



LUND UNIVERSITY

Development of Models, Methods, and Materiel for Deep Tissue Imaging using Light, Ultrasound, and Spectral-Hole Burning

Hill, David

2022

Document Version:

Publisher's PDF, also known as Version of record

[Link to publication](#)

Citation for published version (APA):

Hill, D. (2022). *Development of Models, Methods, and Materiel for Deep Tissue Imaging using Light, Ultrasound, and Spectral-Hole Burning*. Atomic Physics, Department of Physics, Lund University.

Total number of authors:

1

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Development of Models, Methods, and Materiel for Deep Tissue Imaging using Light, Ultrasound, and Spectral-Hole Burning

by David Hill



LUND
UNIVERSITY

Thesis for the degree of Doctor of Philosophy

Thesis advisors: Prof. Stefan Kröll, Dr. Lars Rippe, Dr. Johannes Swartling

Faculty opponent: Dr. François Ramaz

To be presented, with the permission of the Faculty of Engineering of Lund University, for public criticism in the Rydberg lecture hall (Rydbergssalen) at the Department of Physics on Friday, 20th of January 2023 at 09:15.

| | | |
|---|--|-------|
| Organization LUND UNIVERSITY Division of Atomic Physics Department of Physics P.O. Box 118 SE-221 00 LUND, Sweden | Document name DOCTORAL DISSERTATION | |
| | Date of disputation 2023-01-20 | |
| | Sponsoring organization Spectracure AB | |
| Author(s) David Hill | | |
| Title and subtitle Development of Models, Methods, and Materiel for Deep Tissue Imaging using Light, Ultrasound, and Spectral-Hole Burning | | |
| Abstract <p>The medical imaging technique called “ultrasound optical tomography” (UOT) entails light scattered in tissue that is given spatial information by interacting with an ultrasound field, thus allowing for molecular information to be imaged with ultrasonic resolution. This thesis discusses how UOT can be implemented, with most attention put toward the implementation using spectral-hole-burning filters (SHF) in rare-earth-ion-doped crystals. The thesis further details a first principle computational model – the sequential Monte Carlo (SMC) model – used to predict and assist UOT imaging.</p> <p>Several experimental UOT trials on tissue phantoms were conducted, the first investigating the imaging depth capabilities of UOT using SHF and the validity of the SMC model in homogeneous tissue. The second trial was performed to validate the SMC model in heterogeneous media and investigate methods using the SMC model for improving the acquired image. The final experimental trial was conducted on tissue phantoms mimicking real breast tissue with tumours, investigating how the method would perform <i>in vivo</i> when imaging at the optical wavelength 794 nm with SHFs in stoichiometric $\text{Tm}^{3+}:\text{LiNbO}_3$, which displayed a 40 dB suppression of the carrier light which does not interact with the ultrasound. This final experimental trial was also a first test of a transportable UOT system using this wavelength which can be moved to a clinical setting.</p> <p>UOT using SHF could, in an investigated <i>in silico</i> scenario, image 2 cm deeper than the other UOT implementation methods discussed. The SMC model could accurately predict the strength of the UOT signal from arbitrary simulation parameters. Using the SMC model, the inverse problem for the optical absorption could be solved in the fraction of the time required by other models. $5 \times 5 \times 5 \text{ mm}^3$ large “tumours” could be imaged in 4.2 cm thick breast tissue phantoms. UOT using spectral-hole-burning filters is concluded to be a promising way forward for UOT. Furthermore, the SMC model is concluded to be a promising model for assisting UOT imaging in future endeavours, be it with SHF or another method for implementing UOT.</p> | | |
| Key words: Biomedical imaging, ultrasound optical tomography, ultrasound-modulated optical tomography, acousto-optical tomography, Monte Carlo light transport, spectral-hole-burning filters, rare-earth-ion-doped crystals, stoichiometric lithium niobate | | |
| Classification system and/or index termes (if any): | | |
| Supplementary bibliographical information: | Language English | |
| ISSN and key title: 0281-2762, Lund Reports on Atomic Physics, LRAP 588 | ISBN 978-91-8039-482-6 (print) 978-91-8039-481-9 (pdf) | |
| Recipient's notes | Number of pages 234 | Price |
| | Security classification | |

Distribution by (name and address)

I, the undersigned, being the copyright owner of the abstract of the above-mentioned dissertation, hereby grant to all reference sources permission to publish and disseminate the abstract of the above-mentioned dissertation.

Signature 

Date 2022-12-03

Development of Models, Methods, and Materiel for Deep Tissue Imaging using Light, Ultrasound, and Spectral-Hole Burning

by David Hill



LUND
UNIVERSITY

A doctoral thesis at a university in Sweden takes either the form of a single, cohesive research study (monograph) or a summary of research papers (compilation thesis), which the licentiate student has written alone or together with one or several other author(s).

In the latter case the thesis consists of two parts. An introductory text puts the research work into context and summarises the main points of the papers. Then, the research publications themselves are reproduced, together with a description of the individual contributions of the authors. The research papers may either have been already published or are manuscripts at various stages (in press, submitted, or in draft).

Cover: “Impressionist painting of a person probing another persons chest with a super-nova” DALL-E × David Hill

Funding information: The thesis work has been performed in collaboration with and co-supervision of the industry partner Spectracure AB, with the financial support of the Knut & Alice Wallenberg foundation, WACQT, and Vinnova

pp. i-121 © 2023 David Hill

Paper I © 2019 Optica Publishing Group under the Open Access Publishing Agreement

Paper II © 2021 Optica Publishing Group under the Open Access Publishing Agreement

Paper III © 2021 The Authors

Paper IV © 2022 Optica Publishing Group under the Open Access Publishing Agreement

Paper V © 2022 The Authors

Faculty of Engineering, Department of Physics

ISBN: 978-91-8039-482-6 (print)

ISBN: 978-91-8039-481-9 (pdf)

ISSN: 0281-2762

Lund Reports on Atomic Physics, LRAP - 588

Printed in Sweden by Media-Tryck, Lund University, Lund 2023



Media-Tryck is a Nordic Swan Ecolabel
certified provider of printed material.
Read more about our environmental
work at www.mediatryck.lu.se

MADE IN SWEDEN 

*Dedicated to my grandparents
Barbro, Joan, Rolf, and David*

Contents

| | |
|--|-----|
| Abstract | iii |
| List of publications | v |
| Populärvetenskaplig sammanfattning | vii |
| List of abbreviations | ix |
| List of recurring variables | xi |

I Summary of Research Output

| | |
|--|-----------|
| 1 Introduction | 1 |
| 1.1 Diagnostic imaging in tissue | 1 |
| 1.2 Diagnostic imaging using light | 3 |
| 1.3 Outline of thesis | 5 |
| 2 Light in tissue | 7 |
| 2.1 Scattering | 7 |
| 2.2 Absorption | 9 |
| 2.3 Modelling light transport in scattering media | 10 |
| 2.3.1 Light-diffusion approximation | 12 |
| 2.3.2 Monte Carlo modelling | 13 |
| 2.4 Ultrasound optical tomography | 14 |
| 3 Modelling ultrasound optical tomography | 19 |
| 3.1 Ultrasound modelling | 20 |
| 3.2 The different models of light-ultrasound interaction | 21 |
| 3.3 The sequential Monte Carlo model | 23 |
| 3.3.1 Monte Carlo generation of optical paths | 23 |
| 3.3.2 Optical signal in the absence of an ultrasonic field | 25 |
| 3.3.3 Coherent light-ultrasound interaction in homogeneously absorbing media, Paper II | 28 |
| 3.3.4 Heterogeneous absorption, Paper III | 31 |
| 3.3.5 Image formalism, Paper V | 32 |
| 3.3.6 Additional numerical considerations | 33 |
| 3.4 The inverse problem | 36 |

| | | |
|----------|---|------------|
| 4 | Detecting frequency-shifted diffuse light with spectral-hole burning | 39 |
| 4.1 | Properties of the rare-earth-ion-doped crystal | 40 |
| 4.2 | Spectral-hole-burning filters | 43 |
| 4.2.1 | The slow-light effect | 48 |
| 4.3 | Rare-earth-doped crystals for UOT | 50 |
| 4.3.1 | Praseodymium-doped yttrium orthosilicate | 51 |
| 4.3.2 | Thulium-doped stoichiometric lithium niobate | 53 |
| 5 | Comparison of methods for detecting diffuse frequency-shifted light | 61 |
| 5.1 | Overview of the methods that were compared to SHF in Paper IV | 63 |
| 5.1.1 | Single-shot and digital off-axis holography | 63 |
| 5.1.2 | Photorefractive crystal detection | 64 |
| 5.1.3 | Speckle contrast | 66 |
| 5.2 | Quantitative contrast-to-noise comparison of detection schemes | 66 |
| 5.3 | Qualitative comparison of detection schemes | 68 |
| 6 | Ultrasound optical tomography using spectral-hole burning filters | 73 |
| 6.1 | UOT signal strength in homogeneous media | 74 |
| 6.1.1 | Experimental setup | 74 |
| 6.1.2 | Absolute signal comparison (Paper I) | 77 |
| 6.1.3 | Investigating UOT tagging efficiency (Paper II) | 79 |
| 6.1.4 | Lessons from single point UOT measurements | 82 |
| 6.2 | Imaging of heterogeneous phantoms and image-reconstruction, Paper V | 84 |
| 6.2.1 | Experimental setup | 86 |
| 6.2.2 | Results and discussion | 87 |
| 6.2.3 | Lessons from first UOT imaging experiments | 91 |
| 6.3 | UOT in breast-tissue phantoms with tumour-like inclusions | 93 |
| 6.3.1 | Experimental setup | 95 |
| 6.3.2 | Images of realistic breast-tumours | 98 |
| 7 | Outlook and conclusion | 101 |
| | Author contributions | 105 |
| | Acknowledgements | 107 |
| | References | 108 |

Abstract

The medical imaging technique called “ultrasound optical tomography” (UOT) entails light scattered in tissue that is given spatial information by interacting with an ultrasound field, thus allowing for molecular information to be imaged with ultrasonic resolution. This thesis discusses how UOT can be implemented, with most attention put toward the implementation using spectral-hole-burning filters (SHF) in rare-earth-ion-doped crystals. The thesis further details a first principle computational model – the sequential Monte Carlo (SMC) model – used to predict and assist UOT imaging.

Several experimental UOT trials on tissue phantoms were conducted, the first investigating the imaging depth capabilities of UOT using SHF and the validity of the SMC model in homogeneous tissue. The second trial was performed to validate the SMC model in heterogeneous media and investigate methods using the SMC model for improving the acquired image. The final experimental trial was conducted on tissue phantoms mimicking real breast tissue with tumours, investigating how the method would perform *in vivo* when imaging at the optical wavelength 794 nm with SHFs in stoichiometric $\text{Tm}^{3+}:\text{LiNbO}_3$, which displayed a 40 dB suppression of the carrier light which does not interact with the ultrasound. This final experimental trial was also a first test of a transportable UOT system using this wavelength which can be moved to a clinical setting.

UOT using SHF could, in an investigated *in silico* scenario, image 2 cm deeper than the other UOT implementation methods discussed. The SMC model could accurately predict the strength of the UOT signal from arbitrary simulation parameters. Using the SMC model, the inverse problem for the optical absorption could be solved in the fraction of the time required by other models. $5 \times 5 \times 5 \text{ mm}^3$ large “tumours” could be imaged in 4.2 cm thick breast tissue phantoms. UOT using spectral-hole-burning filters is concluded to be a promising way forward for UOT. Furthermore, the SMC model is concluded to be a promising model for assisting UOT imaging in future endeavours, be it with SHF or another method for implementing UOT.

List of publications

This thesis is based on the following publications, referred to by their Roman numerals:

- I Characterization and modeling of acousto-optic signal strengths in highly scattering media**
A. Bengtsson, **D. Hill**, M. Li, M. Di, M. Cinthio, T. Erlöv, S. Anderson-Engels, N. Reistad, A. Walther, L. Rippe, S. Kröll
Biomedical Optics Express, vol. 10, no. 11, pp. 5565-5584 (2019)
- II Acousto-optic interaction strengths in optically scattering media using high pressure acoustic pulses**
D. Hill, A. Bengtsson, T. Erlöv, M. Cinthio, S. Kröll
Biomedical Optics Express, vol. 12, no. 6, pp. 3196-3213 (2021)
- III First principle simulation package for arbitrary acousto-optic interaction in scattering materials**
D. Hill, A. Bengtsson, T. Erlöv, M. Cinthio, S. Kröll
European Conferences on Biomedical Optics (ECBO) 2021, OSA Technical Digest, proceeding paper ES1C.6
- IV Comparison of contrast-to-noise ratios of different detection methods in ultrasound optical tomography**
A. Bengtsson, **D. Hill**, K. Shortiss, L. Rippe, S. Kröll
Biomedical Optics Express, vol. 13, no. 9, pp. 4834-4850 (2022)
- V Sequential Monte Carlo modelling of ultrasound optical tomography: image prediction and absorption reconstruction in optically diffuse media**
D. Hill, A. Bengtsson, K. Shortiss, M. Cinthio, J. Powers, L. Rippe, S. Kröll
Manuscript in preparation, estimated submission: December 2022

All papers are reproduced with permission of their respective publishers.

Populärvetenskaplig sammanfattning

“En bild säger mer än tusen ord” går ju ordspråket, och i dagens samhälle lever vi med mer bilder än någonsin i mänsklighetens historia. Vi omges av bilder på alla möjliga objekt: galaxer, semesterresor, bakterier eller den där matbilden du smällde av med din telefon senaste gången du var på restaurang.

Inte minst inom sjukvården möts vi av bilder i många olika former. Otaliga har väntat med bävan på bilder från en magnetröntgenkamera eller har efter en ultraljudsundersökning låtit sig glädjas av de första bilderna på sina barn. Men ett exempel på en fråga som de mest använda medicinska apparaterna idag; röntgen, magnetröntgen och ultraljud, inte kan svara på är: vilken färg har den där knölen inuti min höft? Anledningen till att vi inte kan svara på den frågan (utan att dra fram skalpellen) är att vävnad sprider ljus.

Som exempel på detta fenomen ber jag dig som läser detta att dra fram din telefon, slå på lampan och trycka din tumme mot den. Det du då kommer se är att hela fingret lyser upp i ett rött sken. Detta eftersom det är främst det röda ljuset som tar sig igenom då blåare färger både sprids och absorberas mer i vävnad. Något du däremot inte kommer att se är en skugga som avslöjar att din tumme innehåller ett ben. Denna information förloras på grund av att ljuset sprids så mycket att det på under en millimeter bytt riktning och börjat färdas i en helt annan riktning. Denna process raderar ut all information om var ljuset varit någonstans, och därför också skuggan av ditt ben. Så även om vävnaden, som du ser där med tummen på lampan, är *genomlyslig* så är den inte *genomskinlig*.

På sätt och vis är det tur att vävnad inte är genomskinlig, alla kanske inte vill kunna se kaffet rinna ner i magen på chefen under morgonmötet. Däremot hade det medicinska värdet varit otroligt: tänk, då hade man bara kunnat titta på en persons hjärta för att leta efter den missfärgning av muskelväggen som föregår en hjärtinfarkt. Eller kolla på den där knölen i din höft och se att det bara är en envis fettknöl. För det är det som är styrkan bakom bilder tagna med ljus: de har en otrolig förmåga att skildra medicinsk information. Syresatt blod ser helt enkelt väsentligt annorlunda ut från icke syresatt blod och fettknölar ser väsentligt annorlunda ut från tumörer.

Denna avhandling handlar om hur man ska ta optiska bilder i vävnad men samtidigt kringgå problemet med spridning, alltså det som fick skuggan av benet i din tumme att skingras. Detta görs med hjälp av ett trick där man märker ljus som passerat en speciell region inuti vävnaden. Om platsen märkningen där sker på absorberar allt utom grönt ljus kommer mer av det märkta ljuset vara grönt. Om man bara tittar på detta märkta ljus, kommer märkningsplatsen därför se grön ut. Genom att flytta på platsen där märkningen sker kan vi skapa en bild oavsett hur många gånger det märkta ljuset spridits.

Det som krävs för denna metod är då ett sätt att märka ljus inuti vävnaden, men också kunna flytta på regionen märkningen sker i på ett säkerställt sätt. Detta kan man göra med en ultraljudspuls. Till skillnad från ljus, sprids ultraljud väldigt lite i vävnad och färdas på ett förutsägbart sätt. Man kan därför skicka en ultraljudspuls till regionen man vill märka, och väl där, lyser man upp vävnaden under en kort tid, så kort att ultraljudspulsen bara hinner röra sig en väldigt liten distans under upplysningen. Eftersom vävnad är genomlyslik kommer en del av det ljuset som vävnaden lyses upp med också att nå till denna plats och interagera med ultraljudet. Denna interaktion ändrar färgen på ljuset med mindre än 0.000001% vilket, om än smått, går att mäta.

I denna avhandling görs denna mätning med hjälp av kristaller som normalt absorberar både det märkta och omärkta ljuset men som går att manipulera till att släppa igenom det märkta ljuset. Alltså, genom att först märka ljuset i vävnaden med ultraljud, samla upp allt ljus som lämnar vävnaden för att sedan leda det genom dessa manipulerade kristaller så återstår enbart det märkta ljuset och vi kan härleda färgen på vävnaden i regionen där ultraljudspulsen är. Denna metod kallas "ultrasound optical tomography" (UOT) som, helt enkelt, gör vävnad genomskinlig (men tursamt nog: inte under morgonmötet).

List of abbreviations

| | | |
|-----|---|---|
| AOM | – | Acousto-optic modulator. |
| ASE | – | Amplified spontaneous emission. |
| AWG | – | Arbitrary waveform generator. |
| CNR | – | Contrast-to-noise ratio. |
| DOT | – | Diffuse optical tomography. |
| NA | – | Numerical aperture. |
| SHF | – | Spectral-hole filtering. |
| SMC | – | Sequential Monte Carlo. |
| SNR | – | Signal-to-noise ratio. |
| UOT | – | Ultrasound optical tomography. |
| TLS | – | Tagged Light Simulator (implementation of the SMC model). |
| | | |
| exp | – | The exponential function. |
| ln | – | The natural logarithm. |
| sgn | – | The sign function. |

List of recurring variables

| | |
|-------------------------|---|
| n | – Index of refraction. |
| c, c_0 | – Phase velocity of light and speed of light in vacuum. |
| μ_a | – Absorption coefficient. |
| $\mathbf{\mu}_a$ | – Vector absorption coefficient. |
| μ'_s, μ_s | – Reduced and normal scattering coefficients. |
| g | – Scattering anisotropy. |
| D | – Diffusion constant. |
| D_m | – Length of the path of the photon with index m . |
| \mathbb{D} | – Photon path length matrix. |
| x, y, z | – Cartesian coordinates. |
| \mathbf{r} | – Spatial vector or coordinate. |
| f_C | – Frequency of the carrier light. |
| f_T | – Frequency of the tagged light. |
| Δf | – Zero centred frequency. $\Delta f = f - f_C$. |
| f_{US} | – Frequency of ultrasound wave. |
| P_1 | – Pressure wave (see Sect. 3.1). |
| \mathbf{A} | – Acoustic amplitude (see Sect. 3.1). |
| F, \mathbf{F} | – Scalar and vector SMC forward models. |
| T_C | – Transmission of carrier through crystal filter. |
| T_F | – Crystal filter transmission. |
| α_B | – Background absorption of inhomogeneous line. |
| τ_1 | – Radiative lifetime (of the excited state). |
| τ_2 | – Coherence time (of the excited state). |
| τ_h | – Lifetime of hyperfine ground state. |
| Γ_{hom} | – Homogeneous linewidth. |
| Γ_{inhom} | – Inhomogeneous linewidth. |

Part I

Summary of Research Output

Chapter 1

Introduction

Taken in its historical context, the medicine of today cannot be described as anything other than amazing. Never have so many people been able to survive complications that only a few decades ago would have been death sentences. This evolution of the medical trade is the result of the steady improvement of not only treatment methods but also diagnostics (i.e. the methods used to identify a patient's ailment). Today's diagnostic arsenal contains all manner of such methods, from blood analysis to electrocardiography. This thesis is about one such potential diagnostic method being developed by the quantum information research group in Lund, a method within the even narrower sub-field of optical diagnostic imaging.

1.1 Diagnostic imaging in tissue

Diagnostic imaging is the process of producing an image of the interior or exterior of a body, depicting some part or functionality of the imaged tissue that in turn can be used to draw diagnostic conclusions. The first such method for imaging the interior of a body, was the x-ray photograph. This method uses a part of the electromagnetic spectrum which interacts minimally with the body's soft tissues, allowing it to resolve the bones underneath. The first "medical x-ray" (see Fig. 1.1a) was of Anna Bertha Ludwig's hand and taken by her husband, the first Nobel laureate in physics, Wilhelm Conrad Röntgen.

Since then, many more diagnostic imaging methods have joined x-ray imaging, all with their own strengths and weaknesses. In the case of x-ray imaging, its weakness is that it does not image soft tissue very well, and that x-rays may damage genetic material, a typically unwanted outcome. This risk is avoided in medical




Figure 1.1: Examples of medical images taken with different techniques. (a) An x-ray image of the hand of Anna Bertha Ludwig (*anno* 1896). (b) An ultrasound image of a foetal heart [1]. (c) A magnetic resonance imaging (MRI) cross-section image of a head [2].

ultrasonography (depicted in Fig. 1.1b), where images of soft tissues are created by effectively measuring the relative difference in elasticity using sound waves. The weakness is here the opposite to x-rays: hard tissues, such as bones, stops any propagation of the sound used to create the image, limiting the use of medical ultrasonography in regions such as inside the skull.

Superior in many ways to the two previous methods is MRI, which poses no risk of damaging genetic material and allows for high-resolution imaging of soft tissues anywhere – even inside the skull (as shown in Fig. 1.1c). The method works by inducing an oscillation of the macroscopic magnetisation of a spatially select group of protons. This oscillating magnetisation in turn generates a radio wave, which can travel freely through the body and be detected with an external antenna. Not only can the density of protons be discerned by the strength of this radio wave, but measuring the decay time of this oscillation¹ further yields functional information of the tissue in which the protons are embedded (e.g. if it is cancerous or not). The weakness of MRI is instead that it is very difficult² to generate the by necessity strong and well-defined magnetic fields required for this spatially selective oscillation in magnetisation to be possible. It is also a slow method, compared to x-ray and ultrasound, as a full three-dimensional image can take over an hour to capture. This poses a challenge for using MRI in regions where the tissue moves around substantially, such as the chest.

¹ This can be done by measuring either the free induction decay or using spin echo techniques. An optical analogue to the latter – photon echo spectroscopy – is explored further in Section 4.3.2.

² Read: expensive.



Figure 1.2: Different tissue samples illustrating contrast available in optical images. (a) Vials with oxygenated (left) and deoxygenated (right) blood [3]. (b) A large cell carcinoma of the lung [4]. (c) A liver tumour [4].

1.2 Diagnostic imaging using light

In comparison to these methods, imaging with light³ has strong potential for diagnostic imaging. This is primarily due to the large amount of optical contrast (i.e. colour difference) exhibited by tissue or by matter in general. In simple terms, one can differentiate and classify a sample of tissue by simply looking at it, as when differentiating between oxygenated and deoxygenated blood (Fig. 1.2a) or classifying tumorous tissues (Fig. 1.2b-c).

This contrast available to optical images is a consequence of the energy of individual light quanta (i.e. photons) being overrepresented in the electronic transitions of atoms and molecules. A photon which energy matches such a transition thus has a chance of being absorbed, a process discussed in Section 2.2 and Chapter 4. A medium of molecules then appears green if its spectrum of energy transitions lacks the colour green (i.e. all but the green colours are absorbed). In contrast to ultraviolet and x-ray photons (the other parts of the electromagnetic spectrum where photons match atomic transitions), light photons generally do not pose any risk to genetic material and thus can be used much more freely in medical applications. The medical limit of light is instead set by the thermal effects of light in tissue (i.e. how much light that can be sent into tissue before it is burned⁴).

Thus, what is resolved with light is different atoms or molecules, allowing for light to provide the same functional imaging of tissue types as MRI. However, where MRI can only differentiate between tissues with two parameters (i.e. proton density and the decay time of the generated oscillation), different wavelengths of light will interact differently with different types of molecules. In Figure 1.3, the

³ “Light”, in this thesis, refers to the visible and near-infrared parts of the electromagnetic spectrum, wavelengths ~ 400 nm to ~ 1.4 μ m.

⁴ This limit in the near-infrared being 300 mW/cm² (average radiation) [5].

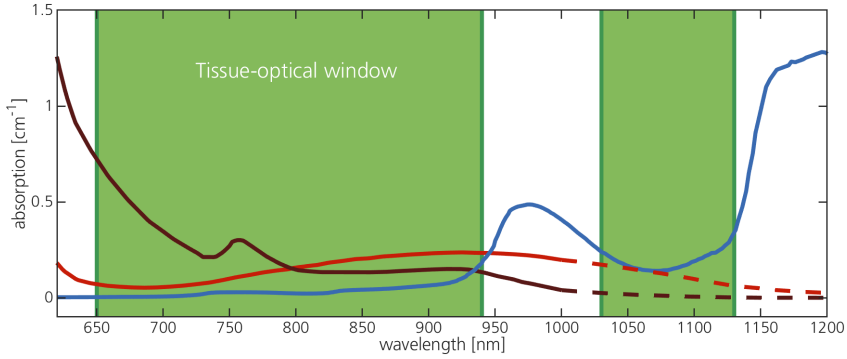


Figure 1.3: Major contributions of the absorption of tissue. The dark red line is the absorption of deoxygenated blood diluted to 3.6%, and the bright red line is the absorption of oxygenated blood diluted to 3.6% [6, 8]. The dashed lines indicate points at which the absorption has been extrapolated to higher wavelengths based on [9]. The blue line is the absorption of water. The green area indicates the position of the tissue-optical window, a spectral range where light can penetrate deep into tissue.

largest absorption contributions in muscle tissue⁵ can be seen, where oxygenated and deoxygenated haemoglobin in blood represent the majority of the absorption in the wavelength range up to ~ 900 nm, after which water becomes the dominant absorber. Blood is furthermore highly absorbing at wavelengths < 600 nm (i.e. the green, blue and violet light), which is why blood is red: in effect, all other colours are absorbed. Furthermore, the higher absorption of deoxygenated blood in the 600–800 nm range explains why it is a darker red than oxygenated blood, as seen in Figure 1.2a. Figure 1.3 further shows the so called “tissue-optical window”. This is a wavelength range of light which is minimally absorbed in tissue. The lower limit of the window is set by the strong absorption of blood, and the upper limit is set by water.

The spectra of blood and water in Figure 1.3 can be seen as a molecular fingerprint. Other molecules exhibit similar fingerprints, and this allows light to differentiate between many different molecules. Even if two molecules have the same absorption for one wavelength, such as oxygenated and deoxygenated blood near 800 nm, the so-called “isosbestic point”, they will certainly not have the same for all others. This multispectral avenue available to light is why optical imaging contains so much information – an image may “speak more than a thousand words”, but an image with colour speaks more clearly. Compared to the large and strong magnetic fields required by MRI, light is also more easily generated and detected – all one needs, technically, is a lamp and a camera. However, there is one glaring issue with using light in tissue: tissue heavily scatters light. So, while tissue is translucent in the tissue-optical window, it is far from transparent, as light is scattered to such an extent that classical imaging with cameras is limited to depths of roughly

⁵ Estimated as tissue with $\sim 4\%$ blood volume [6] and $\sim 80\%$ water volume [7].

1 mm.⁶ This is far from the practically limitless imaging depths of x-ray and MRI and the decimetres available to ultrasonography.

1.3 Outline of thesis

This thesis details one method in which spatial information is encoded in light using an ultrasound pulse, and in doing so, one can circumvent the issue of light-scattering in tissue. Chapter 2 gives an overview of how light interacts with tissue and the principles of the scattering-circumventing technique known as “ultrasound optical tomography” (UOT). Chapter 3 delves more deeply into the different ways in which light and ultrasound coherently interact in tissue and a first-principle model for how images using UOT are formed is presented. Chapter 4 discusses how the tagged signal is separated and detected using spectral holes prepared in rare-earth-ion-doped crystals, as well as the properties of these spectral holes and crystals that are relevant for UOT. In Chapter 5, this detection method is compared with others presented in the literature, both quantitatively and qualitatively.

Finally, in Chapter 6, an experimental UOT journey is presented, from the first UOT signals detected in work underpinning this thesis, to the construction and preliminary testing of a transportable UOT imaging system. Chapter 7 concludes this thesis, where the future of UOT using spectral holes is examined.

⁶ Which is why, when you look at your hand or forearm, you can see some of your subcutaneous veins (located 0.2–2 mm under the skin [10]) but no deeper lying structures.

Chapter 2

Light in tissue

To obtain an overview of how light interacts with tissue, the term “tissue” must first be defined. A specific type of tissue (e.g. muscle, adipose, nervous) is a collection or structure of cells which perform a certain function. These cells can be of different types and of different sizes and shapes: from the nerve cell, which can stretch up to a metre long, to the large fat cells of up to 300 μm in diameter and the blood cells of $\sim 7 \mu\text{m}$ in diameter. These cells are in turn comprised of smaller elements, such as organelles and proteins, which are of the order of $\leq 0.1 \mu\text{m}$. The smallest such element of interest to this thesis is the molecule. As mentioned in the previous chapter, the absorption of light is due to interactions with these molecules (or atoms). A brief overview of this process is provided in Section 2.2. When impinging on a piece of tissue however, light does not interact with it only via absorption. It also scatters.

2.1 Scattering

The scattering process between light and tissue can be either elastic or inelastic.¹ The inelastic processes include Raman scattering, where the energy shift is used in Raman spectroscopy to, among other things, determine the temperature of gases [11, 12]. Raman spectroscopy is also used to study the structures of molecules, which in turn can be used to determine pathogens or diseased tissue, though in shallow lying regions such as the skin [13]. Dominating the scattering in tissue, however, is the elastic scattering processes. One such process is Rayleigh scattering, where light scatters off particles that are of the same size as (or smaller than) the

¹ For elastic scattering, the photon energy is preserved, which is not true for inelastic scattering.

wavelength of the light. Rayleigh scattering arises when the light causes charges in these particles to oscillate, and in turn becoming small radiating dipoles which emit light with the same frequency as the incoming light. The strength of this interaction – and thus also the rate of the Rayleigh scattering – is proportional to λ^{-4} , where λ is the optical wavelength.²

As the cells themselves are ~ 10 – 100 times larger than the wavelength of light, it is not Rayleigh scattering alone that constitutes scattering in tissue. Mie scattering³ also plays a role, as it describes how light scatters against particles larger than the wavelength. This scattering process can be explained by examining the collection of cells constituting tissue as a spatial variation of the refractive index, n , where the refractive index is defined as the constant which relates the phase velocity in a medium c with the speed of light in vacuum c_0 . In effect,

$$c = \frac{c_0}{n}. \quad (2.1)$$

As shown in Figure 2.1a, a wave of light impinging on such a collection of cells is diffracted due to the cell membranes, nuclei, and cytoplasm all having different indices of refraction. As the wave propagates deeper into the tissue, the wave front quickly becomes less and less structured. The continuous diffraction that light experiences due to the tissue thus leads to an initially structured wave front (e.g. a Gaussian beam or a plane wave as in Fig. 2.1a) quickly being reduced to a random wave front. However, while random, this wave front is in the wave optics regime completely deterministic, as this scattering process – for a given distribution of cells – is fully described by Maxwell’s equations. The arbitrary shape and positions of cells, however, make solving these equations demanding, limiting the volumes where these full-wave descriptions of light in tissue are accurate to a few mm^3 [15].

Luckily, quantum mechanics teaches that light can be viewed simultaneously as both a wave and a particle. This duality is important, as a particle passing through a medium of higher and faster velocities will not have its direction changed, but a “wave packet” will, as it is allowed to interfere with itself, and thus diffract according to the same mechanics as before. A simplification (depicted in Fig. 2.1b) is then that each such wave packet, or photon, bounces around inside the tissue at discrete points (i.e. scattering points), and between these points the photons propagate along straight paths. This simplification allows for light propagation in much

² This strong dependence on the wavelength is the main reason why the sky is blue and sunrises are red: in the former case, blue light from the Sun scatters on air molecules and propagates down to the earth’s surface, and in the latter case, only the red light from the Sun is seen, as the bluer colours have scattered out into space or down to Earth.

³ Or an analogue to Mie scattering, as technically, Mie theory only concerns scattering from spherical particles [14].




Figure 2.1: Different models of light in tissue. (a) Light viewed as a wave, with a wave front that is scattered by the cells, organelles, and molecules of tissue. (b) Light viewed as wave packets or photons which can be seen as “bouncing” around inside the tissue. Each photon can be seen to contribute locally to the global wave front, and as such, the two models are equivalent outside the medium.

larger volumes to be accurately described by reducing the light-tissue interaction to only the scattering events and propagation between scattering events. Moving onto the next section however, it is important to remember that the “bouncing photon” description in Figure 2.1b, fails to fully depict the wave nature of the photons.

2.2 Absorption

Another thing that quantum mechanics teaches is that the energy of atoms can only adopt discrete energies. These are the intrinsic energies of the electronic orbitals, which in turn are the stationary solutions to the Schrödinger equation for electrons trapped in a given atomic potential. These orbitals, and thus their energies, are defined by quantum numbers. Of these, the principal and orbital angular momentum quantum numbers are used to define the electronic shell structure, and by extension, the periodic table – topics that return in Chapter 4.

In the Bohr model, light may interact with an atom, causing it to transition from one discrete energy state (i.e. orbital) to another. For two states with energies E_1 and E_2 ($E_1 < E_2$), such an interaction can occur with a single photon of frequency f , when

$$hf = E_2 - E_1 \quad (2.2)$$

where h is Planck’s constant. If the atom is in the state with energy E_1 , it can absorb a photon and “excite” to the state with energy E_2 . Similarly, if the molecule is the state with energy E_2 , it may emit a photon, either spontaneously or by stimulation, and “de-excite” to the state with energy E_1 .

In the case of only two energy states, the molecule can – after absorbing a pho-

ton – only de-excite to the lowest energy state (i.e. the ground state) by emitting an identical photon again. In reality, atomic and molecular systems contain many more states than two, and after absorbing a photon, it may de-excite to the ground state via many more mechanisms, such as by de-exciting to a state with intermediate energy and emitting two photons which energies sum to the energy of the absorbed photon. An excited molecule may also de-excite by interacting with its surrounding medium, transmitting its internal energy to kinetic energy in itself and its neighbours. This is essentially a transformation of the optical energy in the photon to heat, and when light is absorbed in tissue, heat is what the majority of the absorbed-light energy turns into. This generated heat is in turn cooled by diffusion and the flow of blood.

The transition between the same two states, as described by (2.2), may be achieved by light in a span of frequencies. This region of frequencies, the so-called “transition linewidth”, is discussed in more detail in Chapter 4. In general the less interaction a molecule has with its surroundings, the sharper these transition linewidths. Furthermore, the more complex the molecule (i.e. the more constituent atoms), the more energy states exist due to additional rotational and vibrational degrees of freedom. This is why the absorption spectra of blood in Figure 1.3 is a more or less continuous graph, as the majority of the absorption stems from the complex haemoglobin molecule, which when solved in a liquid interacts with other molecules more strongly than it could as a gas. In contrast, gases have by comparison almost discreet linewidths which become even narrower as the gas pressure decreases [16, 17].

Even if the condition in (2.2) is fulfilled, it is not guaranteed that the molecule interacts with the photon and transitions to a higher energy. This interaction probability is defined as the oscillator strength of the transition, and can be described by the quantum mechanical selection rules [16, 17].

2.3 Modelling light transport in scattering media

Regarding the two major ways in which light interacts with tissue, it is interesting to a physicist to see how this can be modelled mathematically. The model generally used for scattering in tissue is the radiative transfer equation (RTE). First formulated to analyse how energy is transferred inside stars [18], the RTE – it was quickly realised – could also be used to model energy transport in other types of media, such as tissue.

The RTE is not examined in detail here, but in short, it is fundamentally an energy-

conservation equation which balances the amount of power in a light beam along the direction \mathbf{k} . The losses to this beam per unit length are absorption and scattering away from \mathbf{k} into another direction \mathbf{k}' . The sources are emission and light in other beams propagating along \mathbf{k}'' which scatter into \mathbf{k} [14].

Unlike the full-wave models based on Maxwell's equations and the wave description of light [15, 19], the RTE treats light as stochastic beams of energy where only the average energy flow is deterministic. It uses global parameters such as the absorption rate μ_a [cm⁻¹] for the average amount of light absorbed per unit length. Similarly, the scattering rate⁴ μ_s [cm⁻¹] determines the average amount of light per unit length that is scattered away from \mathbf{k} . The transmission T of light that is neither absorbed nor scattered in a beam which propagates a distance s follows Bouguer-Beer-Lambert's law [20]

$$T = \exp(-\mu_t s) \quad (2.3)$$

where $\mu_t^{-1} = (\mu_s + \mu_a)^{-1}$ is the total attenuation rate.

The fraction of the scattered power that enters the direction \mathbf{k}' is determined by a scattering phase function $p(\theta)$, where the deflection angle θ is defined as $\cos \theta = \mathbf{k} \cdot \mathbf{k}'$, assuming $|\mathbf{k}| = |\mathbf{k}'| = 1$. Figure 2.2 shows the geometry of scattering of \mathbf{k} to \mathbf{k}' , as well as an example of the scattering phase function $p(\theta)$. In addition to a deflection in the plane of \mathbf{k} and \mathbf{k}' , the beam is also scattered in the plane perpendicular to \mathbf{k} , defining the azimuthal angle ϕ with its corresponding scattering phase function $p(\phi)$.⁵ A final useful global property that can be derived from the scattering phase function is the scattering anisotropy g . This parameter is defined as

$$g = 2\pi \int_0^\pi p(\theta) \cos \theta \sin \theta d\theta = \langle \cos \theta \rangle \quad (2.4)$$

or the cosine of the average deflection angle. The μ_s and g define the more broadly used reduced scattering coefficient $\mu'_s = \mu_s(1 - g)$. This lumped property relates to the rate at which a beam is scattered in a direction other than forward. It is also the property most-often measured in tissue spectroscopy studies, as separating μ_a , μ_s , and g is experimentally difficult, compared to only measuring μ_a and μ'_s . In this thesis, μ_a and μ'_s are measured using time-of-flight spectroscopy [21], but examples of experimental setups which can be used to measure all three parameters can be found in [22].

In general, μ_a depends on the density of the absorbing particles and the oscillator strength of the transition. Similarly, g and μ_s generally depend on the density

⁴ μ_a and μ_s are also referred to as the absorption and scattering coefficients, respectively.

⁵ This phase function is most commonly the uniform random distribution between 0 and 2π .




Figure 2.2: (a) The geometry of a scattering event in the RTE, defining the deflection angle θ and the azimuthal angle ϕ . (b) A polar plot of the probability density (i.e. scattering phase function) for scattering into the angle θ defined in (a). The function depicted in (b) is the Henyey-Greenstein function for $g = 0.3$ [23].

of the scattering particles, though g additionally depends on the structure of the scattering particles.

2.3.1 Light-diffusion approximation

If $\mu_a \ll \mu'_s$, something that is generally true for light in the tissue-optical window [6], the light propagation is said to be scattering dominated. In this case, the RTE can be substantially simplified by viewing light as a criss-cross of many beams going in every direction. For a source function S with a slow time variation, this leads to the slow amplitude diffusion approximation of the RTE used in **Paper I**. Here, the light fluence Φ is described by the light-diffusion equation [14]

$$(\mu_a - D\nabla^2)\Phi = S \quad (2.5)$$

where $D = 1/[3(\mu_a + \mu'_s)]$ is the diffusion constant. The quantity which describes the flow of light is the power flux \mathbf{J} , defined from Fick's law [14, 24, 25]:

$$\mathbf{J} = -D\nabla\Phi. \quad (2.6)$$

The slow time variation requirement of S can be translated quantitatively to that (2.5-2.6) are valid if the relative time t variation of the flux

$$\epsilon = \frac{||\mathbf{J}(t + t_\epsilon) - \mathbf{J}(t)||}{||\mathbf{J}(t)||} \quad (2.7)$$

fulfils

$$\frac{n}{c_0(\mu_a + \mu'_s)}\epsilon \ll t_\epsilon \quad (2.8)$$

for some time of interest t_ϵ (e.g. the rise time of the source).

Confusingly, Φ and \mathbf{J} have the same unit (W/m^2) and both quantities relate to a power flow. To differentiate between the two quantities, Φ can be seen as a sort of density of light, though it is technically the optical power going into and out of a given point. If one wants to calculate the power flowing through a surface, it is instead \mathbf{J} that should be turned to, where the power through the area with normal vector \mathbf{n} is

$$\iint_{\text{Area}} \mathbf{J} \cdot \mathbf{n} dA. \quad (2.9)$$

The main draw of using the diffusion approximation is the comparative ease with which (2.5) can be solved, as for some geometries even analytical solutions exist. One such geometry is the point source in a homogeneous and infinite medium, which is counterintuitively a very useful result. For such a source

$$S = P(t)\delta(\mathbf{r} - \mathbf{r}_{\text{source}}) \quad (2.10)$$

where $P(t)$ is injected power at time t and δ is the Kronecker delta function, the solution to (2.5) becomes

$$\Phi(\mathbf{r}, t) = \frac{P(t)}{4\pi D|\mathbf{r} - \mathbf{r}_{\text{source}}|} \exp\left(-\sqrt{\frac{\mu_a}{D}}|\mathbf{r} - \mathbf{r}_{\text{source}}|\right), \quad (2.11)$$

where \mathbf{r} is the spatial coordinate and $\mathbf{r}_{\text{source}}$ the position of the point source. Using this solution, solutions to geometries with infinite boundaries can be constructed by applying computational tricks such as domain mirroring, further detailed for a semi-infinite slab in **Paper I** and Chapter 6.

More complex problems with spatially varying absorption and scattering rates can be computed with the finite element method (FEM), using standard FEM solvers. This makes the light-diffusion model easy to use and analyse, though the drawbacks are that $\mu_a \ll \mu'_s$ must hold and also that the model has inaccuracies when portraying boundary effects.

2.3.2 Monte Carlo modelling

In the Monte Carlo solution to the RTE, a probabilistic method is employed by randomly propagating single computational photons.⁶ This random propagation

⁶ Note the distinction between a “real” photon and a “computational” photon. In general, there are many, many more real photons than what can be reasonably simulated in a given problem formulation, thus a computational photon is often seen as a group of co-propagating real photons.

obeys the stochastic processes described by μ_s , μ_a , $p(\theta)$, and $p(\phi)$. A single such computational photon is initialised somewhere based on the modelled source and then travels a certain distance s before encountering a scattering event. The probability density function $p(s)$ for a computational photon travelling this distance s is in turn proportional to (2.3), but for $\mu_t = \mu_s$, that is $p(s) \propto \exp(-\mu_s s)$. The fraction of the computational photon that survives this propagation distance is equal to T in (2.3), but this time $\mu_t = \mu_a$, that is $T = \exp(-\mu_a s)$. This transmission is multiplied by the W weight of the photon, which in pseudo-code can be written as $W = W \times T$. After propagating, the computational photon is then scattered into a new direction according to $p(\theta)$ and $p(\phi)$ and then propagates another distance s , again sampled using the scattering probability density function, and it then has its weight reduced even further, and so on. This continues until the photon weight is below some threshold, whereupon it is terminated. To avoid biasing the modelling, a technique called “roulette” can be employed, where there is a 1 in 10 chance that the computational photon survives termination and its weight instead increases to $W = W \times 10$. This overview is only one example of how Monte Carlo modelling can be implemented; and in Chapter 3, a more in depth description of how Monte Carlo modelling of a medium without absorption (i.e. no weights and no roulette) can be found.

Similar to other problems studied with the Monte Carlo method, this method of modelling requires many random samples, corresponding to many generated computational photons. Thus, compared to those of the light-diffusion model, solutions using the Monte Carlo method are not deterministic but instead sometimes require billions of photons to be launched in order for the solution to converge within a wanted error bound. For this reason, it is a computationally slower method. The Monte Carlo solution is however much more versatile than the diffusion approximation, as it has no issue with boundaries and it is accurate outside the regime $\mu_a \ll \mu'_s$. Furthermore, quantities such as Φ may be reverse-engineered by counting the photons that enter a given region. As such, the Monte Carlo method for solving the RTE has become the “gold standard” for modelling light propagation in scattering media and is often the model which simpler models such as the light-diffusion equation are validated against.

2.4 Ultrasound optical tomography

With an understanding of light-tissue interaction, a return to how medical imaging can be performed with light is warranted. Using only light emitted and detected at different points, it is possible to probe the regions between the sources and detectors. Compiling numerous measurements, which each using different source

detector pairs, it is possible to compute the underlying absorption and scattering distributions. This is essentially the method of diffuse optical tomography (DOT), a diagnostic imaging technique already adopted in some areas of medicine [26]. The issue with DOT is that its resolution is, as a rule of thumb, equal to 20–30% of the imaging depth [27]. It also requires good physical access for placement of the sources and detectors around the entire piece of imaged tissue in order to obtain a good transillumination.

As stated in Chapter 1, the basic issue encountered by DOT and optical imaging in tissue in general is that the scattering erases any information about where the light has been. Light which has passed through a highly absorbing part of the medium is therefore indistinguishable from light that has not. In DOT, this is circumvented by using the models in Section 2.3 to calculate the most probable paths of the light through the tissue, and from this conclude where the light has been. To circumvent this issue entirely requires either that the tissue is made non-scattering or that the spatial information is somehow encoded in light which has passed through a specific point inside the tissue. Looking only at this “tagged” light: if the tagging occurred in a region with high absorption, then little of the tagged light would be detected, and vice versa. This tagging can be performed with the aid of an ultrasound wave. This imaging method has been given various names in the literature, including “acousto-optical tomography” (AOT) [28–30], “ultrasound-modulated optical tomography” (UMOT) [31–35], and the preferred denotation in Lund and the one used in this thesis, “ultrasound optical tomography” (UOT) [36, 37].

The principle of UOT is shown in Figure 2.3. Here, light is propagating inside the tissue simultaneously insonified⁷ with an ultrasound wave. Light which passes through the insonified volume is then frequency shifted via the acousto-optic effect, which is discussed in more detail in Chapter 3. For a frequency of the injected light – in this thesis, referred to as the “carrier frequency” f_C and “carrier light” – this generates sidebands of tagged light with frequencies of $f_T = f_C + nf_{US}$, where n is an integer ($n \neq 0$) and f_{US} is the frequency of the ultrasound wave. Please note that in this thesis, the subscripts C and T are recurrently used to separate between the carrier and tagged light. As ultrasound waves are minimally scattered in tissue, the position of the insonified volume may be known in advance. By looking only at this ultrasonically tagged light (methods of which are discussed in Chaps. 4–5), information about the absorption in the insonified volume can be obtained. By moving the insonified region around, information about other regions may be gathered and used to create an image of the absorption.

In this thesis, the UOT method is subdivided into two modes. First, there is the

⁷ “To insonify” is “to flood an area or an object with carefully-controlled sound waves, typically as a part of sonar or ultrasound imaging”.




Figure 2.3: Overview of ultrasound optical tomography (UOT). An ultrasound pulse propagates inside the tissue as it is illuminated. Part of the carrier light (green) is frequency shifted by the ultrasound, giving rise to the tagged sidebands (blue and red). The collected light then contains both the carrier and the tagged light; and using a frequency selection device, the red sideband is selected. (a) The sideband signal has been captured from a first few ultrasound lines and an image has begun to be formed. (b) A few more lines have been captured and the ultrasound is passing over highly absorbing tissue; thus, no tagged light is collected (i.e. all carrier light at the ultrasound is being absorbed). This yields a shadow in the image. (c) The last signal required to form a full image has been captured and the imaging sequence process now begins anew.

reflection mode, as depicted in Figure 2.3, where the light input, collection, and ultrasound transducer are all placed on the same surface and in close proximity to one another. Second, there is the transmission-mode geometry, where the light input and output are placed on opposite sides of the investigated tissue (the position of the ultrasound in this mode is arbitrary). The pros and cons of the two modes are discussed throughout the thesis, but in short, the close proximity of all components in the reflection mode allows for this geometry to be incorporated into a handheld probe as a general-purpose triage tool. The transmission-mode geometry, in contrast, has the benefit of keeping the ratio of carrier to tagged light approximately constant. In the reflection mode, the carrier light is constant, while the tagged light is heavily dependent on the depth of theinsonified volume. This generally puts a greater demand on the method used to discriminate between the tagged and carrier light for the reflection mode, compared to the transmission mode.

Regardless of the mode used, the obtained UOT image contains spectral information, but with the resolution of the ultrasound. As ultrasound pulses may be less than a millimetre long, the potential resolution is on the same mm scale and consistent at different imaging depths, allowing UOT to surpass DOT for imaging over a centimetre into tissue. As the tagged light also traverses other parts of the tissue, the model-based methods for computing images in DOT may also be

used in UOT. Where DOT is confined to a static set of sources and detectors, however, UOT can flexibly change the position of its source by moving the ultrasound. UOT thus both offers the potential for both higher resolution and more flexible optical imaging of tissue, compared to DOT.

Chapter 3

Modelling ultrasound optical tomography

Model : a system of postulates, data, and inferences presented as a mathematical description of an entity or state of affairs

— Merriam-Webster Dictionary

Before discussing the experimental methods and concepts tied to UOT in Chapters 4-6, it is necessary to first detail the first-principle model used throughout the thesis to estimate the strength of UOT signals and simulate UOT images. The model, designated as the sequential Monte Carlo (SMC) model, was first proposed in [38] and has since been expanded to include pulsed ultrasound (**Paper II**), heterogeneous absorption (**Paper III**), and an image formalism (**Paper V**). These developments are detailed in chronological order in Section 3.3 and their experimental validations are presented in Chapter 6.

While the SMC model is the main topic of this chapter, its description is prefaced by a brief overview of the theoretical description of ultrasonic fields in Section 3.1. The predecessors of the SMC model, as well as alternate models of the light-ultrasound interaction, are then discussed in Section 3.2. Finally, in Section 3.4, a brief overview is given of how the SMC model can be used to determine the underlying and medically interesting absorption distribution in tissue by solving the so-called “inverse problem” (i.e. using the same methods for image creation in DOT). The entirety of the SMC model is implemented in MATLAB in the open-source code package named “Tagged Light Simulator” (TLS) [39].

3.1 Ultrasound modelling

As “ultrasound optical tomography” starts with “ultrasound”, so must the modelling. The pressure P in a medium is defined as the ambient pressure P_0 plus some spatially and time-varying field, $P_1(\mathbf{r}, t)$, that is $P = P_0 + P_1$. The scenario of interest in this thesis is when $P_0 = 1$ atmosphere and P_1 is a wave generated by a piston (i.e. an ultrasound transducer) driving an oscillating movement at the medium boundary that propagates into the medium. This pressure wave is interconnected with a movement of the particles in the medium, called the “acoustic amplitude” \mathbf{A} . The pressure and this movement are, in an incompressible and non-dissipative fluid, described by Euler’s equations [40] as

$$\nabla \cdot \mathbf{v} = 0 \quad (3.1)$$

$$\rho \left[\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla) \mathbf{v} \right] + \nabla P = 0 \quad (3.2)$$

where \mathbf{v} is the fluid velocity and ρ the density. For a solution to (3.1–3.2), the acoustic amplitude is calculated as follows:

$$\mathbf{A}(\mathbf{r}, t) = \int_{-\infty}^t \mathbf{v}(\mathbf{r}, \tau) d\tau. \quad (3.3)$$

A wave solution of these equations can be derived by assuming that $P_1(\mathbf{r}, t) \ll P_0$ as the nonlinearity exhibited by (3.2) can be linearised in this regime, yielding $\mathbf{v} = \mathbf{K}(\rho C)^{-1} P_1$ for a plane wave P_1 where C is the ambient speed of sound and \mathbf{K} is the propagation direction of P_1 . However, as will be seen, the light-ultrasound interaction increases with the pressure; and thus, to achieve a strong UOT signal, the highest medically permissive pressure is desirable. These pressures are often in the order of a few MPa, which is >10 atmospheres. This means that these waves produce negative “pressures”¹ (i.e. less than a vacuum), which is possible as the ultrasound-induced movement of the medium creates regions where the vapour pressure acts more unidirectionally. However, this is only a meta-stable state: if the conditions which allow this negative pressure to exist are upheld for too long, or if the negative pressure is too low, the surrounding medium will vaporise and form small bubbles. This is “cavitation”, which in tissue leads the rupture and destruction of cells. Unsurprisingly, it is therefore the magnitude of the negative pressure and the length of time it is maintained that sets the medical safety limit. For the peak negative pressure P_- , the figure of merit is the mechanical

¹ Using the generalised description of pressure, where it is defined as the force created by particle collisions acting on an infinitesimal sphere. Due to phase transitions of matter, this definition allows the pressure to be negative [40].

index $MI = \frac{P_- [\text{MPa}]}{\sqrt{f_{US} [\text{MHz}]}}$, where the medical limit is $MI \leq 1.9$, according to the FDA [41].

In the high-pressure regime, linearisation breaks down, and in its place, more exotic phenomena begin to occur. One of these phenomena in the high-pressure wave regime is that the peaks of the pressure wave propagate at faster velocities than the troughs. For an initial harmonic and continuous waveform, the faster positive pressure propagation warps the waveform towards a sawtooth shape, and in doing so, transfers power from the fundamental ultrasound frequency f_{US} to higher-order harmonics $n f_{US}$, where n is an integer. This transfer continues until the peaks approach the troughs and a shock wave is formed. **Paper II**,² investigated the time scales and ultrasound propagation distances where the spectral components of measured high-pressure ultrasound pulses can be assumed constant. Here, *ansatz* solutions to (3.1-3.2) were also investigated together with the limits of the low-pressure approximation $\mathbf{v} = \mathbf{K}(\rho C)^{-1} P_1$.

3.2 The different models of light-ultrasound interaction

With a description of the ultrasound in place, modelling of the coherent light-ultrasound interaction in turbid media is possible. This modelling is not new and has accompanied the field since the conception of UOT [30, 35, 38, 42–52]. With the avenues of light-ultrasound interactions in scattering media mapped out in [30, 42, 43, 52], the models for UOT were split into subcategories: those using the light-diffusion approximation for the optical transport, and those using the Monte Carlo model. The first diffusion-based model was presented in [44] and has since been used in multiple investigations [45–50]. The diffusion models can be further subcategorised into those that attempt to model the light-ultrasound interaction using first-principle modelling [45–47, 52], and those that take the shortcut of describing the tagged and carrier fields using the coupled light-diffusion equations (see Sect. 2.3.1) [48–50]; in effect,

$$\begin{cases} (\mu_a - D\nabla^2)\Phi_C &= S_C \\ (\mu_a - D\nabla^2)\Phi_T &= \kappa(\mathbf{r})\Phi_C. \end{cases} \quad (3.4)$$

The shortcut approach is the one used in **Paper I**. This is flexible and easily applied but has the drawback of requiring the *a priori* κ , which is the amount of light converted from the carrier field Φ_C to the tagged field Φ_T at point \mathbf{r} . This shortcut model can only be used if κ is measured first, as in **Paper I**, or by simplifying the

² Especially in Appendix A of **Paper II**

ultrasound field to a ball and estimating κ using Raman-Nath theory [36, 48, 49, 53]. Thus, the shortcut model is a poor choice if one wants to investigate the limits of UOT resolution or the design space of the ultrasound pulse without experimental feedback.

The other subcategory of diffusion-based models addresses this issue by using the light-diffusion model to find the probability density of path lengths and how these path lengths interact with an ultrasound. This approach to modelling the UOT signal has historically been simplistic, as these distributions can only be known analytically for specific geometries. Furthermore, a thorough examination of the literature yielded that only continuous plane wave ultrasound [43–46, 54], and plane waved pulsed ultrasound [47] has been investigated with this model. While the path distributions may be calculated numerically for more complex geometries, experimentally used ultrasound pulses are usually complex, with fully three-dimensional ultrasonic pulse envelopes and frequency contents exhibiting multiple harmonics. Implementing such complex ultrasound field descriptions in the first-principle diffusion model is seemingly challenging.

In contrast, the Monte Carlo-based model [35, 38, 51, 52] is a direct interpretation of the light-ultrasound interactions on individual computational photons. As shown in both the literature and this thesis, the Monte Carlo-based models accurately predict the light-ultrasound interaction for arbitrary acoustic and optical fields. The main drawback with these models, however, is that each such photon’s path through the medium and its interaction with the ultrasonic field must be simulated individually, a task that is vastly more computationally expensive than using the light-diffusion approximation. Historically, this has been additionally worsened by the fact that the light-ultrasound interaction is simply tacked onto the Monte Carlo light propagation, thus computational effort is wasted on photons which are not detected. This is addressed in the SMC model by effectively swapping computational time for greater memory usage. The comparatively slow computational speed of the Monte Carlo-based models has also been a major drawback when using the model for the inverse problem, discussed further in Section 3.4. As will be shown, this computational slowness can be addressed for a subset of inverse problems using the SMC model.

In short, if computations are sequenced in the Monte Carlo model, the computational expensiveness can be delegated outside the inverse problem, allowing for rapid reconstruction of absorption distributions – more rapid even than solutions to the inverse problem using the shortcut-diffusion models. This benefit was of course not yet known when UOT was first started to be modelled in the presented thesis work. A question is then why Monte Carlo-based models were chosen over the diffusion models. The question posed in **Paper II** is how the pulse shape affects

the UOT signal, a question the shortcut-diffusion model is inadequate to answer. A Monte Carlo-based model was pursued over the first-principle diffusion model simply because the latter, at the time of writing **Paper II**, seemed less intuitive and required a larger amount of theoretical legwork. At the time of writing this thesis, this opinion has not changed.

3.3 The sequential Monte Carlo model

3.3.1 Monte Carlo generation of optical paths

The SMC model begins with the generation of computational photons in the medium which then propagate according to Monte Carlo light transport. This generation is done without absorption included, and thus, all initialised photons must terminate by exiting the computational domain. As each photon is terminated, its scattering points define a path through the medium, from the injection point to the point at which the photon leaves the domain. To avoid overwhelming the memory capability of the used machine, paths which do not terminate in a predefined detection area are discarded.

There are many software packages that can perform the actual computation of the paths [55–57]. These packages each have their own strengths and weaknesses, with none really meeting the needs of the SMC model without heavy reworking of its source code – the needs of the SMC model being no termination of photons mid-flight and the saving of each terminated photon’s path. As such, a new software package, Multiphoton, was specifically written to be compatible with the SMC model. This package was validated against the CUDAMCML package [57].

As each computational photon requires several MBytes of memory to store the paths during flight, running on the GPU³ with thousands of parallel computations incurs large overheads, as when the GPU memory fills up it requires constant offloading transfers to slower memory caches (at least on the machines that have been used in the presented thesis work). Thus, Multiphoton, in contrast with modern Monte Carlo simulators, runs on the CPU⁴ instead of the GPU.

As briefly outlined in Section 2.3.2, in Monte Carlo light transport [58], computational photons are injected in a source distribution that then propagates through a computational domain, a domain which may be homogeneously or heterogeneously scattering. In Multiphoton, the medium is heterogenised by the definition

³ GPU: “graphical processing unit” – a computer’s graphics card.

⁴ CPU: “central processing unit” – a computer’s processor.

of a set of planar faces \mathbf{X} , \mathbf{Y} and \mathbf{Z} in the x , y and z Cartesian directions, where the first and final face in each set defines the external boundaries of a cuboid-shaped computational domain. If each set contains an integer $E_{x,y,z}$ number of faces, this domain contains $V = (E_x - 1)(E_y - 1)(E_z - 1)$ cuboid voxels, where a voxel denotes a three-dimensional pixel (i.e. “volume pixel”). The scattering coefficient μ_s and scattering anisotropy factor g can, in this model, be defined as the three-dimensional matrices $\mathbf{\mu}_s$ and \mathbf{g} , where each element corresponds to the scattering and anisotropy in the voxel between its six surrounding faces in x , y and z (see Fig. 3.1). The azimuthal angle ϕ is defined as a uniformly random variable $\phi \in]0, 2\pi]$. The deflection angle is most commonly calculated using the Henyey-Greenstein scattering phase function [23], which takes the scattering anisotropy g as one of its input parameters. When computing a path using the Monte Carlo method, it is not the actual phase function $p(\theta)$ that is of interest, but the function $\theta(\chi)$ – in effect, the function which transforms the uniformly random variable $\chi \in]0, 1]$ to a θ that obeys $p(\theta)$. For the Henyey-Greenstein scattering phase function, this function is defined as

$$\cos \theta(g) = \frac{1}{2g} \left\{ 1 + g^2 - \left[\frac{1 - g^2}{1 - g - 2g\chi} \right] \right\}. \quad (3.5)$$

After a scattering event, the new propagation direction $\mathbf{k}' = (k'_x, k'_y, k'_z)$ is a function of the previous direction \mathbf{k} , ϕ and θ [58], that is

$$\begin{aligned} k'_x &= \frac{\sin(\theta)}{\sqrt{1 - k_z^2}} [k_x k_z \cos(\phi) - k_y \sin(\phi)] + k_x \cos(\theta) \\ k'_y &= \frac{\sin(\theta)}{\sqrt{1 - k_z^2}} [k_y k_z \cos(\phi) + k_x \sin(\phi)] + k_y \cos(\theta) \\ k'_z &= -\sin(\theta) \cos(\phi) \sqrt{1 - k_z^2} + k_z \cos(\theta) \end{aligned} \quad (3.6)$$

which, when k_z is close to one, can be exchanged for

$$\begin{aligned} k'_x &= \sin(\theta) \cos(\phi) \\ k'_y &= \sin(\theta) \sin(\phi) \\ k'_z &= \text{sgn}(k_z) \cos(\theta) \end{aligned} \quad (3.7)$$

to avoid numerical errors. Between two scattering events, the photon traverses a subset of voxels with indices \mathbf{I}_i , travelling a distance s_i in each voxel and obeying [59]

$$\ln(\chi) = \sum_i \mathbf{\mu}_s(I_{x,i}, I_{y,i}, I_{z,i}) s_i \quad (3.8)$$




Figure 3.1: (a) A single or (b) many path(s) generated with Algorithm 1 (implemented in Multiphoton in TLS [39]) in a medium with six voxels. The computational photons were injected in a point source in the upper-left (green triangle) and detected in a circular detection area in the bottom-right (red circle). In the centre-left voxel, the scattering is high. In the bottom-left voxel, the scattering is low, while the other voxels have intermediate scattering. The black lines indicate the position of the faces of the six voxels, with three faces along x , two faces along y , and four faces along z .

where the $\chi \in]0, 1]$ is again a uniform random variable. An example of how these criteria are fulfilled – and how simulated photon transport without absorption can be realised – is shown in Algorithm 1.

To simulate the light-ultrasound interaction, Algorithm 1 is repeated until M paths terminating in a predefined detection area have been generated, discarding paths which do not. In this way, a total of Z computational photons are initialised. Thus, barring absorption, the transmission from the injection point to the detection area is M/Z . Examples of paths generated this way can be seen in Figure 3.1.

3.3.2 Optical signal in the absence of an ultrasonic field

The optical source is defined by both its spatial distribution and its temporal distribution (i.e. its electric-field amplitude in space and time).⁵ In time, this field is assumed to be a coherently oscillating field subject to some amplitude modulation and can thus be described by the real part of the complex phasor,

$$\mathcal{E}(t) = E_o(t) \exp(i2\pi f_C t) \quad (3.9)$$

where $E_o(t)$ is the complex amplitude and f_C is again the optical carrier frequency. The power is assumed to divide equally over all Z initialised computational photons, thus the field applied along each path is \mathcal{E}/\sqrt{Z} . The division by \sqrt{Z} is required, as the intensity $I \propto |\mathcal{E}|^2$. If the speed of light c , and the average path distance for a photon in the tissue $\langle D \rangle$ yields a time $t_\epsilon = \langle D \rangle / c$ which fulfils

$$E_o(t) \approx E_o(t + t_\epsilon) \quad (3.10)$$

⁵ The magnetic-field component of light is ignored because, compared to the electric-field component, it interacts very weakly – if at all – with tissue.

Algorithm 1 Pseudo-code algorithm for generation of a single propagation path in a heterogeneously scattering media with scattering matrix $\boldsymbol{\mu}_s$ and scattering anisotropy matrix \mathbf{g} , whose elements correspond to the properties of voxels with indices $\mathbf{I} = [I_x, I_y, I_z]$.

- 1: Initialise an empty data structure \mathbf{R} and initialise the counting variable $j = 1$.
 - 2: Initialise a computational photon by drawing the position \mathbf{r} and direction \mathbf{k} from a defined source distribution (e.g. a collimated top hat).
 - 3: From \mathbf{r} , calculate the current voxel index \mathbf{I} .
 - 4: **while** (\mathbf{r} is inside the computational domain) **do**
 - 5: $\mathbf{R}(j) = \mathbf{r}$.
 - 6: $j = j + 1$.
 - 7: $\xi = \ln(\chi)$, χ is uniformly random in $]0, 1]$.
 - 8: **while** ($\xi > 0$) **do**
 - 9: Calculate the distance s to the next voxel \mathbf{I}' in direction \mathbf{k} .
 - 10: **if** [$\boldsymbol{\mu}(\mathbf{I})s > \xi$] **then**
 - 11: $s = \xi / \boldsymbol{\mu}(\mathbf{I})$.
 - 12: $\xi = 0$.
 - 13: **else**
 - 14: $\xi = \xi - \boldsymbol{\mu}(\mathbf{I})s$.
 - 15: $\mathbf{r} = \mathbf{r} + \mathbf{k}s$, $\mathbf{I} = \mathbf{I}'$.
 - 16: $\mathbf{k} = \mathbf{k}'\{\mathbf{k}, \varphi, \theta[\mathbf{g}(\mathbf{I})]\}$.
 - 17: Calculate the position \mathbf{r}' where the photon exited the computational domain.
 - 18: $\mathbf{R}(j) = \mathbf{r}'$, export \mathbf{R} .
-

any time-delay effects due to paths having different lengths can be ignored.⁶ For the optical fields examined in this thesis, (3.10) is assumed true.⁷ The expectation value of the total detected field phasor, \mathcal{E}_{det} , is the sum of all path's individual electric field phasors \mathcal{E}_m :

$$\mathcal{E}_{\text{det}}(t) = \sum_{m=1}^M \mathcal{E}_m(t) \quad (3.11)$$

$$\mathcal{E}_m(t) = \frac{E_o(t)}{\sqrt{Z}} \exp \left[i(2\pi f_C t + \varphi_m) - \frac{D_m \mu_a}{2} \right] \quad (3.12)$$

where D_m is the length of path m , and φ_m is the accrued phase of this path. The absorption is here included according to the Bouguer-Beer-Lambert law: in effect (2.3) but $\mu_t = \mu_a$. For a medium with homogeneous refractive index, n_0 , the phase φ_m is defined as $\varphi_m = k_C n_0 D_m$, k_C being the optical wave number of the carrier in vacuum. From this, the expectation value of the power spectral density at the detection area \mathcal{I}_{det} is defined as

$$\mathcal{I}_{\text{det}}(f) = \frac{1}{2} c_0 \epsilon_0 \left| \mathcal{F} \left[\sum_{m=1}^M \mathcal{E}_m(t) \right] (f) \right|^2 \quad (3.13)$$

where ϵ_0 is the vacuum permittivity and $\mathcal{F}(\cdots)(f)$ denotes the Fourier transform. If each path is subject to many scattering events, the phases φ_m are expected to be uncorrelated and uniformly distributed. This, together with the linearity of the Fourier transform, allows for the summation in (3.13) to be broken out of the $|\cdots|^2$ brackets for large values of M [38, 60], yielding

$$\mathcal{I}_{\text{det}}(f) = \frac{1}{2} c_0 \epsilon_0 \sum_{m=1}^M |\mathcal{F}[\mathcal{E}_m(t)](f)|^2. \quad (3.14)$$

By introducing the zero centred frequency $\Delta f = f - f_C$, (3.14) can be further manipulated to obtain the slightly simpler expression:

$$\mathcal{I}_{\text{det}}(f) = \frac{c_0 \epsilon_0}{2Z} \sum_{m=1}^M \left| \mathcal{F} \left\{ E_o(t) \exp \left[i\varphi_m - \frac{D_m \mu_a}{2} \right] \right\} (\Delta f) \right|^2 \quad (3.15)$$

where the factor $\exp(i2\pi f_C t)$ present in (3.12) has been removed using the frequency translation properties of the Fourier transform.

⁶In essence the same criteria as the slow time variation requirement discussed in Section 2.3.1.

⁷The validity of this assumption can be examined by putting some typical values into this criteria, for example the electric-field rise time of 1 μs and $\langle D \rangle = 1$ m. This yields the relative error between the two sides of (3.10) as $\sim 3 \times 10^{-3}$.

Equation (3.15) describes the temporal power spectral density over the entire detection area. Since the light emitted from tissue is not spatially coherent, the light in the detection area is distributed in a so-called “speckle” pattern of spatially varying intensities. As done in **Paper IV** and Chapter 5, the intensity of a single speckle in this speckle pattern can be estimated as $\mathcal{I}_{\text{speckle}}(f) = (A_{\text{speckle}}/A_{\text{det}})\mathcal{I}_{\text{det}}(f)$, where A_{det} is the detection area and A_{speckle} the speckle area.

The light-ultrasound interaction is now described by evaluating how φ_m varies in time in the presence of an ultrasonic field. Before delving deeper into this topic however, a second assumption is required in addition to the criteria in (3.10): the ultrasound-induced changes in the scattering and absorption rates can be neglected. This interaction is incoherent and does not yield any frequency shift. Instead, due to the modulation of the scattering rate, the intensity is modulated as light moves more slowly or more quickly through media with higher and lower scattering rates, respectively – differences which in turn are induced by the compression and rarefaction of the ultrasonic field. This incoherent effect has been studied theoretically in the context of imaging [61, 62] but, to the best of the author’s knowledge, has not yet been measured experimentally. This lack of experimental results are explained by this incoherent interaction being very weak [52], and if it has been experimentally measured, it is not evident in the literature due to the general lack of distinction between the coherent and incoherent light-ultrasound interactions in turbid media. Due to its apparent weakness, the interaction has been routinely ignored in the literature when modelling coherent light-ultrasound interactions [35, 38, 52] and will, without further motivation, be ignored here as well.

3.3.3 Coherent light-ultrasound interaction in homogeneously absorbing media, Paper II

Each of the M paths generated using Algorithm 1 is defined by a set of $J_m + 1$ points: $\mathbf{R}_m = \{\mathbf{r}_{m,0}, \mathbf{r}_{m,1}, \dots, \mathbf{r}_{m,J_m}\}$. The first coherent light-ultrasound interaction in turbid media that is modelled is that these points move around their equilibrium⁸ according to the acoustic amplitude \mathbf{A}

$$\mathbf{r}_{m,j}(t) = \mathbf{r}_{m,j} + \mathbf{A}(\mathbf{r}_{m,j}, t). \quad (3.16)$$

⁸ Except the first and last, as these points are not scattering points but only indicate the points at which the computational photons entered and left the computational domain.




Figure 3.2: Depiction of how an optical path is affected by the passage of a single cycle ultrasound pulse: dark green, high pressure; light green, low pressure. (a) The ultrasound has yet to interact with the path. (b) The leftmost scattering point is dragged along by the high pressure, changing the path slightly. Simultaneously, the part of the path in the low-pressure region experiences a lower index of refraction, while the high-pressure region experiences a higher index of refraction, as indicated by the thinner and thicker lines. (c) The leftmost scattering point has returned to equilibrium, while it is the rightmost scattering point's turn to be dragged by the high pressure, which again modulates the optical path. The movement and oscillating index of refraction cause a frequency shift in the light which travels along this path.

This leads to the free-path length between two subsequent scattering points

$$L_{m,j}(t) = \|\mathbf{r}_{m,j}(t) - \mathbf{r}_{m,j-1}(t)\|, \quad \sum_{j=1}^{J_m} L_{m,j}(t) = D_m(t) \quad (3.17)$$

varying in time. In the presence of an ultrasound field, φ_m thus oscillates due to a total oscillation of the optical-path length.

The second coherent light-ultrasound interaction in scattering media arises from the ultrasound induced compression and rarefaction of the medium. This modulates the refractive index n , which in turn yields a phase modulation. The interaction between the refractive index and the pressure is called the piezo-optic effect and is quantitatively described as

$$n(\mathbf{r}, t) = n_0 \left[1 + \frac{\partial n}{\partial P} P_1(\mathbf{r}, t) \right] \quad (3.18)$$

where $\frac{\partial n}{\partial P}$ is the piezo-optic coefficient. For the simulations of light-ultrasound interaction in real and phantom tissues presented in this thesis, $\frac{\partial n}{\partial P}$ is assumed to be the same as water (i.e. $\frac{\partial n}{\partial P} = 1.466 \times 10^{-10} \text{ Pa}^{-1}$ [63]).

The two coherent light-ultrasound interactions are visualised in Figure 3.2. Com-

binning them yields the time-varying phase $\varphi_m(t)$ for path m as

$$\varphi_m(t) = k_C \sum_{j=1}^{J_m} \int_0^{L_{m,j}(t)} \mathbf{n}[\mathbf{x}_{m,j}(t, l), t] dl \quad (3.19)$$

$$L_{m,j}(t) = \|\mathbf{r}_{m,j}(t) - \mathbf{r}_{m,j-1}(t)\| \quad (3.20)$$

$$\hat{\mathbf{k}}_{m,j}(t) = [\mathbf{r}_{m,j}(t) - \mathbf{r}_{m,j-1}(t)] / L_{m,j}(t) \quad (3.21)$$

$$\mathbf{x}_{m,j}(t, l) = \hat{\mathbf{k}}_{m,j}(t)l + \mathbf{r}_{m,j-1}(t) \quad (3.22)$$

where L is the free-path length, \mathbf{k} the path direction, \mathbf{x} the free path length coordinate, and $l \in [0, L_{m,j}(t)]$.

Inserting (3.19) into (3.15) now yields the power spectral density spectrum in the presence of an ultrasound pressure pulse $P_1(\mathbf{r}, t)$ with accompanying acoustic amplitude $\mathbf{A}(\mathbf{r}, t)$. As the perturbed total path distance $D_m(t) = D_m + \Delta_m(t)$, $\Delta_m(t) \ll D_m$, by neglecting the perturbation $\Delta_m(t)$, the absorption can be extracted outside the Fourier transform which in turn allows the expression to be simplified to the following

$$\mathcal{I}_{\text{det}}(f) = \frac{c_0 \epsilon_0}{2Z} \sum_{m=1}^M |\mathcal{S}_m(\Delta f)|^2 \exp(-D_m \mu_a) \quad (3.23)$$

$$\mathcal{S}_m(\Delta f) = \mathcal{F} \{E_o(t) \exp[i\varphi_m(t)]\}(\Delta f) \quad (3.24)$$

Some other resulting spectra $|\mathcal{S}_m|^2$ are seen in Figure 3.3, where simulation parameters are the same as those used in **Paper III**, discussed in the next subsection. While some paths can be seen to interact substantially with the ultrasound, it is important to note that most do not. In **Paper II**, the tagged and carrier signals are estimated by the power in a given frequency band f_a – f_b , in effect

$$F = \int_{f_a}^{f_b} [\mathcal{I}_{\text{det}}(f)] df \quad (3.25)$$

where F is defined as the forward model for the UOT measurement of the sideband in f_a – f_b .

3.3.4 Heterogeneous absorption, Paper III

To expand the model to describe heterogeneous absorption requires spatial information to be encoded in (3.23–3.24). Returning to the Bouguer-Beer-Lambert




Figure 3.3: Spectra of four paths perturbed by an ultrasound pulse. (a) Spectra of a path that did not interact with the ultrasound – this is what the majority of the spectra look like. (b–d) Spectra of paths that increasingly interact with the ultrasound – to the point that in (d), the second order sideband is the dominant frequency.

law (2.3), in a passage through two consecutive regions (A and B) with absorption coefficients μ_A and μ_B , the total transmission T through the regions can be seen as the transmission first through A and then B, in effect:

$$T = \exp(-D_A \mu_A) \exp(-D_B \mu_B) = \exp[-(D_A \mu_A + D_B \mu_B)]. \quad (3.26)$$

In the discretisation of the medium performed in the Monte Carlo generation of the paths, an absorption is assigned to each of these voxels (i.e. $\mu_a \rightarrow \mathbf{\mu}_a$). As outlined in (3.26), the transmission of path m can be written as

$$T_m = \exp\left[-\sum_{I_x} \sum_{I_y} \sum_{I_z} \mathbf{D}_m(I_x, I_y, I_z) \mathbf{\mu}_a(I_x, I_y, I_z)\right] \quad (3.27)$$

where \mathbf{D}_m is a three-dimensional matrix containing the total length of path m in each of the V voxels, again defined by the set of faces \mathbf{X} , \mathbf{Y} , and \mathbf{Z} . $\mathbf{\mu}_a$ is the matrix of the same size but instead containing the absorption in each voxel. \mathbf{D}_m is calculated by ray-tracing each path and noting the total distance travelled in each voxel. As indicated in (3.26), if a path returns to a voxel with index \mathbf{I} , the additional distance travelled in this voxel can simply be added to the current value in $\mathbf{D}_m(\mathbf{I})$. Furthermore, if a coarser voxel geometry is desired after calculating \mathbf{D}_m , a reduced \mathbf{D}'_m can be calculated by joining voxels together and summing their corresponding elements in \mathbf{D}_m . This allows for arbitrary voxel geometries to be created or the absorption to be linked in spatially disconnected areas. This second property can be used for faster simulation of the signals acquired with some *a priori* information, such as an ultrasound image, where different types of tissues can be classified and superimposed on the simulated geometry. Similarly, a more finely resolved absorption geometry can be found for the same paths by adding more faces in the sets \mathbf{X} , \mathbf{Y} , and \mathbf{Z} . The SMC model is updated to account for heterogeneous-based absorption by simply exchanging the factors $\exp(-D_m \mu_a)$ in (3.23) with the corresponding T_m from (3.27). Furthermore, the structure of




Figure 3.4: Signal from the UOT scan example for different inclusion absorptions $\mu_{a,\text{inc}}$ and pulse y_n positions. (a) The simulated first sideband forward model F_n . (b) The same data is shown as in (a), but with each value normalised with $F_n(\mu_{a,\text{bcg}})$, (i.e. normalised to the forward model when the medium is homogeneous, highlighted in red in [a]). Reprinted with permission from **Paper III** © The Authors.

the SMC model allows for computationally cheap editing of the absorption, as neither the spectra \mathcal{S}_m nor \mathbf{D}_m needs to be recalculated when μ_a is changed. This is a significant advantage, as calculating these entities entails a considerably larger computational expense than the final step of summing all spectra in (3.23).

This modelling of the heterogeneous absorption was first presented in **Paper III**. Here, the tagged signal was simulated in a reflection-mode UOT setup where the tissue is illuminated when the ultrasound is at a fixed depth of $z = 2$ cm, but with varying lateral positions y_n ($n = 1, 2, \dots, N$) such that the forward model F in (3.25) could be evaluated multiple times with the ultrasound at the Cartesian coordinates $\mathbf{r}_n = (0, y_n, 2)$ [cm], denoting each such evaluation with the subscript n (i.e. $F \rightarrow F_n$). The ultrasound pulse P_1 was a 6 MHz centre-frequency pulse with a 1.5 mm full width at half max pressure envelope. The medium was divided into two regions, one a $5 \times 5 \times 5$ mm³ large voxel at $(x, y, z) = (0, 0, 2)$ cm, and the other a background voxel making up the rest of the $12 \times 12 \times 12$ cm³ large medium. The medium had a homogeneous scattering with $g = 0.7$ and $\mu'_s = 5$ cm⁻¹. The background voxel had $\mu_{a,\text{bcg}} = 0.2$ cm⁻¹ but the inclusion absorption $\mu_{a,\text{inc}}$ was successively changed from 0.2 cm⁻¹ to 0.6 cm⁻¹. The light was injected as a point source at $(-1.5, 0, 0)$ cm and detected in a circular area with 1 cm diameter and centre at $(1.5, 0, 0)$ cm.

The results of this simulation, as published in **Paper III**, can be seen in Figure 3.4a. To more easily discern the effect of the inclusion, in Figure 3.4b, each simulation of different inclusion absorptions $\mu_{a,\text{inc}}$ is divided by the simulation $\mu_{a,\text{inc}} = \mu_{a,\text{bcg}}$.

3.3.5 Image formalism, Paper V

Instead of the cumbersome three-dimensional matrices \mathbf{D}_m and $\mathbf{\mu}_a$, a formalism using vectors is preferred from a simplicity standpoint, transforming \mathbf{D}_m to the row vector \mathbb{D}_m with V elements, such that $\mathbb{D}_{m,v} = \mathbf{D}_m[\text{key}(v)]$, where the function “key” maps the voxel number v to the coordinate index \mathbf{I} . With the same treatment, $\mathbf{\mu}_a \rightarrow \mathbb{\mu}_a$, though here $\mathbb{\mu}_a$ is defined as a column vector. Now defining \mathbb{D} as the $M \times V$ sized matrix, where each row m of \mathbb{D} is equal to \mathbb{D}_m , the transmission column vector $\mathbf{T} = (T_1, T_2, \dots, T_M)^T$ is defined as

$$\mathbf{T} = \exp(-\mathbb{D}\mathbb{\mu}_a). \quad (3.28)$$

where \exp here denotes the element-wise application of the exponential function.

This simplifies the treatment of the absorption as (3.28) is essentially (3.27) but for calculating all M transmissions simultaneously. But how to simplify the light-ultrasound interaction? For measurement n and path m , a spectrum $\mathcal{S}_{n,m}(\Delta f)$ can be calculated according to (3.24). The index $n = [1, 2, \dots, N]$ is added here to differentiate between N different measurements performed in the same tissue and optical input and output: in effect the same paths can be used to simulate light-ultrasound interaction for different ultrasound configurations. Examples of such configuration differences, or measurements, are illumination of tissue with an ultrasound pulse at different positions, as in **Paper V**, or illumination of tissue insonified with ultrasonic plane waves with different propagation angles [64].

The m -th path’s signal strength of measurement n , $a_{n,m}$, can now be defined as follows:

$$a_{n,m} = \int_{f_a}^{f_b} |\mathcal{S}_{n,m}(\Delta f)|^2 df. \quad (3.29)$$

Basically, $a_{n,m}$ is the area under a sideband in Figure 3.3, or the contribution of path m to the total tagged signal contained between the frequencies f_a - f_b . For a given measurement n , the forward model F_n calculated using (3.25), can be expressed as

$$F_n = \frac{c_0 \epsilon_0}{2Z} \sum_m^M a_{n,m} T_m = \mathbf{A}_n \exp(-\mathbb{D}\mathbb{\mu}_a) \quad (3.30)$$

where the row vector \mathbf{A}_n is defined as $\mathbf{A}_n = \frac{c_0 \epsilon_0}{2Z} [a_{n,1}, a_{n,2}, \dots, a_{n,M}]$. If many measurements are performed (i.e. an image is taken), the image (i.e. the vector) forward model \mathbf{F} is defined as

$$\mathbf{F} = \mathbf{A} \exp(-\mathbb{D}\mathbb{\mu}_a) \quad (3.31)$$

where \mathbf{A} is the matrix where each row n corresponds to the vector \mathbf{A}_n . This image or matrix formulation is the final form of the SMC model for light-ultrasound interaction, and as well as being simple, has other properties beneficial for solving the inverse problem discussed in Section 3.4.

3.3.6 Additional numerical considerations

There has already been discussion of some numerical considerations, such as the Monte Carlo path-generation algorithm and how the discretisation of the medium is performed by sectioning it into V voxels. However, this geometry of cuboid voxels is a choice left to the implementer, as tetrahedron-shaped voxels could be used just as well, as in the VALOMC Monte Carlo simulator [55]. In addition to the previously discussed computational aspects of the SMC model, two other numerical aspects are of importance: (1) how the Fourier transform when calculating $\mathcal{S}_{n,m}(\Delta f)$ is implemented and (2) how the line integral in (3.19) is implemented.

Starting with the Fourier transform, it is preferably implemented using the fast Fourier transform (FFT) by first discretising the time into S points, such that $t \rightarrow t_s = s/f_{\text{sample}}$ where $s = 0, 1, \dots, S - 1$. The integer S and the sampling frequency f_{sample} are then chosen to be sufficiently large that the time of interest is covered by t_s and aliasing of significant sidebands in $\mathcal{S}_{n,m}(\Delta f)$ is avoided. To make the FFT algorithm as efficient as possible, S should be chosen as a power of two. Furthermore, for energy conservation (c.f. Parseval's theorem), an additional factor $1/f_{\text{sample}}$ must be included in the FFT. Overall, the numerical implementation of (3.24) becomes

$$\mathcal{S}_{n,m}(\Delta f) \rightarrow \mathcal{S}_{n,m}(\Delta f_s) = \frac{1}{f_{\text{sample}}} \text{FFT}\{E_o(t_s) \exp[i\varphi_{n,m}(t_s)]\}(\Delta f_s) \quad (3.32)$$

where $\Delta f_s = (s/S - 1/2)f_{\text{sample}}$ is again centred around zero.

To complete this calculation, the optical envelope E_o and the phase $\varphi_{n,m}$ need to be sampled S times. The optical envelope is given by whatever light source being simulated, (e.g. a Gaussian pulse or constant amplitude).⁹ The phase, however, must be evaluated at each time point t_s . The movement of the scattering points $r_{m,j}(t_s)$ is calculated using (3.16), where \mathbf{A} is calculated from a pressure field P_1 using (3.1–3.3). P_1 is in turn either an analytic fit of a measured ultrasound pulse,

⁹For constant amplitude illumination, which has not yet been implemented in TLS [39], instead of using FFT, the short-time Fourier transform (STFT) can be applied [65]. This yields a spectrogram, instead of the single-time spectrum of the FFT. Here, some care needs to be placed on the choice of window function though, as this represents a trade-off between the spectrogram containing non-physical sidelobes or preserving energy conservation.

as in **Papers II** and **V**, or a hypothetical field (e.g. the Gaussian pressure pulse in **Papers III–IV**). From the calculated scattering positions $\mathbf{r}_{m,j}(t_s)$ is calculated from which $L_{m,j}(t_s)$ and $\hat{\mathbf{k}}_{m,j}(t_s)$ follow from (3.20) and (3.21), respectively. The free-path length coordinate $\mathbf{x}_{m,j}(t_s, l)$ in (3.22) is discretised into Q segments, such that $l \rightarrow l_q = qL_{m,j}(t_s)/Q$, where $q = 0, 1, \dots, Q$. The integrals in (3.19) are then estimated numerically using the trapezoidal method, in effect:

$$\int_0^{L_{m,j}(t_s)} \mathbf{n}[\mathbf{x}_{m,j}(t, l), t] dl \rightarrow \frac{L_{m,j}(t_s)}{2Q} \sum_{q=1}^Q \{ \mathbf{n}[\mathbf{x}_{m,j}(t_s, l_{q-1}), t_s] + \mathbf{n}[\mathbf{x}_{m,j}(t_s, l_q), t_s] \}. \quad (3.33)$$

where the refractive index \mathbf{n} is defined in (3.18). In **Paper II**, it is shown that the effect of the movement of the scattering points can be neglected and still achieve a good agreement with experimental data. The light-ultrasound interaction can then be simulated using only the refractive index change, which is a faster computation than when including the effects of \mathbf{A} . Furthermore, as the pressure wave P_1 used to calculate \mathbf{n} can be directly interpolated, the absence of \mathbf{A} removes the need to fit analytical expressions to measured pulses. Using an interpolated P_1 in (3.1–3.3) to calculate \mathbf{v} will, due to numerical errors, yield residual velocities in the wake of the ultrasound pulse, which would be unphysical.

The discretisation of l this way leads to a non-uniform step length in each of the terms of (3.19), where the median step-length is $\ln(2)/(\mu_s Q)$. Choosing Q thus entails relating this median step-length to the wavelength of the simulated ultrasound. In the example in **Paper III**, the wavelength of the ultrasound pulse is $\sim 25 \mu\text{m}$. For the scattering properties used in that example, the choice of $Q = 16$ yields a median step length an order of magnitude shorter than this.

Accelerating computations

The speed of the simulations is dependent on choosing as small as reasonably possible values for f_{sample} , S , and Q . What follows is an outline of a practical method for choosing these values. First, for the problem that is examined, generate a few hundred paths and calculate the spectra $\mathcal{S}_{n,m}(\Delta f)$, using high values for f_{sample} , S , and Q . Then discard the paths which show no sideband interaction. Then further discard all but the five paths with the smallest non-zero sideband peaks and also the five paths with spectra containing the largest sideband peaks (i.e. the paths with the five largest and five smallest $a_{n,m} \neq 0$ values). Re-calculate $\mathcal{S}_{n,m}(\Delta f)$ for these 10 paths with decreasingly large f_{sample} , S , and Q – each time examining

the resolution in comparison with the first simulation. When reaching a point at which these spectra begin to change dramatically, increase the values of f_{sample} , S , and Q slightly above this point and run the simulation for the examined problem in full (i.e. a large M). Prior to this, double-check that the found values of f_{sample} , S , and Q do not breach the Nyquist criteria for the examined sidebands, that t_s covers the time of interest, and that the median step-length is short in comparison with the ultrasound wavelength.

Another way of accelerating the simulation is to realise that not all free-path coordinates $\mathbf{x}_{m,j}$ interact with the ultrasound. Those free paths which have no interaction whatsoever may, in fact, be removed from the summation in (3.19), as they only add a time-constant term to the phase, which does not affect the spectral contents of $|\mathcal{S}_{n,m}|^2$. By setting an estimated distance \mathcal{L} , past which a path is not expected to interact with the ultrasound, a smaller subset of scattering points $\mathbf{R}'_m \in \mathbf{R}_m$ may be used instead by discarding all scattering points in \mathbf{R}_m that are further away than \mathcal{L} . This distance can be defined to a specific point (e.g. the position of the ultrasound at the time the tissue is illuminated by a light pulse), or the distance to a line (e.g. the path along which the ultrasound travels). Calculating whether a scattering point $r_{m,j}$ or $\mathbf{x}_{m,j}$ is within \mathcal{L} is a simple exercise in linear algebra¹⁰ that is significantly less computationally complex than (3.32) and (3.33).

3.4 The inverse problem

The SMC model has been referred to as a “forward model” because it takes a set of input parameters (e.g. μ_a and P_1) and projects this forward to produce an estimated signal value. This is useful when estimating things such as possible signal strengths in UOT, as done in **Paper IV**. However, a forward model can also be used to solve the so-called “inverse problem”. Here, instead of estimating the signal levels from a given set of input parameters, one attempts to estimate the input parameters from a given set of signal levels. For UOT, this input parameter is either the scattering matrix μ_s or the absorption matrix μ_a , as all other input parameters (e.g. P_1 , E_o) or the position of the light detection can be identified from other sources.¹¹

The inverse problem is usually defined using a scalar valued objective function $\Omega(\mathbf{F}, \mathbf{b})$, which compares the forward model $\mathbf{F}(\mu_s, \mu_a)$ with, for example, an experimentally obtained image \mathbf{b} . This objective function is then defined as having

¹⁰ This is implemented in `interaction_prediction.m` in the tagged light simulator package.

¹¹ Technically, the anisotropy g is also an unknown, but since it is quite consistently equal to ~ 0.9 in tissue [66], it is not normally solved for in the inverse problem.

a global minimum of $\mathbf{b} \equiv \mathbf{F}(\bar{\boldsymbol{\mu}}_s, \bar{\boldsymbol{\mu}}_a)$, where the goal of the inverse problem is to find $\bar{\boldsymbol{\mu}}_s$ and $\bar{\boldsymbol{\mu}}_a$. The simplest version of this goal function is $\Omega = \|\mathbf{F} - \mathbf{b}\|$, which is also the goal function used in **Paper V**. However, usually $N \ll V$, and thus the inverse problem is grossly underdetermined, which in the case of UOT means there are many equivalent $\bar{\boldsymbol{\mu}}_s$ and $\bar{\boldsymbol{\mu}}_a$. This can be somewhat alleviated by the regularisation of the objective function by, for example, limiting the values that the input parameters can take, penalising deviations from a mean value, or penalising high spatial gradients of the input parameters. For UOT, regularisation thus allows some preconceived notions to be forced on $\bar{\boldsymbol{\mu}}_a$ and $\bar{\boldsymbol{\mu}}_s$ – for example: $\bar{\boldsymbol{\mu}}_a, \bar{\boldsymbol{\mu}}_s$ should be larger than zero, $\bar{\boldsymbol{\mu}}_a, \bar{\boldsymbol{\mu}}_s$ should not differ too much from their mean values (which can be measured using DOT), or $\bar{\boldsymbol{\mu}}_a, \bar{\boldsymbol{\mu}}_s$ should be spatially smooth. In addition to solving some issues, regularisation introduces others, as excessive use of it may impose too many of these preconceived notions upon $\bar{\boldsymbol{\mu}}_a, \bar{\boldsymbol{\mu}}_s$, which – in extreme cases – means that $\bar{\boldsymbol{\mu}}_a$ and $\bar{\boldsymbol{\mu}}_s$ can be whatever solution the user wants. Too little regularisation, however, and the inverse problem continues to have equivalent solutions.

These solutions are typically found by iterating between values of $\bar{\boldsymbol{\mu}}_a$ and $\bar{\boldsymbol{\mu}}_s$ and evaluating Ω for these values. In this iteration process, the forward model is in turn evaluated multiple times, and thus the speed at which the inverse problem can be solved is determined by the speed of the forward model that is used. As a result, the SMC model is both ill- and well-suited to the UOT inverse problem, as calculating the forward model for a new $\bar{\boldsymbol{\mu}}_s$ is very computationally expensive, given that both the \mathbf{D} and \mathbf{A} matrices must be recalculated. A change in the absorption, however, only requires a new evaluation of (3.27) for a new $\bar{\boldsymbol{\mu}}_a$. The resolution of the inverse problem to reconstruct $\bar{\boldsymbol{\mu}}_a$ in phantom tissue is explored in **Paper V** and is further discussed in Section 6.2 of this thesis.

Chapter 4

Detecting frequency-shifted diffuse light with spectral-hole burning

A key aspect of performing a UOT measurement is the separation of the frequency-shifted (i.e. tagged) light from the unshifted carrier. Although there are many techniques for separating the tagged light from the carrier (some of which are discussed in Chap. 5), the filtering technique of greatest importance to this thesis – and which produced all the experimental results presented in Chapter 6 – is the spectral-hole filtering method, using inorganic crystals doped with rare-earth ions. This is a quite cumbersome designation, and as such, this method is often referred to as simply “spectral-hole filtering” (SHF) or, when the slow-light effect (to be discussed later) is used, “slow-light filtering”.

As the name suggests, this method employs rare-earth elements that, in low concentrations, are doped into inorganic crystalline materials. The rare earths are a group of elements (highlighted in Fig. 4.1a), consisting of the 15 lanthanides (i.e. La–Lu in Fig. 4.1), with yttrium and scandium typically included as well. The first rare earths were discovered from an analysis of the mineral gadolinite mined in Ytterby, Sweden – a town which gave the names to four of the rare earths: ytterbium, yttrium, erbium, and terbium. Despite the nomenclature “rare”, these elements are fairly common in the earth’s crust, where even the two “rarest” (Tm and Lu) are 200 times more abundant than gold. However, they do not tend to concentrate into ores (i.e. pure, as in Fig. 4.1c), making extraction expensive [67].

The other necessary ingredient for creating a spectral-hole filter is inorganic crys-




Figure 4.1: (a) The rare earths' position in the periodic table. (b) A silicon crystal being grown by the Czochralski process. (c) Pure thulium. (d) A polished $\text{Tm}^{3+}:\text{LaF}_3$ crystal.

talline material. This exists in abundance, with sapphire being one example. When used for application in science or technology, these crystals are made synthetically, primarily using the Czochralski or Bridgman-Stockbarger methods. The Czochralski method (depicted in Fig. 4.1b) involves slowly drawing a seed crystal out from a melt of the desired crystalline components. The cooling melt deposits onto the seed crystal, growing a rod with the same crystal structure [68]. The resulting rod of crystalline material (i.e. boule) is then cut and polished to the desired crystal shape, such as the $\text{Tm}^{3+}:\text{LaF}_3$ crystal in Figure 4.1d.

4.1 Properties of the rare-earth-ion-doped crystal

Doping a crystal is as simple as adding a small amount of the dopant (e.g. Tm in the case of Fig. 4.1d) to the melt. As the boule grows, some ions in the ordinary crystalline structure will be replaced by the dopant ions. This process is made easier if the dopant ion is of similar size and has similar chemical properties to the replacement ion, as in the case of $\text{Eu}^{3+}:\text{Y}_2\text{SiO}_5$, where one rare earth (Y) is replaced by another (Eu). As the dopant levels are usually low (less than 1%), the macroscopic properties of the crystal (e.g. index of refraction, thermal properties, and hardness) remain largely unaffected. The dopant ions themselves, however, act as a trapped dense gas due to the special properties of the rare-earth electron configuration.

The electron configuration of the lanthanides can be written as $[\text{Xe}] 4f^a 5d^b 6s^2$ for some integers, $a \geq 1$ and $b \geq 0$, where a $4f$ electron is always removed when a lanthanide is ionised [69]. This valence $4f$ orbital, however, has a radial distribution inside that of the filled $6s$, $5s$, and $5p$ orbitals. These outer orbitals effectively act as a Faraday cage for the $4f$ orbital, shielding it from the environment. The crystal




Figure 4.2: (a) Energy level structure of the 4f orbital of a dopant ion, with levels typically denoted with the Russell-Sanders notation $^{2S+1}L_J$, where S is the total electron spin, L the total orbital angular momentum, and J the total electronic angular momentum. Energy separation of different levels are on the order of 100 THz, the crystal field split on the order of 100 GHz and the hyperfine split on the order of 10 MHz. (b) Each ion absorbs with an homogeneous linewidth. Due to the disrupted crystal lattice, the crystal field split in (a) will differ for different ions. The total ensemble of all ions in the crystal absorbs over a large frequency range.

field also disturbs the orbital of the ion,¹ such that the normally disallowed dipole transitions within the 4f shell become weakly allowed (i.e. transitions within the levels in Fig. 4.2a). While seemingly not hugely significant, transitions within the shielded 4f shell – together with the spatial density of ions in the crystal – make rare-earth-doped crystals a unique system for optical manipulation.

To start with the shielded 4f–4f transitions, as the transition is only weakly allowed, it is noted that the excited state in any such transition is comparatively long-lasting, with radiative lifetimes τ_1 . This gives the potential of very sharp homogeneous linewidths Γ_{hom} for the optical transition, as

$$\Gamma_{\text{hom}} = \frac{1}{\pi\tau_2} \geq \frac{1}{2\pi\tau_1}. \quad (4.1)$$

Equality between the centre and right-hand side of (4.1) is achieved when the coherence time of the excited state τ_2 fulfils $\tau_1 = 1/2\tau_2$.² If this relation is true, Γ_{hom} is said to be “lifetime limited”.

Although the 4f shell is shielded, τ_1 is still reduced due to acoustic quanta (i.e. phonons), which bridge the gaps between levels in Figure 4.2a, causing more rapid

¹ This can be seen as the crystal field mixing in an opposite parity configuration to the wave function which describes the electron cloud (e.g. mixing in 5d into the 4f configuration).

² Another commonly used notation in the literature for the radiative and coherence state lifetimes is T_1 and T_2 .

decay from the excited state. The phonon-induced decay can be heavily reduced by cooling the crystal to a temperature at which the crystal lattice will support very few such phonons (i.e. <4 K). This cooling is achieved by placing the crystal in a cryostat. Even when efforts are made to cool the crystal, the transition is still often not lifetime limited due to other effects – such as spin flips – which limit the coherence of the upper state. Nonetheless, due to the shielding of the 4f orbital and the weakly allowed 4f–4f transition, some of the narrowest optical linewidths ever measured in a solid were of a 4f–4f transition in a rare-earth-doped crystal (specifically $\text{Er}^{3+}:\text{Y}_2\text{SiO}_5$ and $\text{Eu}^{3+}:\text{Y}_2\text{SiO}_5$ [70, 71]).

The reduction of phonons also increases the lifetime of the lowest-lying hyperfine states,³ as the ions are otherwise very well-protected from the environment due to the shielded 4f orbital. This leads to some materials having exceptionally long hyperfine lifetimes (e.g. $\text{Eu}^{3+}:\text{Y}_2\text{SiO}_5$, where a lower limit of several weeks has been measured [72]). This means that if an ion is placed in a specific ground state, it will stay there for a long time, given that the crystal remains cold. As the hyperfine states are typically only separated by ~ 10 MHz, they are even at a few K equally likely to be populated in the Boltzmann distribution [73].

The next section details why it is worth the trouble of using cryostats to achieve a narrow Γ_{hom} and long hyperfine lifetimes. Before that, there is one more property of the rare-earth ion-crystal that is important to mention: the inhomogeneous broadening. As the dopant ion is not of the exact size of the atom it is replacing, a slight irregularity is introduced to the crystal lattice (as seen in Fig. 4.2b). This irregularity produces a spatially heterogeneous crystal field, which in turn slightly affects the energy of the crystal field splitting. Thus, two ions in the lattice will have slightly different transition energies due to its position, while still preserving its homogeneous linewidth properties. When the crystal is examined as a whole, the absorption of all the dopant ions builds up a ~ 10 GHz-wide inhomogeneous absorption profile (Γ_{inhom}) for a given optical transition. Due to the density of the ions, the absorption of the inhomogeneous profile can be incredibly strong, where attenuations of over 100 dB are possible in crystals that are only 1-2 cm long.

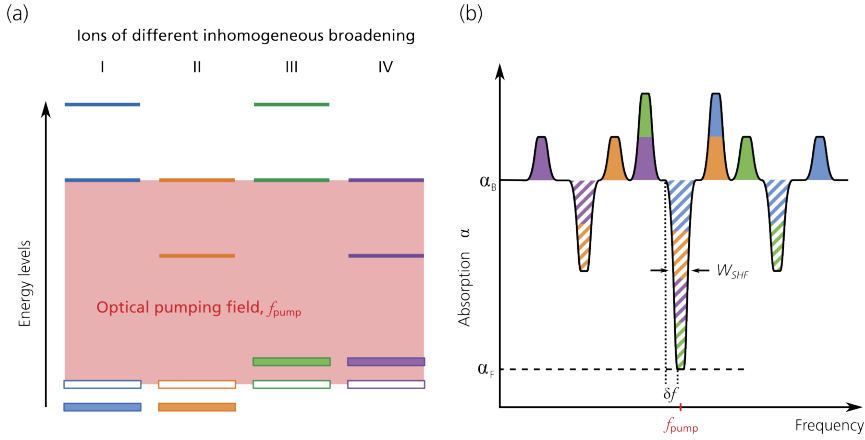


Figure 4.3: Overview of the spectral-hole burning in a 2×2 -level system. As the inhomogeneous broadening is larger than the hyperfine split, all transitions between the two ground states and two excited states are driven by the pumping field in different ion classes I–IV, as seen in (a). This gives rise to a main spectral-hole at f_{pump} in (b), as well as several side-structures. The anti-holes with more absorption arise from the other hyperfine state, now hosting up to twice the number of ions. The side holes arise from the empty states and now cannot be excited to the other excited state either. The colours in (b) indicate which ions in (a) are involved in the different structures.

4.2 Spectral-hole-burning filters

Properties of a spectral hole

A spectral hole is created by optically pumping ions between the hyperfine ground states in which energy separation is hf_{pump} (see Sect. 2.2) and where f_{pump} is the pumping frequency. Optical pumping is a process in which the ion is repeatedly excited by light and the excited ions then have a chance to de-excite to any of the other ground states. The probability of the ion de-exciting to a specific ground state is the “branching ratio”. For the example system given in Figure 4.3a, where the ground state and excited state have two hyperfine states each, an equal probability of de-exciting to any ground state would give a branching ration of 1/2 for every excited state. In reality, a spin projection in the excited and ground states is often strongly preserved, granting a branching ratio skewed heavily toward de-excitation to the original ground state of the ion [74]. This means that optical pumping requires many cycles in order to create a spectral hole.

Due to the inhomogeneous broadening, the same optical wavelength drives all four possible transitions between the ground states and the excited states (though in different ions). As seen in Figure 4.3b, pumping at one frequency thus gives rise

³ Meaning the lowest crystal field and hyperfine states in Figure 4.2a.

to additional structure in the inhomogeneous profile outside of the main spectral hole. These structures are called “side holes” or “anti-holes”, depending on whether they have less or more absorption than the equilibrium absorption. For a system in which the excited state and ground state have different hyperfine splitting, there are $2(n_g - 1)$ side holes, n_g being the number of hyperfine ground states. The number of anti-holes is more complex, and the number of anti-holes there are for a given n_g and number of excited states n_e is given in Table 4.1.

Table 4.1: Number of anti-holes produced by a single optically pumped spectral hole for the number of excited states n_e and ground states n_g in the used transition, assuming arbitrary splitting of said states. Some anti-holes may be degenerate if the splitting of the ground state or upper state are the same, or if multiple levels are equidistant.

| $n_g \setminus n_e$ | 1 | 2 | 3 | 4 | 5 |
|---------------------|----|----|-----|-----|-----|
| 2 | 2 | 6 | 14 | 26 | 42 |
| 3 | 6 | 18 | 42 | 78 | 126 |
| 4 | 12 | 36 | 84 | 156 | 252 |
| 5 | 20 | 60 | 140 | 260 | 420 |

When the ions are pumped in a frequency region, a wider spectral-hole (or filter) can be created. This spectral-hole filter has a few properties, outlined in Figure 4.3b and set by the material properties. The maximum width W_{SHF} of the spectral filter is set by the difference between the minimum frequency separations in the excited and ground hyperfine states. The filter falloff δf – or “sharpness” – is set by the Lorentzian profile of the optical transition; thus, $\delta f \propto \Gamma_{\text{hom}}$. The transmission $T(f)$ at any point on the profile is given by Bouguer-Beer-Lambert’s law:

$$T(f) = \exp[-\alpha(f)L] \quad (4.2)$$

where $\alpha(f)$ is frequency dependent absorption⁴ and L is the crystal length. It is here useful to define the filter transmission $T_F = T(f_{\text{pump}})$ and the carrier transmission, $T_C = T(f_C)$. The equilibrium or background absorption α_B is determined by the density of dopant ions and the oscillator strength of the targeted transition. The internal filter absorption $\alpha_F = \alpha(f_{\text{pump}})$ is set, in part, from the absorption of ions with centre frequencies outside the hole but whose Lorentzian absorption profile stretches into the hole. α_F is also contributed to ions that have not been optically pumped. This remaining population depends on a mix of the efficiency in the optical pumping, the branching ratio, and the hyperfine-state lifetime.

⁴ For historical reasons, α is the symbol used when referring to absorption in these contexts.

Temporal evolution of a spectral-hole

These dependencies can be explored by examining the following rate equations for a two ground-state system and a single excited state:

$$\begin{cases} E' = -\frac{1}{\tau_1}E \\ N' = \frac{R}{\tau_1}E + \frac{1}{\tau_h}(M - N) \\ M' = \frac{1-R}{\tau_1}E - \frac{1}{\tau_h}(M - N) \end{cases} \quad (4.3)$$

where N is the population of the pumped state, M the other hyperfine-state population, and E the excited-state population ($N + M + E = 1$). τ_h is the hyperfine-state lifetime and R is the branching ratio. The population evolution after pump pulse k is given by solving (4.3), with the initial conditions

$$\begin{cases} E_k(t_{0,k}) = (1 - \zeta)E_{k-1}(t_{f,k-1}) + \zeta N_{k-1}(t_{f,k-1}) \\ N_k(t_{0,k}) = (1 - \zeta)N_{k-1}(t_{f,k-1}) + \zeta E_{k-1}(t_{f,k-1}) \\ M_k(t_{0,k}) = M_{k-1}(t_{f,k-1}) \\ N_0 = M_0 = 0.5, E_0 = 0 \end{cases} \quad (4.4)$$

where ζ is the population transfer efficiency in the optical pumping pulse, $t_{0,k}$ is the start time of the evolution after pump pulse k , and $t_{f,k}$ is the time point just before the $k+1$ pumping pulse. This model for the pumping is quite simplistic as it assumes instantaneous transfer between the states, that the transfer efficiency is independent of the number of ions in N and E , and that no off-resonant transfer occurs between the M and E populations. The model also only accounts for a single ion class in the crystal lattice. As the same optical pulse targets different transitions, ζ will differ depending on the ion class. In practice, this means that some ion classes will be more difficult to optically pump than others.

Despite these issues, (4.3) and (4.4) can be used to understand the evolution of the filter during pumping and how the filter evolves after the last pumping pulse. $\tau_{\text{pump}} = t_{f,k} - t_{0,k}$ is the time between optical pumping pulses and should be sufficient to perform the transfer pulse and wait for the excited state to empty. As a rule of thumb, $\tau_{\text{pump}} \approx 4\tau_1$ to allow for a short and coherent transfer pulse and for the excited state to empty before the next transfer. With N determined, α follows from $\alpha \propto N$. A crystal with the background absorption α_B gives $\alpha = 2\alpha_B N$ (the 2 in the expression coming from that the background absorption is given at the equilibrium population $N = 0.5$). The transmission through a crystal is given by (4.2). As an example, the transmission of a filter from $\zeta = 0.05$ and 10^4 pumping pulses can be studied from a solution to (4.3) and (4.4). Jumping ahead a bit, one sees that the parameters $\tau_1 = 150$ us, $\tau_h = 30$ s, and $R = 0.9$ are taken




Figure 4.4: Rate-equation solutions to the spectral hole-filter transmission, T_F , for different parameters in (4.3) and (4.4). The blue traces in each subfigure are created using the same parameters.

to be similar to the 794 nm transition in $\text{Tm}^{3+}:\text{LiNbO}_3$ used to create the filters in Section 4.3.2. The impact of individually tuning ζ , R , and τ_h is shown in the subfigures of Figure 4.4. The results show that a smaller R , larger ζ , and longer τ_h are all desirable for both creating filters faster and ensuring that the filter is long lived.

All three parameters affect the maximum transmission. As the proportional changes of R and ζ in Figures 4.4a and 4.4b are the same, the similarity between the two indicates that a high R can be compensated by increasing the transfer efficiency in each pumping cycle. However, due to the high optical density of the crystal, initially very little of the population is transferred with each pulse, as the pumping “digs” its way through the crystal. While this digging can be made faster by increasing the optical power used, this also causes power broadening [75, 76], which increases δf in Figure 4.3. Thus, decreasing R is preferable. This may be achieved by introducing an electromagnetic radio frequency field over the crystal which matches the hyperfine splitting of the excited state. If this field is driven strongly enough, one can achieve a 50:50 population distribution of the excited ions by effectively “scrambling” the excited state. This, in turn, forces a branching ratio of 0.5 (in the case of a two level hyperfine state). A more sophisticated approach is to flip the excited state populations using a radio wave π -pulse – as done in [77] – which effectively inverts the branching ratio. An appropriately oriented magnetic field of the right magnitude may be used in a similar way to either remove or increase the preservation of the spin projection in the transition, which in turn changes the branching ratios [74, 77, 78].

There is no similar trade-off for τ_h . A longer hyperfine lifetime is always preferable, especially as it increases the window in time in which the hole has high transmission after it has been created. This lifetime can be extended by decoupling a transition mechanism between the different hyperfine states. A common way of

doing this is to apply a magnetic field, which has the added benefit of narrowing Γ_{hom} by reducing interaction between the ions in the lattice [79–81]. To achieve the longest lifetimes, a specific alignment of the magnetic field in relation to the crystal axes is often required [78]. Furthermore, τ_h is often not uniform for all ions trapped in the crystal lattice. This has the effect that the hole transmission initially decreases rapidly and then more slowly.

Designing a spectral filter-creation sequence

What now might be apparent is that the task of finding the optical pumping sequence which generates good filters in a specific rare-earth-doped crystal can be arduous. What follows is a general outline of the process used in this work to find this sequence. First, one must find the shape of the desired transfer pulse. Secant hyperbolic pulses are preferable, being robust in transferring population from the ground to the excited state [82, 83]. To find the correct pulse shape, one can either solve the Bloch equations for some light pulse or, if a stronger description of the ion-crystal interaction is required, solve the Lindblad master equation for the same pulse [16, 84]. This yields much faster feedback when testing pulse parameters than what is achievable in the lab, thus making it a good place to start.

Once a pulse shape has been decided, some experimental trial and error is needed to obtain a good filter. There are two parameters left to optimise: the peak intensity of the pulse (i.e. the Rabi frequency) and the number of repetitions of the transfer pulse. The goal is to reach the maximum transmission while using the lowest number of repetitions and limiting the power broadening. In this way, the best possible filter is prepared in the least possible time. The maximum transmission can be deduced by starting with a sequence of many high-intensity pulses. The optimal pulse sequence can then be found by walking down both the number of repetitions and the peak intensity to shorten the filter creation and narrow δf .

Cycling between probing and optical pumping

With the optical pumping sequence in place, one can use the generated filter intermittently for UOT measurements. Extension of the timeline of the blue traces in Figure 4.4 provides a visualisation of the time periods in which the crystal can be used for filtering in between the re-pumping of the filter. In Figure 4.5, a 300 ms window for UOT measurements is provided between each filter re-pumping. During this time, the filter transmission falls from $\sim 80\%$ to $\sim 50\%$. As a single UOT frame takes <10 ms to capture, the transmission drop in this example filter would not have to be taken into account for a single frame, but some compensation

would be prudent when comparing multiple frames taken within a given probing window.




Figure 4.5: Rate-equation solution to the spectral-hole filter transmission, T_F , over several optical re-pumping and probing windows, where the latter is highlighted in grey.

The ratio of time available for probing to time dedicated to re-pumping is referred to as the “duty cycle” of the UOT filter. Obviously, this ratio should be as high as possible. However, since the optical medical safety limit is 300 mW/cm^2 , the loss of available probing time due to filter re-pumping may be compensated by using higher-intensity pulses during the probing, such that the average signal acquired per second is the same. The only drawback to this is a lopsided frame-rate (i.e. the frames are acquired unevenly in time).

4.2.1 The slow-light effect

There is one more property of the spectral-hole filters which demands attention, and that is the so-called “slow-light effect”. As light propagates inside the spectral hole, it interacts off-resonantly with the ions outside the spectral hole. This interaction drives the ions on the same frequency as the incoming field, and the ions in turn generate a field which oscillates 90° out of phase with the incoming field. The total electric-field propagating inside the hole is now the sum of the incoming field and the generated field. As the latter is small in amplitude, compared to the first, the total field looks like the incoming field but with a phase shift opposite to the propagation. This backward shift happens for each unit length that the incoming field propagates inside the crystal, which aggregates to the entire field slowing down. The closer the incoming field is to the edge of the hole in frequency, the more strongly the ions will be driven and thus the stronger the generated field will be. This leads to a larger retardation for light, with frequencies closer to the edge of the hole.

This describes how monochromatic light becomes retarded inside the crystal and propagates with the phase velocity. However, light is only ever truly monochromatic in theory. What is interesting to consider is how a change in amplitude, an example of information, propagates inside such a structure. The speed of this amplitude change is said to propagate with the group velocity v_g and can be found by dividing up the light into its Fourier components: v_g not only depends on the index of refraction n , but is also affected by a component depending on the derivative dn/df ,

$$v_g = \frac{c_0}{n + f \frac{dn}{df}} \quad (4.5)$$

where c_0 and n is again the the speed of light in vacuum. Derivations of the group velocity can be found in undergraduate textbooks (e.g. [85]), and the article on Wikipedia includes excellent animations that clearly describe the phenomenon [86].

Regarding the spectral hole, as the strength of the induced field sets the monochromatic retardation (i.e. n) a stronger induced field equates to a higher n . As the ions have such narrow linewidths, a small change in frequency strongly affects the driven field when close to resonance and thus n . This equates to a very high dn/df for light inside the hole, yielding a very slow v_g . Using the Kramers-Kronig relations, the group velocity can quantitatively be described as [87]:

$$v_g = \frac{\pi^2 W_{\text{SHF}}}{\alpha_B}. \quad (4.6)$$

In a filter with $W_{\text{SHF}} = 1$ MHz and $\alpha_B = 40 \text{ cm}^{-1}$, (4.6) shows that the filtered light propagates with $v_g \approx 10^{-5} c_0$. In a 1 cm long crystal, this equates to a 4 μs delay of the filtered light, compared to any light outside the hole that is not absorbed. Thus, the tagged light can be filtered from the carrier both by the normal absorptive filtering, but also by time-gating the detected signal.

In addition, the strong slow-light effect achievable in rare-earth crystals also has a potential application for laser stabilisation, which is an important component of optical atomic clocks [88]. The vision is that the traditional mechanical issues with stabilising a laser to a cavity can be avoided by instead using an optically long but physically short cavity. This research is ongoing in the quantum information group in Lund, where a proof of concept was recently presented in [87].

Filtered UOT signal

After the SHF, the signal is acquired on a photodetector. An example can be seen in Figure 4.6a, showing data traces acquired from light which has either been




Figure 4.6: Example of filtered UOT signal acquired in **Papers I–III**, using a 1 μs optical probe pulse. (a) Depiction of the filtered signal with the ultrasound on (blue) and with the ultrasound off (red). (b) Depiction of the same signal, but with the probe frequency outside the inhomogeneous line, thus the carrier and tagged light are not separated.

tagged or not in the setup used in **Papers I–III** (see Sect. 6.1). This is to be compared to Figure 4.6b, which shows the signal using the same experimental parameters – with the exception that the frequency of the probing light lies outside the inhomogeneous line and thus nothing is absorbed. Of the pulse shown in Figure 4.6b, which is mainly comprised of carrier light, a small fraction leaks through, seen as the peak at $t = 0$ in Figure 4.6a. This leakage can still be separated from the filtered tagged light due to the slow-light effect, which in the $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ filter (see Sect. 4.3.1) retards the tagged light by more than 5 μs .

4.3 Rare-earth-doped crystals for UOT

Other than burning the filter, the main challenge when using rare-earth-doped crystals as a filter is finding a suitable ion doped in a suitable host-crystal material. The first step is to identify which ion to target, based on its available 4f–4f transition wavelengths. In Figure 4.7, the theoretical crystal field levels for all 4f terms for all the lanthanides doped into LaF_3 [90] are plotted together with the tissue-optical window from Figure 1.3. As the crystal field will not shift the terms by more than ~ 10 nm, this can act as a guideline as to which ions would be suitable for UOT filtering. Once an ion has been selected, a suitable host material must be found. This involves performing an empirical and resource-intensive spectroscopic study of the desired ion when doped into different crystals. Such a search was one subject of the thesis by Alexander Bengtsson [37]. Contrarily, in the current thesis, only the properties of the crystals used for the UOT imaging presented in Chapter 6 (i.e. $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ and $\text{Tm}^{3+}:\text{LiNbO}_3$) will be discussed.

A third promising crystal ($\text{Tm}^{3+}:\text{LaF}_3$) has been studied extensively; but due to equipment issues, this material has yet been used in UOT. For details of this ma-




Figure 4.7: Calculated energy of the 4f terms of the lanthanide ions doped into LaF_3 . Each term is in turn comprised of several Stark levels, indicated by the thickness of the term in the plot. Since the spin-orbit and electrostatic potential will not be noticeably affected by the change of host material, the energy of the terms is expected to be roughly the same in another host, so long as the lanthanides are also triply ionised in the other host. Diagrams such as this are referred to as “Dieke diagrams”, named after the first such diagram presented in [89]. The plot was generated from the calculated results presented in [90].

terial, again, see the thesis of Alexander Bengtsson [37].

4.3.1 Praseodymium-doped yttrium orthosilicate

Praseodymium-doped yttrium orthosilicate ($\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$) was the crystal used for the experimental results in **Papers I–III** and **V**. In those papers, the $^3\text{H}_4-^1\text{D}_2$ transition was used, which is outside the tissue-optical window, thus making this transition inappropriate for imaging in tissue. Nevertheless, this transition in $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ has been extensively studied [91, 92] and the quantum information group in Lund has accumulated substantial experience in manipulating this material. Thus, $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ was initially a given test bed for developing UOT in Lund.

In $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$, the Pr^{3+} ions replace Y in the crystal lattice in one of two possible sites. This gives rise to two inhomogeneous lines for the $^3\text{H}_4-^1\text{D}_2$ transition, one at 605.977 nm and one at 607.934 nm. The work presented in the current thesis only concerns the first of these, and all following mentions of the transition are referring to that.

The crystal field splits the ground state into several Stark levels and, doing so, quenches the orbital angular momentum – making each Stark level a singlet state. This also cancels out the first-order hyperfine interaction, though there is a second-order interaction. The crystal field thus gives rise to a hyperfine split of each such




Figure 4.8: (a) Illustration of the hyperfine states of the excited and ground states involved in the 606 nm transition in $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ [91, 92]. (b) Illustration of the ground state and excited state involved in the 794 nm transition of stoichiometric $\text{Tm}^{3+}:\text{LiNbO}_3$ in the presence of a magnetic field.

Stark level, where each hyperfine level is doubly degenerate due to Pr^{3+} having an even number of electrons [93]. In the ground state of praseodymium, three such states exist due to its only stable isotope having a nuclear spin of $5/2$. These states are illustrated in Figure 4.8a, where they are denoted by their hyperfine quantum numbers $m_I = \pm 1/2, \pm 3/2, \pm 5/2$. Similarly, the lowest-lying Stark level in the excited state $^1\text{D}_2$ is also comprised of three hyperfine levels. The coherence time τ_2 of the $^1\text{D}_2$ state at 606 nm in $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ is in the order of 100 μs , giving the optical transition a linewidth on the order of kHz [94].

A final property of $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ in regards to UOT filters is that the transition dipole moment of the $^3\text{H}_4 \rightarrow ^1\text{D}_2$ transition is tilted in relation to the principal axes of birefringence in Y_2SiO_5 . These orthogonal axes are denoted \mathbf{D}_1 , \mathbf{D}_2 , and \mathbf{b} , where the transition dipole moment is practically orthogonal to \mathbf{b} and angled 75° from \mathbf{D}_1 towards \mathbf{D}_2 . If light is propagating along \mathbf{b} , this misalignment means that only the minimal absorption polarisation, which turns out to be \mathbf{D}_1 , is stable. Thus, even if the polarisation were initially along the axis of highest absorption (i.e. \mathbf{D}_2), it would rotate towards and propagate with the polarisation along \mathbf{D}_1 [95]. For the $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ crystal used in **Papers I-II** and **V**, this polarisation rotation occurs after approximately 5 mm propagation along \mathbf{b} , after which the polarisation is mainly along \mathbf{D}_1 . Thus, light polarised along \mathbf{D}_1 after passage through the crystal is comprised of carrier light, which can be removed using a polariser set to transmit polarisation along \mathbf{D}_2 , as was the case in **Paper I** and **II**.

The polarisation dependence of the absorption in $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ comes with another dilemma, in that half of the intensity collected from the tissue must be removed in order to achieve high-contrast filters. However, this is preferably avoided to reach the maximum possible imaging depth in UOT and can be avoided in crystals where the absorption is polarisation-independent. This can be achieved if the transition dipole moments of the doped ions are equally distributed in a plane

Table 4.2: Properties of 0.05%Pr³⁺:Y₂ Si O₅, as reported in [91, 94]. Properties may vary due to boule-dependent differences.

| α_{D_1} [94] | α_{D_2} [94] | T_1 [91] | Γ_{hom} [94] |
|-------------------------------|----------------------------|-------------------|------------------------------|
| $3.6 \pm 0.5 \text{ cm}^{-1}$ | $47 \pm 5 \text{ cm}^{-1}$ | 164 μs | 2.4 kHz ($\mathbf{B} = 0$) |

inside the crystal.

In **Papers I–III and V**, a $10 \times 10 \times 12 \text{ mm}^3$ large 0.05%Pr³⁺:Y₂ Si O₅ crystal was used for filtering where the long axis was parallel to **b**. For these filters, $T_F \approx 60\%$. For collimated light, $T_C \approx -45 \text{ dB}$; and $T_C \approx -32 \text{ dB}$ for diffuse light (attenuation including the $\sim 5 \mu\text{s}$ slow-light effect in Fig. 4.6). This decreased filter suppression is expected to arise from a lower polarisation purity when polarising the diffuse light prior to SHF, as a similar drop in suppression for scattered light compared to collimated light was not seen in the Tm³⁺:LiNbO₃ filter.

4.3.2 Thulium-doped stoichiometric lithium niobate

Lithium niobate as a host material

One host-crystal structure which allows for planar symmetry of the transition dipole moments is lithium niobate. This can achieve polarisation-independent absorption when the to be filtered light propagates along the crystal c-axis. Moreover, Figure 4.7 suggests that Tm has up to four transitions within the tissue-optical window. The transitions at $\sim 700 \text{ nm}$ and $\sim 800 \text{ nm}$ are especially well-suited for imaging. The $^3\text{H}_6\text{--}^3\text{H}_4$ transition at 800 nm is close to the isosbestic point of blood, thus an image at that wavelength would be mainly an image of the blood volume. For the other transition, $^3\text{H}_6\text{--}^3\text{F}_3$ at 700 nm , oxygenated and deoxygenated blood have the largest differences in absorption. Combining the two wavelengths thus allows for blood volume-independent imaging of the tissue oxygenation.

An issue with Tm is that the only stable and naturally occurring isotope, ^{169}Tm , has a nuclear spin of $1/2$. Similar to Pr³⁺:Y₂ Si O₅, the first-order hyperfine interaction is quenched, and thus due to the nuclear spin, there is only one non-degenerate hyperfine level. This means that spectral holes cannot be generated in the hyperfine levels of Tm³⁺ without breaking the degeneracy of the hyperfine quantum numbers $m_I = \pm 1/2$. This degeneracy may be broken by the application of a magnetic field, splitting the hyperfine state according to the Zeeman effect [16] (depicted in Fig. 4.8b). As discussed in Section 4.2, applying this field along a specific crystal axis direction may further increase τ_h .

The ${}^3\text{H}_6\text{--}{}^3\text{H}_4$ transition has been studied in congruent $\text{Tm}^{3+}:\text{LiNbO}_3$,⁵ and the properties are reported in [96, 97]. From these results and correspondence with one of the authors (Charles Thiel⁶), it does not seem possible to create high-transmission spectral holes in this material due to the significant overlap of the crystal field levels for inequivalent Tm^{3+} sites. The ${}^3\text{H}_6\text{--}{}^3\text{H}_4$ in $\text{Tm}^{3+}:\text{LiNbO}_3$ does, however, display a very strong oscillator strength compared to Tm^{3+} in other hosts and a narrow linewidth, which would allow sharp and high attenuation filters in relatively short crystals.

As LiNbO_3 does not contain any similar ion with the same ionisation, the doping process is not as simple as Tm replacing one specific atom in the crystal lattice. Instead, Tm replaces both Nb and Li in the lattice to varying degrees, depending on the concentration of Tm [98], though it most commonly replaces Li [99]. As Tm^{3+} replaces Li^{2+} or Nb^{5+} in LiNbO_3 , the doping leads to a nonlocal charge compensation for the Tm ion, which may be the cause of a less Lorentzian-shaped inhomogeneous broadening.

The significant lack of Li in the lattice of congruent LiNbO_3 is believed to increase the crystal lattice strain and thus induce a larger inhomogeneous broadening. This lack of Li is addressed in stoichiometric LiNbO_3 , and thus it was hypothesised that better dopant properties could be found in this host. Based on unpublished data showing Eu doped into stoichiometric LiNbO_3 (information obtained through communication with Charles Thiel *et al.*), this hypothesis seemed plausible. To investigate this further, Kovács László⁷ was contacted, and who subsequently and graciously lent the quantum information group at Lund University a stoichiometric $0.005\%\text{Tm}^{3+}:\text{LiNbO}_3$ with the dimensions $4.7 \times 4.7 \times 6.0 \text{ mm}^3$ such that a spectroscopic survey could be performed. In this crystal, the c-axis was parallel to the 6.0 mm long side. As of the present writing, this survey is still ongoing, though the results are sufficiently promising that the crystal was intermittently used during this survey as a filter for a UOT experiment.

Inhomogeneous broadening of stoichiometric $\text{Tm}^{3+}:\text{LiNbO}_3$

The inhomogeneous line of stoichiometric $\text{Tm}^{3+}:\text{LiNbO}_3$ – or in this case, the more aptly named “profile” (displayed in Fig. 4.9a) – was measured by László *et al.* using a Fourier transform spectrometer. This absorption spectra has been partially confirmed by transmission spectroscopy measurements performed in Lund and

⁵ The term “congruent LiNbO_3 ” being used to describe the crystalline material with the same structure as the theoretical structure of LiNbO_3 , but which has a significant lack of Li.

⁶Montana State University, Bozeman.

⁷Wigner Research Centre for Physics, Budapest.




Figure 4.9: Inhomogeneous profile of ${}^3\text{H}_6\text{--}{}^3\text{H}_4$ in stoichiometric $0.005\%\text{Tm}^{3+}\text{:LiNbO}_3$. (a) The relative absorption at 8 K, measured by László *et al.* The grey area in a depicts the region where the absorption was measured in Lund at 3 K, the results of which are shown in (b). The maximum measurable absorption in (b) was limited to 21 cm^{-1} due to thickness of the crystal and the amplified spontaneous emission (ASE) of the laser, hence the “flat” central peaks. The red arrow in (b) indicates the position at which the filter in Figure 4.10 was created.

displayed in Figure 4.9b. Compared to the congruent $\text{Tm}^{3+}\text{:LiNbO}_3$ spectrum reported in [97], stoichiometric $\text{Tm}^{3+}\text{:LiNbO}_3$ has more resolved peaks, which may indicate less overlap of the crystal field levels. The spectral hole used in the UOT measurement described in Section 6.3 was created at the highlighted peak, as this was the position on the inhomogeneous profile that – at the time surveyed – showed the sharpest and highest transmission holes, while still providing high carrier attenuation.

Though the peak at $\sim 794.3\text{ nm}$ has higher absorption,⁸ the hole burning properties were not favourable, as the peak achievable filter transmission was $T_F \approx 20\%$ and the filter edges were not very sharp. As there is a small hint in Figure 4.9b of two peaks at $\sim 794.3\text{ nm}$, the limited transmission could, like the congruent crystal, be the result of overlapping crystal field levels.

At the highlighted peak in Figure 4.9b, much sharper filters could be created with $T_F = 75\%$. One such filter is depicted in Figure 4.10a₁, where a 4.5 mT field is applied along the crystal c-axis. As the ground-state Zeeman splitting for this field angle was measured to 1.3 MHz/mT , this field also place the anti-holes at $\pm 6\text{ MHz}$ from the main hole, yielding additional carrier attenuation for a UOT signal generated by a 6 MHz ultrasound. Fitting solutions of (4.3) to the decay rate of the filter yielded $\tau_h = 31\text{ s}$, though this was not a good fit ($R^2 = 0.74$). By splitting the ion population into two classes, each with different lifetimes τ_{h1} and τ_{h2} , a better fit ($R^2 = 0.998$) to the model was found when $\tau_{h1} = 6.8\text{ s}$, $\tau_{h2} = 67\text{ s}$, and the population split was 38:62 (i.e. a larger population with the lifetime τ_{h2}).

⁸In this experiment, the maximum absorption of this peak could not be measured by our setup due to the ASE of the lasers outside the inhomogeneous profile. The ASE suppression was roughly 60 dB for this laser (M-squared SolsTiS).




Figure 4.10: (a₁) Single filter for collimated light. (a₂) Single filter but with scattered light and pumped orthogonally to the filtered light propagation. (b) Double filter for collimated light where absorption at $\Delta f = 0$ has been increased using the anti-holes of spectral holes outside the trace.

That a single hyperfine lifetime, τ_h , could not accurately describe the filter-decay rate is attributed to circumstantial differences for ions in the lattice, which in turn affect their individual lifetimes. Such differences could be an ion's distance to its nearest Tm^{3+} neighbour or its unique local host-crystal structure [100]. In addition to these measurements on the marked spot in Figure 4.9b, photon echo spectroscopy was performed as well to study the homogeneous linewidth Γ_{hom} .

Photon echo spectroscopy

In short, photon echo spectroscopy works similar to the example of MRI given in Chapter 1, but instead of a macroscopic magnetic dipole oscillation generating a radio wave, it is instead a macroscopic electric dipole oscillation which generates an optical field. A short optical $\pi/2$ -pulse induces a coherent superposition of the ground state and the excited states, generating many small oscillating dipoles.⁹ Initially, these dipoles are in phase, generating a macroscopic dipole oscillation (i.e. the free induction decay). This signal quickly disappears however, as the amplitude of the macroscopic dipole rapidly becomes zero. This happens because the individual dipoles – the sum of which yields the macroscopic one – are generated by ions from different parts of the inhomogeneous line, and thus oscillate with different frequencies, which in turn makes an individual dipole de-phase in relation to the others as time moves on. A π -pulse sent a time τ after the first, flips the direction of this de-phasing by inverting the phase of the superposition states. A time τ after this second pulse, the ions still in this coherent superposition again

⁹ In the case of MRI, the ground state of hydrogen is split into two states by an external magnetic field, where the lower state is slightly more populated. A radio-wave $\pi/2$ -pulse into this ensemble of protons generates two π -phase-shifted superposition states – one from the protons in the upper state and one from the protons in the lower state. Due to the population imbalance, the amplitude of the sum of these two states is larger than zero, thus a radio wave is emitted as the total state evolves.

come in phase. This leads to an “echo” being emitted a time 2τ after the first pulse, as the individual dipoles are again in phase. The strength of this echo as a function τ can be used to measure coherence lifetime, τ_2 , of the excited state. More in depth discussions of photon echo spectroscopy can be found in [93, 101].

Using this technique and (4.1), the homogeneous linewidth at the marked position in Figure 4.9b was estimated to $\Gamma_{\text{hom}} \approx 10$ kHz and the excited-state lifetime was concluded to be similar to the 160 μs measured in congruent $\text{Tm}^{3+}:\text{LiNbO}_3$ [97].

UOT filter in $\text{Tm}^{3+}:\text{LiNbO}_3$

The filter generated at this position is used in the measurements in Section 6.3, where the filter properties in Figure 4.10a₂ were measured for diffuse light. The increased attenuation at 6 MHz in Figure 4.10a₂ (compared to Fig. 4.10a₁) is attributed to the optical pumping in Figure 4.10a₂ being performed orthogonally to the filtered light propagation and the lack of complete overlap between the filtered and pumping fields inside the crystal.

Using the filter’s anti-hole to increase the carrier absorption (as in Fig. 4.10a) requires a low probability of direct de-excitation to the opposite nuclear spin state. Otherwise, ions excited by the carrier produce radiation with the same frequency as the tagged light when they de-excite, creating a “false signal”. This can be avoided by using another magnetic field strength in which the anti-hole is placed away from the carrier. This leaves the filter with less carrier suppression. This loss in suppression can then be addressed by pumping an additional spectral hole, which produces an anti-hole at the carrier frequency. This additional absorption can be increased again by pumping two spectral holes at a higher and lower frequency to produce anti-holes at the carrier frequency. Such a filter can be seen in Figure 4.10b, where – in addition to the two stacked anti-holes at $\Delta f = 0$ – two spectral holes have been prepared for filtering both first order sidebands generated from a $f_{\text{US}} = 6$ MHz ultrasound. This filter is otherwise created at the same position on the inhomogeneous profile as Figure 4.10a, but with a 19 mT field.

A drawback of using a filter such as that in Figure 4.10b is that it takes longer to create than a filter such as the one in Figure 4.10a, which in turn means a slower image-acquisition rate. To put this into context, the filter in Figure 4.10a took ~ 100 ms to prepare while the filter in Figure 4.10b took four times longer, but only allowed for twice the signal to be detected (disregarding any false signals). If the filter can only be used for ~ 300 ms due to decay of the hyperfine states, this additional filter-creation time seriously reduces the image-acquisition speed.




Figure 4.11: (a) Transmission of the carrier when either matching the antihole to carrier (full blue) or not (dashed red). The smaller axis in (a) is a zoom at the start of the trace. (b) The normalised fluorescence spectra, unfiltered (full blue) or after an interference filter (dashed red). The spectral resolution of the spectrometer used was <0.3 nm.

There is thus a trade-off to take into account for a single-hole filter: place the anti-hole off the carrier frequency, yielding less attenuation (i.e. more noise), or place the anti-hole at the carrier, risking a false signal (i.e. more noise). The preferable option is effectively determined by the probability of direct de-excitation from the excited state to the opposite nuclear spin state. As the difference in attenuation for the carrier is 15 dB when placing the anti-hole on/off the carrier, the additional noise in the two cases is roughly equal for such a probability of 3%. A lower probability means that placing the anti-hole on the carrier is preferable, a higher indicates that placing it off the carrier is preferable.

Fluorescence analysis of $\text{Tm}^{3+}:\text{LiNbO}_3$

To investigate this trade-off, the filter in Figure 4.10a₂ was studied using a $20\ \mu\text{s}$ square optical pulse with a frequency of $\Delta f = 0$. This pulse acted as a synthetic carrier in the imaging experiment discussed in Section 6.3.

The result of “carrier light” hitting the filter is seen in Figure 4.11a. The decay in the trace is attributed to false-signal fluorescence from the excited $^3\text{H}_4$ state to the other nuclear spin ground state. However, the fluorescence may also be from the excited state to a higher lying $^3\text{H}_6$ crystal field Stark level. As these higher Stark levels are not populated, the fluorescence at this transition is not absorbed by the crystal but is close enough in wavelength to be detected within the used detector wavelength bandwidth.¹⁰

An examination of the two traces in Figure 4.11a, reveals that the fluorescence tails are almost identical, while direct transmission is considerably higher when placing

¹⁰ 350-850 nm, GaAs photomultiplier tube.

the anti-hole off the carrier. This indicates that the fluorescence is predominantly to a higher-lying Stark level and that de-excitation probability favours placing the anti-hole on the carrier frequency, as in Figure 4.10.

Figure 4.11b shows the fluorescence spectrum of the used $\text{Tm}^{3+}:\text{LiNbO}_3$ at 3 K, captured in a direction orthogonal to the excitation light. The first peak located at 794.2 nm is direct fluorescence or scattering, while the second peak shifted 0.5 nm is attributed to fluorescence to a Stark level. This agrees with [97], which found the first Stark level in congruent $\text{Tm}^{3+}:\text{LiNbO}_3$ 8 cm^{-1} above the ground state. The third peak, roughly 1 nm from the first peak, is attributed to absorption of the second peak in other ions at other crystal sites (i.e. other positions in the inhomogeneous line; see Fig. 4.9) and subsequent re-emission to the close Stark level. The wider peak at $\sim 804 \text{ nm}$ is attributed to another Stark level.

The fluorescence in Figure 4.11a is mainly attributed to the third peak in Figure 4.11b as the two previous should be largely absorbed in the inhomogeneous profile.¹¹ This fluorescence can be filtered using an interference bandpass filter [102]. However, the central-transmission wavelength λ_T of such filters is dependent on the angle of incidence θ of the light, as

$$\lambda_T = \frac{\lambda_0}{n} \sqrt{n^2 - \sin^2(\theta)} \quad (4.7)$$

where λ_0 is the normal-incidence central-transmission wavelength of the filter and n is the refractive index of the spacer layer of the filter [102]. As the light that is to be filtered is diffuse, it will have a distribution of normal incidences. This means that even using an infinitely sharp interference filter, it is not possible to completely separate the first peak in Figure 4.11b from the fluorescence peaks. As seen in the filtered spectra in Figure 4.11b, the last peak may however be eliminated completely. This was accomplished using an interference filter with a peak transmission of 57%, $\lambda_T = 795.6 \text{ nm}$, and a transmission full-width at half-max of 10.1 nm.

¹¹ The first two peaks are visible in Figure 4.11 due to the orthogonal excitation, thus any fluorescence does not need to traverse the entire crystal absorption.

Chapter 5

Comparison of methods for detecting diffuse frequency-shifted light

The use of spectral-hole filtering (SHF) is far from being the only method for separating the carrier from the tagged components in the detected light, and a myriad of methods has been investigated and used since the birth of UOT research 30 years ago [103]. These methods can generally be divided into two categories: spectral and temporal.

One example of a spectral method is SHF. This is defined as applying frequency selection in frequency space by some optical component and then detecting the unfiltered light. This frequency selection can be absorptive, as in SHF, or involve the selection of optical modes (e.g. using optical cavities). One of the first demonstrations of UOT involved an example of the latter, with the frequency-shifted light matched to a Fabry-Perot cavity [42]. However, as the light emitted from tissue occupies many spatial light modes, very little light will couple to the cavity (they have low *etendue*) and thus the detected power of the tagged light will be low. This is to a large extent solved in SHF, as the frequency selection of a spectral hole is independent of the optical spatial modes. Instead, the only spatial light mode selection in SHF is that of the optics and the crystal itself – where both can be designed to have large *etendue*.

Temporal methods base the tagged-light measurement on an examination of a beat frequency of the light intensity due to the mixing of two optical frequencies. The simplest case is just to examine the field on a single-element photodetector, where

the intensity I_{det} is

$$I_{\text{det}} \approx I_C + I_T + 2\sqrt{I_C I_T} \cos(2\pi f_{\text{US}} t + \Delta\omega) \quad (5.1)$$

where I_C and I_T are the intensities of the carrier and tagged light, and $\Delta\omega$ is the relative phase of the tagged light to the carrier. The strength of the intensity modulation at the frequency f_{US} is thus proportional to $\sqrt{I_T}$. The approximation in (5.1) holds for $f_C \gg f_{\text{US}}$ and the equation further requires that the phase $\Delta\omega$ is constant. As the tissue moves, however, the second criteria is challenged in a process called “speckle decorrelation”, where movement of tissue causes $\Delta\omega$ to change. This leads to a fundamental limitation in temporal methods, which need to acquire their signal in the short time before speckle decorrelation kicks in, which is $\sim 100 \mu\text{s}$ for light propagation through only a few mm *in vivo* [104]. This is technically not very limiting, as the ultrasound has gone through the tissue many times over in this time-frame. It does, however, limit how averaging is performed, as each signal must be acquired individually.

To reach their full potential, temporal methods also require that multiple detectors are used where the area per detector is minimised. This can be seen by examining (5.1) again, but with the replacements $I_C \rightarrow M\bar{I}_C$, $I_T \rightarrow M\bar{I}_T$ where M is now the number of speckles (i.e. the regions where the emitted light is spatially coherent) on the detector and where \bar{I}_C and \bar{I}_T denote the average intensity of the carrier and tagged light in a speckle (i.e. the temporal analogue to $\mathcal{I}_{\text{speckle}}$, which was introduced in Sect. 3.3.2). As the shot noise is proportional to the square root of the total signal on the detector, increasing the number of speckles on the detector yields no increase in signal-to-noise ratio (SNR), as both signal and noise (in the shot-noise regime) scales with \sqrt{M} . Reducing the size of each detector below the speckle size generally does not yield any benefit, as measurements on the same speckle are completely correlated; thus, this does not improve the SNR.

In simple terms, spectral methods are mainly limited by *etendue*, as light must be guided through the spectral element. On the other hand, temporal methods are mainly limited by speckle decorrelation, as detectors can be placed directly on the tissue. In reality, there is a mix of the two limitations for any technique. Quantitatively, different methods can be compared in terms of their respective achievable contrast-to-noise ratios (CNR). This was the aim in **Paper IV**, where three temporal methods are discussed and compared to SHF.

We define the CNR as the difference in tagged signal I_T between tissue voxels A and B, divided by total noise I_{noise} :

$$\text{CNR} = \frac{|I_{T,A} - I_{T,B}|}{I_{\text{noise}}}. \quad (5.2)$$

For SHF detection of both first-order sidebands, the signal in the number of detected photons is $S\eta_{\text{det}}(2\mathcal{N}_{\text{T}} + T_{\text{C}}\mathcal{N}_{\text{C}})$, where S is the area from which light is collected at the tissue and η_{det} the detection efficiency. \mathcal{N}_{T} and \mathcal{N}_{C} are the numbers of tagged first sideband and carrier photons per unit area at the collection point. As SHF can be realised with low noise detectors that may operate in the single photon regime, shot noise may be assumed (i.e. that the noise is the square root of the average signal¹). This yields the CNR for spectral-hole filtering as follows:

$$\text{CNR}_{\text{SHF}} = \sqrt{2S\eta_{\text{det}}} \frac{|\mathcal{N}_{\text{T},\text{A}} - \mathcal{N}_{\text{T},\text{B}}|}{\sqrt{\mathcal{N}_{\text{T},\text{A}} + \mathcal{N}_{\text{T},\text{B}} + T_{\text{C}}\mathcal{N}_{\text{C}}}}. \quad (5.3)$$

For quantitative comparison, numbers must be assigned to the variables η_{det} and T_{C} (\mathcal{N}_{C} and \mathcal{N}_{T} are discussed in Sect. 5.2). Here, $\eta_{\text{det}} = 7\%$ is assumed from a 15% detector quantum efficiency² and a numerical aperture $\text{NA} = 0.71$ of the optics (yielding a $\text{NA}^2 = 50\%$ collection efficiency). The transmission for the carrier T_{C} is set to either -30 or -80 dB.

What follows is a description of the three other prevalent methods that were compared to SHF in Paper IV and their respective expressions for the CNR. As mentioned in **Paper IV**, other detection methods exists, where one particularly promising method not discussed here being laser feedback detection [105].

5.1 Overview of the methods that were compared to SHF in Paper IV

5.1.1 Single-shot and digital off-axis holography

In off-axis holography [106], the emitted light field exiting the tissue, which contains tagged light, is polarised and overlapped with a planar reference field, propagating at an angle with the frequency $f_{\text{C}} + \eta f_{\text{US}}$, $\eta \in \mathbb{Z}$. The resulting spatial interference pattern is captured on a camera with a sufficient integration time that the beating term between the carrier light and the reference field is averaged out. The method can thus be used to measure either the carrier or an arbitrary sideband intensity by selecting the corresponding optical frequency of the reference field. The capture of the interference pattern can in turn be done in a “single-shot” regime (i.e. the integration time is one to a few ultrasound periods) and thus the

¹ Assuming that \mathcal{N}_{C} is equal when probing voxels A and B.

² The efficiency of an average photomultiplier tube.

method is not sensitive to speckle decorrelation. The captured interference pattern is then subject to a digital two-dimensional Fourier transform, after which the strength of the carrier can be discerned by examining the spectral intensity at the spatial frequency that corresponds to the spatial frequency of the expected interference pattern. To adequately resolve the spatial frequency of the interference pattern, more than one camera pixel is assigned to each speckle, where the signal is maximised when a speckle on the detector is $n = 4$ pixels wide [106]. If the angle of the interference pattern is unknown, $n = 16$ pixels must thus be assigned to each speckle. Based on the SNR presented in [106] for a camera with N_{px} pixels, the following expression for the off-axis holography CNR was presented in Paper IV:

$$\text{CNR}_{\text{HOL}} = \sqrt{N_{\text{px}}} \frac{\eta_{\text{det}} |\overline{\mathcal{N}}_{\text{T,A}} - \overline{\mathcal{N}}_{\text{T,B}}|/n}{\sqrt{4\eta_{\text{det}} (\overline{\mathcal{N}}_{\text{T,A}} + \overline{\mathcal{N}}_{\text{T,B}})/n + 2}}. \quad (5.4)$$

Compared to that in (5.3), the signal in (5.4) is the average photons per speckle $\overline{\mathcal{N}}_{\text{T}}$. The detection efficiency was estimated to $\eta_{\text{det}} = 30\%$ from a pixel quantum efficiency of 60% and a 50% loss in power due to the polariser. As the number of photons per speckle is preserved, the NA does not affect η_{det} . The two variations of the setup used in the comparison for this detection method used either a 1 MPx camera or a 50 MPx camera.

5.1.2 Photorefractive crystal detection

Instead of mixing two fields and capturing the interference pattern with a camera, it is possible to mix the emitted field and a plane reference field in a photorefractive material. An interference pattern is formed inside the photorefractive material, which results in the formation of a grating in the index of the refraction of the material due to the photorefractive effect. Part of the reference field will diffract off this grating in the direction of the emitted field, with an efficiency of η_{pr} , replicating the emitted fields wavefront. The signal is then integrated by the detector, removing any beating terms. From here, there are many schemes for deducing the tagged-light intensity by measuring either the tagged sidebands [31, 107, 108] or the untagged carrier [33, 109, 110].

As an example of the former case [108], the frequency of the reference field is set as $f_{\text{C}} + f_{\text{US}}$ and the interference pattern replicates the tagged wavefront. Flipping the phase of the ultrasound by π will also flip the phase of the optical field by π . However the diffracted field “remembers” the previous phase, due to the limited response time of the photo-refractive material and thus the tagged field, and the diffracted field will now destructively interfere on the detector. The difference

in intensity between constructive and destructive interference on the detector is then $\Delta I \simeq 4\eta_{\text{pr}}I_t$. This method requires continuous ultrasound to keep the grating active and to continually frequency shift light in the tissue. The continuous ultrasound might seem to be a drawback, but any apparent loss in resolution is circumvented to great success and good image-acquisition speeds by measuring the UOT signal using plane waves at different angular frequencies, with spatial information collected via spatial Fourier transform [64, 111].

The detection of the tagged light by measurement of the carrier works slightly differently, and in the scheme chosen for the comparison in **Paper IV**, pulsed ultrasound can be used. The photorefractive grating is now constructed in the absence of an ultrasound pulse. As an ultrasound pulse then travels through the medium, the energy conservation in the tagging process means that the only difference in the intensity on the detector is the difference in the interference term between the diffracted field and the emitted field. This yields the intensity difference of $\Delta I \simeq 2\eta_{\text{pr}}I_T$ [108].

While the signal is theoretically higher when detecting the sideband, the scheme using the carrier measurement was chosen for the comparison because it was more flexible to pulsed ultrasound (which was used in the simulated case study) and because experimental implementations of the two techniques have yielded similar SNRs (i.e. they seem experimentally equivalent). Thus, the following CNR expression for this photorefractive method was presented in **Paper IV**:

$$\text{CNR}_{\text{PR}} = \sqrt{2S\eta_{\text{det}}} \frac{\eta_{\text{pr}} |\mathcal{N}_{\text{T,A}} - \mathcal{N}_{\text{T,B}}|}{\sqrt{\mathcal{N}_{\text{C}}(1 + \eta_{\text{pr}})^2}}. \quad (5.5)$$

This method depends on having large crystals with high η_{pr} , but there is an issue with sourcing materials for which the photorefractive grating can be generated in the time that the emitted field is spatially coherent (i.e. within the tissue correlation time³). The to the author known fastest crystals are reported to have response times of a few ms [31, 112], which makes them ill-suited for *in vivo* applications. For the comparison, it was assumed that this issue had been solved and a crystal with a fast response time – either $\eta_{\text{pr}} = 10\%$ with $\eta_{\text{pr}} = 35\%$ – had been sourced of sufficient size to accommodate for all the collected light. $\eta_{\text{det}} = 40\%$ was estimated from an 80% detector quantum efficiency⁴ and with an NA = 0.71 of the optics.

³ A response time that is too fast (i.e. <10 μs) is also unwanted as the “memory” of the diffracted wavefront would be too short.

⁴ For example, a large photo diode.

5.1.3 Speckle contrast

As indicated by (5.1), each speckle in the speckle pattern will “blink” with the frequency of the ultrasound. Imaging this speckle pattern and integrating it for some time will thus produce an image of a speckle pattern that is either “smoothed” or “blurred” when compared to a speckle pattern with no blinking speckles. The contrast (i.e. the magnitude of the blurring) can thus be used to determine the intensity of the tagged light [113–116]. This contrast measures the tagging of light into all sidebands simultaneously. Truncating it to the first-order sidebands, the contrast \mathcal{C} can be written as [117]:

$$\mathcal{C} = 1 - \frac{2\mathcal{N}_T}{\mathcal{N}_C}. \quad (5.6)$$

Equation (5.6) is the contrast in the ideal case (i.e. where the speckle pattern is fully developed and each speckle is uncorrelated). From this, the CNR can be written as

$$\text{CNR}_{\text{SC}} = \sqrt{2N_{\text{px}}} \frac{|\mathcal{N}_{T,A} - \mathcal{N}_{T,B}|}{\mathcal{N}_C}. \quad (5.7)$$

where the noise of the contrast is estimated by assuming a normal distribution for \mathcal{C} , with the standard deviation estimated to $\sqrt{1/N_{\text{px}}}$, using the algorithm in [118]. This noise estimation only accounts for the stochastic variation in each pixel and not for any noise associated with CCD or CMOS cameras that become relevant for the low signal and fast readout rate required by rapid UOT imaging. As such, (5.7) should be seen as the upper limit of what is achievable for the speckle contrast. This is supported by the generally lower contrasts reported in the experiments.

As shown by (5.7), any detection efficiency in the CNR exists in both the numerator and denominator and is thus cancelled out. The only equipment dependent variable for speckle contrast is thus N_{px} , which – as in digital holography – was set to either 1 or 50 MPx.

5.2 Quantitative contrast-to-noise comparison of detection schemes

Using the expressions for the CNR of the four methods, a numerical comparison of the three methods was made. The fraction of tagged and carrier photons that were collected from an area $S = 1 \text{ cm}^2$ was calculated using the SMC model described in Chapter 3 for the geometry depicted in Figure 5.1. Assuming that 900 voxels




Figure 5.1: Simulation domain for calculating tagged signals when probing regions A and B, where the latter had a larger absorption than the background. (a) The domain shown in the plane of the light input and output locations. (b) The domain shown in the image plane (which is orthogonal to the plane in a) with the ultrasound pulse probing region A. (c) The domain shown in the image plane with the ultrasound pulse probing the absorbing region B. Reprinted with permission from [Paper IV](#) ©2022 Optica Publishing Group.

are probed per second, the total input signal allocated per voxel per second was $\sim 10^{15}$ photons. This number was derived from the medical maximum of 300 mW allowed for the 1 cm^2 optical input area and 800 nm wavelength [5]. The speckle diameter used to determine the average photons per speckle was 600 nm, which is determined by both the wavelength and $\text{NA} = 0.71$ [119]. The reduced scattering coefficient used was 5 cm^{-1} based on average parameters for muscle tissue from [6, 120], and $g = 0.9$ was used [6]. The background absorption was set to $\mu_a = 0.2 \text{ cm}^{-1}$. The region A was set to a $3 \times 3 \times 3 \text{ mm}^3$, with a different absorption than the background. Region B was set 5 mm from region A at the same depth. In the case of homogeneous absorption, the optical fluence in regions A and B would be identical due to the symmetry.

Based on the simulated tagged and carrier strengths, the CNR could be calculated for different depths and detection schemes. The results of this are seen in Figure 5.2a, while the depth at which the $\text{CNR} = 1$ is for the different schemes is seen in Figure 5.2b.

Starting with the obvious conclusion from Figure 5.2, SHF strongly outperforms the other techniques, with Figure 5.2b showing that SHF can be used to image at twice the depth of the others. At low depths, the -30 dB and -80 dB filters yield the same CNR, as the noise is dominated by the tagged photons. At the 2 cm depth, the two traces diverge, as the untagged noise becomes significant for the -30 dB filter. The -80 dB filter continues along the beneficial noise scaling, where




Figure 5.2: (a) CNR as a function of depth for an inclusion with $\mu_a = 0.4 \text{ cm}^{-1}$ and background $\mu_a = 0.2 \text{ cm}^{-1}$. (b) The imaging depths at which $\text{CNR} = 1$. I is the SHF method using a $T_C = 10^{-3}$ or $T_C = 10^{-8}$. II is photorefractive detection using a grating efficiency $\eta_{pr} = 0.1$ and $\eta_{pr} = 0.35$. III is the digital holography method, with a $N_{px} = 1 \text{ Mpx}$ or $N_{px} = 50 \text{ Mpx}$ camera. IV is the speckle contrast method using a $N_{px} = 1 \text{ Mpx}$ or $N_{px} = 50 \text{ Mpx}$ camera. $\text{CNR} = 1$ indicates the point at which the two regions cannot be differentiated. Adapted with permission from **Paper IV** © Optica Publishing Group.

the imaging depth is ultimately limited by a lack of detected tagged photons.

The clear underperformer is the speckle contrast method. As the setup is similar – in terms of equipment requirements – to digital holography, the findings suggest that the digital holography method is almost always preferable to the speckle contrast method. The comparison between the two holography methods is more nuanced, however. Although the photorefractive method is slightly superior according to the results presented here, the inclusion of noise terms – that in the present analysis are neglected – may shift the CNR of photorefractive method down to the point that off-axis holography becomes superior. As argued in **Paper IV** however, these noise terms (e.g. crystal defects which scatter light and beam fanning [121, 122]) may still in the best case scenario be dominated by the shot noise.

5.3 Qualitative comparison of detection schemes

As shown by the quantitative comparison in the previous section, SHF is superior to any of the discussed techniques in both the current and the here conceived future setups. While SHF yields the best CNR in this limited comparison, the method has other pros and cons which may make it more or less attractive than the other filtering techniques. Table 5.1 provides an overview of some of the softer traits of the different techniques. These traits, while to some degree subjective in their evaluation, are as important for the comparison of the techniques as the results in Section 5.2. The first trait in Table 5.1 is the geometry of the setup's effect on the achievable CNR. While a general-use UOT device would follow a “reflec-

tion geometry” such as the one in Figure 5.1, “transmission geometries” – where the light input and output are on opposite sides of the tissue – could be used in some medical applications such as breast imaging. In the reflection geometry, the number of carrier photons would remain largely unchanged for imaging different depths, as the majority of the detected light flows close to the surface. In the transmission setup, however, the losses for the carrier and the tagged light both scale with the thickness of the tissue being imaged. Thus, as only the CNR of off-axis holography (5.4) is independent of \mathcal{N}_C , only this detection method is independent of the geometry (in respect of noise from the carrier). The other techniques instead favour transmission geometries, as \mathcal{N}_C is smaller here in relation to \mathcal{N}_T .

Another important trait to consider is the complexity and cost of the equipment. Here, SHF stands out, as it requires both cryostats and laser sources with practically no frequency drift.⁵ These are considerable upfront costs, which limit the accessibility of SHF for many research groups. Conversely, the other methods require simpler light sources that are narrow to the extent that the tagged and carrier light are spectrally separable and that the laser does not drift over the period of a single measurement. The equipment which replaces SHF’s cryostats and crystals in the other techniques (i.e. cameras and the photorefractive materials) are comparatively much cheaper as well.

As illustrated by Figure 4.7, SHF is also very limited in available operational wavelengths. SHF is joined here by the photorefractive method, however, as the photorefractive materials can be wavelength-selective. This must be compared to the wavelength independence of the cameras for the other techniques. However, SHF is the only one of the discussed techniques that is insensitive to speckle decorrelation. For the camera techniques, this should theoretically not be a major issue, as the signal acquisition is performed quickly. However, this is not the case for the photorefractive method, where speckle decorrelation in relation to the material response time is the largest issue. It is worth pointing out that speckle decorrelation has, to the author’s knowledge, only been investigated *in vivo* at depths of a few millimetres [104]. The impact of speckle decorrelation at depths of several centimetres may thus be an unknown factor.

Finally, in relation to scaling up the amount of captured signal, SHF and the photorefractive method provide more straightforward paths. Here, one needs to increase the size of the collection optics and the size of the crystals (rare-earth-ion-doped and photorefractive crystals alike). This requires more power in the pumping of both techniques, as larger crystals mean more ions to pump and larger refractive index gratings. Scaling up the detection area also requires a larger single-

⁵ In SHF, if the laser drifts, the hole will become wider and less deep, yielding a higher carrier leakage while simultaneously giving less signal.

element detector for these techniques. Increasing the detected area (or *etendue*) of the camera techniques requires the same increase in size in the collection optics (if any are used), and a higher pixel count camera is also required. This then increases the amount of data generated per signal acquisition. This puts additional strain on the data analysis, which scales with N_{px} for the speckle contrast method, and $N_{\text{px}} \ln(N_{\text{px}})$ for the digital holography method.⁶

⁶ Assuming that two-dimensional fast Fourier transform is used.

Table 5.1: Qualitative differences between UOT detection schemes. Reprinted with permission from Paper IV © Optica Publishing Group.

| | Spectral hole burning | Photorefractive detection | Single-shot off-axis holography | Speckle contrast |
|---|---|---|---|---|
| Transmission/reflection imaging geometries? | Favours transmission (transmission alleviates requirement for high contrast filter) | Favours transmission | Independent | Favours transmission |
| Limiting expense | Cryostat and frequency stabilised light source (high cost) | Photorefractive material and vibration isolation (medium cost) | High pixel count and high frame rate camera (medium cost) | High pixel count and high frame rate camera (medium cost) |
| Technological advances which will make technique realisable | Crystal filters and cryostats | Photorefractive materials | Camera technology | Camera technology |
| Laser frequency stability requirement | High | Low | Low | Low |
| Optical wavelength flexibility | Limited to where suitable rare-earth ion-doped crystals exist | Limited to where materials exhibit photorefractive effect (Te:SPS has a broad wavelength range) | Not limited to specific wavelength | Not limited to specific wavelength |
| Sensitivity to speckle decorrelation | Insensitive | Sensitive | Sensitive | Sensitive |
| Data generation per frame | Low | Low | High (high bandwidth data stream required) | High (high bandwidth data stream required) |
| Etendue increases with: | Larger crystal sizes | Larger photorefractive material sizes | Increasing number of pixels (also increases data generated) | Increasing number of pixels (also increases data generated) |

Chapter 6

Ultrasound optical tomography using spectral-hole burning filters

Tomography: a method of producing a three-dimensional image of the internal structures of a solid object by the observation and recording of the differences in the effects on the passage of waves of energy impinging on those structures.

— Merriam-Webster Dictionary

With the spectral-hole filtration method in place, it is now possible to discuss the details of how one acquires an actual UOT image. As outlined in Section 2.4, image acquisition entails gathering the sideband intensity generated in an image plane, defined by the ultrasound propagation paths inside the tissue. The depth of the image point is determined by the deterministic propagation speed of the ultrasound. As an ultrasound propagates at $\sim 1.5 \text{ mm}/\mu\text{s}$, the depth of the image point may simply be determined by measuring the time difference between the image-point acquisition and the emission of the ultrasound pulses. For lateral resolution, the ultrasound source itself must somehow be manipulated. For a single-element transducer, this is achieved by mechanical movement of the transducer head [34, 106], while for a linear array transducer, the lateral resolution can be achieved from parallel lines of ultrasound pulses, as in [114]. In the case of a phased array transducer, the ultrasound pulse can be steered in an arbitrary direction.

Section 6.1 discusses the experimental setup and results in Papers I–III, where a medical ultrasound machine was used in its pulsed Doppler mode. This generated the equivalent of a single-element ultrasound, which was used to capture UOT signals in homogeneous media. In Section 6.2, the experimental details of Paper V are discussed, where the same ultrasound machine was used but with a linear array transducer. This yielded lateral resolution and enabled imaging in a two-dimensional plane of simple but heterogeneous media. In Section 6.3, the most recent experiments are presented, demonstrating UOT using SHF at a tissue-appropriate wavelength and with a transportable setup. In the experiment, a research ultrasound machine with a linear array transducer was used to perform UOT imaging on breast tissue, phantoms with phantom breast-tumour mimicking inclusions.

6.1 UOT signal strength in homogeneous media

In Papers I–III, the framework for performing a UOT measurement was studied, using the $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ filter in Section 4.3.1. The aim of the first experiment to be discussed, presented in Paper I, was to both verify the calculations in [36] and explore the technical limitations of the absolute signal strength achievable in UOT.

6.1.1 Experimental setup

The setup for these experiments can be seen in Figure 6.1a. The laser system was a Coherent CR-699 dye laser, operating with rhodamine 6G at 606 nm. The laser was locked to an ultra-low-expansion reference cavity using the Pound-Drever-Hall technique [123], achieving an optical linewidth of a few Hz. The $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ crystal was kept at the 2.17 K (the lambda point of helium) in a helium bath cryostat.¹ The same laser was used to create the filter and probe the tissue, and the optical path for each task was switched using a mechanised flip mirror with a 0.4 s switching time.

Double-pass acousto-optic modulator (AOM)

Light pulses were generated from the continuous wave laser source using an acousto-optic modulator (AOM) setup in a double-pass configuration. An AOM, or Bragg cell, is a small transparent crystal attached to a small piezoelectric transmitter. Aligning an optical beam through this crystal and driving the transmitter with

¹ Oxford Instruments Spectromag.




Figure 6.1: Illustration of the setups for the (a) acoustic field strength and (b) tagging efficiency measurements used in **Papers I–III**. *i*: motorised translation stage, *ii*: hydrophone probe *iii*: ultrasound (US) transducer. The US field was measured in the plane of the US focus (fixed *y* position). *iv*: Beam-path selector for filter creation or probing, *v*: Intralipid phantom with US, *vi*: waveguide and spectral-hole filter, *vii*: PMT detector with shutter. Reprinted with permission from **Paper I** © Optica Publishing Group.

an RF frequency,² f_{AOM} , generates a sound wave, which in turn – via the acousto-optic effect – generates a travelling index of refraction inside the crystal. The incident beam diffracts from the grating in a new direction and is frequency shifted by the frequency of the sound wave (i.e. $f_{\text{diffracted}} = f_{\text{incoming}} + f_{\text{AOM}}$) [124]. By changing f_{AOM} and the driving RF power, an AOM may thus be used as both an intensity modulator and a frequency shifting element, a device allowing us to first generate a filter and then probe the tissue with light shifted in frequency by f_{US} .

An issue with using only a single pass through an AOM is that the diffracted angle of the beam is dependent on f_{AOM} , which causes problems if one wishes to couple different frequencies of light into the same optical fibre, for example. This can be solved with the double-pass configuration, where – in short – the diffracted beam is mirrored back onto the AOM and, aligned correctly, the twice-diffracted beam has no angle-dependence on f_{AOM} [37, 125]. A double-pass AOM thus allows for arbitrary frequency shifting within twice the working range of the AOM (i.e. $f_{\text{diffracted}} = f_{\text{incoming}} + 2f_{\text{AOM}}$), while simultaneously allowing for arbitrary intensity-modulation via modulation of the sound wave amplitude in the AOM crystal. In our experimental setup, this ultrasound wave is produced by an arbitrary waveform generator (AWG), a typically networked device that can be programmed to create the many sequences of pulses required to generate the spectral holes for SHF and probe the tissue. The AWG can also be used to generate the TTL pulses used to control the position of the flip mirror.

² Typically 50-200 MHz.

Table 6.1: Parameters for the setup used in the experiments and for an improved setup. Reprinted with permission from *Paper I* ©2019 Optica Publishing Group.

| Parameter | Symbol | Experiment | Improved |
|---|---------------------------|---------------------------|---------------------------|
| Reduced scattering coefficient | μ'_s | 6.1 cm^{-1} | - |
| Absorption coefficient | μ_a | 0.008 cm^{-1} | - |
| Photons per input pulse | \mathcal{N}_{in} | 8.1×10^{10} | 2.5×10^{13} |
| Laser probe FWHM | τ | $1.0 \text{ }\mu\text{s}$ | $1.0 \text{ }\mu\text{s}$ |
| Probe repetition rate | - | 1.25 kHz | 25 kHz |
| Light guide coupling efficiency | η_{LG} | 20% | 20% |
| Light guide aperture area | A_{LG} | 0.79 cm^2 | 0.79 cm^2 |
| Transmission from light guide output to PMT | T_{optics} | 0.05% | 20% |
| Slow light filter transmission | T_{F} | 60% | 100% |
| Detector quantum efficiency | QE_{det} | 1.7% | 15% |
| Tagging factor (+1st order sideband) | κ | 0.026 cm^2 | 0.026 cm^2 |
| Number of sidebands used | N_{side} | 1 | 2 |
| Total setup efficiency | η_{tot} | 10^{-6} | 0.6% |

Experimental sequence and beam paths

With the flip mirror set to the probe path, the light illuminates a tissue-mimicking phantom, from which light is collected using a liquid-core light guide.³ The phantoms were $7 \times 7 \times L \text{ cm}^3$, where L is the thickness of the phantom along the z -direction (L was varied during the experiments). This light is guided to the cryostat housing the $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ filter, which hits the detector after filtration. The detector was a photomultiplier tube (PMT).⁴ Its quantum efficiency was in turn measured using a heavily attenuated HeNe laser and a power meter. While optically pumping the filter, the detector was closed using a shutter.⁵

In the filter-creation path, light illuminates the crystal counter-propagating the probe path. The setup required a beam splitter to allow for both detector and burn-path input on the same side of the cryostat. For this a 50:50 beam splitter was chosen. For the filter preparation, the filter centre frequency (see Fig. 4.3) was set to the tagged-light frequencies: $f_{\text{pump}} = f_{\text{C}} + f_{\text{US}}$ (for measuring the first-order tagged light), $f_{\text{pump}} = f_{\text{C}} + 2f_{\text{US}}$ (for measuring the second-order tagged light), or the carrier frequency $f_{\text{pump}} = f_{\text{C}}$ (for measuring the untagged light), where f_{C} is again the carrier frequency used when performing the UOT measurement.

The generated filter was the same as the one discussed in Section 4.3.1 and had a $T_{\text{F}} = 60\%$ and, when measuring tagged light, a $T_{\text{C}} \approx -32 \text{ dB}$ for diffuse light.

³ Rofin Australia Pty Ltd, 0.59 numerical aperture and 1 cm diameter.

⁴ Hamamatsu, R943-02, $10 \times 10 \text{ mm}^2$ photocathode effective area.

⁵ Vincent Uniblitz VS14.

The ultrasound source was the Phillips Epic-7 system, with the X5-1 phased array transducer used in **Paper I** and both the X5-1 and L12-3 linear array transducers used in **Papers II–III**. The ultrasound pulses generated by these transducers were measured using a needle hydrophone in a water tank, depicted in Figure 6.1. The pulses were measured in a plane 3.5 cm from the surface of the transducer. The phantoms were created using a recipe based on [126], with Intralipid 20%, water, and agar. The phantoms’ optical properties were measured using time-of-flight spectroscopy [21] to $\mu'_s = 6.1 \text{ cm}^{-1}$ and $\mu_a = 0.008 \text{ cm}^{-1}$. As the phantoms were more than 95% water and the acoustic velocity in the phantom was approximately that of water, the difference in the other acoustic properties was assumed to be negligible. The ultrasound pulse in the centre of the phantom in Figure 6.1a could thus be assumed to be identical to that measured in Figure 6.1b.

With the flip mirror set to the probe arm, a 1 μs full-width at half-max Gaussian light pulse with peak intensity of $\sim 25 \text{ mW}$ was synchronised to illuminate the phantom tissue when the ultrasound pulse had reached the centre of the tissue phantom. The timing of the optical pulse to the ultrasound position was set by delaying the ultrasound pulse TTL marker output before sending it to trigger the AWG, which generated the optical pulse. This delay was implemented using a digital delay box,⁶ allowing the timing to be changed quickly and without having to reprogramme the AWG sequence each time.

Finally, all loss terms in the UOT setup were identified and quantitatively measured. These loss terms are shown in Table 6.1.

6.1.2 Absolute signal comparison (Paper I)

The experimental measurements were compared to a quantitative analytic model, based on the diffusion approximation detailed in **Paper I**, Appendix C. As shown by the parameters in Table 6.1, this model predicted the number of detected carrier photons, \mathcal{N}_C , and tagged photons, \mathcal{N}_T :

$$\mathcal{N}_{C,T} = \mathcal{N}_{\text{in}} N_{\text{side}} \eta_{LG} T_{\text{optics}} T_F \mathcal{T}_{C,T} \quad (6.1)$$

where the carrier \mathcal{T}_C and sideband \mathcal{T}_T transmissions are the model-predicted ratios of the carrier or sideband light hitting the light guide over the input power, thus \mathcal{T}_C and \mathcal{T}_T are proportional to the light guide area A_{LG} . Using the light-diffusion approximation, \mathcal{T}_C and \mathcal{T}_T were calculated using the fluxes $J_{\text{out},C,T}$ at the light guide, calculated using Fick’s law (2.6) from the optical fluences $\Phi_{C,T}(\mathbf{r})$ inside

⁶ Stanford Research Systems, Inc. Model DG535.




Figure 6.2: Positions of sources and drains in the slab diffusion model. Each source outside the medium was generated by inverse mirroring of the two sources inside the medium over the extrapolated boundaries. Green denotes the carrier and blue the tagged. Reprinted with permission from **Paper I** © Optica Publishing Group.

the slab. These fluences were modelled analytically, using the point source solution (2.11) and inversely mirroring (source \leftrightarrow drain) the original carrier and tagged-light point sources over extrapolated boundaries. The offset of these extrapolated boundaries from the real boundaries was equal to $2AD$, where D is the diffusion constant and $A = 2.35$ for an air-water boundary [127]. The field inside the domain of interest (i.e. the slab) was then the sum of all the sources and drains generated by the inverse mirroring, whose positions are depicted in Figure 6.2. The original source for the carrier-light fluence Φ_C was located at a depth of $1/\mu'_s$ in the tissue (i.e. the depth where an isotropic source estimates a thin, collimated beam hitting the surface of the medium). The original tagged-light fluence Φ_T source was located at the ultrasound pulse position at the centre of the phantom. In contrast to the carrier source, the magnitude of the tagged-light source is not a simulation constant but rather scales with $\kappa\Phi_c(\mathbf{r} = \mathbf{r}_{US})$: in effect, it is modelled using the shortcut diffusion model (3.4), where κ is a parameter that defines the ultrasound interaction. Thus, using the diffusion model, $\mathcal{T}_T \propto \kappa$.

Figure 6.3a shows the experimental results and the results of this model together with the predicted carrier and tagged-light signal. The model for the carrier is entirely without fitting parameters, and the model for the tagged-light is fitted using κ as the only free parameter. In addition to the diffusion-model fit to the data presented in **Paper I**, Figure 6.3a also contains results for $\mathcal{T}_{C,T}$ calculated using the SMC model with the pulse in Figure 6.4. These traces are entirely first-principle and do not require any fitting parameter κ for the tagged signal.

With the theory being very well-matched to the experimental results, we can start to examine ways of finding and then improving the absolute limit of imaging with UOT. The improvements estimated to be possible at the time of writing **Paper I** are presented in Table 6.1. These parameters were used (with 6.1) to estimate achievable signal levels when imaging through breast tissue ($\mu'_s = 11 \text{ cm}^{-1}$, $\mu_a = 0.05 \text{ cm}^{-1}$) and muscle tissue ($\mu'_s = 5 \text{ cm}^{-1}$, $\mu_a = 0.2 \text{ cm}^{-1}$), the result of which can be seen in Figure 6.3b.

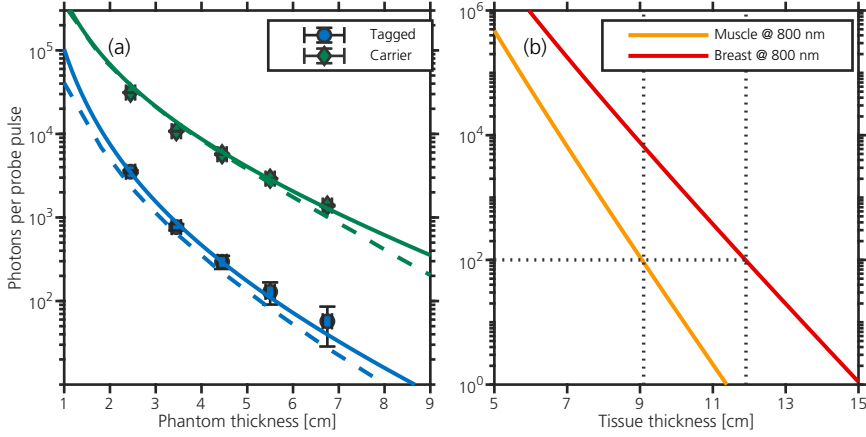


Figure 6.3: (a) Measured tagged and carrier signals compared to the diffusion model and the SMC model. The phantoms had $\mu'_i = 6.1 \text{ cm}^{-1}$ and $\mu_a = 0.008 \text{ cm}^{-1}$. Signal values were obtained by averaging 1,000 probe pulses. The vertical and horizontal error bars represent two standard errors and our estimated phantom thickness accuracy, respectively. The full lines correspond to the diffusion model and the dashed to the SMC model. (b) The diffusion-model predicted number of tagged photons incident on the detector per probe pulse for an optimised UOT transmission-mode setup probing muscle and breast tissue. The ultrasound depth from which tagged photons are generated is set to half the tissue thickness. Adapted with permission from **Paper I** © Optica Publishing Group.

In the three years since the publication of **Paper I**, no evidence indicating that any of the optimal experimental parameters are unrealisable have come forth, except perhaps $T_F = 100\%$, which is very challenging to achieve in a strongly absorbing crystal ($\alpha_B L > 18$ for a crystal length L). The estimation that $T_{\text{optics}} = 20\%$ may actually be slightly pessimistic, as $T_{\text{optics}} = 9.4\%$ was achieved in the setup in Section 6.3, and this was with a crystal area one quarter of what the optics were designed for. Instead, ray-trace simulations for these optics together with the correct crystal size predict that $T_{\text{optics}} \geq 50\%$.

6.1.3 Investigating UOT tagging efficiency (Paper II)

The acoustic setup in Figure 6.1 was used to measure ultrasound pulse shapes with 1.6 MHz and 3.5 MHz centre frequencies. These pulses were generated with the X5-1 phased array transducer and the L12-3 linear array transducer, respectively. The transducers were used to produce pulses with three and two different pulse lengths. In Figure 6.4, one such pulse is depicted with its spectral components displayed.

It was found that this pressure pulse and the others measured could be adequately




Figure 6.4: Characterisation of the X5-1 4mm ultrasound pulse, propagating toward positive y . (a) The pulse in the $x = 0$ plane, where the dashed lines correspond to the fitted envelope. (b) The frontal profile of the pulse in the xz -plane, where the dashed line corresponds to where the pulse is evaluated in (a), (c) and (d). (c) The projection of the pulse on the plane $y = -7$ mm. (d) is a top-down view of the pulse. (e) The normalised Fourier transform of the pulse centreline with $f_1 = 1.6\text{MHz}$. The peaks of this spectrum correspond to the values a'_n , from which $a_n = a'_n / \sum a'_n$ is calculated. (f) The a'_n coefficients for different peak pressures. Reprinted with permission from **Paper II** © Optica Publishing Group

modelled with the template

$$P_1(\mathbf{r}, t) = \Gamma(x-x_0, Ct-y+y_0, z-z_0) \sum_{n=1}^N a_n \sin[nK(Ct-y+y_0) + \varphi_n] \quad (6.2)$$

where the pressure envelope Γ is defined as follows:

$$\begin{aligned} \Gamma(x, y, z) = \Gamma_0 \exp\left(-\frac{y^2}{2\sigma_y^2}\right) & \left[R_1 \text{sinc}^2\left(\frac{x}{\sigma_{x1}}\right) \text{sinc}^2\left(\frac{z}{\sigma_{z1}}\right) \right. \\ & + R_2 \exp\left(-\frac{x^{N_x}}{N_x \sigma_{x2}^{N_x}}\right) \exp\left(-\frac{z^2}{2\sigma_{z2}^2}\right) \\ & \left. + R_3 \exp\left(-\frac{x^2}{2\sigma_{x3}^2}\right) \exp\left(-\frac{z^{N_z}}{N_z \sigma_{z3}^{N_z}}\right) \right]. \end{aligned} \quad (6.3)$$

In (6.2), $K = 2\pi f_{\text{US}}/C$ is the wave number with C being the speed of sound. The acoustic harmonics in the ultrasound pulses were truncated to the first 5 (i.e. $N = 5$). As $\sum a_n = 1$ and $\varphi_n \equiv \pi/2$, Γ_0 in (6.3) represents the peak pressure of the pulse if the constraint $\sum R_i = 1$ is added. The parameter R_i then represents the relative size of the terms within the large square brackets in (6.3). These terms represent a central lobe and “wings” in the x and y directions of the pulse, as seen in Figure 6.4, which are most adequately described by the super Gaussian function.




Figure 6.5: Tagged fractions as a function of peak pressure for the five investigated pulses. Full blue line is the SMC-simulated tagged fraction from the fitted US pulses. The shaded areas around each line depict the standard errors of the means of the simulations. The red dots depict the mean values, and the bars the standard deviations of the experimental measurements. Reprinted with permission from **Paper II** © Optica Publishing Group

The sharpness of the falloff of the wings was set by fitting the even integers N_x and N_y . The fitting parameters of all measured pulses can be found in Appendix B of **Paper II**. With these fitted pulses, the SMC model could be validated and the dependencies of κ investigated.

To compare the model to the data with minimal experimental impact, the tagged fraction η_n of sideband n can be used. This quantity is defined as

$$\eta_n = \frac{\mathcal{N}_{T,n}}{\mathcal{N}_{\text{ref}}} \quad (6.4)$$

where $\mathcal{N}_{T,n}$ is the number of detected tagged photons in the n -th sideband and \mathcal{N}_{ref} is a measurement of \mathcal{N}_C at $P_1 \equiv 0$. The total tagging efficiency η_{all} of all sidebands can similarly be defined as

$$\eta_{\text{all}} = \frac{\mathcal{N}_{\text{ref}} - \mathcal{N}_C}{\mathcal{N}_{\text{ref}}} \quad (6.5)$$

where \mathcal{N}_C is again the detected carrier photons. This quantity is insensitive to linear experimental losses, as these are applied equally to the nominator and denominator. The carrier and tagged signals were measured for the five different pulse shapes at different peak pressures in a 3.8 cm-thick phantom. From these measurements, η_{all} , η_1 , and η_2 were calculated, and the results – as published in **Paper II** – are presented in Figure 6.5, together with the respective SMC-model-predicted counterparts.

Figure 6.5 shows that the tagged fraction of the sidebands are well predicted by the SMC model, but the total tagged fraction is underrepresented by η_{all} . This discrepancy has been the focus of much internal discussion, and the exact source has not yet been verified. As the UOT Monte Carlo model has been validated in the literature many times [38], the discrepancy is attributed to unidentified nonlinear experimental parameters which are not cancelled out when calculating η . One such possible parameter is a loss or saturation term that is $\sim K$ for strong signals (i.e. \mathcal{N}_{C}) and ~ 0 for small signals (i.e. \mathcal{N}_{T}). Such a loss term can originate in a reduced photocathode sensitivity at higher light-intensities for the PMT detector. This effect was not calibrated for, as the measurement of the PMT quantum efficiency was performed at low signal levels. This loss term leads to an increase in both η_n and η_{all} , though the change in η_n is negligible in comparison to the error of the experimental results. Another possible cause of the discrepancy is the low amount of absorption in the filter, which is saturated for high signals (i.e. \mathcal{N}_{C}) but not low signals (i.e. \mathcal{N}_{T}). This would entail a partially nonlinear T_{F} . As T_{F} in Table 6.1 was measured with high signals, this loss term should not affect the carrier measurements. In the low signal regime, this absorption is not saturated, which would lead – in the first approximation – to a linear loss of the signal, with a proportionality constant $\mathcal{A}_{\text{unsat}}$. These two nonlinear parameters transform the terms in (6.4) and (6.5) to

$$\begin{aligned}\mathcal{N}_{\text{ref}} &\rightarrow \mathcal{N}_{\text{ref}} - K \\ \mathcal{N}_{\text{C}} &\rightarrow \mathcal{N}_{\text{C}} - K \\ \mathcal{N}_{\text{T}} &\rightarrow (1 - \mathcal{A}_{\text{unsat}})\mathcal{N}_{\text{T}}.\end{aligned}\tag{6.6}$$

Compensating the data with fitted K and $\mathcal{A}_{\text{unsat}}$ yields the results shown in Figure 6.6 for a K equivalent of a 30% reduction in radiant sensitivity at high intensities and $\mathcal{A}_{\text{unsat}} \approx 20\%$. Furthermore, these compensations applied to the data in Figure 6.3 still fit the model, with the caveat that K must decrease with the signal level and that κ becomes 20% larger. However, this discussion of the discrepancy is only that: a discussion. Further experimental investigation is necessary before any conclusions can be drawn.

6.1.4 Lessons from single point UOT measurements

The results presented in this section so far are the first ones acquired during the underpinning thesis work, and some of the first ever captured in Lund. It may therefore be an understatement to say that many experimental insights arose during this time that was to be of real importance in future experimental work. One




Figure 6.6: Tagged fractions as a function of peak pressure for the five investigated pulses, where the experimental values have been adjusted for an hypothesised low signal level absorption of 20%, which is saturated at high signals, and a 30% drop in photocathode sensitivity for high signal levels. Full blue line is the SMC-simulated tagged fraction from the fitted US pulses. The shaded areas around each line depict the standard errors of the means of the simulations. The red dots depict the mean values, and the bars the standard deviations of the experimental measurements.

such insight concerns the use of specular-reflecting film on the input and output surfaces of the phantom to increase boundary reflectivity. To accommodate for the input and output of the light, a 1.6 mm diameter hole was cut for the beam on the input side and a 1.2 cm diameter hole was cut for the light guide. Reflecting films can, for instance, be found in LCD screens, where they are at the back of the screen to increase the forward radiance of the generated light. The impact of the films on UOT can be understood by returning to the diffusion model depicted in Figure 6.2, where the variable $A \propto (1 - R)^{-1}$ (R is the reflectivity of the bound) sets the location of the extrapolated bound. Thus, as $R \rightarrow 1$, all drains are removed and only the original sources affect the signal, which has a positive effect on the signal level.

The use of these films was not reported in **Papers I–III**, as their use in UOT was included in a patent application, though all the experiments which produced the data in Figure 6.3a were repeated using reflecting film on the output and input surfaces. The impact of reflecting film on the signal is seen in Figure 6.7, where the number of tagged photons without film⁷ is compared to the number of detected photons with film. The net result is an increase in the average signal level by a factor of 2.8. Furthermore, the film has an effect when placed on both the input and output side, which can be understood by again examining the placement of the extrapolated boundaries in Figure 6.2. In qualitative terms, placing reflecting film on the input forces the carrier flux into the medium, allowing more light to

⁷ The same data shown in Figure 6.3a.




Figure 6.7: The impact of reflecting film on a transmission measurement where blue is with film and red is without film. (a) The tagged signal as recorded through a 6.8 cm thick phantom. (b) The number of photons as a function of phantom thickness. (c) The Lund University logo (to the left) mirrored in a square piece of reflecting film (to the right).

be tagged, while film on the output side prevents light exiting the medium except into the light guide. As the inclusion of reflecting film was not intrusive to the measurement and the impact was significant, reflecting films were used to boost signal levels in all future experiments.

Another important lesson worth reiterating is that it was possible to determine, with confidence, all experimental losses – making it in turn possible to predict absolute signal levels with the presented models. The many loss terms recall the saying, “many brooks make a strong river”, as improving the total signal level requires many small, individual improvements. Many of the measured parameters in Table 6.1 were also used for the signal level appreciation in **Paper IV**, and an understanding of the losses was vital for performing the absolute signal reconstruction in **Paper V**.

6.2 Imaging of heterogeneous phantoms and image-reconstruction, Paper V

After a single-point measurement had been acquired, the next natural experimental step was full-image capture. As discussed at the start of this chapter, this required the ultrasound pulse to have the ability to quickly perform lateral movements. This was achieved using the same ultrasound machine (Phillips EPIC7) used in **Papers I–III**. With the help of Phillips engineer Jeff Powers, a specific UOT sequence using the L12-3 transducer was programmed into the EPIC7. This sequence emits ultrasound pulses at 6 MHz in 30 equidistant and parallel lines, where the pulses were designed to be as “narrow” as possible to achieve good lateral resolution. This




Figure 6.8: Two of the phantoms created for imaging. (a) The moment before the inclusion was covered by another layer of phantom material. (b) A finished phantom is viewed from the side, and the end of the including rod can be seen as well.

pulse was again characterised using the setup in Figure 6.1b) and the pulse could be fitted using (6.2) and (6.3), with the fitting parameters given in Appendix A of Paper V.

Heterogeneous tissue phantom creation

The ultrasound source, once characterised, was used to image heterogeneously absorbing phantoms. These phantoms were made following the same recipe as in Papers I–III [126] but were created in two stages. First, a mould with side dimensions $9.5 \times 9.5 \text{ cm}^2$ was used to cast the bottom half of the phantom. After this layer had solidified, a previously prepared rod of the inclusion material was placed on top of it. To create a contrast in the image, the phantom material of this rod contained some concentration of Indian ink, which is strongly absorbing but weakly scattering [128]. A second layer of phantom material was then cast in the mould, covering the inclusion rod. This procedure allowed for the creation of phantoms with heterogeneous inclusions at various depths. The rod shape of the inclusion was chosen to make the alignment of the image easier and insensitive to translation along one dimension. Figure 6.8 provides a snapshot taken during the creation of one such phantom, with a finished phantom where the end of the rod-shaped inclusion can be seen from the side. The scattering and absorption of all parts of the phantom were measured using time-of-flight spectroscopy [21]. For the inclusions, the measurement was performed in a bulk phantom, from which the inclusion was cut. The background properties were measured on the phantoms after we had performed the UOT measurements in a region away from the inclusion.




Figure 6.9: Experimental setup in Paper V. An AWG controls the pulsed light source (LS). One of two paths is selected by the AWG using a motorised flip mirror (FM). With the FM in, the beam from the LS enters the “probe path”, where light enters the tissue phantom and is collected with a light guide (LG). Inside the phantom, the light interacts with pulses from the ultrasound transducer (UST) in the image plane (IP). The light from the LG is imaged through the crystal onto a photomultiplier tube (PMT). A linear polariser (LP) sets the polarisation parallel to the crystal D2-axis. With the FM out, the beam enters the “filter-preparation path”. Here, the beam passes through a beam expander (BE) and illuminates the crystal directly where the closed shutter protects the PMT from stray light-fields. At the end of each b-mode sequence, the ultrasound source (USS) emits a TTL pulse, which is delayed in the tuneable delay box (DB) before triggering the AWG sequence. The AWG also controls the shutter.

6.2.1 Experimental setup

These phantoms were imaged in a reflection-mode setup: in effect, the light input, light collection, and ultrasound source were all placed on the same tissue surface. This configuration is interesting, as it is how a general UOT probe would be constructed. The total experimental setup is depicted in Figure 6.9. The setup contained some of the same components as the setup discussed in Section 6.1, but with some key differences. The first such difference was the direction of the optical pumping beam path. This previously propagated in the opposite direction to the light from the tissue, but here illuminates the crystal perpendicularly to the direction of the light from the phantom. This allowed for the losses from the beam splitter used to be removed entirely from both experimental paths. Another improvement of the losses was the use of a smaller cryostat⁸ with a larger numerical aperture. However, this cryostat had less cooling power than the previous one, which meant that the crystal could only be kept at <6 K. Nonetheless, the increased optical access overshadowed this loss in cooling power, as the transmission through the optics improved from $T_{\text{optics}} = 0.05\%$ to $T_{\text{optics}} = 6\%$.

Another addition to the setup was a small magnetic field to increase T_h . This

⁸ Oxford instruments Optistat CF-V, a helium flow cryostat.



Figure 6.10: Geometries of the SMC simulations used to compare the reflection-mode imaging experimental results. (a) The geometry used when comparing the experimental results directly in Figure 6.11. (b) The voxel geometry used to reconstruct the absorption in Figure 6.12.

~ 10 mT field along the crystal D_1 axis was generated using two coils set in an almost Helmholtz configuration (the width of the cryostat was slightly larger than the radii of the used coils). The magnetic field strength and homogeneity at the crystal was instead estimated using a magnetic-field simulation tool [129]. Overall, this effort yielded $T_F \approx 80\%$, a slight improvement on the previous $T_F \approx 60\%$. The carrier attenuation for this filter was weaker, with $T_C \approx -27$ dB. This lower carrier suppression was mainly due to a lesser slow-light effect arising from a wider filter, $W_{\text{SHF}} = 2$ MHz. In turn, this led to both direct leakage (due to the rounded corners of the filter) and a diminished slow-light effect.

Synchronisation of the light and ultrasound pulses was performed as in Section 6.1 (i.e. using a digital delay box). This time, however, the TTL output used to trigger the light pulse was markers of the first of the 30 ultrasound pulses emitted by the ultrasound, a so-called “frame marker”. After being triggered by a frame marker, the AWG was programmed to generate 30×25 optical pulses. These optical pulses were separated in time by the ultrasound pulse-repetition rate $\mathcal{F}_{\text{US}} \approx 7.5$ kHz. For every 30th pulse, however, the next optical pulse was delayed by $0.5 \mu\text{s}$ or $0.6 \mu\text{s}$ in order for the next 30 probed points to be performed at a slightly larger depth. The zero-depth delay was measured by imaging the edge of the phantom – in effect, a delay where the UOT image was abruptly cut in half due to a lack of tagged light before an ultrasound pulse entered the medium. The measurements of zero depth also gave information about T_C and the noise, as image points taken at times where the ultrasound pulse had not yet entered the tissue could only be comprised of carrier and stray light hitting the detector. In total, each image consisted of $N = 30 \times 25$ image points and took ~ 100 ms to capture.

6.2.2 Results and discussion

Comparison of the experimental and simulated images

One of the captured images can be seen in Figure 6.11, where – in addition to the actual experimental image – a comparison using the SMC model and showing the

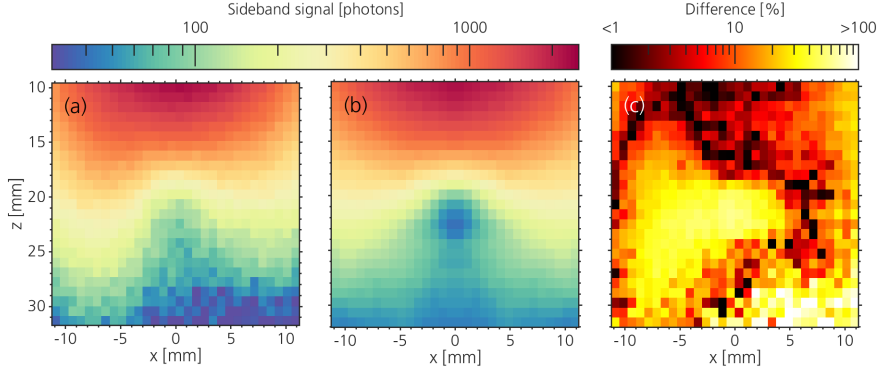


Figure 6.11: Simulation vs experiment for a phantom with $\mu'_s = 4.5 \text{ cm}^{-1}$, $\mu_a = 0.024 \text{ cm}^{-1}$ in the background and $\mu_a = 2.4 \text{ cm}^{-1}$ in the inclusion. The inclusion was at $z = 21 \text{ mm}$, $x = 0 \text{ mm}$ and had the dimensions $5 \times 5 \times 80 \text{ mm}^3$. (a) Experimental image. (b) Projected SMC image. (c) Percentage difference between experiment and simulation, mean difference = 21%. The signal in (a–b) is the incident number of photons on the detector.

difference between the two images is also presented. The modelled signal values were estimated using (6.1), with $\mathcal{T}_T \rightarrow \mathbf{F}$ for the given imaging sequence, \mathbf{F} , defined in Section 3.3.5. The voxel geometry for the simulation with $V = 2$ voxels is seen in Figure 6.10a, where one voxel is the same as the inclusion in the phantom and the other, irregular voxel is everything else in the phantom (i.e. the background).

The comparison in Figure 6.11 shows an average error of 21%, with a large discrepancy in the centre (the location of the inclusion) and the bottom-right. As the difference is dependent on perfectly matching the experimental image to the simulated image, a small misalignment would result in an error. Furthermore, in order to get the light guide, light input, and transducer as close as possible, the transducer and light guide were held with mechanical claws that were as little in the way as possible. They did, however, introduce some difficulties to the aligning of the light input and light guide to the centre of the transducer. Such a misalignment would again introduce differences between the experimental and simulated images. Failing to align all three components along the centre of the transducer would allow a non-symmetric light fluence in the image, which could explain the artefact in the lower-right corner.

With these caveats out of the way, the agreement between the simulated and experimental images is good, considering it was achieved without fitting to the data and using only *a priori* information. However, in this quest to image the inclusion, Figure 6.11a is somewhat lacking, as the size and shape of the inclusion remains open to interpretation. This ambivalence is due to the fact that the fluence of light is exponentially decreasing along z in the image, causing a fluence-distribution




Figure 6.12: Reconstructed μ_a with resulting image $\eta\mathbf{F}(\mu_a)$, using experimental image of a phantom with $\mu'_s = 4.5 \text{ cm}^{-1}$, $\mu_a = 0.024 \text{ cm}^{-1}$ in the background and $\mu_a = 2.4 \text{ cm}^{-1}$ in the inclusion. The inclusion was at $z = 21 \text{ mm}$ $x = 0 \text{ mm}$ and had the dimensions $5 \times 5 \times 80 \text{ mm}^3$. (a) Reconstructed absorption image, $\mu_a = 0.025 \text{ cm}^{-1}$ outside the image. (b) SMC image obtained using the reconstructed absorption. (c) Percentage difference between the experimental and SMC images, with a mean difference of 7.9%. The signal in (b) is the incident number of photons on the detector.

artefact. This artefact is – to a lesser degree – also visible along the x dimension of Figure 6.11a, again due to the reduced light fluence away from $x = 0$. Furthermore, Figure 6.11a gives no details about the absolute absorption of the imaged structure.

Reconstruction of underlying absorption

One solution to both of these issues is to use the SMC model to reconstruct the underlying absorption by solving the inverse problem. The simulation voxel geometry on which this was attempted can be seen in Figure 6.10b, where the number of voxels is $V = N + 1$ (with N being the number of measurement points). This geometry confines the absorption to one large background and irregular voxel, as well as N rods of the same width and height position as the pixels in Figure 6.11 and the same length as the inclusion.

The reconstruction was performed by solving the inverse problem using the objective function

$$\Omega(\mu_a) = \|\eta\mathbf{F}(\mu_a) \oslash \mathbf{b} - 1\|_2^2 \quad (6.7)$$

where $\mathbf{b} = [b_1, \dots, b_n, \dots, b_N]^T$ denotes the vectorised experimental image, such that the elements in \mathbf{F} correspond to the experimental image points in \mathbf{b} , and $\eta = \eta_{\text{LG}} T_{\text{optics}} T_{\text{F}}$ are the losses incurred between the light guide and detector. The \oslash symbol in (6.7) denotes the Hadamard division [130], which is performed

to normalise the error function and simplify the use of standard solvers already implemented in many programming languages (e.g. MATLAB).

The minimisation of (6.7) was performed using the Levenberg-Marquardt algorithm [131] implemented in MATLAB's optimisation toolbox [132]. This algorithm can be sped up by providing the analytic Jacobian \mathbf{J} of the expression inside the norm in (6.7), in effect,

$$\mathbf{J} = -\mathbf{A}_b [\mathbf{D} \circ \overline{\exp}_V(-\mathbf{D}\boldsymbol{\mu}_a)] \quad (6.8)$$

where each row in \mathbf{A}_b is \mathbf{A}_n/b_n (see Chap. 3), \circ denotes the Hadamard product (i.e. the element-wise multiplication of two matrices of the same size [130]), and $\overline{\exp}_V$ denotes the V times column-wise repetition of \exp , making $\overline{\exp}_V(-\mathbf{D}\boldsymbol{\mu}_a)$ a matrix with $M \times V$ elements, M being the number of paths and V the number of voxels (see Sect. 3.3.5).

The experimental image in Figure 6.11a was used to reconstruct the absorption in two minutes,⁹ and the resulting image is seen in Figure 6.12a. This was considerably faster than the several hours reported in [50], which used the shortcut-diffusion model for the light-ultrasound interaction. The reconstructed absorption accurately reflects the general size and absorption of the inclusion. The error between the reconstructed image (Fig. 6.11b) and the experimental image (Fig. 6.11a) can be seen in Figure 6.12c, which is considerably less erroneous than Figure 6.11c, where the average error is now 7.9% (down from 21% in Fig. 6.11c). The absorption of the inclusion is slightly lower than expected, at $\sim 1.9 \text{ cm}^{-1}$, in contrast to the $\sim 2.4 \text{ cm}^{-1}$ measured on the inclusion material prior to the phantom creation.

Fast correction of the fluence-distribution artefact

If two minutes is too long to wait to image the relative absorption, an image of the inclusion can be acquired more quickly using another approach. By simulating a UOT image taken in a homogeneous medium,¹⁰ one obtains the reference image \mathbf{F}_{ref} . If the fluence-distribution artefact stays roughly the same in the absence of the inclusion, an image proportional to the absorption is obtained as $\mathbf{b} \oslash \mathbf{F}_{ref}$. This can be used to remove the fluence-distribution artefact. Furthermore, a database of \mathbf{F}_{ref} 's, calculated using different homogeneous medium properties, can be prepared to allow an operator to quickly test the different references. Figure 6.13 shows this image refinement method for three of the imaged phantoms, each with

⁹ PC specification: 32 GB RAM, Intel i7 4.20 GHz CPU.

¹⁰ Using the SMC model or any other model estimating the UOT signal.




Figure 6.13: Fluence distribution artefact removal using \mathbf{F}_{ref} of images in three phantoms where the dashed white lines depict the position of the inclusions. (a) Image of the same phantom imaged in Figures 6.11–6.12. (b) Image in a phantom with $\mu'_s = 7.1 \text{ cm}^{-1}$, $\mu_a = 0.024 \text{ cm}^{-1}$ in the background, and $\mu_a = 0.9 \text{ cm}^{-1}$ in the inclusion. The inclusion was at $z = 13 \text{ mm}$, $x = 0 \text{ mm}$ and had the dimensions $5.5 \times 5.5 \times 70 \text{ mm}^3$. (c) Image in a phantom with $\mu'_s = 7.1 \text{ cm}^{-1}$, $\mu_a = 0.024 \text{ cm}^{-1}$ in the background, and $\mu_a = 0.6 \text{ cm}^{-1}$ in the inclusion. The inclusion was at $z = 14 \text{ mm}$, $x = 0 \text{ mm}$ and had the dimensions $5 \times 5 \times 90 \text{ mm}^3$. The image in (a) is slightly longer, as each line here was separated by $0.6 \mu\text{s}$ compared to $0.5 \mu\text{s}$ for (b–c).

different inclusion absorptions and background-scattering, where \mathbf{F}_{ref} was generated using the measured background-scattering coefficient and absorption.

6.2.3 Lessons from first UOT imaging experiments

Fluence-distribution artefact

As with the single-point measurement, this first imaging experiment came with some lessons that were, in hindsight, obvious. The first one was the magnitude of the effect of the fluence-distribution artefact, which was strong enough that the inclusion was not clearly visible even in the raw image taken of an entirely black inclusion. This artefact can, as discussed, be circumvented via treatment of the images. However, these treatments carry some risk of allowing misinterpretations of the data. Thus, it is preferable to address the fluence in the actual measurement. In the reflection-mode setup, the fluence-distribution artefact along z cannot be addressed experimentally, but along x , it can be addressed by manipulating the shape of the light input and collection areas.

Image-acquisition speed

One detail that came to the fore during these measurements was the image-acquisition speed. As stated earlier, each image consisted of $N = 25 \times 30 = 750$

image points, where each image point was captured with an individual light-ultrasound pulse pair. The limiting factors were the average light-intensity and ultrasound pulse-repetition frequency, \mathcal{F}_{US} , where the latter was more limiting in terms of imaging speeds. Thus, the time taken to capture a single image was $N/\mathcal{F}_{\text{US}} = 100$ ms for the settings used. This could be shortened somewhat by probing the same ultrasound pulse with multiple light-pulses and increasing \mathcal{F}_{US} , although here hard limits quickly emerge, as light pulses cannot have too small separations if the slow-light effect is to be used, and \mathcal{F}_{US} can only be increased to 15–20 kHz before multiple ultrasound pulses can be seen propagating simultaneously in the tissue.

Another issue exacerbated by measuring the pairs of light-ultrasound pulses concerns the requirements of the data treatment. In the experiment, the signal from the detector was captured using an oscilloscope,¹¹ with each image consisting of 750 traces. These traces were transferred to a computer using TCP/IP over a dedicated ethernet cable. Due to the speed of the TCP/IP protocol (~ 10 MByte/s) and some pre-processing performed by the oscilloscope, this transfer took 4 s and the imaging and filter-creation sequences took a total of 300 ms. Even without this transfer wait, the imaging would include a 2×0.4 s long pause during each cycle, required by the flip mirror for switching between the two paths in the setup. In total, the possible duty cycle of the experiment was $\sim 33\%$, but due to the 4 s wait time, it was experimentally $\sim 2.5\%$. This limited data-transfer speed could be addressed by the use of a PCIe-connected capture card to load the captured data directly into the computer memory. Even the older PCIe 3.0 x4 version has a 4 GByte/s throughput capacity, easily allowing for a card capturing at 500 MS/s in four channels simultaneously. The switching between the two paths could be addressed by exchanging the flip mirror for a device with minimal mechanical movement. Examples of fast-switching devices include a Pockels cell and polarising beam splitter, a MEMS fibre-optic switch, and an AOM.

Implementation of these measures could potentially reduce the amount of dead time in the imaging cycle, including the time required to open the shutter protecting the detector from stray light-fields during the filter-preparation sequence. The time required to open and close the shutter.¹² was 4 ms

Reflecting film

Again, a piece of reflecting film was used in these experiments, this time cut to accommodate the ultrasound transducer, light guide, and input light. This film

¹¹ Teledyne Lecroy, WaveRunner HRO 66Zi.

¹² Vincent Uniblitz VS14.




Figure 6.14: UOT image in the phantom with $\mu'_s = 5.5 \text{ cm}^{-1}$, $\mu_a = 0.024 \text{ cm}^{-1}$ in the background and $\mu_a = 7.8 \text{ cm}^{-1}$ in the inclusion. The inclusion was at $z = 42 \text{ mm}$, $x = 0 \text{ mm}$ and had the dimensions $4.5 \times 4.5 \times 70 \text{ mm}^3$. (a) UOT image without any reflecting film. (b) Photo of reflecting film with holes for the transducer, light guide, and light input. (c) UOT image using the reflecting film.

– and the effect it had on the image – can be seen in Figure 6.14. As in the transmission case, the film increased the signal by a factor of approximately three. In the regime in which the CNR is dominated by the shot noise, this equates to an increase in CNR of a factor $\sqrt{3}$. As the CNR falls exponentially, this effectively only increases the depth at which the $\text{CNR} = 1$ by 2–3 mm (as seen in Fig. 5.2b). In conclusion, since including the film in a setup is so unobstructive, the benefit – though small – still outweighs the cost.

6.3 UOT in breast-tissue phantoms with tumour-like inclusions

So far, the UOT experiments discussed have all used optical wavelengths which have poor penetration in tissue (i.e. wavelengths outside the tissue-optical window). Moving to another wavelength required, first, a suitable filter material. This material was the $4.7 \times 4.7 \times 6.0 \text{ mm}^3$ stoichiometric $\text{Tm}^{3+}:\text{LiNbO}_3$ crystal discussed in Section 4.3.2. The UOT experiments also made use of cryostats dependent on cryogenic liquids. Such systems pose both logistical and safety challenges, as the cryogenic liquids and subsequent gases need to be handled. These challenges are exacerbated when moving to a clinical setting for clinical testing of the UOT.

The experiments discussed in this section were designed to explore these issues in preparation for our UOT experiment to move into a clinical setting. The experiments tested a fully moveable setup (i.e. built on wheels), using a “cryogenic-free” cryostat (i.e. a closed-cycle cryostat), and operating with a wavelength in the tissue-optical window. This system was then tested again on phantoms, though this time

the phantoms were made with the aim of representing breast tissue, with a small inclusion indicating malignant cancerous tissue.

Tumorous breast-tissue phantoms

The phantoms were prepared similarly to those in Section 6.2. In a cuboid mould with a base area of $9 \times 9 \text{ cm}^2$, a layer of healthy-breast-tissue-mimicking phantom material was cast. Once this was solidified, cuboid pieces of phantom tissue mimicking cancerous tissue were placed on top. This “cancerous” phantom was made in the same way as the “healthy”, but with other concentrations of Indian ink and Intralipid. Over this, another layer of healthy phantom tissue was cast. This second layer was as thick as the first, plus the height of the tumorous inclusions. The “tumours” were thus placed in the centre of the phantom.

The absorption and scattering parameters of the healthy and tumorous tissue can be seen in Table 6.2. To remain as close to a future measurement on real tissue as possible, these values were based on *in vivo* measurements. For bulk healthy-breast tissue, *in vivo* parameters are easily sourced in the literature [133, 134]. This is not the case for *in vivo* cancerous tissue, where the literature is more sparse. The measurements closest to 794 nm, the wavelength required by the $\text{Tm}^{3+}:\text{LiNbO}_3$ filter, of *in vivo* cancerous tissue were 811 nm [135] and 811 nm [135], which report very different μ'_s for both the healthy and cancerous tissues. The reported μ'_s of the healthy tissue in [135, 136] differs from the average reported in the literature [133, 134], while the μ_a reported by these sources is comparable to the literature average. Interestingly, the ratios of healthy and cancerous μ'_s in [135, 136] are the same.

Table 6.2: Properties used to create the healthy- and cancerous-tissue phantom types, based on [134, 136].

| | $\mu'_s [\text{cm}^{-1}]$ | $\mu_a [\text{cm}^{-1}]$ |
|------------------|---------------------------|--------------------------|
| Healthy tissue | 8.1 | 0.041 |
| Cancerous tissue | 12 | 0.12 |

The healthy tissue in Table 6.2 is based on the 786 nm data reported in [134], which was taken from a large number of women. The cancerous μ'_s is the healthy scattering parameter, scaled with the ratio of cancerous and healthy μ'_s reported in [135, 136]. The cancerous μ_a was taken directly from [136]. Incidentally, the ratio of the healthy and cancerous tissue μ'_s is approximately the same in *ex vivo* measurements as in the *in vivo* measurements (of which there are more sources in the literature [6, 137, 138]), though the values of μ'_s and μ_a for healthy tissue




Figure 6.15: Two pictures of the same phantom. (a) Phantom under construction the moment before the “tumour” inclusions were covered by another layer of “healthy” phantom tissue. (b) The same phantom after it had been used for imaging, sliced along the red dashed line in (a).

can differ widely from the *in vivo* counterparts. This difference can be explained by the physiological processes that change the scattering and absorption that start immediately after the tissue is removed (e.g. dehydration and deoxygenation) or by the changes incurred due to some samples being frozen prior to the optical measurement.

To ensure consistent accuracy in the optical properties of the phantoms, it was necessary to perform the time-of-flight spectroscopy during the mixing of the phantom, before they were cast in the mould. A phantom with three “tumour” inclusions is depicted in Figure 6.15.

6.3.1 Experimental setup

As in Section 6.1, the phantoms were imaged in a transmission-mode geometry. Many of the same components as in the previous experiments were employed, but with some updates. The ultrasound was a Verasonics system,¹³ using the linear array L7-4 transducer with a 6 MHz centre-frequency ultrasound pulses. The light source was a continuous wave titanium-sapphire laser¹⁴ frequency locked using the Pound-Drever-Hall technique [123] to a vibration isolated and temperature-stabilised reference cavity,¹⁵ which has a drift of ~ 1 Hz/s. This stabilised source was aligned through a double-pass AOM setup, with optical pulses generated using an AWG.¹⁶ The output from the double-pass AOM was coupled into a polarisation-maintaining fibre and routed to the experimental table depicted in Figure 6.16. The available output power of this laser is considerably greater than

¹³ Verasonics Vantage 32LE – Research ultrasound platform.

¹⁴ M-squared SolsTiS pumped with a Lighthouse Photonics Sprout system at 15 W.

¹⁵ Menlo Systems GmbH, ORS-Cavity.

¹⁶ Agilent N8241A.




Figure 6.16: Setup used for imaging breast tissue phantoms.

that of the laser used in the previous experiments, and as such, approximately 200 mW could be delivered to the phantom (a factor 10 improvement). An engineered diffuser¹⁷ was used to evenly distribute this probe light over a $\sim 1 \text{ cm}^2$ area at the phantom.

The mechanised flip mirror used previously to swap between probing and optical pumping paths was exchanged for a dual AOM setup. When these AOMs are turned on, the positive first-order diffracted beam of the first AOM is aligned into the second AOM. Here, the first negative order is selected; thus, if both AOMs operate with the same frequency, the optical frequency before and after the two AOMs will be unchanged. This switch can be activated within 1 μs and is thus considerably faster than the 0.4 s required by the mechanised flip mirror. One drawback concerns efficiency. For the optical path to the phantom, the losses are minimal; but for the filter-preparation arm through the AOMs, transmission is roughly 30%. This is low and could be improved; but as only $\sim 30 \text{ mW}$ was required for the filter preparation, this level of efficiency was deemed sufficient.

Another new component is the cryostat shown in Figure 6.17. This is a custom-designed closed-cycle cryostat made by MyCryoFirm, with large optical access and 2" windows and a 3" diameter bore along the probe beam path. This allowed for 2" and high-numerical aperture lenses¹⁸ to be mounted inside the cryostat, yielding $T_{\text{optics}} = 9.4\%$, which is larger than in the previous experiments, even though the crystal area is a quarter of that of the $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ crystal used in Sections 6.1 and 6.2.

To prevent unfiltered light reaching the detector, all entrances to the cryostat (other

¹⁷ Viavi Solutions, Inc., EDS-2-08168-A.

¹⁸ Thorlabs ACL5040U-B.




Figure 6.17: (a) Cryostat with vacuum enclosure removed, making the magnet coil and central bore visible. (b) Cryostat placed in the setup in Figure 6.16. The crystal is viewed from the side and the probe arm is incoming from the right-hand side. (c) Crystal mount inside the cryostat, viewed from the same direction as the probing light. The solid-copper shield with a small square window in the centre of (c) blocks any light going around the crystal. The lower right of (c) depicts a three-dimensional model of the crystal (purple) in its mount (orange) and shield (grey transparent).

than the light guide) were made light-tight. For the filter-creation path, this required a mechanised shutter that was closed during the probing part of the sequence.¹⁹ After the shutter to the detector had been closed, the light was filtered by the bandpass filter²⁰ discussed in Section 4.3.2. This had the added benefit of also filtering out any other light. As such, the setup could be run with dimmed room lights, without any increase in detector dark counts, something that was not previously possible (actual measurements were, as a precaution, carried out with room lights switched off).

To stop any light from the light guide circumventing the crystal, a sample mount was designed with a shield to stop any light from passing around the crystal. The sample mount can be seen in Figure 6.17c. To monitor the temperature of the crystal, two thermal sensors were used. The first was placed at the bottom of the sample mount and the second was placed on a part of the mount only connected to the cryostat cold finger via the crystal. From this, the crystal could be seen to maintain a temperature of 2.8–3.1 K while imaging and 2.8–3.4 K during the filter-creation sequence.

The imaging sequence consisted of a frame of 21 parallel ultrasound lines emitted with 1 ms separation. The timing of the light and ultrasound pulses was implemented programmatically, using the AWG to trigger the ultrasound pulse, thus making the delay box obsolete. This separation allowed the fluorescence still leak-

¹⁹ Thorlabs SHB1T.

²⁰ Quantum Design GmbH, 795FS10-50.

ing through the bandpass filter to subside. Otherwise, a separation of 200 μs was possible without any risk of overheating the ultrasound transducer. The pulses were measured to a peak negative pressure of above -3 MPa, which – for the 6 MHz centre frequency – yielded a mechanical index below the medical safety limit (see Sect. sec:US).

Before each filter-preparation sequence, a total of four frames were captured. More frames could have been used, as the T_F only fell by 14 percentage points during this time. To limit the time overhead associated with the pre-processing of the oscilloscope and transfer-speed bottleneck, the signal was continuously averaged 100 times on the oscilloscope before being transferred. The total acquisition time of these 4×100 frames was 25 s, which is roughly the time that a patient can be expected to hold still and hold their breath. With the ~ 100 ms acquisition time, the UOT duty cycle is 25% with a (lopsided) 16 frames/s.

To provide rapid experimental feedback, in addition to images of the embedded inclusions, a reference image of a part of the phantom with no inclusion was also acquired. This reference image was then used in a similar manner to \mathbf{F}_{ref} in Section 6.2.2 to remove the fluence-distribution artefact.

In addition to the UOT images, ultrasound images of the same area with the inclusions – which were visible on the ultrasound due to air bubbles trapped between the first phantom-tissue layer and the placed inclusions – were also captured.

6.3.2 Images of realistic breast-tumours

Figure 6.18 shows images of 25 s and 250 s image-acquisition times, taken in the 4.2 cm-thick phantom in Figure 6.15, and divided by a reference image. Accompanying each UOT image is a corresponding ultrasound image.

Decreasing the noise and increasing the probing power would further increase the contrast-to-noise ratio, making the images clearer in a shorter image-acquisition time. Further increasing the UOT duty cycle by acquiring more frames before re-preparing the filter is also possible, though this would require some type of filter-decay compensation – something not performed for the results presented in Figure 6.18. A power increase in the current setup is highly possible, as the laser-to-tissue transmission was 10% and could be increased to 50% by optimised optical components.

Attempts were also made to image similar inclusions in a 5 cm-thick phantom; though here, the noise was too strong to acquire sufficient contrast. Since the noise was dominated by the fluorescence, future imaging using the $\text{Tm}^{3+}:\text{LiNbO}_3$

filter at 794 nm should address this issue. One interesting possibility would be filtering using a rubidium gas cell, as Rb has a sharp transition at 795 nm that could be tuned by varying the pressure of the gas [139].

Even with the correction using the reference image, the fluence distribution imposes itself onto the image. In contrast to the reflection geometry discussed in Section 6.2, the fluence distribution in the transmission mode is more shapable in both image axes by collecting and injecting light over larger areas. There was an attempt in the presented experiment to increase the size of the injection area, but due to the restricted collection area, an increase of more than 1 cm² had little effect on the fluence-distribution artefact.

Though the results here are limited in terms of depth of analysis, these images nonetheless demonstrate that UOT can image medically relevant contrasts in medically relevant tissue volumes. This result is therefore encouraging for the future development of UOT.




Figure 6.18: UOT images of one of the $< 5 \times 5 \times 5 \text{ mm}^3$ large inclusions in the phantom in Figure 6.15. In the first column the ultrasound image is acquired over 25 s, (i.e. averaged from 400 captured frames). The second column is acquired during over 250 s (i.e. averaged from 4,000 frames). In both columns, the image is the result of a capture- image divided by a reference image. The last column is the ultrasound B-mode image of the same area as the UOT images in each row. The ultrasound image is performed with the largest possible gain and emitted pulse intensity, and the inclusions are believed to be visible in the ultrasound images only due to trapped air between the two layers of phantom tissue.

Chapter 7

Outlook and conclusion

“The most important step a man can take. It’s not the first one, is it?
It’s the next one. Always the next step”.

— Brandon Sanderson, *Oathbringer*

Ultrasound optical tomography has been a developing technique now for almost 30 years, and the use of spectral-hole filters may represent the final push that the technique needs to become clinically viable. The results in Section 6.3 thus enable a hopeful outlook to the research project, as the equipment and methods presented in this thesis have proven to be transportable and to work in scenarios directly translatable to clinical application.

That said, the results in Section 6.3 also show that further improvements are needed before UOT using SHF is sufficiently optimised to warrant clinical trials. The results in Section 6.3 also show the impact of the ever-present fluence-distribution artefact, a problem that will require increasing attention as the project moves forward into more toward heavier implementation of image analysis tools.

Experimental outlook

At the time of writing, experimental development is moving forward with the quantum information group at Lund University focusing on applying UOT to mammography imaging. In the ongoing work to increase the “many brooks” of the signal level, a new detector with above 20% quantum efficiency is being acquired and efforts to optimise the transfer efficiency from laser head to tissue are being

made. Removing the fluorescence in $\text{Tm}^{3+}:\text{LiNbO}_3$ is another goal that is close at hand and which, when achieved, will allow for UOT frames to be taken an order of magnitude more quickly, without risking oversaturation of the detector. Overall, these fairly straightforward improvements will allow for major improvements in the signal and contrast-to-noise ratio achieved in the same amount of imaging time as in Section 6.3.

Looking further into the future, the implementation of multiple wavelengths presents an interesting avenue of development. Instead of relying on a reference image produced by a simulation or image of a known “blank” region, as in sections 6.2-6.3, the multispectral properties of light could be utilised instead. Imaging tissue oxygenation would be made possible by dividing an image taken at 690 nm, for example, by an image at the isosbestic point of ~ 790 nm. This operation would thus remove the effects of blood volume and, to an extent, scattering. This will require both high suppression filters and stabilised sources at both of these wavelengths. The former is close at hand, with the $\text{Tm}^{3+}:\text{LaF}_3$ crystals investigated in the thesis of Alexander Bengtsson [37], where high suppression filters at 689.5 nm are demonstrated. The latter requirement for multispectral imaging is further away and presents greater technical difficulty. This difficulty may be partially solved by the ever-developing laser industry: as high-powered and cheap diode lasers in near infrared spectrum become more prevalent, it may be possible to move away from the complex Ti:Sapphire system used in Section 6.3.

Future of the sequential Monte Carlo model

With multiple wavelengths come more opportunities and challenges in image analysis. It remains to be seen whether the SMC model will be the correct tool with which to meet these future challenges. It is a safe bet, however, that some type of model will be required, and the SMC model has so far proved promising in accuracy, flexibility, and speed – though of course, this speed is only acquired if considerable time is first spent computing the matrices presented in Section 3.3.5.

The SMC model still requires some development, as it is not currently well-optimised for simulating continuous illumination of the tissue as the ultrasound pulse propagates. Here, fundamental changes will be required. One such change is that, instead of producing a power spectrum in the pulse interaction, a spectrogram (i.e. a time-varying spectrum) is to be calculated. Another change will concern the handling of the interaction distance used to speed up simulations, which – as discussed in Section 3.3.6 – is implemented as the distance to a fixed point or line and not the actual current position of the ultrasound pulse (i.e. the interaction distance must be defined larger than the actual size of the ultrasound due to it propagating).

Also regarding implementation, there is no doubt that many improvements can be made by simply moving from MATLAB to some other, more efficient programming language, such as C. Another avenue of promising development would be to integrate the model with an acoustic simulator, such as the k-Wave toolbox [140].

However, even in its current state, the SMC model is already being used by others in the quantum information group at Lund University to analyse UOT imaging setups. A master's student, Henrik von Friesendorff, is pursuing one such study [141], using the model for a study on computer tissue phantoms from the radiological department at Lund University Hospital [142]. Another such study, headed by Kevin Shortiss, post-doc, is investigating imaging through the skull and inside the brain. A personal hope is that the TLS source code [39] will spread outside of its crucible and provide a benefit to the community at large, be it for analysing experimental setups or for simulating clinical use cases.

Another personal hope is that you, dear Reader, having reached this point, feel that you have become ever-so-slightly wiser. For it has been a true privilege to have been able to share with you the research in this thesis, and to have had your attention until now.

The End

Author contributions

Paper I: Characterization and modeling of acousto-optic signal strengths in highly scattering media

Alexander Bengtsson did most of the experimental planning with the assistance of me, Lars Rippe and Stefan Kröll. The experimental UOT data was gathered by me and Alexander where Lars and Stefan K. provided guidance during the experiments. I also planned and performed the ultrasound characterisation with the help of Magnus Cynthio and Tobias Erlöv. I constructed the light diffusion model describing the carrier and tagged light intensities. The tissue phantom characterisations which were performed by Meng Li and Mengqiao Di with the assistance of Nina Reistad. Stefan Andersson-Engels helped with the writing of the manuscript and discussions on light in scattering media. Alexander wrote the majority of the manuscript and I wrote part of it with the input of the other authors.

Paper II: Acousto-optic interaction strengths in optically scattering media using high pressure acoustic pulses

I planned most of the experiment (which was a continuation of the experiment in **Paper I**) with the assistance of Alexander Bengtsson and Stefan Kröll. Alexander and I performed the measurements together. I planned and performed the ultrasound pulse characterisation with the aid of Magnus Cynthio and Tobias Erlöv, who also provided input on the acoustic modelling. I further performed the theoretical model development, performed the data analysis and wrote the manuscript with the input of the other authors.

Paper III: First principle simulation package for arbitrary acousto-optic interaction in scattering materials

The paper was a continuation of **Paper II** with the same author contributions. I wrote the manuscript and performed the additional theoretical modelling and *in silico* study.

Paper IV: Comparison of contrast-to-noise ratios of different detection methods in ultrasound optical tomography

Alexander Bengtsson did the literature based analysis of the contrast-to-noise ratios of the different techniques, to which I, Lars Rippe, Stefan Kröll and Kevin Shortiss contributed in discussions. I wrote and performed the simulation used to quantitatively compare the different techniques based on the theory work by Alexander. Alexander wrote the manuscript to which I and the other authors provided feedback. Kevin, Stefan and I edited the manuscript during the review process.

Paper V: Sequential Monte Carlo modelling of ultrasound optical tomography: image prediction and absorption reconstruction in optically diffuse media

I planned and performed the UOT experiments together with Alexander Bengtsson and Kevin Shortiss with guidance from Lars Rippe and Stefan Kröll. I planned the ultrasound pulse characterisation which me and Kevin Shortiss performed together. Magnus Cinthio and Jeff Powers programmed the ultrasound source that was used. I developed the matrix formulation of the sequential Monte Carlo model and formulated how it could be applied to the absorption reconstruction. I performed the data analysis and am writing the manuscript with input from the other authors.

Acknowledgements

Adam Smith notes in his *The Wealth of Nations*, that a woolen-coat is “the produce of the joint labor of a great multitude of workmen” and I cannot help but feel that the same can be said of a Ph.D. While I do not wish to diminish my own contribution to the work presented in this thesis, none of it would have been possible without the immeasurable toil of people whom I will never know, nor will ever know me. To those people, thank you.

This thesis would also not have been possible without the many people who has co-travelled with me on this arduous journey that is concluding with the words in front of you. One such person is Stefan Kröll – whose door has always stood open and who has met my many questions with kindness, patience, and – above all – great answers. Stefan, I am so very grateful that I have had you as my supervisor and to have been part of the research group that you have created. To my co-supervisor Lars Rippe: your inventiveness and explosive enthusiasm has created many great discussions – thank you for changing my mind on so many topics. To my other co-supervisor Johannes Swartling: your collectedness and calm during even the most trying periods has been a true inspiration and I am grateful for all the trust you have placed in me on behalf of Spectracure for pushing this project forward.

I am also grateful for my other colleagues, chief among these being Alexander Bengtsson, with whom I have spent so many days with in the lab and who through discussion has contributed to so much of the work presented in this thesis. Alexander, your methodological and practical approach to any problem never ceased to inspire and, because of you, I think I will always have some ice-hockey analogies close at hand. To Kevin Shortiss, thank you for your enthusiasm and levity in the lab (even when things looked dire) and for volunteering yourself to be the first user of the tagged light simulator. To Chunyan Shi, thank you for teaching me the ways of the lab and helping me plan and run my first ever experiments. To the present and former members of the Quantum Information group:

Andreas Walther, Adam Kinos, Mohammed Alqedra, Hafsa Syed, Abdullah Abdelatief, Safi Rafie-Zinedine, David Gustavsson, Marcus Lindén, Sebastian Horvath, Akvile Zabaliute-Karaliune, Ibtisam Dahir Ahmed, Henrik von Friesendorff, and Johannes Renders - thank you all for creating a great atmosphere to do research and for coffee breaks where no subject was taboo.

My thanks also go out to everyone else at atomic physics for the many nice lunch breaks and the always present offer to help with spare equipment or expertise. Thanks especially to the administrative staff: Åke Johansson, Jane Nilsson and Maria Algotsson – you make the research possible.

I would also like to thank my many collaborators. To Magnus Cynthio, Tobias Erlöv and Jeff Powers, for your many contributions on ultrasound sources. To Emilie Krite-Svanberg, for your invaluable medical expertise and for teaching me what my work can actually be used for. To Nina Reistad, for helping us with the creation and validation of phantom tissues. To Stefan Andersson-Engels and your team in Cork, for the many great discussions on the future of UOT and for sharing your tissue-optics expertise.

I would also like to acknowledge the people at Spectracure: Marcelo, Elin, Mirjam, Dennis, Rebecca, Simon, Kristina, Matilda, Stefan, Marija, Sofia, Ulrika, Christofer, and Johan. Thank you all for always making me feel welcome and as a part of the team. To Björn Thomasson especially, thank you for sharing your office with me once every week and for all the help you provided when I started at Spectracure.

And last but not least, to my parents Helen and Douglas, my sister Joy, my extended family, my friends and my partner Johanna. Thank you all for supporting and believing in me, and especially for making me laugh when I've needed it.

References

- [1] Nevit Dilmen, “Ultrasound images of fetal heart.”
https://commons.wikimedia.org/wiki/File:Pregnancy_ultrasound_110303135229_1358380.jpg (Creative commons BY-SA 3.0).
- [2] Ptrump, “One frame of an MRI scan of the head.”
https://commons.wikimedia.org/wiki/File:MRI_Scan_General_Illustration.jpg (Creative commons BY-SA 4.0).
- [3] Yui Kubo, “Arterial blood and venous blood model.”
<https://www.flickr.com/photos/yu-kubo/6305735712/> (Creative commons BY-NC-SA 2.0).
- [4] Mikael Häggström, “Two examples of tumors at gross pathology.”
<https://commons.wikimedia.org/wiki/File:Tumors.jpg>.
- [5] “American national standard for the safe use of lasers,” ANSI Z136.1, American National Standards Institute, Inc. (1993).
- [6] S. L. Jacques, “Optical properties of biological tissues: a review,” *Physics in Medicine & Biology*, vol. 58, no. 11, p. R37 (2013).
- [7] H. Mitchell, T. Hamilton, F. Steggerda, and H. Bean, “The chemical composition of the adult human body and its bearing on the biochemistry of growth,” *Journal of Biological Chemistry*, vol. 158, no. 3, pp. 625–637 (1945).
- [8] S. Prahl, “Tabulated molar extinction coefficient for hemoglobin in water.”
<https://omlc.org/spectra/hemoglobin/summary.html>.
- [9] J. M. Murkin and M. Arango, “Near-infrared spectroscopy as an index of brain and tissue oxygenation,” *British Journal of Anaesthesia*, vol. 103, no. suppl_1, pp. i3–i13 (2009).

- [10] C. M. Goh, R. Subramaniam, N. M. Saad, S. A. Ali, and F. Meriaudeau, "Subcutaneous veins depth measurement using diffuse reflectance images," *Optics Express*, vol. 25, no. 21, pp. 25741–25759 (2017).
- [11] W. M. Tolles, J. W. Nibler, J. McDonald, and A. B. Harvey, "A review of the theory and application of coherent anti-Stokes Raman spectroscopy (CARS)," *Applied Spectroscopy*, vol. 31, no. 4, pp. 253–271 (1977).
- [12] S. Kröll, M. Aldén, P.-E. Bengtsson, and C. Löfström, "An evaluation of precision and systematic errors in vibrational CARS thermometry," *Applied Physics B*, vol. 49, no. 5, pp. 445–453 (1989).
- [13] D. Pappas, B. W. Smith, and J. D. Winefordner, "Raman spectroscopy in bioanalysis," *Talanta*, vol. 51, no. 1, pp. 131–144 (2000).
- [14] L. V. Wang and H. I. Wu, *Biomedical Optics: Principles and Imaging*. John Wiley & Sons (2012).
- [15] J. Lundgren, *Design of Functional Structures and Measurement Techniques for Electromagnetic Waves*. PhD thesis, Lund University (2021).
- [16] C. J. Foot, *Atomic Physics*. Oxford University Press (2005).
- [17] R. C. Hilborn, "Einstein coefficients, cross sections, f values, dipole moments, and all that," *American Journal of Physics*, vol. 50, no. 11, pp. 982–986 (1982).
- [18] J. H. Jeans, "The equations of radiative transfer of energy," *Monthly Notices of the Royal Astronomical Society*, vol. 78, no. 1, pp. 28–36 (1917).
- [19] B. J. Hunt, *The Maxwellians*. Cornell University Press (2005).
- [20] P. Bouguer, *Essai d'optique, sur la gradation de la lumiere*. Claude Jombert (1729).
- [21] T. Svensson, E. Alerstam, D. Khoptyar, J. Johansson, S. Folestad, and S. Andersson-Engels, "Near-infrared photon time-of-flight spectroscopy of turbid materials up to 1400 nm," *Review of Scientific Instruments*, vol. 80, no. 6, p. 063105 (2009).
- [22] S. A. Prahl, "Everything I think you should know about inverse adding-doubling," <https://omlc.org/software/iad/manual.pdf> (2011).
- [23] A. N. Witt, "Multiple scattering in reflection nebulae. I - a Monte Carlo approach," *The Astrophysical Journal Supplement Series*, vol. 35, pp. 1–6 (1977).

- [24] A. Fick, “Über diffusion,” *Annalen der Physik*, vol. 170, no. 1, pp. 59–86 (1855).
- [25] J. Philibert, “One and a half century of diffusion: Fick, Einstein,” in *Diffusion Fundamentals: Leipzig 2005*, vol. 1, ch. 1, pp. 8–18, Leipziger Universitätsverlag (2005).
- [26] Y. Yamada and S. Okawa, “Diffuse optical tomography: present status and its future,” *Optical Review*, vol. 21, no. 3, pp. 185–205 (2014).
- [27] P. Välisuo, “Optical methods for assessing skin flap survival,” in *Biophotonics for Medical Applications*, ch. 12, pp. 331–346, Woodhead Publishing (2015).
- [28] J. Gunther and S. Andersson-Engels, “Review of current methods of acousto-optical tomography for biomedical applications,” *Frontiers of Optoelectronics*, vol. 10, no. 3, pp. 211–238 (2017).
- [29] B.-C. Forget, F. Ramaz, M. Atlan, J. Selb, and A.-C. Boccara, “High-contrast fast Fourier transform acousto-optical tomography of phantom tissues with a frequency-chirp modulation of the ultrasound,” *Applied Optics*, vol. 42, no. 7, pp. 1379–1383 (2003).
- [30] M. Kempe, M. Larionov, D. Zaslavsky, and A. Genack, “Acousto-optic tomography with multiply scattered light,” *Journal of the Optical Society of America A*, vol. 14, no. 5, pp. 1151–1158 (1997).
- [31] F. Ramaz, B. Forget, M. Atlan, A.-C. Boccara, M. Gross, P. Delaye, and G. Roosen, “Photorefractive detection of tagged photons in ultrasound modulated optical tomography of thick biological tissues,” *Optics Express*, vol. 12, no. 22, pp. 5469–5474 (2004).
- [32] G. Yao and L. V. Wang, “Theoretical and experimental studies of ultrasound-modulated optical tomography in biological tissue,” *Applied Optics*, vol. 39, no. 4, pp. 659–664 (2000).
- [33] P. Lai, X. Xu, and L. V. Wang, “Ultrasound-modulated optical tomography at new depth,” *Journal of Biomedical Optics*, vol. 17, no. 6, pp. 1 – 6 (2012).
- [34] S.-R. Kothapalli, S. Sakadzic, C. Kim, and L. V. Wang, “Imaging optically scattering objects with ultrasound-modulated optical tomography,” *Optics Letters*, vol. 32, no. 16, pp. 2351–2353 (2007).
- [35] T. S. Leung and S. Powell, “Fast Monte Carlo simulations of ultrasound-modulated light using a graphics processing unit,” *Journal of Biomedical Optics*, vol. 15, no. 5, pp. 1 – 7 (2010).

- [36] A. Walther, L. Rippe, L. V. Wang, S. Andersson-Engels, and S. Kröll, “Analysis of the potential for non-invasive imaging of oxygenation at heart depth, using ultrasound optical tomography (UOT) or photo-acoustic tomography (PAT),” *Biomedical Optics Express*, vol. 8, no. 10, pp. 4523–4536 (2017).
- [37] A. Bengtsson, *Material and technique development for ultrasound optical tomography using spectral hole burning filters*. PhD thesis, Lund University (2022).
- [38] Y. Huang, M. Cua, J. Brake, Y. Liu, and C. Yang, “Investigating ultrasound–light interaction in scattering media,” *Journal of Biomedical Optics*, vol. 25, no. 2, p. 025002 (2020).
- [39] D. Hill, “Tagged light simulator.”
<https://bitbucket.org/qilund/tagged-light-simulator>.
- [40] A. Pierce, *Acoustics: An Introduction to Its Physical Principles and Applications*. Springer International Publishing (2019).
- [41] “Marketing clearance of diagnostic ultrasound systems and transducers,” fda-2017-d-5372, U.S. Food and Drug Administration (2019).
- [42] W. Leutz and G. Maret, “Ultrasonic modulation of multiply scattered light,” *Physica B: Condensed Matter*, vol. 204, no. 1–4, pp. 14–19 (1995).
- [43] G. Mahan, W. Engler, J. Tiemann, and E. Uzgiris, “Ultrasonic tagging of light: theory,” *Proceedings of the National Academy of Sciences*, vol. 95, no. 24, pp. 14015–14019 (1998).
- [44] L. Wang, “Mechanisms of ultrasonic modulation of multiply scattered coherent light: an analytic model,” *Physical Review Letters*, vol. 87, p. 043903 (2001).
- [45] J. G. Hoskins and J. C. Schotland, “Acousto-optic effect in random media,” *Physical Review E*, vol. 95, no. 3, p. 033002 (2017).
- [46] S. Sakadžić and L. V. Wang, “Ultrasonic modulation of multiply scattered coherent light: an analytical model for anisotropically scattering media,” *Physical Review E*, vol. 66, no. 2, p. 026603 (2002).
- [47] S. Sakadžić and L. V. Wang, “Modulation of multiply scattered coherent light by ultrasonic pulses: an analytical model,” *Physical Review E*, vol. 72, no. 3, p. 036620 (2005).

- [48] J. L. Hollmann, R. Horstmeyer, C. Yang, and C. A. DiMarzio, “Diffusion model for ultrasound-modulated light,” *Journal of Biomedical Optics*, vol. 19, no. 3, p. 035005 (2014).
- [49] M. Jang, H. Ruan, B. Judkewitz, and C. Yang, “Model for estimating the penetration depth limit of the time-reversed ultrasonically encoded optical focusing technique,” *Optics Express*, vol. 22, no. 5, pp. 5787–5807 (2014).
- [50] C. Dupuy, *IMAGE RECONSTRUCTION FOR ACOUSTO-OPTICS: TOWARDS QUANTITATIVE IMAGING*. PhD thesis, Université Paris Sciences et Lettres (2017).
- [51] S. Powell and T. S. Leung, “Highly parallel Monte-Carlo simulations of the acousto-optic effect in heterogeneous turbid media,” *Journal of Biomedical Optics*, vol. 17, no. 4, pp. 1 – 12 (2012).
- [52] L. Wang, “Mechanisms of ultrasonic modulation of multiply scattered coherent light: a Monte Carlo model,” *Optics Letters*, vol. 26, pp. 1191–3 (2001).
- [53] N. N. Nath and C. V. Raman, “The diffraction of light by high frequency sound waves: Part 1,” in *Proceedings of the Indian Academy of Sciences – Section A*, vol. 2 (1936).
- [54] F. J. Chung, J. G. Hoskins, and J. C. Schotland, “Coherent acousto-optic tomography with diffuse light,” *Optics Letters*, vol. 45, no. 7, pp. 1623–1626 (2020).
- [55] A. A. Leino, A. Pulkkinen, and T. Tarvainen, “ValoMC: a Monte Carlo software and MATLAB toolbox for simulating light transport in biological tissue,” *OSA Continuum*, vol. 2, no. 3, pp. 957–972 (2019).
- [56] J. Jönsson and E. Berrocal, “Multi-Scattering software: part I: online accelerated Monte Carlo simulation of light transport through scattering media,” *Optics Express*, vol. 28, no. 25, pp. 37612–37638 (2020).
- [57] E. Alerstam, T. Svensson, and S. Andersson-Engels, “CUD-AMCML: user manual and implementation notes.” <https://www.atomic.physics.lu.se/fileadmin/atomfysik/Biophotonics/Software/CUDAMCML.pdf> (2009).
- [58] S. A. Prahl, “A Monte Carlo model of light propagation in tissue,” in *Dosimetry of Laser Radiation in Medicine and Biology*, vol. 10305, pp. 105–114, SPIE (1989).

- [59] L. Wang, S. L. Jacques, and L. Zheng, “MCML—Monte Carlo modeling of light transport in multi-layered tissues,” *Computer Methods and Programs in Biomedicine*, vol. 47, no. 2, pp. 131–146 (1995).
- [60] J. W. Goodman, *Statistical Optics*, ch. 2.9. Wiley (2015).
- [61] H. Ammari, E. Bossy, J. Garnier, L. Nguyen, and L. Seppecher, “A reconstruction algorithm for ultrasound-modulated diffuse optical tomography,” *Proceedings of the American Mathematical Society*, vol. 142, no. 9, pp. 3221–3236 (2014).
- [62] G. Bal and J. C. Schotland, “Inverse scattering and acousto-optic imaging,” *Physical Review Letters*, vol. 104, no. 4, p. 043902 (2010).
- [63] W. Riley and W. Klein, “Piezo-optic coefficients of liquids,” *The Journal of the Acoustical Society of America*, vol. 42, no. 6, pp. 1258–1261 (1967).
- [64] J.-B. Laudereau, A. A. Grabar, M. Tanter, J.-L. Gennisson, and F. Ramaz, “Ultrafast acousto-optic imaging with ultrasonic plane waves,” *Optics Express*, vol. 24, no. 4, pp. 3774–3789 (2016).
- [65] S. A. Fulop and K. Fitz, “Algorithms for computing the time-corrected instantaneous frequency (reassigned) spectrogram, with applications,” *The Journal of the Acoustical Society of America*, vol. 119, no. 1, pp. 360–371 (2006).
- [66] W.-F. Cheong, S. A. Prahl, and A. J. Welch, “A review of the optical properties of biological tissues,” *IEEE Journal of Quantum Electronics*, vol. 26, no. 12, pp. 2166–2185 (1990).
- [67] G. B. Haxel, J. B. Hedrick, , and G. J. Orris, “Rare earth elements—critical resources for high technology,” USGS Fact Sheet 087-02, U.S. Department of the Interior, USGS (2002).
- [68] P. Friedrich, “Methods for bulk growth of inorganic crystals: crystal growth,” in *Encyclopedia of Condensed Matter Physics*, pp. 262–274, Elsevier (2005).
- [69] Wikipedia, “Periodic table - ionisation energy,” (2022). https://en.wikipedia.org/wiki/Periodic_table#Ionisation_energy.
- [70] T. Böttger, C. W. Thiel, R. L. Cone, and Y. Sun, “Effects of magnetic field orientation on optical decoherence in $\text{Er}^{3+}:\text{Y}_2\text{SiO}_5$,” *Physical Review B*, vol. 79, p. 115104 (2009).

- [71] R. W. Equall, Y. Sun, R. L. Cone, and R. M. Macfarlane, "Ultraslow optical dephasing in $\text{Eu}^{3+}:\text{Y}_2\text{SiO}_5$," *Physical Review Letters*, vol. 72, pp. 2179–2182 (1994).
- [72] F. Könz, Y. Sun, C. W. Thiel, R. L. Cone, R. W. Equall, R. L. Hutcheson, and R. M. Macfarlane, "Temperature and concentration dependence of optical dephasing, spectral-hole lifetime, and anisotropic absorption in $\text{Eu}^{3+}:\text{Y}_2\text{SiO}_5$," *Physical Review B*, vol. 68, p. 085109 (2003).
- [73] S. A. C. McDowell, "A simple derivation of the Boltzmann distribution," *Journal of Chemical Education*, vol. 76, no. 10, p. 1393 (1999).
- [74] K. B. Hill and A. E. Craig, "Rare-earth-based spectral memories: material implications," in *Advanced Optical Data Storage*, vol. 4988, pp. 7 – 14, SPIE (2003).
- [75] N. Vitanov, B. Shore, L. Yatsenko, K. Böhmer, T. Halfmann, T. Rickes, and K. Bergmann, "Power broadening revisited: theory and experiment," *Optics Communications*, vol. 199, no. 1-4, pp. 117–126 (2001).
- [76] E. S. Maniloff, F. R. Graf, H. Gygax, S. B. Altnern, S. Bernet, A. Renn, and U. P. Wild, "Power broadening of the spectral hole width in an optically thick sample," *Chemical Physics*, vol. 193, no. 1-2, pp. 173–180 (1995).
- [77] J. G. Bartholomew, R. L. Ahlefeldt, and M. J. Sellars, "Engineering closed optical transitions in rare-earth ion crystals," *Physical Review B*, vol. 93, no. 1, p. 014401 (2016).
- [78] C. Venet, B. Car, L. Veissier, F. Ramaz, and A. Louchet-Chauvet, "Deep and persistent spectral holes in thulium-doped yttrium orthosilicate for imaging applications," *Physical Review B*, vol. 99, p. 115102 (2019).
- [79] Y. Sun, C. Thiel, R. Cone, R. Equall, and R. Hutcheson, "Recent progress in developing new rare earth materials for hole burning and coherent transient applications," *Journal of Luminescence*, vol. 98, no. 1-4, pp. 281–287 (2002).
- [80] S. R. Hastings-Simon, M. Afzelius, J. Minář, M. U. Staudt, B. Lauritzen, H. De Riedmatten, N. Gisin, A. Amari, A. Walther, S. Kröll, *et al.*, "Spectral hole-burning spectroscopy in $\text{Nd}^{3+}:\text{YVO}_4$," *Physical Review B*, vol. 77, no. 12, p. 125111 (2008).
- [81] R. Yano, M. Mitsunaga, and N. Uesugi, "Nonlinear laser spectroscopy of $\text{Eu}^{3+}:\text{Y}_2\text{SiO}_5$ and its application to time-domain optical memory," *Journal of the Optical Society of America B*, vol. 9, no. 6, pp. 992–997 (1992).

- [82] M. Silver, R. Joseph, C.-N. Chen, V. Sank, and D. Hoult, “Selective population inversion in NMR,” *Nature*, vol. 310, no. 5979, pp. 681–683 (1984).
- [83] I. Roos and K. Mølmer, “Quantum computing with an inhomogeneously broadened ensemble of ions: suppression of errors from detuning variations by specially adapted pulses and coherent population trapping,” *Physical Review A*, vol. 69, no. 2, p. 022321 (2004).
- [84] D. Manzano, “A short introduction to the Lindblad master equation,” *AIP Advances*, vol. 10, no. 2, p. 025106 (2020).
- [85] B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics*, ch. 5. Wiley series in pure and applied optics, Wiley Interscience (2007).
- [86] Wikipedia, “Group velocity,” (2022).
https://en.wikipedia.org/wiki/Group_velocity.
- [87] S. P. Horvath, C. Shi, D. Gustavsson, A. Walther, A. Kinos, S. Kröll, and L. Rippe, “Slow light frequency reference cavities—proof of concept for reducing the frequency sensitivity due to length fluctuations,” *New Journal of Physics*, vol. 24, no. 3, p. 033034 (2022).
- [88] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P. O. Schmidt, “Optical atomic clocks,” *Reviews of Modern Physics*, vol. 87, no. 2, p. 637 (2015).
- [89] G. H. Dieke and H. Crosswhite, “The spectra of the doubly and triply ionized rare earths,” *Applied Optics*, vol. 2, no. 7, pp. 675–686 (1963).
- [90] W. Carnall, G. Goodman, K. Rajnak, and R. Rana, “A systematic analysis of the spectra of the lanthanides doped into single crystal LaF_3 ,” *The Journal of Chemical Physics*, vol. 90, no. 7, pp. 3443–3457 (1989).
- [91] R. Equall, R. Cone, and R. Macfarlane, “Homogeneous broadening and hyperfine structure of optical transitions in $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$,” *Physical Review B*, vol. 52, no. 6, p. 3963 (1995).
- [92] M. Nilsson, L. Rippe, S. Kröll, R. Klieber, and D. Suter, “Hole-burning techniques for isolation and study of individual hyperfine transitions in inhomogeneously broadened solids demonstrated in $\text{Eu}^{3+}:\text{Y}_2\text{SiO}_5$,” *Physical Review B*, vol. 70, no. 21, p. 214116 (2004).
- [93] R. Macfarlane and R. Shelby, “Coherent transient and holeburning spectroscopy of rare earth ions in solids,” in *Modern Problems in Condensed Matter Sciences*, vol. 21, pp. 51–184, Elsevier (1987).

- [94] Y. C. Sun, “Rare earth materials in optical storage and data processing applications,” in *Spectroscopic Properties of Rare Earths in Optical Materials*, ch. 7, pp. 379–429, Tsinghua University Press and Springer-Verlag Berlin Heidelberg (2005).
- [95] M. Sabooni, A. N. Nilsson, G. Kristensson, and L. Rippe, “Wave propagation in birefringent materials with off-axis absorption or gain,” *Physical Review A*, vol. 93, no. 1, p. 013842 (2016).
- [96] C. Thiel, Y. Sun, T. Böttger, W. Babbitt, and R. Cone, “Optical decoherence and persistent spectral hole burning in $\text{Tm}^{3+}:\text{LiNbO}_3$,” *Journal of Luminescence*, vol. 130, no. 9, pp. 1598–1602 (2010).
- [97] Y. Sun, C. W. Thiel, and R. L. Cone, “Optical decoherence and energy level structure of $0.1\%\text{Tm}^{3+}:\text{LiNbO}_3$,” *Physical Review B*, vol. 85, p. 165106 (2012).
- [98] L. Xing, W. Yang, and J. Lin, “Impact of crystal defect on upconverted white-light emission in lanthanide-doped LiNbO_3 single crystal,” *Optical Materials*, vol. 84, pp. 215–220 (2018).
- [99] A. Lorenzo, H. Jaffrezic, B. Roux, G. Boulon, and J. García-Solé, “Lattice location of rare-earth ions in LiNbO_3 ,” *Applied Physics Letters*, vol. 67, no. 25, pp. 3735–3737 (1995).
- [100] H. Syed, A. Kinos, C. Shi, L. Rippe, and S. Kröll, “Microscopic model of spin flip-flop processes in crystals doped by rare-earth ions,” *Physical Review B*, vol. 106, p. 184109 (2022).
- [101] C. W. Thiel, T. Böttger, and R. Cone, “Rare-earth-doped materials for applications in quantum information storage and signal processing,” *Journal of Luminescence*, vol. 131, no. 3, pp. 353–361 (2011).
- [102] S. Svanberg, *Atomic and Molecular Spectroscopy: Basic Aspects and Practical Applications*, ch. 6. Springer (2004).
- [103] F. A. Marks, H. W. Tomlinson, and G. W. Brooksby, “Comprehensive approach to breast cancer detection using light: photon localization by ultrasound modulation and tissue characterization by spectral discrimination,” in *Photon Migration and Imaging in Random Media and Tissues*, vol. 1888, pp. 500–510, SPIE (1993).
- [104] M. M. Qureshi, J. Brake, H.-J. Jeon, H. Ruan, Y. Liu, A. M. Safi, T. J. Eom, C. Yang, and E. Chung, “In vivo study of optical speckle decorrelation time

- across depths in the mouse brain,” *Biomedical Optics Express*, vol. 8, no. 11, pp. 4855–4864 (2017).
- [105] K. Zhu, B. Zhou, Y. Lu, P. Lai, S. Zhang, and Y. Tan, “Ultrasound-modulated laser feedback tomography in the reflective mode,” *Optics Letters*, vol. 44, no. 22, pp. 5414–5417 (2019).
 - [106] Y. Liu, R. Cao, J. Xu, H. Ruan, and C. Yang, “Imaging through highly scattering human skulls with ultrasound-modulated optical tomography,” *Optics Letters*, vol. 45, no. 11, pp. 2973–2976 (2020).
 - [107] M. Gross, F. Ramaz, B. Forget, M. Atlan, A. Boccara, P. Delaye, and G. Roosen, “Theoretical description of the photorefractive detection of the ultrasound modulated photons in scattering media,” *Optics Express*, vol. 13, pp. 7097–7112 (2005).
 - [108] M. Gross, M. Lesaffre, F. Ramaz, P. Delaye, G. Roosen, and A. Boccara, “Detection of the tagged or untagged photons in acousto-optic imaging of thick highly scattering media by photorefractive adaptive holography,” *The European Physical Journal E*, vol. 28, no. 2, pp. 173–182 (2009).
 - [109] T. W. Murray, L. Sui, G. Maguluri, R. A. Roy, A. Nieva, F. Blonigen, and C. A. DiMarzio, “Detection of ultrasound-modulated photons in diffuse media using the photorefractive effect,” *Optics Letters*, vol. 29, pp. 2509–2511 (2004).
 - [110] L. Sui, R. A. Roy, C. A. DiMarzio, and T. W. Murray, “Imaging in diffuse media with pulsed-ultrasound-modulated light and the photorefractive effect,” *Applied Optics*, vol. 44, pp. 4041–4048 (2005).
 - [111] M. Bocoum, J.-L. Gennisson, J.-B. Laudereau, A. Louchet-Chauvet, J.-M. Tualle, and F. Ramaz, “Structured ultrasound-modulated optical tomography,” *Applied Optics*, vol. 58, no. 8, pp. 1933–1940 (2019).
 - [112] S. Farahi, G. Montemezzani, A. A. Grabar, J.-P. Huignard, and F. Ramaz, “Photorefractive acousto-optic imaging in thick scattering media at 790 nm with a $\text{Sn}_2\text{P}_2\text{S}_6\text{:Te}$ crystal,” *Optics Letters*, vol. 35, no. 11, pp. 1798–1800 (2010).
 - [113] J. Li, G. Ku, and L. V. Wang, “Ultrasound-modulated optical tomography of biological tissue by use of contrast of laser speckles,” *Applied Optics*, vol. 41, pp. 6030–6035 (2002).

- [114] L. J. Nowak and W. Steenbergen, “Reflection-mode acousto-optic imaging using a one-dimensional ultrasound array with electronically scanned focus,” *Journal of Biomedical Optics*, vol. 25, no. 9, pp. 1 – 14 (2020).
- [115] R. Zemp, S. Sakadžić, and L. V. Wang, “Stochastic explanation of speckle contrast detection in ultrasound-modulated optical tomography,” *Physical Review E*, vol. 73, p. 061920 (2006).
- [116] R. J. Zemp, C. Kim, and L. V. Wang, “Ultrasound-modulated optical tomography with intense acoustic bursts,” *Applied Optics*, vol. 46, pp. 1615–1623 (2007).
- [117] S. Resink and W. Steenbergen, “Tandem-pulsed acousto-optics: an analytical framework of modulated high-contrast speckle patterns,” *Physics in Medicine & Biology*, vol. 60, no. 11, p. 4371 (2015).
- [118] D. D. Duncan, S. J. Kirkpatrick, and R. K. Wang, “Statistics of local speckle contrast,” *Journal of the Optical Society of America A*, vol. 25, no. 1, pp. 9–15 (2008).
- [119] Y. Liu, P. Lai, C. Ma, X. Xu, A. A. Grabar, and L. V. Wang, “Optical focusing deep inside dynamic scattering media with near-infrared time-reversed ultrasonically encoded (TRUE) light,” *Nature Communications*, vol. 6, no. 1, pp. 1–9 (2015).
- [120] G. Alexandrakis, F. R. Rannou, and A. F. Chatziioannou, “Tomographic bioluminescence imaging by use of a combined optical-PET (OPET) system: a computer simulation feasibility study,” *Physics in Medicine & Biology*, vol. 50, no. 17, p. 4225 (2005).
- [121] J.-B. Laudereau, *Acousto-optic imaging: challenges of in vivo imaging*. PhD thesis, Université Pierre et Marie Curie-Paris VI (2016).
- [122] H. Zhang, P. Hemmer, P. Wang, S. Rokutanda, M. Yamamoto, and L. V. Wang, “Ultrasound-modulated optical tomography using four-wave mixing in photorefractive polymers,” in *Photons Plus Ultrasound: Imaging and Sensing 2008: The Ninth Conference on Biomedical Thermoacoustics, Optoacoustics, and Acousto-optics*, vol. 6856, pp. 524–531, SPIE (2008).
- [123] R. Drever, J. L. Hall, F. Kowalski, J. Hough, G. Ford, A. Munley, and H. Ward, “Laser phase and frequency stabilization using an optical resonator,” *Applied Physics B*, vol. 31, no. 2, pp. 97–105 (1983).
- [124] B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics.*, ch. 19. Wiley Interscience (2007).

- [125] E. A. Donley, T. P. Heavner, F. Levi, M. Tataw, and S. R. Jefferts, “Double-pass acousto-optic modulator system,” *Review of Scientific Instruments*, vol. 76, no. 6, p. 063112 (2005).
- [126] R. Cubeddu, A. Pifferi, P. Taroni, A. Torricelli, and G. Valentini, “A solid tissue phantom for photon migration studies,” *Physics in Medicine & Biology*, vol. 42, no. 10, p. 1971 (1997).
- [127] M. Keijzer, W. M. Star, and P. R. Storch, “Optical diffusion in layered media,” *Applied Optics*, vol. 27, no. 9, pp. 1820–1824 (1988).
- [128] P. Di Ninni, F. Martelli, and G. Zaccanti, “The use of India ink in tissue-simulating phantoms,” *Optics Express*, vol. 18, no. 26, pp. 26854–26865 (2010).
- [129] Pádraig, “Magnetic field of circular coils using elliptical integrals, MATLAB central file exchange,” (Accessed mar. 2021). <https://www.mathworks.com/matlabcentral/fileexchange/43394-magnetic-field-of-circular-coils-using-elliptical-integrals>.
- [130] Wikipedia, “Hadamard product (matrices),” (Accessed oct. 2022). [https://en.wikipedia.org/wiki/Hadamard_product_\(matrices\)](https://en.wikipedia.org/wiki/Hadamard_product_(matrices)).
- [131] J. J. Moré, “The levenberg-marquardt algorithm: implementation and theory,” in *Numerical analysis*, pp. 105–116, Springer (1978).
- [132] The MathWorks, Inc., “Matlab and optimization toolbox release 2021b,” (2022). <https://se.mathworks.com/help/optimization/>.
- [133] J. L. Sandell and T. C. Zhu, “A review of in-vivo optical properties of human tissues and its impact on PDT,” *Journal of Biophotonics*, vol. 4, no. 11-12, pp. 773–787 (2011).
- [134] T. Durduran, R. Choe, J. Culver, L. Zubkov, M. Holboke, J. Giammarco, B. Chance, and A. Yodh, “Bulk optical properties of healthy female breast tissue,” *Physics in Medicine & Biology*, vol. 47, no. 16, p. 2847 (2002).
- [135] B. J. Tromberg, O. Coquoz, J. B. Fishkin, T. Pham, E. R. Anderson, J. Butler, M. Cahn, J. D. Gross, V. Venugopalan, and D. Pham, “Non-invasive measurements of breast tissue optical properties using frequency-domain photon migration,” *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, vol. 352, no. 1354, pp. 661–668 (1997).

- [136] S. Fantini, S. A. Walker, M. A. Franceschini, M. Kaschke, P. M. Schlag, and K. T. Moesta, "Assessment of the size, position, and optical properties of breast tumors in vivo by noninvasive optical methods," *Applied Optics*, vol. 37, no. 10, pp. 1982–1989 (1998).
- [137] T. L. Troy, D. L. Page, and E. M. Sevick-Muraca, "Optical properties of normal and diseased breast tissues: prognosis for optical mammography," *Journal of Biomedical Optics*, vol. 1, no. 3, pp. 342–355 (1996).
- [138] N. Ghosh, S. Mohanty, S. Majumder, and P. Gupta, "Measurement of optical transport properties of normal and malignant human breast tissue," *Applied Optics*, vol. 40, no. 1, pp. 176–184 (2001).
- [139] J. S. Six, T. Hughes-Riley, K. F. Stupic, G. E. Pavlovskaya, and T. Meersmann, "Pathway to cryogen free production of hyperpolarized krypton-83 and xenon-129," *PLOS ONE*, vol. 7, no. 11, p. e49927 (2012).
- [140] B. E. Treeby and B. T. Cox, "k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave fields," *Journal of Biomedical Optics*, vol. 15, no. 2, p. 021314 (2010).
- [141] H. von Friesendorff, *Evaluating imaging techniques for ultrasound optical tomography in breast tissue*. MSc thesis, Lund University (2023).
- [142] P. R. Bakic, C. Zhang, and A. D. Maidment, "Development and characterization of an anthropomorphic breast software phantom based upon region-growing algorithm," *Medical Physics*, vol. 38, no. 6 Part 1, pp. 3165–3176 (2011).