100 calibration	121
2.1	Distance scale	122
2.2	White balance	124
2.3	Imaging system noise	124
2.4	Linearization and HDR merging	128
2.5	Dark frames and bad pixel filtering	134
2.6	Flat-fields	137
3	Measurements	139
3.1	Relative radiance measurements	139
3.2	Local thickness measurements	145
3.2.1	Local thickness of the optical images	145

3.2.2	Local thickness of the SEM images	153
3.3	Autocorrelation measurements	157
3.3.1	Autocorrelation of the optical images	158
3.3.2	Autocorrelation of the SEM images	160
4	Conclusions	165
	References	168
A	Uncertainty of the HDR merging	177
B	Tables of relative radiance measurements	179
C	Spectroscopic measurements	183
C.1	Sample fabrication	183
C.2	Measurement setup	184
C.3	Measurements	185

List of acronyms

This section contains a list of acronyms used in this thesis and their meaning.

AFM	Atomic force microscopy	MOSFET	Metal-oxide-semiconductor field-effect transistor
BF	Bright-field	MWNT	Multi-wall carbon nanotube
CCD	Charge-coupled device	NA	Numerical aperture
CD	Compact disk	NaN	Not a number
CMOS	Complementary metal-oxide-semiconductor	OD	Optical density
CMYK	cyan, magenta, yellow, key	PC	Principal component
CNT	Carbon nanotube	PCA	Principal component analysis
CPU	Central processing unit	RAM	Random-access memory
CRF	Camera response function	RAW	Raw image file format
DF	Dark-field	RGB	red, green, blue
DVD	Digital versatile disc	RIE	Reactive-ion etching
FF	Flat-field	RMSE	Root mean square error
FoV	Field of view	SD	Standard deviation
HDR	High dynamic range	SEM	Scanning electron microscopy
HMD	Half-maximum distance	SNR	Signal-to-noise ratio
HSL	hue, saturation, lightness	sRGB	Standard red, green, blue color space
HSV	hue, saturation, value	SWNT	Single-walled carbon nanotube
ICRF	Inverse camera response function	UV	Ultraviolet
IR	Infrared	UW	Unweighted
KDE	Kernel density estimation	W	Weighted
LDR	Low dynamic range	WD	Working distance
LED	Light-emitting diode		

Part I

Light and microscopy

1 Introduction

The main subject of interest in this research is the digital imaging of a deposition of carbon nanotubes (CNTs) on a silicon or quartz chip, and the processing of the acquired images. The CNT depositions are a result of drying of a water droplet that was pipetted on a chip with spin coated CNTs on it. The deposition of the prepared samples resembles a coffee ring stain with some special optical properties. The stain is formed from CNTs that are arranged fully two-dimensionally and have two visually distinct zone. The imaging of these zones has motivated the development of digital image processing, which is described in this thesis.

One trend that is undoubtedly evident in the progress of science and technology in the past decades is the ever decreasing length scale at which we view and manipulate materials. We can observe the nanoscale world, for example, via different electron microscopy methods[p. 7][1], atomic force microscopy[p. 1–3][2] and more recently via helium ion microscopy[p. 3][3]. We can fabricate materials in the nanoscale using top-down methods, such as different techniques of lithography[4], laser ablation[5], laser pyrolysis[5] and arc discharge, or bottom-up approach, for example chemical vapor deposition[4], sputtering[5], pulsed laser deposition[6] and spinning[7]. In electronics the scaling of the MOSFET gate length has scaled down from the 10 micrometer range since 1971 to some 5 nanometer range by 2020[8]. This has enabled the production of affordable consumer electronics, increasing computational power and decreasing power consumption. Overall the main driving force behind nanotechnology are the various interesting and different properties materials exhibit, in comparison to their bulk counterpart, when one or more of their dimensions are shrunked down to the nanoscale and the applications these properties can enable[7].

The interesting properties of nanoscale materials arise from the fact that quantum mechanical phenomena are important at this scale[7]. Among the most exciting nanomaterials are the CNTs. In 1991 Iijima reported a method to produce hollow helical microtubules formed from graphitic carbon and he was the first to identify these structures as carbon nanotubes[9]. Synthetic low-dimensional allotropes of

carbon, such as different types of fullerenes, graphene and nanotubes, have been a subject of considerable research interest in the 2000s [10–12]. In the case of CNTs, the main drive behind the research interest are the various extraordinary physical properties, such as high electrical and thermal conductivity and strong elasticity[13]. Carbon nanotubes can be thought of as a graphene sheet rolled into a cylinder. A single such cylinder is a single-walled carbon nanotube (SWNT). The carbon nanotubes used in the samples in this research are multi-walled carbon nanotubes (MWNT), which consists of multiple nested cylindrical layers. CNTs can also be thought of as a 1D material due to its nanometer range diameter and very high aspect ratio.

In the nanoscale world another, relatively new, interesting group of materials are metamaterials. In their review article Kumar et al. describe that metamaterials get their name based on their different physical properties that are seldom encountered in their natural bulk material counterparts[14]. They state that the curious properties generally arise from the carefully engineered physical size, geometry, orientation and arrangement of features. Some well-known types of metamaterials are those with a negative refractive index for electromagnetic[15] or acoustic waves[14] and those with anisotropic thermal properties[16].

In fact, the CNT deposits of this work might exhibit similar optical phenomena as those of 2D metamaterials. In the two distinct zones of the CNT depositions, the distances between the CNTs vary in the wavelength range of visible light, possibly giving rise to the observed optical phenomena due to interference. This acts as an initial hypothesis for the samples' optical properties. One of the objectives in this thesis work is to make observations about the samples mainly with an optical microscope paired with a digital camera, whose output images will be further processed for quantifiable observations. Additionally scanning electron microscopy, atomic force microscopy and spectroscopic measurements are utilized. The primary objective is to enable the calibration of a digital camera in order to acquire images that can act as quantifiable measures of the imaged scene's radiance. Part I of this thesis is focused on the physical fundamentals of light and optical microscopy and the CNT depositions, while part II is focused on the digital imaging and image processing. Finally part III will act as the experimental section of this research, focusing on the measurements and methods introduced in parts I and II.

With all this work there are multiple core questions that we aim to find an

answer to. Can a consumer digital camera be utilized as a radiance sensor and at what level of accuracy? What kind of processing and corrections are needed and what are the limiting factors of the accuracy? What can we deduce on the CNT samples' properties based on the processed images? To our knowledge there are also certain new approaches developed in this work. In part II, section 3.2 we'll introduce a method to measure a camera's linearity based on a fundamental relationship, which is nothing new by itself. However the approach to combine this measure of linearity with a series of differently exposed images of a static scene to formulate an optimization problem for solving the response of a camera is to our knowledge a novel approach. This is discussed in part II, section 3.3.1. We also establish a robust method for analyzing the uncertainty of our image processing pipeline. This includes an analytical equation for the uncertainty of a high dynamic range based on merging multiple normal images, which is something we are not aware has been established in the literature and is discussed in part II, section 2.4. Finally on the subject of image correction, in part II, section 5.2 we propose a very simple, yet effective method to experimentally produce a flatfield image by mean stacking a large quantity of images with local details. This approach is something that we are not aware of having been utilized previously. Generally this issue is an easy one, but the specific conditions of darkfield microscopy and imaging complicate things at least in the sense of experimentally approached correction.

2 Fundamentals of light and light–matter interactions

Nanoscale materials tend to have curious interactions with light due to the wave properties of light and the wavelength-scale structures of nanomaterials and also the particle properties of light and the possible excitations between light and matter. In general the scattering and diffraction of light is a broad subject and the phenomena typically have rigorous mathematical frameworks that are beyond the scope of this research. Here we will approach scattering and diffraction with the context of the nanotube deposition samples in a more qualitative sense.

2.1 Light as an electromagnetic wave

2.1.1 Basic properties

The wave nature of light can be represented as the oscillations of an electric field and magnetic field in mutually perpendicular angles and whose oscillations are perpendicular to the direction of propagation. Mathematically, for a single frequency or wavelength, light can be represented by a wave equation

$$A = A_0 e^{i\omega t - i\theta}, \quad (1)$$

in which A_0 is the real amplitude, ω is the angular frequency, t is time and finally θ is the initial phase[17, p. 642]. The exponential term overall represents the phase of the wave.

The frequency f , or equivalently angular frequency $\omega = 2\pi f$, of light is an invariant property. The frequency is related to the wavelength of light by

$$f = \frac{v}{\lambda}, \quad (2)$$

in which v is the velocity of the wave motion's propagation in a particular medium[17, p. 643]. In the case of light, wavelength is typically used as a common descriptor for

light. In these cases typically what is being referred to is the wavelength in vacuum λ_0 , from which the wavelength in a particular medium be determined via

$$\lambda = \frac{\lambda_0}{n}, \quad (3)$$

in which n is the refractive index of the medium[18, p. 16].

2.1.2 Coherence

Another set of properties that arises from the wave nature of light are its coherence and chromacity. Coherence refers to the phase relationship between beams of light and chromacity refers to the range of contributing frequencies[18, pp. 19–20]. These are important properties in terms of interference phenomena, which we'll get into in the next section. Two beams of light are said to be coherent if there is a persisting relationship between the phases of each beam. A measure of this is the coherence time, which defines a time interval over which the relationship between the beams persists[19, p. 316]. It is approximately given by

$$\Delta t \approx \frac{1}{\Delta f}, \quad (4)$$

in which Δf is the variation of frequencies in the beams. From this one can also define a coherence length, which is the distance traveled within the coherence time. Therefore coherence length is give by

$$\Delta l = \frac{c}{\Delta f}. \quad (5)$$

and can be thought of as the length over which a wave is sinusoidal and can be reliably predicted[19, p. 316]. Table 1 shows the coherence lengths for some light sources. Two highlights from the table are the white light and the special He–Ne laser, which each have a coherence length of under two mean wavelengths and even up to the order of 10^{13} wavelengths respectively, illustrating especially well the high coherence of the laser.

As an idealization, a beam of light is said to be monochromatic if it is composed of only a single frequency. However, in practical terms this is an impossible case as it would require a pulse of infinite duration[18, p. 36]. In reality even a supposedly monochromatic beam, such as a laser beam, is composed of a distribution, though

a narrow one at that, of frequencies.[18, p. 36] Typically a beam is said to be polychromatic if it is composed of multiple reasonably narrow distributions of frequencies. White light refers to a broad continuous spectrum of contributing frequencies.

Source	$\bar{\lambda}$ (nm)	$\Delta\lambda$ (nm)	Δl
Thermal IR	1000	≈ 4000	≈ 25000 nm
Mid-IR	4000	≈ 2000	≈ 8000 nm
White light	550	≈ 300	≈ 900 nm
Mercury arc	546.1	≈ 1.0	≤ 0.03 cm
Kr ⁸⁶ discharge lamp	605.6	$1.2 \cdot 10^{-6}$	0.3 m
Stabilized He-Ne laser	632.8	$\approx 10^{-6}$	≤ 400 m
Special He-Ne laser	1153	$8.9 \cdot 10^{-11}$	$15 \cdot 10^6$ m

Table 1. Approximate coherence lengths Δl and mean wavelengths $\bar{\lambda}$ and wavelength ranges $\Delta\lambda$ of the wavetrain for given light sources. Reproduced from[19, p. 318].

2.1.3 Polarization

Another interesting property of light is its polarization. In the case of light polarization refers to the direction in which the electric field of propagating light oscillates[19, p. 47]. Polarization is a particularly interesting property of light, as it can lead to different interactions of light and matter depending on the polarization and the material in question. For unpolarized light the oscillation direction of the electric field varies randomly over time. Light can, however be also partially polarized. For partially polarized light one polarization direction is more common than other directions.

Polarization can be represented as two mutually perpendicular components by a rotation ϕ with respect to some arbitrary axis as

$$A_x = A_0 \cos \phi \quad \text{and} \quad A_y = A_0 \sin \phi, \quad (6)$$

in which A_x represents the x component and A_y the y component of the amplitude A_0 [18, p. 51] The most simple case of polarization is plane polarized light, for which the electric field oscillates in a single non-rotating plane. For plane polarized light the phase difference between the two components is zero, or more precisely $0 + n\pi$, where n is an integer. When the phase difference between the components is $\pi/2 +$

$n\pi$, the light is circularly polarized. These two cases can be thought of as special cases of the more general elliptical polarization, in which the phase difference doesn't belong in either of the previously mentioned cases.[20, pp. 163–164]

In some materials the electronic configuration can be oriented and therefore direction-dependent interactions between polarized light and the material in question can occur. Such materials are called anisotropic, while materials that don't exhibit directional behaviour are called isotropic. The most common directional phenomena between polarized light and anisotropic materials are birefringence and dichrocity.[18, p. 56] Of these the more common birefringent materials exhibit a direction-dependent variation in refractive index, leading to double refraction. The birefringent material causes an incident ray of light to split into two separate rays.[21, p. 142] In dichroic materials components of light plane-polarized in a particular direction are absorbed, while components plane-polarized in perpendicular with the previous component are transmitted. Dichroic materials are commonly used to produce linearly polarized light.[20, p. 167] Some other interactions of polarized light with matter is optical activity, which refers to the rotation of the direction of polarization for plane-polarized light, optical rotary dispersion, which is wavelength-dependent variation in optical activity, and circular dichroism, which is the variation of absorption of left-hand and right-hand circularly polarized light.[18, p. 56]

Another source of polarization is polarization by reflection. When light is reflected from a boundary of two media with different refractive indexes, the proportion of light that is refracted and reflected depends on the polarization of the incident light, discounting light incident normal to the boundary. For light that is plane-polarized perpendicular to the plane of incidence (s-polarization), the components' intensities are[18, p. 62]

$$I_{s,\text{refl}} = I_0 \left(\frac{\sin(i - r)}{\sin(i + r)} \right)^2 \quad (7)$$

and

$$I_{s,\text{refr}} = I_0 \left(\frac{2\cos i \sin r}{\sin(i + r)} \right)^2. \quad (8)$$

For light plane-polarized parallel to the plane of incidence (p-polarization) the proportions are given by[18, p. 62]

$$I_{p,\text{refl}} = I_0 \left(\frac{\tan(i - r)}{\tan(i + r)} \right)^2 \quad (9)$$

and

$$I_{\text{p,refr}} = I_0 \left(\frac{2 \cos i \sin r}{\sin (i + r) \cos (i - r)} \right)^2. \quad (10)$$

These equations are valid for dielectric materials, metallic materials have typically a different and greater reflectance.[18, p. 62]

2.2 Light as a particle

While the wave model of light can be used to describe a wide variety of phenomena, there are also phenomena that it doesn't describe. In certain situations light also acts as a particle. When light passes through material, the electric field associated with it disturbs the electronic configurations of atoms and molecules. This interaction can lead to absorption of quanta, exciting the electronic energy levels of the absorbing system. While if the quanta is not absorbed, they are instead delayed while they temporarily disturb the electronic configurations of the medium. This represents the slower propagation speed in a medium and thus the medium has a higher refractive index n .

2.2.1 Photoelectric effect

Probably the most well-known phenomenon in which the interactions of light and matter disagrees with the classical electromagnetism is the photoelectric effect. The photoelectric effect refers to the phenomenon when a material emits electrons upon light hitting said material[19, p. 54]. This phenomenon acts as the basis for photoelectric sensors such as digital cameras.

Knight gives a concise description of the photoelectric effect in his book[17, p. 1266–1268]. Electrons occupy specific discrete energy levels in atoms and molecules with an associated binding energy. When a photon with a larger energy than the binding energy of an electron's state hits the system, the electron can be emitted with a kinetic energy of

$$E_k = hf - E_B, \quad (11)$$

in which E_k is the kinetic energy of the emitted electron, E_B is the binding energy of the electron and h is the Planck constant. Basically an electron can be emitted only if the incident photon's energy is sufficiently large, while any excess energy contributes to the emitted electron's kinetic energy. Materials have a number of

varying atomic bindings and generally the maximum kinetic energy of an emitted electron is[17, p. 1266–1268]

$$E_{k,\max} = hf - W, \quad (12)$$

in which W is the minimum energy required to emit an electron, typically referred to as the work function. The work function is characteristic to a given material.

2.2.2 Blackbody radiation

Blackbody radiation describes the electromagnetic radiation emitted by matter as a function of temperature, representing the transformation of a object's internal energy into electromagnetic energy[19, p. 80]. A perfect black body is an idealization, it is an opaque body that absorbs all radiation hitting it and is in thermal equilibrium with its environment[22, p. 689]. It emits a continuous spectrum that is characteristic to the body's temperature. The energy emitted at a wavelength λ by a body at temperature T is given by

$$E_\lambda = \frac{8\pi hc}{\lambda^5} \frac{1}{\exp(hc/kT\lambda) - 1}. \quad (13)$$

Blackbody radiation can be used as an approximation to describe the emissions of astronomical objects or evaluate the temperature of a glowing hot object[23, p. 48]. At room temperatures most of the blackbody emission occurs in the infrared range. Blackbody radiation also explains the continuous emission spectrum of incandescent lamps[19, p. 80].

2.3 Simple lenses

A basic building block for many optical systems are simple lenses, which contain only two refracting surfaces[19, p. 157]. Combining simple lenses, a more complex compound lens can be created, which act as objectives, condensers and other optical components in optical microscopes, cameras and other optical devices. A simple lens is essentially just a single thin lens element with two refracting surfaces. The converging and diverging effects of these lenses can be deduced by applying Snell's

law to the geometry of the lens

$$n_1 \sin \theta_1 = n_2 \sin \theta_2, \quad (14)$$

which relates the refractive indices n of each medium and the angle θ of the incident and refracted light relative to the surface normal[17, p. 757].

Simple lenses come in two varieties separated based on their effect on an incident collimated beam of light[p. 56][21]. Positive lenses cause these parallel beams to converge upon emerging from the lens, while negative lenses cause the light beams to diverge. Positive lenses are convex and form a real image and negative lenses are concave and do not form a real image.




Figure 1. Various types of positive and negative lenses. Reproduced from [21, p. 56].

The lens equation describes the relationship between a lense's focal distance f and an object's and its image's distances a and b respectively. The equation is typically presented in the form

$$\frac{1}{f} = \frac{1}{a} + \frac{1}{b} \quad (15)$$

and the lens system has a magnification factor M

$$M = \frac{b}{a} = \frac{f}{a - f}, \quad (16)$$

from which one can notice some special cases for the magnification factor for certain

values of a [21, pp. 59–60]. Figure 2 illustrates the threshold values for a that yield different magnifications of virtual and real images. The main takeaway is that for




Figure 2. Image formation at increasing object distances a relative to the lens. Modified from [21, p. 60]

object distances $a < F$ a magnified virtual image is formed and for $a > F$ a real image is formed. For $a = F$ the light beams from the object propagate in parallel and no image is formed. Furthermore for the real images with $2F > a > F$ the resulting image is magnified and for $a > 2F$ the image is demagnified, while the point $a = 2F$ results in an unaffected image.

2.4 Interference and diffraction

Probably one of the first things that comes to mind for physicists from the term diffraction are the patterns formed by monochromatic coherent light passing through one or more small slits. The difference in optical paths causes phase differences between rays of light and due to interference gives rise to the maxima and minima observed in a diffraction pattern [21, p. 89–92]. Diffraction in general refers to the phenomenon of lightwaves being scattered and bent by obstacles [21, p. 79]. In more precise terms, in this research we are interested in the far-field or Fraunhofer diffraction. In the case of the optical microscope, the whole image is formed from a diffraction pattern [21, p. 94]. A point in the specimen plane is imaged as an Airy disk in the image plane [21, p. 83], and the viewed image can be thought of as the combination of Airy disks for all points in the specimen plane.

2.4.1 Diffraction gratings

Hecht describes diffraction gratings as a repetitive array of diffracting elements, either apertures or obstacles[p. 488][19]. They are intriguing optical devices and it follows from the previous statement that they can be made from either transmissive or reflective materials. For both types of gratings diffraction arises from the equally spaced grooves that are produced on the surface of the material. Some common everyday objects that acts as diffraction gratings are the different optical discs, such as a CD, DVD and Blu-Ray discs. The optical paths taken by light are shorter at the grooves than at the untouched areas of the material. For a diffraction grating, be that a transmissive or reflective, we have the well-known grating equation[18, p. 46]

$$d(\sin \theta_i - \sin \theta_m) = m\lambda, \quad (17)$$

from which we can solve for the angle of diffraction maxima, which yields

$$\theta_m = \sin^{-1} \left(\sin \theta_i - \frac{m\lambda}{d} \right). \quad (18)$$

In equations 17 and 18 θ_i and θ_m are the angle of incidence and angle of maxima, respectively, relative to the grating normal. d is the grating spacing, λ is the wavelength of the light and finally m is an integer marking the diffraction order. $m = 0$ is the zeroth order diffraction, that is specular reflection, $m = 1$ and $m = -1$ mark the first order diffractions and so on.

From equation 17 you can see that different wavelengths are diffracted to different angles. When white light hits a diffraction grating, each diffraction order will have a continuous spectrum and these spectra of each diffraction order might overlap and make it difficult to observe the different orders. Additionally for diffraction to occur, the spacing of the features of a grating must be larger than the wavelength of the incident light.

2.4.2 Thin film interference

Another phenomenon in which interference plays a central role is thin film interference. Thin film interference refers to the interference effects when light propagates from one medium into another medium with a thin layer of a third medium in between[17, p. 684–685]. A common everyday example is how you can see different colors in a puddle of water with oil floating on top of it, depending on the viewing angle and thickness of the oil layer.

The phenomenon occurs due to how light behaves at a boundary of two mediums. The Slayters give a brief description of this phenomenon by following the beams of light propagating into a thin film[18, p. 40], which we can also follow here. A beam of light is typically split into a reflected and refracted portion. The refracted beam then hits the second boundary at the bottom of the thin layer and is again split into a reflected and refracted beam. The once-refracted and once-reflected beam then again splits when coming back up to the upper boundary of the thin layer. From this beam, the refracted component can then interfere with a beam that is reflected at the first boundary. There will also exist portions of the original beam that have been reflected more than once from the bottom boundary. The different optical paths introduce a phase difference in the beams. Additionally if a beam is reflected from a material of a higher refractive index, it experiences a phase-shift of π . [17, p. 685] The condition for constructive interference is given by[18, p. 40]

$$m\lambda = 2nt \cos r, \quad (19)$$

while the destructive interference condition is given by

$$(1/2 + m)\lambda = 2nt \cos r. \quad (20)$$

In equations 19 and 20 t is the film thickness, r is the angle of refraction relative to the surface normal and m is an integer for the order of interference. In these cases the interference will be incomplete as the intensity of the beams that have undergone reflections and refractions is decreased[18, p. 40].

2.5 Models of scattering

The scattering behavior of light is dependent on a number of things. The scattering event can be either elastic or inelastic. Krishna et al. describe these two processes as follows[24]. In the case of elastic scattering the internal energy of the scattering particle and the wavelength of the light both remain unchanged. In the case of inelastic scattering the wavelength of the light and the internal energy of the scattering light both change due to excitation. Furthermore the size of the scattering particle and the wavelength of the incident light, each independently, lead to different behaviors, as we'll soon see.

In the context of this research, the most relevant distinction for scattering is the size of the scattering particle and the different models used for different particle sizes. There are three domains based on the particle size with respect to the wavelength of the incident light. The size parameter

$$\alpha = \frac{\pi D}{\lambda} \quad (21)$$

is used to determine these zones[25]. In equation 21 D is the diameter of the particle and λ is the wavelength of the incident light. Using the size parameter α , three domains can be identified.

On the larger side of things is the limit $\alpha \gg 1$ for geometric scattering. In this limit the scattering particle is far larger than the wavelength of the incident light, exceeding about 10 wavelengths[19, p. 93]. As the name suggests, in the geometric domain the scattering can simply be modeled by the geometry of the event.

On the smaller side of things is the limit $\alpha \ll 1$ for Rayleigh scattering, which has a strong λ^{-4} [p. 89][19]. In this case the size of the scattering particle is much smaller than the wavelength of the incident light. Probably the most common example of Rayleigh scattering is the blue hue of the atmosphere in daytime. This form of scattering is predominantly elastic in nature and its maximum intensities are in the forward and back directions[19, p. 89].

Finally in the middle of the limits is the $\alpha \approx 1$ range for Mie scattering[25]. In this range the scattering particle is roughly the same size as the wavelength of the incident light. Some common examples of Mie scattering is the color of clouds and milk. The scattering is approximately equally strong for different wavelengths, leading to the white color of the previous two examples.[19]

2.6 Reflectance, transmittance and absorbtance

Stepping back from the micro scale of scattering and diffraction, another set of interesting optical properties of a material are their reflectance R , transmittance T and absorbtance A . Höpe gives a compact definition for these in his work[26]. Reflectance is measures the proportion of a reflected radiant flux against an incident radiant flux. Similarly transmittance measures the proportion of a transmitted radiant flux against an incident radiant flux. Finally absorbtance measures the radiant flux that is absorbed by a material, that is the proportion of the incident flux that isn't transmitted or reflected. For a given material, these values are generally a function of the wavelength, angle and polarization of the incident light and, barring inelastic inelastic scattering and fluorescence, should sum to unity[26]. These material properties are basically the macro scale expression of the scattering and diffraction phenomena occurring in the micro scale on the material's surfaces and inside the material. These properties can be measured as a function of wavelength, angle of incidence and polarization state in a goniometric setup. In such a setup the spectra of the transmitted and reflected light from a sample are measured and compared against the spectrum of the incident light.

2.6.1 Filters

Transmittance and absorbtance are especially important properties for filters, which can be used to modify the spectrum of incident light. Filters come in various types and have different interactions with the incident light. Probably the most basic filter type is the neutral density filter, which applies a flat attenuation of the light's intensity over the spectrum[21, p. 46]. Neutral density filters are characterized by their absorbtance or transmittance, or by optical density, which is defined as $OD = \lg(1/T)$.

An important type of filter are the edge filters, which either absorb or transmit light beyond a certain threshold wavelength. The first one is the short pass filter, which absorbs light with a longer wavelength than the filter's threshold and the second is the long pass filter, which inverses the behavior and absorbs light with a shorter wavelength than the threshold[21, p. 46]. They are characterized by their cut-off and cut-on wavelengths, at which they transmit 50% of the incident light. A high quality filter has high transmission in the pass range, a low transmission

in the block range and a steep curve at the cut-off between peak transmission and minimum transmission[21, p. 46].

The two different edge filter types can be combined to create a band pass or band block filter, which blocks or passes are certain narrow portion of the spectrum. These filters are characterized by their central wavelength and full width at half maximum transmittance or absorbance for band pass and band block types respectively.[21, p. 46]

Polarizers can be used to filter out unwanted polarization states of incident light. In microscopy polarizers are typically used in pairs, with the second one being called an analyzer with the method being known as polarization microscopy[21, p. 153]. The polarizer first transmits only a desired polarization state and the analyzer can then be used to observe the intensity of the transmitted light at different analyzer rotations. Maximum transmission is obtained when the polarizer and analyzer's polarization directions are parallel, while minimum transmission is obtained when the polarization directions are perpendicular. The intensity transmitted through an analyzer is given by the law of Malus

$$I = I_0 \cos^2 \phi, \quad (22)$$

in which ϕ is the orientation of the analyzer relative to the polarizer[21, p. 141]. In the case of unpolarized light incident on a polarizer, the transmitted intensity is just the average value of equation 22 over all angles ϕ , giving[17, p. 1147]

$$I = \frac{1}{2} I_0. \quad (23)$$

3 The optical microscope

The optical microscope is the oldest type of microscope, with the two-lens compound arrangement dating back to 1590[18]. Modern optical microscopes utilize two sets of lenses to achieve spatial resolutions of 350 nm in air and up to 240 nm in immersion oil[21]. The optical microscope's magnification is quite modest, being limited by the resolving power of light, compared to modern electron microscopes and atomic force microscopes. It does have the advantage of being much easier to use and as such doesn't require as much of a training as charged particle microscopes. Additionally it's considerably less costly and basically doesn't impose any additional requirements for the facilities, such as complex maintenance, cooling and vibration isolation stages.

Optical microscopy can be performed by observing either the light transmitted through a sample or the light reflected by a sample. The two methods require a different microscope construction with the reflected light microscope being the slightly more complex one, as the light path in such microscope is partially a two-way path. In the context of this work we are mainly interested in the reflected light microscope but we'll also make some notes on transmitted light microscopes.

3.1 General structure

One doesn't necessarily have to choose just a reflected or a transmitted light type microscope, but instead microscopes with support for both are readily available from manufacturers. These microscopes basically just mash the two microscope types together about the common section that is the objective, binocular head and eyepieces. Figure 3 shows the construction of such a microscope and its important components. The reflected and transmitted configurations have their separate light sources in their own lamp housings. A collector lens is placed in the opening of the lamp housing, as the name suggests, to gather and pass light to the optics farther in the optical column. The light then passes through apertures and diaphragms to align the beam so that all light paths are parallel to each other. This is also the portion of the optical column in which filters are placed, for example polarizers or

color filters. The next part is the mirror block that reflects the aligned light towards the objective and the sample. The objective then focuses the light on the sample and the reflected light is passed into the binocular head and finally focused at the eyepieces for viewing.




Figure 3. Diagram of an optical microscope with both reflected and transmitted light support naming the important components of the microscope. The yellow beam illustrates the path of light from each lamp housing. Reproduced from [27]

As one can observe in figure 3, the main difference between the transmitted and reflected types is in the focusing optics in the upper part of the microscope. For a transmitted light microscope the optical components that focus the light on the sample and guide the transmitted light into the binocular head are separate. For the reflected light type the same optical components that focus the light on the sample also have to focus the reflected light into the binocular head. Additionally when the aligned light beam is reflected towards the sample at the mirror block, the transmitted light type has a mirror used for this. The reflected light type on the other hand must use a half mirror so that the aligned light beam can be reflected towards the sample and the reflected light from the sample can pass through to the

binocular head.

3.2 Magnification

3.2.1 Compound lenses and infinity-corrected optics

Practically all modern optical microscopes are of the compound type. The term compound microscope arises from the fact that at its most basic configuration a compound microscope uses at least two lenses together to achieve the final magnified image. These two lenses are the objective lens near the viewed object and the eyepiece lens which is used to view the final image, whose positions can be seen in figure 3. The propagation of the light beams from the specimen object all the way to the final image is illustrated in figure 4, showing how the compound lenses work together in the most basic configuration. In short, the functionality of the setup can be described as follows[21, pp. 1–2]. The object is placed in front of the objective outside of the focal length of the lens so that a real image is formed in front of the ocular lens. The ocular is placed in a position at which the intermediate real image is within the focal length of the ocular causing the light beams passing through the ocular to diverge. An observers eye forms a real image at the retina from the divergent beams and the object is perceived as a magnified virtual image behind the object. The magnifications of the two lenses add multiplicatively meaning an objective lens with a 50x magnification together with a 10x magnification eyepiece produce a final image of 500x magnification.

The basic compound lens setup is designed assuming that the sample is just outside of the focal distance of the objective when in focus. The intermediate image is then formed at a fixed point and therefore the microscope has fixed distance between the nosepiece opening and the ocular. This distance is the mechanical tube length of the microscope and it was standardized at 160 mm before infinity corrected optics became standard equipment[21, p. 15]. Murphy and Davidson recount the issues surrounding the fixed mechanical tube lengths as follows[21, p. 15]. For a fixed-length mechanical tube any optical accessories, such as filters, illuminators and polarizers for example, causes the mechanical tube length to effectively increase. This in turn introduces aberrations into the optical system. Infinity optics solve this problem with an objective which is designed to project the image to infinity, that is the light beams emerging from the objective propagate in parallel. These parallel




Figure 4. Beam propagation from the specimen object to the eye. The objective lens collects light from the specimen and forms a real image in front of the eyepiece. The eyepiece then forms a real image on the retina together with the eye. The final image is perceived as a greatly magnified virtual image behind the specimen. a) Shows a finite optical train, while b) shows an infinity-corrected optical train and the portion in which the beams propagate in parallel. Reproduced from [21, p. 16].

beams are then focused to form the intermediate image by an additional tube lens, also known as a telan lens. The working principle, as described by the authors, is that any optical components that are designed to produce a 1x magnification maintain incident parallel beams as parallel beams as they pass through the component. With this design a variety of optical accessories can be added into the light path without shifting the focal point of the intermediate image.

3.2.2 Lens aberrations

While a basic compound microscope can indeed be constructed with just two lenses as in figure 4, a high-quality microscope has an objective, which contains multiple lenses to correct for aberrations. Furthermore they can be designed for a specific advanced imaging technique. Using a simple lens to magnify an image incurs a variety of aberrations on the final image that increase in severity with magnification. Figure 5 illustrates the main lens aberrations and why they occur. The lens aberrations, if left uncorrected, severely limit the magnification at which a clear image can be formed. Indeed the goal in designing an objective lens is to limit the size of aberrations to

under the diffraction limited blur[28]. The objective barrels of optical microscopes come in a wide variety in terms of magnification and the magnitude of aberration correction, as correcting aberrations is a trade-off between which corrections are the most important for given application. A more robust discussion and the quantitative analysis of lens aberrations is beyond the scope of this work.

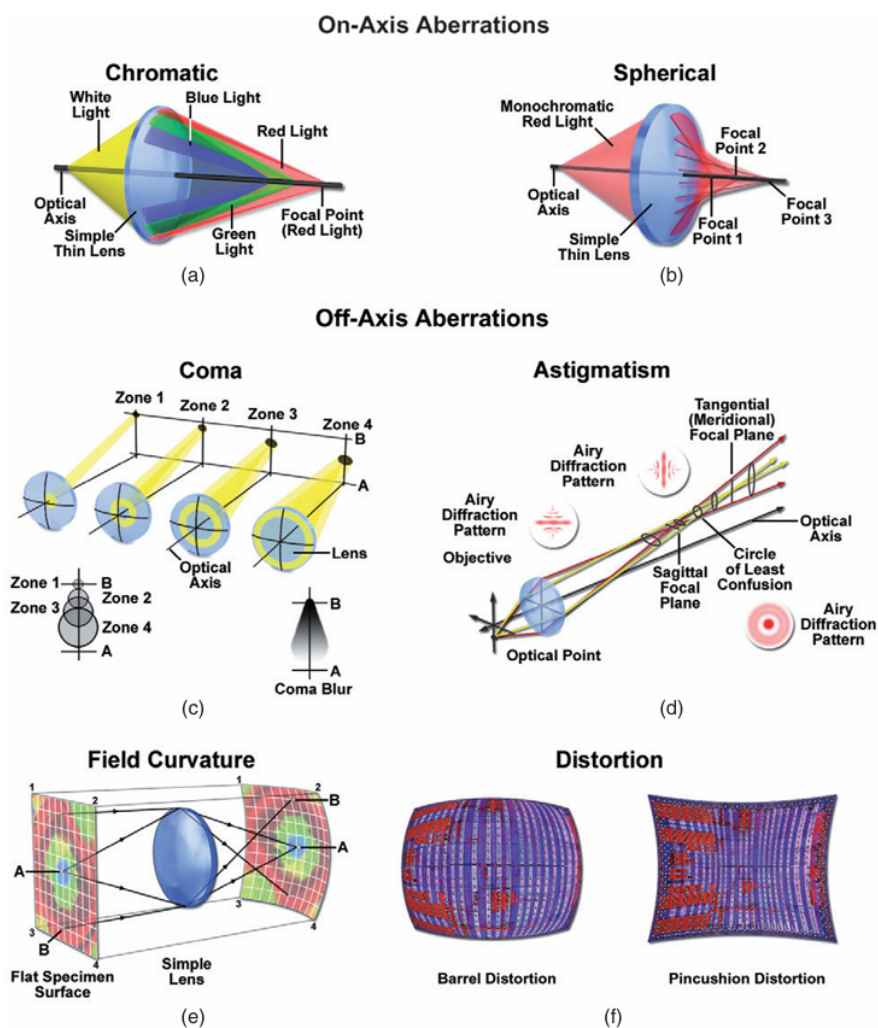


Figure 5. Principle aberrations of a simple lens. a) Light of varying wavelengths focuses at different distances. b) Incident rays parallel to the optical axis refract more the further away from the optical axis they are. c) Rays from an off-axis source focus at different locations depending on the point at which they pass through the lens. d) Rays from an off-axis source passing through the vertical and horizontal diameters have different foci. e) The image plane is in reality curved and projecting it on a plane causes more distortion the further from the true image plane the projected point is. Reproduced from [21].

3.2.3 Objectives

Microscope objectives are characterized by their magnification, immersion medium, working distance (WD), numerical aperture (NA) and type [21, p. 70]. Most objectives are designed to be used in air but in the case of samples under a cover slip, a higher resolving power can be achieved by using a medium with a refractive index close to that of the cover slip's glass. By creating a volume of medium with an uniform refractive index between the cover slip and objective, the distortive effect caused by refraction of light propagating from mediums of different refractive indices can be avoided. The intended medium of an objective is marked on the objective barrel as shown in figure 6. Typically a special type of transparent oil with a refractive index of 1.515 is used but there are also special objective designs, such as water immersion objectives for biological samples in an aqueous medium [21, p. 68].




Figure 6. The markings of an objective barrel. Reproduced from [21, p. 70].

The working distance refers to the distance between the front lens element of the objective, while the sample is in focus, and either the surface of the cover slip on the sample or the surface of the sample itself if no coverslip is used. Typically an increasing magnification requires a decreasing WD [29] and for oil immersion objectives, the WD is restricted to small distances in order to maintain a volume of the immersion liquid between the sample and objective. Special objectives with a

long WD are also designed to enable a wider range of focus. These objectives can be used to view thick samples with features at long distances from the cover slip or sample surface[21, p. 68]. These objectives are prove very useful when the sample needs to be manipulated in some way with an external object, like a micropipette or an external light source for example.

The numerical aperture is a measure of the angle over which the objective can collect light from the sample. The NA is a key parameter in objectives and it affects the spatial resolution and contrast of the image. NA is defined as[21, p. 104]

$$NA = n \sin \theta \quad (24)$$

in which n is the refractive index of the medium and θ is the half angle of the cone formed by sample light accepted by the objective. The image brightness B is defined as [21, p. 71]

$$B \propto (NA/M)^2 \quad \text{in transmitted light mode,} \quad (25)$$

$$B \propto NA^4/M^2 \quad \text{in reflected light mode.} \quad (26)$$

The ultimate resolution of an optical microscope is limited by the diffraction of light caused by the aperture of the objective. In the image an Airy disk pattern is formed by a point in the specimen plane. The radius of the Airy disk pattern in a small angle approximation is given by[21, p. 84]

$$r = 1.22 \frac{\lambda}{2n \sin \theta}, \quad (27)$$

or in terms of numerical aperture

$$r = 1.22 \frac{\lambda}{NA}. \quad (28)$$

Typically the resolution of an optical system doesn't reach the diffraction limit and is instead limited by the previously discussed lens aberrations[18, p. 48].

Microscope objective come in three varieties in terms of differing magnitudes of aberration correction. The first type are achromats, which are red and blue (656 nm and 486 nm) corrected for chromatic aberration, while spherical aberration is corrected for yellow-green (540 nm) light[21, p. 65]. These objectives are considerably

less expensive than the other two types and they excel in monochromatic light. [21, 29] They are best suited at low magnifications, at most 40x [21]. Additionally they have the least amount of lens elements making them a good choice for polarization and differential interference applications [29].

The second type are the semiapochromats, which are made from materials that have low color dispersion and high transmission, allowing for easy chromatic and field curvature correction[21, p. 66]. They are corrected for chromatic and spherical aberrations at blue, green and red wavelengths. These objectives are also a great choice for UV-light, polarization and differential interference applications due to their high transmission and color correction.

The third type are the apochromats, which are the most highly corrected objectives[21, p. 66]. They are corrected for chromatic aberration at four wavelengths: red, green, blue and violet. Additionally the wavelengths between the fully corrected ones have a reasonably good color correction[29]. Spherical aberration correction is applied for green and blue light. The wide range of corrected wavelengths makes them the best choice for working in white light and colors [21, 29]. Apochromat objectives have the largest number of lenses in them, making them less favorable for polarization and differential interference applications [29]. The apochromats are the most expensive of the three types of objectives mentioned here.




Figure 7. Cross sections of achromat, fluorite and plan apochromat objectives. The cross sections illustrate how the number of lenses and the complexity of an objective increases with increasing degree of aberration correction. Reproduced from [21, p. 65].

3.2.4 Oculars

As noted earlier, the oculars work together with the objective to form the final magnified image on the retina of the viewers eye using the intermediate image in the microscope. The ocular's magnification should be some 500-1000 times the objective's NA in order to make the most use of the objective's resolving power while not causing empty magnification[21, p. 72]. Oculars are characterized by their magnification and field of view, which are marked on the eyepiece. For example 10x/0.5 would indicate a 10x magnification and 0.5mm FoV. Oculars have various special designs from additional corrections for curvature of field or chromatic aberration to providing additional features such as a wider FoV or additional spacing for the viewer to wear glasses[18, 21].

3.3 Light delivery

In order to view the specimen at a proper contrast and resolution, the specimen must be illuminated in an even fashion and the light source must be at a proper intensity and correctly focused and aligned. The light source is of course an adjustable lamp, typically a tungsten filament lamp, a type of arc lamp or a halogen lamp. In the case of multiple different available light sources one must also consider the emission spectrum of the lamp for the task at hand.

3.3.1 Light sources

The light sources used in microscopes come in a wide variety with very different emission spectra. Some lamps have distinct very intense emission lines at given wavelengths, some have very neatly continuous spectra and others have both characteristics in varying degrees. Some typical examples of different types of lamps used as light sources are different arc lamps, incandescent lamps, LED lamps and metal halide lamps.

Incandescent lamps are a type of lamp which a lot of us are familiar with in terms of household usage. Incandescent lamps are a common choice for a light source used in microscopes and they are typically tungsten-halogen lamps[30]. They are thermal radiators and therefore the emission spectrum is a continuous spectrum resembling a blackbody radiator as can be seen in figure 8d. Only a small portion of the emitted energy, about 15-20%, is in the visible spectrum with most being in the infrared

range.[30] With increasing voltage the emission is brighter and the spectrum shifts towards shorter wavelengths, making its color balance variable[21, p. 36]. They also provide a bright illumination applicable for even polarized light applications, they are inexpensive and paired with color filters they can provide even illumination.

The second prominent type of light source are arc lamps. These lamps are very bright, varying between 10-100 times brighter than incandescent lamps, making them the lamps with highest luminance and radiance of continuously operating lamps[21, 31]. Arc lamps work by igniting an arc between electrodes with a high voltage followed by maintaining an arc with a lower voltage. Two varieties of arc lamps are used: mercury vapor arc lamps and xenon arc lamps. For the mercury arc lamp only a third of the emission is within the visible spectrum and about half in the UV range[31]. For the xenon arc lamp about one fourth of the emission lies between 400-700 nm wavelengths with about 70% being in wavelengths longer than 700 nm[32]. Mercury arc lamps exhibit multiple distinct emission lines, as can be seen in figure 8 a). The lines can be up to 100 times brighter than the continuous portion of the spectrum[31]. The xenon arc lamp has a nearly continuous spectrum in the visible wavelength range but has complex emission lines in the IR region, shown in figure 8 b). The arc lamps are best used in fluorescence microscopy and with samples that scatter very little light into the microscope objective[31]. The increased brightness provided by the arc lamps does however come with a drawback in being more inconvenient than other types of light sources. The arc lamps have a much shorter lifetime, are harder to align properly and the plasma arc itself flickers and moves between the electrodes causing some instability in the illumination[31]. They are also more expensive than incandescent lamps and require a special power supply due to the high voltage required to ignite the plasma arc initially. Finally the light intensity can't be regulated with an adjustable power supply[21].

On the more modern side of light sources we have the LED lamps. LEDs are semiconductor p-n junctions which emit photons and as such they have, in comparison to the other light sources, a quite narrow emission bandwidth[21, p. 39]. The bandwidth is typically some 20-50 nm wide making it an order of magnitude smaller than incandescent lamps or arc lamps, but still separate from lasers with their 1 nm bandwidths. Figure 8e shows the spectra of different color LEDs. Their advantages over the other lamps are the fact that they have very high degrees of spatial and temporal stability, they have the longest lifetimes and can be switched

on and off at high frequencies[33].

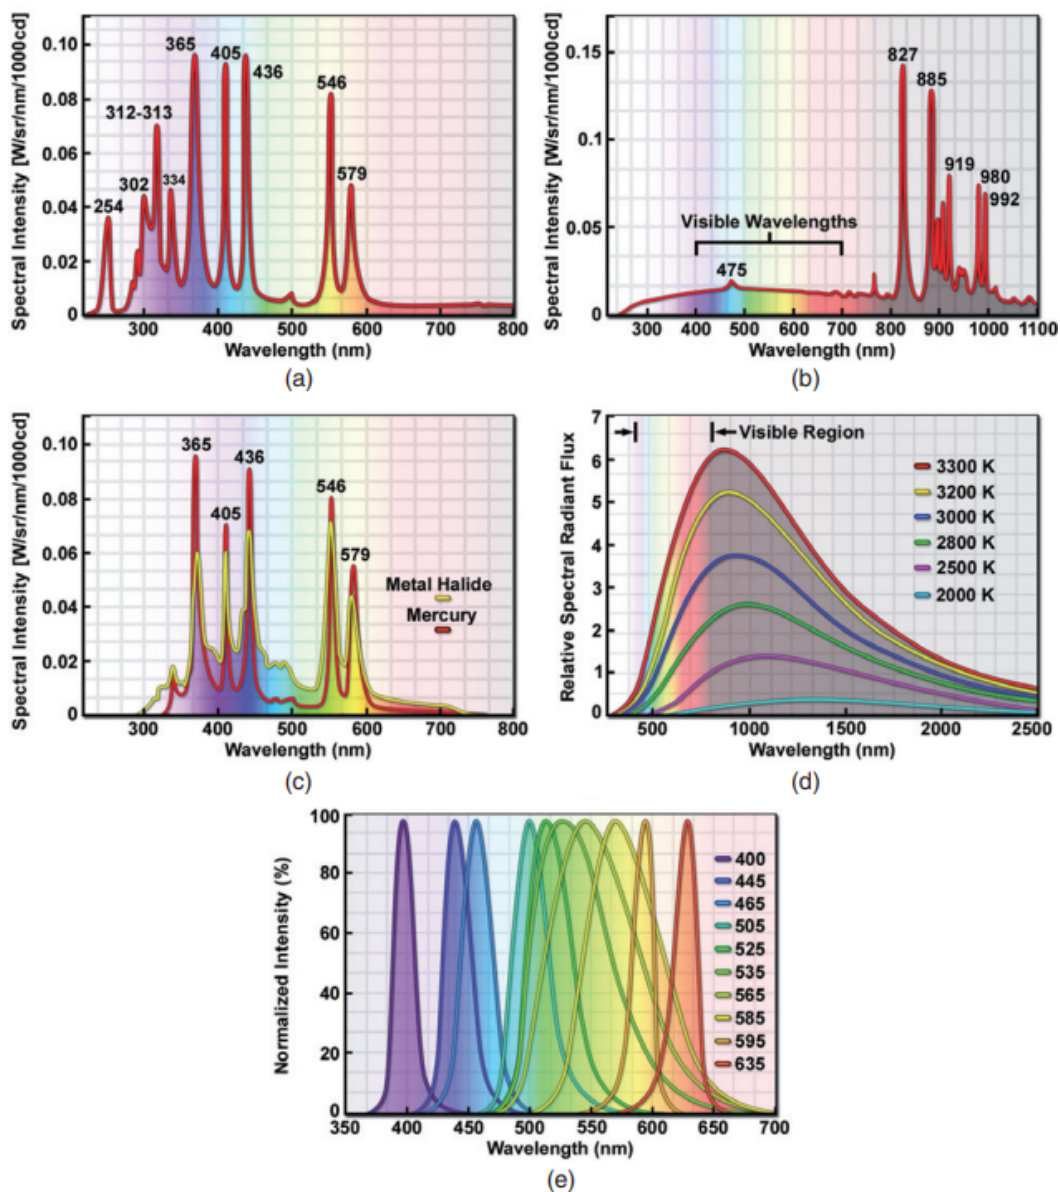


Figure 8. Intensity spectra in various units for different light sources. a) Mercury arc lamp, most of the emissions occur in the IR and UV ranges not shown in the figure. b) Xenon arc lamp, it has similar peaks in the IR range as the mercury arc lamp in the visible range. c) Metal halide lamp, it has very similar emission peaks to the mercury arc lamp. d) Tungsten filament lamp, it has a continuous spectrum with a varying white balance depending on the driving current. e) LED lamps, LEDs can be designed for a specific narrow bandwidth. Modified from [21].

3.3.2 Filters in optical microscopes

The light source type determines the spectrum of the light available but filters can be used to modify the initial spectrum. Not all wavelengths are made equal for all applications and the different types of filters introduced in part I, section 2.6.1 can be used to modify the emission spectrum. Neutral density filters are the most basic type of filter and can be used to attenuate the intensity of light. Color and interference filters can be used to block and pass certain wavelengths.

The neutral density filters are especially useful to pair with arc lamps, whose emission intensity can't be adjusted[21, p. 46]. The attenuation in neutral density filters are based on either absorption or reflection. Reflecting filters must be oriented so that the reflecting side is pointed towards the light source, while absorbing filters don't have a preferred orientation[21, p. 46]. Reflective filters must also be handled with care as the reflecting surface is coated with evaporated metal in order not to scratch the surface.

Davidson and Murphy divide edge filters into two categories[21, pp. 47–48]. These are the colored glass filters and interference filters. They describe that colored glass filters are made with glass that contains rare earth transition elements or colloidal colorants. Interference filters on the other hand are made by stacking multiple layers of glass coated with transparent metal salts as a dielectric material. They state that the thickness of these layers is designed to be one half or one quarter of wavelength in thickness in order to facilitate constructive interference for specific wavelengths[p. 47][21]. The colored glass filters work via absorption and the authors note that as such their transmission properties can change due to heat. They also highlight that the colored glass filters are more resistant to scratching and chemicals than the interference filters, similar to the absorbing and reflecting neutral density filters respectively. A last feature of the colored glass filters, that we'd like convey from Davidson and Murphy's work, is that they are insensitive to the angle of the incident light, are less expensive and have a longer lifetime than the interference filters. The interference filters do however typically have a steeper cut-in and cut-off curves[21, p. 47].

The different edge filters are very useful for situations where a portion of the spectrum needs to be blocked. For example in fluorescence microscopy only the specific wavelength range emitted by a fluorescent material is desirable to observe[21, p. 45]. Considering the nature of the different filter types, it is evident that edge filters

are the only way to separate specific ranges of wavelengths for direct observation or image capturing.

Polarizers are very useful, special type of filter that can produce polarized light[21, p. 136]. They can be used in observing any optically anisotropic properties of a material and are fundamental in polarized microscopy, in which the contrast in the image is formed by the differences in intensity for different polarization states[21, p. 152]. Modern polarizers are created from dichroic materials as polaroid sheets, in which the oriented macromolecules absorb light plane-polarized in the filter's absorption direction. We discussed dichroicity briefly in section part I, 2.1.3.

3.3.3 Collector and condenser lenses

The light emitted by the lamp in the lamp housing must be focused into a even collimated beam. The collector lens at the opening of the lamp housing serves this purpose, collecting as much light as possible from the light source. Additionally the collector lens is used to focus an image of the light source in a specific position to achieve Köhler illumination[21, p. 9]. Köhler illumination is discussed more in part I, section 3.4.1 The lens that is used to focus light from the illuminator onto a small area on the sample is the condenser lens. As can be seen from equation 25, the image brightness has a M^{-2} relation to magnification, meaning brightness must increase as the square of magnification in order to maintain a constant brightness under increasing magnification. This makes the condenser lense's role very important in illumination. In a transmitted light microscope the condenser is a separate lens underneath the sample, while in a reflected light microscope the objective acts also as a condenser lens for light propagating down towards the sample.

Like the objectives and oculars, condenser also have different degrees of aberration correction in their designs. The very common Abbé condenser for example suffers from chromatic and spherical aberrations and curvature of field[21, p. 71]. Better designs such as the aplanatic condenser corrects for spherical aberration and curvature of field, while the aplanatic-achromatic corrects for even chromatic aberration in addition[21, pp. 71–72].

3.3.4 Diaphragms

As seen in figure 3, the beam column has two diaphragms, the aperture and field diaphragms, in the light path. The first one in the light path is the aperture diaphragm, which controls the amount of light passed from the collector lens further into the light path. These diaphragms together allows the user to adjust the contrast and resolution of the image[21, p. 9]. They have an optimum setting for maximizing these parameters; too open an aperture causes glare and a poor contrast, while a too closed an aperture causes increased diffraction and poor resolution.

3.4 Contrast modes

There are a variety of illumination and contrast formation techniques with each of them based on a different interaction between light and the sample. The various interactions provide a different contrast formation mechanism and thus in the final image a different contrast, highlighting details unique for each technique. While there are a large variety of techniques, probably the most common examples are the bright field and dark field illumination techniques. In these illumination methods the contrast originates from the absorbtance of light and scattering of light respectively. Other examples are cross-polarized light and phase contrast illumination, with the contrast originating from the rotation of polarized light and interference of varying path lengths of light respectively.

3.4.1 Bright field illumination

The most basic bright field (BF) lighting method is the critical illumination method, in which an image of the light source is focused on the specimen. Naturally this means that the image of the light source is visible in the final image leading to uneven illumination and reduced image quality. Instead basically all scientific optical microscopes apply Köhler illumination, whose design was invented by August Köhler and published in 1893[21, p. 1]. In Köhler illumination the idea is to ensure even lighting on the specimen plane and completely defocus the image of the light source at the retina. This is achieved by specific placement of lenses. A diagram of both critical and Köhler illumination methods is shown in figure 9.

In Köhler illumination there is a collector lens near the light source housing, which as the name implies, is used to collect and direct as much light as possible




Figure 9. Diagram of Köhler and critical illumination setups. a) The conjugate field planes, note the focal points at the sample stage and the retina. b) The conjugate aperture planes, the focused image of the light source is defocused on the sample stage and at the retina. c) Critical illumination setup, the image of the light source is focused both at the sample stage and retina. The noticeable difference between the setups is that in Köhler illumination there is an additional collector lens and diaphragm in front of the light source. Modified from [21, p. 8].

toward the specimen. Adjacent to the specimen is a condenser lens, which is used to specifically focus most of the light coming from the collector lens on a small area on the specimen. Additionally in order to avoid stray light from polluting the final image and to maximize beam alignment there are diaphragms placed along the beam path.

As noted earlier, the contrast formation mechanism in BF illumination is the absorption of light in the sample. Be it a reflected or transmitted light microscope, a portion of the light is absorbed by the sample. The absorbance is dependent on the material of the sample, the angle at which light hits the surface at a point on the sample and especially in the case of transmitted light, the thickness of the sample. At least some of these properties must vary on the sample in order for any observable contrast to arise in the final image.




Figure 10. Bright field image of standard paper at 20x magnification.

The BF method is the simplest illumination method and is thus typically utilized in most optical microscopes. The illumination method forms a bright background in the image and the varying amount of absorbance gives the final image its contrast from white fully reflected light to black completely absorbed light. It follows that this makes BF illumination a poor choice for very homogeneous samples and great at viewing samples with lots of surface detail in reflected light microscopy. BF illumination is especially problematic in thin biological samples in which the sample has barely any absorbance. In such a case samples can be stained to provide color and contrast[21, p. 115]. Figure 10 shows an image of paper at 20x magnification under bright field illumination. Some detail can be resolved from the overlapping fibers of the paper, but overall there is little sense of depth in the image due to the bright background.

3.4.2 Darkfield illumination

The darkfield (DF) illumination technique is likely the second to most common technique of illumination. The basic idea of DF illumination is to have the illuminating light hit the sample at an oblique angle so that only light that is scattered by the sample enters the objective, while light that is directly transmitted or reflected misses the objective lens entirely due to the oblique angle[21, p. 130]. The name for the technique arises from the fact that the image has a dark background and any details in the image are bright due to the light scattered by the sample material.

The typical approach to allow light to hit the sample only at oblique angles is to block the central portion of the beam with a darkfield condenser annulus[21, p. 130]. For a transmitted light microscope this is achieved with a disc shaped patch stop at the center of the beam allowing only an outer ring of light to pass towards the condenser lens. For a reflected light microscope the mirror block uses a fully silvered mirror with an elliptical hole in the center to only direct an outer ring of light towards the objective[34].

Darkfield illumination excels in showing small details in samples that would be hidden under bright field illumination due to the bright background. In DF mode a greater sense of depth is given to the images as the background is dark. Figure 11 shows an image of paper at 20x magnification in DF mode. From the image it is evident that a greater sense of depth is achieved, as multiple layers of overlapping fibers can be resolved from the image. The fibers appear more transparent in the DF image and less detail can be resolved at the surface of individual fibers than in BF mode.

3.5 Olympus BX53M optical microscope

In this research an Olympus BX53M optical microscope with a frame supporting both reflected and transmitted light observation was utilized. Of these only the reflected light observation was utilized. The light source used was a LED light with a reported color temperature of 5700 K, corresponding to a white light with a slight blue hue. Objectives of various magnifications supporting DF illumination and transmitted and reflected observation were utilized in this research and they are listed in table 2. All the objectives are planar achromats. The microscope is additionally equipped with a fixed polarizer and rotatable analyzer allowing polarized microscopy. The




Figure 11. Darkfield image of standard paper at 20x magnification.

microscope is also equipped with an Olympus SC100 CMOS color camera, though we'll get more into this camera in part III.

Magnification	NA	WD (mm)	Max FoV (mm)
5x	0.10	12.0	4.40
10x	0.25	6.5	2.20
20x	0.40	1.3	1.10
50x	0.75	0.38	0.44
100x	0.9	0.21	0.22

Table 2. Properties of the utilized objectives



Figure 12. A photograph of the Olympus BX53M microscope utilized in this research.

3.6 Beyond the diffraction limit

In this research in addition to utilizing the optical microscope, the atomic force microscope (AFM) and the scanning electron microscope (SEM) were also utilized. These are very useful tools in observing the nanoscale world that is beyond the diffraction limit of conventional optical microscopes and they require at least a brief discussion before delving into AFM and SEM images in the following sections.

SEM utilizes the interactions of electrons and matter and magnifications between tens of times to hundreds of thousands of times can be achieved[1, p. 19]. Maximum resolutions typically are in the nanometer range but sub nanometer can also be achieved[1, p. 10]. SEMs have various applications ranging from image formation based on different signals to different X-ray spectroscopy methods[1, p. VII]. In the case of this research the main function of SEM is to simply provide structural information about the samples that would be impossible to achieve with the resolution of an optical microscope. Emphasis is given on the ability to observe large areas at a

high resolution.

The AFM utilizes a micrometer scale tip attached to a cantilever to measure its deflection as it is scanned over a sample surface[2, pp. 8–10]. This enables the measurements of a material’s mechanical properties as well as image formation at resolutions down to a sub nanometer range[2, pp. 3, 255]. AFMs are useful at obtaining topographic data of sample surfaces, especially in terms of measurable height or depth[2, p. 5]. In this research AFM images are utilized in a similar fashion as the SEM images; a means to achieve information otherwise unobtainable via optical microscopy, but with an emphasis given on the height data of small areas.

By taking careful note of the position at which an AFM or SEM image is taken, it is possible to match them with optical images of the the same position. From these matched image pairs it is easy to take note of the resolution limit of optical microscopes. Such an image pair of a silicon chip with fabricated electrodes and MWNTs is shown in figure 13. The imaged sample is shown just to demonstrate the resolution limit of an optical microscope and how small particles are seen in a DF optical image, otherwise the sample bears no relevance to this research.




Figure 13. An optical DF image and AFM peak force error image of a silicon chip with fabricated electrodes and deposited MWNTs. The particles are tagged with a letter and a number and are matched between the two images. From the image pair it is clear to see that the individual MWNTs are too small to be observed in the optical image.

4 Samples and basic inspection

As was stated in the introduction, the main interest in this thesis is the imaging of the special kind of CNT depositions, that have two optically distinct zones. These zones can be observed in figures 14 and 15. Of the two zones, the outer one appears visually dimmer in DF images, while the inner zone appears brighter. In this section we briefly describe the fabrication of the samples and their characterization with microscopic imaging and spectroscopic techniques, that are straightforward in a modern laboratory such as the Nano Science Center of the University of Jyväskylä. This section will conclude part I. In part II we will concentrate on developing methods of digital image processing of optical microscopy data to obtain more information of the sample. The developed methods will finally be utilized in part III for quantitative measurements.




Figure 14. BF images captured of one of the CNT deposition samples used in this research. a) 5x magnification, b) 20x magnification. In BF mode it is nigh impossible to see detailed structure at the top of the web like formations.




Figure 15. DF images captured in the same positions of the same sample as in figure 14. a) 5x magnification, b) 20x magnification. DF mode gives a greater sense of depth in the case of this sample, as details on top of the web like formations can be seen.

4.1 Sample fabrication

The three main samples of this work, that have both zones, were produced by Esa Hyyryläinen in 2018, within his M.Sc. thesis work, according to the following procedure. Firstly, silicon chips were cleaned and the surface made hydrophilic with an O₂ plasma treatment via RIE. Secondly a layer of carbon nanotubes, immersed in dichloroethane, were deposited on the chips' surfaces by spin coating. This leaves a uniform and thin, practically two-dimensional, layer of nanotubes on the chip. Then in an environment with controlled humidity a droplet of water is pipetted on the chip. The formation process of the deposition is illustrated in figure 16 along with a photograph of the main samples. As the pipetted drop expands it sweeps away nearly all carbon nanotubes that are not parallel to the flow direction of the water. As a hydrophobic material, the swept nanotubes then float to the surface of the drop. The water drop is then allowed to evaporate, which causes the surface area of the drop to shrink. This leaves less and less area for the floating CNTs as the drop evaporates and eventually when the surface area is small enough in relation to the area taken by the CNTs, a deposition with two optically distinct zones is formed. The deposition takes a curious form that resembles a dried coffee stain. The samples utilized in this research are listed in table 3 alongside with some basic information and DF and BF images. In the appendix A are some more details about the sample fabrication process in the context of the spectroscopic measurements.




Figure 16. Illustration of the sample fabrication process from the spin coated chip to the dried deposition and a photograph of the main samples. On the silicon substrates the depositions' resemblance with a dried coffee stain is evident. On the quartz chip the deposition appears more subdued. a) Shows the initial spin-coated surface. b) shows the pipetted water droplet expanding on the chip surface and sweeping material on top of the droplet. c) shows as the droplet dries by evaporation and leaves a shrinking area for the swept material. d) shows the final deposition after the water has completely evaporated.

4.2 Sample inspection with different microscopy techniques

There are various observations that can be made from basic optical, SEM and AFM imaging. In the case of the optical images as they are, we can mostly make qualitative observations of the samples, but quantitative analysis will require further processing of the image data and is left for part III.

4.2.1 Optical images

The optical images are formed from the combination of red, green and blue color channels. The intensity of the images of each channel are directly proportional to the irradiance the camera sensor receives, that is how many photons impinge the sensor per time unit. By controlling the exposure time the brightness of the image can be adjusted and the color channels can also be observed separately. We'll get

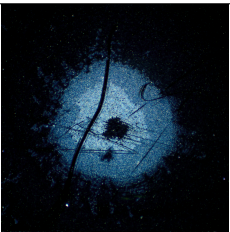
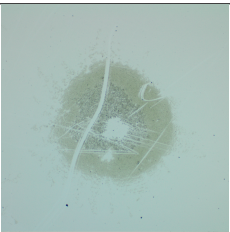
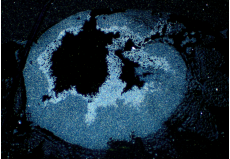
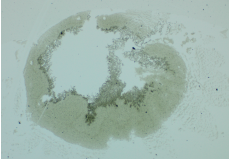
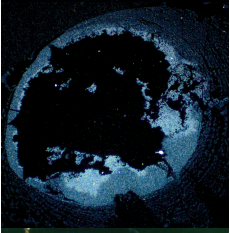
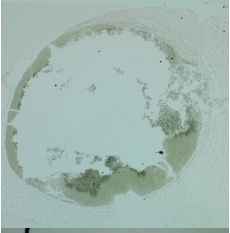
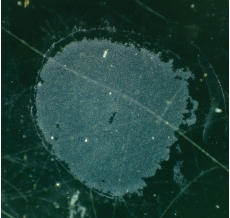
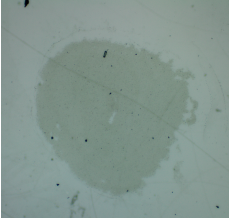
No.	By	Date	Substrate	DF	BF
1	EH	29.6.2018	Silicon		
2	EH	29.6.2018	Silicon		
3	EH	29.6.2018	Silicon		
4	SV	20.1.2023	Quartz		

Table 3. Table of the utilized samples listing their number, creator (EH = Esa Hyyryläinen, SV = Sami Voutilainen), fabrication date, substrate material and DF/BF images.

more into the details of digital imaging in part II.

There are a couple of features in the optical images that are likely to first pique a viewer's interest in the DF and BF images. As we showed in figures 14 and 15, when we observe the sample with an optical microscope in BF and DF modes, firstly it shows that two visually distinct zones have formed: an outer and an inner ring. In BF mode the outer circle appears brighter than the inner one and vice versa in DF mode. Secondly the striking blue hue of the DF images commands attention. The blue hue can in reality be slightly misleading without further inspection, but we'll discuss that more a bit later.

First we'll take notice on the differences of the BF and DF mode images. On one hand the BF mode images excel at showing where MWNT and amorphous carbon

material are, which can be seen as dark material against a very bright unobscured substrate. This distinction isn't as clear in the DF images due to a peculiar blue glow observed at the edges of deposited material. The DF mode on the other hand very clearly shows more detail in the deposition's structure and gives a sense of depth in the deposition.

Something interesting can be noticed in the DF images when capturing an image at a short exposure time. In figure 17 there are two DF images taken at a longer and a shorter exposure time and two BF images, all taken at the same position. This is where the blue hue can be observed to be slightly misleading. The shorter exposure DF image's values are scaled up for easier viewing. Firstly in the BF image we can take note of where the deposited material lies. Then observing these positions in the bright DF image, we can notice that the sample has a strong blue hue as seen in previous DF images. Additionally in this closeup we can observe that there is indeed a blue glow along the circumference of the deposited material. This glow is weaker in the darker DF image, in which we can observe that at the actual positions of the deposited material, there is a clear green hue. This could suggest that the deposited material's actual color is more of a green hue, while the blue color could originate from interactions of the deposited material and the silicon substrate. Utilizing an analyzer and polarizer pair in BF mode provides additional insight. The third image in figure 17 shows the same position captured in BF mode with an analyzer that is rotated to filter out the light that is reflected from the substrate. This shows a similar blue hue as in the DF images, which implies that the blue hue isn't dependent on the light's angle of incidence relative to the sample. We'll return to these ideas again in part III with quantitative tools.

Figure 18 shows the same sample as in figures 14 and 15 with additional SEM images taken from both of the sample's zones. The thin mostly straight strips of varying thickness are MWNTs that are free of clutter. Most of the MWNTs are covered in amorphous carbon particles that accumulate on the MWNTs. These amorphous particles are mostly seen aggregated into clusters and attached to the MNTNs. Very few amorphous carbon particles are seen independently. The SEM images also show a visibly clear difference in the amount of material in each zone. This is further illustrated in figure 21, which shows images of the dim and bright zones respectively. The key feature we want to highlight is that the bright zone SEM image has 3000x magnification, while the dim zone SEM image has 8000x




Figure 17. Optical images of sample No. 2 taken at 50x magnification. a) DF image with 30 ms exposure, b) 5 ms exposure with values multiplied by three for viewing, c) BF image with analyzer–polarizer pair for filtering reflected light from substrate, d) regular BF image.

magnification, and yet the images look relatively similar to each other. Considering that in DF mode an optical microscope image should be brighter at areas with more non-flat features, the sparser area being the brighter one is counter intuitive.

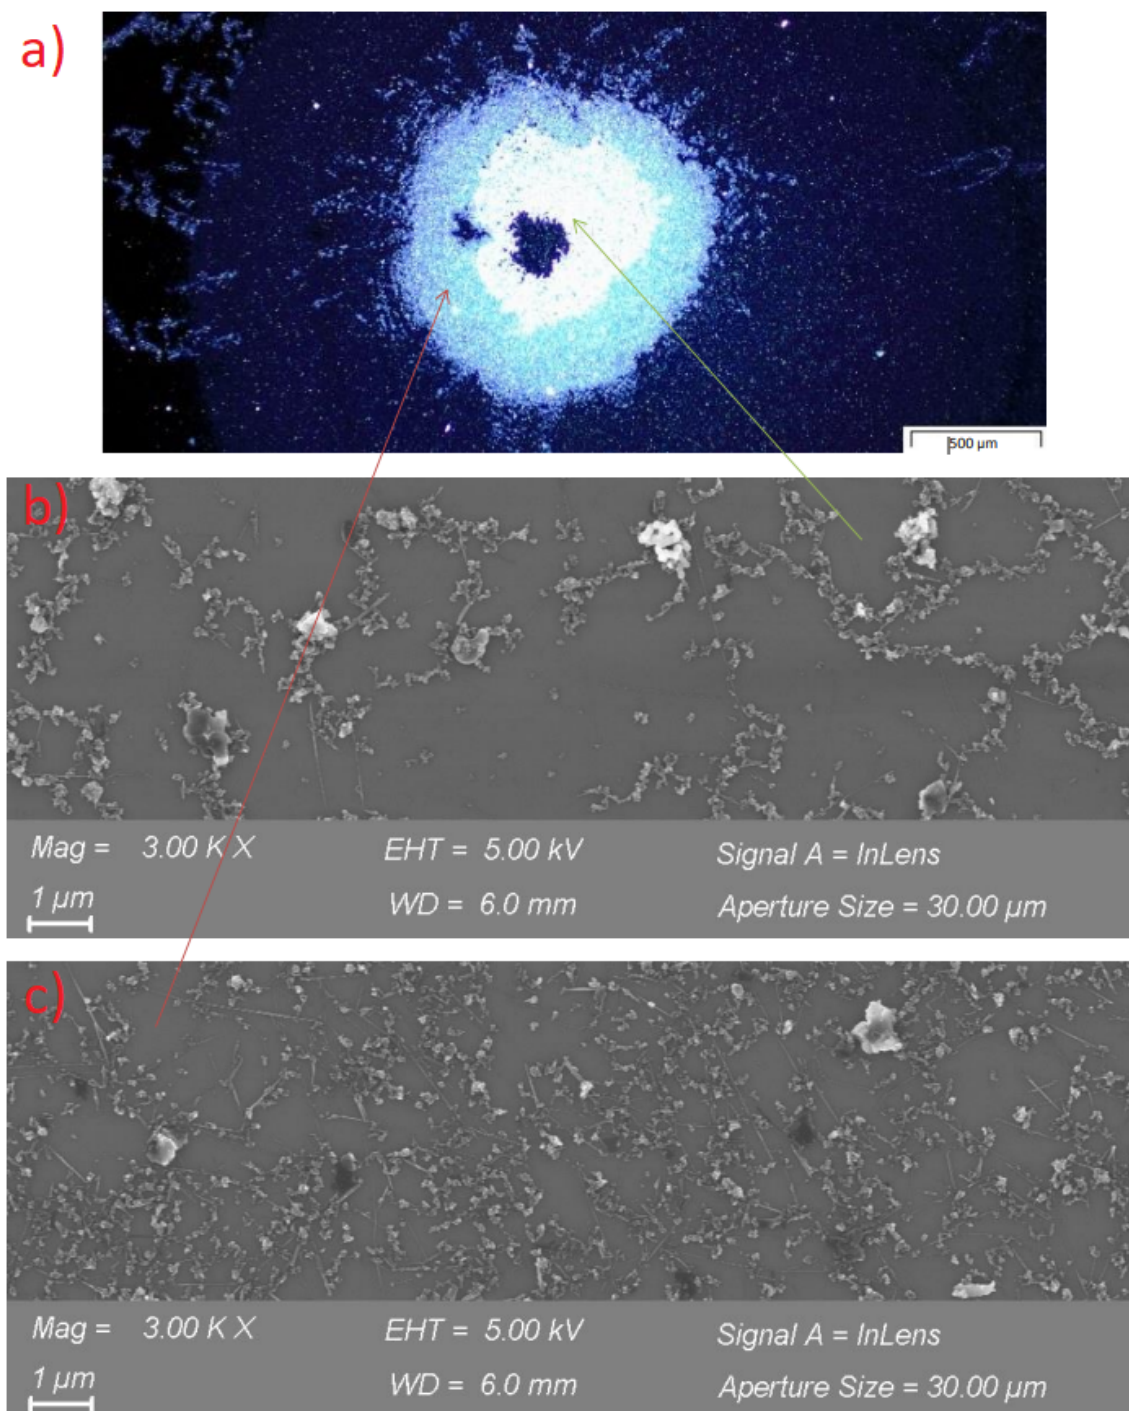


Figure 18. a) DF optical image showing at which positions the SEM images in b and c were taken. b) A SEM image of the brighter zone, showing a sparser distribution of material than in c. c) A SEM image of the dimmer zone, showing a denser distribution of material than in b. In b and c some clutter free MWNTs can be seen, but most are covered in amorphous carbon particles. Images a–c are originally captured by Esa Hyyryläinen.

The sample surface is formed by a network of randomly oriented MWNTs and amorphous carbon particles. While the individual particles and MWNTs are in random orientations, in general on the surface a reasonably ordered structure is formed from the network of MWNTs. One hypothesis for the MWNT deposition's behaviour in visible light is that the MWNT network acts as a diffraction grating. The mean space between features varies in the visible spectrum's wavelength range and it could be a deciding factor in why the sparsely populated area appears brighter than the more densely populated one.

Additionally as noted before, for diffraction to occur, the spacing of the features of a grating must be larger than the wavelength of the incident light. This could suggest that in the more densely populated area of the sample only the shorter wavelengths of the visible spectrum can be diffracted, while in the sparsely populated area a larger portion of the visible spectrum can be diffracted.

Considering that the subject of interest in this research are MWNTs, whose diameter typically ranges from 2 to 20 nm and length from few μm to tens of μm and larger amorphous carbon particles, the scattering of light from the sample's different objects could have behaviours corresponding to geometric, Rayleigh and Mie scattering. Along the tube axis the length scale is typically clearly above the Mie scattering domain. In the radial direction however we are dealing with length scales that are in the Rayleigh and Mie scattering domains.

4.2.2 AFM images

AFM images can be formed using various signals of an atomic force microscope. The most typical of these is the height sensor signal. The intensity of the image maps to height values that are recorded from the deflecting AFM tip, as the tip scans over a sample surface. Another signal type we showed in figure 13 is the peak force error signal. It is a feedback signal and doesn't yield any real information apart from making an aesthetically pleasing image that gives a sense of three dimensions.

From the height sensor AFM images we can observe and quantify the height, and therefore the diameter, of the MWNTs on the sample surfaces. The AFM images show especially well that the samples are indeed quite two-dimensional in nature. Some AFM images and their corresponding positions in an optical image are shown in figure 19. From these images it can be observed that most of the material in the depositions lies in the height range of tens of nanometers.




Figure 19. AFM height sensor images of sample no. 2 taken from different points in the sample. a) is taken inside the dim zone, b) is taken also in the dim zone, but closer to the bright zone and finally c) is taken at the interface of the bright zone and empty chip.

The quantification of the nanotube's diameters would be quite difficult using the AFM images taken of the main samples of this research. For that task we utilize AFM images taken from a sample fabricated by Esa Hyyryläinen for his own master's thesis research. The sample was fabricated in the same manner and using the same MWNT suspension as the main samples of this research, but has a cleaner evacuation zone where the pipetted water droplet has swept away the spin coated MWNT material. In the evacuation zone are chains of MWNTs and amorphous carbon. These chains are an useful location for quantifying the MWNTs' diameters, as there are fewer amorphous carbon particles around the MWNTs than in the main deposition we're interested in. Three AFM images were captured from such a chain and from these the height of a number of MWNTs was determined using Nanoscope

Analysis software and the step tool. The tool can be used to mark a box in the image, plotting the height sensor data in the direction of the dividing line of the box, while the direction perpendicular to the divider is used to average the height sensor data. The AFM images were first processed by utilizing the 2nd order flattening tool, which fixes bow and tilt in AFM images by fitting a second order polynomial at each scan line. The heights were then determined as the difference between the maximum height observed at the nanotube, from which the average height of the surrounding substrate was subtracted. Figure 20 shows the three AFM images, an optical image of the chain and finally table 4 lists the heights determined at each of the numbered locations.

Location #	a) (nm)	b) (nm)	c) (nm)
1	8.05	22.60	10.60
2	7.26	19.30	7.51
3	31.00		3.07
4	29.90		13.40
5	9.02		3.05
6	11.80		10.50
7	16.10		8.48
8	5.55		4.50
9	6.52		11.80
10	10.90		
11	10.60		
12	7.91		
13	2.45		
14	20.40		
15	5.34		
16	26.50		

Table 4. Diameters of the MWNTs determined from the AFM images a-c in figure 20 at the numbered locations. The mean diameter is 12.00 nm and standard deviation is 8.02 nm.

The MWNTs' diameters are in the very typical range as reported in literature, with the smallest here being 2.45 nm and the largest being 31 nm. Some uncertainty is naturally associated with the height measurements. The height measurement using the step tool yields slightly different results depending on the size of the averaging box and its location, with the results varying over a range of couple nanometers. In this research this is acceptable, as highly precise diameter measurements aren't something we are mainly interested in here. Additionally for the case of the largest nanotube




Figure 20. AFM height sensor images of a sample fabricated by Esa Hyyryläinen. The optical image shows the junction about which the AFM images were taken. The numbered boxes mark the locations at which heights and therefore the diameters of the MWNTs were determined. The heights are reported in table 4.

diameters, a MWNT's structure can flatten, skewing the height measurement towards smaller heights than in reality. While height statistics could be inferred in a similar manner for the amorphous carbon particles, we'll leave that for part III of this thesis. There we'll utilize a more powerful method to quantify the sizes of the particles and distances between them.

4.2.3 SEM images

Like with AFM images, SEM images can be formed from various signals of the scanning electron microscope. The signal used to form the images in this research are secondary electrons. These are low energy electrons and therefore they originate only from the first few nanometers of a sample surface. This gives the images a very high resolution.

The SEM images provide a slightly different view into the samples. Using the SEM images it is very straightforward to quantify the length of the MWNTs. SEM is the better option in measuring distances over AFM, as the AFM tip slightly distorts lateral features as the tip scans over sample geometry. Figure 21 shows SEM images from the dim and bright zones, with yellow lines and numbers used to mark them. Statistics with histograms, fitted curves and descriptive values are shown in figure 22 for the length measurements. The MWNTs were identified and a segmented line was overlaid manually on top of them for the length measurement using ImageJ software.

From the dim zone 300 nanotubes were identified, while plenty more still remain in the image. The longest single nanotube was $3.76 \mu\text{m}$ and the shortest fragment was $0.146 \mu\text{m}$. The mean and standard deviation of the tube lengths is $0.861 \mu\text{m}$ and $0.595 \mu\text{m}$ respectively. Similarly 375 nanotubes were identified from the bright zone with the single longest and shortest tubes being $5.278 \mu\text{m}$ and $0.146 \mu\text{m}$ respectively. The mean nanotube length was $0.946 \mu\text{m}$ and standard deviation was $0.797 \mu\text{m}$.

In the 8000x magnification SEM image one pixel is equal to about $0.0068 \mu\text{m}$ and the error due to placing the nodes of a segmented line shouldn't be greater than 5 pixels, or equivalently $0.034 \mu\text{m}$. Most lines consist of only two nodes, leading to an approximately $0.05 \mu\text{m}$ measurement uncertainty by the uncertainty propagation law. Some very curved lines can consist of up to 20 nodes, which would lead to a measurement uncertainty of approximately $0.16 \mu\text{m}$. Similarly for the 3000x magnification one pixel is equal to about $0.017 \mu\text{m}$ leading to a $0.24 \mu\text{m}$ error for two nodes and up to $0.76 \mu\text{m}$ for 20 nodes. In the end the uncertainty of individual MWNTs' lengths in most cases is small compared to the statistical variation of the lengths and to the uncertainty introduced by ambiguous nanotube ends.

The 5 pixel error is a fair estimation of the error for manually placing the nodes. However, rather than the uncertainty of node placement, a greater source of uncertainty in the images is how well the individual nanotubes' ends can be identified. The nanotubes used in the length calculations were chosen so that their ends were as

clear to spot as possible. Nevertheless, some more or less ambiguous nanotubes were included in the calculation. If a nanotube entered an amorphous carbon cluster and it wasn't clear that the tube continues through it, the tube's end was estimated to be at the center of that particle. Very large particles were avoided altogether. With multiple nanotubes of similar diameter inside or under a cluster of amorphous carbon particles, it is impossible to perfectly determine if the nanotube truly terminates or if it just bends inside and continues. Likely the distribution of the nanotubes' lengths should tend slightly more towards greater lengths than what the measured distribution represents, but only slightly.

Overall the histograms and their log-normal fits in figure 22 are very similar and suggest that there is no selectivity for the zones in terms of what kind of deposited material they consist of. The only difference is the density of the deposited material, which we'll get more into in part III. Something else we have noticed in the SEM images is that typically the longest nanotubes tend to have a smaller diameter, while the shorter fragments tend to have a larger diameter. This is expected behavior as the bending modulus of MWNTs is inversely proportional to their diameter, i.e. MWNTs with smaller diameters can bend more without irreversible deformations[35]. At least one relevant step before depositing the material on the chips, is the sonication of the MWNT solution. This likely causes the thicker nanotubes to fracture into the smaller pieces we observe in the SEM images.

We have made various observations from the optical images, AFM images and SEM images and what is seen are two very clearly distinct zones with different brightnesses and amount of deposited material. We would like to get some insight into the depositions' nature by analyzing the digital camera data obtained from the optical images. The analysis so far is, however, mostly only qualitative since the output of a digital camera is nonlinear. This makes quantitative observations basically useless as any numerical values derived from the nonlinear pixel values of an image are skewed.

Overall the key question, around which work of this section is centered, is whether consumer-grade digital cameras can be utilized as radiance sensors. We will be systematically building a model for the image formation process, which in turn helps us determine the necessary processing and correction steps for the image processing pipeline. With the required steps established, we can then overcome the issues of the imaging system by proper processing of image data, which will be the main subject

of part II. Finally in part III, we will apply this knowledge and analyze the optical images we already presented and get an answer whether consumer digital cameras can be utilized as radiance sensors or not, what is their accuracy and what limits their accuracy.




Figure 21. SEM images of a) the dim zone and b) bright zone. Note the different magnifications in each image. Due to the bright zone having less material, the length measurement data was gathered over a larger area.

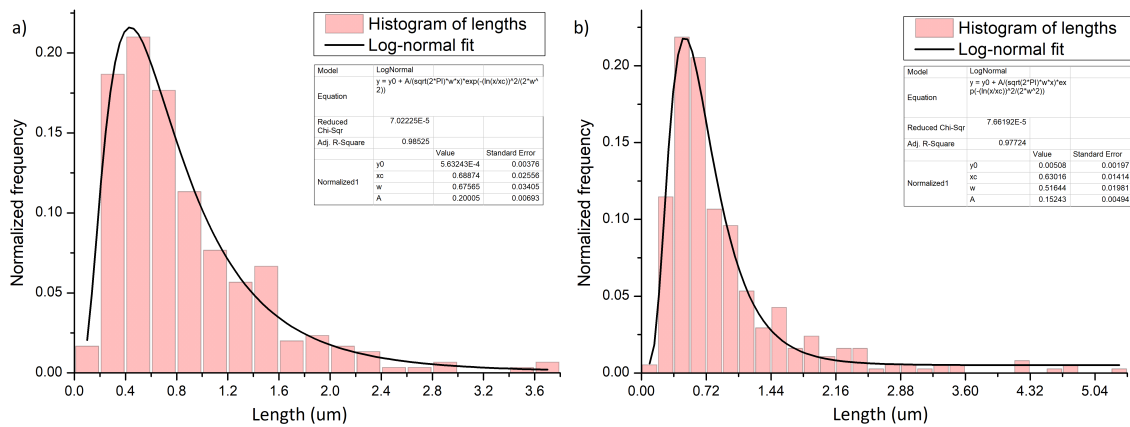


Figure 22. Histograms of the nanotubes observed in the SEM images in figure 21 a) and b) respectively. For each histogram a least squares error fit of a log-normal distribution is also plotted with the inset showing the fit parameters.

Part II

Digital images

1 Introduction to digital imaging

The charged particle microscopes and AFM provide a higher resolving power than an optical microscope, however they naturally can't directly provide information about the optical properties of a subject in the same way an optical microscope paired with a digital camera can. The history of digital imaging is intrinsically tied to the history of semiconductors. The first digital image sensor was the charge-coupled device (CCD) which was invented in 1969[36], while the first commercial digital cameras came to the market in late 1980s[37]. Since then digital cameras have become an everyday device, originally as a dedicated device and later as a feature that nearly all smartphones have.

Digital camera sensors come in two types: the previously mentioned CCD and an complementary metal-oxide-semiconductor (CMOS) sensor, both based on the metal-oxide-semiconductor field-effect transistor (MOSFET)[36, 38]. They both generate charge via photo electric effect, but the differences come in how the generated charge is quantified. CCD devices transport the charge accumulated to a row of pixels into one corner of the chip and read out. The voltage signal is then amplified and converted to a digital value by an analog-to-digital converter (DAC). CMOS sensors have their amplification and readout circuitry at each pixel of the sensor. On one hand CCD sensors have lower noise and a higher sensitivity to light than CMOS sensors. On the other hand CMOS sensors consume much less power and are less costly to manufacture than CCD sensors. The performance of CMOS sensors has rapidly improved over time and in most situations they are sufficient, but CCD sensors are still preferred in some highly demanding applications, astronomy being a prime example.[36]

In this research alongside the optical microscope, a CMOS camera has been utilized. The charge accumulated in the CCD and CMOS elements is mostly linear as it is proportional to the irradiance the elements are subjected to[36, p. 173]. Based on this one might expect that the output of the image sensor is linear. However in reality this is seldom the case, as commercial camera's typically apply nonlinear

operations to the sensor output in order to present an aesthetically pleasant image optimized for human vision[36, 39]. Especially when the sensor's exposure is either very short or very long with respect to the illumination, the sensor can behave in a nonlinear fashion.[39] In opposition to commercial digital cameras, more specialized digital camera's can maintain a linear response. Essentially the linearization aspect of this research is something that can be avoided with device choices, however the work herein still provides tools to analyze the linearity of such devices first-hand. The other image processing methods and corrections are something that can be generalized into most digital cameras. In the end one key aspect of this research is to enable the calibration and linearization of a nonlinear digital camera's response to enable quantitative irradiance measurements from such digital images.

Optical microscopes can be equipped with spectrometers to provide precise measurements of irradiance and its spectrum[40], but it does require a whole separate device attached to the microscope, possibly replacing a digital camera. These devices are naturally a lot more expensive in comparison to a digital camera, which can be readily observed by browsing the catalogs of spectroscopy and camera manufacturers. Also, as examples, there have been some studies on the subject of utilizing digital cameras over spectrometers for measuring the optical properties of human teeth restorations[41–43] and orthodontic elastomeric ligatures[44], with one of the key motivations being the cost-effectiveness of digital cameras. The referenced studies find that digital cameras are equally or nearly equally reliable for color analysis. While these studies don't delve into the issue of linearity with respect to exposure time, they do demonstrate the possibility of utilizing digital cameras for accurate measurement of colors.

With a linearized digital camera a user can use multiple exposures to construct a high dynamic range (HDR) image of a scene, which can provide quantitative measurements of relative irradiance for each color channel of a camera[39]. Based on the image formation process in digital cameras[45], we'd also argue that if further corrections are made that eliminate or at least minimize certain artifacts in the imaging system, an estimate of the original scene radiance can be computed in a relative scale. We will utilize the established methods in part III to obtain HDR images of the CNT depositions with the aim to gain even more insight into their nature.

Digital images, their acquisition, processing and editing are each a large topic of

their own. Part II of this thesis approaches these subjects from the perspective of utilizing an inherently nonlinear camera. While most of what is discussed herein can be applied to digital imaging in general, it should be noted that also a large variety of machine vision cameras exist, meaning that especially the issue of linearization can be avoided via proper choice of device. One motivation of this research was making best of what we have and the author's curiosity if even a "normal" consumer camera could be utilized in accurate quantitative measurements. Even though the linearization methods proposed here are not fully needed for machine vision cameras, we do propose that the methods can still be utilized to quantify and measure the linearity of such a camera first hand. Further more we argue that the HDR merging methods and imaging system artifact corrections are generally applicable to any digital camera, due to their fundamental approach.

An emphasis is maintained on the application of digital imaging for optical microscopy. In the sections to follow we'll briefly go over the process of image acquisition and how it can be modeled. With established models a camera's linearity, noise and artifact sources can be understood and any undesired system properties can be mitigated. All the work culminates into the full image processing pipeline and how we'll utilize our camera to produce estimates of original scene radiances and what kind of measurements we'll perform in part III.

2 Digital image acquisition

Digital cameras' CCD and CMOS sensors are semiconductor devices that in essence map the irradiance E at a sensor's pixel array to brightness values B in the final image[45]. The device's sensor is a two-dimensional array of wells, in which charge is generated by the photoelectric effect[36, pp. 159–166]. The charge is then transferred and amplified, and the electrical signal is finally sampled and quantized into the bits of a digital image's pixels. Grossberg and Nayar describe that the image acquisition process can be split into two transformations starting from a scene radiance L [45]. The first transformation s is a linear transformation that maps the scene radiance L into irradiance at the image sensor E . The second transformation is f , which is a nonlinear mapping of the sensor irradiance E into the final image brightness B . s basically models the optics of the imaging system, such as vignetting and fixed-pattern artifacts, while f models the nonlinear response of the imaging system. This process is visualized in the flow diagram shown in figure 23.

In this research we are going to work through the flow diagram in reverse order. Starting from image brightness, we'll recover the camera response function (CRF), or more precisely the inverse camera response function (ICRF). This can then be utilized to linearize the image brightness and together with known exposure times to obtain image data in terms of irradiance[39, 45]. This will be the subject of part II, section 3. After the data has been linearized and merged into a single HDR image, we will apply linear corrections that reverse the optical effects of the imaging system. This finally yields data that acts as an estimate of the original scene radiance. The actual image acquisition and initial data handling acts as the first building block in our image processing pipeline.




Figure 23. Simple flow diagram of the two transformations encompassing the image acquisition process, in which s maps the scene radiance L to irradiance at the sensor E , followed by f , which maps E into image brightness B . s is a linear transform, though it can vary spatially, while f is a nonlinear transform, but spatially uniform. Reproduced from [45].

2.1 Images as data

In terms of handling images digitally, images basically carry data that encodes, depending on the format, a captured image’s brightness B and color into a two-dimensional array of bytes. How this data is encoded changes between different color spaces. The two key properties for the encoding are its bit depth and color space. Bit depth determines how many distinct values the range $[0, 1]$ can be divided into[36, pp. 14–15] and therefore determines the precision and memory cost for an image’s pixel. Color space is a quite broad concept, but simply put it can be defined as an typically three-dimensional geometrical model in which coordinates represent a color and the vector basis represents their color components[36, p. 84].

For a typical sRGB image the bit depth is eight and there are three color channels[36, p. 15], meaning three bytes representing 8-bit unsigned integers, one byte for each color channel, carrying the brightness information. With similar logic, for a single channel image a bit depth of 32 means the numbers represent floating-point values. On display devices the colors are created with a color space based on the additive RGB color model[36, p. 416]. In printed media the subtractive CMYK color model is used to create colors[36, p. 15]. These two models are usually represented as three-dimensional spaces in Cartesian coordinates by assigning the three colors of each model on an axis[46]. In computer graphics and image processing there are also the HSL or HSV color models, which are based on the hue, saturation and value or lightness[36, p. 413]. These color spaces are commonly represented in a cylindrical coordinate system[46].

In the scope of this research we utilize images with bit depths of both 8 and 64 and start from the sRGB color space. We do not explicitly move away from the sRGB color space, and are using the blue, green and red color channels throughout this work, however we will perform all computations and store HDR images with

64-bit channels. Something that must be noted is that the sRGB color values are not linear[47] and as such are problematic to do calculations with. Allen and Tiantaphillidou describe that the reasoning for this nonlinearity is due to the RGB values being gamma-expanded to place more datapoints in the darker values instead of placing them evenly[36, p. 390]. This in turn is due to human vision being able to more easily notice color differences in darker colors than in bright ones. Without the gamma-expansion one would notice bands in an image. The terms gamma expansion and gamma compression refer to the process of applying a simple power law to the brightness values of an input image in an image system[36, p. 377]. The name arises from the the presentation of the general form of the power law in such contexts, where γ is used as the symbol for the exponent in the power law. The sRGB color space standard defines the transfer function

$$C_l = \begin{cases} \frac{C_s}{12.92} & , C_s \leq 0.04045 \\ \left(\frac{C_s+0.055}{1.055}\right)^{2.4} & , C_s > 0.04045 \end{cases} \quad (29)$$

in which C_l is a linear RGB value and C_s is an gamma-expanded RGB value, with 2.4 being the gamma value[47]. The same standard also defines the inverse transform as

$$C_s = \begin{cases} 12.92C_l & , C_l \leq 0.0031308 \\ 1.055C_l^{(1/2.4)} - 0.055 & , C_l > 0.0031308. \end{cases} \quad (30)$$

Both the C_l and C_s values are in the [0,1] range.

Basically whenever we do calculations with the image brightness values, we must transform the nonlinear values into linear values. This simply follows from the fact that with nonlinear values any arithmetic operations will be skewed. Using equation 29 to transform to linear values upon loading an image might suffice for some cameras, but it is more likely for manufacturers to utilize device-specific transfer functions that are different from the sRGB transfer function[48]. In our case the linearization of images will be performed by using an inverse camera response function, which we'll get into in part II, section 3. The transform to linear values is something that has to be done, because otherwise the values are basically weighted and would therefore skew the results of any operations, while the transform back into nonlinear space is something that we can choose to do or not to do depending on the situation. Naturally all calculations will be performed with 64-bit channels to

maintain numerical accuracy.

Finally in equations we will use the notation (x,y) for images, to denote that they are arrays of values. Furthermore the values contained in the array are basically vectors of three components, one for each color channel. The color channels are always handled independently of each other and when referring to pixel values of an image in a given domain, we are referring to a three-component value. If a single color channel is referenced, we'll use subscripts R, G and B for red, green and blue channels respectively. The images we handle in this research come in three domains. The typical eight-bits per channel image that a camera outputs is in the brightness domain B with 256 discrete values between $[0, 1]$ for each channel. Then we have the exposure domain X for linear images, also with 256 discrete values in the $[0, 1]$ range. Starting with images in the irradiance domain E , the values can technically have any non-negative real value. In reality the values' precision will be limited at least by the numerical precision of floating point numbers and more practically by the uncertainties of the imaging system. Finally we have the radiance domain image L , which constitutes the best estimate for the radiance of the original scene, with the same considerations as the irradiance domain images.

2.2 Noise

Digital cameras are prone to various different types of noise that reduce the signal-to-noise ratio (SNR). The signal produced by the physical sensor of a digital camera is directly proportional to the number of photons detected by the sensor[39]. The number of arriving photons follows Poisson statistics and as such the signal has an associated variance over multiple measurements[49]. This is known as photon shot noise.

Digital camera sensors as semiconductor devices generate charge even in a perfectly dark environment due to thermal excitation[49]. Dark current is independent of the measured signal level but increases with rising temperature. The actual magnitude of the charge generated follows Poisson statistics as each thermal excitation event is independent, just like the arrival of photons for photon shot noise. Furthermore the magnitude and variance of dark current and defects between pixels cause fixed-pattern noise and pixel response non-uniformity[36, p. 442].

There are also additional noise sources and artifacting that affect the image quality. Allen and Tiantaphillidou describe in their book various types of noise and

their origin[36, pp. 442–443]. The surrounding electronics that are used to measure the number of electrons in the imaging sensor add readout noise. The amplifiers that boost the voltage signal add noise into the signal. Also the digitization process of the voltage values into a digital number level causes small loss in the accuracy of the measured values. Finally resetting of the photoelement to a reference level adds noise due to variation on each reset, circuit clock signals and outside interference can also add noise to the signal[36, pp. 442–443].

For a static scene the effect of random noise can be mitigated in a statistical sense by capturing multiple images of the same scene and calculating a pixelwise mean of the images. All the noise sources in a digital imaging process are symmetric and have a Gaussian profile[50]. This means that we can compute a mean image to reduce the effects of noise. The nonlinear mapping of most digital cameras causes some skewness in the noise profiles[50], but we'll discuss its effect on the mean image in part II, section 3.4. Finally, acquiring a mean for a pixel's value over N images gives a SNR of[36, p. 442]

$$\text{SNR} = \frac{N}{\sqrt{N}} = \sqrt{N}. \quad (31)$$

With equation 31 it is easy to see that the original SNR of a single image decreases by a factor of $1/\sqrt{N}$ and therefore has diminishing returns with an increasing number of images. For example four images cuts the SNR into 50% of the original value and 10 images cuts it to 31.6%. Due to these diminishing returns you must consider how many images do you have the time and memory to capture.

2.3 Mean values

Normally calculating the mean of a data set is a trivial endeavor, but in the case of image formats with limited values we argue that it requires a little more consideration with respect to the final result and how it is interpreted. If we are only interested in calculating the mean of the pixel values in a given set, then there are no issues. However, issues start to appear if we are to assign some physical meaning to the pixel values, and we need to take care with the pixel values at the edges of the data range near zero and one. At zero the sensor is in cutoff and just from the one image, we can't tell if the irradiance at the particular pixel causes the sensor to barely be at zero or if it is far below zero. Similarly at one the sensor is saturated and we can't really tell if the sensor is barely saturated or far above one.

With this consideration in mind, areas of interest in a single image should be at the middle values of the sensor's range to increase accuracy, which can typically be adjusted with a camera's exposure time. In this research one thing we will be investigating is a camera's linearity and we will be comparing the brightnesses of images with vastly different exposure times. In those cases, to ensure physically meaningful comparisons, we propose to use the following approach. We'll only utilize the pixels, which in both of the two images under comparison are sufficiently within the sensor range. Noise adds another consideration in this, so that in such comparisons, while the values will have the range $[0,1]$ we argue that only a range $[0 + \delta, 1 - \delta]$ should be utilized. We propose that δ is a value that reflects the noise characteristics of the utilized camera and which we can determine empirically. δ might not necessarily be the same for the opposite edges of the range.

Finally in this research we'll be dealing with two different means. The first is used to calculate the mean image of multiple images to increase measurement accuracy. Nearly all images utilized in this research are acquired by capturing multiple images and taking the pixel-wise mean of them. This can basically be interpreted as a temporal mean. The temporal mean is so commonly used here that we won't be indicating it in notation, but throughout this text all captured images can be assumed to be temporal means, unless stated otherwise. The other mean that will be also utilized is the spatial mean of an image's pixel values. Considering an image $I(x,y)$ with X columns and Y rows of pixels, the spatial mean is calculated as[51]

$$\bar{B} = \frac{1}{X \cdot Y} \sum_{x=0}^{X-1} \sum_{y=0}^{Y-1} B(x,y) = \langle B(x,y) \rangle. \quad (32)$$

Here it should also be noted that the (x,y) notation is dropped from the spatial mean of an image, as the value is naturally no longer a function of position in the image. We'll also occasionally use the notation ΣB to indicate a sum of all the pixels in an image.

2.4 Uncertainty values

Using mean images yields us a better accuracy and enables us to use some measure of variability for uncertainty. In the case of the mean images, the measurement uncertainty shall be described by the standard error, whose estimate is given by[52]

$$\delta B(x,y) \approx \frac{\sigma_B(x,y)}{\sqrt{N}}, \quad (33)$$

in which $\delta B(x,y)$ denotes the uncertainty of an image and σ_B is the pixel-wise standard deviation of the N images used to determine B . Considering that we are calculating the pixel-wise mean of a static scene, the spread of the pixel values for a particular pixel will be quite small, meaning the resulting standard error value will be a lot smaller than the actual pixel value. This in turn means that of the $[0, 1]$ range, the uncertainty image will only utilize a very small portion. Therefore we either choose to save the resulting uncertainty value in a more precise format, that is 64-bit channels. Different scaling solutions, for example to 16-bit or 8-bit ranges, could be used to save storage capacity but this would lead to a slight precision loss and thus we'll simply store directly in the 64-bit format.

Another thing that has to be taken into consideration for calculating the variance image is the actual RAM of a computer and numerical stability. The usual way of calculating the variance for a given set of sampled values $\{x_0, \dots, x_{N-1}\}$ would utilize the equation to calculate the unbiased sample variance for a discrete random variable:

$$s^2 = \frac{1}{N-1} \sum_{i=0}^{N-1} (x_i - \bar{x})^2, \quad (34)$$

in which s^2 is the sample variance, N is the number of samples and \bar{x} is the mean value of the set[52]. The problem with applying this equation directly to an algorithm is that it clearly requires two passes through the data set due to using the mean value for each element in the sum. This is obviously problematic for large sets, high bit depth images and computers with a small RAM.

In this research we will utilize the Welford algorithm for calculating the mean and variance of an image set[53]. The Welford algorithm can be used to determine the mean and variance of a set for each additional element, enabling one-pass calculation.

The values that are updated for each new element x_n in the set are

$$\begin{aligned}\bar{x}_n &= \bar{x}_{n-1} + \frac{x_n - \bar{x}_{n-1}}{n}, \\ M_{2,n} &= M_{2,n-1} + (x_n - \bar{x}_{n-1})(x_n - \bar{x}_n), \\ s_n^2 &= \frac{M_{2,n}}{n-1}.\end{aligned}$$

These equations yield us a very simple, efficient and accurate algorithm for computing an acquired image $B_A(x,y)$ and standard error image $\delta B_A(x,y)$ for an image set[54]. While memory isn't most likely an issue with a spatial mean and spatial variance of an image, the Welford algorithm can similarly be used to calculate those too, but instead of iterating through a set of images $\{B_0(x,y), \dots, B_{N-1}(x,y)\}$, we'll just iterate through all the pixels in an image $\{(x,y) \mid x \in [0,X] \text{ and } y \in [0,Y]\}$.

While using the pixel-wise mean image and standard error image based on an image stack is the sensible and a rather obvious way to approach uncertainties, the equivalent is not seemingly so obvious for the spatial mean and its standard error. In their article Santo et al. determine the standard uncertainty of a spatial mean by computing the spatial mean of the associated standard uncertainty image[51]. The authors do not provide further reasoning for this approach, but it is an approach that we agree with. We propose the following reasoning for this method. Naively in this case one might approach the issue by just computing the standard error for the population of pixels that is used to compute the spatial mean. This approach however does not yield a relevant measure. The standard deviation of this population of pixels rather measures the spread of pixel values of a given scene. Unless the scene is uniform across these pixels, then they are not samples of the same distribution in the same manner the pixels at the same position in a stack of images for a static scene are. Therefore the logical approach for computing the standard error of a spatial mean is to compute the spatial mean of the corresponding standard error image at the same pixel positions. This measures the typical uncertainty at those pixel positions and yields a sensible estimate for the uncertainty of the spatial mean.

3 The camera response function

Starting to work backwards from the image brightness towards scene radiance, we'll first need to reverse the nonlinear response of the camera. This yields us data in terms of image exposure X , from which we can determine the irradiance E when the exposure time is known. Then we can apply corrections to reverse the linear effects of the system, finally yielding data in terms of scene radiance L .

As we noted earlier in part II, section 2.2, the number of photons arriving at a camera sensor's pixel follow Poisson statistics. Therefore with a large enough flux, the number of arriving photons in a given time follows a Gaussian distribution. Furthermore as noted, the charge accumulated to a sensor's pixel is proportional to the irradiance it experiences and is mostly linear[36, p. 173]. Therefore under constant lighting conditions the accumulated charge is also linear with respect to time. From this it should follow that if exposure time is doubled, then the pixel's brightness should also double if the lighting is kept the same – provided that the sensor is not saturated. In most digital cameras this is not the case as they apply a nonlinear mapping to the sensor's outputs[39].

This nonlinearity means that even when maintaining all conditions the same, apart from exposure time, you cannot accurately compare the brightness values between two images captured even with the same camera. Furthermore even within a single image the brightness between different pixels cannot be accurately compared directly due to the nonlinear mapping. The recovered inverse camera response function (ICRF) will act as the second building block of our full image processing pipeline.

3.1 Empirical model of the CRF

In this research we will utilize the empirical model of CRFs developed by Grossberg and Nayar and reported in their paper 'What is the Space of Camera Response Functions? Their work begins from the following assumptions:[45]

- The camera response function is the same for all the sensor's pixels. This means that the CRF can be modeled independently, with the artifacts of the imaging process handled by a separate function. This makes the CRF a function of only one variable $f(X) = B$, in which B is the image brightness.
- The CRF maps the exposure values X to image brightness B with the camera's response ranging from a minimum value B_{\min} to a maximum value B_{\max} . The response's units are arbitrary and are normalized to $B_{\min} = 0$ and $B_{\max} = 1$.
- The CRF is a monotonically increasing function so that it is invertible and maps B_{\min} to a minimum detectible exposure X_{\min} and X_{\max} to a maximum detectible exposure E_{\max} . Additionally the minimum and maximum exposures are also normalized to 0 and 1 respectively.

The theoretical space of CRFs is then defined as

$$W_{\text{RF}} := \{f | f(0) = 0, f(1) = 1, f \text{ increases monotonically}\}, \quad (35)$$

based on the assumptions. As a side note, Grossberg and Nayar note that any function f of exposure may be thought of as a finite-dimensional vector, whose components are the values of f sampled at fixed points in the exposure domain[45]. While certainly a helpful way of visualizing the space of the camera functions, as the authors note, it is also the method which we will use to handle the camera response functions. We work with 8-bit images and therefore are limited to 256 data points, so the CRF will be represented by a 256-dimensional vector.

Grossberg and Nayar continue by sampling the response function at P points yields the vector $(B_1, \dots, B_P) = (f(X_1), \dots, f(X_P))$ for which the last point in the exposure domain will be set as $X_P = 1$ [45]. Now based on the assumptions, if f is a valid response function, it is normalized so that at exposure 1 the brightness is $B_P = f(1) = 1$. They argue that all response vectors lie in the plane W_1 , formed by response vectors, whose last component B_P is 1. Taking another arbitrary valid

response function f_0 , their difference $h := f - f_0$ must lie in a plane W_0 , which is parallel to W_1 and intersects with the origin. From this it follows that any response function can be expressed as $f = f_0 + h$ with f_0 being an arbitrary response function and h is an arbitrary vector that lies in the plane W_0 [45].

The authors also deduce, based on set theory and the assumption of the response function being monotonically increasing, that the monotonic response functions form a cone Λ [45]. Therefore the theoretical space of response functions is the intersection of the cone Λ and plane W_1 . The authors finally deduce that the mean of any set of CRFs is also a valid CRF.

Finally Grossberg and Nayar formulate the following general approximation model for CRFs:

$$\tilde{f} = f_0(X) + \sum_{n=1}^M c_n h_n(X), \quad (36)$$

in which f_0 is some arbitrary base response function, h_n are an arbitrary choice of basis vectors for the vector space W_0 and c_n are the coefficients of the model[45]. This is a general model for an approximation and it doesn't give any information about what the best choices for a base response function and basis vectors are. Grossberg and Nayar approach the choice of basis by employing response functions of real-world imaging systems. The authors also explicitly note that all of this is also directly applicable to the ICRF due to the fact that any valid CRF is an invertible function and therefore have all the same properties[45].

In their paper Grossberg and Nayar report that they have a database that contains a total of 201 response functions measured from photographic films, digital cameras and synthetic gamma curves[45]. They use principal component analysis (PCA) on their database to find a base response function f_0 and a basis h_n . PCA can be thought of as a method of reducing a data set's dimensionality, while keeping as much of the original information as possible[55]. So in this case the original data is the finite-dimensional sampling of each response function in the database at points $0, \dots, P - 1$. The basis vectors h_n for the empirical model are then determined by selecting the eigenvectors corresponding to the M largest eigenvalues of the covariance matrix, whose elements are[45]

$$C_{m,n} = \sum_{p=0}^{P-1} (g_p(X_n) - f_0(X_n))(g_p(X_m) - f_0(X_m)). \quad (37)$$

In equation 37 the sum indexes the p response functions in the database, with g_p being the p th response function and f_0 is the mean of all the database's response functions.

The CRF maps the exposure values of a scene to the brightness values observed in the final image captured by a camera. In this research we are interested in the reverse case, that is finding the exposure of a scene from a given image based on the brightness values. This means we are interested in the inverse camera response function (ICRF). The CRF is an invertible function and therefore the ICRF has all the same properties as the CRF. In the following sections we will go over some methods of obtaining an ICRF and introduce our method, which is based on the empirical model of Grossberg and Nayar.

An even better set of basis vectors could be derived from the data set created by Chen et al. reported in their article 'Analyzing Modern Camera Response Functions'. They produce a similar data base to that of Grossberg and Nayar, but with cameras released after 2013[56]. The data set would likely form a set of more appropriate basis vectors for our case. Unfortunately, to our knowledge the authors haven't made the database public.

3.2 Measuring linearity

Generally ignoring the effects of noise, saturation and cut-off, when comparing two images of the same scene, whose only difference is exposure time, the following relationship holds true

$$\frac{g(B_s(x,y))}{t_s} = E(x,y) = \frac{g(B_l(x,y))}{t_l}, \quad (38)$$

in which g is the camera's ICRF and the indexes s and l denote the image with the shorter and the longer exposure respectively[57]. Essentially this equation just tells us that in the static case the camera sensor measures the same irradiance in both images and that the exposures of each image are linearly related to each other by their exposure times.

We propose two ways to utilize equation 38 for measuring linearity. One way we approach this problem is to capture a series of differently exposed images and determine the exposure time ratio between each unique pair. Then for a pair of images, we'll choose the image with the shorter exposure time as a reference. The

ratio of exposure times can be used to scale down the longer-exposure image to match the exposure time of the reference image. Then the difference between the images can be calculated by subtracting the scaled image from the reference image, i.e.

$$A(x,y) = g(B_s(x,y)) - \frac{t_s}{t_l}g(B_l(x,y)), \quad (39)$$

where we denote the difference as $A(x,y)$. We'll call this difference as the absolute disparity image. This relation can easily be solved from equation 38. In the ideal case without the effects of noise, saturation and cut-off and a perfectly known ICRF, the disparity image would be zero everywhere. If the ICRF is incorrect or if there is noise, then the disparity is not zero for all pixels. Naturally for a real system both of these are true. Equation 39 yields the disparity in an absolute scale in terms of pixel values, as the name suggests, but we can also determine the difference in an relative scale by dividing the disparity image by the scaled down image. This yields what we'll call the relative disparity image

$$R(x,y) = \frac{g(B_s(x,y)) - rg(B_l(x,y))}{rg(B_l(x,y))}, \quad (40)$$

in which $r = t_s/t_l$.

Utilizing the disparity images we propose the following simple metric for measuring the linearity of a imaging system. We'll calculate the spatial mean and sample standard deviation for $R(x,y)$, yielding us a value $\bar{R} \pm \sigma_{\bar{R}}$. Then for a set of images we'll calculate this value for all unique pairs in the set to plot the data with \bar{R} on the y-axis and the exposure time ratio t_s/t_l on the x-axis. An absolute and relative scale examples of such plots for the Olympus SC100 camera are in figure 24. Some of the images utilized in calculating the data points are shown in figure 25. The closer to zero all the spatial means of the disparity images are, the more accurately the ICRF is known and the less noise there is.

Figure 24 shows how each color channel behaves for the Olympus SC100 camera. The relative scale shows how the disparity increases with a larger difference in exposure times between a pair of images, while curiously the absolute scale shows the highest disparity in the middle of the plotted exposure ratio scale.

For a single image pair the linearity can be observed in a more robust manner by binning the values of the disparity image and plotting a histogram. The histogram then directly shows the spread of the disparity values. Additionally from a single

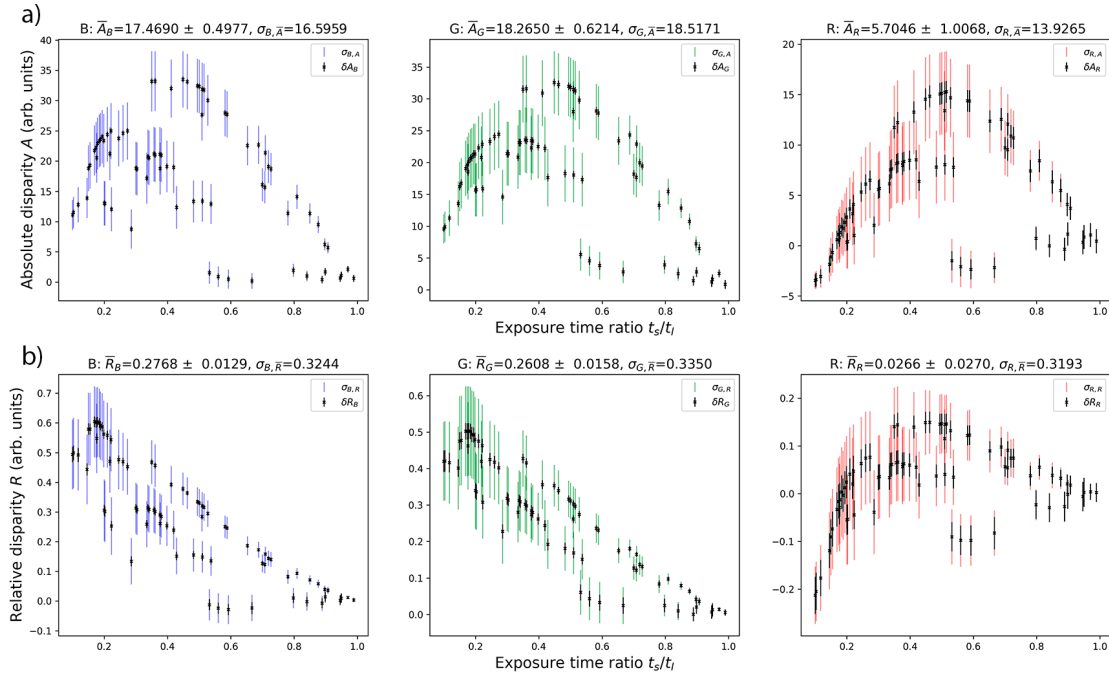


Figure 24. Plots of the pairwise spatial means of the absolute and relative disparity images, for each color channel of Olympus SC100. On the x-axis is the ratio of the short exposure time against the long exposure time for a given image pair. The colored bars represent the sample standard deviation of the disparity values and are divided by five for better visualization. The black bars represent the uncertainty of disparity. Additionally on top of the graphs are the mean values and standard deviation of the plotted data points. Values outside the range $5 \leq B \leq 250$ have been rejected from the measures.

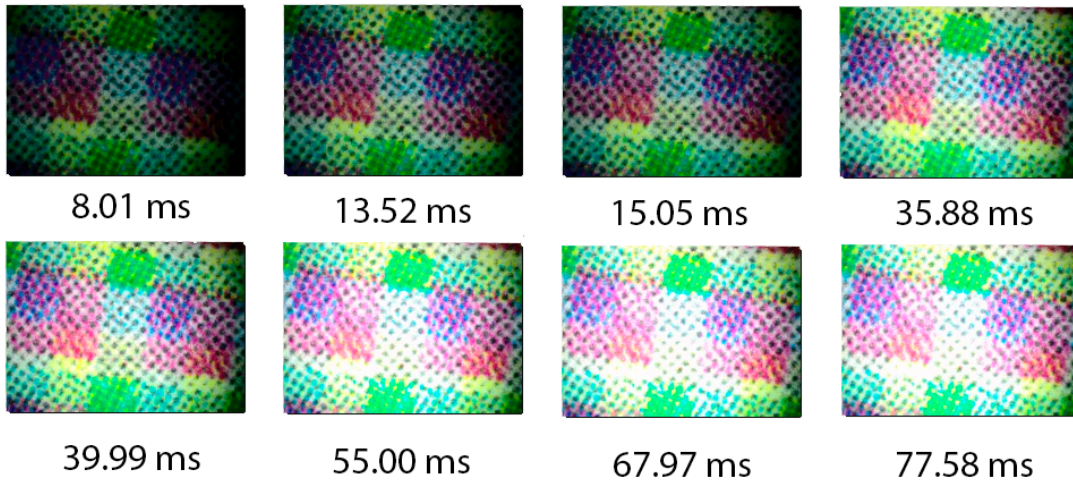


Figure 25. Some of the images utilized in calculating the data points used in the disparity plots show in figure 24.

disparity image one can observe the locations of specific disparity values and compare their locations to the original images. The relative disparity images for each color channel for the 40 ms and 100 ms image pair of sample No. 1 and histograms of their values are shown in figure 26. The histograms show that there are very few pixels that are even close to matching linearly in the pair of images, i.e have disparity of zero. Furthermore it can be seen that the red channel is the most linear of the three, while blue has the highest nonlinearity. From the disparity images it can be seen that the most nonlinear values are at locations of high brightness in the source images.

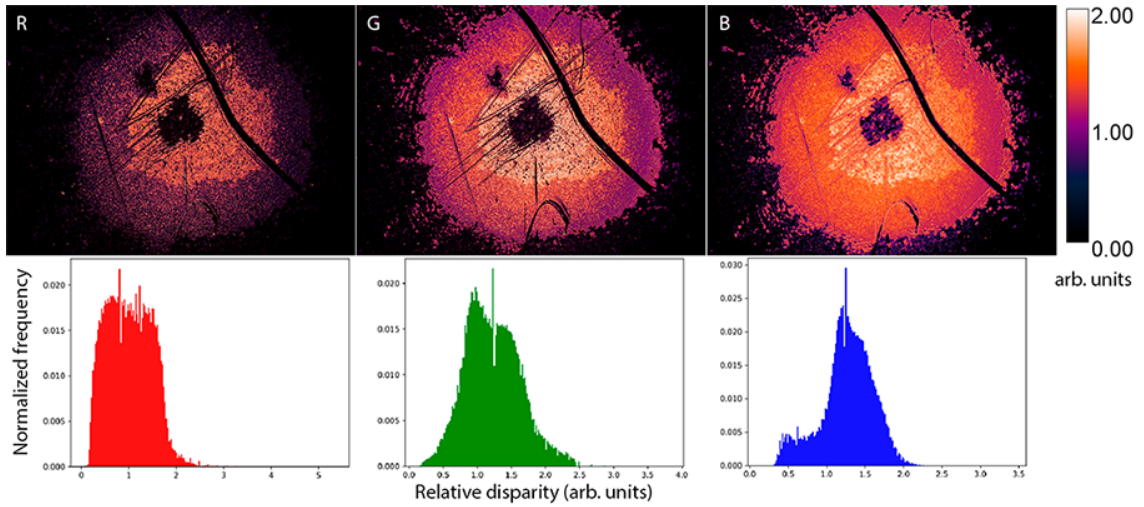


Figure 26. Relative disparity images of the 40 ms and 100 ms pair of images of sample No. 1 in figure 25. The images' look up tables are constrained to $[0, 2]$ range for viewing. Images with high linearity should have disparity values close to zero.

3.3 Recovering the ICRF

In literature there are various approaches to recovering the CRF or ICRF of a camera. Some methods work with a single image and others require multiple images of a static scene with different exposure times. Some methods recover an invertible function, while some don't. In this section we'll establish a method to recover the ICRF based on the linearity measure we proposed in part II, section 3.2.

3.3.1 Recovery from exposure series

We'll approach the recovery of the inverse camera response function by utilizing the linearity measure introduced in part II, section 3.2. In precise terms, we minimize the mean of the sum of all pixels in a disparity pair for an ordered set Ω consisting of N images $B_n(x,y)$ captured at increasing exposure times t_n . Something that should also be noted is that we need to minimize the distance from zero, therefore each subtraction operation should rather be a difference operation. Strictly speaking the mean of all disparity pairs' pixel sums is not necessary for minimization, only the total sum is important, but it makes it easier to observe and compare the magnitude of the disparity pixel sums for sets with a different number of images. At this point the reader might question why we utilize the spatial sum and not the spatial mean in this formulation. We noticed that utilizing the spatial sum instead of the spatial mean produces better results. The subject of minimization is the ICRF, so with the preceding description in mind the objective function can be formulated as

$$f(g; \Omega) = \left(\sum_{i=0}^{N-1} \sum_{j=i+1}^{N-1} \Sigma |R_{i,j}| \right) (C(N,2))^{-1}, \quad (41)$$

in which $C(N,2)$ denotes the number of possible pairs of images out of a set of N images and $R_{i,j}$ is the relative disparity image of a pair of images i and j . Expanding the objective function further yields

$$f(g; \Omega) = \left(\sum_{i=0}^{N-1} \sum_{j=i+1}^{N-1} \sum_{x=0}^{X-1} \sum_{y=0}^{Y-1} \frac{|g(B_i(x,y)) - r_{i,j}g(B_j(x,y))|}{r_{i,j}g(B_j(x,y))} \right) (C(N,2))^{-1}. \quad (42)$$

Finally by utilizing the empirical model of camera response functions we can transform this problem into finding the values of the coefficients of the principal components. We'll utilize the first five principal components in this research, like in the work of Grossberg and Nayar[45].

There is also one modification we'll make Grossberg and Nayar's model. We'll change the base function g_0 : instead of using the mean ICRF g_0 determined from the database, the base function will be set in the general form $g_0 = B^\gamma$, in which γ is a non-negative real number. It will be an additional parameter in the optimization along with the coefficients of the principal components. We add this new optimization parameter on the basis that the individual CRFs in the database vary greatly, and

argue that this general exponential is a very simple method to improve the fit with only a small increase in the computational cost.

The objective function 42 is straightforward to implement into an optimization algorithm. As this objective function utilizes a series of images of increasing exposures, the amount of processed data can quickly grow very large. To facilitate faster optimization, one might choose to sample a smaller number of pixels instead of the whole image. We will introduce a parameter in our algorithm that can be adjusted to utilize a smaller sample of pixel positions from the image stack. Calling back to the discussion of the physical meaningfulness of image brightness values in part II, section 2.3, we'll need to consider that here. We have chosen to reject pixel values within 5 values of the data range edges based on the noise characteristics of the utilized camera, which will be shown in part III, section 2.3. While not explicitly shown in equation 42 or equations 39 and 40, only values $5 \leq B \leq 250$ are utilized and values falling outside are set to NaN.

Overall this problem formulation is inspired by the work of Matsushita and Lin, reported in the article 'Radiometric Calibration from Noise Distributions'[50]. This formulation is also very similar to the one used by Akyüz and Genctav in their article 'A reality check for radiometric camera response recovery algorithms' to measure the accuracy of several ICRF recovery algorithms[48]. In their work Akyüz and Genctav utilize a very similar measure to equation 42. There are two key differences in these two formulations.

The first is that in our formulation a weighting function is not explicitly used and instead we implicitly use a rejection threshold of 5 pixel values from the data range edges, while Akyüz and Genctav use an explicit weighting function. The bigger difference is that in their work the authors utilize essentially what we established in this work as the absolute disparity, while we utilize the relative disparity. In the present work we noticed that utilizing the relative linearity measure produced more accurate results over the absolute measure and chose to focus mainly on it. At the time of formulating our optimization algorithm we had not yet encountered Akyüz and Genctav's article and were independently developing our approach. In the end the similarities are not surprising as both approaches start from the very fundamental relationship in equation 38.

3.4 The ICRF and temporal mean

As established in part II, section 2.2, the temporal mean is an important tool to mitigate noise in images in situations, where a static scene can be imaged over time. At this point the reader might have begun to wonder about the accuracy of the temporal mean when the acquired images of a camera are nonlinear. This is an important aspect to consider. As we know at this point, each acquired image used in the temporal mean's calculation is nonlinear. The value of the temporal mean is consequently skewed. This could be corrected by applying the ICRF to each image used in the temporal mean. This poses an issue with the ICRF recovery method we developed in part II, section 3.3.1, which directly optimizes the linearity measure. In the case of the mean stacked images that are used to solve the ICRF, there isn't really anything we can do about this issue in terms of linearity. However, we argue that in the case of static scenes it is reasonably justified to ignore this issue.

While pixel values that are far away from each other are useless to compare, values that are sufficiently close to each other can still be reasonably compared. We reason that if the slope, or that is the derivative, of the ICRF doesn't change too much in a short range, the pixel values are close enough to linear in a short range. For smooth ICRFs this condition should be quite accurate and therefore there shouldn't be too much of an issue with recovering the ICRF via directly optimizing the linearity measure for mean stacked images, or generally mean stacking images without linearizing them first. For a particular camera we can determine the standard deviation of the pixel values for each distinct brightness level in order to get a measure how widely the pixel values tend to spread and therefore can observe the range in which the approximation holds true. In a way we'd argue that this approach is as if we implicitly used a first order Taylor expansion for an unknown function.

After recovering the ICRF, any temporal mean could be calculated by first applying the ICRF to each frame, as that step doesn't take too much computation time. One could argue, in a strict sense, that the application of the ICRF that is being iterated in the optimization algorithm should be applied to each frame of the image stack making up the temporal mean images. This would mean that for each image a temporal mean should be calculated hundreds of times. This would be computationally so costly that the small inaccuracy caused by applying the iterated ICRF to only the final temporal mean image is the only viable choice.

4 High dynamic range images

Everyone who has used a camera has taken notice of how adjusting the exposure time can change the contrast of the image. Overexposed images are overall bright and details in the brighter portions are obscured due to saturation, while details in dark areas can be seen. On the other hand underexposed images are overall dark and details in the darker areas are hidden due to minimal camera response, while details in brighter areas can be seen as the camera’s response isn’t saturated. This isn’t a problem for your own vision, as the human eye has a far wider dynamic range than a typical digital camera. Dynamic range refers to the range of the smallest and largest possible values that a display device can present or a capture device can acquire in a single exposure[36, p. 5]. As we have noted earlier in part II, section 2.1, most devices handle RGB images as three 8-bit depth channels and thus they are capable of only 256 different levels per channel. High dynamic range (HDR) imaging can be used to surpass this limitation by using more than 8 bits per channel[36, p. 242]. This will be the second-to-last building block of the full image processing pipeline.

A HDR image can be constructed with regular low dynamic range (LDR) images with the help of known exposure times and a linearized response[39]. Capturing a static scene at different exposure times, so that every pixel of the image is captured at least once safely within the sensor range enables us to construct a HDR image. Naturally each image has to be first linearized with the ICRF[39]. As we’ve defined in part II, section 3, the ICRF maps the image from brightness domain B to exposure domain X . So we’ll denote images that have been linearized as $X(x,y)$.

With the images in a linear domain, the HDR image can be determined as a linear combination of the images captured at different exposure times. In general this process can be expressed as

$$E_{\text{HDR}}(x,y) = \frac{\sum_{n=0}^{N-1} w(\cdot)X_n(x,y)/t_n}{\sum_{n=0}^{N-1} w(\cdot)}, \quad (43)$$

in which $w_n(\cdot)$ are weights[58]. Essentially this is a general form of a weighted average, where the weights are left implicit due to the various forms they appear in

literature. We choose a very similar weighting function as in the work of Robertson et al. – a Gaussian function[59]. The weighting function is a normal distribution centered about the middle pixel value. We’ll also leave out the normalizing multiplier to set the value of the weighting function to 1 at the middle pixel value. The weight is given by

$$w(B) = \exp\left(-\frac{1}{2}\left(\frac{B - \mu}{\sigma}\right)^2\right), \quad (44)$$

with $\mu = 0.5$ and $\sigma = \sqrt{15}/30$. The form of the Gaussian isn’t exactly the same as in the work of Robertson et al. as our parameters lead to a slightly more tightly centered Gaussian. The parameters were chosen by experimenting with different values and with the aim to strongly minimize the weight of values close to the data range edges. The purpose of the weighting function is to give more weight to pixel values that are towards the center of the pixel value range. This is done because they are a more reliable measure of the exposure than pixel values at the range’s edges, as stated in part II, section 2.3. With this weighting function the pixel values at the data range’s edge, for example, get a weight of about 5.53E-4, while pixels with values 0.2 and 0.8 get a weight of about 0.067.

Naturally we use the base images $B_n(x,y)$ in the weighing function, as they are the actual measured values, rather than the linearized images $X_n(x,y)$. The weighting function is an important choice in terms of the accuracy of the final HDR image. In literature there are various choices utilized. Debevec and Malik in their work utilize a simple hat function[39], Robertson et al. use a Gaussian function[59], which is the approach we have adapted here.

With this choice of weight the HDR image can be computed as

$$E_{\text{HDR}}(x,y) = \frac{\sum_{n=0}^{N-1} w(B_n(x,y))X_n(x,y)/t_n}{\sum_{n=0}^{N-1} w(B_n(x,y))} \quad (45)$$

from which one can observe the intuitive trend that pixels with shorter exposures imply a greater irradiance and vice versa for longer exposures. This establishes a relative irradiance scale in the image[39], where unity becomes assigned to pixels with certain combinations of original brightness values and exposure times. The

uncertainty of the HDR image is given by

$$\delta E_{\text{HDR}} = \sqrt{\sum_{n=0}^{N-1} \left(\left(\frac{w'(B_n)g(B_n) + w(B_n)g'(B_n)}{\sum_{n=0}^{N-1} w(B_n)} - \frac{w'(B_n)w(B_n)g(B_n)}{\left(\sum_{n=0}^{N-1} w(B_n)\right)^2} \right) \frac{\delta B_n}{t_n} \right)^2}. \quad (46)$$

, whose derivation is discussed in appendix A and where we omitted the pixel notation (x,y) for the images. We utilized the uncertainty propagation law[52] for the N independent variables to derive equation 46.

5 The linear effects

With the ability to linearize image brightness, we can now focus on correcting the linear effects of the imaging system. As shown in the flow diagram in figure 23, the correction of the optical effects is something that the camera typically does right after acquiring the raw sensor data. In our case we disable as many of the camera's own operations as possible, opting to do the corrections ourselves for a better understanding and quantification of the operations. Before delving deeper into this section, we'll note that while we write the equations in terms of exposure domain images, the equations are equally applicable to irradiance domain images, as both are linear domains. The work in this section will act as the final building block in the full image processing pipeline.

When capturing a scene with a camera, the resulting image has various artifacts in it caused by different physical phenomena[51]. For quantifying radiance from an image, this poses an obvious problem in terms of getting accurate results. Knowing the origin and the effect of these artifacts on the images provides us the understanding and methods to correct them to an extent. In this section we'll go over the important artifacts, their sources and how they are corrected. Generally the image capturing process of a CCD or CMOS camera can be expressed as

$$X_A(x,y) = X_{\text{BIAS}}(x,y) + X_{\text{THERM}}(x,y) + r(x,y)X_0(x,y), \quad (47)$$

in which $X_{\text{BIAS}}(x,y)$ is the camera noise that isn't dependent on temperature, while $X_{\text{THERM}}(x,y)$ is the thermally dependent noise, $r(x,y)$ is the response matrix that accounts for the non-uniformities of the camera response and finally $X_0(x,y)$ is the true image, which has no contributions from the bias or thermal signals or non-uniformity[51].

5.1 Dark current and hot pixels

Even in a pitch black environment charge is generated in the wells of the sensor's pixels[36, 51]. This causes the image to have pixel values larger than zero, when zero would be the conventionally expected value. Sensor imperfections can cause further related artifacts, Allen and Tiantaphillidou give a concise description for these[36, p. 171]. They describe some pixels as hot pixels, which generate charge at a faster rate than others and saturate easily. It follows that they appear as brighter pixels than their typical neighbors. Similarly the authors describe some pixels as dark pixels, that have a lower response than their typical neighbors[36, p. 171]. It follows that they appear as darker than most pixels. Large differences in charge generation rate can leave behind fixed patterns in the image. The amount of charge generated not only increases with rising temperature but naturally also with increasing exposure times. The effect of dark current can be mitigated by using calibration images[51]. By capturing an image with the camera shutter closed you acquire an image which is ideally only affected by the dark current. By averaging over an ample amount of such images you can create a dark frame calibration image at the same exposure time as the image to be corrected. The effects of dark current can then be easily corrected by subtracting the dark frame from the acquired image, that is

$$X_C(x,y) = X_A(x,y) - X_D(x,y), \quad (48)$$

in which $X_C(x,y)$ is the corrected image[51] and $X_D(x,y)$ is the dark frame containing both the contribution of X_{BIAS} and X_{THERM} [51]. The step represented by equation 48 is usually handled by consumer cameras by an analogue pre-processor[36, p. 274]. The thermal component is valid for only the specific temperature and exposure time combination and a standard approach is to acquire a mean stacked master dark frame at the same operating temperature and exposure time[60]. There are also methods to scale the thermal component[60] and even handle corrections for consumer cameras after image acquisition[61].

In the case of this research, we'll utilize the camera's PC software interface to perform the dark frame subtraction, as the obtained data is not in a RAW format. Performing the subtraction outside would either reduce the usable data range for us or the values would need to be scaled to utilize the full range again, which would introduce a lot of uncertainty. The software can be set to capture a dark frame

calibration image, which is subtracted from a captured image within the software before the user receives an output image. The software performs the dark frame calibration with the camera sensor blocked and the process takes about a minute. The user can output the calibration image from the software, which is an 8-bit gray scale image. Most likely this means that the software holds a proper RAW image within, which it utilizes in the dark frame calibration before any demosaicing or any nonlinear operations take place. It is difficult to determine in certain terms if the software has different dark frames for different exposure times or just one at a long exposure time, which it then scales depending on the exposure time of the image under correction. Regardless of these details, the dark frame calibration performed by the software yields high quality images with no contribution of the bias signal X_{BIAS} . In the output images only some hot pixels can be observed at long exposures, suggesting that only the bias signal is subtracted.

To deal with the hot pixels left in the image we'll utilize a filtering process. We can utilize manually captured dark frames, that have not been subjected to internal dark frame subtraction, to map out positions of hot pixels by looking for pixels with values above a given threshold. Following that we can apply a filter on those pixel positions in a subject image to replace their value with a more reliable value. We utilize a 7x7 pixel median filter. We'll utilize a simple global threshold, but there are also implementations to detect hot pixels that utilize a localized approach[61].

5.2 Non-uniformity

Imagine a perfectly flat evenly illuminated white scene, in which every point radiates at the same intensity. Ignoring the previously mentioned noise, capturing an image from such a scene should ideally result in an image in which all the pixels have the same value. In reality the acquired image has a darkened periphery and possibly other dark smudges or halos. The radial darkening is known as vignetting. Basically the farther the pixels are from the center of the image, the darker the pixels are. This poses an obvious problem when one tries to quantify brightness or threshold an image accurately. More precisely, there are three possible sources for vignetting. Ray gives an effective description of these in his book as follows[62, pp. 131–133]. Natural vignetting is caused simply by the geometry of the optics and sensor and the illumination intensity is proportional to $\cos^4\theta$, in which θ is the angle of incidence of the light hitting the sensor. Mechanical vignetting is caused by physical obstructions

from a lens intruding into the field of view, lens shades or filters and their holders. Optical vignetting is caused by an effective reduction in aperture diameter for oblique rays. In addition to these the previously mentioned hot and dark pixels can leave fixed patterns in images and finally spatial variations in the quantum efficiency of the sensor can overlay a spatially random, but temporally fixed, pattern on the image[63].

These non-uniformities are represented in equation 47 by the response matrix $r(x,y)$, which can be determined experimentally[51]. By capturing a set of images of an uniform evenly lit surface and averaging the set, a flat-field image $X_{FF}(x,y)$ is produced. With an evenly lit uniform surface we can reasonably assume that each pixel of the sensor should ideally be hit by the same number of photons and therefore the image shows the systematic non-uniformities only. The correction can be performed by first calculating the spatial mean of the flat field image, which can then be used to express the response matrix [51]. The response matrix is given by[51]

$$r(x,y) = \frac{X_{FF}(x,y)}{\overline{X_{FF}}}. \quad (49)$$

Solving for the true image $X_0(x,y)$ in equation 47 yields

$$X_0(x,y) = \frac{X_A(x,y) - X_{BIAS}(x,y) - X_{THERM}(x,y)}{r(x,y)}.$$

Substituting equation 48 and taking note of the considerations of the bias frame and dark frame 49 then yields the final expression for the true image in terms of images that can be experimentally determined

$$X_0(x,y) = \frac{\overline{X_{FF}}}{X_{FF}(x,y)} [X_A(x,y) - X_D(x,y)]. \quad (50)$$

Finally applying the bad pixel filtering to $X_0(x,y)$ yields a properly corrected image. Something that should be noted is that in the case of short exposures and visible light, the contribution of the thermal component can be very small and may be ignored[51]. We have also noted from our own dark frames that bad pixel filtering can be ignored for short exposures, at least for the camera used in this research.

As noted earlier in part II, at the end of section 5.1, we utilize the dark frame calibration of the camera software in this research. In terms of equation 50 this means

that the term related to dark frames $X_D(x,y)$ is dropped, as the acquired image $X_A(x,y)$ will already be corrected for it. So the final equation for the corrections in our case will be

$$X_0(x,y) = \frac{\overline{X_{FF}}}{X_{FF}(x,y)} X_A(x,y). \quad (51)$$

5.3 Practical considerations of the flat-field image

The capturing conditions of an evenly lit uniform surface can be difficult to achieve in practice. While the microscope does provide a very high quality illumination, the smaller the distance scale, the harder it is to find an uniform surface that doesn't show local detail. Additionally the flat-field image should be captured not only for each objective, but also in both BF and DF illumination. Considering the nature of DF illumination we discussed in part I, section 3.4.2, if a flat and formless uniform surface is used, the image will be completely dark. This leaves us with two conflicting conditions.

This problem can however be overcome by building a composite image. In their article Brady and Legge determine the vignetting of a camera by translating a small object across the image diagonal and capturing images of it at different positions[64]. They argue that these images can be composited to provide a flat-field image. In their case this only provides data along the diagonal in their case and assumes symmetric non-uniformity. The same idea could be used here with a microscope by translating a selected portion of a surface over the whole image plane. This process could be made conveniently automatic with a motorized microscope stage, but in the case of this research our microscope isn't equipped with one. This process would require an exorbitant effort by hand as it has to be done for each objective in both BF and DF mode and it wouldn't work for asymmetric defects in the image.

In our case we propose to use an altered version of the above method. We use a piece of paper as a subject and captured an image series of the paper as the microscope stage is constantly moved and therefore the subject. We argue that by capturing enough images like this, on average, different features of the paper should have been imaged at different positions equally, while the system non-uniformities remain static. Additionally the movement of the microscope stage causes motion blur in the image, further obscuring local details of the surface. Mean stacking the image series should then yield an image that approximates a flat-field image reasonably

well. While this is a more approximate method than imaging a single object on the diagonal, it also provides information over the whole image area and therefore doesn't make assumptions about the symmetry of the non-uniformities.

Another point worth taking into consideration is the spatial mean of the flat-field image \bar{X}_{FF} and how it is used in the correction. Looking at equation 50, you can see that the spatial mean and the flat-field image $X_{FF}(x,y)$ together are used as a kind of normalizing map. The term darkens pixel locations that are brighter than the spatial mean and brightens those that are darker. This is generally the best approach when the lighting itself is truly uniform and the image might have both brightening and darkening non-uniformities.

In the case of the optical microscope used in this research, the system has no brightening non-uniformities and the lighting itself is at its best in the center of the image, especially in the case of DF mode, whose lighting is less uniform than in BF mode. In this case, considering that all the non-uniformities in the images do is darken them, then why darken some portions further with the correction by averaging over the whole image area? With these considerations in mind for this case, heuristically speaking, an image that is corrected for vignetting and other non-uniformities should be at least as bright as the original image at all locations. Taking \bar{X}_{FF} from the whole image means that the value is affected by the darkening non-uniformities. We argue that if the value was instead taken from a central portion of the image where the illumination is the most ideal, it would yield a physically more accurate value. This train of thought is naturally problematic in a situation in which the optical system has significant defects in the center of the image, but that is not the case for our system.

5.4 Uncertainty

Now that the systematic effects of the imaging system are known and can be corrected using equation 50, we can also quantify the uncertainty in the final true image. The uncertainty can be evaluated by applying the uncertainty propagation law into

equation 50, yielding[51]

$$\begin{aligned} \delta X_0^2(x,y) = & \delta \bar{X}_{FF}^2 \left(\frac{X_A(x,y) - X_D(x,y)}{X_{FF}(x,y)} \right)^2 + \delta X_D^2(x,y) \left(\frac{\bar{X}_{FF}}{X_{FF}(x,y)} \right)^2 \\ & + \delta X_{FF}^2(x,y) \left(\bar{X}_{FF} \frac{X_A(x,y) - X_D(x,y)}{X_{FF}^2(x,y)} \right)^2 + \delta X_A^2(x,y) \left(\frac{\bar{X}_{FF}}{X_{FF}(x,y)} \right)^2. \end{aligned} \quad (52)$$

In equation 52 each δX term represents, in order, the uncertainties of the final image $X_0(x,y)$, the bias frame $X_B(x,y)$, the flat field image $X_{FF}(x,y)$ and the subject image X_A . The uncertainty images are defined as in equation 33 and the uncertainty of the flat-field mean \bar{X}_{FF} is defined as the spatial mean of $\delta X_{FF}(x,y)$, as was discussed in part II, section 2.4.

For the case in our research we need also the uncertainty of the correction equation without the dark frame term. Applying the uncertainty propagation law to equation 51 yields

$$\begin{aligned} \delta X_0^2(x,y) = & \delta \bar{X}_{FF}^2 \left(\frac{X_A(x,y)}{X_{FF}(x,y)} \right)^2 + \delta X_A^2(x,y) \left(\frac{\bar{X}_{FF}}{X_{FF}(x,y)} \right)^2 \\ & + \delta X_{FF}^2(x,y) \left(\bar{X}_{FF} \frac{X_A(x,y)}{X_{FF}^2(x,y)} \right)^2. \end{aligned} \quad (53)$$

In equation 53 $\delta X_A^2(x,y)$ carries the uncertainty information of the acquired image and dark frame subtraction together. For the exposure images $X_A(x,y)$ this is simply evaluated as the uncertainty image $\delta X_A(x,y)$ of an averaged image stack after linearization. For the irradiance domain image $E_A(x,y)$, $\delta X_A(x,y)$ is simply replaced by $\delta E_A(x,y)$ as computed via 46.

6 Other image-based measurements

Accurate calculations with the brightness values of images, at least in the sense of assigning physical meaning to them, requires linearity. There are, however, some calculations that don't need such calibrations and can in fact be done without too much additional processing. Though the benefit of linearity and especially the expanded bit depth of 64 in the case of the HDR images does afford more control and precision.

6.1 Autocorrelation

An interesting property in signal processing is the autocorrelation function, which for a signal measures how similar at time t it is to itself at time $t + \delta t$ [65]. In the most simple case of one-dimensional functions, the autocorrelation of $f(u)$ is given by[66]

$$\gamma(x) = f \star f = \int_{-\infty}^{\infty} f(u)f(u \pm x)du. \quad (54)$$

The value of the autocorrelation is equal for both signs of x in equation 54 and has its maximum at zero[65]. These both are a logical result as naturally the function f is similar to the unshifted version of itself and a shift of one of the functions in a given direction in the convolution operation can be written as a shift of the other function in the other direction[65, p. 63]. Autocorrelations are simpler to compute for signals with a zero mean value and it is common for the autocorrelation to be normalized by the signal's standard deviation so that its values will fall in the range $[-1,1]$. [65] In this work we will subtract the mean value from the image and normalize using the standard deviation.

Positive autocorrelation means the signals at the two points in time are similar, zero means no correlation and finally negative means anticorrelation[65]. Autocorrelations can be generalized to more dimensions and therefore applied to images, where it can be used to measure periodicity in images, which in turn can be used to measure how ordered features in an image are. In the case of images, which have discrete and bounded data instead of continuous and infinite, the autocorrelation

can be calculated by

$$\gamma(a,b) = \sum_x^{X-1} \sum_y^{Y-1} I(x,y)I(x-a,y-b), \quad (55)$$

in which X and Y are the number of columns and rows in the image respectively, while a and b represent the distance from the corresponding x and y positions[67].

Autocorrelations in the case of images can be calculated in various ways. The true 2D spatial autocorrelation can be used to output an array of values, which can also be interpreted as an image, as can be seen from equation 55. Figure 27 shows an example of a source image and its 2D autocorrelation. The 1D autocorrelation can be applied to different rotations in the image, such an example is shown in figure 28. The 2D spatial autocorrelation function can also be radially averaged to produce results similar to the 1D autocorrelation. These two are actually identical to each other in a case where the image is isotropic[67]. The autocorrelation is rarely calculated via the definition given in equation 55, but instead Fourier transforms are utilized. This reduces the processing time from $O = n^4$ to $O = n \log n$ [67].




Figure 27. The top image shows the logarithm of a relative radiance image and the bottom shows the autocorrelation image of the relative radiance image. The autocorrelation values are mapped to display values from the range $[-0.05, 0.2]$ to $[0, 255]$ for viewing. The original image has obvious periodic features that result in strong autocorrelation in certain directions and at certain distances.

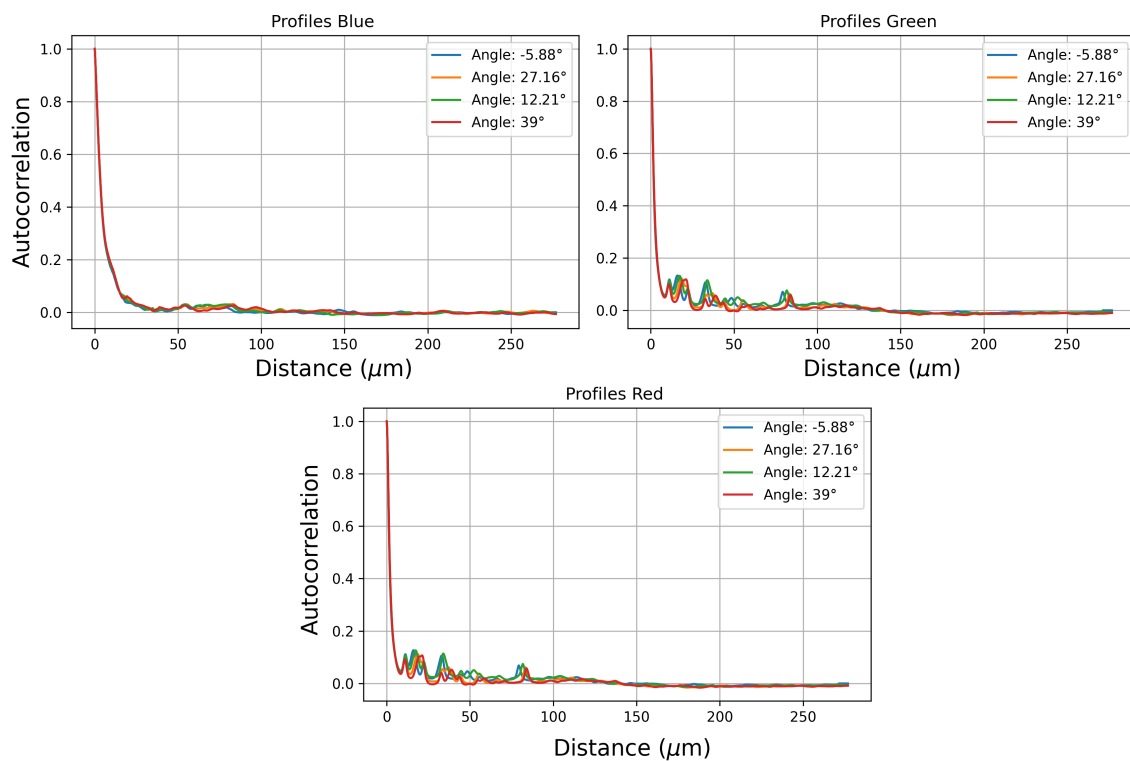


Figure 28. Line profiles plotting the autocorrelation values of the image shown in 27 in indicated directions. Positive angle is measured in counterclockwise direction from horizontal right. The example profiles show some significant peaks in the green and red channels, while the blue channel decays to insignificant autocorrelation quite quickly.

6.2 Local thickness

Local thickness acts as a measure of distance between features of interest or thickness of those features in an image. The algorithm that is used to determine local thickness here is applicable to 2D and 3D images. In the case of 3D data we would talk about spheres and voxels, while in our case of 2D data we'll talk about circles and pixels. First an image must be thresholded into a binary image so that the features of interest are set to one and the space between them to zero. Following the thresholding, for each high pixel a circle of maximal radius is then fitted inside the region so that it contains the point and fits inside the region. The radius of that circle is then set as the value of that pixel.[68]

The algorithm that calculates the local thickness unsurprisingly does not utilize the definition directly. In this research we'll utilize the open source Local Thickness plugin of ImageJ, developed by Dougherty[68]. The algorithm is also relatively complex and any in depth analysis is beyond the scope of this research. Based on Dougherty's work we can paraphrase a short description[68]. The algorithm first calculates a distance map, which sets a pixel's value to the diameter of the largest circle centered on that pixel and fits inside the region. The distance map is then transformed into a distance ridge, which removes redundant datapoints from the distance map. Finally using the remaining non-redundant circles, for each high pixel the distance ridge is scanned to find the largest circle that contains that pixel. The documentation and a link to the GitHub page of the plugin is available on ImageJ's webpage[68].

Dougherty's work is largely based on algorithms discussed in Hildebrand and Rügsegger article[69], Saito and Toriwaki's article[70], and more loosely based on Remy and Thiel's article[71]. Local thickness calculations, especially their 3D version, are a very useful post processing method in any application that works with 3D volumes, such as computed tomography.

7 Full image processing pipeline

With all the pieces of the image handling system in hand, the full pipeline is visualized in figure 29. Basically operations and analysis can be performed for any of the images after the linearization via the camera's solved ICRF. For the sensor exposure estimate $X(x,y)$ one has to take note of any values that were too close to the data range edges in the original brightness domain image $B(x,y)$. After the HDR merging, with sufficient data there is no need to worry about the data range edges. If every pixel was imaged at least once sufficiently far away from data range edges, then the irradiance domain image $E(x,y)$ is a reliable estimate as is also the radiance domain image $L(x,y)$, based on our discussion in part II, sections 2.4 and 5.

The bad pixel filter is also something that can be skipped when the utilized images don't have very long exposure times. This is due to the effect of the dark current being relatively insignificant at short exposures, as outlined in part II, section 5.1. Exposure domain images $X(x,y)$ can be used to compare pixel values within a single image even across channels – with the caveat of the data range edges, fixed-pattern effects and wavelength response of the imaging system. Though this will still be less reliable than utilizing a HDR image. The point of the linearization is to get rid of the nonlinearities that invalidate arithmetic operations and therefore comparisons of pixel values and therefore enable the accurate merging of differently exposed images into a single HDR image.

With a proper number of differently exposed source images, the irradiance domain images $E(x,y)$ and the radiance domain images $L(x,y)$ get rid of the data range edge problems still present in the exposure domain images $X(x,y)$. While the exposure domain image samples a possibly small portion of the scene radiance, the irradiance and radiance domain images can encode the full dynamic range of the scene in a relative scale, effectively getting rid of data range edges[39]. We argue that it is precisely due to the fact that the HDR images encode the full dynamic range, that comparisons across color channels can be made. The caveat of cross-channel comparisons is the wavelength response of the imaging system, which means that comparing relative radiance images obtained on different devices cannot be made at

least directly.

The autocorrelation, local thickness and relative radiance measurements we'll be performing in part III don't require much additional processing. For the autocorrelation we subtract the spatial mean from the input image and divide by its standard deviation in order to normalize the resulting autocorrelation image. After that the image is ready for the autocorrelation computation. For the relative radiance measurements we need to possibly manually select a rectangular area for analysis and we'll manually select a threshold value to segment the features of interest from the images. The formed binary mask can then be used to compute spatial statistics to measure relative radiance values for features of interest. The local thickness measurements have a similar workflow in that we'll need to segment the features of interest. Similarly we might need to manually select a rectangular area in an image and we'll also use thresholding to select features of interest. For the local thickness measurement of the relative radiance images we can use the relative radiance measurements to guide the choice of threshold value. In the case of local thickness the binary mask is also slightly improved by running a single binary close operation on it to minimize the amount of artificial gaps and small specks in the mask, after which the image is ready for the local thickness algorithm. Figure 29 also shows the preprocessing used on the input images for the different measurements.

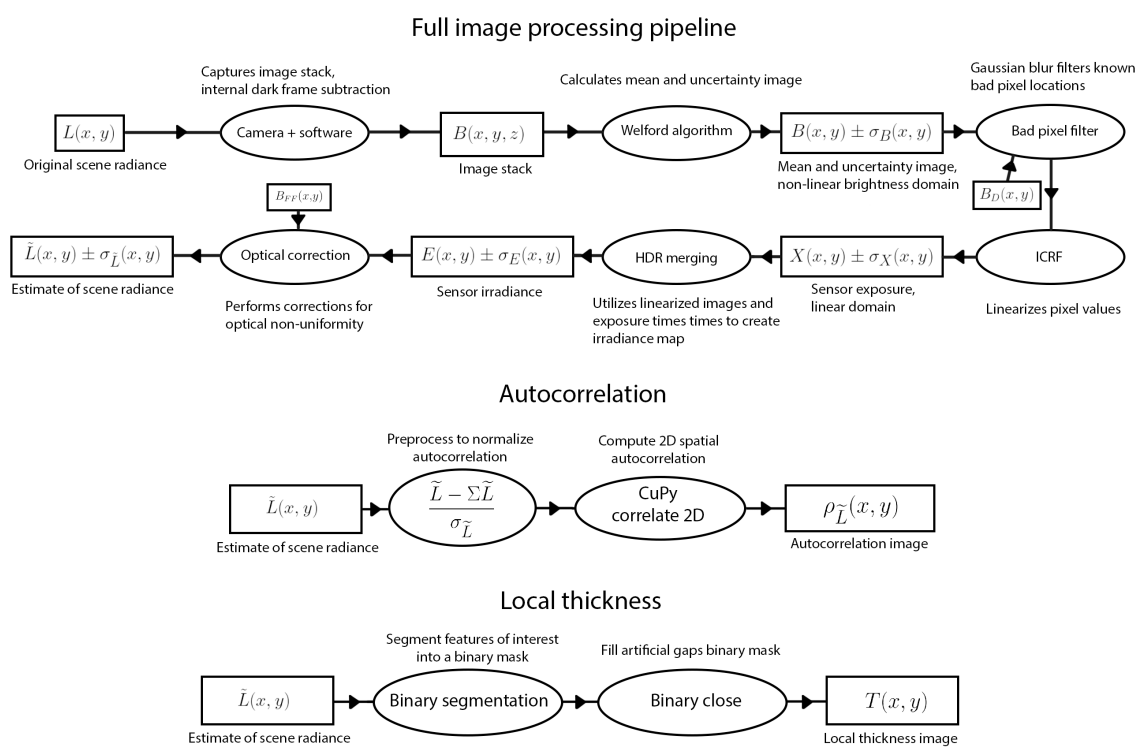


Figure 29. Flow diagram of the image handling pipeline. And possible subsequent autocorrelation and local thickness processing.

Part III

Measurements and results

The last part of this thesis is focused on applying the knowledge and methods we've discussed in parts I and II. Firstly in section 1 we'll go over our algorithm for solving the ICRF, how it performs and find the optimal parameters for it. A reader interested in only the results can skip this section and jump right into the calibration of the Olympus SC100 camera.

Section 2 is focused on the calibration of the utilized digital camera. More specifically its distance scale calibration, white balance, dark frames, flat-fields and linearization. Secondly we'll discuss and analyze how well the linearization and correction methods perform. We finish the section with the construction of HDR images, that are the relative radiance images. We'll observe the quality and linearity of them by constructing two relative radiance images with only one difference – their illumination intensity. We'll utilize an analyzer to modify the illumination, while keeping other conditions static. The analyzer causes the unpolarized light's intensity to be cut in half, as was discussed in part I, section 2.1.3. This gives us a method to quantify the linearity of the images and finally answer the question, if commercial digital cameras can be used as radiance sensors.

Following the calibrations, we are finally ready for the measurements related to the samples in section 3. The first is the relative radiance measurements, with which we will gain insight into how the relative radiances change when moving from the bright zone to the dim zone. Secondly we will utilize local thickness measurements for both the relative radiance images and SEM images. These will help us determine the typical sizes of features of interest and the distances between them. These will be critical in terms of explaining the optical properties of the MWNT depositions. Finally we'll perform autocorrelation measurements to observe if the samples have periodic or anisotropic features. These will also give additional measurements for the typical particle sizes in the images. These will also be performed for both the relative radiance and SEM images. The measurement results will be analyzed with the aim of gaining insight into the nature of the MWNT depositions. We are aiming to gain answers to the question of why the samples look as they do and, finally concluding this thesis.

1 Solving the ICRF

Before getting into properly linearizing the images of interest, we'll test our exposure series method of solving the ICRF. After the ICRF is solved we'll be ready to linearize images and deal with the linear effects of the imaging system and to construct the estimates of a scene's radiance.

1.1 The minimization algorithm

The first five principal components for each color channel were calculated by utilizing the database of camera response functions. While the database contains also precalculated principal components, these were calculated for all color channels together. We opted to calculate the principal components for each channel separately by ourselves. This was done by first inverting the functions from the database and separating them based on their color channel. Then a mean ICRF and a covariance matrix for each channel was calculated, followed by determination of the principal components.

The algorithm we chose to minimize the objective function introduced in part II, section 3.3.1, is the differential evolution algorithm and more specifically the implementation in the Python SciPy library. The differential evolution algorithm is a population-based, stochastic, evolutionary algorithm that can be utilized for global optimization. The algorithm initializes a population of candidate solutions and at each iteration mutates each candidate by mixing them with each other based on different strategies[72].

We chose this algorithm for its simplicity, efficiency and because the algorithm doesn't utilize gradients of the objective function. The differential evolution algorithm in SciPy utilizes various parameters to control the process. We'll briefly go over the most important of them and report our choice of parameter values. Probably the most important parameter is the strategy. We chose to utilize the current-to-best/1 strategy, which is a strategy that compromises between fast convergence and exploration of the parameter space[72]. The next two important parameters are the

mutation and recombination constants. These control, respectively, the search radius and how likely it is for a candidate solution to progress into the next generation of the population. For the mutation constant we chose to utilize values in the range $[0, 1.95]$ and dithering, which randomly changes the mutation constant on a generation by generation basis. For the recombination constant we set its value to 0.4 so that convergence is sufficiently slow in order to not prematurely terminate the algorithm. The initial population is formed by using SciPy's 'sobol' method, which aims to maximize the coverage of the parameter space for the population.

All in all, we utilized the differential evolution algorithm to iterate the coefficients for each principal component and the exponent of the base function, which were utilized to construct the ICRF, as expressed in equation 36. As stated earlier, the optimization can be performed by either utilizing the mean ICRF or an exponential, whose power is an extra parameters subject to optimization, as stated in part II, section 3.3.1. Additionally two constraints were added; one to maintain the iterated ICRF in $[0,1]$ range and one to maintain a continuously increasing ICRF.

Finally the database of camera response functions includes also various gamma curves as possible response functions. In the optimization algorithm the user can choose whether to include these in the determination of the principal components or not.

1.2 Optimal parameters for the linearity measure

There are a number of parameters, whose impact on the linearity measure and therefore the solved ICRF we need to determine. The first is the number of sampled pixels. The computation time increases very quickly as the number of pixels is increased, while it should yield a more reliable linearity measure due to the increased number of data points. The second parameter is the rejection threshold for including pixel values in measuring linearity. Essentially the rejection threshold value describes how many pixel values at the 8-bit data range's both edges are rejected. This parameter's effect shouldn't be too large in terms of computation speed, but is very crucial in terms of accuracy. Increasing the rejection threshold allows the minimization process to focus more on the most reliable middle values of the data range, which should lead to a more accurate ICRF. On the other hand, in terms of the final HDR-merged image, increasing the rejection range means that more input images are required in order to ensure that all pixels on each color channel are

imaged at least once within the accepted pixel value range.

Readers might also be wondering how the effect of these parameters can be measured if the underlying minimization algorithm is stochastic in nature. The differential evolution algorithm's implementation in SciPy provides the option for seeded minimization. This means that we can provide a seed for the process allowing for repeatable minimizations. We will utilize an integer seed with value 1995.

We'll first inspect the effect of the number of sampled pixels. In figure 30 are shown plots of the mean disparity values for the utilized data set against the number of sampled pixels. The rejection threshold has been set to 5 pixel values at both the lower and upper range of values, i.e. only values $5 < B < 250$ are accepted. From the mean absolute disparity plot we can see that there is only small differences upon varying the number of utilized pixels. The relative disparity plot shows a bit more differences over the mean absolute disparity plot. The relative disparity plot suggests that using at least a ratio of over $1E-04$ would be sufficient, at least when accepting pixel values in the range $5 < B < 250$. Above this ratio the values are reasonably consistent for different ratios. The Olympus SC100 camera has a bit over 10.5 million pixels in an image, so using over 1000 pixel positions seems sufficient. With stricter rejection thresholds more pixel positions should be included, as more pixels will be rejected.

The mean standard deviation of both the absolute and relative disparities is very consistent over the range of utilized pixels, especially the relative one. This means that the spread of disparity values is consistent over different pixel sampling sizes and that the underlying pixel values are samples of the same distribution. This is an expected result, as taking multiple sufficiently large samples from the same population should yield a similar sample standard deviation for each sample.

The second varied value is the rejection threshold, for which mean disparity values for the data set are plotted against the threshold value in figure 31. For these the ratio of utilized pixels was set to about $3E-3$, meaning about 30000 pixel positions were utilized. This number naturally decreases based on the rejection threshold. The most noteworthy of these plots are especially the mean relative disparity and the mean standard deviation of relative disparity. The figures show that the mean relative disparity and the mean standard deviation of relative disparity decreases for an increasing rejection threshold. This can be attributed to two things. The first is that the pixel values at the data range's edges are the most unreliable due to

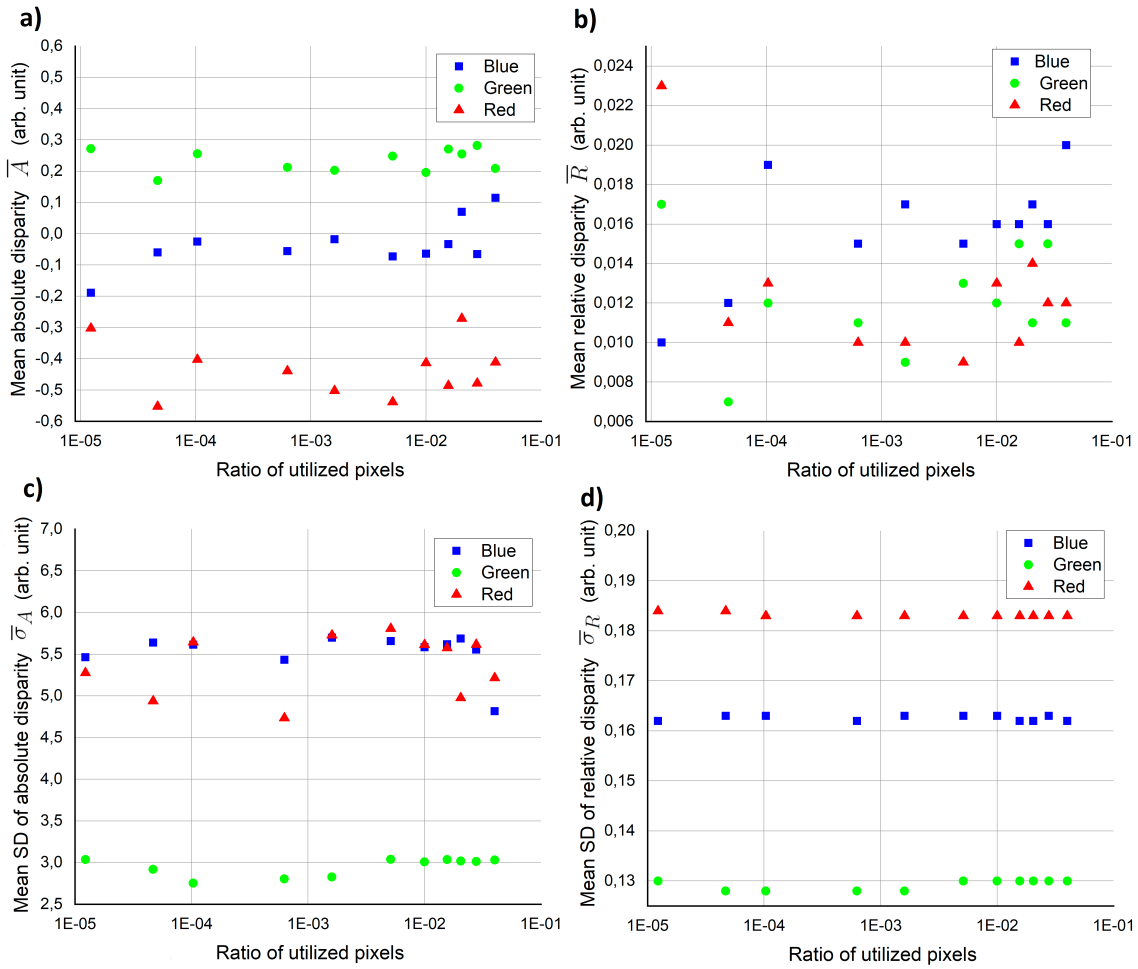


Figure 30. The mean a) absolute disparity, b) relative disparity, c) standard deviation of absolute disparity and d) standard deviation of relative disparity for each color channel as functions of ratio of utilized pixels.

cut-off and saturation. The second is that the minimization algorithm can essentially focus on minimizing the disparity values for a smaller range of pixel values. Later in part III, section 2.3 we'll discuss the noise of the imaging system. In figure 39 it can be seen that overall the standard deviation for each pixel value level is quite consistent over the data range for the blue and green channels, staying at about the 4 to 6 range. The red channel is more noisy with up to a peak standard deviation of over 12. Looking at the mean relative disparity plot in figure 31b), increasing the rejection threshold to about one standard deviation decreases the disparity a lot. Increasing it further still decreases it, but yields diminishing returns. The changes in mean absolute disparity and the standard deviation of absolute disparity seem to flatten above a rejection threshold of 4. All in all these suggest that a rejection

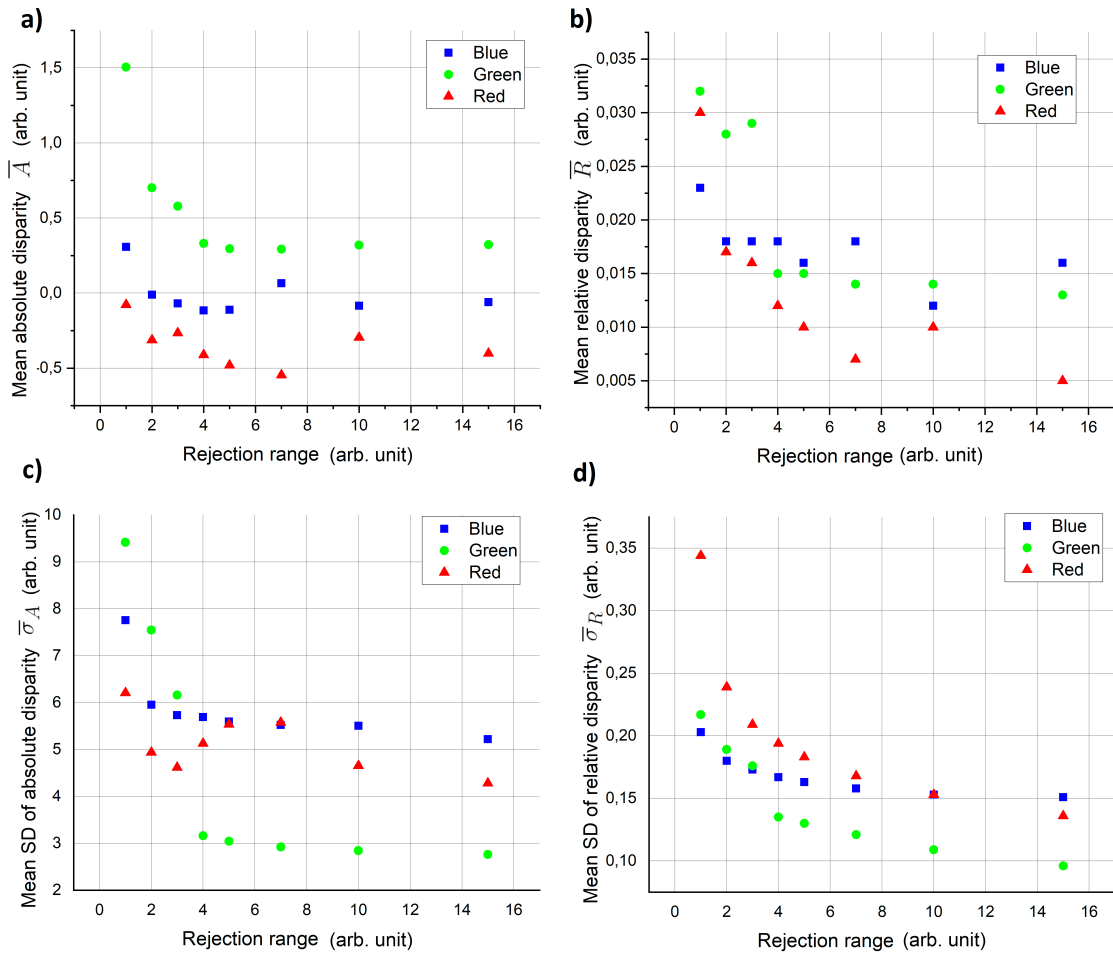


Figure 31. The mean a) absolute disparity, b) relative disparity, c) standard deviation of absolute disparity and d) standard deviation of relative disparity for each color channel as functions of the pixel value rejection range.

threshold of 5 is sufficient. Using too large of a rejection threshold means that a lot more input images are needed, while the benefits it yields are small. A large rejection threshold can even become very problematic, as images with very different exposure times won't have any pixel positions in both images at which the pixel values in each image fall within an acceptable range. This means that no data could be obtained from such an image pair in terms of computing the linearity measure.

In the end the parameter choices are a compromise between effort, the best results and computation time. In figure 32 are plotted the computation times for solving the ICRF with the parameters utilized in the two previous cases. The increase of computation time for increasing the number of utilized pixels is reasonably linear and entirely expected. What isn't as expected however is the increase in computation

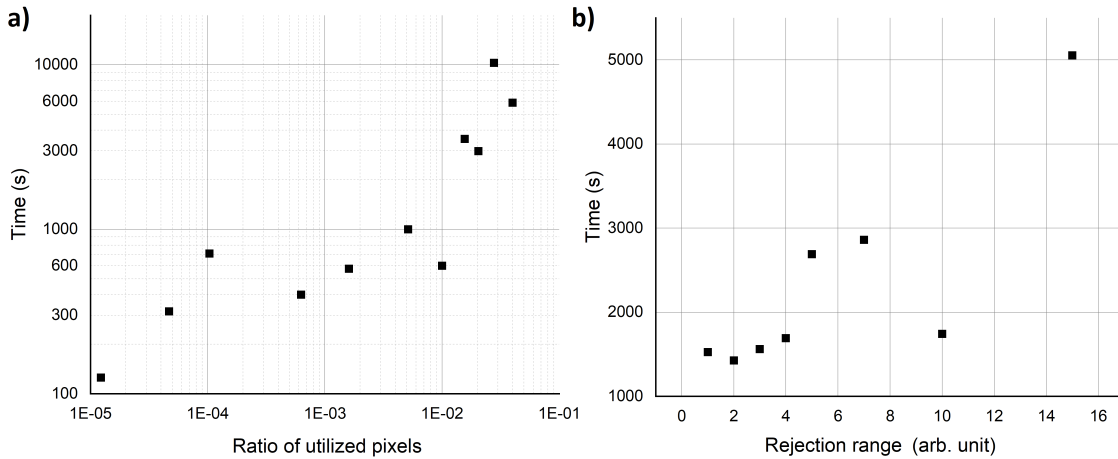


Figure 32. Computation time plotted against a) the ratio of utilized pixels and b) the rejection range. Computations were performed on a Ryzen 5 5600x CPU @3,7 GHz.

time for increased rejection thresholds, which too seem to increase linearly with an increasing rejection threshold. One thing that can be improved is the fact that the algorithm could be implemented so that it would use a dedicated GPU. As most of the operations regarding the images are element-wise operations that can be run in parallel, using a GPU could decrease the computation time by orders of magnitude.

1.3 Optimal principal components and variability of minimization

In addition to finding the optimal parameters for the linearity measure, we also need to determine the optimal principal components. The precomputed principal components in Grossberg and Nayar’s database include gamma curves and there is no separation for color channels, though the CRFs in the database do report the color channel. We can compute different principal components based on how we utilize the database. The choice of including or ignoring the gamma curves and whether to split the CRFs based on color channels or not can each be a beneficial or detrimental choice in terms of finding the most accurate ICRF. We’ll observe these choices’ effects on the result by running the minimization process for each combination of decisions, yielding four different sets of results. We’ll also use this as an opportunity to observe the variability of the resulting ICRF when the minimization is unseeded. This is done by running the minimization for each of the four different sets of principal

components 20 times and observing the differences.

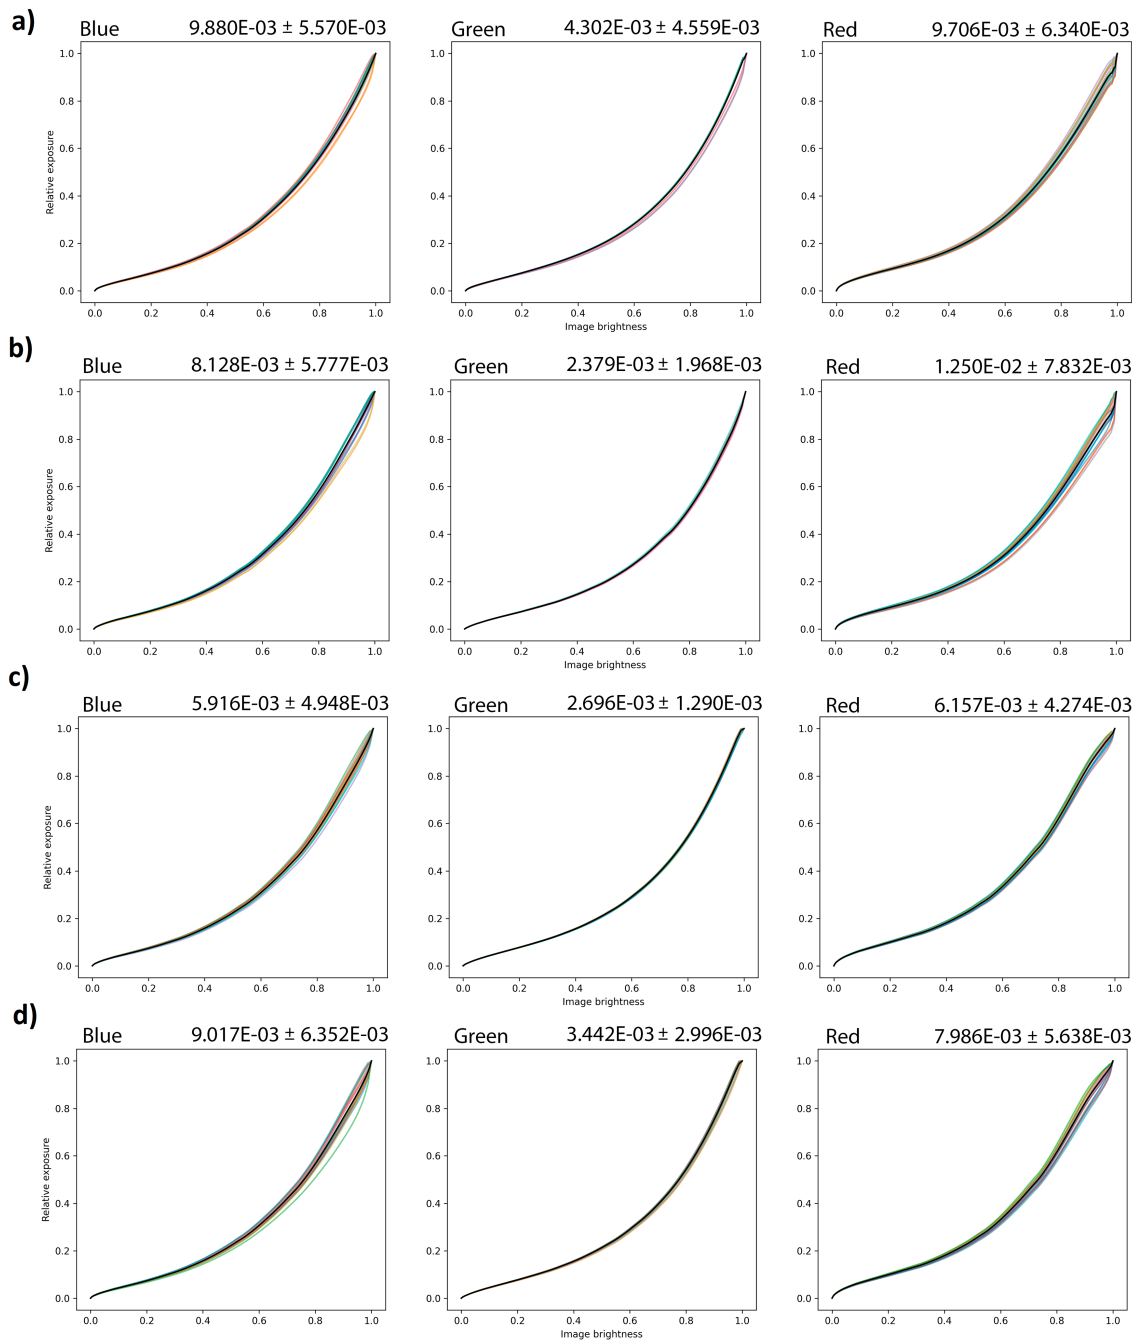


Figure 33. The ICRFs of each color channel of Olympus SC100 camera solved with the same parameters for 20 repeats. The colored curves represent an individual solution and the opaque black curve represents the mean ICRF of the 20 solutions. Also shown is the mean RMSE and its associated standard deviation, computed by comparing each individual ICRF to the mean ICRF. a) Gamma curves not included, color split used. b) Gamma curves included, color split used. c) Gamma curves not included, no color split. d) Gamma curves included, no color split. A fraction of about $1.23E-5$ of all pixel positions were sampled with a rejection threshold of 5.

From figure 33 we can observe that repeatedly solving the ICRFs with same parameters yields results that are quite close to each other. The largest mean RMSE value occurs on the red channel for the principal component set with gamma curves included and color split enabled. The RMSE values' units can be converted to pixel values by multiplying with 255, which yields for the largest RMSE value about 3.19, while the smallest is about 0.61. The RMSE values indicate that for most cases the variability is small and most solutions are clustered in a small region of the solution space. All in all this means that despite the stochastic nature of the underlying differential evolution algorithm, the process is quite consistent at finding a good solution. In theory, if there was no noise and if the model was a perfect fit to the data, the solved ICRF would be the same on each attempt. However due to the stochastic nature of the algorithm, noise and imperfections between the model and reality, there are multiple solutions that yield about the same measure of linearity. With these in mind, it is not surprising that the red channel has the most variability as it is also the noisiest channel.

As for which set of principal components result in the most accurate ICRF, we need to observe how the linearity measure behaves for different sets of components. For each set of principal components, we determine the mean and the standard deviations for each channel's mean absolute and relative disparities and their associated standard deviations. The values are reported in table 5. For recap, the mean absolute disparity measures for each channel the average difference from expected linear values for a set of images based on their exposure times. This value has an associated standard deviation that measures how spread out the differences from expected values are. Solving multiple ICRFs with the same parameters, computing these measures for each channel and finally determining their mean and standard deviation results in very helpful quantities. These values help us choose the optimal set of components out of the different sets. Additionally these help us further to observe how much the variability of the ICRFs output by the optimization algorithm affects the linearity measures.

In table 5 the bold values indicate the best value for each channel and each set of principal components. The best value is simply the one that is the closest to zero, as it corresponds to a set of components that most consistently results in the ICRF that yields the highest measure of linearity. The most valuable quantity is the mean of the mean relative disparity, followed by the standard deviation of the mean relative

Measurand	PC set	Absolute disparity		Relative disparity	
		Mean	SD	Mean	SD
Blue mean	a	-3.370E-02	7.703E-02	1.580E-02	1.910E-03
	b	-4.665E-02	6.133E-02	1.385E-02	1.500E-03
	c	-6.236E-02	5.443E-02	1.455E-02	1.300E-03
	d	-3.204E-02	8.736E-02	1.548E-02	1.930E-03
Blue SD	a	5.351E+00	1.522E-01	1.623E-01	4.702E-04
	b	5.506E+00	1.872E-01	1.624E-01	5.026E-04
	c	5.424E+00	1.406E-01	1.626E-01	5.096E-04
	d	5.381E+00	2.013E-01	1.625E-01	5.108E-04
Green mean	a	3.171E-01	2.488E-02	1.885E-02	1.530E-03
	b	2.768E-01	2.128E-02	1.475E-02	6.387E-04
	c	3.092E-01	1.529E-02	1.555E-02	6.710E-04
	d	3.089E-01	1.609E-02	1.557E-02	7.878E-04
Green SD	a	2.997E+00	6.875E-02	1.298E-01	7.864E-04
	b	2.876E+00	3.021E-02	1.285E-01	5.130E-04
	c	3.074E+00	3.030E-02	1.281E-01	2.942E-04
	d	3.058E+00	4.841E-02	1.280E-01	1.842E-04
Red mean	a	-2.412E-01	1.096E-01	2.445E-02	2.090E-03
	b	-2.170E-01	9.541E-02	2.445E-02	1.880E-03
	c	-2.351E-01	6.946E-02	2.418E-02	1.560E-03
	d	-2.383E-01	6.490E-02	2.426E-02	1.360E-03
Red SD	a	5.520E+00	2.127E-01	1.839E-01	3.663E-04
	b	5.497E+00	2.510E-01	1.840E-01	5.620E-04
	c	5.942E+00	1.303E-01	1.841E-01	2.132E-04
	d	5.891E+00	1.878E-01	1.840E-01	2.085E-04

Table 5. Statistics of the computed linearity measures for the 20 ICRFs solved for each principal component set. The principal component (PC) set letters correspond to the subfigures a-d in figure 33. The bold values are the values that are the closest to zero for each measurand.

disparity.

Out of all the sets of PCs set b has the most bold values. This also coincides with the priorities of these quantities with both the blue and green channels' means of mean relative disparity belonging to set b. Further analyzing the resulting means of mean relative disparity via Welch's t-test for statistically significant differences we find the following. Firstly on the blue channel half of the resulting pairs have statistically significant differences, and one of the insignificant pairs just barely fails the test with its value 0.06344 being only slightly over 0.05. Secondly, on the green channel only one pair has statistically insignificant differences between their

Ch	a-b	a-c	a-d	b-c	b-d	c-d
B	9.632E-04	1.903E-02	5.861E-01	1.177E-01	3.390E-03	6.344E-02
G	3.357E-11	2.645E-09	2.488E-09	3.238E-04	5.569E-04	9.281E-01
R	1.000E+00	6.431E-01	7.314E-01	6.197E-01	7.109E-01	8.573E-01

Table 6. Results of Welch’s t-test for each channel under the null hypothesis that $\mu_1 - \mu_2 = 0$. The values are computed using the mean of mean relative disparity for each pair of principal component sets. Bold values mark larger than 0.05 probabilities, indicating statistically insignificant difference.

values. The red channel is very monotone and there are no statistically significant differences on it between the different PC sets. All in all this means that on the red channel the different sets of PCs result in very similar ICRFs, while on the blue and green channels the choice has significant impact on the linearity measure. Therefore judging by the PCs effects on the mean relative disparity, set b is the optimal choice.

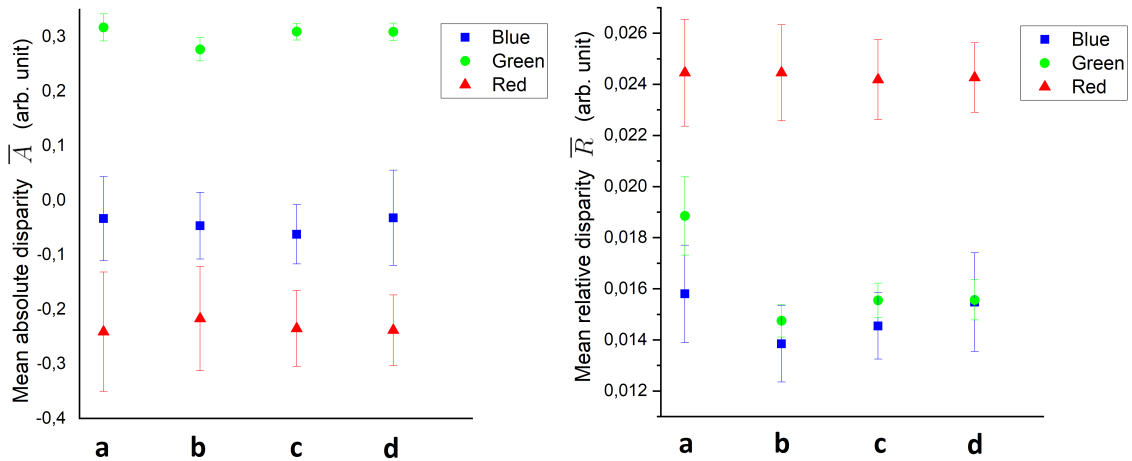


Figure 34. Plots of the mean of mean relative disparity for each set of PCs a-d and for each color channel, based on the data reported in table 5. The bars represent the standard deviation of the mean relative disparity values. From the plot it is relatively easy to see which pairs of PC sets result in statistically significant differences.

One doesn’t necessarily have to use the same set of PCs for all channels. Instead different sets could be chosen for different channels. In the case of our Olympus SC100 camera, we can use set b for all channels, as the red channel shows no significant variations for different PCs. Likely one of the most important causes of this is the high noise on the channel. Most likely the choice of optimal PCs is very device dependent.

1.4 Comparison to other data sets

One way to verify the performance of these linearization algorithms is to run them on an existing data set and compare the resulting ICRF to published results. For this we require a data set in which both the original image data and a determined ICRF are both publicly available. We found one quite robust data set related to an article by Engel, Usenko and Cremers, in which both an ICRF and the original image data are available to download. The camera model they utilize in their research is the uEye UI-3241LE-M-GL monochrome, global shutter CMOS camera by IDS.[73]

From their data sets we chose three sets in which a static scene was imaged at various exposure times. For each exposure time about ten frames were captured, so we could use them to determine a mean frame for the exposure-based calibration. As the camera name suggests, the camera outputs 8-bit gray scale images and the authors report that they utilize jpeg compression. The lossy jpeg compression is suboptimal in terms of maximizing the accuracy of results.

We run the optimization algorithm for each of the three chosen data sets with a small modification. We'll utilize the principal components with no color split and duplicate the monochrome image for each color channel as it is. This way we can essentially simultaneously run three attempts at solving the ICRF and then take the mean of each solution. The resulting ICRFs, solved from the three different sets of images, are shown in figure 35 alongside with the ICRF provided by Engel et al. The RMSE between our ICRFs and the provided ICRF shows that they match very closely, indicating that our algorithm provides an accurate ICRF.

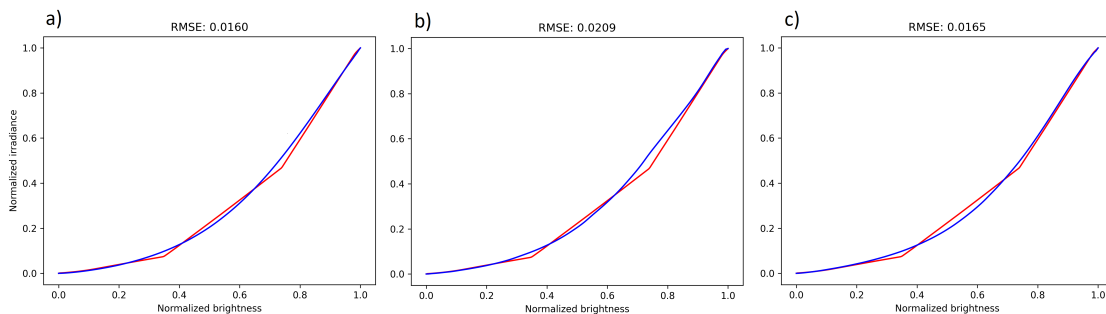


Figure 35. Plots of the ICRFs solved by our algorithm (red) based on the different data sets published by Engel et al. and the ICRF provided by the data sets' authors (blue). a) Narrow gamma sweep 3, b) Wide gamma sweep 1 and c) Wide gamma sweep 2. Above the plot are the computed RMSE values between our ICRFs and the one provided in the data set.

2 Olympus SC100 calibration

The photographs used in this research were all taken using a Olympus BX53M optical microscope paired with a Olympus SC100 CMOS color camera, whose specifications are provided in table 7. The image acquisition software was Olympus Cell Sens Entry on a Windows 7 PC. All subject images and dark frames were acquired by capturing a set of 67 images, which were used to create and store a mean image and a scaled variance image, as was described in part II, sections 2.3 and 2.4 respectively. The image pairs act in calculations as the mean and standard error of a set of measurements.

In the camera's settings all unnecessary algorithms and adjustments were turned off. This includes internal automatic white balance, gamma adjustment, sharpening filtering and shading correction. Internal dark frame subtraction is used for subject images, but for their uncertainty images the dark frame subtraction is turned off. The camera doesn't have a single ISO gain value but instead employs a color channel specific gain value, which we'll utilize to fix the white balance of the camera.

Chip size	Type 1/2, 3 Inch
Effective area	6.12 mm x 4.59 mm
Pixel size	1.67 x 1.67 μm
Max resolution	3480 x 2748 pixels
Image type	24-bit RGB
Read-out speed	5-36 MHz
Interface	USB 2.0
Camera mount	Standard C-mount

Table 7. Specifications of the Olympus SC100 CMOS color camera.

2.1 Distance scale

The microscope images' scale was calibrated for the 50x magnification and then calculated for the lesser magnifications. A silicon chip with electrodes fabricated on it was utilized as a sample and the distance between electrodes was determined from an AFM image. The electrodes were fabricated at $7 \mu\text{m}$ intervals and the mean distance and standard error of the mean were determined. Similarly for optical images of the same chip, the mean distance and standard error of the mean were determined in pixels for the electrodes' distances to each other. The distance scale is then established by dividing the mean distance in μm by the mean distance in pixels, with the standard error being determined using the error propagation law. Figure 36 shows the utilized AFM image and optical image, with lines showing the locations at which the profiles of the AFM image were obtained from. Using the line profiles obtained from both the AFM and optical images, the middle points of each peak were utilized as the location of the electrodes. The middle point was determined as the mean value between the two edges of a peak. This process establishes a distance scale for the 50x objective of the microscope and this value can then be simply scaled for the other magnifications. Table 8 shows the distance scales along with their standard error. The distance scale value for the 50x objective matches closely to the theoretical value that can be computed as physical pixel size divided by the total magnification. The camera mount applies a reported 0.35x magnification, so the theoretical value is $0.0945 \mu\text{m}/\text{px}$. The theoretical value doesn't quite match the experimentally determined value, falling some 3.4 standard errors outside of it. Still, the values are reasonably same and we will utilize the experimentally determined distance calibration in this work.

Magnification	Distance per pixel ($\mu\text{m}/\text{px}$)	Standard error ($\mu\text{m}/\text{px}$)
50x	0.0923	0.0009
20x	0.2309	0.0022
10x	0.4617	0.0044
5x	0.9235	0.0088

Table 8. The distance scale and its standard error for the different objectives for Olympus BX53M microscope and Olympus SC100 camera pairing.




Figure 36. AFM and optical images of an silicon chip with electrodes fabricated at about $7 \mu\text{m}$ separations of each other. The lines in the mark the line profiles that were utilized to determine the mean distances of the electrodes. The red box in the optical image marks approximately the area of the AFM image.




Figure 37. Profile of the pixel values of line 3 in the AFM image shown in figure 36. The location of the peaks corresponds to the electrodes on the chip. The peak location is determined as the mean value of the edges of the peaks. Similar methodology was used for all electrode distances for both the AFM and optical images.

2.2 White balance

The white balance of the camera was adjusted by viewing the transmitted light source of the microscope without any obstructions in the way through the 5x objective. The transmitted light source is of the same type as the reflected light source. Using a mean stacked image, a spatial mean about the brightest point of the image was computed for each of color channel and the camera's white balance set by adjusting the individual channels' gains until the means of each channel was as close to each other as possible. The closest matching spatial means were: $\bar{B}_r = 247.713$, $\bar{B}_g = 247.996$ and $\bar{B}_b = 247.738$ found by adjusting the camera's channel gains to red: 1.95, green: 1 and blue: 1.01. It is highly likely that the linearization for most sets of ICRFs leads to a shift in the white balance of images. Figure 38 shows the image that was utilized to set the white balance before and after linearization, indeed showing a changed white balance. In the linearized exposure domain the values changed to $\bar{X}_r = 197.851$, $\bar{X}_g = 238.0811$ and $\bar{X}_b = 243.347$. This means that the color ratios changed from about 0.333 to 0.291, 0.334 to 0.350 and 0.333 to 0.358 for red, green and blue channels respectively. The shift is quite modest, but certainly visibly noticeable. While it might be tempting, these ratios before and after linearization can't be used to correct the white balance. Instead some method of tying the ICRFs of each channel would be the correct approach to handle the white balance post-linearization. This is an aspect that we don't include in the present work. While it would be beneficial to be able to accurately quantify the relative radiances of the channels with respect to each other, with the methods we have developed here, we can only quantify the relative radiance values on individual color channels. The color channels relation to each other is left on a qualitative level of being able to say that one is brighter than the other, but no more than that.

2.3 Imaging system noise

A simple method of visualizing the noise of the imaging system is to plot a heat map of a pixels true value against its observed distribution. Basically for each pixel in a stack of images, we assign its mean value as an estimate of its true value. Then we collect a histogram of the pixel values that make up our estimate and plot the histogram. Repeating this process for all pixel values yields a plot along whose diagonal is the true value of a pixel, while any straying off of the diagonal is due to




Figure 38. On the left-hand side is the image that was utilized to set the white balance for the Olympus SC100 camera. On the right-hand side is the image after linearization, showing a shift in the color balance.

noise. Noise profiles with CellSens dark frame subtraction off and on are shown in figures 40 and 41 respectively.




Figure 39. Plots of the standard deviation of each mean-based dark subtracted noise profile. The red channel is considerably more noisy than the blue and green channels, especially at lower end of the range.

Finally we can utilize the subtracted noise data to determine a lower and an upper threshold for the pixel values that are rejected in the linearity measurements for being too close to the data range edges. In figure 39 we have determined and plotted the standard deviation of each noise profile. The edgemoost standard deviation values are once again something we need to be wary of, as they are smaller than they should be in reality due to sensor cut-off and saturation. These rejection thresholds will be utilized in all linearity measurements.

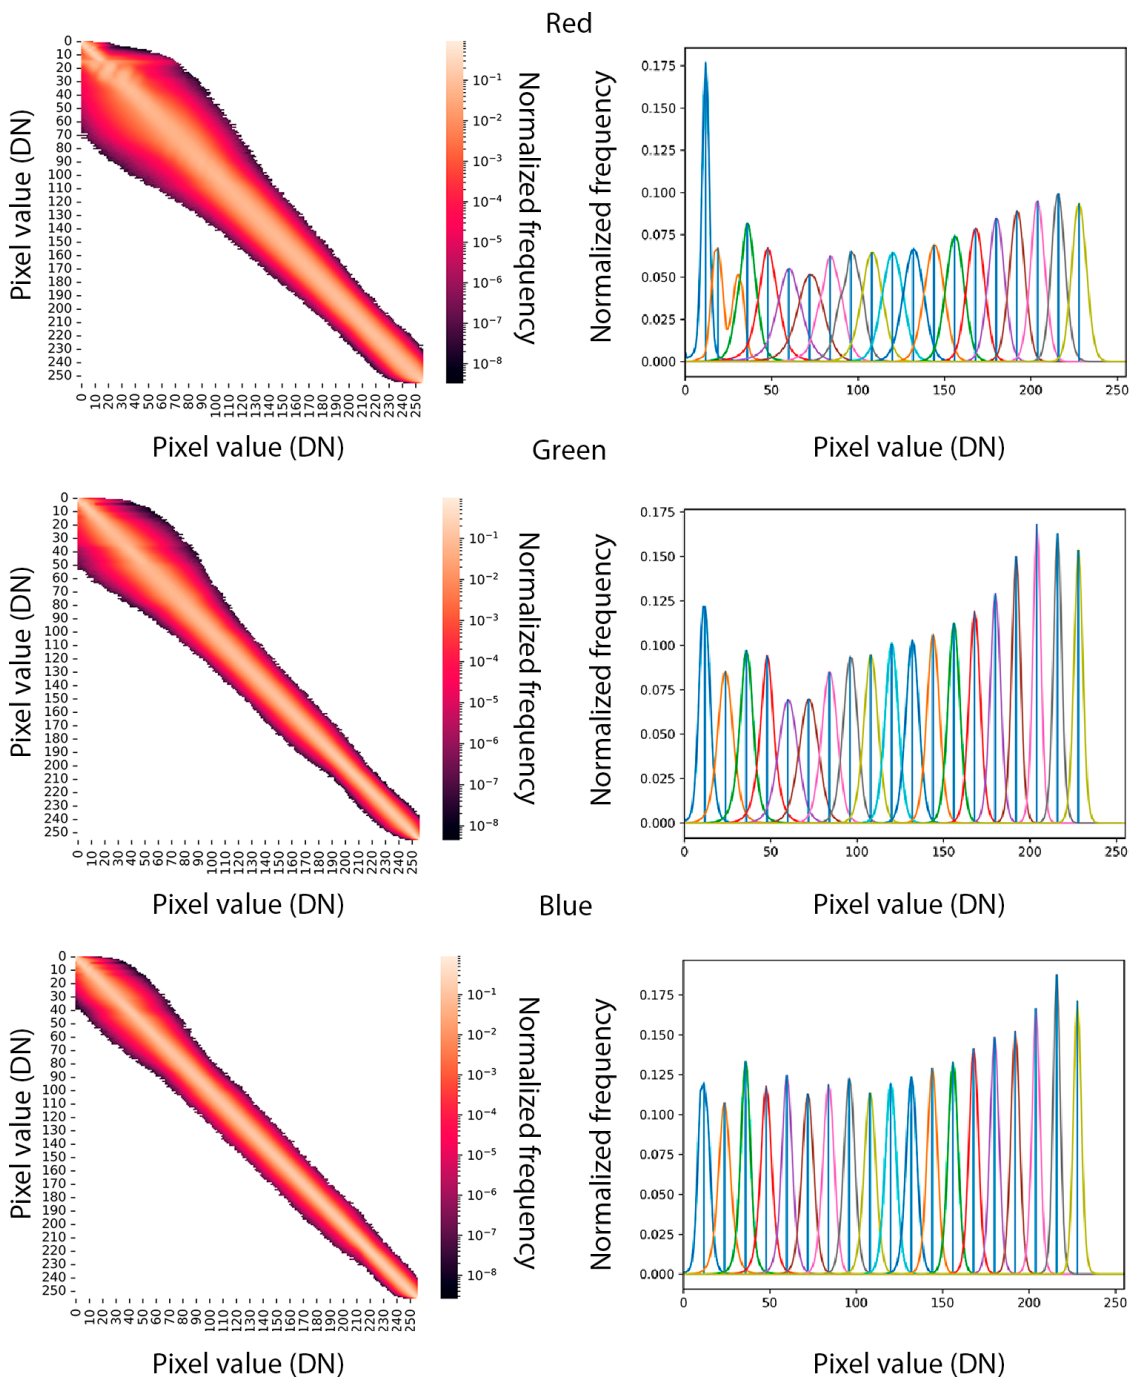


Figure 40. Heat maps and plots of select noise profiles for the red, green and blue color channels based on setting the mean of each pixel stack as its true value. These profiles were produced with dark frame subtraction turned off. The profiles are measured in digital numbers (DN), that is in the 8-bit integers directly observed in the images.

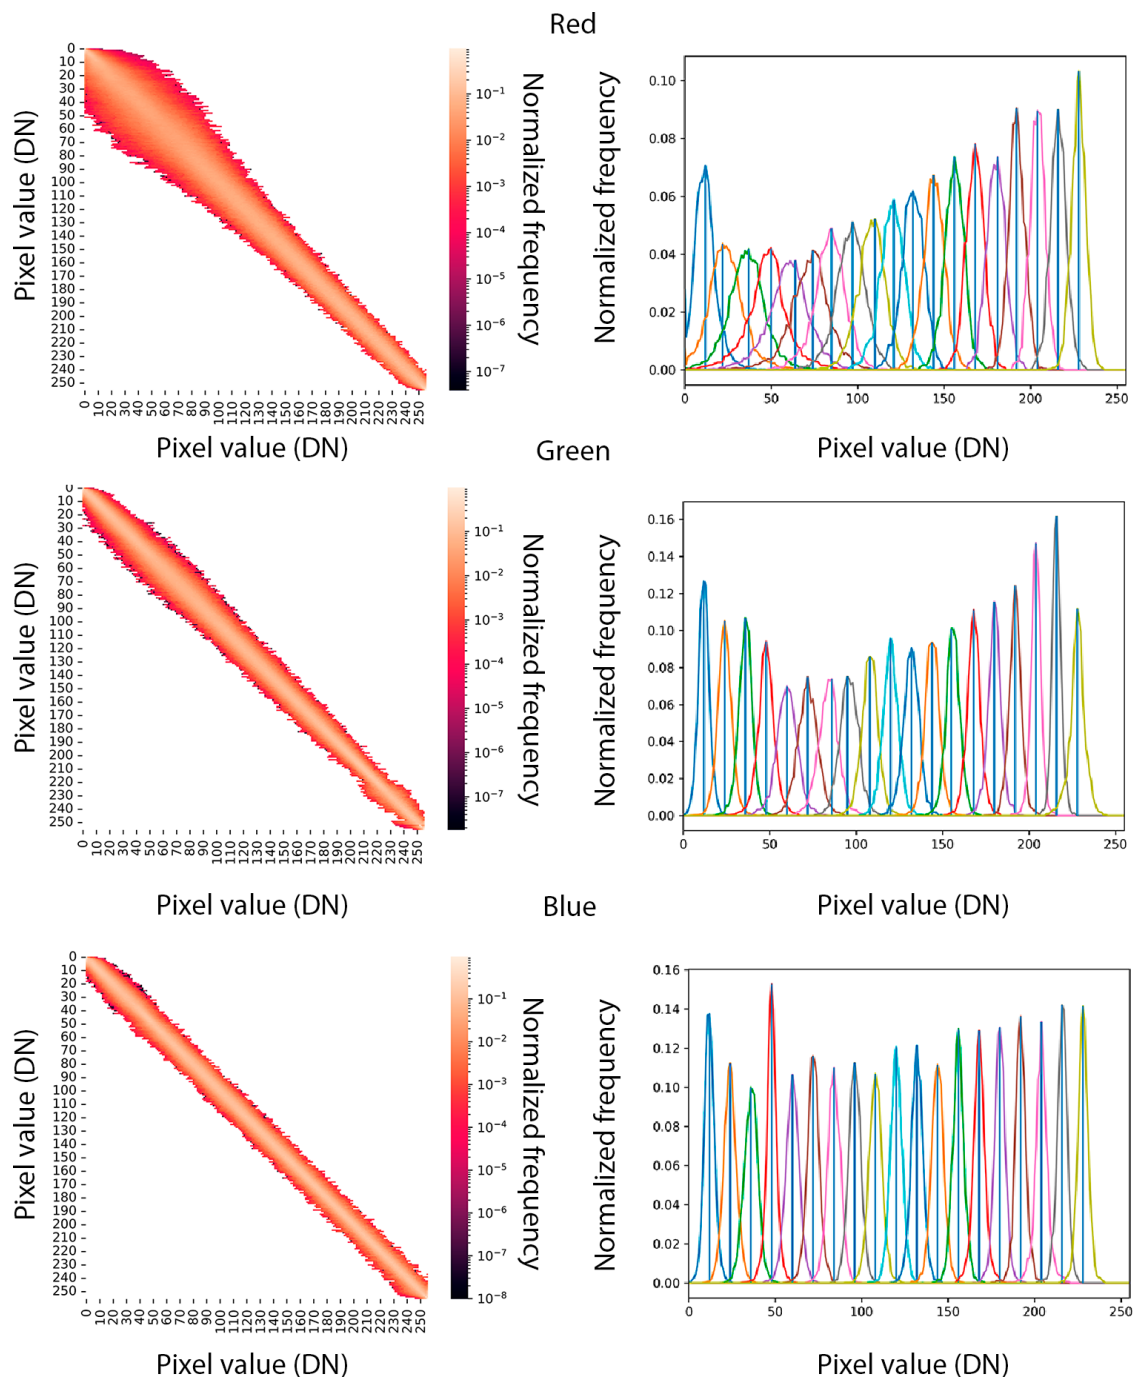


Figure 41. Heat maps and plots of select noise profiles for the red, green and blue color channels based on setting the mean of each pixel stack as its true value. These profiles were produced with dark frame subtraction turned on. The profiles are measured in digital numbers (DN), that is in the 8-bit integers directly observed in the images.

2.4 Linearization and HDR merging

The ICRF for the Olympus SC100 camera was solved by utilizing the differential evolution algorithm with the parameters stated in part III, section 1.1. For the objective function the rejection threshold was set to 5 and the images were sampled at a fraction of about $4.46\text{E-}3$ of all pixels. The principal components that were utilized is the set b , which as described in part III, section 1.3, includes the gamma curves and utilizes a color split for the CRFs from Grossberg and Nayar's database. The exposure series used to solve the ICRF is the same as shown earlier in figure 25, consisting of a total of 13 images captured at exposure times between 8.01 ms and 79.96 ms. The solved ICRF is shown in figure 42. Figure 43 shows the disparity plots for the camera following linearization of the same images that were used to compute the linearity plots in figure 24. When comparing these to the disparity plots of the base images it becomes very noticeable how nonlinear the device is on its own. We can finally verify the performance of the HDR merging process by comparing the pixel values of given sites within and across two HDR images of the same scene, merged from images captured under different known illumination intensities. Figure 45 shows two HDR images, who are merged from exposure series with the only difference being the illumination intensity. In the first image the illumination is half that of the second, ensured by a polarizer placed in the path of unpolarized light. We'll divide the two HDR images into equally large sites and compute a mean radiance value for the site with an associated uncertainty. In this case, we shall divide them into 900 sites, basically a 30×30 grid.




Figure 42. The solved ICRFs for each color channel of the Olympus SC100 camera. The red and green channels behave in a quite similar fashion being quite nonlinear, while the blue color channel is the most linear.

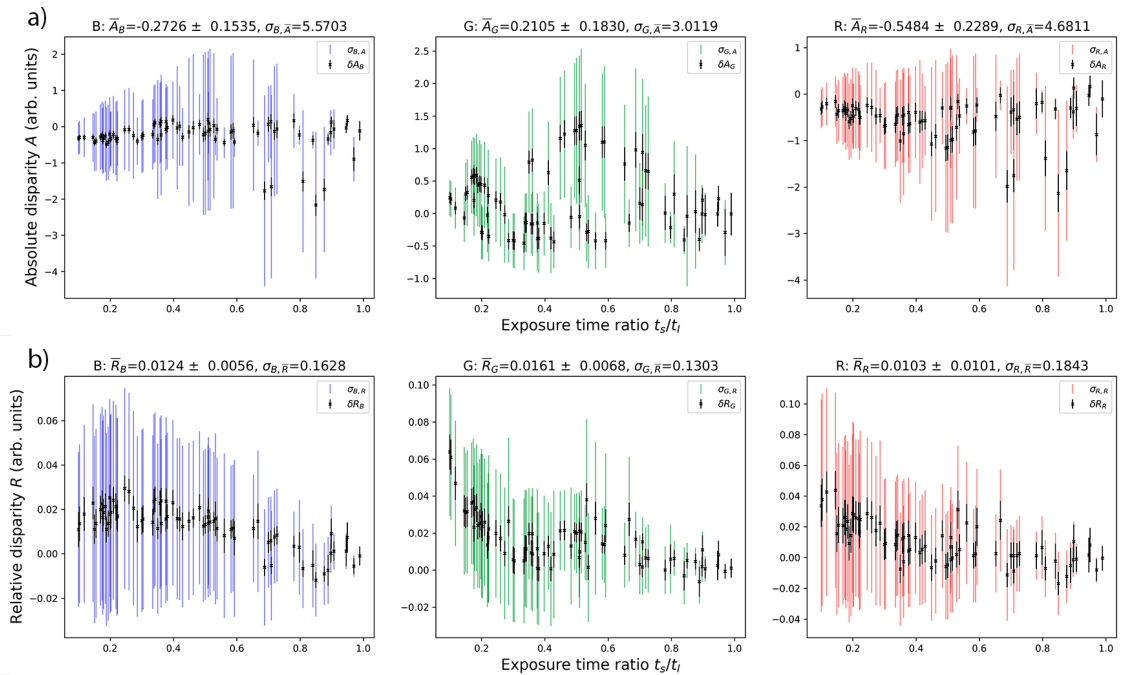


Figure 43. Plots of the absolute and relative disparities for the Olympus SC100 camera after solving for its ICRF. The points represent mean disparity for a pair of images and the colored bars represent the standard deviation of disparity and are divided by five for better visualization. The black bars represent the uncertainty of disparity. Above the plots are computed the mean and standard deviation for the data points in each plot.

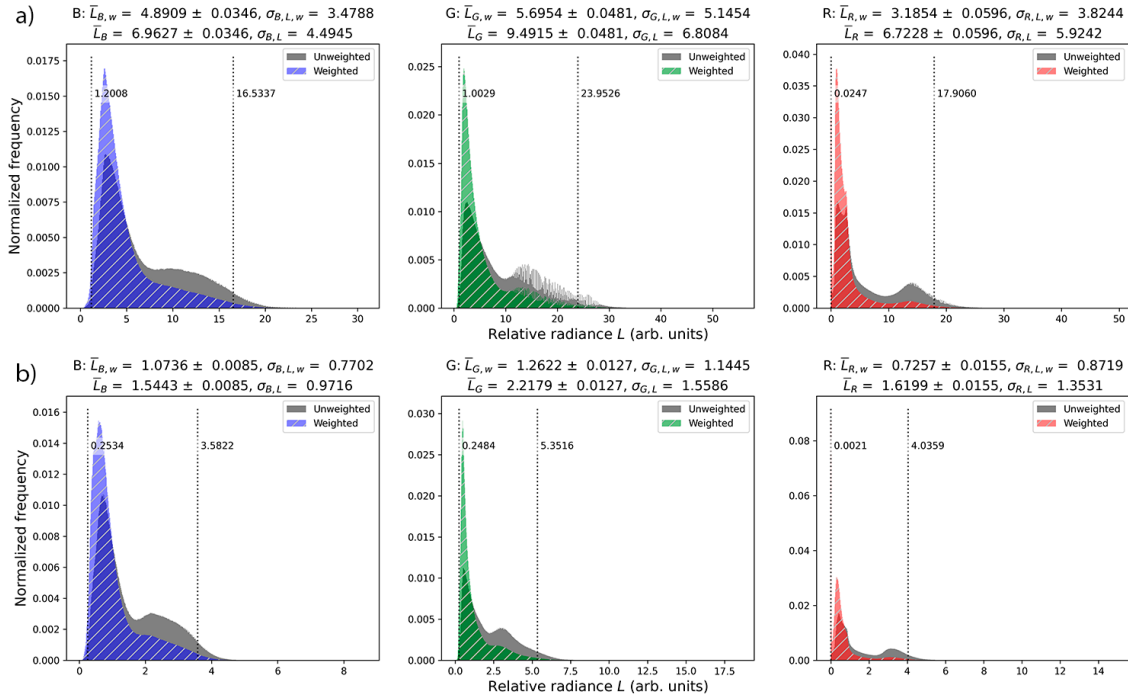


Figure 44. Unweighted and weighted histograms of each channel's relative radiance values in the HDR images, whose LDR scaled images are illustrated in figure 45. The uncertainties of the relative radiance values were utilized as weights. a) and b) correspond to the images in which an analyzer was utilized and not utilized respectively. Above each plot are reported the weighted and unweighted mean relative radiances with their uncertainty and the standard deviation of the relative radiance values. The vertical lines indicate the relative radiance value range that was linearly mapped to LDR range in figure 45c) and d).




Figure 45. Images of a piece of paper with color stripes printed on it, 5x DF objective, left-hand side with an analyzer and right-hand side without an analyzer. a) and b) are normal images captured at 39.99 ms exposure time. A total of seven and six exposures were utilized for each case. c) and d) are HDR images, scaled to LDR by linearly mapping the values between the 1st and the 99th percentiles of the HDR images' weighted histograms. e) and f) map the full range of the HDR images back to LDR. The HDR-merged images a-f illustrate well the effect of flat-field correction and amount of detail that would be lost in LDR images due to cut-off and saturation. c) and d) show only minute differences in the highlights. The brightness in f) is slightly more even across the image than in e) due to a wider range of relative radiance values being scaled back to LDR, resulting in a lower contrast.




Figure 46. Unweighted and weighted histograms and KDEs of the cross-image ratios of mean relative radiance for each of the 900 sites and for each color channel. The weighted histograms and KDEs utilize the uncertainty of the mean relative radiance values as weights.

Firstly we'll look at how the same site's mean radiance values compare to each other across the two HDR images. We expect the ratio of the mean radiance values of each site to be consistent for each color channel. Basically we expect that the distribution of the ratios of each site's mean radiance across the two images is closely clustered about a mean value. In figure 46 are plotted a histogram and a kernel density estimate (KDE) for the ratio of the mean radiances for the 900 sites across the two images.

From figure 46 we can observe that the ratios of relative radiance values in the two images are quite closely clustered, indicating two things. Firstly it indicates that the images indeed in a relative scale record the same information of the scene. Secondly it indicates that the images' relative radiances are related to each other by a scalar factor. Intuitively one might have expected the ratio to be about 0.5, but as stated in part II, section 7, each channel's radiance values for each image are related to absolute radiance by an unknown scalar factor – and thus by extension to each other. For our research this result is already enough, but if one desires, it is possible to determine a scalar factor that can be used to convert the relative radiance values to absolute physical units. The relative radiance values within an image can be compared directly, but comparisons across different images requires normalizing the relative radiance values by the maximum value out of all channels. This ensures that white balance is maintained.

Secondly we'll make a comparison in two steps. We'll first compare the mean relative radiance between two sites in one image by computing the ratio of the mean relative radiances. Then we compute the ratio for the same pair of sites in the




Figure 47. Unweighted and weighted histograms and KDEs of the relative difference between the ratios of mean relative radiance for each unique combination of zones in an image. The weighted histograms and KDEs utilize the uncertainty of the mean relative radiance values as weights.

Method	Ch	Cross-image zone ratios			Pair zone ratios		
		Mean	Uncertainty	SD	Mean	Uncertainty	SD
UW	B	0.223	0.002	0.010	0.0145	0.0002	0.0631
	G	0.238	0.002	0.011	0.0115	0.0002	0.0645
	R	0.248	0.004	0.029	0.0123	0.0003	0.1574
W	B	0.222	0.002	0.010	0.0189	0.0002	0.0647
	G	0.236	0.002	0.010	0.0160	0.0002	0.0653
	R	0.243	0.004	0.024	0.0399	0.0003	0.1742
KDE	B	0.224	0.002	0.010	0.0189	0.0002	0.0675
	G	0.239	0.002	0.011	0.0160	0.0002	0.0682
	R	0.254	0.004	0.025	0.0399	0.0009	0.1818

Table 9. Mean, standard deviation of mean and standard deviation for the cross-image zone ratios and pair zone ratios computed for each channel with different methods. UW = unweighted, W = weighted and KDE = kernel density estimate.

other image. Finally we can compare the ratios obtained from the two images in a similar fashion to equation 40, though this time there is no need for scaling. This can then be done for all possible unique combinations of sites in the images. Figure 47 shows the histograms and KDEs for each channel computed from the 404550 combinations of sites. This result acts as a proof and quantification that the two differently illuminated HDR images encode the same information in relative terms at a given level of accuracy and that comparing relative radiances is valid. One just has to keep in mind that if the illumination intensity is not the same, then the relative radiances are related to each other by an unknown scalar factor.

2.5 Dark frames and bad pixel filtering

In our image pipeline we opt to utilize the dark frame subtraction feature of the CellSens Entry camera software. Manually captured dark frames are used only to create the bad pixel maps to filter the pixel values at the bad positions. While the bias subtraction step could be done manually, this would leave us with a smaller data range for the 8-bit channels. We'd be subtracting the noise floor from all pixel values and are left with that much smaller usable data range. Dead pixel filtering could be done in a similar fashion by acquiring flat very bright, but not saturated, images and then mapping out pixel positions, whose value is too low compared to the image's mean brightness. For our camera we found that dead pixel mapping is not required.

The nature of the camera software's dark subtraction is a black box for us, but we can make an educated guess at how it works. Firstly upon choosing dark frame calibration in the software, the user is prompted to block the camera from all light. The user can then start the process, which takes a few minutes. Then the user can output the dark frame in the software. This gives out an gray scale 8-bit image in which hot pixels are still present. The two factors suggest that the camera might capture a mean stacked image that is not demosaiced and is instead subtracted from an acquired image before demosaicing. Likely the dark frame is captured only for the shortest exposure time, meaning that the process corrects for the bias current. This is further suggested by the fact that hot pixels are still present in images with long exposure times, while the bias current is not.

Figure 48 shows the ratio of pixel positions, whose value is greater than the given threshold as a function of exposure time for dark frames. The dark frames in a) have been acquired without internal dark frame subtraction and instead a bias frame has been manually subtracted from them before counting the hot pixels. In b) the dark frames have been acquired with the internal dark frame subtraction enabled. Figure 48b) clearly shows that the internal dark frame subtraction is excellent at removing the bias and reasonably well removes lot of the hot pixels, as is evident by comparing the ratio of hot pixel positions between a) and b). Though at increasing exposures more and more hot pixels pass through the software's internal dark frame subtraction. This is illustrated in figure 49, which shows close ups of two points in a 5x DF HDR image of sample 3, scaled to LDR by linearly mapping relative radiance range [0,

0.025] to [0, 255]. The first row corresponds to images with no bad pixel filtering and the second row corresponds to filtering based on the dark frames subject to internal subtraction. Finally the third row corresponds to filtering based on the manually acquired dark frames. This clearly shows that filtering for hot pixels is required to increase the accuracy of relative radiance values at dark locations and that the best way to map these positions is to utilize dark frames that aren't subject to the software's internal subtraction. The bad pixel filter causes some minor artifacts in the image, that are best illustrated in the third column of images in figure 49. In the filtered images on the second and third row some slightly darker pixels have appeared on bright particles, these dark pixels are not present in the image on the first row. This happens because a small feature that appears moderately bright in an long-exposed image happens to lie on a known hot pixel location. That pixel location's value is then replaced by the median of its neighborhood, which due to the small size of the feature has a lower value than the pixels constituting the feature. For larger features this doesn't happen because the median filter fits sufficiently well within the feature.

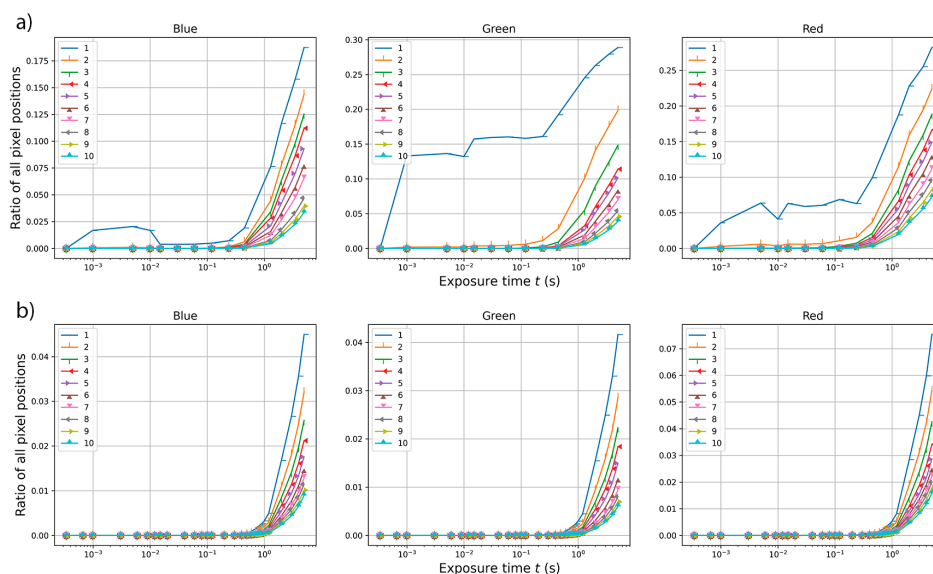


Figure 48. For each channel's dark frames the ratio of pixel positions with value equal to or greater than the given threshold value as a function of exposure time. a) is based on manually handled dark frames, while b) is based on images that were subject on internal subtraction.




Figure 49. Close ups of a 5x DF HDR image of sample 3 scaled to LDR by linearly mapping relative radiance range $[0, 0.025]$ to $[0, 255]$. In the first row of images no bad pixel filtering has been employed. In the second row bad pixel filtering has been done based on manually acquired dark frames, that have been subject to CellSens software's internal dark frame subtraction. Finally the third row shows images, whose bad pixel filtering is based on manually acquired dark frames without the software's internal dark frame subtraction. At the centers of the green circles are indicated some of the darkening artifacts caused by the filtering process in bright parts of the image.

2.6 Flat-fields

The flat-field images used to correct the vignetting and fixed pattern artifacts were produced by the method we discussed in part II, section 5.2. For each objective in both BF and DF modes, sets of 67 to 200 images were utilized in linearizing and mean stacking images of a piece of paper kept in constant motion during the acquisition. This method yielded rather good flat-field images, some of which are shown in figure 50 along with a brightness plot along the top-left to bottom-right diagonal.

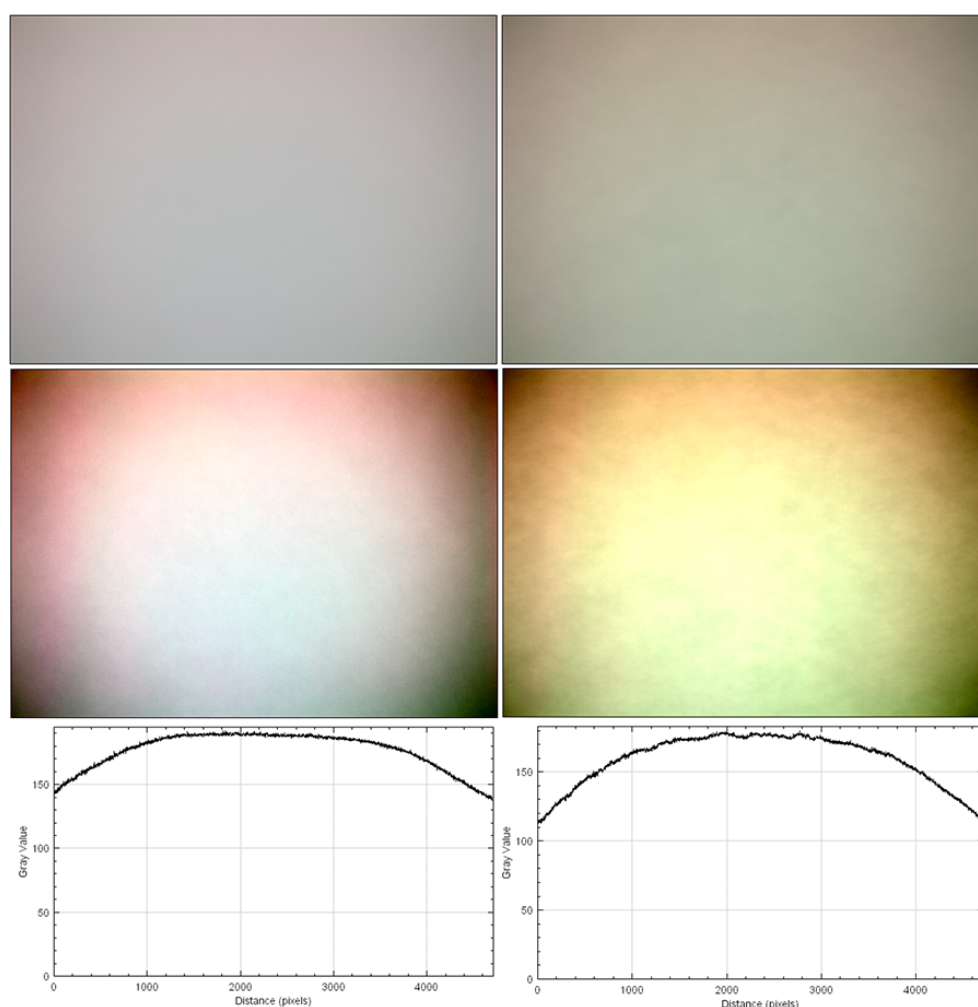


Figure 50. On the left-hand side is the constructed and linearized flat-field image of the 20x BF objective and on the right-hand side similar image for the 20x DF objective. The second row shows the same images with contrast adjusted to maximum to better illustrate the non-uniformities. Bottom row shows plots of image brightness from the contrast-adjusted images along the top-left to bottom-right diagonal.

From the flat-field images we can observe that there is indeed considerable vignetting happening. Furthermore the vignetting is not symmetric nor is it the same for each different color channels at the same locations. The phenomenon is stronger when using DF mode, though strictly speaking it isn't only vignetting. Instead in DF mode the light field isn't as uniform as in BF mode, which is the main source of the darker fringes in the image. Similar features are observed for all the other objectives.

For the flat-field images the quantification of a standard error is a very different problem than for the other images. This is due to the flat-field image being constructed from multiple images of different subjects, rather than from multiple images of the same subject. The image is essentially a simulated one. This means the normal way of calculating a standard error of the mean for the stack doesn't provide any relevant information. We approach this problem by utilizing the collected noise data of the imaging system. This way we can virtually assign for each pixel value a standard error based on the noise data as if the pixel value was something that was measured directly. If we had a normal mean-stacked flat-field image to compare this simulated one to, we'd find that the standard error across the image would be about the same overall. As we assign to each pixel an uncertainty from aggregate data, we have no specific information for each pixel, which means that some pixels get assigned a greater uncertainty than they have in reality, while others get assigned a lesser one. The overall behavior of the standard error evens out over a large number of pixels.




Figure 51. 50x DF HDR images of sample 3's dim zone, scaled back to LDR by linearly mapping relative radiance values $[0, 20]$ to $[0, 255]$. On the left-hand side no flat-field correction has been applied, while on the right-hand side it has been utilized. Note especially the difference in the corners of the images.

3 Measurements

With the image correction and linearization methods developed and analyzed, we can now get into the image-based measurements. As a first step we'll put the linearization methods to full use to construct HDR images to analyze the radiance of the different features of the samples on each color channel. Following this we'll get into the local thickness and autocorrelation measurements.

3.1 Relative radiance measurements

We can utilize the relative radiance domain images to determine how bright the different features in the images are with respect to each other within and across color channels. This gives us some insight about what interactions of light and matter could give rise to the appearance of the MWNT depositions. We approach this by identifying interesting features in various HDR images constructed of the different samples. In the images we either use thresholding or manual selection of areas, or a combination of both, to compute spatial means and standard deviations for relative radiance values for a site. Both unweighted and weighted values are computed, with the corresponding pixel locations in the standard error image being the weights. We can accurately compare values that originate from images with the same magnification and illumination.

For the images captured utilizing the 50x objective we first manually select a rectangular zone from the image that clearly contains only the desired bright or dim zone. Then we refine the selection further by selecting only pixel locations, whose relative radiance is above a given threshold. A sensible threshold value would be one that ensures that no background of the clearly empty silicon chip is included. A common threshold value was determined for all channels by observing from multiple 50x magnification relative radiance images the spread of relative radiance values at locations of the sample material on each channel. A threshold of $L = 2.5$ was selected as it worked well at segmenting the sample material from the empty chip. For images utilizing the 5x and 10x objectives we only utilized a manual rectangular selection as

the spatial resolution at those magnifications is quite low for judging what is and isn't empty chip space within the bright and dim zones of the depositions.

Spatial means for relative radiance values were computed from multiple relative radiance images and multiple sites for both the bright and dim zones. These values are reported for each site in the tables in appendix B with both weighted and unweighted values. From these values further mean values were computed for both zone types at each magnification. Unweighted values are reported in table 10, while weighted ones are reported in table 11. In both tables are also computed the ratios of the lower magnifications relate against the 50x magnification and the ratios of the dim zone against the bright zone.

Firstly one can note that the weighted and unweighted zone-wise means of relative radiance values are quite closely matched to each other. The weighted values are consistently slightly lesser than the unweighted ones. This is expected behavior, as the inverse standard error weighing gives preference for smaller relative radiance values. The larger relative radiance values indicate either a very bright spot, a long exposure time or a combination of both. Regardless these spots tend to have a larger uncertainty associated with them, causing the weighted values to be slightly lesser than their unweighted counterparts. Similarly the uncertainty for each channel behaves differently, with the red channel's uncertainty being proportionally the largest compared to the associated value over the other channels.

The information of greatest importance in the zone-wise mean relative radiance tables is in the last section, that compares the values of the dim zone against the bright zone. As we know, the illumination intensity is not the same across different magnifications. This can be readily seen from equation 24. In the case of our objectives the illumination intensity is such that qualitatively speaking the 10x objective results in the brightest image, 5x objective in slightly darker and the 50x objective by far the darkest. This means that the relative radiance values in images utilizing different objectives can't be compared directly, as the relative radiances are related to each other by an unknown scalar factor. The important part is that we can instead compare how the relative radiance values of the bright zone and dim zone relate to each other across objectives. Here the results in both the unweighted and weighted cases indicate the same thing. The relative radiance of the dim zone is about 40% to 70% of that in the bright zone depending on the color channel in question and the objective utilized.

Bright zone					
Mag	Ch	Zone-wise mean	Uncertainty	Ratio to 50x	Uncertainty
50x	B	12.1824	0.0344	-	-
	G	10.2838	0.0343	-	-
	R	4.8427	0.0278	-	-
10x	B	10.5608	0.0278	0.8669	0.0045
	G	9.1814	0.0303	0.8928	0.0053
	R	4.0566	0.0242	0.8377	0.0099
5x	B	5.0987	0.0147	0.4185	0.0096
	G	4.0404	0.0137	0.3929	0.0121
	R	1.9531	0.0111	0.4033	0.0200
Dim zone					
50x	B	7.8041	0.0217	-	-
	G	7.0402	0.0226	-	-
	R	3.0484	0.0174	-	-
10x	B	5.7818	0.0155	0.7409	0.0052
	G	3.9854	0.0132	0.5661	0.0081
	R	1.6663	0.0094	0.5466	0.0146
5x	B	2.6176	0.0067	0.3354	0.0113
	G	1.8208	0.0061	0.2586	0.0180
	R	0.8081	0.0046	0.2651	0.0304
Ratio of dim to bright					
50x	B	0.6406	0.0025	-	-
	G	0.6846	0.0032	-	-
	R	0.6295	0.0051	-	-
10x	B	0.5475	0.0021	0.8546	0.0074
	G	0.4341	0.0020	0.6341	0.0099
	R	0.4108	0.0034	0.6525	0.0191
5x	B	0.5134	0.0020	0.8014	0.0327
	G	0.4506	0.0022	0.6583	0.0502
	R	0.4137	0.0033	0.6572	0.0821

Table 10. Unweighted mean values for the spatial mean relative radiance for each color channel at different magnifications, computed from the spatial mean relative radiance values across multiple sites. Also reported are the ratios of the zone-wise mean values against the 50x magnification (the 50x section is omitted as it is unity.) For each value the uncertainty is reported as the standard error. The bottom contains the ratio of the values of the dim zone against the bright zone.

Bright zone					
Mag	Ch	Zone-wise mean	Uncertainty	Ratio to 50x	Uncertainty
50x	B	11.1244	0.0344	-	-
	G	6.8134	0.0343	-	-
	R	2.8576	0.0278	-	-
10x	B	10.3749	0.0278	0.9326	0.0044
	G	7.6065	0.0303	1.1164	0.0057
	R	3.2383	0.0242	1.1332	0.0108
5x	B	5.0620	0.0147	0.4550	0.0093
	G	3.6119	0.0137	0.5301	0.0119
	R	1.6954	0.0111	0.5933	0.0198
Dim zone					
50x	B	7.2611	0.0217	-	-
	G	4.6138	0.0226	-	-
	R	1.5665	0.0174	-	-
10x	B	5.6340	0.0155	0.7759	0.0052
	G	3.3013	0.0132	0.7155	0.0088
	R	1.1885	0.0094	0.7587	0.0179
5x	B	2.5805	0.0067	0.3554	0.0112
	G	1.6795	0.0061	0.3640	0.0168
	R	0.6682	0.0046	0.4265	0.0306
Ratio of dim to bright					
50x	B	0.6527	0.0028	-	-
	G	0.6772	0.0048	-	-
	R	0.5482	0.0081	-	-
10x	B	0.5430	0.0021	0.8320	0.0068
	G	0.4340	0.0024	0.6409	0.0086
	R	0.3670	0.0040	0.6695	0.0171
5x	B	0.5098	0.0020	0.7810	0.0293
	G	0.4650	0.0025	0.6866	0.0353
	R	0.3941	0.0037	0.7189	0.0569

Table 11. Weighted mean values for the spatial mean relative radiance for each color channel at different magnifications, computed from the spatial mean relative radiance values across multiple sites. Also reported are the ratios of the zone-wise mean values against the 50x magnification (the 50x section is omitted as it is unity.) For each value the uncertainty is reported as the standard error. The bottom contains the ratio of the values of the dim zone against the bright zone.

For the moment let us focus on the unweighted results. The 50x objective result is the one that most likely measures the radiance ratios of the actual deposition material across the zones due to the use of thresholding to minimize the effect of the empty chip on the result. The 5x and 10x objectives on the other hand measure the radiance ratio of the whole system, including the effect of the empty chip. This statement is supported by the ratio of zone-wise means dropping a reasonably similar amount for the 50x objective for each channel, while for the 5x and 10x objectives the drop is more pronounced on the green and red channel. While the objectives likely have slightly different transmission for different wavelengths, the difference isn't large and doesn't come into play for the comparison of the zones' relative radiances as the comparisons are made objective-wise. The different transmissions however do affect the values where we compare the dim-to-bright ratios across the different objectives. The dim-to-bright ratio for the 50x objective cannot be accurately compared to those of the 5x and 10x objectives due to the thresholding, but the ratio for the 5x and 10x objectives should logically be the same. Comparing the ratios for each color channels does indeed show that they match closely to each other, though not within the uncertainty bounds. We can think of three reasonable explanations for the discrepancy. Firstly the mean relative radiance values for the 5x and 10x objective images originates from values across three sites but only one sample each. Secondly the uncertainty of the relative radiance values takes into account only the propagation of the uncertainty in the brightness of the original images from which the HDR images were merged. The linearization itself is assumed to be accurate. Thirdly, as stated previously, the objectives might have different transmission profiles.

In table 12 are reported the color ratios for the zone-wise mean relative radiances. These were computed by first determining the color ratio for each site independently and then computing the mean color ratio for sites corresponding to the bright and dim zones. Similarly to earlier, both weighted and unweighted values are computed. Comparing the bright zone values to the dim zone, we can see that both the unweighted and weighted results indicate the same. The proportion of the blue channel's radiance is slightly larger in the dim zone than in the bright zone. Overall the relative radiance measurements suggest that the actual deposition material reflects the green light most strongly, then blue and finally red. This result coincides with the reflectance spectra measured for the quartz chip sample, that are presented in appendix C.

	Bright zone			
	Unweighted	Uncertainty	Weighted	Uncertainty
B	0.4490	0.0017	0.5075	0.0022
G	0.3766	0.0016	0.3434	0.0018
R	0.1744	0.0010	0.1491	0.0011
	Dim zone			
B	0.4799	0.0017	0.5396	0.0021
G	0.3632	0.0014	0.337	0.0016
R	0.1569	0.0008	0.1234	0.0009

Table 12. The mean ratios of each color channel’s relative radiance w.r.t the sum of all channels’ relative radiances.

Finally we should consider which of the computation methods more accurately reflects reality; the unweighted or weighted mean relative radiance. Looking at the unweighted and weighted dim-to-bright ratios, we can see that they are mostly closely matched – except for the red channel. For example the unweighted dim-to-bright ratio for the red channel using the 50x objective changes from 0.6295 to 0.5482 going to weighted, which is not nearly as closely matched with the ratios of the blue and green channels anymore. While typically inverse weighting by uncertainty can be beneficial for the accuracy of inferred values, in this case it is problematic in terms of determining the mean relative radiance. The largest relative radiance values for the red channel have an inherently large uncertainty. This wouldn’t be a problem if all the large values didn’t have also a larger uncertainty. This likely skews the mean value lower than it should be. This problem arising from the weighing is likely caused by the fact that we utilized the absolute values of the uncertainties in the inverse uncertainty weighing. This issue probably wouldn’t have occurred if we were to use relative uncertainties instead of absolute uncertainties. With relative uncertainties the large values get disproportionately small weights due to their absolute uncertainties being naturally larger than for small values.

3.2 Local thickness measurements

The local thickness calculations are an effective way to quantify the sizes of features of interest in images and also their distances from each other. Here we compute some local thickness values for the relative radiance images and SEM images at chosen locations.

3.2.1 Local thickness of the optical images

While we can't determine the sizes of any nanoscale features in the optical images, we can observe and quantify the typical thicknesses and separations of the microscale features of the samples. Using the irradiance domain images we can, relative to normal image, more accurately threshold the images of the MWNT deposition samples and determine how far apart the light reflecting materials of interest are from each other in the images and how large they are. The uncertainty of the local thickness computations for the relative radiance images originate from two sources. We'll set a given value as the threshold and compute the local thickness for it. We can then utilize the uncertainty of the relative radiance images to determine a window for the threshold value and perform the measurement for the lower and upper values of the this window, choosing the maximum deviation as the uncertainty. This results in a local thickness image and an associated uncertainty image. The local thickness image's pixel values are estimates of the features' size in pixels. We can further compute these in units of μm by utilizing the calibrated size of the pixels as determined in part III, section 2.1. The uncertainty of the pixel size can be just added in quadrature to the uncertainty originating from the threshold value. This approach is more of a black-box approach, but a more robust approach for the uncertainty of local thickness computations is beyond this work.

The previous relative radiance measurements come in handy when determining a threshold to use for segmenting the features of interest from the relative radiance images. The local thickness computations are performed only for images utilizing the 50x objective. We compute the local thickness values for each channel separately, using the previously calculated mean relative radiance values for each zone and each channel. Those values themselves were computed by rejecting the background of the empty chip from the computation and thus act as excellent values for segmenting the sample material from the images. The pixels whose relative radiance is greater

than the channel's threshold value are used to compute the size of the features, while pixels whose value falls below the threshold are used to compute the distances between features. We'll call these the local thickness and local separation respectively. Figures 52, 53 and 54 show examples of the original radiance image, binary mask, local thickness and local separation for blue, green and red channels respectively.




Figure 52. Different computed images for the blue channel showing S01P01. Top-left shows the relative radiance of the blue channel in a log scale. Top-right is a binary mask using the given value as a threshold. Bottom-left shows the computed local thickness using the binary mask, while bottom-right shows the local separation using the same mask.




Figure 53. Different computed images for the green channel showing S01P01. Top-left shows the relative radiance of the green channel in a log scale. Top-right is a binary mask using the given value as a threshold. Bottom-left shows the computed local thickness using the binary mask, while bottom-right shows the local separation using the same mask.




Figure 54. Different computed images for the red channel showing S01P01. Top-left shows the relative radiance of the red channel in a log scale. Top-right is a binary mask using the given value as a threshold. Bottom-left shows the computed local thickness using the binary mask, while bottom-right shows the local separation using the same mask.

Figure 55 shows histograms of the local thickness and local separation values for each channel based on the local thickness and local separation images shown in figures 52, 53 and 54. The histograms include the values of all non-zero pixels. Zeros aren't included because this would skew the measures, after all we are interested in




Figure 55. Histograms of the local thickness and local separation values of each channel for S01P01 with non-zero values included.

the sizes and separations of the features, not the not-features so to say. The mean values, their uncertainties and standard deviations for the local thickness and local separation were computed for given sites. These statistics are reported in table 13.

Taking a look at the local thickness values in table 13, the first thing we take note of is that the mean thickness of the features is consistently larger in the bright zone than in the dim zone. This is the most pronounced on the blue channel, while the difference on the green and red channels is more modest. Rather unsurprisingly, there is considerable overlap in the sizes of observed features judging by the standard deviations of the observed local thickness values, but the differences in the mean values are still significantly different even for the green and red channels.

Jumping over to the local separation values, we can see that there isn't really a consistent trend from one zone to the other. The overlap of separation values is much stronger across the zones than for the local thickness values and there are large differences from one sample to the next for both zone types, much more than for the local thickness values. From these local separation statistics the strongest statement we can make is that the separations are likely quite similar across the zones.

Comparing the local thickness and local separation values to each other, we can see that in the dim zones there is a consistent trend of the separation values being

greater than the thickness values. In the bright zone there isn't such a trend, as S03P02 breaks that consistency. Effectively the scale of the thickness and separation values with respect to each other reflects the fraction of the surface area that the segmented features occupy. Similar thickness and separation values suggest a coverage of close to 50%, while a larger thickness over separation tips this fraction to larger values and vice versa for larger separation over thickness. For example in the binary mask images for S01P01 in figures 52, 53 and 54 we can maybe estimate that the black areas, corresponding to features of interest, cover slightly less surface area than the white. This is then reflected in the thickness and separation values shown for S01P01 in table 13.

Finally let's take a look at how the thickness and separation values behave across the three color channels sample-wise. For the thickness values we can see that typically the blue channel has the largest values, followed by red and finally green. For the separation values there are different trends for the dim and bright zones. For the dim zone the largest separation is for red, followed by blue and green. For the bright zone the order changes to blue, red and green. Overall from these observations we can note that the green features are the smallest and most closely packed across both zones. The differences for the sizes of the green and red features are quite small and we've noticed from then relative radiance images that areas with bright green features typically coincide with the bright red features. Something that is not taken into account here is the chromatic aberration. It is clear that chromatic aberration causes different wavelengths of light to come into focus at different distances. Most likely the green channel is the one that is the best focused one, which would mean that both the red and blue channels would be slightly out of focus. The analysis here doesn't take that into account, but it is likely that both the red and blue features would be slightly smaller if the images were captured so that each channel had been imaged separately with maximal focus.

Overall, based on the previous considerations, the green channel likely represents the features of interest most accurately. Based on the different mean values measured for the different sites and their associated uncertainties, this suggests that the mean local thickness of the features of interest in the dim zone is in the range 0.58–0.61 μm for the silicon chip samples, while their local separation is in the range 0.79–0.99 μm . In the bright zone these ranges are 0.75–0.95 μm and 0.54–1.12 μm . This indicates that there is preference in how the material settles in the deposition, with it forming

Dim zone							
Site	Ch	Thickness (μm)			Separation (μm)		
		Mean	Uncert.	SD	Mean	Uncert.	SD
S01P04	B	0.668	0.007	0.480	1.533	0.015	0.729
	G	0.590	0.006	0.397	0.974	0.010	0.346
	R	0.671	0.007	0.433	2.06	0.03	0.75
S02P01	B	0.761	0.008	0.466	1.118	0.011	0.558
	G	0.592	0.006	0.308	0.93	0.01	0.37
	R	0.662	0.007	0.317	1.80	0.02	0.66
S03P01	B	0.870	0.009	0.522	0.982	0.010	0.522
	G	0.605	0.006	0.328	0.806	0.008	0.327
	R	0.651	0.007	0.352	1.54	0.02	0.60
S04P01	B	0.819	0.008	0.561	1.107	0.011	0.582
	G	0.673	0.007	0.426	1.019	0.010	0.435
	R	0.638	0.007	0.413	1.082	0.011	0.496
Bright zone							
S02P02	B	1.116	0.011	0.666	1.519	0.015	1.093
	G	0.758	0.008	0.390	0.810	0.008	0.511
	R	0.870	0.009	0.509	0.979	0.010	0.623
S03P02	B	1.418	0.014	0.834	0.813	0.008	0.685
	G	0.815	0.008	0.406	0.550	0.006	0.329
	R	1.019	0.010	0.588	0.676	0.007	0.426
S01P01	B	1.087	0.011	0.596	1.525	0.015	0.922
	G	0.937	0.010	0.434	1.188	0.012	0.751
	R	0.952	0.010	0.498	1.258	0.013	0.769

Table 13. Computed mean, uncertainty of the mean and standard deviation values for the local thickness and separation of features in the relative radiance images.

thicker features in the bright zone than in the dim zone. The separations between the features suggest similar separations, though the bright zone likely has more variation in this.

3.2.2 Local thickness of the SEM images

Using the SEM images we can get a measure of how large the typical structures making up the sample material are and how far apart they tend to lie from each other. For the SEM images we don't have an uncertainty image, nor a similarly calibrated pixel size like for the optical images. The SEM images do carry pixel size information, which we'll use to convert the results to micrometers, but for the uncertainty we just choose to set a flat 5% uncertainty to the size of the pixels. The SEM images are typical 8-bit single channel images for which we found that the threshold value 110 works well at segmenting the particles from the background. Additionally a single binary close operation was performed on the mask to better match the mask with the desired features. The operation first dilates the mask by one pixel and then erodes it by one pixel, closing holes and gaps in the features while trying not to add any extra size to features. Otherwise the analysis is similar to the one performed on the relative radiance images.

Figures 56 and 57 show examples of the original SEM image and its local thickness and local separation images for sample 01 in the dim and bright zones respectively. The histograms visualizing the non-zero data based on these images are shown in figure 58. Similar thickness and separation computations were performed for multiple SEM images of different sites in sample 01 at 1000x, 3000x and 8000x magnifications. The mean values, their uncertainties and standard deviations for the local thickness and local separation for each site are compiled in table 14.

The results shown in table 14 are likely the most important in terms of explaining the optical properties of the samples. First, looking at the mean local thickness values observed at sites in each zone, we note that the local thickness of the material is significantly larger in the bright zone than in the dim zone. The 3000x results and the single 8000x result are very similar, but the 1000x has a clearly different value. The 1000x result should be disregarded as the resolution of the image w.r.t features in the image is low, resulting in imprecise values. Similarly on the local separation side we can see that the mean separation increases significantly in the bright zone. These results indicate that in the bright zone the mean local thickness of the particles is in the range 0.12–0.15 μm and the separation is in the range 0.51–0.60 μm . In the dim zone these values are in the range 0.06–0.10 μm and 0.17–0.34 ranges respectively.

Considering that local thickness measures radius, we can convert to diameter by doubling the values, yielding an estimate for the mean particle diameter and the

Dim zone							
Site	Magnification	Thickness			Separation		
		Mean	Uncert.	SD	Mean	Uncert.	SD
S01SD01	3000x	0.062	0.004	0.059	0.24	0.02	0.16
S01SD02	8000x	0.060	0.003	0.061	0.178	0.009	0.097
S01SD03	3000x	0.089	0.005	0.054	0.27	0.02	0.17
S01SD04	3000x	0.098	0.005	0.080	0.32	0.02	0.19
S01SD05	3000x	0.098	0.005	0.080	0.29	0.02	0.18
Bright zone							
S01SB01	3000x	0.137	0.007	0.100	0.56	0.03	0.33
S01SB02	1000x	0.27	0.02	0.15	0.80	0.05	0.43
S01SB03	3000x	0.144	0.008	0.133	0.54	0.03	0.32
S01SB04	3000x	0.126	0.007	0.084	0.57	0.03	0.32

Table 14. Computed mean, uncertainty of the mean and standard deviation values for the local thickness and separation of features in the SEM images.

mean distance between particles. This would put the values for the bright zone as 0.24–0.30 μm and 1.02–1.20 μm and for the dim zone as 0.12–0.20 μm and 0.34–0.68 μm . Most importantly this puts the mean distance between particles in the bright zone above the wavelength range of the visible spectrum, while in the dim zone the range covers almost the full wavelength range of the visible spectrum. As we know from part I, for diffraction to occur the feature spacings must be larger than the wavelength of incident light. This suggests that in the bright zone most of the features can readily diffract the full visible spectrum, while in the dim zone features can only diffract a portion of the visible spectrum – if at all.

As for the particle sizes, these values would put them according to the size parameter 21 comfortably about the Mie scattering domain in both zones. Looking back at the unweighted relative radiance measurements and especially those made on images utilizing the 50x objective, we saw that the different channels dimmed reasonably similarly moving from the bright to the dim zone. Similarly the computed color ratios for each zone in table 12 show that the ratios are reasonably similar to each other across the zones, with a slightly stronger preference towards blue in the dim zone over the bright. This makes sense if one thinks of the total scattering being caused by a portion of Mie scatterers and Rayleigh scatterers. Moving from the bright zone to the dim zone, the proportion of Mie scatterers reduces, while the proportion of Rayleigh scatterers increases, i.e. the mean particle size reduces. This finally leads to a slight preference towards blue wavelengths in the dim zone.

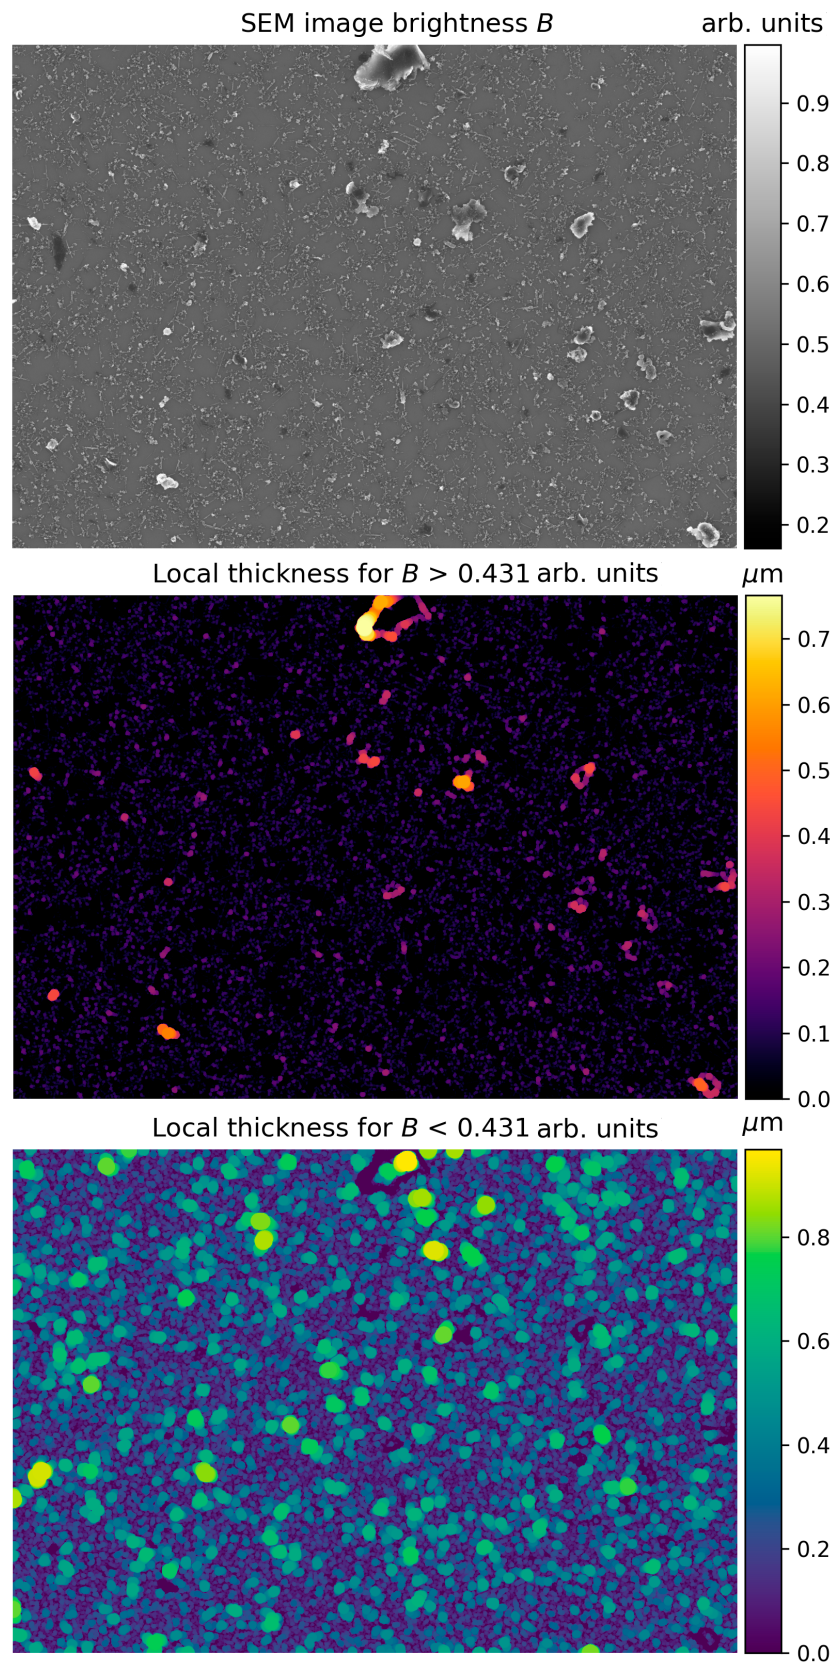


Figure 56. SEM image from S01 dim zone and its local thickness and local separation images.




Figure 57. SEM image S01SB04 from the bright zone and its local thickness and local separation images.




Figure 58. Histograms of the local thickness and local separation images based on the SEM images in the bright and dim zones. a) and b) are based on the images of the bright zone in figure 57, while c) and d) are based on the images of the dim zone in figure 56. On top of the histograms are reported the mean value, uncertainty of the mean and standard deviation for each data set.

3.3 Autocorrelation measurements

The autocorrelation measurements can give us insight into how ordered the structures in the images are. We'll perform these on both the irradiance domain images and the SEM images. The autocorrelation of images was computed utilizing the correlate function of CuPy-cuda12x, so that their values were normalized to range $[-1, 1]$. This was done for each channel by subtracting the spatial mean and then dividing by the spatial standard deviation. The autocorrelation measurements have slight overlap with the local thickness measurements in that the typical size of features

can be also determined from the size of the central peak of an autocorrelation image. This can be further utilized to determine preferential orientations by finding the directions corresponding to the largest and smallest autocorrelation values. The autocorrelations were computed for selected regions of the bright and dim zones of the samples and the regions were cropped so that their width and height were equal.

From the autocorrelation images we find the directions of maximum and minimum autocorrelation with the following process. We sample the directions from -90deg to 90deg in intervals of one degree. For each direction a line profile, whose length is half of the square region's width is sampled via linear interpolation. Then for the line profile the area under the curve is used to assign the maximum direction as the direction with the largest area under the curve, with similar approach used for the minimum direction. Also a radial mean profile is computed by averaging over all the sampled line profiles. Finally a half-maximum distance is determined for the maximum, minimum and mean line profiles. For the optical images each channel is handled separately.

3.3.1 Autocorrelation of the optical images

From the optical images we can determine the typical scale of features visible at each magnification, any preferential orientations and possible periodicity. Figure 59 shows the autocorrelation images of each channel for S01P01. Viewers should take note of the colorbar's scale, what the images show is the central peak that goes up to autocorrelation value one. Moving outwards from there it quickly drops to very low levels. The autocorrelation images show basically just white noise, indicating that beyond the initial peak, no significant autocorrelation is observed.




Figure 59. Autocorrelation images of S01P01 for each color channel.

Figure 60 shows the line profiles of the autocorrelation images in the directions of maximum and minimum autocorrelation and additionally the radial mean autocorrelation. Due to the logarithmic scale of the x-axis, the autocorrelation values on the y-axis don't quite go to one. This is because the zero distance is not included in the logarithmic x-axis. This is more pronounced in the line profile of the SEM image's autocorrelation shown in figure 62. The optical image's line profiles we can more readily observe that there just about only noise after the initial peak. On the blue channel there seems to be just a small positive bump in the mean profile at about the $10 \mu\text{m}$ mark. On the blue channel we can also see what might be initially a weird half-maximum distance value. The half-maximum distance for the minimum profile is greater than for the mean profile. This is simply caused by the method we used to determine the maximum and minimum directions. As stated earlier, we utilized the area under the curve to determine the directions. This can lead to such a situation.

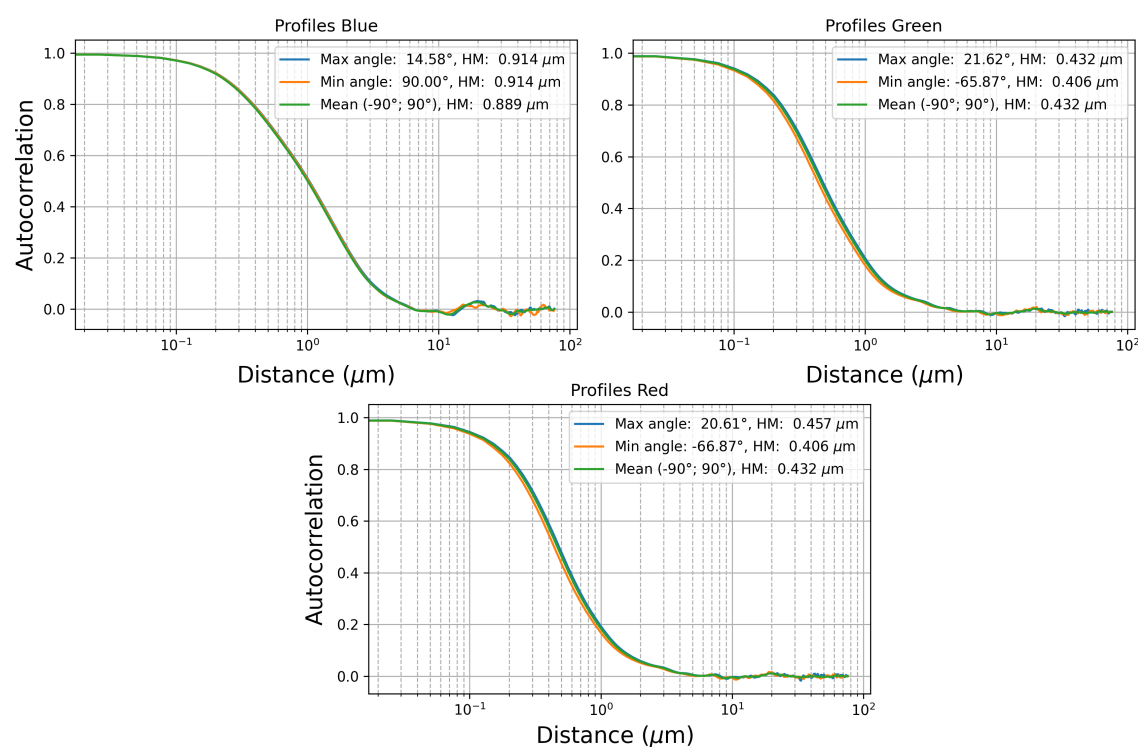


Figure 60. Line profiles of the autocorrelation images shown in figure 59. The profiles are plotted in the directions of maximum and minimum autocorrelation and also a radial mean profile is computed. For each profile the distance to the half-maximum value is computed.

Similar autocorrelation images and line profiles were determined for the same sites that were utilized in the local thickness measurements. The maximum and

minimum directions and the half-maximum distances are compiled in table 15. Based on our observations of the autocorrelation images and half-maximum distance values in table 15, there is practically no emergent periodicities or anisotropy in the relative radiance images. The differences between the half-maximum distances in the minimum and maximum directions correspond to at best one pixel, most of them even less. This indicates that the relative radiance images are very isotropic for all channels. Comparing the half-maximum distances to the local thickness values in table 13 shows that the values for the blue channel are quite similar, while for the green and red channels the half-maximum distances are slightly smaller than what the corresponding local thickness values are. This is caused by the fact that the features are not completely uniform, but that their relative radiance values fall off smoothly as we move away from the center of a feature. The autocorrelation falls off more strongly and leads to a smaller estimate for the typical radius of features if measured by the half-maximum distance. In the blue channel's case our threshold relative radiance and the half-max distance of autocorrelation happen to match, but for the other channels there are two possibilities. Either the local thickness measure overestimates the feature size or the autocorrelation half-max distance underestimates the feature size. The truth of these measurements likely falls somewhere in between these. The important part is that they are still reasonably closely matched and show the same trends from channel to channel and from dim zone to bright zone.

3.3.2 Autocorrelation of the SEM images

In the nanoscale images obtained via SEM we can determine if the webs of nanotubes and amorphous carbon form a periodic structure and if there are any preferential orientations for them. Like for the optical images, figures 61 and 62 show an example autocorrelation and line profile image for S01SB04, whose original SEM image was shown in figures 57.

The results for the different sites are compiled in table 16. As was for the optical autocorrelation results, the results for the SEM images are very similar. The results indicate that the images are very isotropic and that no significant autocorrelation is observed beyond the initial peak. Comparing to the local thickness values we encounter the same trend as in the case of the optical images. The values are somewhat smaller than for their local thickness counterparts, but the trend from dim zone to bright zone persists. Like with the optical image half-maximum distances,

Dim zone						
Site	Ch	Max angle	Min angle	Max HMD	Mean HMD	Min HMD
S01P04	B	83.97	-90.00	0.708	0.680	0.623
	G	90.00	-69.89	0.368	0.368	0.340
	R	90.00	-75.92	0.396	0.396	0.368
S02P01	B	90.00	-90.00	0.608	0.608	0.568
	G	9.55	49.78	0.365	0.365	0.365
	R	11.56	49.78	0.405	0.405	0.405
S03P01	B	63.85	-90.00	0.593	0.593	0.537
	G	49.78	-90.00	0.367	0.339	0.311
	R	53.80	-90.00	0.395	0.395	0.367
S04P01	B	83.97	-86.98	0.649	0.649	0.608
	G	76.93	-85.98	0.487	0.446	0.405
	R	77.93	-80.95	0.446	0.446	0.405
Bright zone						
S01P01	B	14.58	90.00	0.914	0.889	0.914
	G	21.62	-65.87	0.432	0.432	0.406
	R	20.61	-66.87	0.457	0.432	0.406
S02P02	B	61.84	-85.98	0.680	0.635	0.612
	G	59.83	-65.87	0.363	0.340	0.317
	R	32.68	-71.90	0.385	0.385	0.363
S03P02	B	70.89	-55.81	0.568	0.549	0.529
	G	16.59	-90.00	0.353	0.333	0.313
	R	90.00	-90.00	0.392	0.392	0.353

Table 15. Autocorrelation statistics for each of the relative radiance images' channels. Directions of maximum and minimum autocorrelation and the distance to half-maximum autocorrelation in the maximum and minimum directions and for the mean radial autocorrelation.

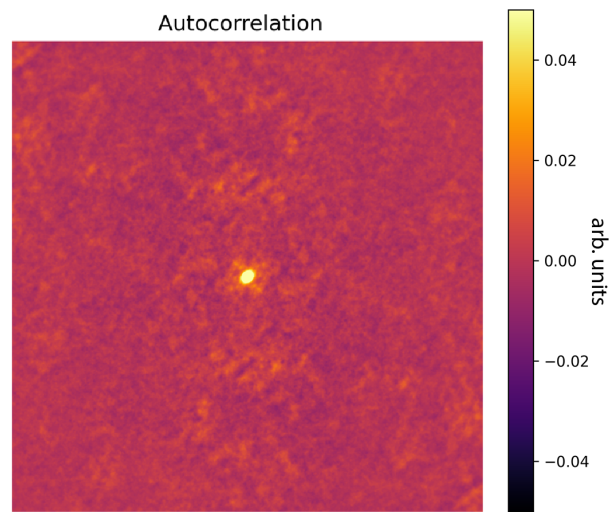


Figure 61. Autocorrelation image of S01SB04.




Figure 62. Line profile of the autocorrelation image shown in figure 61. The profiles are plotted in the directions of maximum and minimum autocorrelation and also a radial mean profile is computed. For each profile the distance to the half-maximum value is computed.

these likely slightly under estimate the radii of the particles. Most likely the true values of the mean particle radius lies somewhere between the values observed here and those determined by the local thickness measurements. Regardless that doesn't really change our the deductions we made earlier based on Mie scatterers and Rayleigh scatterers. With the typical particle sizes observed here, we still don't quite fall into the Rayleigh scattering domain in terms of the mean particle size.

Dim zone					
Site	Max angle	Min angle	Max HMD	Mean HMD	Min HMD
S01SD01	-7.54	67.88	0.02	0.02	0.02
S01SD02	-22.63	85.98	0.016	0.016	0.016
S01SD03	80.95	-50.78	0.02	0.02	0.02
S01SD04	90	-90	0.024	0.024	0.024
S01SD05	3.52	90	0.024	0.024	0.024
Bright zone					
S01SB01	75.92	-90	0.032	0.032	0.029
S01SB02	25.64	-47.77	0.051	0.051	0.051
S01SB03	45.75	-56.82	0.032	0.032	0.032
S01SB04	-13.58	90	0.028	0.028	0.028

Table 16. Autocorrelation statistics for the SEM images. Directions of maximum and minimum autocorrelation and the distance to half-maximum autocorrelation in the maximum and minimum directions and for the mean radial autocorrelation.

4 Conclusions

In this research we have done a wide variety of things. We opened with an inspection of the MWNT depositions via optical, AFM and SEM images. We formulated an objective function based on the expected linearity of a digital camera, which can be minimized to solve a camera's ICRF and linearize its response, inspired by the work of Matsushita and Lin[50]. We implemented an image processing pipeline to produce HDR images that function as best estimates for a scene's radiance, taking Debevec and Malik's work[39] as a starting point, while making additions and modifications based on other works and of our own. We put our method of solving the ICRF and processing images into HDR images under scrutiny. Finally we put the image processing into use and performed relative radiance, local thickness and autocorrelation measurements on the optical and SEM images.

The algorithm for solving the ICRF yielded quality results, producing a very closely matched result with the external data set and based on what we observed with the disparity plots for Olympus SC100. It'd be interesting to see how the algorithm fares for a wider variety of digital cameras. It is highly likely that there are many parameters, whose optimal state is device dependent. Which principal components to utilize, which rejection thresholds to set and what kind of parameters to use in the differential evolution solver. Furthermore it would be interesting to see how different global optimization methods would work with the problem. Especially because the differential evolution algorithm is relatively slow due to its stochastic nature.

The HDR merging produced images of high quality and accuracy in terms of linearity. As we observed in part III, section 2.4, the ratios of the zone-wise mean relative radiance from the image with an analyzer against the image without the analyzer were quite normally distributed. The unweighted mean ratios were 0.223, 0.238 and 0.248 with SDs of 0.010, 0.011 and 0.029 for the blue, green and red channels respectively. This showed that the images are highly linear and that each channel's relative radiances are related to absolute radiance by an unknown scalar factor. We also observed how the fixed-pattern correction and bad pixel filtering

worked in the HDR-merged images. The mean-stacked images of a piece of paper in different positions worked well at correcting the non-uniformities in the HDR-merged images. In terms of vignetting, the experimental approach likely works better than utilizing a theoretical equation, as the experimental approach also notes the possible different behaviors of each channel and non-symmetric behavior.

Finally using the relative radiance images we measured how the relative radiance values changed for each channel when moving from a sample's dim zone to the bright zone. For the areas that minimize the contribution of the silicon chip, we determined that the relative radiance drops to about 64% for blue, 68% for green and 63% for of their value in the bright zone. For the total relative radiance, which includes the contribution of the empty chip, we determined a drop to about 51% or 55% for blue, 45% or 43% for green and 41% for red, based on images utilizing the 5x and 10x objectives respectively. The precise numbers for these are reported in table 10. The key takeaway is that in the 50x objective images, in which the background of the empty chip was segmented away, the drop is more uniform across the channels than in the 10x and 5x objective images, in which the blue channel drops much less than the green and red channels.

In the local thickness measurements of the optical images we wanted to best segment the parts that are most likely to correspond to the sample material. We utilized the measured zone-wise mean relative radiance values to segment the sample images. Naturally on the basis of the optical images we don't get to measure the actual dimensions of the nanoscale material, but we do get to measure the typical dimensions of the bright features in each zone. With the green channel being likely the most in focus and the sample material itself reflecting most strongly in the green wavelength range, we found that in the optical images the mean diameter of features was about 1.16–1.22 μm in the dim zone and 1.50–1.90 μm in the bright zone. The mean distance between features was found to be 1.58–1.98 μm and 1.08–2.24 μm in the dim and bright zones respectively. This shows that the features of interest are typically larger in the bright zone than in the dim zone, while the distances between features are reasonably similar for both zones. The absolute values of these results would change slightly with different threshold values for the binary mask. However if the same basis for the threshold value was used for both zones, the trend should logically remain the same.

The local thickness measurements of the SEM images are the most important

in terms of explaining the optical properties of the MWNT depositions. We found that the mean particle diameter was 0.24–0.30 μm and 0.12–0.2 μm for the bright and dim zones respectively. The mean distances between particles were found to be 1.02–1.20 μm and 0.34–0.68 μm .

The autocorrelation measurements showed us that there were no significant periodic features in the samples in either of the relative radiance images or SEM images. The samples also turn out to be highly isotropic. The half-maximum distances observed for both the optical and SEM images reflect the results of the local thickness measurements. The differences are minor and likely caused by the half-maximum distance underestimating the size of features.

From these results we constituted the following explanation for the optical properties of the depositions. Due to the typical distances between particles in the bright zone, it can cause diffraction to occur over the full range of the visible spectrum, while the dim zone can cause diffraction for only a portion or not at all in the range of the visible spectrum. Diffraction does, however, require periodic features and as we noted based on the autocorrelation measurements, there are no strong periodicities in the deposited material. If any diffraction phenomena occur, they are likely very localized and randomly oriented. The more important part in explaining the optical properties are the particle sizes in each zone. The observed particle sizes mean that in both zones a large portion of the particles are Mie scatterers, resulting on their part in a uniform response across the visible spectrum. A portion of scatterers on the smaller size in dimensions act as Rayleigh scatterers, which is stronger for shorter wavelengths, i.e. blue. The proportion of Mie scatterers falls when moving from the bright zone to the dim, leading to a slight preference in the blue intensity when compared to the bright zone. Overall most of the intensity difference between the bright and dim zone are likely simply caused by the differences in typical particle size for each zone, as scattering intensity increases strongly with particle size. Finally the silicon chip adds its contribution at close distances to particles. We've noted that at all locations where there is material on the chip, a slight glow surrounds it, which is most intense on the blue channel. This likely occurs due to incident light reflecting from deposited material to the silicon chip and then towards the objective. What causes the selectivity in the particle sizes across the zones is at this point unknown and is a potential subject for a research of its own.

Overall the analysis has shown the potential of HDR images as a tool in optical

microscopy. It enables quantitative analysis related to brightness and opens powerful ways to visualize samples, whose scattered light intensity has a very large dynamic range. Additionally it makes segmentation of images based on pixel values much easier in comparison to LDR images. The image analysis methods used here were overall relatively simple and it would be interesting to see how more complicated analyses might benefit from HDR imaging. The linearity measure we established here could also act as a tool to analyze the linearity of even supposedly linear cameras. While we can't say that any random digital camera could be calibrated to be a radiance sensor, we have shown one example of this and it would be interesting to see how well this approach would work on a wider selection of cameras.

Finally in terms of uncertainty analysis we established a robust method of estimating the uncertainty of the relative radiance images. We did however perform the analysis with the assumption that the solved ICRF itself is accurate. In our analysis of the pair of images with no analyzer and analyzer we saw that the expected linearity is not perfect. The SD of the values in table 9 is multiple times larger than the uncertainty. Likely a large portion of the spread measured by the SD values originates from the imperfect linearization, because the standard deviation should be much closer to the uncertainty values, were the linearization perfect. It is also most probable that in the case of our camera there exists no ICRF that could, within the bounds of uncertainty established from the original mean stacked 8-bit sRGB input images, perfectly linearize the image and also fulfill the requirement of being invertible. The quality these results in the end is based off of how accurate the empirical model of camera response functions is. For example if a camera utilizes a piecewise function to map exposure values to image brightness, it would slightly throw off the quality of linearization. In the HDR merging the weighing function we chose alleviates these problems by giving the most weight for brightness values in the middle of the brightness range. However in the case of perfect linearization no weighing function would be required. In the end the uncertainty of the linearization itself is not large enough to throw off the analysis we performed and would instead just increase the bounds of uncertainty. Though in a certain sense it is beneficial to see how only the uncertainty of the original input images carries on through the image processing and analysis, as then we can identify what spread is caused by the imperfection of the linearization and what is caused by the original measured uncertainty.

References

- [1] G. Michler. *Compact Introduction to Electron Microscopy*. Springer, 2022. ISBN: 978-3-658-37363-4.
- [2] B. Voigtländer. *Atomic Force Microscopy*. 2nd. Springer, 2019. ISBN: 978-3-030-13653-6.
- [3] G. Hlawacek and A. Götzhäuser. *Helium Ion Microscopy*. NanoScience and Technology. Springer, 2016. ISBN: 978-3-319-41988-6.
- [4] A. Biswas et al. "Advances in top-down and bottom-up surface nanofabrication: Techniques, applications & future prospects". In: *Advances in Colloid and Interface Science* 170 (2012), pp. 2–27. DOI: 10.1016/j.cis.2011.11.001.
- [5] S. Kumari et al. "A comprehensive review on various techniques used for synthesizing nanoparticles". In: *Journal of Materials Research and Technology* 27 (2023), pp. 1739–1763. DOI: 10.1016/j.jmrt.2023.09.291.
- [6] Z. Liu et al. "Wafer-scale synthesis of two-dimensional materials for integrated electronics". In: *Chip* 3 (2024). DOI: 10.1016/j.chip.2023.100080.
- [7] S. Ahire et al. "The Augmentation of nanotechnology era: A concise review on fundamental concepts of nanotechnology and applications in material science and technology". In: *Results in Chemistry* 4 (2022). DOI: 10.1016/j.rechem.2022.100633.
- [8] R. Ratnesh. "Advancement and challenges in MOSFET scaling". In: *Materials Science in Semiconductor Processing* 134 (2021). DOI: 10.1016/j.mssp.2021.106002.
- [9] S. Iijima. "Helical microtubules of graphitic carbon". In: *Nature* 354 (1991), pp. 56–58. DOI: 10.1038/354056a0.
- [10] M.F. De Volder et al. "Carbon Nanotubes: Present and Future Commercial Applications". In: *Science* 339 (2013), pp. 535–539. DOI: 10.1126/science.1222453.

- [11] Y. Li. "The Quarter-Century Anniversary of CarbonNanotube Research". In: *ACS Nano* 11 (2017), pp. 1–2. DOI: 10.1021/acsnano.7b00232.
- [12] E. Randviir, D. Brownson, and C. Banks. "A decade of graphene research: production, applications and outlook". In: *Materials Today* 17 (2014), pp. 426–432. DOI: 10.1016/j.mattod.2014.06.001.
- [13] B. Kaushik and M. Majumder. *Carbon Nanotube: Properties and Applications*. In: *Carbon Nanotube Based VLSI Interconnects*. New Delhi: Springer, 2014.
- [14] K. Rakesh et al. "Overview on metamaterial: History, types and applications". In: *Materials Today: Proceedings* 56 (2022), pp. 3016–3024. DOI: 10.1016/j.matpr.2021.11.423.
- [15] N. M. Litchinitser et al. "Chapter 1 Negative refractive index metamaterials in optics". In: ed. by E. Wolf. Vol. 51. *Progress in Optics*. Elsevier, 2008, pp. 1–67. DOI: 10.1016/S0079-6638(07)51001-2.
- [16] Q. Zhang and Y. Sun. "Anisotropic thermal expansion based on a novel metamaterial". In: *International Journal of Mechanical Sciences* 268 (2024). DOI: 10.1016/j.ijmecsci.2024.109024.
- [17] R. Knight. *Physics for Scientists and Engineers: A Strategic Approach with Modern Physics, 3rd edition*. Pearson, 2014. ISBN: 978-1-292-02078-5.
- [18] E. Slayter and H. Slayter. *Light and Electron Microscopy*. Trumpington Street, Cambridge CB2 1RP, United Kingdom: Cambridge University Press, 1992.
- [19] E. Hecht. *Optics, 5th edition*. Pearson, 2017. ISBN: 978-0-133-97722-6.
- [20] A. Lipson, S. Lipson, and H. Lipson. *Optical Physics, 4th edition*. Cambridge University Press, 2011. ISBN: 978-0-521-49345-1.
- [21] D. Murphy and M. Davidson. *Fundamentals of Light Microscopy and Electronic Imaging: 2nd edition*. Trumpington Street, Cambridge CB2 1RP, United Kingdom: John Wiley and Sons, 2012. DOI: 10.1002/9781118382905.
- [22] A. Zangwill. *Modern Electrodynamics*. Cambridge University Press, 2012. ISBN: 978-0-521-89697-9.
- [23] I. Morison. *Introduction to Astronomy and Cosmology*. NanoScience and Technology. Wiley, 2008. ISBN: 978-0-470-03334-0.

- [24] R. Krishna, T. Unsworth, and R. Edge. "Raman Spectroscopy and Microscopy". In: *Reference Module in Materials Science and Materials Engineering* 162 (2016). DOI: 10.1016/B978-0-12-803581-8.03091-5.
- [25] S. Wolf and N. Voshchinnikov. "Mie scattering by ensembles of particles with very large size parameters". In: *Computer Physics Communications* 162 (2004), pp. 113–123. DOI: 10.1016/j.cpc.2004.06.070.
- [26] A. Höpe. "Chapter 6 - Diffuse Reflectance and Transmittance". In: *Spectrophotometry*. Ed. by T. A. Germer, J. C. Zwinkels, and B. K. Tsai. Vol. 46. Experimental Methods in the Physical Sciences. Academic Press, 2014, pp. 179–219. DOI: <https://doi.org/10.1016/B978-0-12-386022-4.00006-6>.
- [27] M. Abramowitz and M. Davidson. *Reflected Light Microscopy Anatomy*. URL: <https://micro.magnet.fsu.edu/primer/anatomy/reflected.html> (visited on 01/17/2022).
- [28] E. Dereniak and T. Dereniak. *Geometrical and trigonometric optics*. The Edinburgh Building, Cambridge CB2 8RU, United Kingdom: Cambridge University Press, 2008.
- [29] G. Slunder and J. Nordberg. *Digital Microscopy, Chapter 1 – Microscope Basics*. Academic Press, 2013, pp. 1–10. DOI: 10.1016/B978-0-12-407761-4.00001-4.
- [30] M. Davidson. *Tungsten-Halogen Incandescent Lamps*. URL: <http://zeiss-campus.magnet.fsu.edu/articles/lightsources/tungstenhalogen.html> (visited on 06/18/2024).
- [31] M. Davidson. *Fundamentals of Mercury Arc Lamps*. URL: <http://zeiss-campus.magnet.fsu.edu/articles/lightsources/mercuryarc.html> (visited on 06/18/2024).
- [32] M. Davidson. *Fundamentals of Xenon Arc Lamps*. URL: <http://zeiss-campus.magnet.fsu.edu/articles/lightsources/xenonarc.html> (visited on 06/18/2024).
- [33] M. Davidson. *Fundamentals of Light-Emitting Diodes (LEDs)*. URL: <https://zeiss-campus.magnet.fsu.edu/articles/lightsources/leds.html> (visited on 06/18/2024).

- [34] R. Rottenfusser, E. Wilson, and M. Davidson. *Illumination and the Optical Train*. URL: <https://zeiss-campus.magnet.fsu.edu/articles/basics/opticaltrain.html> (visited on 04/09/2022).
- [35] B. Yakobson and P. Avouris. *Mechanical Properties of Carbon Nanotubes*. Vol. 80. 2001, pp. 287–327. ISBN: 978-3-540-41086-7. DOI: 10.1007/3-540-39947-X_12.
- [36] E. Allen and S. Tiantaphillidou. *The Manual of Photography, 10th edition*. Elsevier Ltd., 2011. ISBN: 978-0-240-52037-7.
- [37] K. Tatsuno. "Current trends in digital cameras and camera-phones". In: *Science & Technology Trends* 18 (2006), pp. 35–44.
- [38] E. R. Fossum. "The Invention of CMOS Image Sensors: A Camera in Every Pocket". In: *Proceedings of the SMTA Pan Pacific Microelectronics Symposium 2020*. 2020, pp. 4335–4343. DOI: 10.23919/PanPacific48324.2020.9059308.
- [39] P. Debevec and J. Malik. "Recovering high dynamic range radiance maps from photographs". In: *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques* (1997), pp. 369–378. DOI: 10.1145/258734.258884.
- [40] A. Barbara et al. "Microscope spectrometer for light scattering investigations". In: *Applied Optics* 49 (2010). DOI: 10.1364/AO.49.004193.
- [41] D. Anand et al. "Shade selection: spectrophotometer vs digital camera – a comparative in-vitro study". In: *Annals of Prosthodontics & Restorative Dentistry* 2 (2016).
- [42] E. Cal et al. "Application of a digital technique in evaluating the reliability of shade guides". In: *Journal of Oral Rehabilitation* 31 (2004), pp. 483–491. DOI: 10.1111/j.1365-2842.2004.01197.x..
- [43] F. Jarad, M. Russell, and B. Moss. "The use of digital imaging for colour matching and communication in restorative dentistry". In: *British Dental Journal* 199 (2005), pp. 43–49. DOI: 10.1038/sj.bdj.4812559.
- [44] V. Bhandari et al. "Reliability of Digital Camera over Spectrophotometer in measuring the Optical Properties of Orthodontic Elastomeric Ligatures". In: *The Journal of Indian Orthodontic Society* 48 (2014). DOI: 10.5005/jp-journals-10021-1252.

- [45] M. Grossberg and S. Nayar. "What is the Space of Camera Response Functions?" In: *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.* (2003). DOI: 10.1109/CVPR.2003.1211522.
- [46] A. Hanbury. "Constructing cylindrical coordinate colour spaces". In: *Pattern Recognition Letters* 29 (2008), pp. 494–500. DOI: 10.1016/j.patrec.2007.11.002.
- [47] "Amendment 1 - Multimedia systems and equipment - Colour measurement and management - Part 2-1: Colour management - Default RGB colour space - sRGB". In: *IEC 61966-2-1:1999/AMD1:2003* (2003).
- [48] A. Akyüz and A. Genctav. "A reality check for radiometric camera response recovery algorithms". In: *Computers & Graphics* 37 (2013), pp. 935–943. DOI: 10.1016/j.cag.2013.06.003.
- [49] R. Berry and J. Burnell. *The Handbook of Astronomical Image Processing*. Willmann-Bell, Inc., 2005. ISBN: 9780943396828.
- [50] Y. Matsushita and S. Lin. "Radiometric Calibration from Noise Distributions". In: *2007 IEEE Conference on Computer Vision and Pattern Recognition* (2007). DOI: 10.1109/CVPR.2007.383213.
- [51] M. Santo et al. "Standard uncertainty evaluation in image-based measurements". In: *Measurement* (2004), pp. 347–358. DOI: 10.1016/j.measurement.2004.09.011.
- [52] H. Ku. "Notes on the Use of Propagation of Error Formulas". In: *Journal of Research of the National Bureau of Standards. Section C: Engineering and Instrumentation* 70C (1966), p. 263. DOI: 10.6028/jres.070C.025.
- [53] B. Welford. "Note on a Method for Calculating Corrected Sums of Squares and Products". In: *Technometrics* 4 (1962), pp. 419–420. DOI: 10.1080/00401706.1962.10490022.
- [54] E. Schubert and M. Gertz. "Numerically Stable Parallel Computation of (Co-)Variance". In: *SSDBM '18: Proceedings of the 30th International Conference on Scientific and Statistical Database Management* 10 (2018), pp. 1–12. DOI: 10.1145/3221269.3223036.

- [55] I. Jolliffe and J. Cadima. "Principal component analysis: a review and recent developments". In: *Philosophical Transactions Royal Society A* 347 (2016). DOI: 10.1098/rsta.2015.0202.
- [56] C. Chen, S. McCloskey, and J. Yu. "Analyzing Modern Camera Response Functions". In: *2019 IEEE Winter Conference on Applications of Computer Vision* (2019). DOI: 10.1109/WACV.2019.00213.
- [57] S. Mann and M. Ali. *Chapter 1 - The Fundamental Basis of HDR: Comparametric Equations*. Academic Press, 2016. ISBN: 9780081004128. DOI: 10.1016/B978-0-08-100412-8.00001-2.
- [58] O. Gallo and P. Sen. *Chapter 3 - Stack-Based Algorithms for HDR Capture and Reconstruction*. Academic Press, 2016. ISBN: 9780081004128. DOI: 10.1016/B978-0-08-100412-8.00003-6.
- [59] M. Robertson, S. Borman, and R. Stevenson. "Estimation-theoretic approach to dynamic range enhancement using multiple exposures". In: *Journal of Electronic Imaging* 12 (2003), pp. 219–228. DOI: 10.1117/1.1557695.
- [60] B. Ma et al. "A new method of CCD dark current correction via extracting the dark Information from scientific images". In: *High Energy, Optical, and Infrared Detectors for Astronomy VI*. Ed. by A. D. Holland and J. Beletic. SPIE, 2014. DOI: 10.1117/12.2055416.
- [61] J. Dunlap, E. Bodegom, and W. R. "Correction of dark current in consumer cameras". In: *Journal of Electronic Imaging* 19 (2010). DOI: 10.1117/1.3358365.
- [62] S. Ray. *Applied Photographic Optics*. 3rd. Focal Press, 2002. ISBN: 978-0-240-51540-3.
- [63] D. Olsen et al. "Radiometric Calibration for AgCam". In: *Remote Sensing* 2 (2010), pp. 464–477. DOI: 10.3390/rs2020464.
- [64] M. Brady and G. Legge. "Camera calibration for natural image studies and vision research". In: *Journal of the Optical Society of America A* 26 (1 2009), pp. 30–42. DOI: 10.1364/JOSAA.26.000030.

- [65] J. Semmlow. "Chapter 2 - Basic Concepts in Signal Processing". In: *Signals and Systems for Bioengineers (Second Edition)*. Ed. by J. Semmlow. Second Edition. Biomedical Engineering. Academic Press, 2012, pp. 35–80. DOI: <https://doi.org/10.1016/B978-0-12-384982-3.00002-X>.
- [66] R. Heilbronner. "The autocorrelation function: an image processing tool for fabric analysis". In: *Tectonophysics* 212 (1992), pp. 351–370. DOI: 10.1016/0040-1951(92)90300-U.
- [67] C. Robertson and S. George. "Theory and practical recommendations for autocorrelation-based image correlation spectroscopy". In: *Journal of Biomedical Optics* 17 (2012). DOI: 10.1117/1.JBO.17.8.080801.
- [68] B. Dougherty. *Local Thickness*. URL: https://imagej.net/imagej-wiki-static/Local_Thickness (visited on 01/30/2024).
- [69] T. Hildebrand and P. Rügsegger. "Theory and practical recommendations for autocorrelation-based image correlation spectroscopy". In: *Journal of Microscopy* 185 (1997), pp. 67–75. DOI: 10.1046/j.1365-2818.1997.1340694.x.
- [70] T. Saito and J. Toriwaki. "New algorithms for euclidean distance transformation of an n-dimensional digitized picture with applications". In: *Pattern Recognition* 27 (1994), pp. 1551–1565. DOI: 10.1016/0031-3203(94)90133-3.
- [71] E. Remy and E. Thiel. "Exact medial axis with euclidean distance". In: *Image and Vision Computing* 23 (2005), pp. 167–175. DOI: 10.1016/j.imavis.2004.06.007.
- [72] J. Qiang and C. Mitchell. "A Unified Differential Evolution Algorithm for Global Optimization". In: *IEEE Transactions on Evolutionary Computation* (2014). URL: <https://www.osti.gov/biblio/1163659>.
- [73] J. Engel, V. Usenko, and D. Cremers. "A Photometrically Calibrated Benchmark For Monocular Visual Odometry". In: *ArXiv abs/1607.02555* (2016). DOI: 10.48550/arXiv.1607.02555.
- [74] D. Lide. *CRC Handbook of Chemistry and Physics, Internet Version 2005*. Boca Raton: CRC Press, 2005.
- [75] M. Bohr and I. Young. "CMOS Scaling Trends and Beyond". In: *IEEE Micro* 37 (2017), pp. 20–29. DOI: 10.1109/MM.2017.4241347.

- [76] R. Bracewell. *The Fourier Transform and Its Applications, 3rd edition*. McGraw Hill, 2000. ISBN: 0-07-303938-1.

A Uncertainty of the HDR merging

The uncertainty associated with different computed quantities in this research is estimated by applying the uncertainty propagation law to the utilized equations. The uncertainty propagation law states that for a function f of N mutually independent variables x_n with an associated uncertainty δx_n , the uncertainty is given by

$$\delta f(x_0, \dots, x_{N-1}) = \sqrt{\sum_{n=0}^{N-1} \left(\frac{\partial f}{\partial x_n} \delta x_n \right)^2}. \quad (56)$$

First we'll derive the uncertainty of a HDR merged image given by equation 45. First it must be noted that $X_n(x, y) = g(B(x, y))$. First we'll utilize the following substitutions (we momentarily drop the explicit (x,y) notation of images):

$$S = \sum_{n=0}^{N-1} w(B_n), \quad (57)$$

$$G = \sum_{n=0}^{N-1} w(B_n)g(B_n)t_n^{-1}. \quad (58)$$

This allows us to express equation 45 as a simple quotient $E_{\text{HDR}} = G/S$, which can further be expressed as a product of $1/S$ and G to which we can apply the product rule of derivatives. Let us first consider the derivative for the first input image B_0 . We have

$$\frac{\partial E}{\partial B_0} = \frac{\partial}{\partial B_0} \frac{G}{S} = \frac{\partial G}{\partial B_0} \cdot \frac{1}{S} - \frac{G}{S^2} \frac{\partial S}{\partial B_0} \quad (59)$$

from our substitutions. Then we can simply compute the derivatives for G and S , which are

$$\frac{\partial G}{\partial B_0} = w'(B_0)g(B_0)t_0^{-1} + w(B_0)g'(B_0)t_0^{-1}, \quad (60)$$

$$\frac{\partial S}{\partial B_0} = w'(B_0), \quad (61)$$

in which we again utilize the product rule and the fact that the differentiating w.r.t B_0 sets all the other terms in the sums to zero. Substituting these back to equation

59 and reverting the substitution for G and S yields

$$\frac{\partial E}{\partial B_0} = \frac{w'(B_0)g(B_0)t_0^{-1} + w(B_0)g'(B_0)t_0^{-1}}{\sum_{n=0}^{N-1} w(B_n)} - \frac{w'(B_0)w(B_0)g(B_0)t_0^{-1}}{\left(\sum_{n=0}^{N-1} w(B_n)\right)^2}. \quad (62)$$

From this it is straightforward to see that we'll have one of these terms for each input image B_n giving us the uncertainty equation

$$\delta E = \sqrt{\sum_{n=0}^{N-1} \left(\left(\frac{w'(B_n)g(B_n) + w(B_n)g'(B_n)}{\sum_{n=0}^{N-1} w(B_n)} - \frac{w'(B_n)w(B_n)g(B_n)}{\left(\sum_{n=0}^{N-1} w(B_n)\right)^2} \right) \frac{\delta B_n}{t_n} \right)^2}. \quad (63)$$

Here we have assumed that the uncertainty of the images' exposure times is negligible. For our mean stacked images this should be a fair assumption due to the large number of utilized images. The uncertainty analysis would be considerably more complex if the uncertainty of the exposure times was taken into account, as the image brightnesses B_n are directly dependent on their exposure times t_n . In an algorithm the derivatives of the ICRF and the weighting function can be determined using analytical or finite difference methods depending on how they are parameterized. In our case the derivative of the ICRF will be determined via finite difference and the weighting function will be chosen so that an analytical derivative can be utilized.

B Tables of relative radiance measurements

This appendix contains the spatial mean values for selected zones in the samples used in this research. The sample names are written in shorthand, for example S01P01 means sample 1 position 1. The mean relative radiance values are listed for each channel along with the used magnification in the image. Also reported are the standard error for the spatial mean and the standard deviation of the observed relative radiance values. The values are computed via two methods, that are the unweighted and weighted values. For each pixel's relative radiance value the inverse of the standard error value for that pixel was utilized as a weight. Table 17 contains the values computed for selected areas of the samples' dim zones, while table 18 contains the values for areas in the bright zones.

Site	Mag	Ch	\bar{L}	$\delta\bar{L}$	σ_L	\bar{L}_w	$\delta\bar{L}_w$	$\sigma_{L,w}$
S01P04	50x	B	7.1036	0.0349	2.3782	6.5532	0.0349	1.9566
		G	7.0288	0.0392	7.3803	4.4703	0.0392	3.5400
		R	3.0195	0.0297	4.3615	1.4196	0.0297	1.6572
S02P01	50x	B	7.9999	0.0384	2.2864	7.4132	0.0384	1.9289
		G	6.8144	0.0377	6.9299	4.5850	0.0377	3.4337
		R	3.0440	0.0302	4.2276	1.6024	0.0302	1.7379
S03P01	50x	B	8.3088	0.0394	2.2339	7.8170	0.0394	1.9284
		G	7.2775	0.0407	7.1268	4.7862	0.0407	3.6986
		R	3.0818	0.0303	4.0714	1.6775	0.0303	1.6956
S04P01	50x	B	3.8450	0.0184	1.5487	3.4943	0.0184	1.0375
		G	4.0514	0.0221	3.1258	3.3241	0.0221	1.5107
		R	2.0605	0.0208	2.1302	1.5534	0.0208	1.0203
S02F01	5x	B	2.6357	0.0117	0.3173	2.5971	0.0117	0.3173
		G	1.8289	0.0105	0.6761	1.6913	0.0105	0.5643
		R	0.8083	0.0080	0.4179	0.6587	0.0080	0.3155
S02F02	5x	B	2.9268	0.0130	0.3622	2.8834	0.0130	0.3552
		G	2.0003	0.0116	0.7647	1.8418	0.0116	0.6343
		R	0.8575	0.0085	0.4765	0.6963	0.0085	0.3412
S02F03	5x	B	2.2901	0.0100	0.2873	2.2609	0.0100	0.2756
		G	1.6332	0.0098	0.5794	1.5053	0.0098	0.5123
		R	0.7584	0.0074	0.3774	0.6494	0.0074	0.2730
S03F01	10x	B	6.1159	0.0295	0.7668	5.9782	0.0295	0.7540
		G	4.1158	0.0233	2.3396	3.3821	0.0233	1.6304
		R	1.7031	0.0166	1.2083	1.2121	0.0166	0.7952
S03F02	10x	B	5.7178	0.0257	0.8934	5.5791	0.0257	0.8403
		G	3.9932	0.0230	2.4843	3.2386	0.0230	1.6641
		R	1.7059	0.0166	1.3390	1.2077	0.0166	0.8107
S03F03	10x	B	5.5116	0.0249	0.9414	5.3446	0.0249	0.8882
		G	3.8471	0.0222	1.9670	3.2833	0.0222	1.4694
		R	1.5898	0.0155	1.2069	1.1456	0.0155	0.7340
S04F01	10x	B	3.2048	0.0157	0.5745	3.1540	0.0157	0.4545
		G	3.1076	0.0181	1.0224	2.9823	0.0181	0.7167
		R	1.5919	0.0152	0.7052	1.4369	0.0152	0.4584

Table 17. Table of the spatial mean relative radiance for each channel computed at different positions in the samples' dim zones.

Site	Mag	Ch	\bar{L}	$\delta\bar{L}$	σ_L	\bar{L}_w	$\delta\bar{L}_w$	$\sigma_{L,w}$
S01P01	50x	B	11.7790	0.0577	4.0962	10.3773	0.0577	3.9039
		G	9.9186	0.0584	7.6984	6.6219	0.0584	4.8099
		R	4.6918	0.0473	4.7209	2.7692	0.0473	2.5514
S02P02	50x	B	12.3431	0.0608	3.4378	11.5173	0.0608	3.3056
		G	10.5412	0.0599	8.3819	6.8472	0.0599	5.3223
		R	4.9515	0.0489	5.0241	2.8107	0.0489	2.6032
S03P02	50x	B	12.4253	0.0603	3.4850	11.4788	0.0603	3.5172
		G	10.3918	0.0597	7.9284	6.9710	0.0597	4.9905
		R	4.8849	0.0484	4.6575	2.9929	0.0484	2.5463
S02F04	5x	B	5.1489	0.0255	0.4316	5.1386	0.0255	0.4202
		G	4.1056	0.0246	1.1711	3.8363	0.0246	1.1011
		R	1.9713	0.0189	0.7601	1.8024	0.0189	0.6162
S02F05	5x	B	5.0976	0.0263	0.4737	5.0615	0.0263	0.4890
		G	3.8965	0.0229	1.2568	3.4506	0.0229	1.2339
		R	1.8514	0.0185	0.8040	1.5881	0.0185	0.7248
S02F06	5x	B	5.0496	0.0245	0.7091	4.9859	0.0245	0.7261
		G	4.1190	0.0236	1.4733	3.5488	0.0236	1.5090
		R	2.0367	0.0203	0.8768	1.6956	0.0203	0.7730
S03F04	10x	B	9.9452	0.0443	1.3452	9.8088	0.0443	1.3345
		G	8.3857	0.0473	3.5002	7.0166	0.0473	3.2711
		R	3.6709	0.0376	1.5653	3.0359	0.0376	1.2877
S03F05	10x	B	11.1169	0.0513	1.7888	10.8147	0.0513	1.9691
		G	10.0803	0.0578	4.0987	7.9855	0.0578	4.2474
		R	4.5695	0.0474	2.2200	3.4803	0.0474	1.9186
S03F06	10x	B	10.6204	0.0488	1.2878	10.5011	0.0488	1.2730
		G	9.0782	0.0517	3.5389	7.8172	0.0517	3.3355
		R	3.9294	0.0403	1.7721	3.1985	0.0403	1.4669

Table 18. Table of the spatial mean relative radiance for each channel computed at different positions in the samples' bright zones.

C Spectroscopic measurements

A goniometric setup was utilized to measure the reflectance, transmittance and absorbance for two samples, as touched on previously in part I, section 2.6. The first was one of the original two-zone samples produced on a silicon chip, while the second one was the only good one-zone sample produced on a quartz chip. For the silicon chip naturally only a reflectance spectrum was obtained, and for the quartz chip both reflectance and transmittance spectra were obtained, which also enabled the calculation of absorbance.

The original idea was to collect spectra separately for the two different zones in the samples. This idea quickly proved to be unrealistic as the smallest possible beam spot size in the measurement setup was about 1 mm in diameter. The ideal case would be to have the beam spot lay completely on the material surface of interest in all angles. In the case of the samples on silicon substrates, only one of them has a sufficiently large area covered by the MWNT deposition. Even this sample doesn't have enough surface area for measuring the two zones separately and thus separate spectra were not possible to be measured. For the quartz substrate sample the obtained spectrum is basically only on the dim zone as there was no bright zone formed on it, as noted in part I, section 4.1.

C.1 Sample fabrication

A sample with the two visually distinct zones is difficult to reproduce if one desires large areas for the zones. This was attempted within this work on a quartz chip in order to maximize the measurable area for spectroscopic/goniometric measurements described below. Unfortunately, this did not fully succeed.

The hydrophility treatment for a chip surface is important in order to allow the water droplet to spread into a maximal area on the chip. As an alternative to the RIE treatment, a RCA solution treatment was also utilized, which makes the surface even more hydrophilic in comparison to the plasma treatment.

The amount of MWNTs for spin coating were varied by pipetting a different

number of droplets of the nanotube solution on the chip. A higher number of droplets and therefore a larger amount of MWNTs results in a more dense deposition but isn't by itself a deciding factor if a two-zone sample is successfully produced.

Another thing that was varied was how long the pipetted water droplet was allowed to dry. This was adjusted by controlling the humidity of an environmental chamber. Varied attempts were made by changing the duration of how long high humidity of over 90% was maintained and how sharply the humidity is dropped near to zero. The most successful attempts at reproducing the two-zone samples were made by leaving the samples for about an hour in high humidity and then dropping the humidity to almost zero by pumping nitrogen into the chamber. A quick drop in the humidity after a sufficient time in high humidity seems to be important for creating a sample without large gaps in the deposition. Otherwise the sample looks like it was torn apart.

Finally the amount of water pipetted on the chips was also varied, ranging from up to $1\mu\text{l}$ to $4\mu\text{l}$. The amount of water controls from how large of an area the water sweeps away the spin coated MWNT material, with the upper range taking over most of a chip's surface area. This seemingly doesn't affect anything else but the size of the final deposition.

All in all large two-zone samples proved to be very difficult to produce and there seems to be something, for the time being, unknown aspects of high importance in the production. For the goniometric measurements we managed to produce one good one-zone sample. The sample has only the dimmer zone but is quite uniform and doesn't have gaps in the deposition.

C.2 Measurement setup

The measurements were performed with a white light source, whose light was s-polarized. Each spectra was taken as a mean of 5 accumulations with a wavelength range of 320 – 800 nm. The spectrometer's front slit was set to 0.01 mm and the exposure time was set to 0.05 s.

The sample is attached between clamps, which are mounted to a stand that allows us to adjust its position in the plane of the table and rotate it about two axes, one normal to the table and the other normal to the first axis. With these degrees of freedom the sample can be properly positioned perpendicular to the incident light and the rotation axis of the sample can be positioned so that the area of interest

always remains in the incident light beam. The light from the light source is passed through optical fibers and through two irises in order to limit the beam spot size to the area of interest and achieve collimated lighting.

After the light source has stabilized, the initial spectrum is measured. After this the sample is placed in the light beam and the measurement of reflectance and transmittance spectra can begin. reflectance can be measured between 10 – 75 degrees, while transmittance can be measured between 0 – 75 degrees. The transmitted and reflected spectra are measured by rotating the input port of the spectroscope to the position at which its input is maximal. The maximal position is found by constantly measuring the spectrum and manually finding the position at which the spectrum is maximized. At the maximal position the spectrum can then be measured for that particular rotation of the sample. The sample can then be rotated to its next desired angle relative to the incident light and another set of spectra measured.

C.3 Measurements

The first measurement was performed on the silicon substrate sample. With this sample the silicon substrate considerably affects the measured spectra. The substrate's effect was mitigated by determining theoretical reflectance spectra for silicon for each measured angle and subtracting these theoretical spectra from the measured spectra. Additionally the final spectra were FFT filtered in a six data point range to smoothen the shape of the spectra. These spectra are shown in figure 65 and 66. The spectra that include the silicon substrate's effect are shown in figures 63 and 64.

First we can observe the overall shape of the reflectance spectra over the different measured angles in 65 and 66. The general shape of the spectrum appears to be quite consistent over the measured angles with a steep slope in the 400 – 500 nm range, followed by a more gentle slope in the 500 – 800 nm range. This means that the sample material's reflectance is considerably stronger in the green to red colors. In the spectra that include the substrate's effect, we can see that the reflectance is at its strongest in the 300 – 400 nm range. This corresponds to blue colors and matches with what we observe in most of the DF images captured using the optical microscope. This is expected as it matches closely with measured reflectance spectra of silicon in literature and theoretically determined spectra. The stronger reflectance in the green and red colors for the actual deposited material matches with the DF




Figure 63. Goniometric reflectance plots for the silicon substrate sample. reflectance was measured in the 20 – 75 degree range in steps of 5 degrees. This figure shows the plots in the range 20–45.

images captured at short exposure times.

The second measurement was performed on the quartz substrate sample. In this sample's case the quartz substrate is an excellent substrate choice for the goniometric measurement, as its interference with the measurements was minimal in comparison




Figure 64. Goniometric reflectance plots for the silicon substrate sample. reflectance was measured in the 20 – 75 degree range in steps of 5 degrees. This figure shows the plots in the range 50–75.

to the silicon substrate. The final spectra were similarly FFT filtered in a six data point range to smoothen the shape of the spectra.

The transmittance, reflectance and absorbance of the quartz substrate sample are shown in figures 67–69. The transmittance of the sample behaves in an expected




Figure 65. The same spectra as in 63 and 64, but additionally the spectra were normalized by their maximum values, followed by subtraction of theoretical spectra for silicon. This figure shows the plots in the range 20–45.

fashion; it is less than the transmittance of pure quartz and the transmittance decreases over increasing angles. This happens for pure quartz also, but here it decreases faster as at more oblique angles the area of the holes in the deposition are less visible to the incident light. Basically you can think of a circular hole appearing




Figure 66. The same spectra as in 63 and 64, but additionally the spectra were normalized by their maximum values, followed by subtraction of theoretical spectra for silicon. This figure shows the plots in the range 50–75.

elliptical as it is rotated. The transmittance also decreases slowly in the 0 – 50 degree range, followed with a faster decrease in the 55 – 75 degree range.

The reflectance of the quartz substrate sample starts at very low values and slowly increases in the 15 – 50 degree range. In the 55 – 70 degree range reflectance

increases considerably faster and also starts showing stronger and stronger preference for the upper half of the 300 – 800 nm wavelength range. This preference is very similar to that observed in figures 65 and 66 for the silicon substrate sample.

The quartz sample also has some absorbance in a few wavelength intervals. Determined by subtracting the transmittance and reflectance values from unity, as noted in part I, section 2.6. The lowest portion of the measured wavelength range isn't that reliable as the measured counts were in the order of few hundred counts at best, while counts for most of the wavelength range were in the order of few thousands.




Figure 67. Goniometric reflectance (blue), transmittance (red) and absorbance (green) plots for the quartz substrate sample. Transmittance was measured in the 0 – 75 degree range, while reflectance and therefore absorbance was measured in the 15 – 70 degree range. This figure shows the range 0–25.

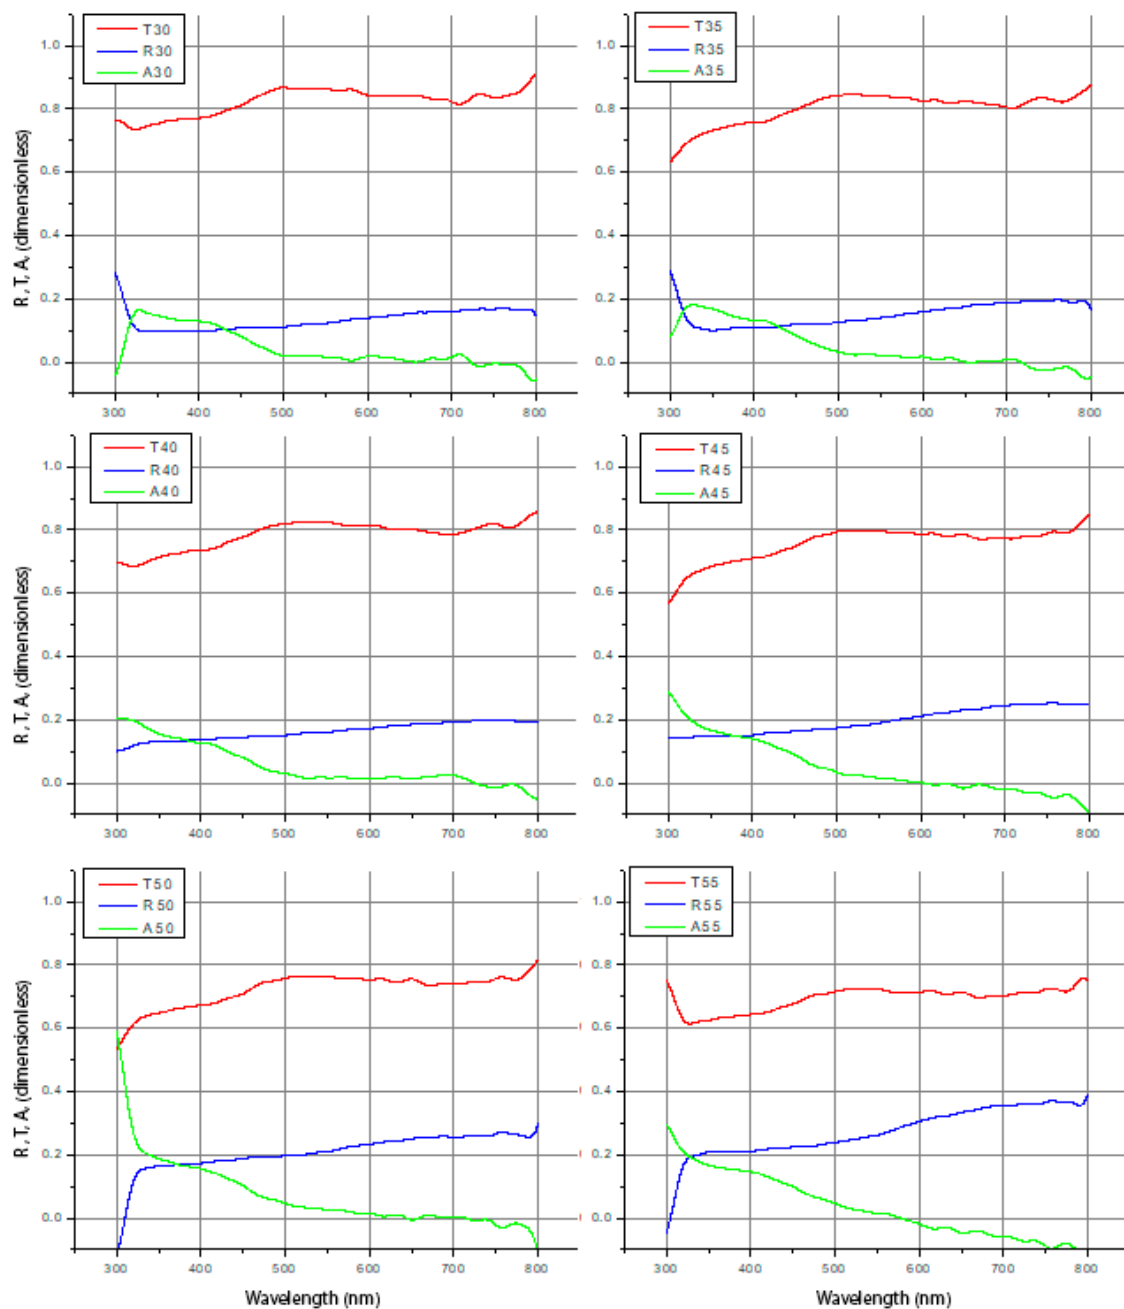


Figure 68. Goniometric reflectance (blue), transmittance (red) and absorbance (green) plots for the quartz substrate sample. Transmittance was measured in the 0 – 75 degree range, while reflectance and therefore absorbance was measured in the 15 – 70 degree range. This figure shows the range 30–55.




Figure 69. Goniometric reflectance (blue), transmittance (red) and absorbance (green) plots for the quartz substrate sample. Transmittance was measured in the 0 – 75 degree range, while reflectance and therefore absorbance was measured in the 15 – 70 degree range. This figure shows the range 65–75.