



Norwegian University of
Science and Technology

Modelling Underwater Hyperspectral Images

Tord Kriznik Sørensen

Master of Science in Physics and Mathematics

Submission date: July 2017

Supervisor: Dag Werner Breiby, IFY

Co-supervisor: Paul Anton Letnes, Ecotone

Norwegian University of Science and Technology
Department of Physics

Abstract

The underwater hyperspectral imager is a new type of imager able to acquire hyperspectral images under water. Common consumer cameras have poor spectral resolution and only captures light in three bands, in the red, green and blue parts of the light spectrum. Hyperspectral cameras, on the other hand, might capture hundreds of spectral bands in both the ultraviolet, visible and infrared regime. This property makes such cameras ideal for automatic mapping and object classification, which is important tools in fields such as environmental mapping.

In this thesis, we have taken the first steps towards a physical model for images acquired with an underwater hyperspectral imager. Light is heavily attenuated in water, and corrections need to be made if measurements from different distances and water bodies are to be compared. The proposed model consists of three sub-models: a model describing a lamp, a model for spectrally dependent attenuation in water and a model for reflectance. The lamp was modelled as a point source positioned behind the real lamp, emitting light in a Gaussian angular distribution. Attenuation in water was modelled using radiative transfer theory, and reflectance using the Oran-Nayar reflectance model. With the proposed model we were able to describe the collimation of the light well, and the assumed distribution as a reasonable first approximation. The model was also partly able to describe the change of shape of the light spectra with attenuation.

We have tested classification in underwater hyperspectral images to see the effect of uneven illumination and attenuation in water on the classification accuracy. It was also used as a tool for evaluating the proposed model. Uneven illumination and attenuation had a significant effect on classification accuracy, and hence need to be accounted for in real world applications. Two approaches were tested for removing these effects: normalisation and estimating reflectances. Normalisation gave a near perfect classification, while estimating reflectances using the proposed model performed worse than classification on the raw spectra. The estimated reflectances had however similar shapes, and it is thus mainly a problem with estimating the correct magnitudes for the illumination. The preprocessing method which seems to have the most potential when classifying different water bodies is to first estimate reflectance, and then normalise. We then circumvent the problem with illumination, and are able to classify between water bodies as long as the attenuation is known.

The proposed model performs within the expectations of a first attempt at modelling images acquired with an underwater hyperspectral imager. The developed simulation- and analysis tools will be useful in further research and field measurements.

Sammendrag

Undervanns hyperspektrale kamera er ny type kamera som tar hyperspektrale bilder under vann. Vanlige kameraer har lav spektral oppløsning og kan kun fange lys i tre spektrale bånd i den røde-, blå- og grønne delen av lysspekteret. Hyperspektrale bilder derimot har hundrevis av spektrale bånd som dekker både den ultrafiolette-, synlige- og infrarøde delen av lysspekteret. Dette gjør at slike kameraer er ideelle for automatisk kartlegning og klassifikasjon av objekter, som er viktig i applikasjoner som miljøovervåkning.

I denne oppgaven har vi tatt de første stegene for å utvikle en fysisk modell for bilder tatt med et undervanns hyperspektralt kamera. Lys blir kraftig dempet når det propagerer i vann. For å kunne sammenligne og gjøre analyser av bilder tatt i forskjellige høyder, og i forskjellige vann, må denne dempningen korrigeres for. Den foreslåtte modellen består av tre deler: en modell for en lampe, en model for spektralt avhengig dempning i vann og en modell for reflektans. Lampen ble modellert som en punktkilde plassert bak den virkelige lampen. Det ble antatt at den sprer lys i en Gaussisk form. Dempning i vann ble modellert ved hjelp av ligninger for lystransport og reflektans ved Oran-Nayar modellen for reflektans. Med den foreslåtte modellen klarte vi å beskrive lysets kollimering og den antatte lysdistribusjonen var et godt første estimat. Den korrigerte også delvis for endringer i de målte spektraenes formendring.

Vi har testet klassifikasjon i undervanns hyperspektrale bilder for å undersøke effekten av ujevn belysning og dempning i vann på klassifikasjonsstyrken. Det ble også brukt som en test av den foreslåtte modellen. Ujevn belysning of dempning fra vann hadde stor effekt på klassifikasjonen og trengs derfor å bli tatt hensyn til. To fremgangsmåter ble brukt for å korrigere vekk disse effektene: normalisering og estimering av reflektanser. Normalisering gav en nesten perfekt klassifikasjon, mens reflektanseestimering gav dårligere klassifisering enn direkte klassifikasjon på de målte spektrene. De estimerte reflektansene hadde i stor grad riktig form, og det er derfor hovedsakelig et problem med den estimerte størrelsen av belysningen. Den preprosseseringsmetoden som virker som den har størst potensial er å først estimere reflektans, for så å normalisere. På den måten unngår vi problemet med feil i korrigeringen av belysningen. Det gjør oss også istand til å klassifisere på tvers av vanntyper og avstander, så lenge vi kjenner til hvordan lys dempes i vann.

Den foreslåtte modellen fungerer innenfor forventningene til et første forsøk på en modell. Simulerings- og analyseverktøyene som er utviklet iløpet av denne mastergraden vil være til hjelp i videre forskning og feltmålinger.

Acknowledgements

First and foremost I would like to thank the people at Ecotone for welcoming me and allowing me to use their equipment. Especially I would like to thank Paul Anton Letnes and Lars Martin Aas who have been my co-supervisors. I really appreciate you giving me an opportunity to write about this topic even though I knew nothing about machine learning, hyperspectral imaging or how light behaves in water before I started. I have learned a lot and I look forward to be working with you.

It has been challenging designing and building experimental setups, and I have got a newfound respect for the time it takes to plan and perform experiments. To the hardware guys Paul Sverre and Håvard, thank you for letting me use your workshop and helping me with the hardware. Not even once did I get a no for an answer. This also applies for the software guru Magne, thank you for helping me with setting up the recordings and teaching me about how to communicate with the imager.

To my supervisors at NTNU, Dag Werner Breiby and Bjørn Kåre Alsberg, thank you for discussions and for honest and thorough feedback.

I would like to thank Trondheim Biological Station (TBS) for letting me use their water tank, lab, and spectrometer for some of the measurements. I would also like to thank Bernard Schartmueller for helping me with the logistics and preparing the equipment at TBS for measurements. That was no doubt a two man job. Unfortunately the tank measurements from TBS could not be used in the end, but a lot was learned from these initial tests. A special thanks goes to Sturla Haltbakk which was a great help during my time there.

Lastly I would like to thank Marmine for letting me use their underwater hyperspectral imager.

Contents

1	Introduction	1
2	Theory	5
2.1	Hyperspectral Imaging	5
2.2	Optical Properties of Water	8
2.3	Reflectance	12
2.4	Principal Component Analysis	14
2.5	Partial Least Squares regression	17
3	Modelling and experimental work	21
3.1	Proposed model for underwater hyperspectral images	21
3.2	Experiments and determination of model parameters	25
4	Results and Discussion	31
4.1	Explaining variances using Principal Component Analysis	31
4.2	Reflectance of test objects	32
4.3	Determining lamp properties	33
4.4	Tank experiment and comparison with simulations	36
4.5	Classification	44
4.6	Note on earlier work	47
4.7	Further work	48
5	Conclusion	51
	Bibliography	52
	Appendix	57
A	Code documentation	57

Introduction

Hyperspectral imaging is a method for taking images with large amounts of spectral information. In contrast to a conventional camera, which has poor spectral resolution and only measured three spectral bands, it captures a nearly continuous spectrum of light in each pixel. Thus a lot more information can be collected from such an image than from images taken by conventional cameras. It is sometimes referred to as "chemical imaging" as it provides information on chemical compositions of the imaged object. Underwater hyperspectral imagers (UHI) is a relatively new type of imager, and still under development. It is essentially a hyperspectral imager inside a watertight container able to travel to great depths.

The UHI is patented and produced by Ecotone [1], which is a spinoff company from the Norwegian University of Science and Technology (NTNU). Following the startup of Ecotone, a UHI was deployed in April 2010 for mapping objects of interest on the seafloor at Hopavågen in Norway [2]. Later, the same prototype was used by scuba divers at the Great Barrier Reef and Shark Bay for the mapping of corals. The first time a UHI was deployed on a remotely operated vehicle (ROV) under water was in April 2012. Its mission was to look at cold water corals in Trondheimsfjorden. Images were taken at sixty to eighty meters depth, and a second survey was later performed down to a depth of four hundred meters. During the last five years, surveys using UHI have been used for bio-geo-chemical mapping of the seafloor in several waters to a maximum depth of 4200 meters [3]. In August 2016, a UHI was used on a mission in collaboration with the Marmine project [4] to test if the technology could be used for mapping deep sea minerals. Lately, there has been research into the applications of UHIs by Trondheim Biological Station (TBS). Ingvild Andersson [5] has looked at the bio-optical diversity in cold water coral habitats, and Aksel Mogstad [6] investigated the spectral characteristics of coralline algae.

Oil and gas companies are required to inspect and report the environmental

impact of their activities. Traditional monitoring methods, used offshore since the 1980s, involve sediment sampling using remotely deployed grabs or coring devices. Sediments are analysed for their physical composition and chemical content as well as biological analysis.

During the past two decades, underwater technology has become more widely available. Subsea habitats are frequently being surveyed using visual methods, typically image and/or video recordings deployed via ROVs. Generally, a biologist will sit in the control room with the ROV operators and register observations directly into a customised database. Typical observations include organisms observed, density of sponges, general seafloor conditions and extent of drill cuttings deposition. For post-drilling surveys, the extent and intensity of drill cuttings are recorded.

The Norwegian Environmental Agency generally requires visual surveys in advance of exploratory drilling in areas regarded as potentially sensitive, such as sponge grounds or coral habitats. A major concern is how reliable visual assessments of sensitive habitats and drill cuttings deposition are. Quantification of sponge density is notoriously subjective, and it is often difficult to pinpoint the border between the deposition zone.

From 2014 to 2016 a project under the Petromaks2 programme [7, 8] explored whether a UHI could be used for environmental mapping of the seabed. Their objectives were to compare UHI recordings and traditional methodology using video image for offshore habitat mapping and monitoring, and to assess impacts of drilling activities on sea floor habitats. They concluded that UHI measurements have the potential to become a fast and cost effective method for environmental mapping. Some benefits compared to the traditional methodology are: automated recordings, automatic and objective classification and monitoring at greater depths than currently possible. They further conclude that there are some challenges which need to be handled before UHIs can be used reliably.

They observed that both the magnitude and shape of the measured light spectra changed considerably with distance under water. These spectral changes worsened the prediction accuracy of classifications. Hence, it is important to acquire a model describing light's propagation in water, and derive an invariant quantity from the measured data, which may be compared across different water bodies.

The reflectance of an object is such an invariant property. It is the amount of light reflected off a surface, and acts like a spectral fingerprint. There are primarily two methods which are used to acquire reflectance from the ocean floor [2]. The first is using a plaque of known reflectance situated within the frame of the hyperspectral image. The sensed reflectance from each pixel in the image can then be estimated by comparison with the measured spectra from the plaque. Such a method is cumbersome when doing remote sensing of the seabed, due to the need

of placing an object on the seabed. The second method is to use radiative transfer theory for how light propagates in water. Measurements of the constituents in the water must then be made simultaneously, which can be done by a device attached to the camera. This information can then be used to model the attenuation of the light beam and estimate the reflectance.

Due to UHIs still being under development, there have been few studies on how to model systems containing them. In previous project work at NTNU, we proposed a simple model for how attenuation in water could be described [9]. It proved to compensate for attenuation effectively, but was from a modelling perspective not complete. It did neither explain the effect of the divergence of light in space, nor the effect of uneven illumination. In the current study, the main aim was to include these factors into the model, and to test classification using estimated reflectances. The purpose of these models is to interpret data by correction for the effects of water. We have also briefly tested Principal Component Analysis for explaining variation in the hyperspectral images. The project work and this master thesis have been the first steps towards modelling systems containing UHIs.

To assess the correctness of the proposed model, experiments were performed with a UHI measuring a scene illuminated by a single lamp. The scene contained four Lego blocks resting on a grey plate. Reflectances of the five components in the scene were measured, as well as the lamp's spectra and emitted distribution. A framework for analysing hyperspectral images in Python was developed, as well as a 3D ray tracing engine for simulations.

Theory

This chapter introduces hyperspectral imaging, optical properties of water including radiometry, and derives the equation governing radiative transfer. Principal component analysis and partial least squares regression is also covered. The section on optical properties in water is mostly based on the theory described in Mobley's excellent book *Light and Water: Radiative Transfer in Natural water* [10].

2.1 Hyperspectral Imaging

Digital images are built up of pixels. Each pixel in a conventional digital image contains the *irradiance* measurement of three spectral *bands* in the red, green and blue regions of the light spectrum. We call such cameras RGB cameras, and the images they acquire RGB images. A band is defined as an extent in the spectral domain centered around a wavelength, and irradiance describes how much incident energy is measured per area and time. In other words, a pixel can be looked upon as a sensor measuring the irradiance emitted, or reflected, of a part of 3D *scene* in the real world. The entire digital image is a 2D representation of that scene.

RGB images are made such that they measure bands which are similar to those human eyes can detect. They are able to describe everything humans see, but there is more information to be gathered from the incoming light by increasing spectral resolution. The resolution can be increased by increasing the number of bands measured and/or making them narrower. Cameras that measure more than three bands are grouped into two classes Multispectral Imagers (MI) and Hyperspectral Imagers (HI). HIs can have hundreds of narrow bands while MIs usually have ten or less broad bands. Since MIs are not used in this thesis, they will not be discussed further. It is more useful to look into the workings of an HI, which is further classified by two aspects [11]: the method of how they acquire spectral, information and the method in which they acquire spatial information.

Methods for acquiring the spatial information include *whisk-broom*, *push-broom* and *framing*. Whisk-broom is a point scanning technique which has a field of view (FOV) of one pixel. To capture the entire scene a camera needs to scan in both the horizontal and vertical direction. A push-broom camera is similar, but scans the entire length in e.g. the horizontal direction in one exposure, and line by line in the vertical. A framing HI uses a 2D FOV that stays stationary while the image is acquired. If a scene is bigger than the FOV, multiple images can be taken to capture the entire scene.

Spectral information can be acquired through *filtering*, *dispersive* or *interferometric* techniques. The filtering HIs take several images with tunable filters that let different bands through for each exposure, and the dispersive variety uses gratings or prisms to spread light of different wavelengths onto a detector. The interferometric methods are based on Fourier-transform spectrometers which consist of two interferometers.

In this thesis, a push-broom HI with dispersive acquisition of spectral information has been used. Figure 2.1a shows the most essential components of such a camera. An HI is moving in the y -direction acquiring an image of a scene. In the figure, the scene is represented by a dark grey rectangle. An objective lens L_1 captures the scene and passes it to a slit S limiting the field of view to the bright grey line of width Δy . The light beam is then passed further to a lens L_2 , collimating the beam before it hit a dispersive component P separating light of different wavelengths. Finally, a lens L_3 collects each wavelength and projects them to specific locations on an array detector. Each line is added to the image as a slice in the x - and λ -plane. The result is a 3D datacube, as illustrated in Figure 2.1b, where each pixel contains a spectrum of the irradiance E .

The primary goal of hyperspectral imaging is to be able to do *classification* and *clustering*. Classification is the operation of finding objects in a dataset matching preknown patterns, and clustering automatically finds objects of similarity without any prior knowledge. These operations are possible since different materials reflect different parts of the light spectrum. The amount of reflected light we call reflectance, and this property is often unique to a given material. Therefore the reflectance is also known as a material's spectral fingerprint. HIs have been used in medical science to diagnose diseases [12], and in food science for quality assessment [13]. Remote platforms, like planes, drones and satellites, have been deployed with HIs for vegetation mapping [14] and toxic cyanobacteria bloom monitoring [15]. Seabed mapping have also been performed [16, 17], but such monitoring is limited to shallow waters for airborne platforms due to light being heavily attenuated in water. If HIs are to be used when monitoring the deep ocean seabed, they need to be underwater close to the seafloor.

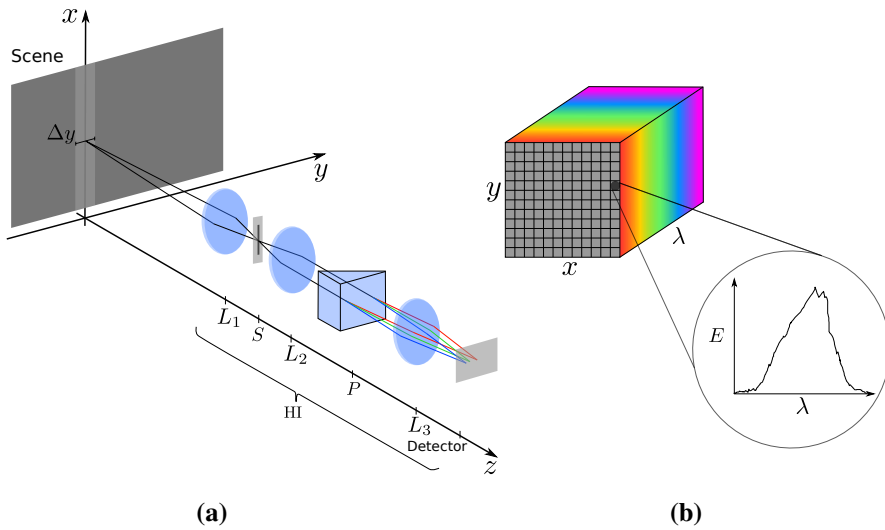


Figure 2.1: (a) Schematics of a push-broom HI with dispersive acquisition. Light is reflected from a scene, and enters the HI through an objective lens (L_1), then it is passed through a slit (S) and through a collimating lens (L_2). The light is then dispersed by a dispersive component (P), and collected through a lens (L_3) onto an array detector. The imager is moving in the positive y -direction. (b) Visualization of a 3D datacube. x and y represent spatial dimensions and λ the spectral.

2.2 Optical Properties of Water

When light travels through a medium, it can either be scattered, absorbed or transmitted. Consider a column of water with volume ΔV and thickness Δr which is illuminated by collimated light with power $\Phi_i(\lambda)$, as illustrated in Figure 2.2. Here λ denotes wavelength. After entering the water column, an amount $\Phi_a(\lambda)$ of the light is absorbed and another amount $\Phi_s(\lambda)$ is scattered. The remaining power, $\Phi_t(\lambda)$, is transmitted. If no photons undergo a change in wavelength during the scattering process $\Phi_i(\lambda) = \Phi_a(\lambda) + \Phi_s(\lambda) + \Phi_t(\lambda)$ due to energy conservation.

Several fundamental quantities can be defined by the above definitions. The absorbance A , which is the fraction of incident power absorbed within ΔV , is defined as:

$$A(\lambda) = \frac{\Phi_a(\lambda)}{\Phi_i(\lambda)}, \quad (2.1)$$

the scatterance $B(\lambda)$ is the fraction of scattered power:

$$B(\lambda) = \frac{\Phi_s(\lambda)}{\Phi_i(\lambda)}, \quad (2.2)$$

and transmittance $T(\lambda)$ is the fraction of transmitted power:

$$T(\lambda) = \frac{\Phi_t(\lambda)}{\Phi_i(\lambda)}. \quad (2.3)$$

Note that these quantities are wavelength dependent. The absorbance $A(\lambda)$ should not be confused with the absorbance $D(\lambda)$, which is defined as

$$D(\lambda) = \log \frac{\Phi_i(\lambda)}{\Phi_s(\lambda) + \Phi_t(\lambda)}. \quad (2.4)$$

These ratios can further be used to find the absorption coefficient $a(\lambda)$ and the scattering coefficient $b(\lambda)$. They describe the absorbance and scattering per unit distance and, with regards to Figure 2.2, are defined as

$$a(\lambda) = \lim_{\Delta r \rightarrow 0} \frac{A(\lambda)}{\Delta r}, \quad (2.5)$$

$$b(\lambda) = \lim_{\Delta r \rightarrow 0} \frac{B(\lambda)}{\Delta r}. \quad (2.6)$$

For convenience, define their sum as the attenuation coefficient $c(\lambda)$,

$$c(\lambda) = a(\lambda) + b(\lambda), \quad (2.7)$$




Figure 2.2: Light with power Φ_i enters a water column of volume ΔV and width Δr . The power is reduced due to scattering Φ_s and absorption Φ_a . The remaining power Φ_t is transmitted through.

which describes the total attenuation of light through water.

In the above, it is assumed that no inelastic scattering occurs. Inelastic scattering, also known as transpectral scattering, is when a photon of wavelength λ_1 is scattered into another wavelength λ_2 . This can happen if a molecule absorbs a photon and at a later stage emits a photon with a different energy than before. The remaining absorbed energy can either be emitted again as another photon or converted to other energy forms. In natural waters transpectral scattering may occur due to interaction with dissolved matter, or due to Raman and Brillouin scattering from the water molecules themselves

2.2.1 Radiometry and Radiative Transfer

Radiometry is the science of measuring the distribution of radiated power in space, and radiative transfer is a theory describing the evolution of radiant energy. This section starts by introducing the concepts of radiance and irradiance, then the radiative transfer equation is derived phenomenologically .

Radiance and Irradiance

The fundamental radiometric quantity in water optics is radiance. Imagine an amount of radiant energy ΔJ incident in a time interval Δt onto a surface ΔA . Let the surface be located at \mathbf{x} and that the energy arrives through a solid angle $\Delta\Omega$ centered at some direction $\boldsymbol{\xi} = (\theta, \phi)$. Also assume that the photon wavelengths are within an interval of length $\Delta\lambda$ centered at λ . An illustration of the scenario is depicted in Figure 2.3a, where a device for measuring radiance is illustrated. If we let the aforementioned Δ -quantities approaches the infinitesimal limit, the radiance can be defined as

$$L(\mathbf{x}, t, \boldsymbol{\xi}, \lambda) = \frac{\partial^4 J}{\partial t \partial A \partial \Omega \partial \lambda}. \quad (2.8)$$

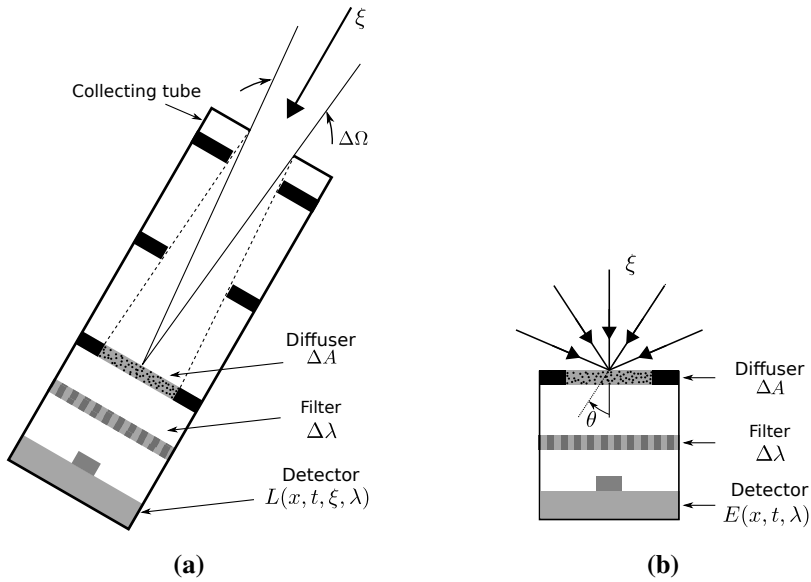


Figure 2.3: Illustration of instruments used to define radiance (a) and irradiance (b). ξ denotes the directions of the incoming light and $\Delta\Omega$ the solid angle which is let through to a detector with area ΔA . Before hitting the detector the light is divided into bands of width $\Delta\lambda$. x is the position of the detector and θ is the angle in which the light hits the detector measuring irradiance. L is the measured radiance and E is the measured irradiance.

L gives information about the entire light field, but in applications such complete information is difficult to acquire and rarely needed. Instead, it is more common to measure irradiances.

By removing the collection tube from Figure 2.3a, we have the situation in Figure 2.3b, and light from all directions can hit ΔA . The irradiance E is then, in the infinitesimal limit, defined as

$$E(\mathbf{x}, t, \lambda) = \frac{\partial^3 J}{\partial t \partial A \partial \lambda}, \quad (2.9)$$

and can be related to the radiance by integrating over all incoming directions in the upper hemisphere Ξ ,

$$E(\mathbf{x}, t, \lambda) = \iint_{\xi \in \Xi} L(\mathbf{x}, t, \xi, \lambda) |\cos \theta| d\Omega. \quad (2.10)$$

The cosine factor is added because a light beam incident on a surface ΔA with an angle θ effectively sees an area of size $\cos \theta \Delta A$.

Radiative Transfer Equations

In the literature, there are several derivations of the radiative transfer equation (RTE) describing radiance. Stamnes [18] derives it from a general Boltzmann equation, Fante [19] from Maxwell's equations and Measures [20] from quantum mechanical descriptors of absorption and scattering. In the following, a more phenomenological derivation is performed, as done by Mobley.

Four processes are necessary and sufficient for writing down an energy balance equation for radiative transfer:

1. Loss and gain of photons by elastic scattering.
2. Loss of photons from absorption.
3. Loss and gain of photons by inelastic scattering.
4. Gain of photons by conversion of non-radiant energy to radiant.

Note that loss and gain do in this context not mean that the energy is lost or created, but rather that it is converted to other forms or scattered from or into the lightbeam. The first three processes have already been discussed in various detail earlier in this chapter. The fourth process can, for example, occur by bioluminescence or other chemical reactions.

Consider a beam of radiant energy that propagates a distance Δx through a volume of water ΔV . If the distance Δx is small, it is reasonable to assume that the change in radiance over Δx due to absorption is proportional to the incident radiance:

$$\frac{\Delta L(\mathbf{x}, \boldsymbol{\xi}, \lambda)}{\Delta x} = -a(\lambda)L(\mathbf{x}, \boldsymbol{\xi}, \lambda). \quad (2.11)$$

In Equation (2.11), $a(\lambda)$ is the absorption coefficient and ΔL is the change in radiance over the distance Δx . Similarly, the loss due to elastic scattering can be written as

$$\frac{\Delta L(\mathbf{x}, \boldsymbol{\xi}, \lambda)}{\Delta x} = -b(\lambda)L(\mathbf{x}, \boldsymbol{\xi}, \lambda), \quad (2.12)$$

where $b(\lambda)$ is the scattering coefficient.

Assume that an incident radiance, with solid angle $\Delta\Omega'$ centered on direction $\boldsymbol{\xi}' \neq \boldsymbol{\xi}$, crosses the light beam. The change of radiance in direction $\boldsymbol{\xi}$ from elastic scattering from light outside the ΔV can then be formulated as

$$\frac{\Delta L(\mathbf{x}, \boldsymbol{\xi}, \lambda)}{\Delta x} = L(\mathbf{x}, \boldsymbol{\xi}', \lambda)\beta(\boldsymbol{\xi}' \rightarrow \boldsymbol{\xi})\Delta\Omega', \quad (2.13)$$

where $\beta(\xi' \rightarrow \xi)$ is the fraction of radiance scattered from direction ξ' to ξ and is called the angular scatterance. To get the full contribution we integrate over all directions:

$$\frac{\Delta L(\mathbf{x}, \xi, \lambda)}{\Delta x} = \iint_{\xi' \in \Xi} L(\mathbf{x}, \xi', \lambda) \beta(\xi' \rightarrow \xi) d\Omega'. \quad (2.14)$$

The two last processes, loss and gain in energy from inelastic scattering and gain from sources, need to be handled for each specific case. The change in radiance will here be modelled with a generic source term $S(\mathbf{x}, \xi, \lambda)$:

$$\frac{\Delta L(\mathbf{x}, \xi, \lambda)}{\Delta x} = S(\mathbf{x}, \xi, \lambda). \quad (2.15)$$

Let Δx become infinitesimal. Then, by the definition of the derivative

$$\frac{\Delta L(\mathbf{x}, \xi, \lambda)}{\Delta x} \rightarrow \frac{dL(\mathbf{x}, \xi, \lambda)}{dx}$$

By adding all the terms the standard form of the Radiative Transfer Equation (RTE) is found:

$$\frac{dL(\mathbf{x}, \xi, \lambda)}{dx} = -c(\lambda)L(\mathbf{x}, \xi, \lambda) + \iint_{\xi' \in \Xi} L(\mathbf{x}, \xi', \lambda) \beta(\xi' \rightarrow \xi) d\Omega + S(\mathbf{x}, \xi, \lambda). \quad (2.16)$$

As before, $c(\lambda)$ is the attenuation coefficient. Although we now have a model for the radiant transfer, there are shortcomings. The RTE presented here has been developed for unpolarised light. Scattering can however induce polarisation, even if the original beam is unpolarised. Nevertheless, the RTE derived give accurate results for many oceanographic applications. Three arguments may explain this fact. First, the particles responsible for scattering in the ocean are usually much larger than the wavelength of light used in underwater remote sensing. Polarisation by scattering is often larger for particles with size smaller than the incoming wavelength. Second, multiple scattering can be significant in water. These multiple scattering events lead to depolarisation of the light. Third, most often irradiances are measured. Polarisation for different directions then tends to average out.

2.3 Reflectance

Reflectance is defined as the amount of light reflected from a surface. A surface with surface normal N is depicted in Figure 2.4. Light is coming from direction V onto the surface, and a measurement is made of the reflected light in direction




Figure 2.4: Definition of the coordinate system used for describing reflectance. Light is incident on a surface from direction \mathbf{V} with angle θ_i to the surface normal \mathbf{N} , and ϕ_i to the x -axis. It then reflects from the surface and is measured by a detector, which is located in direction \mathbf{W} at an angle θ_r to \mathbf{N} , and ϕ_r to the x -axis. The beige area illustrates the area which is imaged by the detector.

\mathbf{V} . \mathbf{V} and \mathbf{W} are defined by the polar angles (θ_i, ϕ_i) and (θ_r, ϕ_r) respectively. θ_i and θ_r denote the zenith angles and ϕ_i and ϕ_r the azimuth angles. The reflectance R can then be described by the irradiance E_i hitting the surface and the measured radiance L_r from direction \mathbf{W}

$$R(\theta_i, \phi_i, \theta_r, \phi_r, \lambda) = \frac{L_r(\theta_i, \phi_i, \theta_r, \phi_r, \lambda)}{E_i(\theta_i, \phi_i, \lambda)}. \quad (2.17)$$

Phenomenologically reflectance can be divided into two categories: specular reflection and diffuse reflection [21]. Specular reflection occurs for glossy mirror like surfaces where the outgoing direction \mathbf{W} is the mirrored direction of \mathbf{V} about the surface normal \mathbf{N} such that $(\theta_r, \phi_r) = (\theta_i, \phi_i + \pi)$. Light is not being reflected into any other directions. Diffuse reflection happens when light hits a rough or granular surface. The light will then reflect into multiple directions due to the microscopic irregularities. An ideal diffuse reflecting surface will reflect irradiance equally in all directions, but its size depends on the incident angle of the light. Such a surface is called a *lambertian surface* [22]. In real world applications there are no pure specular- or diffuse surfaces, however some materials show primarily either specular or diffuse behaviour,

Reflectance can be measured directly using a gonireflectometer [23] which both illuminate a surface, and measure the radiance from it in several directions. Such a measurement is time consuming because for each illuminated direction the entire hemisphere around the surface needs to be measured. Other approaches have been investigated, but they often include expensive and complex equipment with mirrors, projectors etc. An option to avoid such equipment and simplify measurements is to assume that the reflectance follows some parametric model.




Figure 2.5: Illustration of V-cavities used in the Oran-Nayar reflectance model.

According to Mobley [24] it is usual in oceanographic applications to assume lambertian reflectance. The reflected light can then be described by

$$L_r = R_0 E_i \cos \theta_i. \quad (2.18)$$

Here R_0 is the reflectance measured at normal incidence to the surface and E_i is the incoming irradiance. The cosine factor is added since an incoming irradiance at angle θ_i to an area A only sees an area of $A \cos(\theta_i)$. A generalization of the lambertian model was made by Oren and Nayar [25]. They assumed that the irregularities in the surface could be represented by microscopic symmetric V-cavities, as seen in Figure 2.5, and that the walls in those cavities reflected as lambertian surfaces. To represent the random nature of granularities on the surface, they further assumed that the area of each wall followed a gaussian distribution. The Oran-Nayar model for the reflected irradiance can then be written as

$$L_r = E_i R_0 \cos \theta_i [C_1 + C_2 \max[0, \cos(\phi_i - \phi_r)]] \sin \alpha \tan \beta] \equiv E_i R, \quad (2.19)$$

where R has been implicitly defined in the last equality. The parameters C_1 , C_2 , α and β are defined as

$$\begin{aligned} C_1 &= 1 - 0.5 \frac{\kappa^2}{\kappa^2 + 0.57}, \\ C_2 &= 0.45 \frac{\kappa^2}{\kappa^2 + 0.09}, \\ \alpha &= \max(\theta_i, \theta_r), \\ \beta &= \min(\theta_i, \theta_r), \end{aligned}$$

where κ determines the roughness of the surface. Note that when $\kappa = 0$, Equation (2.19) equals Equation (2.18).

2.4 Principal Component Analysis

Principal Component Analysis (PCA) is a method that decomposes a dataset into new variables called latent variables. These latent variables are linear combinations of the original variables and span an orthogonal subspace of the original

variable space. This subspace may thus represent the dataset even though it is of a lower dimension than the original variable space. This is possible because the latent variables are constructed in such a way that they point in the direction of maximum variance [26]. Let us call the space spanned by the latent variables for the latent variable space (*LV-space*). PCA is a dimensionality reduction method since it reduces a dataset of many variables to only a few describing the most variation. Each latent variable describe a part of the dataset called a principal component (PC).

For explaining how PCA works let $X \in \mathbb{R}^{m \times n}$ denote a dataset which consist of m rows and n columns. Let each row be a sample from a measurement and let each column describe a variable, or feature, in that measurement. In a UHI measurement a pixel would correspond to a sample and the wavelengths to variables. Assume the dataset is columnwise mean centered. The dataset can then be written in the form

$$X = TP^T + E = \sum_{i=1}^d \mathbf{t}_i \mathbf{p}_i^T + E. \quad (2.20)$$

$T \in \mathbb{R}^{m \times d}$ is called the score matrix and $P \in \mathbb{R}^{n \times d}$ the loading matrix. \mathbf{t}_i and \mathbf{p}_i are the i th columns of T and P . They are called the scores and loadings of PC i , which is defined as the outer product $PC_i = \mathbf{t}_i \mathbf{p}_i^T \in \mathbb{R}^{m \times n}$. d is the number of dimensions in the *LV-space* and is equal to the number of PCs chosen. This choice will be handled later. $E \in \mathbb{R}^{m \times n}$ is the residual which contains the unstructured information and noise inside X , while T and P contains the structure. The columns in P are of unit length and denote the directions of highest variance, and are hence the unit directions in the *LV-space*. T denotes the projection of X onto the *LV-space* and its rows represents the sample coordinates in that coordinate system.

According to Höskuldsson [27], the latent variable of the i th PC can be derived by finding the \mathbf{p}_i which gives maximum variance in z_i subject to $\mathbf{p}_i \mathbf{p}_i^T = 1$,

$$\begin{aligned} \max_{\mathbf{p}_i \mathbf{p}_i^T = 1} [\text{var}(z_i)] &= \max_{\mathbf{p}_i \mathbf{p}_i^T = 1} [\mathbf{z}_i^T \mathbf{z}_i] \\ &= \max_{\mathbf{p}_i \mathbf{p}_i^T = 1} [\mathbf{p}_i^T X^T X \mathbf{p}_i], \end{aligned} \quad (2.21)$$

This can be done by using Lagrange multipliers. Define the Lagrange function L with a Lagrange multiplier μ as

$$L(p, \mu) = \left(\mathbf{p}_i^T X^T X \mathbf{p}_i \right) - \mu \left(\mathbf{p}_i \mathbf{p}_i^T - 1 \right), \quad (2.22)$$

and solve the optimality conditions

$$\begin{aligned}\frac{\partial L}{\partial \mathbf{p}_i} &= X^T X \mathbf{p}_i - \mu \mathbf{p}_i = 0, \\ \frac{\partial L}{\partial \mu} &= \mathbf{p}_i \mathbf{p}_i^T - 1 = 0,\end{aligned}\tag{2.23}$$

which then reduces to an eigenvalue equation

$$X^T X \mathbf{p}_i = \mu \mathbf{p}_i.\tag{2.24}$$

From statistics, $X^T X \in \mathbb{R}^{m \times m}$ is the covariance matrix of X . Thus, the loadings are the eigenvectors of the covariance matrix. Let the first loading vector, \mathbf{p}_1 , be the eigenvector corresponding to the highest eigenvalue and \mathbf{p}_2 correspond to the next highest and so on. This makes the loading vectors represent the directions of most variance in decreasing order [28].

Variables with large magnitudes tend of influencing the PCA model more than variables with smaller magnitudes [26]. Standardization can be used to ensure that each variable have an equal opportunity to influence the model irrespective of magnitude. Standardization is the combined operation of columnwise mean centering and scaling each column by its standard deviation. Given a dataset X , the standardized dataset \tilde{X} is given as

$$\tilde{X} = (X - \bar{X})S.\tag{2.25}$$

\bar{X} is a matrix whose columns contain the column means, and S is a diagonal scaling matrix with the reciprocals of the column standard deviation on the diagonal.

Choosing the number of PCs

In order to choose the number of PCs needed to describe a dataset, the variance explained by each component can be estimated. This estimate can be found by looking at the ratio of the eigenvalue corresponding to the current PC to the sum of all eigenvalues. Denote the explained variance of PC i by f_i , and the i th eigenvalue of $X^T X$ by μ_i . Then the explained variance is given by [28]

$$f_i = \frac{\mu_i}{\sum_{j=1}^m \mu_j}.\tag{2.26}$$

The variance explained by the first k components, F_k , is then

$$F_k = \sum_{i=1}^k f_i.\tag{2.27}$$

The number of required components depends on both complexity and nonlinearity of the latent structures. A sufficient number of PCs should be included to account for most of the structure in the data.

Figure 2.6 shows an example of an explained variance plot. The explained variance from each component is plotted against the component number. In this example one would choose to keep the two or three first PCs since they explain the majority of the variance.




Figure 2.6: An example of an explained variance plot, f_i denotes the explained variance of PC i .

2.5 Partial Least Squares regression

While PCA is a method to investigate the variance within independent x -variables, the Partial Least Squares Regression (PLS) investigates the connection between independent x -variables and dependent y -variables. PLS builds a regression model between a multivariate sample matrix X and a response matrix Y based on latent structures within both datasets. This is done in such a way as to maximize the covariance between the latent variables [27]. Like PCA it is a dimensionality reduction technique, and hence one needs to choose how many dimensions to reduce to.

Assume a mean centered dataset $X \in \mathbb{R}^{m \times n}$ and a response matrix $Y \in \mathbb{R}^{m \times k}$. X and Y are matrices with m samples, and n and k features respectively. The PLS model can then be written as [29]

$$X = TP^T + E_x, \quad (2.28)$$

$$Y = ZQ^T + E_y, \quad (2.29)$$

$$T = XW(P^TW)^{-1} \equiv XW'. \quad (2.30)$$

$T \in \mathbb{R}^{m \times d}$ and $Z \in \mathbb{R}^{m \times d}$ are called the x - and y - score matrices, $P \in \mathbb{R}^{d \times n}$ and $Q \in \mathbb{R}^{k \times d}$ the x - and y -loading matrices, and $E_x \in \mathbb{R}^{m \times n}$ and $E_y \in \mathbb{R}^{m \times k}$ the x - and y -residuals. $W \in \mathbb{R}^{n \times d}$ are the PLS weights, and W' is the transformation matrix from X to T . d is the chosen number of PLS components, and thus also the number of dimensions in the latent variable spaces.

Analogously to PCA, PLS has an intuitive geometric interpretation. Denote the i th column vectors in the loading matrices for \mathbf{p}_i and \mathbf{q}_i and the i th score column vectors for \mathbf{t}_i and \mathbf{z}_i . We call \mathbf{z}_i and \mathbf{t}_i the latent variables of X and Y respectively, and \mathbf{p}_i and \mathbf{q}_i the latent directions. The LV -subspace in the X -space is spanned by the x -loadings \mathbf{p}_i , and the LV -subspace in the Y -space by the y -loadings \mathbf{q}_i , $i = 1, \dots, d$. The x - and y -scores are the projections from the original feature spaces along \mathbf{p}_i and \mathbf{q}_i onto the LV -spaces.

The latent variables in X can be used as predictors for Y [30]

$$\hat{Y} = TQ^T. \quad (2.31)$$

\hat{Y} is the predicted response matrix. To get the prediction directly from X , define the regression coefficients $B = W'Q^T$, and get

$$\hat{Y} = XB, \quad (2.32)$$

which is the final regression model.

Discriminant analysis using PLS

Discriminant Analysis (DA) is a method for discerning the relationship between predefined groups, or classes, and a set of samples [31]. Given a set of m samples and n features the task is to assign each of these samples into the predefined classes. The class affiliation should be decided purely on the samples themselves. A DA model is created by training it with samples of known class labels. From those training samples a discriminant function separating the different classes in the feature space [32] is constructed, and this separator can be used to assign class labels to new samples.

PLS can be applied to create such a classification method [33]. It is named Partial Least Squares Discriminant Analysis (PLSDA) and creates linear separators. An advantage of this method is that it creates the separator in a latent variable space instead of the original feature space. This can create a larger separation between classes and a more robust classification model.

To train the PLSDA model, create a dummy response matrix $Y \in \mathbb{R}^{m \times k}$, where m is the number of training samples and k the number of classes. The columns of Y describe the class labels for the m samples:

$$y_{j,i} = \begin{cases} 1, & \text{if sample } j \text{ belongs to class } i \\ 0, & \text{otherwise} \end{cases}, \quad (2.33)$$

As an example, consider the following: We have a dataset X with samples belonging to three different classes

$$X = \begin{bmatrix} 1 & 4 & 3 \\ 1 & 5 & 3 \\ 2 & 2 & 1 \\ 2 & 3 & 3 \end{bmatrix},$$

The first two rows are samples from the first class, the third sample from the second class and the last sample from the third class. The dummy matrix would then have the following form

$$Y = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

In practice PLSDA should be used with care when the number of variables are larger than the number of samples. The method is greedy in the sense that it might find that some variables correlate by chance. Validation of the method is hence important. Brereton and Loyd have made a nice review over the pitfalls and dangers of PLSDA in Ref. [34].

The class prediction is obtained by performing the PLS regression $\tilde{Y} = XB$. Each column in \tilde{Y} describes how well a sample in X belongs to the corresponding class. Let \tilde{y}_j denote the j th row in \tilde{Y} . We set the class membership of the j th sample to the column index which corresponds to the largest value in \tilde{y}_j [35].

Validation and choosing PLS components

A simple way of validating the classification model is the macro averaged $F1$ -score which is a combination of recall (REC), and precision (PRE) [36]. Assume

a prediction has been made for some samples. Denote the predicted class labels as \mathbf{g} . \mathbf{g} is a vector where element j contains the predicted class label for samples j and so on. For binary classification REC and PRE are defined as as

$$REC = \frac{tp}{tp + fn}, \quad (2.34)$$

$$PRE = \frac{tn}{tn + fp}. \quad (2.35)$$

$$(2.36)$$

tp is the number of true positive predictions by \mathbf{g} , tn the true negative, fp the false positive and fn the false negative. Binary classification is classification with only two classes. Let PRE_i and REC_i denote the precision and recall for class i , then the $F1$ -score for a multiclass prediction is defined in Equation (2.37) [37].

$$F1 = 2 \sum_{i=1}^k \frac{PRE_i \cdot REC_i}{PRE_i + REC_i}. \quad (2.37)$$

k is the total number of classes. Since the optimal score of REC_i and PRE_i is 1.0, it follows that the optimal $F1$ -score also is 1.0. The number of PLS components can be chosen to optimize the $F1$ -score.

To this end, assume \tilde{X} is a dataset and $\tilde{\mathbf{g}}$ a vector describing known class affiliation for each sample in \tilde{X} . Let X_{train} be a matrix constructed from random samples of \tilde{X} , and let X_{test} be made from the remaining samples. The PLSDA model is trained on X_{train} and applied to X_{test} . The result is a vector \mathbf{g} containing the predicted class labels. Several PLSDA models can now be created with different number of PLS components. Then, the $F1$ -score from each prediction can be compared with the ground truth $\tilde{\mathbf{g}}$. The optimal number of PLS components is the one giving the prediction the best $F1$ -score. The same methodology can be applied to check for random correlations by keeping the PLS components constant and training on different, randomized, test- and training sets.

Modelling and experimental work

This chapter introduces the proposed model for underwater hyperspectral images. The model is divided into three main parts: a model of a lamp, of attenuation in water and a reflectance model. After explaining the model the experiments used to determine model parameters is presented.

3.1 Proposed model for underwater hyperspectral images

The system which we want to model is depicted in Figure 3.1. A UHI and a lamp is placed a height h over a scene. The lamp is positioned a distance l_s beside the UHI, and is facing downwards in direction \mathbf{d}_s . We then assume the lamp emits beams of light to each position x , and denote the direction from the lamp to x for \mathbf{d} . The beam travels a distance s_0 before it hits the scene and is reflected. It then travels the distance s_1 and hits the UHI, which measures the light.

For modelling the lamp it is proposed that a lamp emitting a conelike shape can be modelled as a point source positioned a distance l behind the real light source in direction $-\mathbf{d}_s$, as illustrated in Figure 3.2a. P is the imagined point source, S is the real lamp, and s'_0 is the length from the point source to the position illuminated. The rationale is that we expect the light from a real lamp to behave more collimated than a point source positioned in the same location.

A point source emits light equally in all directions which gives rise to a spherical wave. When this wave propagates in space, its total area increases, but for energy to be conserved the total energy needs to stay constant. Hence the irradiance needs to decrease with the same speed as the increase in area. Let us call this phenomenon for the spatial spread of light. The scenario is shown in Figure 3.2b. A point source is emitting light with power Q . After the light has propagated a distance s , it forms a sphere of radius s . Since the area of a sphere is $A = 4\pi s^2$, the irradiance is $E(s) = Q/(4\pi s^2)$. The drop in irradiance from a distance s to




Figure 3.1: A system where a UHI measures incoming irradiance E . A lamp emits light with spectra Q , and angular distribution Θ centered around the lamp's direction d_s . d is the direction to the illuminated position x , and θ_s is the angle between d_s and d . s_0 and s_1 are the lengths traversed by the lightbeam, and N is the surface normal at position x . When incident light hit the scene, it is incident with an angle θ_i to the surface normal, and the reflected light is measured by the UHI at angle θ_r . h is the height from a scene to the UHI, and l_s is the horizontal distance between the UHI and the light source.




Figure 3.2: a) An illustration of the light source model. s'_0 is the distance from the imagined point source to the position illuminated, and other variables are the same as defined in Figure 3.1. b) Illustration of spatial spread. As light is spread into space, the energy is distributed across larger and larger areas. s denotes the length the light has propagated, and the yellow areas represent areas at different s .

$s + \delta s$ is therefore

$$\frac{E(s + \delta s)}{E(s)} = \left(1 + \frac{\delta s}{s}\right)^{-2}. \quad (3.1)$$

The motivation for including the parameter l was to be able to adjust the collimation of the lamp. We now do a sanity check to see if the model has the appropriate properties. If $l \rightarrow \infty$ the change in irradiance disappears for finite δs and the lamp emits parallel rays and is completely collimated. The spherical wave has such a large radius that the rays seem parallel. So the larger the l the more collimated the light behaves, which is what we want. l should be determined using experimental data.

A lamp does not emit light in all directions like a point source, but rather in a finite set of directions. The orientation of the lamp and angular distribution of the light is included by assuming the light is spread in a gaussian shape

$$\Theta(\theta_s) = e^{-\xi(\theta_s)^2/(2\sigma^2)}, \quad (3.2)$$

where $\xi(\theta_s)$ is the position to the illuminated area relative to the lamp on the x -axis. It can be determined as $\xi(\theta_s) = -h \tan(\theta_s)$ by using trigonometry. σ is the standard deviation and needs to be fitted to experimental data. We have assumed that the angular distribution is only a function of θ_s which lies in the xz -plane. This means that we only model a slice of the entire light cone. Because the UHI and lamp moves with the same speed, only a single slice is measured, and hence we only need to model that particular slice.

Reflection is modelled using the Oran-Nayar model described in Section 2.3. As before θ_i and θ_r in Figure 3.1 are respectively the incident and reflected zenith

angle. The UHI and the lamp are both in the same plane, and hence $\phi_i = 0$ and $\phi_r = \pi$ for all x .

Attenuation of light in water is assumed to be described by Equation (2.16). It is also assumed that there are no internal sources, transspectral scattering and no elastic scattering coming from outside the water. Then the equations for radiance can be solved analytically to become

$$L(s, \lambda) = L_0(\theta_s, \lambda)e^{-c(\lambda)s}. \quad (3.3)$$

$L_0 = Q(\lambda)\Theta(\theta_s)$ is the radiance emitted from the lamp at angle θ_s . The RTE is widely used for compensation for the attenuation in water, and has been used in commercial software like HydroLight [38].

We assume that all light beams which are hitting the area A covered by a pixel are parallel with angle θ_i to the surface normal. Then the irradiance E_r reflected of that area into angle θ_r can be written as

$$E_r(\lambda, \theta_s) = \frac{L_i(\lambda, \theta_s)}{4\pi s_0^2} R(\lambda, \theta_i, \theta_r) A. \quad (3.4)$$

Here $L_i = Q(\lambda)\Theta(\theta_s)e^{-c(\lambda)s_0}$ and is the radiance incident at each point within A . Propagating each reflected ray into the UHI we end up with the proposed model

$$E(s'_0, s_1, \theta_s, \lambda) = \frac{Q(\lambda)\Theta(\theta_s)}{(4\pi s'_0 s_1)^2} e^{-c(\lambda)(s_0+s_1)} R(\theta_s, \lambda) A. \quad (3.5)$$

The UHI is modelled as a point and we assume it captures all the reflected rays from area A . We have written R as a function of θ_s since θ_i and θ_r are given by θ_s if we know the geometry of the seabed. For simplicity we also assume that the area A is constant for the heights we will be using. Equation (3.5) can then be used to simulate an underwater hyperspectral image by simulating a line at a time. It can also be used to estimate reflectances from measurements.

For simulating a UHI measuring a scene a simulation tool has been made along with a tools for analysing hyperspectral images. The simulation tool is a 3D ray tracing engine implemented in Python which simulates a line scanning hyperspectral camera acquiring data of a simulated scene. Lamps with adjustable spectra can be placed anywhere in space and can be set to move with the same speed as the UHI. From the lamps beams of light is traced to each pixel in the scene, reflected and traced into the UHI. The UHI is considered a point detector. Different materials can be added by setting reflectances to each pixel. Height differences and an inclination in the scene, can also be represented, however shadow and boundary effects are neglected. See Appendix A for a short documentation of the code.

3.2 Experiments and determination of model parameters

To assess the proposed model with theory we need to measure a scene with a UHI, measure the lamp properties Q and Θ , and the attenuation coefficient c . This section explains the experiments performed and how the model parameters were estimated.

3.2.1 Measuring a scene

For measuring a scene, the UHI was mounted on a rig inside a tank which was filled with tap water. See Figure 3.3a for a sketch of the side view and Figure 3.3b for the front view of the setup. The dimensions of the tank were $H = 80$ cm, $W = 75$ cm, and $L = 129$ cm. The interior of the tank was painted in a dull, low-reflecting, dark color to reduce reflection from the tank walls. The rig had two motors which could move the camera up and down to change the camera's altitude h over the tank floor and its position along the length of the tank. The motors are illustrated as black boxes in Figure 3.3a. Since the camera is using push-broom technology, a scan along the length of the tank is necessary to acquire a full hyperspectral image of the tank floor. Images were taken at 2.5 cm intervals from $h = 28.5$ cm to $h = 53.5$ cm with an exposure time of 20 ms.

The camera measured the incoming irradiance as digital counts on 12-bit scale and was submerged during all measurements. By using an attenuation meter (Viper, Trios) the attenuation of the water in the tank was measured before the experiment was performed. Because the attenuation measured by the attenuation meter is relative to pure water, we need to add the attenuation coefficient of pure water to the measured coefficient to get the total attenuation coefficient. For the pure water attenuation coefficients we used the values measured by Pope and Fry [39]. The program for using the attenuation meter was written in Python using a modbus and the RS485 communications protocol.

Four Lego blocks in different colours were placed on the tank bottom to form the scene. The Lego blocks were placed in a line as showed in Figure 3.3c on top of a grey plate. The grey plate made for easy modification of the Lego blocks' position and alignment. For calculating the extent of one pixel in each image, lines were drawn every 5 cm on the edges of the grey plate. The Lego blocks were 0.5 cm high and 7 cm long and wide.

The room was kept dark during the measurements. The only light source was a halogen lamp (Osram Decostar 35 Titan 50 W) which was attached to the side of the camera angled at 24° towards the center of the scene. A constant current of 3.3 A and voltage of 8.4 V were given to the lamp for all measurements.

Before starting the measurements, the camera had been calibrated to measure bands between 380 nm and 750 nm. The calibration was performed using an Oriel

Instruments 6034 Mercury-Neon calibration lamp with known sharp spectral lines. After calibration, the noise standard deviation in the irradiance was 0.8 counts. The wavelengths measured had an uncertainty of 1 nm.




Figure 3.3: (a): Side view of the experimental setup. $H = 80$ cm is the tank height, $L = 129$ cm the tank length, and h the camera's distance to the tank floor. (b): Front view of the experimental setup. $W = 75$ cm is the tank width. (c): Top view of the tank floor with the four Lego blocks and the grey plate. The figures are not to scale.

3.2.2 Measuring reflectances

The reflectance of the objects were measured by using a spectrometer (JAZ, Ocean Optics) a reflection probe (premium-grade reflection probe, Ocean Optics) and a light source (HL-2000, Ocean Optics). These components were connected through a fiberoptic cable as seen in Figure 3.4a. The light from the source travels through the cable and is emitted from the tip of the probe shown in Figure 3.4b. The yellow circles are openings where the light is emitted. There is also an opening in the middle which captures reflected light and sends it back through the cable to be measured by the spectrometer. The probe was positioned on top of the Lego blocks and the grey plate in a holder (Reflection Probe Holder, Ocean Optics).

To convert the measured reflected light to reflectance we need to measure a reference object of known reflectance. Assume a measurement of the reflected light, E , has been made of some material and another measurement E_r is made of the reference which has known reflectance R_r . Then, assuming that light attenuates as in Equation (3.5) also for air, we can write

$$\frac{E}{E_r} = \frac{R}{R_r}, \quad (3.6)$$

which can be solved for the unknown reflectance, R . The relation will hold as long as the distance from the probe to the reflectance object and the object of interest are the same. Constant distances were ensured by the holder. The holder had two slots in which the probe could be mounted, one at normal incidence to the surface at $\theta_i = 0^\circ$ and one at $\theta_i = 45^\circ$. θ_i is the incoming angle defined in Figure 2.4.

Each Lego block was rubbed with grinding paper to reduce specular reflection. These rubbed blocks were the same as the ones used in the tank experiment. The spectroscopic measurements were performed at four places on each block and two on the grey plate. A grey 20% reflectance spectralon plaque from Labsphere Inc. was used as the reference object.

To compare the reflectance in air with the reflectance in water the above procedure was repeated for submerged probe and objects. The transition is not completely trivial. According to Fresnel the reflectance changes when the surrounding media changes. For specular surfaces this change is given by Fresnel's equations [40]. The relation is however not applicable to diffuse reflection as from the roughened Lego blocks or from the spectralon. For calculating the underwater reflectances the known reflectance R_r was used without any compensation for the water.




Figure 3.4: a) Setup for measuring reflectances. A light source L is connected to a fiberoptic cable ending in a probe P. The probe emits the light onto the surface of a Lego block B, which reflects it back into the probe. The fiberoptic cable then carries the light into a spectrometer S. b) The tip of the probe. Light is emitted from the yellow circles and is collected by the red circle. (Figure: Ocean Optics).

3.2.3 Measuring lamp properties

Energy spectrum

To measure the energy spectrum of the lamp it was positioned at the bottom of the tank in Figure 3.3, replacing the Lego blocks. It pointed upwards while a measurement by the UHI was made. During the measurement the UHI was 1 cm above the lamp and the tank was not filled with water. We approximate Q as the total radiance captured in the pixels which makes up the lamp.

Collimation and angular distribution




Figure 3.5: a) Sideview of the UHI and striped plaque. The dotted line shows the line which is imaged by the UHI. b) Illuminated striped plaque from the front showing the UHI position as x_{UHI} and the distance to the lamp l_s .

To determine the parameters σ from Equation (3.2) and l from Figure 3.2a, we repurposed the rig above the tank in Figure 3.3. The lamp and the UHI were

placed on the left side facing a striped plaque. The plaque was attached to a base which could be moved back and forth with a motor. A sideview of the setup is shown in Figure 3.5a and the front of the striped plaque in Figure 3.5b. Each stripe had a width of $\delta w = 0.5$ cm. The fixed center of the lamp was $x = l_s = 10$ cm away from the center of the UHI at $x = 0$ in the x -direction and level in the y -direction. We estimate an uncertainty of 0.5 cm in these position measurements. The lamp was directly facing the plaque, with an uncertainty of 3° in its orientation in the xz -plane and 1° in the yz -plane. While the lamp needed to be placed on an improvised platform, the UHI was placed in a pre-existing holder making it face the plaque directly with low uncertainty.

Measurements of the plaque were made for $h = 25$ cm to $h = 100$ cm in 5 cm intervals. All measurements were done in air with no other light sources than the lamp. The UHI and lamp were kept stationary, and the plaque was moved to each position with an uncertainty of 0.05 cm. Since this is a push-broom camera, only one line of the stripe plaque were measured, as illustrated by the dotted rectangle in Figure 3.5b. Let the measured irradiance at position x along the plaque be $E_m(h, x)$ for some wavelength. We assume that E_m follows the proposed model in Equation (3.5) and that $\Theta(h, x)$ is wavelength independent. Then denote $\hat{E}_m(h, x)$ as the maxima normalised irradiance. The measured spatial light distribution $\Theta_m(h_{\text{ref}}, x)$ is then found as

$$\Theta_m(h_{\text{ref}}, x) = \frac{s_0'^2(h_{\text{ref}}, x) s_1^2(h_{\text{ref}}, x)}{s_0'^2(h_{\text{ref}}, x_{\text{max}}) s_1^2(h_{\text{ref}}, x_{\text{max}})} \hat{E}_m(h_{\text{ref}}, x) \Theta(h_{\text{ref}}, x_{\text{max}}) \quad (3.7)$$

where x_{max} is the position where E_m is at its largest and h_{ref} is the distance to the measurement which will be used to fit Θ to Θ_m . The normalization is needed since we neither have the reflectance of the plaque nor a good measurement of the size of Q , which we will come back to later in Chapter 4. The reflectance can in principle be measured like in Section 3.2.2.

By using a least square fit we can determine σ by fitting the gaussian in Equation (3.2) to the measured distribution in Equation (3.7). A problem is that we do not know $\Theta(h_{\text{ref}}, x_{\text{max}})$ from Equation (3.7). The problem is easily solved since Θ should equal 1 at $\theta_s = 0$, and can normalise Θ_m such that it's maximum is 1 before fitting.

Θ was above written as a function of h and x since it is more convenient during the fit and for describing the measured distribution for several heights. To use the fitted Θ on measurements from other heights than h_{ref} , we change variables from x to θ_s . Since x_{UHI} is defined to $x = 0$, and the lamp is l_s beside the UHI, ξ from Equation (3.2) is $\xi = x - l_s = -h_{\text{ref}} \tan \theta_s(h, x)$. Remember that θ_s is the angle between the lamp direction d_s and the direction to the location imaged d . Then we

can write Equation (3.2) as

$$\Theta(h, x) = e^{-h_{\text{ref}} \tan(\theta_s(h, x)) / (2\sigma)^2}. \quad (3.8)$$

Because there is two model parameters to be determined we need one more equation. It can be found from the change in irradiance from one distance h_0 to another distance h_1

$$\frac{\hat{E}(h_1, x_1)}{\hat{E}(h_0, x_0)} = \frac{\hat{E}_m(h_1, x_1)}{\hat{E}_m(h_0, x_0)} \equiv \Delta E(h_0, h_1, x_0, x_1). \quad (3.9)$$

Here x_0 and x_1 are the positions measured in measurements at h_0 and h_1 respectively. By writing the Equation out we get a second order polynomial for l :

$$\begin{aligned} & l^2 \left[s_1^2(h_0, x_0) - \Delta E \Delta \Theta s_1^2(h_1, x_1) \right] \\ & + 2l \left[h_0 s_1^2(h_0, x_0) - \Delta E \Delta \Theta s_1^2(h_1, x_1) \right] \\ & + (h_0^2 + (l_s - x_0)^2) s_1^2(h_0, x_0) - (h_1^2 + (l_s - x_1)^2) \Delta E \Delta \Theta s_1^2(h_1, x_1) = 0. \end{aligned} \quad (3.10)$$

where $\Delta \Theta = \Theta(h_0, x_0) / \Theta(h_1, x_1)$.

Having Equation (3.7) and (3.10) there are at least two approaches for determining the parameters l and σ . The first approach uses equation (3.10) and lets the positions measured be set as $x_0 = l_s$ and $x_1 = l_s$. Then $\Delta \Theta = 1$ and finding l reduces to finding the roots of the polynomial. σ can then later be estimated by Equation (3.7). The second approach is an iterative one. First set l large and fit a gaussian according to Equation (3.7). Then use the σ estimated and use Equation (3.10) to calculate l and put it back into Equation (3.7), and repeat the procedure until σ and l changes minimally.

We have to be cautious when estimating Θ of the lamp in air when it is to be used under water. The housing around the coil inside the lamp is surrounded by air and we expect the light to refract when the media is changed. To account for the refraction between the air inside the lamp and the water outside we apply Snell's law of refraction [40] on θ_s to get the refracted angle θ_s^{water}

$$\theta_s^{\text{water}} = \arcsin \left(\sin(\theta_s) \frac{n_{\text{water}}}{n_{\text{air}}} \right). \quad (3.11)$$

Where $n_{\text{water}} = 1.33$ and $n_{\text{air}} = 1$ is the refractive indices of water and air. The spectral dependence in the refractive indices are ignored.

Results and Discussion

This chapter presents and discusses the main findings in this thesis. We start by briefly looking at PCA as a method for investigating attenuation and illumination, then at the reflectance measurements, and continue by determining the model parameters describing the lamp. Thereafter, a comparison of the tank experiments and simulations will be made, and lastly we look at classification.

4.1 Explaining variances using Principal Component Analysis

A method which was briefly investigated in this thesis was PCA, and we include here a short summary of the result. As seen in Section 2.4, PCA finds the latent variables describing the most variation in a dataset. One could imagine that the effect of illumination and attenuation could be sorted into separate principal components. By using the ray tracer to simulate the tank floor, from the experiment in Section 3.2.1, we created several images taken at different heights. In each image we extracted areas containing the Lego blocks and applied PCA to the combined set of samples. In this way variations with height could be investigated. To only look at the variation in attenuation, we set $\Theta(\theta_s) = 1$ for all θ_s . Later, illumination was added and we looked at variations along the grey plate in the x -direction. Using PCA we were in general only able to tell that there were some variation in the image connected to height and position, but nothing quantitative about the processes causing the variation.

We believe PCA is more useful when investigating correlations in systems where direct modelling is not possible, or situations where one has measurements, but no developed theory. In other words, PCA is more appropriate when little prior knowledge about the system is known, and we want to qualitatively investigate the sources for variation. Since it is unlikely that PCA will be used in future work for modelling illumination and attenuation, we will not focus on it further on.

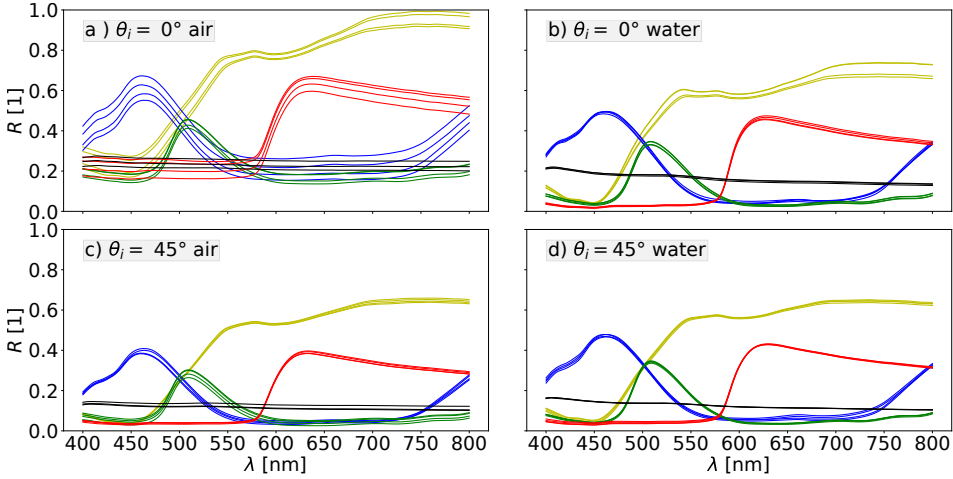


Figure 4.1: Calculated reflectances for air and water for the Lego blocks and the grey plate at 0° and 45° incidence. The color of a graph represents corresponding color of the Lego block and the black lines represents the grey plate.

4.2 Reflectance of test objects

In Figure 4.1 the calculated reflectances from the roughened Lego blocks are shown. They were found by solving for R using Equation (3.6). The yellow line corresponds to the yellow block, the red line to the red block, and so on. Each line represents one measurement. Before calculating the reflectances, the measured reflected light was smoothed using a Savitzky-Golay filter [41]. The reference reflectance R_r was given digitally from the producer (Labsphere), but the wavelengths given and the wavelengths measured from the spectrometer were different. Hence, to calculate the reflectances, we needed to interpolate R_r to match the wavelengths measured, which was performed using a cubic spline. A cubic spline is a piecewise cubic polynomial used for interpolation [42].

For the reflectances measured at $\theta_i = 0^\circ$ in air there are, large variations in magnitude. This is probably caused by two factors: there are local variations in the reflectance caused by uneven grinding, damages, or other contaminations. Additionally these differences are increased since the light is coming in at normal incidence giving rise to specular reflection. Using the mean for each color, and fitting κ in the Oren-Nayar model in Equation (2.19), we get $\kappa \approx 0$. In other words, the Lego blocks reflect much like a lambertian surface with mainly diffuse reflection.

The reflectance for the submerged samples are curious. Both angles of incidence, $\theta_i = 0^\circ$ and $\theta_i = 45^\circ$, gave approximately the same reflectance spectra with

the reflectance at normal incidence tending to be slightly larger. The reflectance measurements for the Lego blocks in water were performed twice to check for experimental errors, however both measurements showed the same results. We expect that there is another systematic error tied to this experiment. A candidate is that we have not accounted for the change in reflectance of the reference object when it is under water. Most likely the change would not affect the angular distribution, but rather the magnitude of the reflectance for each wavelength, and hence is probably not the cause.

Another thing to note is that the reflectances from the Lego blocks tend to be grouped into two separate sets for each of the reflectance plots. It is most prominent for the yellow block. Each Lego block was built up of six smaller blocks. The reflectance measurements were performed on two of these, two locations on each. The grouping could be caused by the fact that those smaller blocks have a small difference in reflectance. The effect was also seen for the second reflectance measurement. We also observe that both the submerged reflectances are very similar to the dry reflectance at 45° .

It is evident that the change of medium affects the reflectance of a diffuse object, especially its angular dependence. It is however not certain how this can be corrected for in a proper way. Therefore, since the measured reflectance is approximately the same at $\theta_i = 45^\circ$ and at normal incidence, we have assumed that reflectance of the Lego blocks is constant in the simulations further described in this thesis. We further assume it is equal to the submerged reflectance at $\theta_i = 45^\circ$. A disadvantage of this approach is that the energy of the simulated system might not be conserved since there is equally emitted light in all directions. In the future, a proper treatment of reflectance should be measured with a gonireflectometer. Optimally such measurements should be done both above and below water for comparison.

4.3 Determining lamp properties

This section presents the results from the experiments determining illumination properties. We will start by presenting the measurement for determining the lamps spectrum Q . Figure 4.2a shows an image of the measurement of the lamp for $\lambda = 615$ nm. Note that it has been normalized such that the colorbar describes irradiance from 0 to 1. We clearly see the reflector surrounding the coil in the middle and notice that the irradiance seems to have a maximum on the reflector and not on the coil. This is probably caused by the coil emitting most light to the sides and not in the forward direction. In Figure 4.2b the measured spectrum is shown. Note that it has been normalized because we were not able to measure the energy emitted from the lamp sufficiently well. Doing simulations using the model

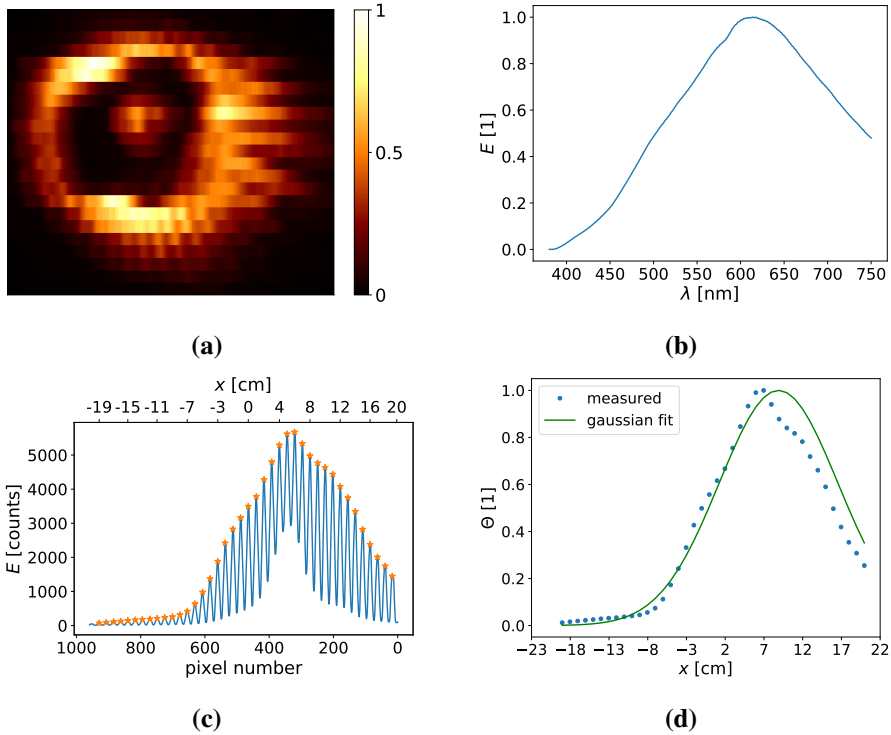


Figure 4.2: a) Image of the lamp represented through one wavelength, $\lambda = 615$ nm. The measurement in each pixel was normalized to the maximum in the image. b) Normalized lamp spectrum. c) Measurement from the striped plaque at $h = 25$ cm in digital counts. The orange dots mark the peaks which define Θ_m d) The measured Θ_m is depicted with blue dots, and the green line is the fitted Θ . c) and d) are plotted for $\lambda = 615$ nm

of the lamp, and comparing with the measurements of the striped plaque and the measurements in the tank, it was clear that the energy was underestimated. We have hence chosen to use the normalized spectrum in the simulations discussed further.

When illuminating the striped plaque, the white stripes appear as peaks when plotting $E(h, x)$, as shown in Figure 4.2c. The physical location of these peaks can be estimated. Assume that the UHI is positioned in the center of the image, and define $x_{\text{UHI}} = 0$. Then the position of the lamp in the image is $x_s = l_s$. Since we know that the width of each stripe is 0.5 cm the distance between each peak (white stripe) is 1 cm. By counting the number of pixels between each peak we can estimate how many pixels there are per centimeter and determine the location of the peaks relative to the UHI. It has been assumed that the peaks are located in the middle of each stripe. For the calculations below, measurements taken at

$\lambda = 615 \text{ nm}$ will be used as it was the wavelength with the strongest signal.

For the iterative estimation of the lamp parameters, l was initiated to 10m and $x_0 = x_1 = l_s = 10 \text{ cm}$ before estimating the initial σ . Remember that l was the position of the point source from Figure 3.2a, and σ was the standard deviation of the lamp's emitted distribution. After three iterations we end up with the values $l = 19 \text{ m}$ and $\sigma = 7.13 \text{ cm}$. For the direct approach of estimating l and σ we set $\theta_s = 0$, which implies $x_0 = x_1 = l_s$, and gives $l = 11 \text{ m}$ and $\sigma = 7.13 \text{ cm}$. This l differs significantly from the iterative method, and we further test for different values of θ_s in the direct method, and x_0 and x_1 in the iterative method. What is found is that there are large fluctuations from $l = 60 \text{ cm}$ to $l = 19 \text{ m}$ in both approaches, whereas σ however stays at a fairly constant value, that is between $\sigma = 7.13 \text{ cm}$ and $\sigma = 7.18 \text{ cm}$. The small variations in σ even though there is large variation in l can be explained by the fact that when $l \gg h$ we get $\Theta_m \approx \hat{E}_m(h, x)$ in Equation (3.7).

The large variations in l are probably caused by the roots in Equation (3.10) being sensitive to errors in $\Delta\Theta$ and ΔE . Since we only have discrete samples of $\hat{E}_m(x, h)$ due to the nature of the acquisition method, x_0 and x_1 is not accurate. This inaccuracy makes $\Delta\Theta$ and ΔE inaccurate. Errors are also introduced in $\Delta\Theta$ by the fitting of Θ itself. Figure 4.2d shows Θ_m and the fitted Θ with $l = 11 \text{ m}$ and $\sigma = 7.13 \text{ cm}$. Note that the peak is not located at $x = 10 \text{ cm}$, which is the lamp's location, but rather about 2 cm to the right of it. This is probably caused by the lamp not facing the plaque directly, but at an angle causing the peak to be shifted to the left. It could also be that Θ should include an offset due to the construction of the lamp. When simulating further on, we will not add such an offset to the model due to the uncertainty involved. In the fit we see that there are absolute errors of up to 0.15. Introducing such an error into $\Delta\Theta$ in the calculation of l gives rise to changes of 500 cm. The problem of determining l accurately from Equation (3.10) hence seems like an ill posed problem. Estimating l and σ from measurements at other heights gives the same qualitative behaviour.

In Figure 4.3 measurements and simulations of the experiment of the striped plaque is plotted for each distance h . The blue dots are the measured $\hat{E}_m(h, x)$, the red line is from a simulation where $l = 19 \text{ m}$, the orange line for $l = 11 \text{ m}$ and the green line for $l = 560 \text{ cm}$. The last l was found by visual optimisation by the criteria that the reduction in irradiance from the simulated peak at $h = 100 \text{ cm}$ to the one at $h = 25 \text{ cm}$ should be the same as the reduction for the measured peaks. Note that all measurements are normalized to the maximum of the measurement at $h = 25 \text{ cm}$, and the same is done for the simulations. This normalization is done because we do not know the energy put into the system, and the irradiances for the simulation and the measurements are thus bound to be different in size. This also makes it easy to compare how much the measured and simulated irradiances drop

between each height. We notice that the reduction in irradiance in the simulation with calculated l s is too low. The model describes the reduction in irradiance fairly accurately for the visually optimised l . The value of l is not really important as such, it is a fictive distance which is used as a tool for describing collimation. Even though l varied much, we always got $l > 0$ when estimating l , which implies that there is need to account for collimation.

The distribution in the simulations evolves in a qualitatively correct way with distance, but there are asymmetries in the measured shape which are not captured by the current model. Looking at the image of the lamp from earlier in Figure 4.2, this should be expected. An effect which we do not expect is the change in location of the peak in the measured irradiance. If the lamp was not facing the plaque directly, but at an angle, we would expect the peak to drift away from the simulated peak. We observe however that from $h = 35$ cm to $h = 55$ cm the peak drifts as expected, but after that it drifts in the opposite direction and seems to overlap with the location of the simulated peak.

4.4 Tank experiment and comparison with simulations

In this section we will present the results from the measurements in the water tank, illustrated in Figure 3.3, along with simulations. We split this section into the spatial- and spectral properties of the model.

4.4.1 Spatial properties of the model

As with the experiment with the striped plaque $x = 0$ cm is defined as the position of the UHI, which again is assumed to be in the middle of the image. $y = 0$ cm is defined to be at the bottom of the grey plate. The black lines drawn onto the grey plate were used to determine the physical length of each pixel and determine their position in each image, in the same way as was done with the white stripes earlier.

On the right hand side of Figure 4.4 is an RGB representation of the hyperspectral images taken at $h = 28.5$ cm, $h = 41$ cm and $h = 53.5$ cm. The simulated images are shown on the left in the same figure. Both have been cropped to show only the region of the image containing the Lego blocks and the grey plate. Qualitatively the images look similar, that is, they have the same qualitative distribution of light and colouring. The green block however looks too dark compared to the measurement, and the location of the peak in the light distribution seems to be slightly off. However, this is expected, since we also observed such an error in with the striped plaque experiment.

For a more quantitative assessment, the normalized irradiance across the grey plate is plotted in Figure 4.5. As in Figure 4.3, all irradiances are normalized to the

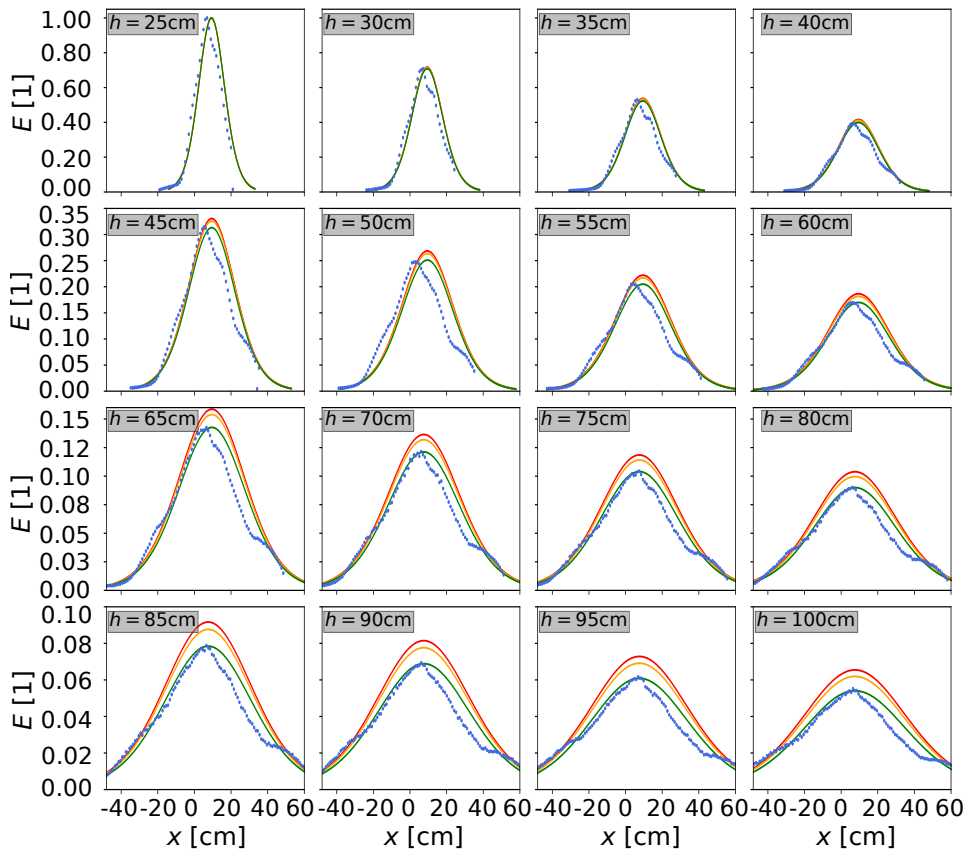


Figure 4.3: Normalized measurement of the plaque at several distances for $\lambda = 615$ nm along the x axis. All plots are normalized to the maximum at $h = 25$ cm. The blue dots show the measured irradiance, and the green, orange and red the simulated irradiance. The values are normalized with respect to the maximum value in the measurement at $h = 25$ cm.




Figure 4.4: RGB representations of the simulations a,c,e) and of the measured data b,d,f).

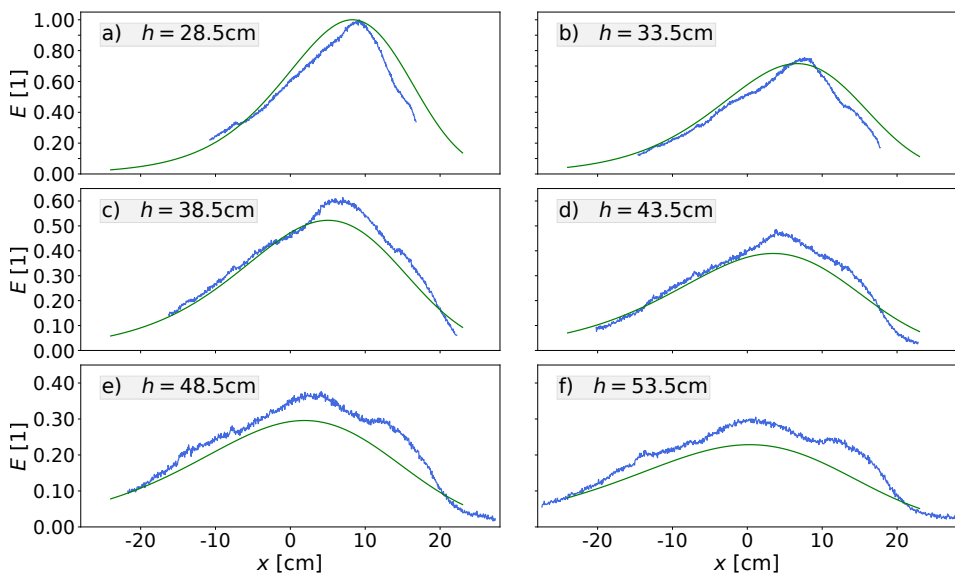


Figure 4.5: Light distribution along a measured line (blue) in the experiment and in a simulation (green) for $\lambda = 615$ nm.




Figure 4.6: Illustration of a light cone hitting the tank floor , and two slices which could be measured by an UHI.

maximum in the first measurement, that is, the one at $h = 28.5$ cm. The blue line is the measured irradiance, and the green line is the simulation. In the simulation $\sigma = 6.3$ cm and $l = 560$ cm have been used, which differ from those used in the earlier discussion. This change was done since the earlier found σ gave a bad fit, which was probably caused by a different slice of the light cone being measured in this experiment. In Figure 4.6, the light distribution from a light cone hitting the tank floor is shown with two slices. Remember that the lamp moves with the UHI and hence the same slice of the light cone is measured for the entire image. Slice 1 is hitting the light cone in the center and is hence wider than slice 2 which is off center. If the UHI measured slice 1, we would expect σ to be larger than if slice 2 was measured. The orientation of the lamp has probably also changed for this experiment. Since the lamp is asymmetric, which may also lead to different illumination conditions.

We notice that the simulated irradiance in Figure 4.5 drops quicker than in the measurement. One could think that this difference comes from an error in l , since such a behaviour was seen when estimating l with the striped plaque. However, modifying this parameter gave no improvements. Which suggests adding the attenuation term has a collimating effect on the model. Another source for error could be that the lamp does not point in the correct direction d_s . We see that the peak of the distributions moves slightly, but that is expected after the measurements with the plaque. We do not see a shift which indicates that the direction is significantly wrong, i.e. that the locations of the simulated peak and the measured peak move further apart with increased height. A third possibility is that the attenuation coefficient c from Equation (3.5) is too large. Redoing the simulation and scaling c down slightly improves the situation, but even removing the attenuation altogether by setting $c(\lambda) = 0$ for all λ did not completely make up for the high dropoff. It is not likely that the error in the attenuation coefficient is large. A fourth

possibility is that there is a systematic error when measuring the heights so that the distance between each image is overestimated. Again redoing the simulation, and this time adjusting the simulated heights, we are able to achieve a reduction in irradiance which matches the measured one for each height. However, to entirely compensate for the increased dropoff, the distance between each measurement had to be reduced by 0.5 cm, which is large compared to the uncertainty in the measurements. Therefore, given that the reduction in irradiance was accurate in air, it is believed that there is an aspect of attenuation in water which is not captured by the proposed model.

4.4.2 Spectral properties of the model

Attenuation in water effects light in two ways. It both dampens and changes its' shape. The change in shape arises due to different wavelengths being attenuated differently. Due to the illumination on the tank floor being uneven, and different for each height imaged, it is difficult to conclude anything specific regarding the magnitude of the attenuation. Hence, only the change in shape could be considered from the performed tank experiment. This was done by normalising and then mean centering the measured spectra before and after correcting for attenuation. By correcting for attenuation we imply using Equation (3.3) to remove the effect of propagation in water. This is done by solving for L_0 , and then use that we have assumed that all rays incident on the area covered by each pixel are parallel. Then the irradiance also attenuate exponentially.

Figure 4.7a-c) shows the measured-, pure water-, and total attenuation. Note that the measured attenuation was smoothed with a Savitzky-Golay filter before adding it to the pure water attenuation to form the total attenuation. The spectra shown in d, h, l, p) are the measured spectra from the yellow, blue, green and red Lego blocks respectively at $x \approx 0$. PCA has been used for noise reduction, and the spectra has been normalised to the maximum from the measurement at $h = 28.5$ cm. e, i, m, q) shows the maximum normalized spectra. From these figures we observe that there is a change in shape, but over the distances measured it is small. For a clearer visual comparison, the spectra are both normalised and mean centered in Figure 4.7f, j, n, r).

Focusing on the modified yellow spectra in Figure 4.7e) and f), we see that the main difference in shape is at wavelengths greater then 600 nm. This harmonizes with the attenuation coefficient in Figure 4.7c). From 480 nm to 600 nm, c changes little compared to the jump at 600 nm. Since we expect exponential damping, such a jump should have noticeable effects on the shape of the spectra. For wavelengths larger than 600 nm, c continues to rise and we observe increased difference in the spectra's shape towards longer wavelengths. The same is observed for the spectra

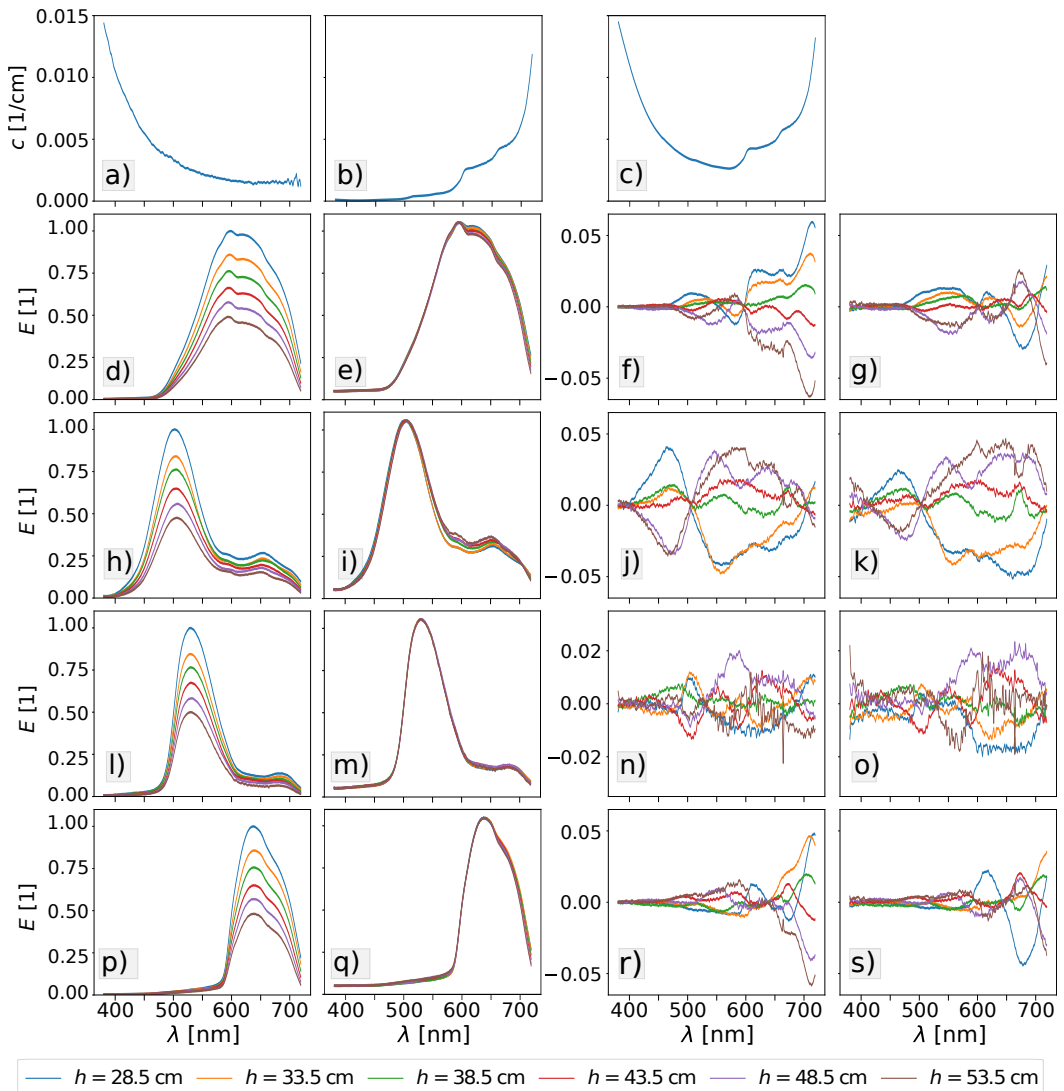


Figure 4.7: a,b,c) measured-, pure water- and total attenuation coefficients respectively. d, h, l, p) Measured spectra from the yellow, blue, green and red Lego blocks respectively at $x \approx 0$ for $h = 28.5$ cm to $h = 53.5$ cm. e, i, m, q) The same spectra normalised to their maximum values. f, j, n, r) The spectra normalised and then mean centered. g, k, o, s) Spectra after correcting for attenuation, normalised and mean centered. The legend applies to figures d-s)

from the red and green block, but the spectra from the blue block shows slightly different behaviour.

When correcting for attenuation, we expect there to be small corrections for the blue and green spectra, since they are located at relatively constant c . However, for the blue spectra there seems to be an overcompensation for longer wavelengths. There the variance increases, while for lower wavelengths it decrease. The situation is reversed for the corrections to the yellow and red spectra. There the variation clearly decreases for longer wavelengths, and slightly increases for shorter. This suggest that the model does not accurately describe the attenuation. It seems it is able to partly correct for attenuation for short distances, but to make any definitive conclusions, measurements taken further apart has to be made.

A comparison between simulated- and measured irradiance spectra from the tank experiment are shown in Figure 4.8. These are from pixels close to $x = 0$ from the image acquired at 28.5 cm. Figure 4.8a) shows the normalised simulated spectra from the Lego blocks and the grey plate, and 4.8b) the normalised measured spectra. These spectra are normalised relative to the maximum in the yellow spectra. The color of each line shows which block it is from and striped lines indicate measured spectra. Other than the simulated blue spectra being slightly higher than it should in the upper figure, the simulation seems to give qualitative good spectra. The same behaviour were seen at images acquired at the other heights and the model represents the relative magnitudes of the spectra well.

Figures 4.8c-f) show comparisons between the spectra from each Lego block in the simulation and the measurement. They are normalised to their maximum value to compare shapes. Directly comparing the simulated spectra and the measured ones, we see a systematic deviation. The simulated red, blue and green spectra are shifted noticeably to lower wavelengths and the yellow spectrum is too wide. Such a shift hints to an error in the calibration of either the UHI, spectrometer and/or the attenuation meter. These calibrations were checked and proved to be correct within 1 nm. By shifting the reflectance such that the measured and simulated spectra overlap, we get a good similarity. Hence it is believed that the shift comes from systematic errors in the reflectance measurements. Similar shifts in the lamps emitted spectra or the attenuation coefficient did not produce any good corrections.

The measured spectra and estimated reflectance for a randomly chosen pixel in each height are shown in the first two columns in Figure 4.9. The estimation is done by solving Equation (3.5) for R and inserting the measured spectra and all of the remaining known variables. Note that the scaling is not correct since we have not accounted for the energy emitted from the lamp, but rather used the normalized emitted spectra. The shown reflectance is hence not the true reflectance, which should be on the interval $[0, 1]$.

To obtain a quantitative measure for how much variation there is within the




Figure 4.8: Comparison between simulated- and measured spectra from the tank experiment. These spectra are taken from pixels close to $x = 0$ from the measurement at $h = 28.5$ cm. Striped lines show the measured spectra and whole lines the simulated.

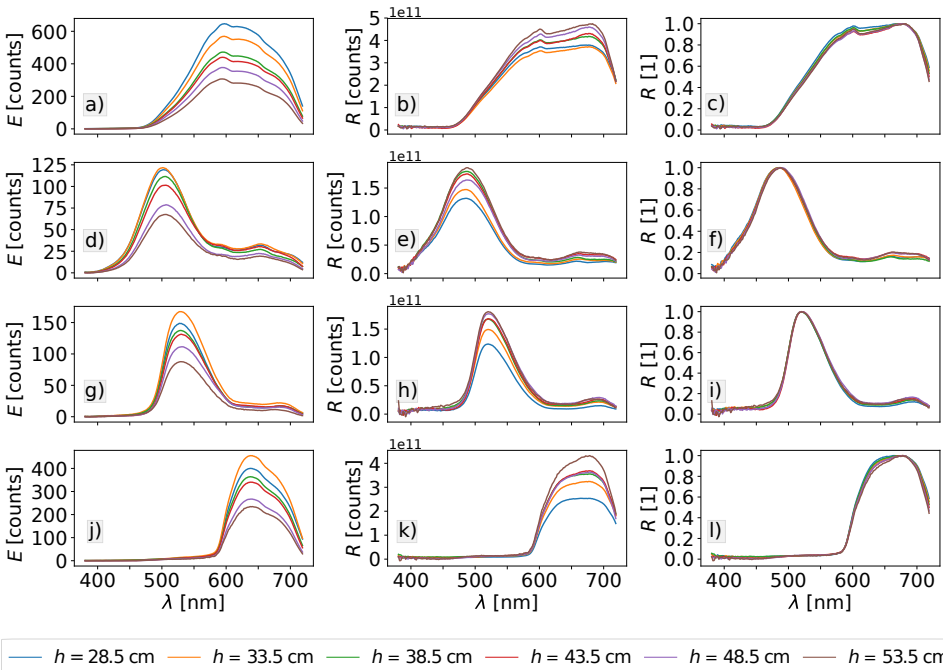


Figure 4.9: Spectra from random pixels on the Lego blocks. a, d, g, j) shows the spectra from the yellow, blue, green and red Lego blocks respectively. b, e, h k) show estimated reflectances using said spectra, and the c, f, i, l) show the normalised estimated reflectances.

spectra from each block, we calculated the variances for each color. The same was done for the estimated reflectances. Since the scalings are different for the measurements and the estimations, we first divided by the respective means. For these pixels, we observe that the relative variance decreases after reflectances have been estimated, but how much depends on both color and wavelength. Decreased variance means a better classification since the samples to be classified resemble the training samples. We see a 70–80% decrease in relative variance for these spectra. However, when estimating reflectance for other random pixels we notice that the variation may also increase. This probably occurs because the lamp model over- and undercompensated illumination. By normalizing the estimated reflectances, we see that the shapes correspond well, as seen in Figures 4.9c, f, i, l).

4.5 Classification

In this section we will investigate the classification of images taken by an UHI, and how the water column and illumination affect classification accuracy. We will also use Equation (3.5) to estimate reflectances, and see if this improves the classification.

Before doing the classification, we need to define which pixels will be used to train the classification models. In total, four models were made for each preprocessing strategy. First we trained three models using the closest measurement at $h = 28.5$ cm. Rectangles were defined within each Lego block which contain the training set, as seen in Figure 4.10. The three training sets are marked by numbers 1–3. Training set 1 contains essentially the whole block, set 2 a portion of the right side of each block and set 3 contains a portion of the left side. Note that training set 1 also contains the pixels of training sets 2 and 3. The purpose of these three training sets is to see how variations in illumination influences the classification. Training set 2 contains the brightest pixels, and training set 3 the dimmest. A classification model was also trained on the image taken at $h = 53.5$ cm. In this case the illumination is fairly uniform compared to the image taken at $h = 28.5$ cm, and only one classification model was made, using a rectangle containing the whole Lego block. Denote this set training set 4, and also denote a model trained by training set i for model i .

Cross validation was performed for all classification models by creating them with 70% of the training set and testing on 30%. All models were performing with good classification within the test set with 7 PLS components.

To calculate the F1 scores, given by Equation (2.37), the training sets containing the entire blocks are regarded as ground truth. Such sets were made for the measurements at all heights. The F1 scores from a classification using the raw spectra can be seen on the left in Figure 4.11. Model 2 clearly gives the worst clas-




Figure 4.10: Training sets 1-3 marked from the measurement at $h = 28.5$ cm marked as rectangles. Model 1 covers the entire Lego blocks, model 2 the bright part of each block and model 3 the dimmer part.

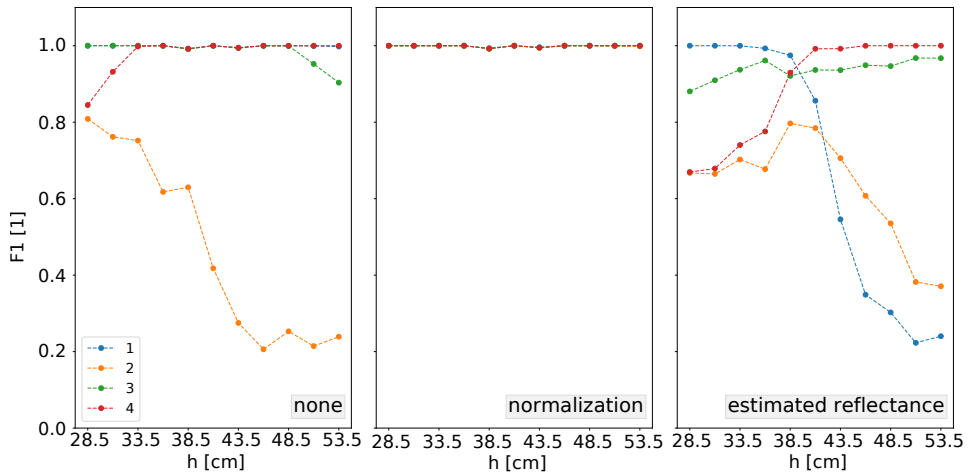


Figure 4.11: F1 scores for classification models 1-4 with different preprocessing strategies: no preprocessing(left), normalization(middle) and estimated reflectances (right).




Figure 4.12: Prediction from a classification with model 2 on the image taken at $h = 28.5$ cm. The colors denote which class the pixels are classified as: yellow, blue, green and red for the Lego blocks and black for the grey plate. The striped rectangles shows the pixels used as ground truth for each class.

sification with a maximum F1 score of 0.8 on the height at which it was trained on. Then the F1 score quickly drops as pixels in other images taken further away are classified. The other models seem to classify fairly well. We note however that the F1 score for model 3 and 4 also drops when used on images taken far away.

The bad prediction of model 3 probably stems from the fact that it is trained on bright pixels. When the model misclassifies it misclassifies pixels to grey plate pixels, as seen in Figure 4.12 for the image at $h = 36$ cm. The black pixels are where the model predicts the grey plate and the other colors are for the different Lego blocks. The stapled rectangles show the ground truth. The reason the Lego blocks are classified as a part of the grey plate is twofold. The spectrum for the grey plate covers all wavelengths as well as being the dimmest of all the classes in the model. This combination makes the dimmer spectra from the Lego blocks look like the grey plate when classifying. In other words, the variation in the spectras' magnitude is not covered by the classification model. This is true both for variances in spectra due to illumination and due to propagation in water. Models 1, 2 and 4 perform better since they better capture the variation in the spectra. It is however expected that if images had been taken farther away they would also perform poorly.

We see that the classification is heavily dependent on the training set, and so far the failed classification is caused by the differences in irradiance between training

samples and the samples to be classified. We have two methods of compensating for such differences, by using the proposed model to estimate reflectances, and normalizing the spectra in each image.

By normalizing the spectrum in each pixel to their maximum value, we get a very good classification as seen from the F1 scores in the middle plot of Figure 4.11. This tells us that the change in shape of each spectrum is small over the distance the light propagates in these measurements. We expect that normalizing the measurements will not work for classification done on images from different water bodies, for water bodies with large c or for measurements where the propagation length is long. This is because each wavelength attenuates differently.

In the plot on the right in Figure 4.11 the F1 scores of a classification using the estimated reflectances is shown. The classification is clearly worse than for the other preprocessing methods. This is tied to the classification models being sensitive to variations in the size of the reflectances and that the illumination model over- and undercompensates, as discussed earlier. By normalizing the estimated reflectances, we get equally good classification as for the normalised raw spectra, which again indicates that the shapes of the estimated reflectances for each block are similar.

4.6 Note on earlier work

Earlier, in project work [9], we made a simpler model for modelling attenuation in water. In this model we neglected the spatial spread of light, and assumed even illumination at all times. It was on the following form

$$E(h, \lambda) = RE_0(\lambda)e^{-c(\lambda)h}, \quad (4.1)$$

where E is the measured irradiance, E_0 is the irradiance emitted from the lamp and h the distance from the UHI to the tank floor. Note that h enters the equation instead of s like in Equation (3.5).

In the project we did not have the equipment available for measuring the reflectance R or the attenuation coefficient c . Hence, we estimated c by using the relative difference between the measured heights h and the measurement at the lowest height h_0

$$c = -h \ln \left(\frac{E(h, \lambda)}{E(h_0, \lambda)} \right). \quad (4.2)$$

Then, by using linear regression, c could be estimated. After estimating c , we mapped the measured irradiances from all heights to h_0 . We then got a good similarity in both the spectra's magnitude and shape.

The results was not reproducible with the current experimental setup. Most likely this is due to the previous work using two lamps instead of one, and therefore acquiring significantly more even illumination. As commented in the earlier work, we were using a subset of the data in which we were applying the model, to estimate c . Which could have led to artificially good results. The effect of small changes in illumination with height, and spatial spread could have been caught unintentionally by the simpler model.

4.7 Further work

This thesis has been quite broad encompassing both an illumination model, a model for propagation in water, and tested classification. Some of the challenges with modelling images taken by a UHI has been uncovered, and in further work these challenges should be handled separately for the best results. Some suggestions follow.

If the magnitudes of the irradiance is to be modelled correctly, the illumination model needs to be improved. The model of the lamp assumes a Gaussian distribution in the plane. An asymmetric shaped function which is more peaked should be tested. Also, the assumption that the center of the lamp and the center of the UHI is in the same plane should be removed. A complete 3D model of the lamp should be made. Then it would be possible to set the parameters of the model once, and thereafter be able to use the lamp in arbitrary positions and orientations relative to the UHI.

For validating a 3D model of the lamp there are two alternatives. One option is to make a bigger striped plaque with striped in both the horizontal and vertical direction, and replacing the push-broom hyperspectral imager with a camera that has a 2D FOV. The same methodology used in this thesis could then be applied to create the 3D model. For a better validation of the model, the experimental setup should be improved. During this thesis, the measurements with the striped plaque, described in Section 3.2.3, were partly improvised. A rig for doing this exact experiment should be constructed with more precision. It should be able to rotate the lamp along its center, and angle it in both the xz -plane and yz -plane.

The other option is to measure angular distribution directly using a goniometer. Also in this scenario, the UHI should be switched with another, more practical, measuring device. As a design suggestion, such a goniometer could relatively cheaply be made with microcontrollers and stepper motors controlling the position of the measurement device. Then, human error would be removed and the time used for taking the measurements could be spent elsewhere.

An effect which has not been the focus of this thesis, but which is going to be a problem in field measurements, is shadows. The lamp was in these measurements

positioned perpendicular to the scanning direction. It might be beneficial to place it in the front of the camera, and thus creating shadows behind the UHI instead of in the measured line. Adding more lamps to the system might also be favourable.

By using the experiment in Section 3.2.1, we were not able to say anything about the magnitudes of the attenuation. We could have done an approximation and found a position in the scene where the illumination is constant over height, but the uncertainty would have been very high. What should be done is to use a collimated light source and measure its emitted light under water at different distances. A natural extension would also be to do that experiment in different water bodies.

Eventually, the attenuation model should be expanded to include backscattering, transpectral scattering, effects of particles etc. Due to the complicated nature of such effects, an analytical solution to the RTE, Equation (2.16), is not possible. The developed ray tracing simulation should be expanded to solve the RTE for each light beam traced. Other more computationally efficient methods might also exist. As mentioned earlier, a commercial software solving the RTE is HydroLight. However, it makes the assumption that the seafloor is flat, and that the lightfield is only dependant on depth. For the usage of UHIs close to the seabed, where the geometry of the seafloor may have a large impact, these assumptions do not hold.

Other classification methods than PLSDA should be tested and tweaked for the specific problem at hand. Especially, methods which is insensitive to scaling would be preferred. While we want to classify spectra on both measured irradiance and shape, we should strive for a balance. Normalizing did prove to be an effective method for classification, but this involves throwing away the information contained within the magnitude of irradiance measured. An approach where different preprocessing strategies and classification methods are combined could lead to improved overall classification. An alternative is to use PLS as a preprocessing step to extract latent variables of most covariance, and then use a classification method like support vector machines [43] to do the classification itself.

Conclusion

This thesis has been the first step towards creating a model for underwater hyperspectral images. We have proposed a model for such a system built up of three submodels: a model lamps, a model for how light attenuates in water, and a model for reflectance. Experiments were performed to determine the validity of the model and to determine model parameters. The intent of the model is to correct for the effects of attenuation, such that analysis of images taken at different distances, and between water bodies can be made.

The lamp was modelled as a point source positioned a distance l behind the real lamp. The rationale is that a real light source is more collimated than a point source. For larger l the light behaves more collimated, and for $l \rightarrow \infty$ it emits parallel rays. When estimating l from data, we found that the equation for finding it was ill-posed. However, by manually optimising l , we were able to describe the reduction of irradiance with distance accurately in air. In water, we experienced a reduction with distance which was larger than expected.

For modelling the angular distribution of the emitted light, we assumed that the lamp emits light in a Gaussian shape. Apart from the measured distribution being slightly asymmetric and more peaked than the assumed one, it is a reasonable first approximation.

Attenuation in water was modelled using the radiative transfer equations for light in water. We assumed there was no transpectral scattering involved, no internal sources, and that no stray light entered the detector. Over the distances imaged, the change in shape of the spectra due to attenuation was small, but noticeable, and the proposed model was able to partly correct for the effect of attenuation.

Reflectance was attempted modelled through the Oran-Nayar reflectance model. The reflectance from the Lego blocks measured above water seemed to follow the model with roughness $\kappa \approx 0$. When the Lego blocks were submerged under water, the reflectance from $\theta_i = 0$ and $\theta_i = 45^\circ$ were essentially similar. Hence, we

observe that the angular dependence of the reflectance changes when submerged.

Comparisons with experiments showed that the simulated spectra were qualitatively similar to the measured ones. We observed however that there were shifts to shorter wavelengths, which are believed to arise from a systematic error in the reflectance measurement.

Classification was performed using partial least squares discriminant analysis. Four training sets were defined, three on the image taken at $h = 28.5$ cm, describing areas of different illumination, and one at $h = 53.5$ cm. The choice of training set proved to be very important when it comes to prediction strength, and the classification method seemed more sensitive when classifying on brighter pixels than dimmer.

We tested two methods for improving classification: normalization and estimating reflectances. To measure if the spectra across images became more similar when estimating reflectances, we took samples from each lego block at several heights and calculated the variance. We saw a decrease in variance across all wavelengths, and qualitatively the shapes of the estimated reflectances were similar to the measured ones. Directly estimating reflectance proved to perform worse than classifying on the spectra themselves. This was probably caused by the lamp model over- and undercompensating for the illumination.

Normalization proved to be the most effective preprocessing step. This was probably due to the classification models being sensitive to variations in irradiance, and that the shape of each spectra did not change significantly. Normalization as a preprocessing method is not sufficient when doing classification in different water bodies, nor when images are taken far apart. This is due to the attenuation being different for different wavelengths. Currently, the preprocessing method which can be applied across different water bodies, and which also gives good classification, is to first estimate reflectances and then normalize. We then circumvent the problem with uneven illumination, and are able to classify between water bodies as long as the attenuation is known.

Several tools for analysing, simulating and processing hyperspectral images have been developed through this thesis, and the software will be useful in further investigation of underwater hyperspectral imaging.

Bibliography

- ¹Ecotone, *Ecotone webpage*, Available at <http://ecotone.com/> (06.07).
- ²G. Johnsen, Z. Volent, H. Dierssen, R. Pettersen, M. V. Ardelan, P. Fearn, M. Ludvigsen, and M. Moline, “Underwater hyperspectral imagery to create biogeochemical maps of seafloor properties”, in *Subsea optics and imaging* (WoodHead Publishing Ltd., Cambridge, 2013) Chap. 20.
- ³G. Johnsen, M. Ludvigsen, A. Sørensen, and L. M. Sandvik Aas, “The use of underwater hyperspectral imaging deployed on remotely operated vehicles – methods and applications”, in *Ifac-papersonline*, Vol. 49, 23 (2016), pp. 476–481.
- ⁴Marmine, *Marmine webpage*, Available at <http://pet.geo.ntnu.no/wordpress/marmine/> (06.07).
- ⁵I. H. K. Andersson, “Bio-optical diversity on cold water coral habitats”, Master’s Thesis, NTNU (2017).
- ⁶A. A. Mogstad, “Spectral characteristics of coralline algae: a multi- instrumental approach, with emphasis on underwater hyperspectral imaging”, Master Thesis’s, NTNU (2017).
- ⁷Petromaks2, *Patromaks2 webpage*, Available at <https://www.forskningsradet.no/prognett-petromaks2/Forside/1253980921324> (06.07).
- ⁸I. M. Hansen, F. Beuchel, J. Bytingsvik, S. Cochrane, R. Pettersen, S. Ekehaug, L. M. Sandvik Aas, I. Eide, C. Ballantine, P. A. Letnes, K. Hopland Sperre, and L. Tassara, *New technology and methods for mapping and monitoring of seabed habitats*, tech. rep. (2016).
- ⁹T. K. Sørensen, “Classification of Underwater Hyperspectral Images”, Project work at NTNU (2017).
- ¹⁰C. D. Mobley, *Light and Water: Radiative transfer in natural waters* (Academic Press, 1994).

-
- ¹¹R. G. Sellar and G. D. Boreman, “Classification of imaging spectrometers for remote sensing applications”, *Optical Engineering* **44** (2005).
- ¹²G. Lu and B. Fei, “Medical hyperspectral imaging: a review”, *Journal of Biomedical Optics* **19** (2014).
- ¹³D. Wu and W. Sun, “Advanced applications of hyperspectral imaging technology for food quality and safety analysis and assessment: A review — Part II: Applications”, *Innovative Food Science & Emerging Technologies* **19**, 15–28 (2013).
- ¹⁴Y. Xie, Z. Sha, and M. Yu, “Remote sensing imagery in vegetation mapping: a review”, *Journal of Plant Ecology* **1**, 9–23 (2008).
- ¹⁵K. Randolph, J. Wilson, L. Tedesco, L. Li, D. L. Pascual, and E. Soyeux, “Hyperspectral remote sensing of cyanobacteria in turbid productive water using optically active pigments, chlorophyll a and phycocyanin”, *Remote Sensing of Environment* **112**, 4009–4019 (2008).
- ¹⁶P. R. C. Fearn, W. Klonowski, R. C. Babcock, P. England, and J. Phillips, “Shallow water substrate mapping using hyperspectral remote sensing”, *Continental Shelf Research* **31**, 1249–1259 (2011).
- ¹⁷S. Phinn, C. Roelfsema, A. Dekker, V. Brando, and J. Anstee, “Mapping seagrass species, cover and biomass in shallow waters: An assessment of satellite multi-spectral and airborne hyper-spectral imaging systems in Moreton Bay (Australia)”, *Remote Sensing of Environment* **112**, 3413–3425 (2008).
- ¹⁸K. Stamnes, “The Theory of Multiple Scattering of Radiation in Plane Parallel Atmospheres”, **24**, 299–310 (1986).
- ¹⁹R. L. Fante, “Relationship between radiative-transport theory and Maxwell’s equations in dielectric media”, **71**, 460–468 (1981).
- ²⁰R. Measures, *Laser Remote Sensing: Fundamentals and Applications* (Krieger Pub.Co., 1992), p. 510.
- ²¹B. Lamond, P. Peers, A. Ghosh, and P. Debevec, “Image-based Separation of Diffuse and Specular Reflection using Environmental Structural Illumination”, *Proceedings of IEEE International Conference on Computational Photography (ICCP)* (2009).
- ²²A. Robles-Kelly and C. P. Huynh, *Imaging spectroscopy for scene analysis*, Vol. 1 (2013).
- ²³T. Weyrich, J. Lawrence, H. P. A. Lensch, S. Rusinkiewicz, and T. Zickler, “Principles of Appearance Acquisition and Representation”, *Foundations and Trends in Computer Graphics and Vision* **4**, 75–191 (2007).

-
- ²⁴C. D. Mobley, *Ocean Optics web book*, 2014.
- ²⁵M. Oren and S. K. Nayar, “Generalization of Lambert’s reflectance model”, Proceedings of the 21st annual conference on Computer graphics and interactive techniques, 239–246 (1994).
- ²⁶M. Nordmo, “Data driven analysis in oil and gas operations”, Master Thesis’s, NTNU, 19–36 (2016).
- ²⁷A. Höskuldsson, “The H-principle: New Ideas, Algorithms and Methods in Applied Mathematics and Statistics”, *Chemometrics and Intell. Lab. Syst.* **23**, 1–28 (1994).
- ²⁸I. T. Jolliffe, *Principal Component Analysis*, 2nd ed., Vol. 30, 3 (Springer, 2002), p. 487.
- ²⁹S. Wold, M. Sjöström, and L. Eriksson, “PLS-regression: A basic tool of chemometrics”, *Chemometrics and Intelligent Laboratory Systems* **58**, 109–130 (2001).
- ³⁰H. Martens and M. Martens, *Multivariate Analysis of Quality An Introduction* (John Wiley & Sons, Incorporated, 2001).
- ³¹G. J. McLachlan, “Discriminant Analysis and Statistical Pattern Recognition”, *Discriminant Analysis and Statistical Pattern Recognition*, 1–26 (1992).
- ³²A. C. Rencher, *Methods of Multivariate Analysis Second Edition* (2002).
- ³³B. K. Alsberg, D. B. Kell, and R. Goodacre, “Variable Selection in Discriminant Partial Least-Squares Analysis”, *Analytical Chemistry* **70**, 4126–4133 (1998).
- ³⁴R. G. Brereton and G. R. Lloyd, “Partial least squares discriminant analysis: Taking the magic away”, *Journal of Chemometrics* **28**, 213–225 (2014).
- ³⁵Sergios Theodoridis and K. Koutroubas, *Pattern Recognition*, 4th ed. (Academic Press, 2009), pp. 13–89.
- ³⁶P. Baldi, S. Brunak, Y. Chauvin, C. a. Andersen, and H. Nielsen, “Assessing the accuracy of prediction algorithms for classification: an overview.”, *Bioinformatics* (Oxford, England) **16**, 412–424 (2000).
- ³⁷A. Özgür, L. Özgür, and T. Güngör, “Text Categorization with Class-Based and Corpus-Based Keyword Selection”, *Proceedings of the 20th International Conference on Computer and Information Sciences* **3733**, 606–615 (2005).
- ³⁸S. Scientific., *Features of HydroLight*, Available at <http://www.sequoiasci.com/product/hydrolight/> (06.07).
- ³⁹R. M. Pope and E. S. Fry, “Absorption spectrum (380–700 nm) of pure water II Integrating cavity measurements”, *Applied Optics* **36**, 8710 (1997).
- ⁴⁰Eugene Hecht, “Refraction”, in *Optics* (2014) Chap. 4. The Propagation of Light.
-

-
- ⁴¹A. M. G. Savitzky, “Smoothing and Differentiation of Data by Simplified Least Squares Procedures A”, **36**, 1627–1639 (1964).
- ⁴²J. E. Marsden, L. Sirovich, M. Golubitsky, S. S. Antman, P. Holmes, D. Barkley, M. Dellnitz, and P. Newton, *Numerical analysis in modern scientific computing: An introduction*, 2nd ed. (2003), pp. 218–230.
- ⁴³G. Mountrakis, J. Im, and C. Ogole, “Support vector machines in remote sensing: A review”, *ISPRS Journal of Photogrammetry and Remote Sensing* **66**, 247–259 (2011).

Appendix

A Code documentation

This chapter quickly summarizes the main classes developed in this thesis. In total a couple of thousand lines have been written and several utilities for handling hyperspectral datacubes.

class HypImgs:

Class for loading hdf5 files and has basic visualization tools.

class Datacube:

Container for hyperspectral image cubes. Has the ability to hold one dataset which it can do several operations on. Some operations include PCA, basic statistics and advanced plotting features. Each Datacube instance has the ability to hold several Window instances which describe subparts of the Datacube. When a Window is set all functionality of the Datacube is performed on this Window as if the limited range described by the Window was the entire dataset.

class Window:

Each Window is a subpart of the Datacube and each Datacube contains at least one window which cover the entire dataset contained within it. Additional windows can be added if one want to look at specific sections of the dataset without copying the data over to a new Datacube. In this way a dataset only needs to be loaded once. A Window also holds statistics on the dataset in the specified range.

class Datacubes: Wrapper class around the Datacube container. Allows for convenient storage of several datasets at the same time. Operations which can be performed on a single Datacube is implemented in such a way that they can be

called from a Datacubes instance and be applied to all Datacubes within it.

class PCADatacubes: Performs PCA on a Datacube object and stores loading vectors and other important parameters related to PCA. Since a datacube may not be on the expected form for PCA PCADatacubes does checks to see whether the expected form is given and converts the data to the correct form if not.

class PLSDA: Performs PLSDA on a dataset given number of PLS components. It has functionality to fit the classification model and then predict samples later on. There is also visualization functionality for easy interpretation of the results.

class ClassGenerator: Given a Datacube object and a range it outputs samples and a response matrix which can be used to train a PLSDA model.

class PerformanceEvaluator: Contains functionality for evaluating classification performance and visualizing the result.

Class Simulation: Performs a simulation involving an UHI. It has member variables including all the classes defined below and information on resolution, which parts of the scene should be simulated etc. For each pixel in the scene simulated a lightbeam is propagated, using the Lightpropagator class, from the position of a Lightsource object to the location of the pixel. Then after reflection the lightbeam is propagated to the position of a UHI and the measured irradiance is stored.

Class UHI: Holds the measured dataset in the simulation as a Datacube instance and holds other relevant information tied to the UHI simulated.

Class Lightsource: Holds information of a particular lightsource, including spectra, spatial distribution and model parameters.

Class Lightpropagator: Propagates light from one point in space to another. Holds the propagation model and parameters related to the propagation of light.

Class Bottom: Keeps all information related to the scene simulated. It has information on the geometry and reflectances contained within it as well as functionality for light to interact with the scene.

Class HeightProfileGenerator: Generates predefined height profiles which can be passed on to a Bottom instance.

Class ReflectanceGenerator: Generates predefined reflectance profiles which can be passed on to a Bottom instance.