Frequency Domain High Density Diffuse Optical Tomography For Functional Brain Imaging

by

Matthaios Doulgerakis

A thesis submitted to the

University of Birmingham

for the degree of DOCTOR OF PHILOSOPHY

School of Computer Science
College of Engineering and Physical Sciences
University of Birmingham
March 2020

ABSTRACT

Measurements of dynamic near-infrared (NIR) light attenuation across the human head together with model-based image reconstruction algorithms allow the recovery of three-dimensional spatial brain activation maps. Previous studies using high-density diffuse optical tomography (HD-DOT) systems have reported improved image quality over sparse arrays. Modulated NIR light, known as Frequency Domain (FD) NIR, enables measurements of phase shift along with amplitude attenuation.

It is hypothesised that the utilization of these two sets of complementary data (phase and amplitude) for brain activity detection will result in an improvement in reconstructed image quality within HD-DOT. However, parameter recovery in DOT is a computationally expensive algorithm, especially when FD-HD measurements are required over a large and complex volume, as in the case of brain functional imaging. Therefore, computational tools for the light propagation modelling, known as the forward model, and the parameter recovery, known as the inverse problem, have been developed, in order to enable FD-HD-DOT.

The forward model, within a diffusion approximation-based finite-element modelling framework, is accelerated by employing parallelization. A 10-fold speed increase when GPU architectures are available is achieved while maintaining high accuracy. For a very high-resolution finite-element model of the adult human head with ~600,000 nodes, light propagation can be calculated at ~0.25s per excitation source. Additionally, a framework for the sparse formulation of the inverse model, incorporating parallel

computing, is proposed, achieving a 10-fold speed increase and a 100-fold memory efficiency, whilst maintaining reconstruction quality.

Finally, to evaluate image reconstruction with and without the additional phase information, point spread functions have been simulated across a whole-scalp field of view in 24 subject-specific anatomical models using an experimentally derived noise model. The addition of phase information has shown to improve the image quality by reducing localization error by up to 59%, effective resolution by up to 21%, and depth penetration up to 5mm, as compared to using the intensity attenuation measurements alone. In addition, experimental data collected during a retinotopic experiment reveal that the phase data contains unique information about brain activity and enables images to be resolved for deeper brain regions.

Acknowledgements

I would like to thank all the PhD students and academics in the Medical Imaging Lab and PSIBS, for creating such a rich environment for cross-pollination of ideas, allowing for science to thrive. A particular thanks to Prof. Adam Eggebrecht, Dr Stanislaw Wojtkiewicz and Dr Daniel Lighter for the countless inspiring scientific discussions and their practical help.

A very special thanks to my supervisor Prof. Hamid Dehghani for being the person who made the existence of such an excellent environment possible; for always being supportive and patient; for sharing his knowledge and experience; for guiding without taking control.

Finally, I would like to thank my wife for all her love and support, and my son for helping me getting up in the mornings even when I did not want to.

This work has been funded by the National Institutes of Health under Grant Nos. R01NS090874-08, RO1-CA132750, K01MH103594, and R21MH109775.

Contents

List Of Figures

List Of Tables

List Of Publications

Journal Publications:

Toward real-time diffuse optical tomography: accelerating light propagation modeling employing parallel computing on GPU and CPU.

M. Doulgerakis, A. Eggebrecht, S. Wojtkiewicz, J. Culver, and H. Dehghani.

In *Journal of Biomediomedical Optics (JBO)*, vol. 22, no. 12, pp. 1–11, 2017.

High-density functional diffuse optical tomography based on frequency-domain measurements improves image quality and spatial resolution.

M. Doulgerakis, A. T. Eggebrecht, and H. Dehghani,

In *Neurophotonics*, vol. 6, no. 03, p. 035007, 2019.

Conference Proceedings:

Towards real-time functional human brain imaging with diffuse optical tomography.

M. Doulgerakis, A. Eggebrecht, J. Culver, and H. Dehghani.

In *European Conference on Biomedical Optics, Diffuse Optical Spectroscopy and Imaging VI*, vol. 10412, p. 1041209 , 2017

Improving Functional Diffuse Optical Tomography Reconstruction Quality Utilizing Frequency Domain Measurements.”

M. Doulgerakis, A. Eggebrecht, J. Culver, and H. Dehghani.

In *Biophotonics Congress: Biomedical Optics Congress 2018*, vol 2018, pp. 5–6, 2018.

Information rich phase content of frequency domain functional Near Infrared Spectroscopy”

M. Doulgerakis-Kontoudis, A. T. Eggebrecht and H. Dehghani.

in *Neural Imaging and Sensing 2019*, 2019, vol. 1, no. 1, p. 13.

1 Introduction

The brain is a part of the nervous system, a network of cells specialised in transmitting information composed primarily by neurons. While neurons themselves can not store or process information, the connectivity pattern of the networks occurring within the brain can create specialised neuron hubs, that can perform computations or store information. This fascinating ability is not only responsible for coordinating and controlling physical functions, but it also enables the existence of intelligence and emotions. Evolutionary, the cortex is the most developed part of the human brain; it is a thin layer (1-5 mm) covering the brain surface, where most of the high-level functions - as language processing and vision- occur. While the cortex is only a thin layer, it accounts for almost two-thirds of the mass of the human brain, as its surface is dominated by bulges and fissures, called *gyri* and *sulci*, that greatly extend its area. When a neuron transmits information, in practice, it is transmitting electrical current, by physically altering the configuration of membranes in its cells; this procedure requires energy, which means that the neural tissue needs to metabolise glucose, which requires the presence of oxygen. To accommodate for these metabolic needs, a network of blood vessels is lacing the neural system, providing additional blood in the areas with neural activity, increasing the local blood flow and oxygenation; this phenomenon is called "neurovascular coupling".

1.1 Neuroimaging

Functional brain imaging is the visualization of neuron activations in the brain and has an enormous impact on the study of the brain, helping researchers understand and classify different cerebral areas and the relationships between them. It has been used to

detect, localise and classify brain regions, when the subjects are involved in decision making or social interaction tasks[1,2], understand the interplay of different brain regions when experiencing complex emotions[3–5], and while performing locomotive tasks[6,7]. The association of brain activity with both physical and psychological events is propelling applications in a wide area, such as guiding treatment and monitoring the rehabilitation progress in cases of stroke, depression, or schizophrenia[8–10]. Furthermore, real-time brain imaging can provide concurrent information, which is crucial in many clinical applications, and can also be used to develop brain-computer interfaces with potential in numerous applications: from medical, as a communication and control tool for disabled patients; in daily life, providing hands-free control of computers and machinery; to entertainment, for controller-free gaming[11–14].

There are a plethora of functional brain imaging modalities, but the ones that can provide adequate spatial accuracy, in order to satisfactorily capture brain region activation, depend upon very expensive equipment and require the subjects to be completely still, therefore greatly limiting the practical applications. Diffuse optical tomography (DOT) is an emerging imaging technique with neuroimaging capabilities, and considerable advantages over the other existing modalities: it is non-invasive, requires relatively inexpensive and portable equipment, which is wearable, and uses near-infrared (NIR) light, thus allowing some movement whilst being completely harmless to the subjects, therefore can be used in applications where other imaging techniques are unsuitable.

1.2 Other modalities

Neuron activation can be detected either directly, by measuring the electrical charges transmitted between neurons, or indirectly, by measuring changes in tissue oxygenation that take place to provide the required energy, also related to the blood-oxygen level-dependent (BOLD) effect. Two categories of functional neuroimaging modalities have been developed, the ones that measure the electrical charges and the ones that rely on the haemodynamic effects. The ones that directly measure electrical charges have a superior temporal resolution, in the order of milliseconds, while the ones based on the haemodynamic changes are bound to the vascular signals speed that can induce delays up to 5 seconds.

The most established, non-invasive, brain functional imaging techniques are briefly described in this section, with a better insight given for the case of functional NIR spectroscopy (fNIRS) and functional DOT (fDOT) in the sections 1.3 and 1.4 respectively. A summary of the characteristics of the examined neuroimaging modalities is presented in Table 1.1, which describes the type of measurement, spatial and temporal resolution, and portability. Additionally, the temporal and spatial resolution is compared graphically in Figure 1.1.

Table 1.1. Comparison of Neuroimaging Methods

Neuroimaging Method	Type of measurement	Temporal Resolution	Spatial Resolution	Portable
EEG	Direct	Very High (milliseconds)	Low (cm)	Yes
MEG	Direct	Very High (milliseconds)	High (mm)	No
PET	Indirect	Very Low (seconds-minutes)	High (mm)	No
fMRI	Indirect	Low (seconds)	High (mm)	No
fNIRS	Indirect	Low (seconds)	Low (cm)	Yes
fDOT	Indirect	Low (seconds)	Medium (mm-cm)	Yes

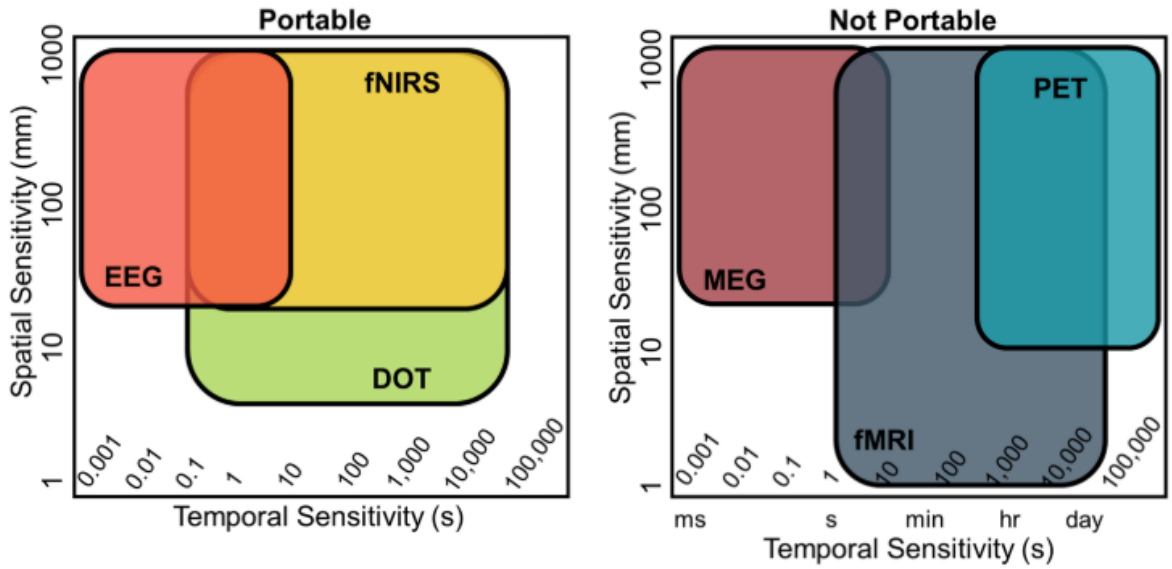

Figure 1.1 Comparison of the spatial and temporal resolution in portable (left) and non-portable (right) neuroimaging methods[15].

It is worth mentioning that there are invasive neuroimaging modalities, that require surgical procedures, such as electrocorticography[16] and intracortical neuron recording[17], but they will not be examined here, as the emphasis of this work is on non-invasive technologies.

1.2.1 Electroencephalography

Electroencephalography (EEG) is a non-intrusive and non-ionizing technique that directly measures differences in electrical current, using electrodes attached on the surface of the head, therefore providing high temporal resolution. However, due to poor conductivity of the intervening tissues, such as skull and cerebral fluid, the signal to noise ratio is poor, resulting in low spatial resolution, to centimetre accuracy[18]. Additionally, it is impossible to have a confident detection of the actual sources of energy, since currents might cross-cancel each other before reaching the surface of the head where the measurements take place[19]. EEG requires relatively low-cost equipment, on the scale of £10,000, and is relatively portable.

1.2.2 Magnetoencephalography

Magnetoencephalography (MEG) is also non-invasive, non-ionizing and provides high temporal resolution as it measures the magnetic fields resulting from the currents as they travel through the neurons. The magnetic fields are not affected by intervening tissues, therefore those systems provide high spatial resolution, to millimetre accuracy[18]. However, the measured signals are extremely small and rapidly attenuate with distance, limiting depth penetration, and very expensive equipment, costing millions of pounds, is required to provide robust measurements. Furthermore, there is the fundamental limitation on the approach: magnetic fields radially oriented to the sensors cannot be measured, therefore limiting the field of view[20].

1.2.3 Functional Magnetic Resonance Imaging

Functional Magnetic Resonance Imaging (fMRI) is non-invasive and non-ionizing, based on the BOLD effect, therefore provides low temporal resolution. The subjects are exposed to a strong static magnetic field and a gradient varying magnetic field. As

oxygenated haemoglobin (HbO_2) is diamagnetic but deoxygenated haemoglobin (HbR) is paramagnetic, the resonance between the two magnetic fields can be correlated to concentrations of oxygenated and deoxygenated haemoglobin, without explicitly measuring them, therefore there are conflicting theories regarding the meaning of the measurements[21–24]. However, the resulting spatial resolution is high, to millimetre accuracy, with very good depth penetration and low noise levels as the measurements are not affected by intervening tissues, making fMRI the golden standard in functional brain imaging. However, the required equipment is very expensive, costing millions of pounds, and not portable, limiting the practical applications, especially when considering that due to the strong magnetic fields there must not be any kind of conductive implants in the subjects.

1.2.4 Positron Emission Tomography

In Positron Emission Tomography (PET) the subject is injected with a radioisotope that is decaying relatively fast, emitting positrons that when they collide with electrons in the body cause a reaction known as annihilation where photons are released in opposite directions and are measured around the subject's head. The measurements are often combined with underlying anatomical information from MRI data, a technique called PET/MRI[25]. PET measures blood flow that can also indicate neuron activation, with a low temporal resolution, similar to BOLD, but, while the measurements are affected by intervening tissues, provides high spatial resolution, to millimetre accuracy. The method is invasive and ionizing, and besides requires relatively expensive non-portable equipment, costing hundreds of thousands of pounds.

1.3 Functional Near-Infrared Spectroscopy

NIR light propagation in biological tissue (typically between 650 – 950 nm) benefits from low absorption, thereby allows sampling of the brain cortex tissue at depths beyond 1 cm from the head surface using measurements on the scalp. Differential, spectral NIR, measurements can recover relative changes in oxy (HbO_2) and deoxy (HbR) haemoglobin concentration that convey information about brain function, akin to the BOLD effect in fMRI, and is known as fNIRS.

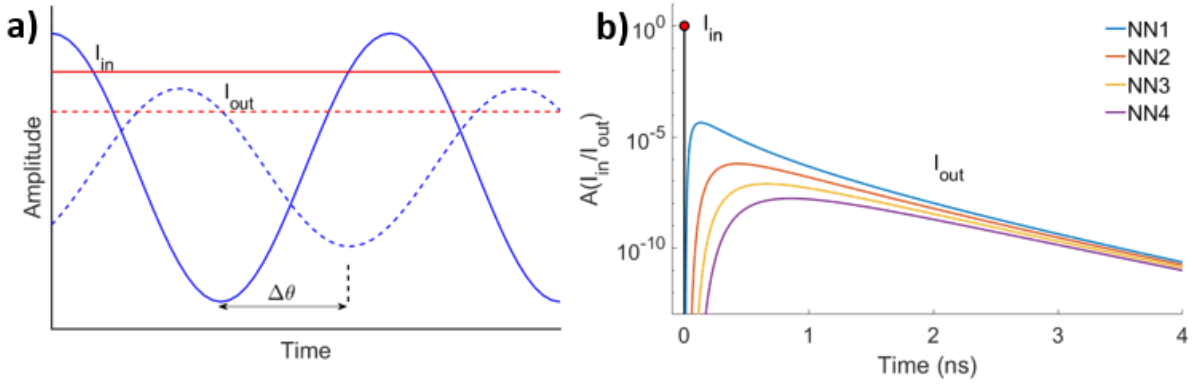

Figure 1.2 Frequency domain and time-resolved data. a) Modulated light (blue line) with RMS intensity level I_{in} (red line) is injected in the tissue; the re-emerging measured light (blue dashed line) has undergone a phase shift $\Delta\theta$ and has reduced intensity level I_{out} (red dashed line), b) Time-resolved data temporal point spread functions for different measurement neighbourhood distances of 1st (NN1) to 4th (NN4).

When continuous wave (CW) NIR light is used, measurements provide access to only light intensity attenuation from the tissue. By contrast, amplitude-modulated NIR light, typically in the range of 100-200 MHz and known as frequency domain (FD) NIRS, enables additional measurements of a phase shift of the NIR light[26]. Measurements of the phase shift are related to the average photon path-length and can be used to directly estimate rather than assume scattering parameters within the tissue, which enables more accurate optical parameter recovery[27]. FD-NIRS has been used in breast optical tomography for characterization and monitoring of lesions[28], joint imaging for arthritis

detection and treatment monitoring[29], infant brain development monitoring[30] and brain trauma assessment[31]. In previous studies utilizing FD NIRS, the phase measurements have typically been disregarded when recovering haemoglobin concentrations[32–34]. This is largely due to the fact that the accurate absolute measurement of phase is not trivial and requires complicated and often time-consuming instrument calibration procedures, as the phase can be affected by several factors such as fibre length, fibre bending, and attachment angle, each of which can present long term drifts and requires re-calibration. Nevertheless, the frequency-domain multi-distance approach has been shown to provide increased sensitivity to cerebral oxygenation using intensity and phase data[35,36], however, the multi-distance approach has been extremely challenging to adopt with an HD-DOT tomographic set-up due to the strong demands on the dynamic range of the high bandwidth detector instrumentation. Specifically, given the logarithmic increase of signal attenuation in tissue as a function of distance, the detector gain needs to be optimised to allow accurate detection from both near and far sources, which is possible in simple setups using only 1 source and 4 detectors: i.e. where each detector only detects signal from one source. However, for an HD-DOT system this becomes more challenging when a single detector will measure photons arriving from multiple sources (at different distances) and thereby affecting the overall quality of the Signal-to-Noise ratio and hence accurate phase measurements.

Alternatively, measurements of pico-second modulated light (time-resolved, or TR) provide a direct measure of the temporal point spread function of the detected photons to account for both tissue absorption and scattering. However, systems with enough channels to provide overlapping measurements, thus achieving DOT, have a low

temporal resolution (75 seconds for a full imaging cycle[26]), though various mathematical models have been proposed to increase the temporal resolution[37,38]. Although TR systems are considered the golden standard within fNIRS, typically only spectroscopic measures are used to provide bulk tissue optical measurements and these systems have not been widely adopted for HD-DOT.

The CW, FD, and TR measurements are all sensitive to haemodynamics and are therefore indirect measures of variations in brain function. An alternative method, known as the event-related optical signal (EROS) or fast optical signal, uses just the phase measurements of FD systems and has been reported to be sensitive to neuronal activity[39]. It has been suggested that this technique does not suffer from the haemodynamic delay effect, and may be sensitive to scatter-related changes within firing neurons. However, in practice this technique requires extensive stimulus repetitions to allow adequate signal averaging to obtain required signal-to-noise, thus cancelling the real-time monitoring potential and resulting in experimental sessions of extremely long duration (approximately 60 minutes for one stimulus[40]) that are impractical for many applications.

To investigate haemodynamic changes in a practical setting, measurements of differential phase changes are much easier to apply and have shown a high correlation with the fMRI BOLD signal, suggesting that phase changes convey information about functional haemodynamic phenomena [41]. This is an expected result as changes in the absorption of a medium will not only affect the total number of detected photons (the measured intensity) but also the average photon diffusion path-length (the measured phase). How a change in absorption within the volume leads to differential changes in

both (logarithmic) intensity and (linear) phase can be understood when considering the distribution of the time of flight of the photons, Figure 1.2. Early arriving photons travel a shorter distance and stochastically sample more superficial tissues while the later arriving photons travel longer distances and therefore tend to both stochastically sample deeper tissues and have a higher probability of undergoing absorption events[42]. Therefore, the phase measurement sensitivity profile to changes in absorption is deeper than the sensitivity of the intensity attenuation measurement because early and late photons contribute equally to the mean time of flight while the total intensity is dominated by early arriving photons[43].

1.4 Functional Diffuse Optical Tomography

The two most established functional imaging techniques, functional MRI (fMRI) and positron emission tomography (PET) have a common set of disadvantages: they are not portable, very expensive equipment is used and usually require constraints on the subject to be completely still, that can be uncomfortable for anyone, more so in cases of claustrophobic patients. Furthermore, fMRI cannot take place when metallic or electronic implants exist, and PET exposes the subjects to ionising radiation, therefore must be avoided in paediatric cases and long-term studies. In contrast, DOT is non-ionizing, requires relatively low-cost equipment which is wearable and therefore allows some movement of the subject, is relatively portable and can be used in clinical configurations where the use of fMRI or PET is not possible, for example as a bedside monitoring tool.

Diffuse Optical Tomography (DOT) is a soft tissue imaging technique that is based on injecting NIR light into a volume and measuring the resurfacing light, therefore often

also referred to as NIR optical tomography. Due to strong scattering, sparse optical NIRS measurements obtain low spatial specificity relative to fMRI. Image quality can be significantly improved by arranging NIR elements in a high-density (HD) grid to provide overlapping measurements that allow tomographic reconstructions, a procedure known as functional diffuse optical tomography (fDOT).

DOT is non-invasive and non-ionizing, and has been used for brain functional imaging[44–48], neonatal brain monitoring[49,50], as a diagnostic tool for breast cancer[51–55], measuring tissue metabolism[32], diagnosing arthritis in small joints[56] and performing trans-rectal endoscopy[57]. Functional DOT (fDOT) is based on the haemodynamic effects, exploiting the distinctive absorption coefficients of oxygenated and deoxygenated haemoglobin in the NIR spectrum, therefore it has a relatively low temporal resolution, of seconds, whilst providing good spatial accuracy from millimetre to centimetre range. Additionally, the size of the equipment allows for bedside monitoring, use in operating theatres and portability, for example an ambulance could be equipped with an HD-DOT system.

1.5 Proposed approach

Almost all reported fDOT neuroimaging studies employ continuous wave (CW) measurements of light level attenuation. Frequency domain (FD) strategies use radio frequency (RF) modulation of the NIR sources to enable additional measurements of phase shift in the signal. To date, FD measurements have not been used in the context of tomographic optical functional brain imaging and the effects of incorporating phase information together with intensity in HD-DOT have yet to be investigated.

It is hypothesised that FD-HD-DOT will improve the imaging quality of HD-DOT, which has already achieved results comparable to fMRI, in terms of resolution and localisation, as well as depth sampling. The potential improvement in reconstructed image quality due to the additional incorporation of phase shift measurements was investigated, which reflect the time-of-flight of the measured NIR light, within the tomographic reconstruction from high-density measurements.

However, computations become an underlying issue for FD-HD-DOT, making the problem computationally intractable for regular desktop computers. Therefore, computational tools employing parallel linear solvers on the graphics processing unit (GPU) and central processing unit (CPU) were developed to accelerate the light propagation modelling in the forward problem; while a memory-efficient algorithm for the formulation of the inverse problem has been proposed, also allowing parallelisation. These tools were evaluated against the standard computational toolboxes to ensure high solution accuracy.

The current state of the art HD-DOT algorithms were expanded, in order to include phase information in the reconstructions, evaluating the effect in image quality, in terms of localisation, resolution and maximum imaging depth, in simulations with experimentally derived noise models. Furthermore, FD data collected in-vivo during standard retinotopic experiments were analysed, to identify if phase conveys neuroactivity information; and finally, performed reconstructions using amplitude only data and amplitude and phase data to assess potential imaging benefits.

The key motivation behind these advances are (1) to develop a computational model that is able to predict photon-migration in tissue in an efficient and timely manner and

(2) to develop an inversion algorithm (inverse-model) that is able to incorporate this additional data-type (phase) without detrimental loss in computational time while (3) demonstrating the benefits of phase information in sampling depth and resolution.

1.6 Thesis Outline

The first chapter of this thesis was an introduction to neuroimaging, including an overview of established approaches with a more detailed consideration of optical imaging with fNIRS and fDOT, including motivation for the proposed FD-HD-DOT approach.

In chapter 2, the computational aspect of HD-DOT is highlighted, delving into the underlying physics and mathematics employed to achieve image reconstruction and emphasizing the need for computational optimisations.

Chapter 3 presents the proposed computational methods leading towards real-time FD-HD-fDOT, featuring parallel computing and algorithm optimizations for the computationally expensive parts of the procedure.

Chapter 4 presents a comparison between continuous wave and frequency domain HD-DOT, performing simulations on 24 adult head models while employing an experimentally derived noise model, concluding on a statistically significant image quality improvement when using the frequency domain approach.

Chapter 5 provides an analysis of in-vivo data and reconstructions, during retinotopic protocols, revealing stimuli-driven information in the amplitude and phase data, and confirming the simulation findings regarding the accuracy of the frequency domain reconstructions.

Finally, chapter 6 concludes on the benefits of FD over CW HD-DOT and the value

of the developed computational tools and discusses possible future work on questions that have arisen through this work.

2 Computational Methods

This chapter provides an overview of the underlying mathematics that directly affect the computational aspects of DOT and emphasises the necessity of parallelization, specifically considering the existing implementation within NIRFAST[58]. DOT can be summarised in two steps: the forward model, which estimates light propagation in tissue providing predictions of light levels at the volume surface, with given underlying optical properties; and the inverse model, which uses the predictions from the forward model in addition to measured light levels to estimate changes to the underlying optical properties.

The most computationally expensive procedure is the solution of large sparse linear systems, involved in estimating light propagation in the forward problem, and the formulation of the sensitivity matrix, known as the Jacobian, in the inverse problem.

2.1 Forward Problem

The first step of the DOT algorithm, the forward problem, is the basis for the image reconstruction, therefore, it must be as accurate as possible, as any errors will heavily affect the solution of the inverse problem. Many approaches have been introduced to model light propagation through a diffuse medium, and most of them can be classified into two categories: analytical[59–61] or numerical[62,63] methods.

The propagation of light through a turbid medium is described accurately by the Radiative Transport Equation (RTE)[64]. In cases where scattering dominates over absorption, as in biological tissue, the Diffusion Approximation (DA) (Eq. 2.1) can be derived from the RTE, which is significantly easier to solve but sacrifices some accuracy.

Analytical methods, calculate the light field by solving the RTE or the DA as can be applied in a given shape medium, using the Born and Rytov approximations[27,65]. These

approaches are computationally fast, but they are limited in modelling only for simple shapes and simple optical properties distributions. In contrast, numerical approaches can model complex, arbitrary geometries and media with complex heterogeneous optical properties distributions.

The accuracy of the numerical solutions is greatly affected by the prior knowledge of the underlying tissue geometry, therefore their maximum potential is reached when combined with input from standard imaging techniques[66,67] or generic atlas models[68]. The numerical approaches for calculating the forward model can be divided into two categories: the ones using Galerkin discretisation, like the Finite Elements Method (FEM)[58], and those using Monte Carlo (MC)[69] approach. The Galerkin method is a technique where a heterogeneous nature problem is split into many smaller parts, creating a non-uniform mesh. Effectively each of the elements can be considered homogeneous and the diffusion equation can be formulated as a linear system, describing simultaneously all elements. The MC approaches also divide the medium into volumetric elements[69,70], and then focus on the path of the photons as they travel through the medium. Traditionally uniform discretisation was used to divide the medium but recently, non-uniform discretization has been achieved as well. While the MC approaches are considered the golden standard in light propagation modelling, as they can solve the RTE, the computational time is rather slow, because millions of photons must be simulated in order to achieve an accurate solution, making the light propagation modelling very time-consuming.

The modelling of light propagation through a medium relies on two parts: the physics to describe photon diffusion and the FEM formulation that discretises the problem and

allows it to be computationally tractable. To describe light propagation through a diffuse medium the diffusion approximation equation is used:

$$-\nabla \cdot \kappa(r)\nabla\Phi(r,\omega) + \left(\mu_a(r) + \frac{i\omega}{c_m(r)}\right)\Phi(r,\omega) = q_0(r.\omega) \qquad \text{Eq. 2.1}$$

where μ_a is absorption coefficient, $q_0(r,\omega)$ is an isotropic source term and $\Phi(r,\omega)$ is the photon fluence rate at position r and frequency ω. $\kappa=1/3(\mu_a+\mu_s')$, where μ_s' is the reduced scattering coefficient. Finally, $c_m(r)$ defined as $c_0/n(r)$ where c_0 is the speed of light in vacuum and n(r) is the refractive index. Mixed-boundary conditions (also known as Robin or index mismatch type III for the FEM model are assumed, to deal with the case where light escapes the medium and does not return[71,72].

The FEM discretization consists of a volume Ω, divided into E elements, usually tetrahedral, joint in N vertex nodes. The fluence $\Phi(r)$ is approximated by a polynomial function:

$$\Phi^h(r) = \sum_{i=1}^{N} \Phi_i u_i \Omega^h \qquad \text{Eq. 2.2}$$

where Ω^h is finite dimensional subspace spanned by $\{u_i(r); i= 1\ldots N\}$ basis functions. The problem then becomes a sparse, well-posed large dimensions linear problem[58,73,74] of the form:

$$M\Phi = q_0 \qquad \text{Eq. 2.3}$$

where specifically, M is:

$$\left(K(\kappa) + C\left(\mu_a + \frac{i\omega}{C_m}\right) + \frac{1}{2A}F\right) \qquad \text{Eq. 2.4}$$

where the matrices $K(\kappa)$, $C\left(\mu_a + \frac{i\omega}{c_m}\right)$, and F have entries defined by

$$K_{ij} = \int_{\Omega} \kappa(r)\nabla u_i(r) \cdot \nabla u_j(r) d^n r \qquad \text{Eq. 2.5}$$

$$C_{ij} = \int_{\Omega} \left(\mu_a + \frac{i\omega}{C_m}\right) u_i(r) u_j(r) d^n r \qquad \text{Eq. 2.6}$$

$$F_{ij} = \int_{\partial\Omega} u_i(r) u_j(r) d^n r \qquad \text{Eq. 2.7}$$

and the source vector q_0 has terms

$$q_{0_i} = \int_{\Omega} u_i(r) q_0(r) d^n r \qquad \text{Eq. 2.8}$$

where the dimension of *M* is *N* by *N*, and q_0 is representing the sources and has dimension *N* by the number of sources, *Q*, of the DOT system.

The resolution of the FEM meshes heavily affects the solution for the forward model, with high-resolution meshes (more vertices) representing more accurately the underlying model. The resolution of MRI data for tissue segmentation is ~1 mm, so a mesh with element volume of ~1 mm^3 should be able to sufficiently describe the details captured by the MRI. Additionally, dividing the volume in higher resolution meshes represent the diffusion equation better, by minimizing the discretization error invoked by the Galerkin discretisation, as the volume of elements tend to zero, the discretization error also tends to zero. However, increasing the mesh resolution results in the requirement of solving a larger linear system (Eq. 2.3) to estimate the light fluence, thus, until now, dramatically increasing the computational cost, a problem that is addressed in chapter 3.1.

2.2 Inverse Problem

The purpose of the inverse problem is the recovery of the optical parameters at each node of the FEM model by exploiting the information contained in the forward model in conjunction with the boundary measurements taken from the surface of the volume. The aim is to find such optical property values, μ, that minimise the difference between calculated, Φ, and measured, Φ_M, boundary data. To this end, a Tikhonov minimisation is employed, with a regularisation parameter, λ, which is defined as the ratio of the variance of the measurement data with respect to the variance of optical properties $\lambda = \sigma^2_{\Phi_M}/\sigma^2_{\mu}$:

$$\min_{\mu}\left\{\sum_{i=1}^{NM}(\Phi_M^i - \Phi^i) + \lambda\sum_{j=1}^{NN}(\mu_j - \mu_{0,j})\right\} \qquad \text{Eq. 2.9}$$

where NM is the number of measurements and NN the number of nodes, μ_0 is an initial guess of the optical properties of a node, that in the case of functional brain imaging is usually assuming values according to prior knowledge of tissue properties. This is achieved by building a sensitivity matrix, the Jacobian, which defines the relationship between changes in the boundary data, $\partial\Phi$, with respect to optical properties, $\partial\mu$, of the FEM nodes, $J = \left(\frac{\partial\Phi}{\partial\mu}\right)$. As functional brain imaging is based on the absorption spectrum of the oxy and deoxy haemoglobin, therefore this work only accounts for absorption only changes, assuming that the scattering remains constant.

For high-resolution meshes, the number of nodes is much larger than the number of measurements (NN>>NM), therefore approaches based on perturbations of nodal optical values to build the Jacobian are not computationally efficient[75].

The approach employed to form the Jacobian is the Adjoint method[76], where the direct fluence Φ^i for each source *i*, and the Adjoint fluence Φ^j_{Adj} for each detector *j* must be calculated. Then for each node location *r*, the product of the direct and the adjoint field is considered, intergraded for each node *N* of the *k* elements that belong to the connectivity list τ of this node; also considering basis function u for each element *k*. The result is a dense matrix of dimensions NN x NM that describes the change of surface measurements with regard to a change of the optical properties of the medium. So, the construction of the Jacobian requires to model light propagation, as in the forward problem, twice; once for sources and once for detectors, a procedure that is accelerated employing parallelization as described in chapter 3.1.

$$J(i,j,r) = \left[\sum_{k|N_k \in \tau(r)} \Phi^i_k u_k(r)\right] \times \left[\sum_{k|N_k \in \tau(r)} \Phi^j_{Adj,k} u_k(r)\right] \qquad \text{Eq. 2.10}$$

The dimensions of the Jacobian, for a single wavelength, in continuous wave mode, is the number of associated source/detector (SD) pairs by the number of nodes of the mesh. This may result in Jacobians of extremely large dimensions as the number of SD pair measurements increases with HD DOT and as discussed earlier, the number of nodes has to be large for functional imaging, for example in the HD DOT system modelled in this work, on the highest resolution mesh the Jacobian would have dimensions about 3,500×600,000. This is further increased when considering the frequency domain spectrally constrained recovery (Eq. 4.6), where the size of the Jacobian is $[4\mathrm{NM} \times 2\mathrm{NN}]$, which, for the system described here, equals 14,000×1,200,000, which would require 135 Gb of memory to be stored, this problem is addressed in chapter 3.1.4.

fDOT is concerned with small dynamic changes that occur during brain activations, therefore, the inverse model can be considered linear. This is possible since the goal in functional imaging is to detect changes from a reference baseline, often referred to as temporal/difference/dynamic imaging. Assuming that such functional changes are small, they will give rise to small changes in the measured data, therefore allowing the use of Rytov approximation for mapping them directly to small changes in optical properties and hence functional changes[77]. More precisely, even though the attenuation of optical signals with respect to the underlying optical properties is non-linear, a small change in signal attenuation can be considered 'locally' linear with respect to a small change in optical properties. When a good assumption of the underlying optical properties distribution is considered (often derived directly from published data for each biological tissue type) and used in conjunction with subject specific models (as derived from MRI or a registered atlas), a functional temporal/dynamic event is mapped directly by using these small linear temporal changes.

A Moore-Penrose pseudoinverse with Tikhonov regularization can be used to find an approximation of the inverse of the Jacobian, $\mathbf{J}^{\#}$, and perform a single-step recovery of absorption.

$$\mathbf{J}^{-1} \approx \mathbf{J}^{\#} = \mathbf{J}^{\mathrm{T}}(\mathbf{J}\mathbf{J}^{\mathrm{T}} + \alpha\mathbf{I})^{-1} \qquad \text{Eq. 2.11}$$

$$\partial\mathbf{x} = \mathbf{J}^{\#}\,\partial\mathbf{y} \qquad \text{Eq. 2.12}$$

The regularisation term is calculated as $\alpha = a * \max\left(\mathrm{diag}(\mathbf{J}\mathbf{J}^{\mathrm{T}})\right)$. The weight of this regularisation parameter provides a tuning between high spatial frequency noise and smoothness in the image. Due to this association with image resolution[78], when running simulations without noise, fine-tuning this parameter can improve image resolution

beyond an experimentally and physiologically meaningful range. In practice, this parameter is dictated by the noise of the measurements[79], therefore all simulations in this work use the same regularisation weight of $a = 0.01$, following previously reported methods[80].

The result is a dense, under-determined linear problem that must be solved to get the update of optical properties on the nodes. A variety of approaches have been proposed to reduce the computational complexity of the Jacobian inversion. Often, a smaller reconstruction basis is defined, e.g. using a second, coarser finite element mesh or a pixel basis for the reconstruction[81,82], or constraining the reconstruction to only the cortex[83] or the brain region[84], to reduce the number of unknowns in the inversion problem therefore also reducing the computational cost. More recently, handling the Jacobian as a sparse matrix[85] has also been proposed.

2.3 Discussion and Conclusions

The emergence of fDOT in clinical settings is strongly related to the advancements in computing within the last two decades. The parallel processing power of GPUs offers a low-cost, easily available solution for the acceleration of linear operations, such as the ones involved to estimate light propagation, which was the computational bottleneck of the forward problem. Additionally, as the volume of data is constantly increasing, programming libraries that handle sparse linear operations have been developed, which further accelerated computations and improved memory efficiency. For the inverse problem, the calculation of the sensitivity matrix involves integrating over the number of nodes times their connectivity, which usually numbers in the millions, resulting to extremely large matrices that are impossible to store in memory.

In chapter 3, computational tools and optimization methods are described that allow the acceleration and computational tractability of DOT , by parallelizing the solution of linear systems in the forward problem while sparsifying and parallelising the creation of the Jacobian in the inverse problem, in order to enable the reconstruction algorithm to be executed in relatively low-cost desktop or laptop computers.

3 Towards real-time optical tomography

Parameter recovery in diffuse optical tomography is a computationally expensive algorithm, especially when used for large and complex volumes, as in the case of human brain functional imaging, whereby the lack of real-time solutions is impeding research and clinical applications. This chapter proposes computational tools for the acceleration and optimization of the forward and the inverse problem, intending to accommodate FD-HD-DOT modelling using a standard desktop or laptop personal computer. The development of these tools is motivated by the desire to have fast and efficient algorithms, based on either subject-specific or atlas-registered models. In each case, an accurate forward model and data inversion matrix needs to be calculated, in a computationally efficient manner, for each subject and experiment. In the subject-specific case, the model needs to be calculated from subject specific information (e.g. as derived from MRI) containing all internal tissue structures. For the atlas-registered case, the atlas (which contains all internal structures from an atlas-based population average) is first registered to the subject's geometry and then used. In either case, although the problem is assumed to be linear for parameter recovery, the mapping function (mapping measured data to optical changes, known as the Jacobian) needs to be calculated for each subject, which itself is derived from the forward problem.

3.1 Forward Problem computational optimisation

The acceleration of the forward model, within a diffusion approximation based finite element modelling framework, is addressed, employing parallelisation to expedite the calculation of light propagation in realistic adult head models. The proposed methodology is applicable for modelling both continuous-wave and frequency-domain

systems with the results demonstrating a tenfold speed increase when GPU architectures are available, whilst maintaining high accuracy. It is shown that for a very high-resolution finite element model of the adult human head with ~600,000 nodes, consisting of heterogeneous layers, light propagation can be calculated at ~0.25 seconds per excitation source.

This work provides tractable solutions to overcome the current computational time and memory limitations arising when dealing with high-resolution FEM meshes and HD source-detector (SD) pairs, both important features for high quality functional DOT (fDOT) brain imaging. Additionally, an extended evaluation is performed, on high resolution meshes with up to ~600,000 nodes, based on a realistic anatomical head model, with five tissue layers, focusing on achieving the desired numerical accuracy for functional brain imaging. Furthermore, support for complex numbers for the cases of frequency-domain simulations is incorporated.

3.1.1 Background

In this chapter, the acceleration of the forward light propagation model is examined, while maintaining numerical accuracy. Specifically, the focus is on the application of DOT for functional imaging of an adult human head, employing the finite element method (FEM) to solve the diffusion approximation (DA) (Eq. 2.1) for modelling of light propagation as implemented within the NIRFAST[58] modelling and image reconstruction software package.

The proposed acceleration approach relies on employing parallel computing to expedite the solution of the forward problem, an option that has recently become popular due to the relatively low cost of GPUs. Most modern CPUs have 4 to 8 logical cores

(processing units) which are very fast (~3GHz), whilst GPUs have hundreds (700-4,000) of cores but are slower (~1GHz); therefore GPUs are suited to solve problems that can be divided into subproblems, which can be solved in parallel on multiple processors. The overall principle is that the CPU and the GPU can cooperate, sharing the load of computations, which with appropriate distribution of tasks, massively accelerates the computing process.

That has attracted the attention of researchers for solving similar problems in medical imaging[86–88]. In DOT the acceleration of the forward model with GPUs has been employed for Monte-Carlo algorithms[89–91], where simulating the behaviour of each photon can be efficiently parallelised, with reported accelerations in the scale of $10^2 - 10^3$. However, millions of photons must be simulated to achieve an accurate solution, so modelling the total fluence for a geometrically large volume and multiple excitation sources still requires time in the order of tens of minutes.

Accelerating the forward model of DOT using GPU parallelization has also been reported with FEM formulation. Specifically, the acceleration of the forward model solution has been proposed for frequency domain simulations[92]. Unfortunately, due to a lack of sparse arithmetic architecture, computationally tractable mesh sizes have been limited to ~9,000 nodes, due to the excessive memory requirements. Solving the forward model employing GPU parallelization, in continuous wave, with parallelization over the nodes and over the excitation sources simultaneously, using a block-based formulation of the forward linear system has also been proposed[93]. However, due to the size of the augmented block matrix, this approach is only applicable in small size meshes, with up to ~2,500 nodes. Using multiple GPUs to solve the forward problem in continuous wave,

for infant brain studies has been suggested[94] but has been only evaluated qualitatively in a homogenous phantom head model, again in meshes with ~9,000 nodes. An approach combining CPU and GPU parallelization was proposed[95] where the imaging domain is decomposed into overlapping subdomains, therefore allowing a high level of parallelization for the forward problem. However, this approach was only evaluated in continuous wave, on a simplified cylindrical geometry with homogenous optical properties where the decomposition in regular overlapping subdomains is a straightforward procedure; while in the case of a complex volume, as the adult head, such decompositions are not a trivial task. Finally, a framework for the solution of the forward problem in continuous wave and frequency domain, accelerated in the GPU, was proposed and evaluated in homogenous cylindrical models with up to 330,000 nodes[96]. This work identified the relation between the error of the iterative solution and the optical properties of the volume, and found the single-precision numerical accuracy ($2^{-23} = 1.19 \times 10{-}7$) insufficient for solving for a wide range of optical properties. Section 3.1.2.5 further expands and explains the topic of numerical accuracy. However, none of the previous work has evaluated the accuracy and computational speed of iterative solvers on anatomically realistic head models, with high-resolution meshes and high density (HD) DOT system.

The work presented in this chapter will highlight the importance of the balance between computational speed and numerical accuracy for computational models in HD-DOT. Specifically, it will be shown that the computational time of any iterative based forward solver can be greatly reduced by increasing the termination threshold. However this practice can lead to numerical inaccuracies as it is spatially dependent with regard

to illumination source locations (i.e. initial boundary conditions). Finally, it will be demonstrated that the computational speed improves when using parallel computing, as expected, particularly when appropriate iterative matrix inversion algorithms and preconditioners are utilised.

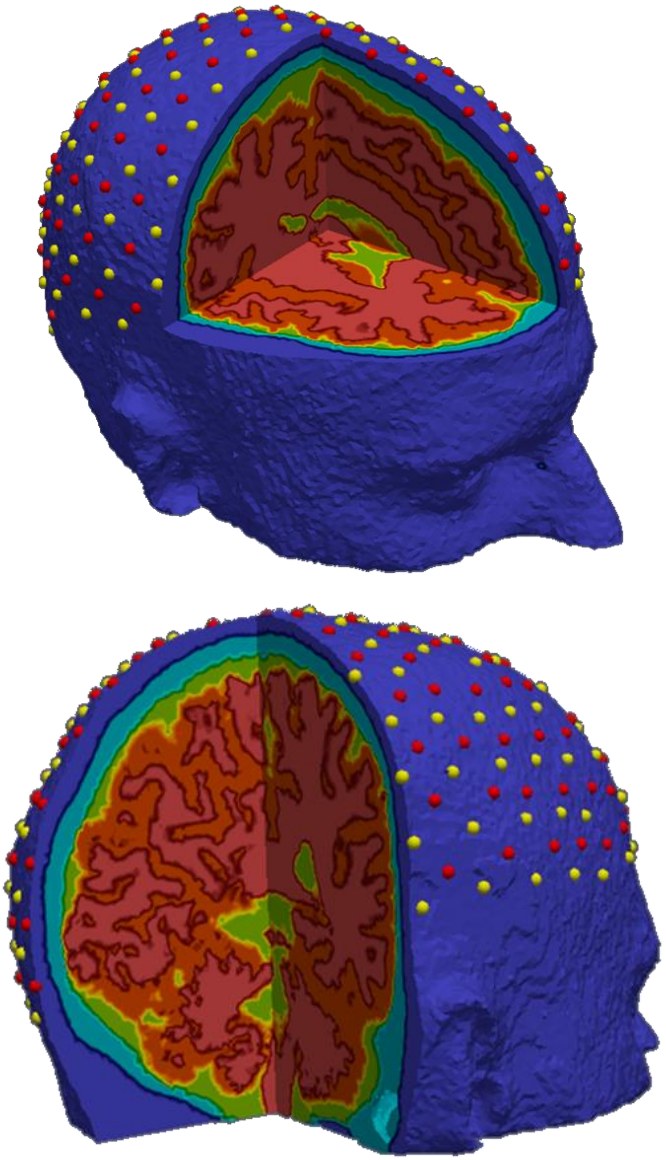

Figure 3.1 The modelled high-density DOT system. 158 sources (red), 166 detectors (yellow), on an adult head model with 5 tissue layers.

3.1.2 Method

The procedure followed in DOT image reconstruction can be summarised by the following consecutive steps: modelling the light propagation through the medium, also known as the forward problem and a parameter recovery process based on the forward model and NIR measurements, also known as the inverse problem.

The first step of the DOT algorithm, the forward problem, is the basis for the application of model-based image reconstruction, therefore, it must be as accurate as possible, numerically and geometrically, as any errors will affect the formulation of the inverse problem.

The accuracy of the numerical solutions of FEM is greatly affected by the prior knowledge of the underlying tissue geometry; therefore the maximum potential is reached when FEM is combined with input from other standard imaging techniques[66,67] or generic atlas models[68]. Mesh generation based on structural images from other modalities, usually MRI, is a well-studied problem. There are existing algorithms that automatically segment tissue layers and create surface-based meshes[97]. In this work the geometry and optical property distribution are dealt with by using subject specific meshes, produced by segmenting MRI scans. Subject specific meshes are used in most of the experiments, with the exception of the experiment described in section 3.2.2.2 where an atlas was used. Optical properties for each tissue layer were assigned according to existing literature.

When the model volume remains constant, but meshes with more elements are created, what effectively changes is the resolution of the mesh. Higher resolution meshes have elements with smaller volume and therefore minimise partial volume effects due to

mesh elements integrating over multiple, segmented tissue regions. The fine complex structures, such as the brain cortex in a head model, can be modelled more accurately with a high-resolution mesh, providing optical properties assigned to each node (or element) that are more likely to represent the underlying baseline optical properties of the tissue at each position.

When the volume is meshed, FEM is employed to formulate a discretised weak form representation of the DA for each node of the mesh. It follows that as the volume of each element tends towards zero (increasing the mesh resolution), the calculated approximation becomes more accurate, therefore typically in fDOT very high-resolution meshes, with up to 600,000 nodes are used. This is a domain size-dependent problem, a smaller volume such as an infant's head will require fewer nodes to achieve elements of sufficiently small volume. Additionally, dividing the volume into a higher resolution mesh minimises the discretization error introduced by FEM. However, increasing the mesh resolution results in the requirement of solving a bigger linear system to estimate the light fluence, which, until now, has dramatically increased computational time. The focus of this work is optimizing numerical approaches employing FEM-DA to estimate light propagation - a well-studied problem that assumes the DA is valid for all tissue properties used[58,73,74]. However, the advancements described herein can be applied to any models (e.g. Radiative Transport Equation) based on discretised approximations.

FEM is a numerical technique where a heterogeneous problem is divided into many smaller parts, creating a non-uniform mesh, consisting of elements defined by connected nodes. The diffusion equation can then be discretised and represented as a set of linear equations, describing simultaneously all the nodes and hence the entire medium. The

problem thus reduces to a sparse, well-posed linear problem of the form: $\boldsymbol{M\Phi = q}$ dimensions N by N, with N denoting the number where **M** is a sparse matrix with of nodes; **q** represents the sources and has dimension N by the number of sources Q of the DOT system, and **Φ** is the photon fluence rate for all nodes for each source, as has dimensions N by Q.

3.1.2.1 Computational problem

The solution of large sparse linear systems involved in the forward model is currently the computational bottleneck of the DOT algorithm. To solve the linear systems arising in the forward modelling, in the form $\mathbf{Ax} = \mathbf{b}$, for **x**, where **A** has dimensions $N \times N$, where N is the number of unknowns, the inverse of **A** must be calculated. However, calculating a matrix inverse is computationally inefficient, therefore a variety of algorithms have been proposed that can solve linear systems without explicitly calculating a matrix's inverse. These algorithms can be either direct, providing an exact solution, or an approximate, usually employing an iterative algorithm. The storage convention used in this work to represent sparse matrices within memory is the compressed row storage, where the non-zero values are stored in one vector, the corresponding column indices in another vector, and only the element indices where row changes have occurred are stored in a third vector. This requires $2N_{NZ} + N + 1$ space in memory for a $N \times N$ matrix with N_{NZ} non-zero entries.

3.1.2.2 Direct solvers

The most popular direct solver is the Gaussian Elimination, also known as row reduction, where the echelon form of **A** is calculated through row operations on the augmented matrix $(\mathbf{A}|\mathbf{b})$. The echelon form of **A** is an upper triangular matrix, making

the solution of the linear system easy through backward substitution. However, Lower Upper (LU) factorization is considered the standard efficient computational approach for the direct solution of a linear system. The LU factorization decomposes **A** into lower (**L**) and upper (**U**) triangular matrices. Substituting **L** and **U** in $\mathbf{Ax} = \mathbf{b}$ gives $\mathbf{LUx} = \mathbf{b}$, and letting $\mathbf{Ux} = \mathbf{Y}$, then $\mathbf{LY} = \mathbf{b}$. Now it is trivial to solve $\mathbf{LY} = \mathbf{b}$ for **Y** through forward substitutions and then solving $\mathbf{Ux} = \mathbf{Y}$ through backwards substitutions. The advantage of LU factorization over the traditional Gaussian elimination is that decomposing **A** into **L** and **U** is independent of **b**, also known as the right-hand-side vector. This allows **L** and **U** to be used for solving for multiple right-hand side vectors. However, this approach has very large memory requirements of $N^2 + N$ and high computational cost, of $\frac{2}{3}N^3$ floating-point operations (FLOPS), to solve a full linear system. A more efficient alternative, that can be used only if **A** is Hermitian, therefore symmetric when real, is the Cholesky factorisation, where **A** is decomposed to $\mathbf{L}\,\mathbf{L}^*$, where * denotes the transpose conjugate operator, requiring $\frac{1}{2}N^2 + N$ memory for a full system. The linear system can be solved as with the LU method, substituting $\mathbf{U} = \mathbf{L}^*$, with computational cost, for the solution of full systems, of $\frac{N^3}{3}$ FLOPS. However, in the case of sparse linear systems, such as resulting from the FEM formulation, memory and computational costs are related to the number of non-zero elements of **A** rather than the size N and, additionally, there are reordering strategies that when applied on sparse matrices allow more sparse factorizations. Specifically, in this work, the approximate minimum degree permutation algorithm was found to produce the most sparse factorizations, therefore was used for all the direct solvers. Nevertheless, factorization approaches rely on forward and/or

backward substitutions to provide a solution, therefore they cannot be efficiently parallelised. In MATLAB when solving linear systems invoking the backslash operator, the Cholesky approach is used when the matrix is Hermitian, otherwise the LU approach is employed. The "spparams"[98] command was used to confirm that all the real linear systems were solved with Cholesky solver and all the complex with LU. The MATLAB backslash operator is considered as the numerical ground truth for the solution of linear systems throughout this work.

3.1.2.3 Iterative solvers

To overcome the computational limitations of direct linear algebraic solvers, a variety of approximate linear solvers have been proposed that can be classified into three general categories: iterative, multigrid, or domain decomposition methods[99]. Multigrid and domain decomposition methods can be very efficiently parallelised, with solving speed not greatly affected by the size of the linear system. However, these methods require additional input parameters (e.g. the range of eigenvalues of the system, restriction and prolongation parameters, smoothing operators) that might be difficult to define and may vary for different systems to efficiently converge to adequate approximations. As such, these methods work best when they are tailored to solve a very specific problem. In contrast, iterative solvers are generic and require little or no additional input from the user, and thus are traditionally chosen for the solution of linear systems describing light propagation.

Iterative approaches approximate a solution vector $\mathbf{x}_n$ and then attempt to minimise the residual $r_n = \|\mathbf{b} - \mathbf{A}\mathbf{x}_n\|$ through n iterations, until r_n is lower than a user-defined residual threshold r_{th}. However, in practice, the termination criterion is defined

relatively as $t_c = \frac{r_n}{r_0}$, where r_0 is the residual after the initial guess, with the initial guess $\mathbf{x}_0$ set usually as a vector of zeros. Using a relative threshold ensures that the iterations will converge with a final r_n usually within the same order of magnitude as the t_c, even when the number of unknowns is very large. Iterative approaches usually work on a projection space for increased computational efficiency. The most established projection scheme is the Krylov subspace, which is based on the Cayley-Hamilton theorem that implies that the inverse of a matrix can be found as a linear combination of its powers. The Krylov subspace generated by a $N \times N$ matrix **A**, and a vector **b** of dimension N, is the linear subspace spanned by images of **b** under the first α powers of **A**.

$$K_\alpha(\boldsymbol{A}, \boldsymbol{b}) = span\{\boldsymbol{b}, \boldsymbol{A}\boldsymbol{b}, \boldsymbol{A}^2\boldsymbol{b}, \dots, \boldsymbol{A}^{\alpha-1}\boldsymbol{b}\} \quad \text{Eq.3.1}$$

This formulation avoids matrix to matrix operations and instead utilises matrix to vector operations which can be very efficiently implemented in parallel architectures. The Krylov subspace is generated while the solver seeks to find the minimum of the projection space. Usually, least squares or gradient-based optimization techniques are employed to solve such problems. There are many proposed algorithms to implement a Krylov space solver but there is no clear conclusion on which one is fastest when the same termination criterion is required[100]. The most popular approaches for the Krylov space gradient optimization is the Conjugate Gradient (CG), but is not guaranteed to work in non-Hermitian linear systems[101]. However, there are appropriate Krylov subspace solvers that can handle non-Hermitian systems with relatively low additional computational cost, like the Biconjugate Gradient Stabilised (BiCGStab).

3.1.2.4 Preconditioners

Iterative solvers do not have robust performance and can be very slow when the condition number of the system is very large. To overcome this, preconditioned versions of the solvers have been developed. Efficient preconditioning can largely reduce the condition number of a linear system leading to a dramatically reduced number of iterations to convergence. The preconditioner **P**, in effect, is changing the geometry of the Krylov subspace to a simpler one, making the solution of the system much easier by providing an approximation of the matrix inverse that is easy to compute and solve. Instead of trying to minimise $\|\mathbf{b} - \mathbf{A}\mathbf{x}_n\|$, the expression to minimise becomes $\|\mathbf{P}^{-1}\mathbf{A}\mathbf{x}_n - \mathbf{P}^{-1}\mathbf{b}\|$, to be effective, the preconditioner **P** must be of much lower condition than **A**. In general, the $P^{-1}A$ product should be as close as possible to identity matrix or in other words $P^{-1} \approx A^{-1}$. It is hard to theorise what consists a good preconditioner, the main diagonal of **A**, also known as Jacobi preconditioner, can be very effective in diagonally dominant systems, however, usually an incomplete factorization of **A** is used; as incomplete LU or incomplete Cholesky (IC) factorization. Recent research on solving linear systems focuses mainly on the choice of efficient preconditioners, emphasizing preconditioners that can be implemented in parallel architectures[102], rather than improving the solvers themselves. This happens because the time within each iteration is greatly reduced due to the high level of parallelism offered by GPUs[103], whilst preconditioning reduces the number of iterations needed to converge. Nevertheless, in practice, choosing the best preconditioner is usually a trial-and-error procedure. There are block-based preconditioners that are favourable for GPU parallelization[104,105]. In practice, the IC with no prior permutation or pivoting scheme was found to be the best preconditioning option for fast convergence of the MATLAB

based iterative solvers, whilst the Factorised Sparse Approximate Inverse (FSAI)[106] was found to be the option that produced the fastest overall result with the CUDA and OpenMP implementations. FSAI is constructed by solving local linear systems for each column of $\mathbf{A}$ to approximate an $\mathbf{A}^{-1}$ with sparsity pattern defined by powers of $\mathbf{A}$. Additionally, a preconditioner inspired by FSAI was implemented, where the local linear systems were solved in parallel and only for the three larger values for each column, achieving similar preconditioning effectiveness whilst reducing the computational time for the construction of the preconditioner. This preconditioner is referred to as "FSAIP" for the rest of this work.

3.1.2.5 Numerical accuracy

The iteration residual r_n, and to an extent, the realization of the termination criterion t_c, is bound to the numerical binary representation precision of numbers that the machine, the programming language, and the employed libraries allow. In modern systems, this is double precision, represented by 64 bits of memory, which in practice can represent numbers with relative differences no smaller than 2^{-52}, this is $\sim 2.22 \times 10^{-16}$, which is the minimum value defined in MATLAB. Any difference smaller than this is lost due to the quantization involved in converting a number that belongs to the real set $\mathbb{R}$ into the binary set $\mathbb{B}_2^{64}$, where $\mathbb{B}_2 = \{0,1\}$. Therefore, requesting termination residual lower than a scale of 10^{-16}, will not result in a more accurate solution, since any additional variation would be under the double-precision quantization bin size of MATLAB and will be rounded to the nearest bin. Apart from the binary rounding errors, when solving a linear system with an iterative solver, the

maximum solving precision that can be achieved is analogous to the condition number of the system.

3.1.2.6 Complex numbers support

The existing open-source libraries that provide low-level functions and primitive data structures for parallel programming support on CPU and GPU are the Open Multi-Processing language (OpenMP)[107] and Compute Unified Device Architecture (CUDA)[108]. However, up to the current version, OpenMP 4.5, does not provide native complex numbers support; and CUDA, whilst it provides support for complex numbers, does not come with high-level mathematical functions such as sparse iterative linear solvers and preconditioners, therefore implementations that allow complex support are not a trivial task. Additionally, open source mathematical libraries that provide iterative solving of sparse linear systems on parallel architectures, as PARALUTION[109] and ViennaCL[110], do not provide complex numbers support. However, when formulating the forward problem for systems operating in the frequency domain, the resulting linear system consists of complex numbers. Nevertheless, there are algebraic schemes that allow a linear system of complex numbers to be represented as an equivalent system of real numbers, solved in the real number domain, and then the solution can be converted back to a complex representation. There are four possible formulations of equivalent real systems as described in[111], the approach chosen for this work is the K1 approach, that is formulated as:

$$\mathbf{A}_\mathrm{c} = (x + yi) \leftrightarrow \begin{pmatrix} x & -y \\ y & x \end{pmatrix} = \mathbf{A}_\mathrm{r} \qquad \text{Eq.3.2}$$

where $\mathbf{A}_\mathrm{c}$ is the complex form and $\mathbf{A}_\mathrm{r}$ the equivalent real representation generalising, the n^{th} dimensional complex linear system $\mathbf{A}_\mathrm{c}\mathbf{x}_\mathrm{c} = \mathbf{b}_\mathrm{c}$ with entries:

$$\begin{pmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,n} \\ a_{2,1} & a_{2,2} & \dots & a_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n,1} & a_{n,2} & \cdots & a_{n,n} \end{pmatrix} \times \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix} \quad \text{Eq.3.3}$$

is equivalent to the real linear system $\mathbf{A}_\mathrm{r}\mathbf{x}_\mathrm{r} = \mathbf{b}_\mathrm{r}$ with entries:

$$\begin{pmatrix} \Re a_{1,1} & -\Im a_{1,1} & \Re a_{1,2} & -\Im a_{1,2} & \cdots & \Re a_{1,n} & -\Im a_{1,n} \\ \Im a_{1,1} & \Re a_{1,1} & \Im a_{1,2} & \Re a_{1,2} & \cdots & \Im a_{1,n} & \Re a_{1,n} \\ \Re a_{2,1} & -\Im a_{2,1} & \Re a_{2,2} & -\Im a_{2,2} & \dots & \Re a_{2,n} & -\Im a_{2,n} \\ \Im a_{2,1} & \Re a_{2,1} & \Im a_{2,2} & \Re a_{2,2} & \cdots & \Im a_{2,n} & \Re a_{2,n} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \Re a_{n,1} & -\Im a_{n,1} & \Re a_{n,2} & -\Im a_{n,2} & \cdots & \Re a_{n,n} & -\Im a_{n,n} \\ \Im a_{n,1} & \Re a_{n,1} & \Im a_{n,2} & \Re a_{n,2} & \cdots & \Im a_{n,n} & \Re a_{n,n} \end{pmatrix} \times \begin{pmatrix} \Re x_1 \\ \Im x_1 \\ \Re x_2 \\ \Im x_2 \\ \vdots \\ \Re x_n \\ \Im x_n \end{pmatrix}$$

$$= \begin{pmatrix} \Re b_1 \\ \Im b_1 \\ \Re b_2 \\ \Im b_2 \\ \vdots \\ \Re b_n \\ \Im b_n \end{pmatrix} \quad \text{Eq.3.4}$$

The equivalent real system has the same sparsity pattern and sparsity factor as the original complex system, however the new system has double the number of unknowns, therefore requires double the computations, and additionally, FEM-DA linear systems in FD are no longer Hermitian, therefore the BiCGStab solver is employed for FD simulations.

In addition, a BiCGStab solver, based in CUDA, operating directly on the complex domain was implemented and used with the developed parallel constrained FSAI version (FSAIP) to solve FD simulations.

3.1.2.7 GPU/CPU parallelization

The proposed approach for accelerating fluence estimation relies on employing efficient libraries for linear algebra operations and performs remarkably faster when

GPU based parallel architectures are available. Over the last decade, the technical advancements in GPUs, and their relatively low cost, has made GPU computing a very attractive option. Specifically, many linear algebra operations can be parallelised very efficiently in GPU architectures[112] while using sparse representations, resulting in massive reductions of computational time. This can be applied to the solution of the forward model, dramatically decreasing the computational time required to estimate the Krylov subspace. Solvers based on libraries that can be used both in CPU, using OpenMP, and GPU, using CUDA, were implemented to guarantee accessibility by all users.

The mathematical library employed to provide efficient implementations of high-level linear algebra operations is PARALUTION[109] which offers a wide variety of linear solvers and preconditioners, supports sparse matrix and vector formats, and allows a high level of abstraction between code and hardware, making the code highly portable and efficiently scalable to the available hardware. The produced CUDA based implementations will retreat to OpenMP if there is no GPU available in the system.

An algorithm to distribute workload between the CPU and GPU was implemented, the workload was distributed by balancing the right-hand side input (sources) between CPU and GPU. Benchmarking tests were performed on all mesh resolutions to define the best workload distribution in each case. However, it was found that in all meshes above 70,000 nodes the solely GPU based solution was faster, whilst with meshes with smaller number of nodes (~50,000), the computational time reduction was less than a second. On the other end, the CPU implementation is faster than the equivalent GPU implementation in meshes with less than 15,000 nodes. This is primarily due to time-

consuming data transfer and device initialisation procedures. Nevertheless, this is depended on hardware, number of right-hand side vectors and complexity of the imaging domain.

MATLAB provides sparse linear solvers on the CPU, that can be easily parallelised over the right-hand side vectors using the parallel computing toolbox. Although, there are overhead data transfers between memories (RAM, CPU cache memory, GPU memory) and between computational threads and memory that do not allow computational accelerations to scale linearly with the number of available computational cores.

3.1.2.8 Experimental set-up

Data from MRI of an adult head was segmented and meshed into thirteen different resolution meshes using the algorithm proposed by Jermyn et al[97], which produces high quality meshes using the Stellar mesh improvement algorithm[113], where volume distribution and the dihedral angles formed between each pair of faces are particularly important, effectively removing flat elements. The modelled DOT instrument is a high-density system with 158 NIR light sources and 166 detectors. Each detector is related to sources in separation distance configurations from 1.3 to 4.6cm, resulting in 3500 associated source-detector pairs. More details about the resolution of the meshes can be found in Table 3.1 and the optical properties for each layer of the anatomical model are described in Table 3.2[114].

Table 3.1 Different resolution meshes based on linear tetrahedral elements for an adult head model.

Number of nodes	Number of elements	Element volume average and standard deviation (mm^3)
50721	287547	9.26 ± 3.43
68481	393863	6.76 ± 2.29
101046	589658	4.51 ± 1.64
139845	821926	3.24 ± 1.18
205568	1215434	2.19 ± 0.78
235564	1395242	1.90 ± 0.68
271135	1609152	1.65 ± 0.59
305058	1813036	1.46 ± 0.52
324756	1931374	1.37 ± 0.49
360777	2149250	1.23 ± 0.43
411567	2454350	1.08 ± 0.37
515837	3084689	0.86 ± 0.29
610461	3656890	0.72 ± 0.24

Table 3.2 Optical properties of tissue layers at 750nm wavelength[114].

Tissue Layer	μ_a (mm^{-1})	μ_s' (mm^{-1})	Refractive Index
Scalp	0.0170	0.74	1.33
Skull	0.0116	0.94	1.33
Cerebrospinal fluid	0.004	0.3	1.33
Gray Matter	0.0180	0.84	1.33
White Matter	0.0167	1.19	1.33

The light propagation model was calculated for all 158 sources in all experiments, in continuous wave mode and in frequency domain mode at a modulation frequency of 100 MHz, for one NIR wavelength of 750 nm. All the experiments were performed on a desktop computer with 16GB of RAM, an Intel Core I7-4790 CPU with 4 physical cores, allowing two threads per core, resulting in 8 logical cores @ 3.6GHz, and an NVIDIA

GTX970 graphics card with 1664 logical cores @ 1050MHz with 4GB dedicated memory.

3.1.2.9 Metrics

It is important to ensure that employing an iterative linear solver will not increase the error of the solution. To this end, the accuracy of the proposed solvers was compared against the direct solution, calculated with the backslash operator in MATLAB. There is no standard way of comparing two matrices, $\mathbf{\Phi}_{\text{ref}}$, for the fluence calculated with a direct solver, and $\mathbf{\Phi}_{\text{ite}}$, for the fluence calculated with an iterative solver, however, the first step for all approaches is taking the difference $\mathbf{\Phi}_{\text{dif}} = |\mathbf{\Phi}_{\text{ite}} - \mathbf{\Phi}_{\text{ref}}|$. The most common metrics to quantify the difference $\mathbf{\Phi}_{\text{dif}}$ are the ones induced from vector norms: the 1-norm $\|\mathbf{\Phi}_{\text{dif}}\|_1$ which is the maximum of the column sums of $\mathbf{\Phi}_{\text{dif}}$, the ∞-norm $\|\mathbf{\Phi}_{\text{dif}}\|_\infty$ which is the maximum of the row sums of $\mathbf{\Phi}_{\text{dif}}$, and the l2-norm $\|\mathbf{\Phi}_{\text{dif}}\|_2$ which is the maximum singular value of $\mathbf{\Phi}_{\text{dif}}$, also known as the spectral norm.

However, those metrics do not provide easily comprehensible quantities, therefore the relative error per node, r, was calculated as:

$$\varepsilon_{rel}(r) = \frac{\left|\boldsymbol{\Phi}(r)_{ref} - \boldsymbol{\Phi}(r)_{ite}\right|}{\left|\boldsymbol{\Phi}(r)_{ref}\right|} \times 100\% \qquad \text{Eq.3.5}$$

This relative error representation is useful for visualisation of the error on the mesh nodes and boundary data and provides more comprehensible numbers than the matrix norms.

3.1.3 Results

The evaluation is performed in one adult head model using a HD-DOT system with 158 sources and 166 detectors (Fig. 1). The behaviour of the solvers is examined under

varying error demands, and in different mesh resolutions, considering the accuracy and the computational time. The focus is on the relation between mesh resolution (and hence problem size), termination criterion, computational time, and solution error. The direct solutions are only possible to calculate up to the 400,000 nodes mesh for continuous wave (CW) and up to 200,000 nodes in frequency domain (FD) systems, due to high memory requirements, so all quantitative comparisons are performed in the subset of the meshes where a direct solution is available.

3.1.3.1 Qualitative and quantitative comparisons

Assuming that the diffusion approximation is valid for the proposed application, and that the quality of the mesh is appropriate, then the error from solving the linear systems should be kept at a minimum, which purely defines the numerical accuracy of the computational model. It is important to note that what is being presented can be applied to any other numerical models, including those based on the Radiative Transport Equation. However, the amount of error that can be afforded in the modelling procedure is dependent on the error tolerance for the application. Figure 3.2 shows the surface fluence when utilizing the CUDA based solver at different termination criteria; the simulated light source is near the back of the head, indicated by the blue dot and arrow.

When high termination criterion is set, the fluence is not estimated for the distant nodes as the solver iterates to a stable solution quickly. The fluence approximately follows an exponential decay through tissue, therefore its value dramatically decreases

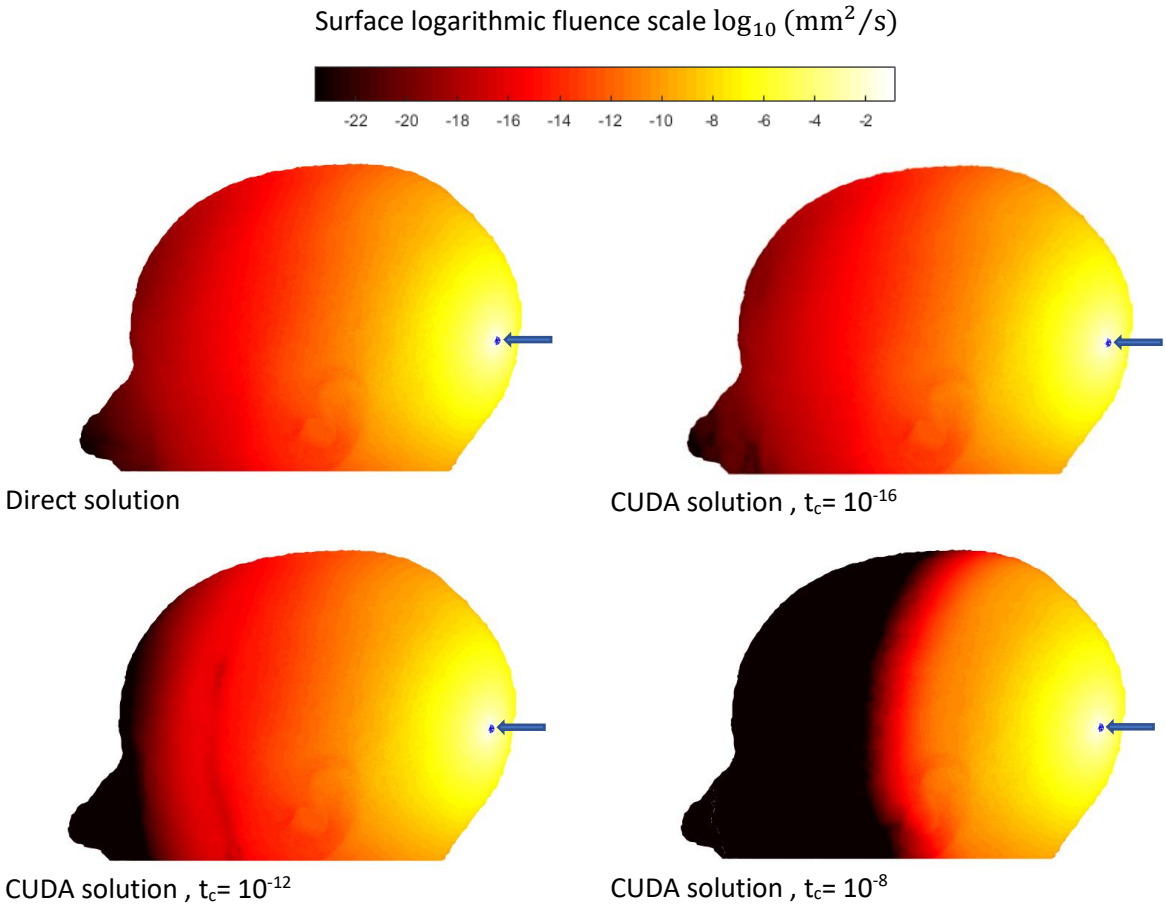

Figure 3.2 Visual comparison of surface fluence whilst using different termination criteria (t_c). Simulation in continuous wave, for one source indicated in blue, in a 400,000 nodes mesh, solving with CUDA CG.

with distance from the source, therefore the required termination criterion is reached while only calculating the solution for the highest fluence values.

As FD simulations provide amplitude and phase information, the errors for each were examined separately. Fig 3.3 and 3.4 provide a quantification of the relationship between distance from the source and the relative nodal amplitude and phase errors arising from solving with an iterative solver. The demonstrated simulation is in the frequency domain, at 100MHz, for a mesh with 200,000 nodes, solving with a CUDA BiCGStab and FSAI. The maximum relative error for all nodes as a function of distance from the source is extracted for different termination criteria.

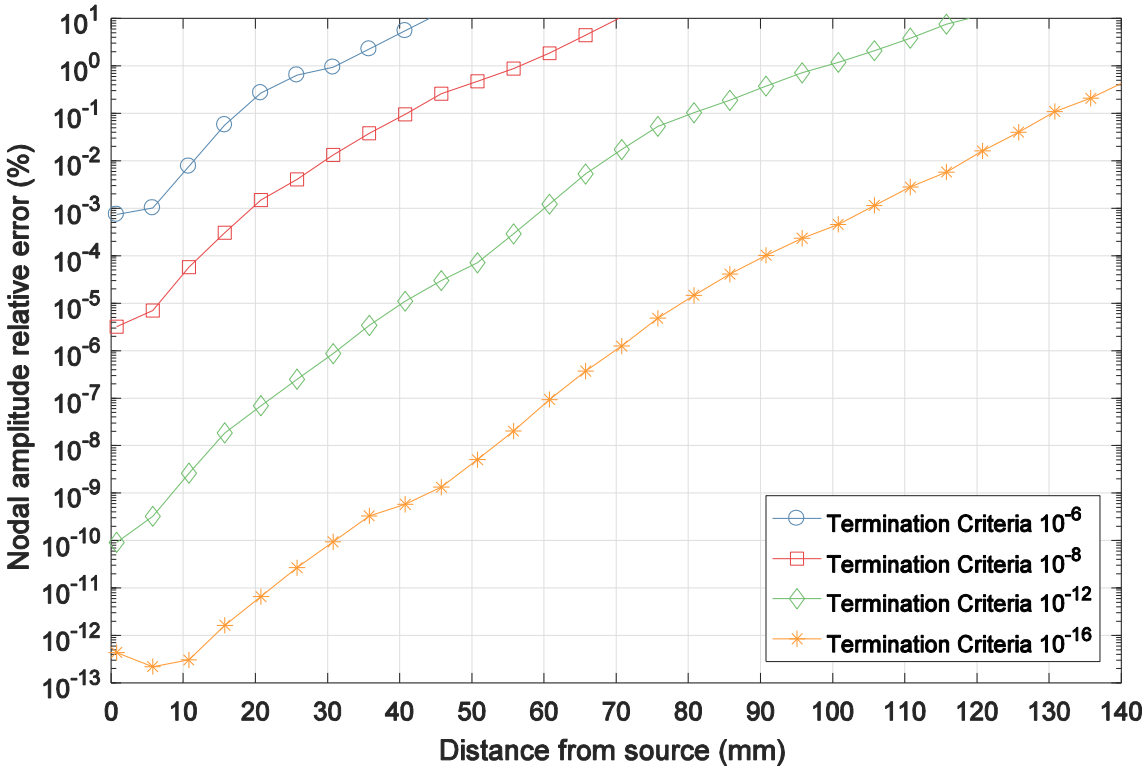

Figure 3.3 Maximum relative amplitude errors per node (Eq.3.5) as a function of distance from the source. Comparison between different termination criteria. Simulation 100MHz frequency, on a 200,000 nodes mesh, solving with CUDA BiCGStab.

As evident, the relative errors are small and located away from the source with low termination criterion of 10^{-16}, but they become larger and manifest nearer the source as the termination criterion rises. Similar results (not shown) were acquired for amplitude errors from continuous-wave simulations. For the lowest termination criterion, the

relative error reduces initially at 5 mm, which could be caused by the limit of the numerical accuracy for this very low termination criterion.

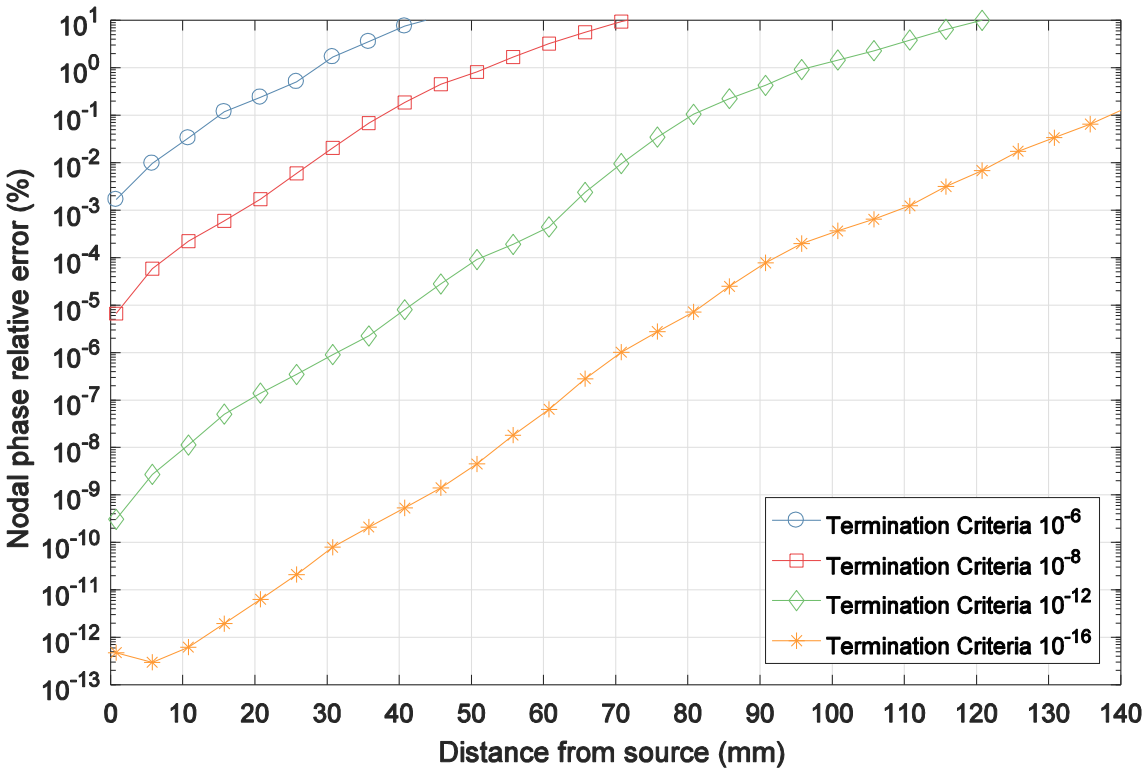

Figure 3.4 Maximum relative phase errors per node (Eq.3.5) as a function of distance from the source. Comparison between different termination criteria. Simulation 100MHz frequency, on a 200,000 nodes mesh, solving with CUDA BiCGStab.

Lower termination criterion will provide smaller numerical errors, but also slower solver convergence, as a larger number of iterations is required. In the modelled DOT system, the maximum SD separation typically considered to acquire boundary data is at 46 mm. The performed evaluation in Figs 3.3 and 3.4 reveal that for anatomically accurate adult head models, the termination criterion can be selected in the range of 10^{-8} or lower, for CW and FD systems, to ensure that minimal error is introduced in the parameter recovery, when acquiring measurements from SD separation distances less than 46 mm. The sensitivity matrix will have an error approximately square of the error of the forward solution. Therefore, termination criterion chosen to be large, a practice often employed to accelerate reconstructions, while the boundary data is measured in

large SD separations, can lead to large errors in the sensitivity matrix and, consequently, large errors in the parameter recovery and image reconstruction.

It is worth noting that since the forward problem is not ill-posed, the approximate solution will always converge towards the true solution, with the termination criterion dictating the maximum numerical distance between the approximation and the true solution.

3.1.3.2 Computational Time Comparisons

The parameters that mainly affect the computational speed of iterative linear solvers are: the size of the problem, which in our case is the number of nodes, the number of right-hand side vectors, which in our case is the number of excitation sources, the termination criteria, and finally the preferred algorithm and implementation. These were the variables for the experiments, while the computing engine remained the same. All the experiments were performed 10 times, the mean time is shown in all figures, while the standard deviation in all cases was small, at around 1 second for CPU implementations and 0.1 seconds for GPU, therefore it is not shown in the figures.

Figure 3.5 shows the computational time for fluence estimation for 158 sources in a 400,000 node mesh as a function of the termination threshold. The direct solver provides an exact solution to the linear system, therefore does not introduce any error. However, is displayed as a horizontal line through all the termination criteria in Figure 3.5, to serve as a point of reference. The GPU based solver yields the best termination criterion to computational time ratio. Employing implementations that do not require much additional time to converge to smaller errors can increase the accuracy of the estimated light propagation model whilst keeping the computational time low.

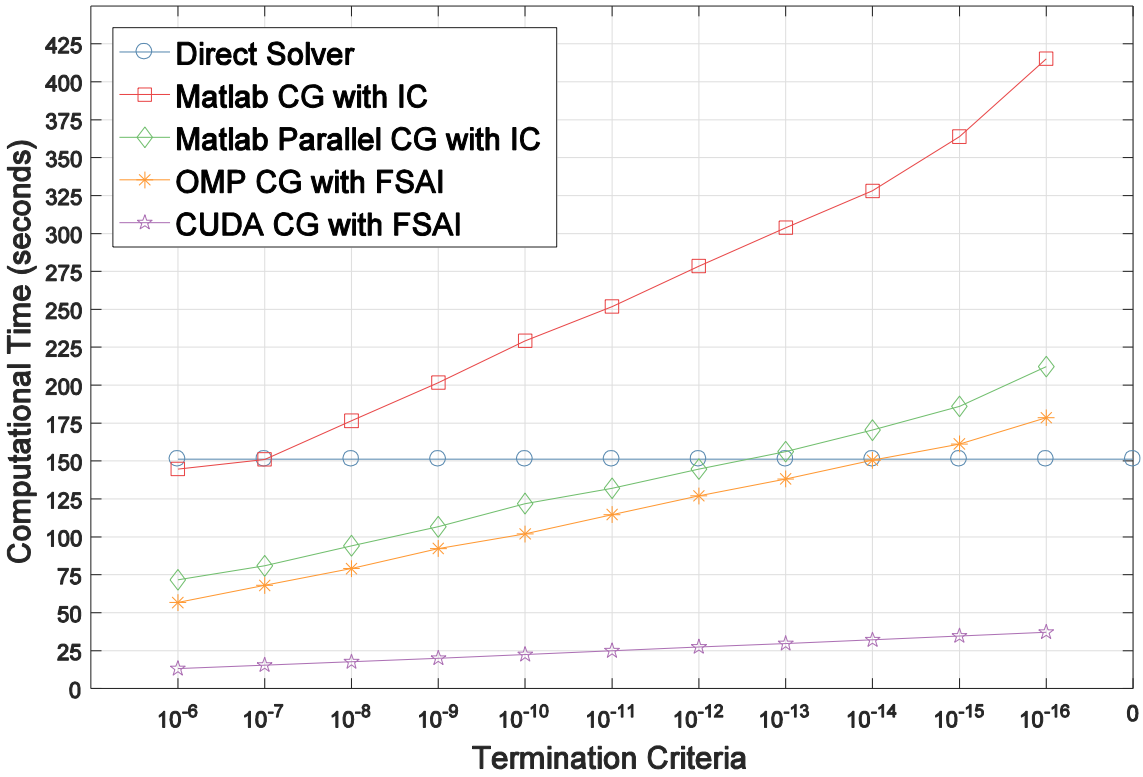

Figure 3.5 Computational time as a function of termination criteria. Comparison between different linear solvers. Simulation in continuous wave, for 158 sources in a 400,000 nodes mesh.

Each source is represented by one right-hand side vector in the linear system resulting from the FEM , and the fluence must be calculated for all sources. To achieve this, the iterative solvers must create the Krylov space under the projections of each right-hand side vector, which, as expected increases the computational cost and therefore the computational time required. Figure 3.6 demonstrates how the number of sources (right-hand side vectors) affects the computational time of the solution, showing that the computational time increases linearly with the number of sources. It is interesting to note that the direct solver, that yields the exact solution relying on Cholesky decomposition followed by forward and backward substitutions, is almost as efficient for each additional source as the GPU based solver. However, the time spent initially for the

factorisation is very large, which in combination with the very high memory requirements as discussed in section 2.3.1, render the direct solver impractical.

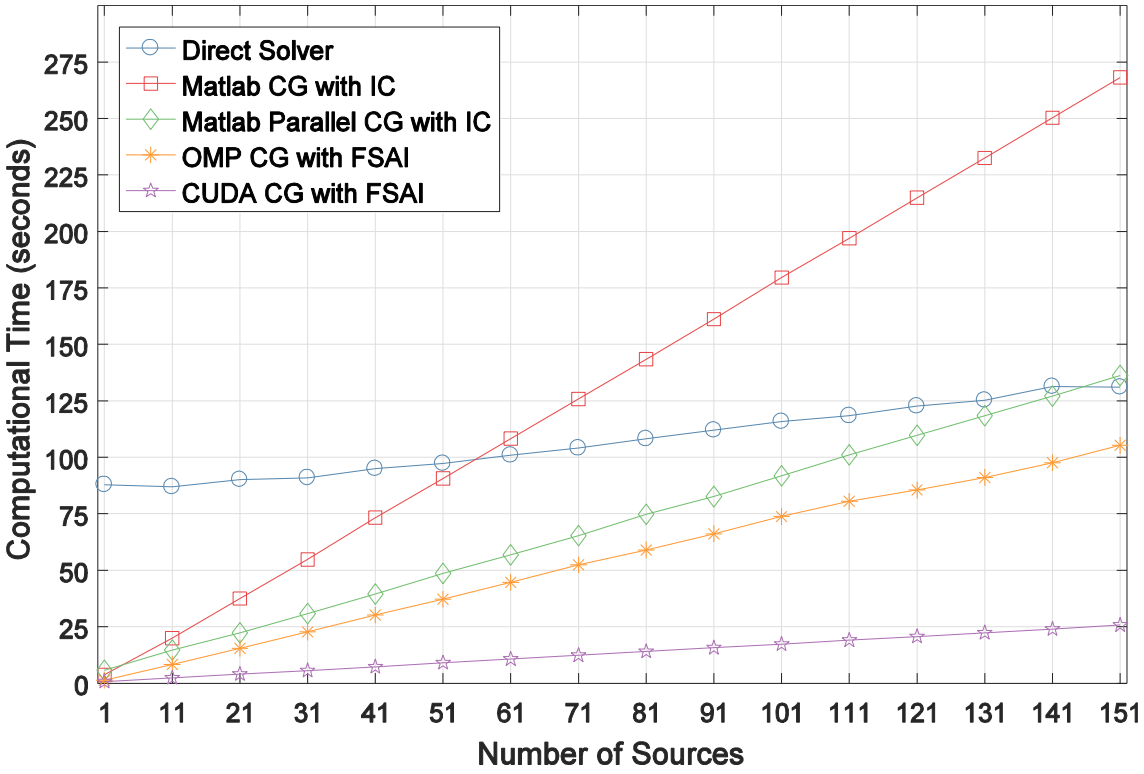

Figure 3.6 Computational time as a function of excitation sources number. Comparison between different linear solvers. Simulation in continuous wave, in a 400,000 nodes mesh, with 10^{-12} termination criterion.

Nevertheless, the factor that affects the computational time the most is the resolution of the mesh. The more nodes a mesh contains, the bigger the linear systems that need to be solved is, therefore more mathematical operations have to be applied to create the Krylov space. Figure 3.7 presents the computational time needed as the mesh resolution increases. The fastest of the solvers is the CUDA based solver, which achieves a computational time of ~42 seconds for calculating the fluence for all 158 excitation sources in a 600,000 node mesh, this is ~0.25 second to calculate the fluence for one source. The CUDA based solver performs almost 11 times faster than the MATLAB based iterative solver without any parallelization, which takes ~460 seconds for the

same calculation. The direct solver can only solve up to systems with 500,000 nodes before the 16GB hardware memory availability becomes an underlying issue.

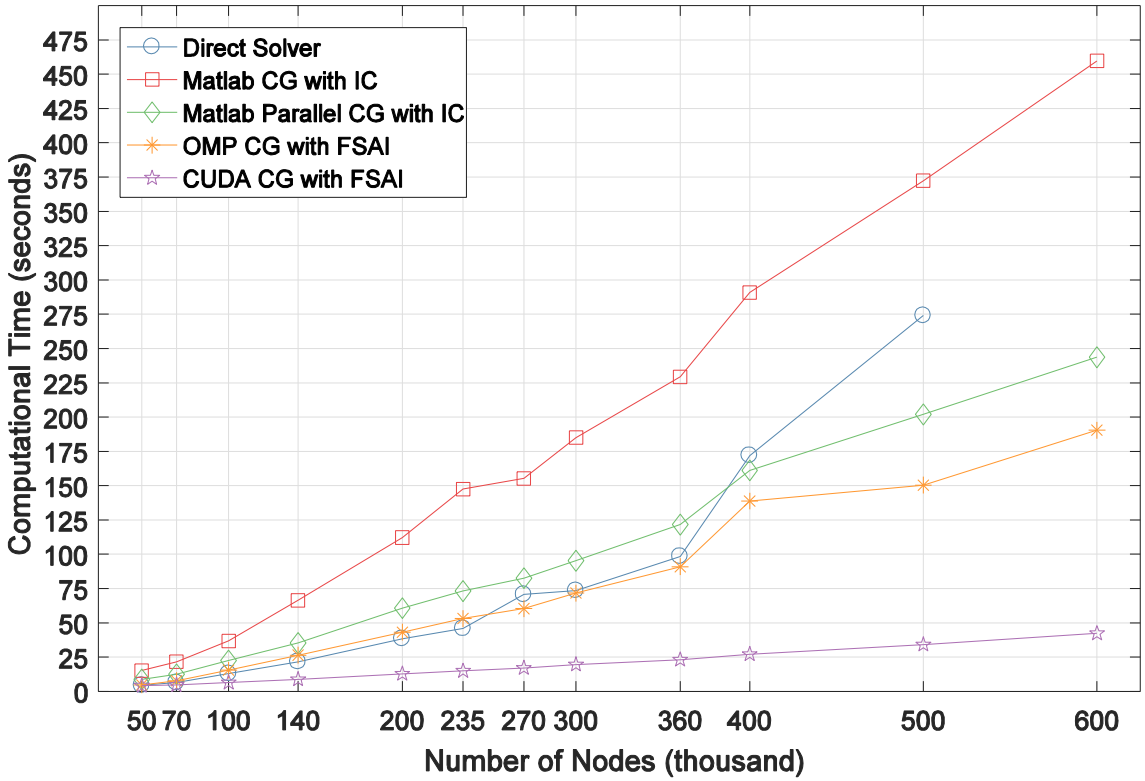

Figure 3.7 Computational time with respect to mesh resolution. Comparison between different linear solvers. Simulation in continuous wave, for 158 sources with 10^{-12} termination criterion.

Figure 3.8 displays the computational time for different mesh resolutions for frequency-domain simulations at a modulation frequency of 100 MHz. The direct solver can only handle up to 200,000 nodes, due to the increased memory requirements for storing complex numbers. The displayed computational time includes the computations to create the equivalent real system and transform the fluence back to complex after solving the system where is necessary.

The direct solver takes 4,612 seconds to calculate the fluence for the 200,000 nodes mesh, however Figure 3.8 was limited to 1,100 seconds to provide a better scale. The direct solver becomes intractable due to the increased memory requirements for complex arithmetic storage and because of the non-Hermitian nature of the FEM matrix, which is also reflected as increased memory and computational requirements for the required LU decomposition (in comparison with the Cholesky for the real cases). A linear system

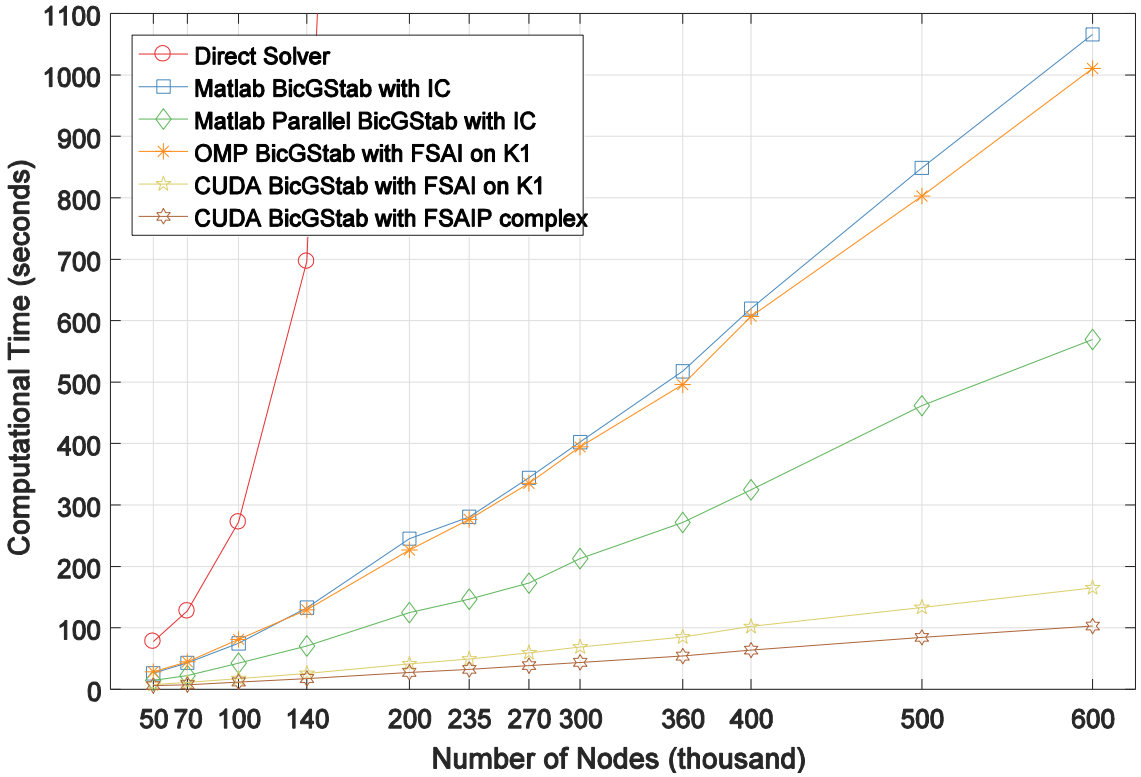

Figure 3.8 Computational time with respect to mesh resolution. Comparison between different linear solvers. Simulation in 100 MHz frequency, for 158 sources with 10-12 termination criterion.

resulting from a frequency domain FEM mesh does not have the same condition number as the same mesh in continuous wave, due to different attenuation coefficients for frequency-modulated light, which makes the FD problem harder to solve, therefore, there is not a direct analogy between their computational costs. Though, one can roughly assume that for a given mesh if a CW solution requires O operations the FD will require $2O$. This is confirmed in our demonstrated results in Figure 3.7 and Figure 3.8.

Furthermore, the “OMP BiCGStab with FSAI on K1” operates on the complex to real transformed (K1) matrix, resulting in double computations in comparison to the “MATLAB BiCGStab with IC” which operates directly on the complex domain. As an approximation, one could assume that if a mesh in CW requires O number of operations, it requires 2O in the complex domain but 4O when the complex to real transformation is used. Also, the Matlab parallel version requires almost half the computational time of the non-parallel Matlab version, and the OpenMP version is slightly faster than the Matlab parallel version when operating in the same space (O). Then it is possible to observe the following: a solution on CW would take T seconds for Matlab non-parallel, T/2 for Matlab parallel, and slightly faster than T/2 for OpenMP (note that all these cases do O operations). In contrast, a solution on FD domain would take 2T for Matlab non-parallel (operates in 2O), T for Matlab parallel (operates in 2O) and slightly faster than 2T for OpenMP (operates on 4O). Furthermore, the implemented complex CUDA version, which operates on 2O, requires approximately half the computational time in comparison with the CUDA on the K1 (4O) space.

3.1.4 Conclusion

The DOT reconstruction algorithm, especially when employed for high-density functional brain imaging, suffers from large computational time, mainly due to solving large sparse linear systems as required for the forward model. In this work the accuracy and computational speed of iterative solvers is evaluated on anatomically realistic head models, with high-resolution meshes for a high density (HD) DOT system. This work provides both a computationally efficient GPU and CPU implementations of stable

linear matrix solvers, based on CUDA and OpenMP, compiled as MATLAB compatible mex files.

The developed solvers allow users to explore new approaches in DOT, that until now have been challenging due to computational inefficiency of the algorithm, namely whole-head high-density tomographic imaging. Simulations of light propagation in tissue that would otherwise take a long time to perform, are now shown to be achievable in a few minutes, forging a path towards real-time DOT. Although the work presented here is based on only one GPU node, the same approach can be applied in systems with multiple GPUs and extended to cloud computing to achieve real-time solutions.

3.2 Inverse Problem computational optimisation

A sensitivity matrix, also known as the Jacobian, is the basis for solving the inverse problem, providing the recovery of spatiotemporal changes of internal optical properties throughout the imaging domain, using temporal derivatives of measurements obtained on the surface of the volume, known as boundary data. However the calculation of the Jacobian is computationally expensive, even using the Adjoint method which is very computationally efficient, therefore different strategies have been proposed to expedite the inverse problem. Reconstruction on a smaller basis with adaptive meshing or coarse voxelization has been used[58,80,115], additionally, parallel computing has been employed[95].

The approach described here, achieves a hundredfold required memory reduction, on large mesh node numbers, and tenfold computational time reduction whilst producing sparse Jacobians where the non-zero elements have identical values with the corresponding elements of the full Jacobian.

3.2.1 Method

The Adjoint method was employed to form the Jacobian (Eq. 2.10). Previous work has proposed the reduction of the reconstruction space by imposing a sensitivity threshold to calculate the Jacobian only in a part of the volume[116]. In more recent work[85] an approach for building sparse Jacobians was suggested, where initially the sensitivity for each SD pair is approximated, then a threshold relative to the maximum sensitivity of each SD pair is selected. Specifically, Eq. 2.10 is simplified by assuming a uniform voxel-wise sampling of the imaging volume, considering only the values of the direct, Φ, and adjoint, Φ_{Adj}, fields, rather than having to integrate over each element individually:

$$J(i,j,r) = \Phi^{i} \times \Phi_{Adj}^{j} \qquad \text{Eq.3.6}$$

In practice, a high-resolution mesh with 3,650,000 tetrahedral elements, achieves a very fine uniform sampling of the volume and with a variation of only ± 0.24 mm^3 (Table 3.1). Assuming this is variation is almost zero, the volume is then uniformly sampled and can be simplified to a multiplication of the two fluences, rather than integration over each volume. This achieves an approximation of adequate accuracy for the suggested thresholding approach, however not accurate enough to achieve image reconstruction.

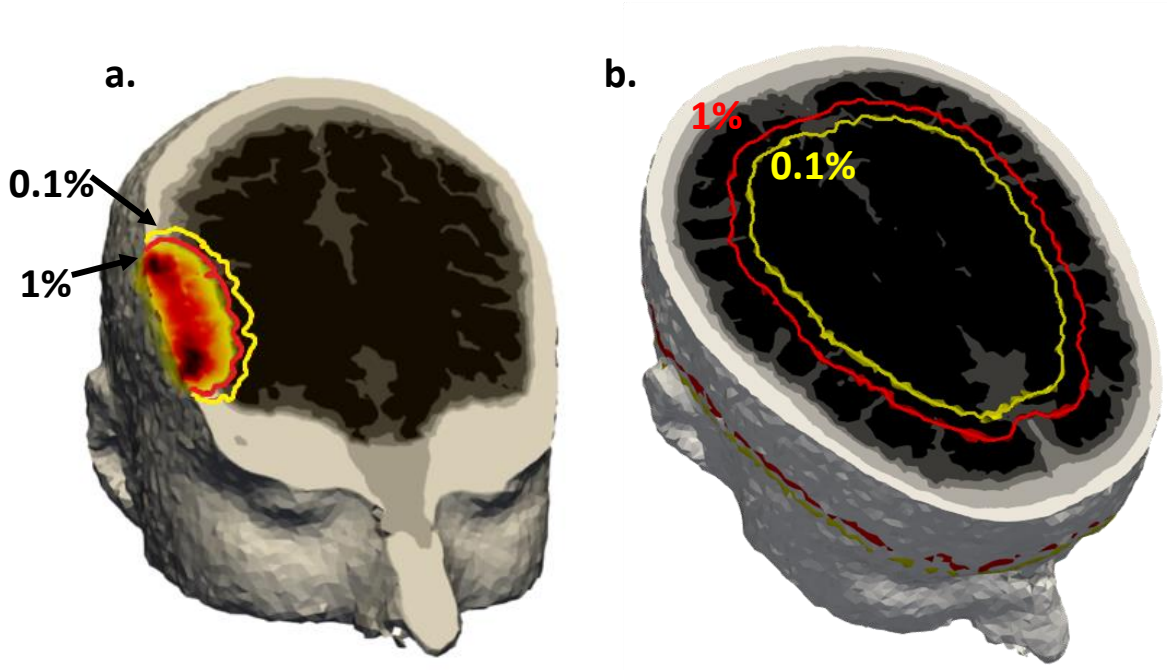

Figure 3.9 Regions defined by different thresholds. a) 0.1% (yellow) and 1% (red) sensitivity threshold contours for a single measurement, b) Representation of the area covered for 0.1% (yellow) and 1% (red) sensitivity threshold for all the measurements.

Essentially, the thresholding defines a region of interest (ROI) with high sensitivity in the modelled volume (Figure 3.9a), where the parameter recovery is meaningful; therefore, the Jacobian is calculated only for the nodes within this ROI. The calculation of the sensitivity of each measurement is performed only for a very small subset of the nodes of the mesh, resulting in a sparse Jacobian, which still samples a significant part of the volume of the head (Figure 3.9b).

3.3 Results

To quantify the improvements in computational speed and memory efficiency the proposed approach was tested on 24 subject specific meshes, across a range of mesh resolutions and sensitivity thresholds, whilst also exploring the benefits of parallelisation for the construction of the Jacobian. Lower sensitivity thresholds produce larger ROIs which, in the case of fDOT, allows deeper nodes to contribute towards the reconstruction

(Figure 3.10). As the median brain cortex depth is set at 12.5cm from the surface of the head, if we assume that all measurements have a chance to measure brain activity, we need to consider a threshold that allows nodes at least up to this depth to be considered for reconstruction.

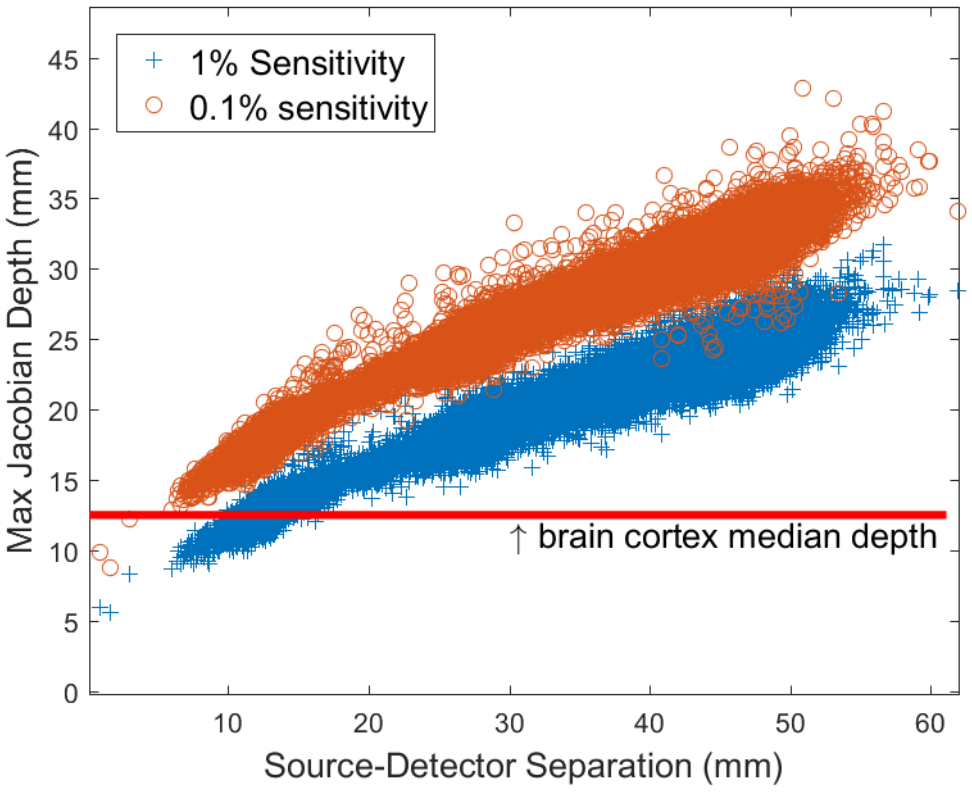

Figure 3.10 Maximum tissue sampling depth for all source-detector pairs and two sensitivity thresholds, for all 24 subjects. The red line indicates the brain cortex median depth.

3.3.1 Computational Comparison

Considering the computational efficiency of the DOT algorithm the sparsity percentage of the Jacobian affects the number of calculations required for creating the Jacobian, but also the number of operations required for image reconstruction which involves regularising and solving the linear system described by the Jacobian. Sparse matrix computations are extremely efficient because only the non-zero entries are stored in memory and later considered for arithmetic operations. Therefore, high sparsity

accelerates the computations during multiple steps of the reconstruction and, furthermore, preserves space in memory, which is often a bottleneck for DOT algorithms.

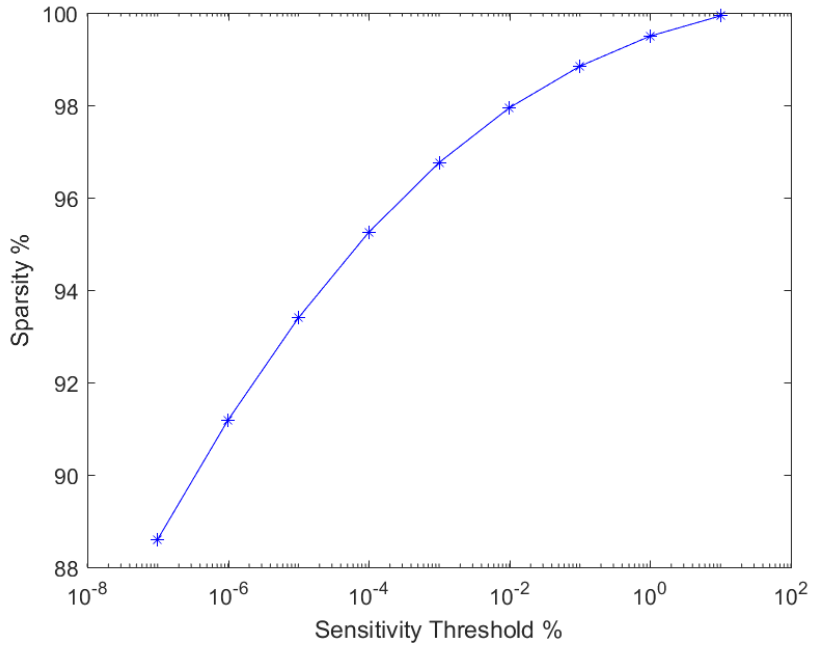

Figure 3.11 Sparsity percentage over different sensitivity thresholds. The 1% and 0.1% thresholds achieve sparsity over 98%.

The sparsity percentage of the Jacobian is directly related to the sensitivity threshold, the size of the volume sampled, and the average source-detector separation, as more distant measurements sample larger part of the volume. However, it is not affected by the number of measurements or the number of nodes of the FEM mesh. Since when selecting a safe threshold, as the 0.1%, there is no loss of reconstruction accuracy, this approach can be incorporated in the standard pipeline of all DOT applications. However, it is more beneficial when employed for imaging of large volumes, where usually reflectance measurements are considered.

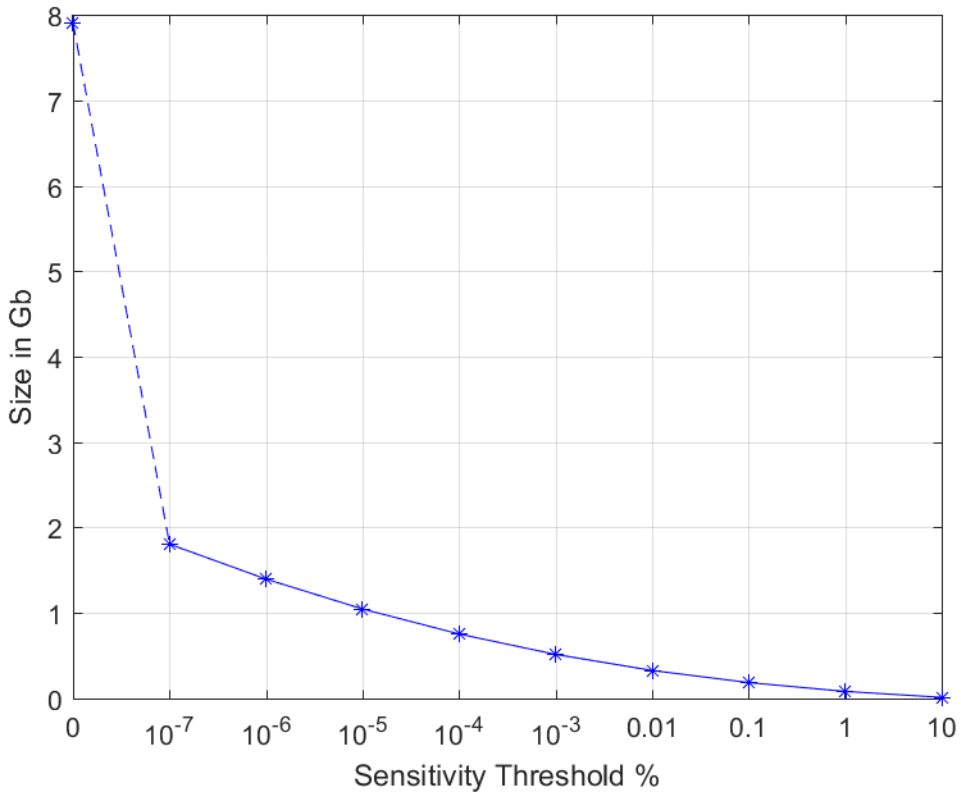

Figure 3.12 Size of the Jacobian in memory (Gb) over different sensitivity thresholds. Estimated for a 300,000 nodes mesh with 3478 measurements.

For a head mesh with 300,000 nodes and 3478 measurements the calculation of the full Jacobian results in a dense matrix of approximately 8Gb size in memory; whilst the 1% threshold results in 0.11 Gb and the 0.1% threshold in a 0.27 Gb (Figure 3.12), a memory improvement of 70 and 30 times, respectively.

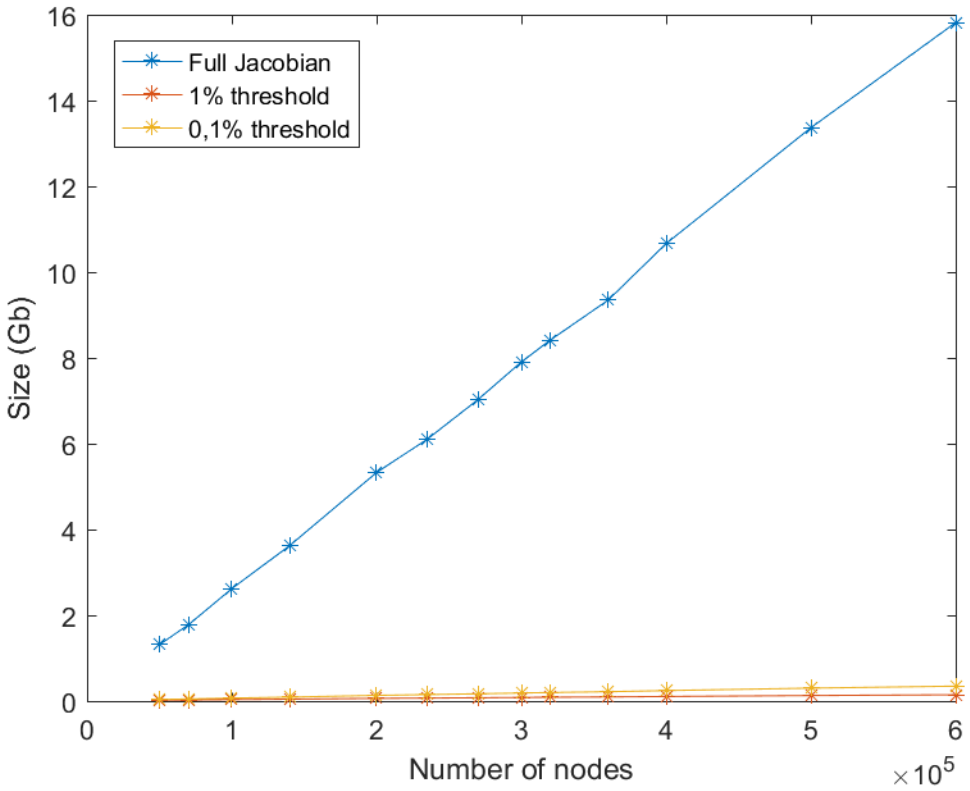

Figure 3.13 Memory size of the full (blue line), 1% (red line) and 0.1% (yellow line) Jacobian over different mesh resolutions.

While the sparsity of a thresholded Jacobian is not affected by the number of nodes of the mesh, the memory size of the Jacobian is direly connected to the number of nodes. Therefore, meshes with a larger number of nodes receive a larger benefit of memory efficiency when applying a sensitivity threshold. A mesh with 600,000 nodes results in a full Jacobian of approximately 16 Gb, while the 1% threshold produces 0.144 Gb and the 1% threshold 0.345 Gb, an improvement of 110-fold and 50-fold respectively (Figure 3.13).

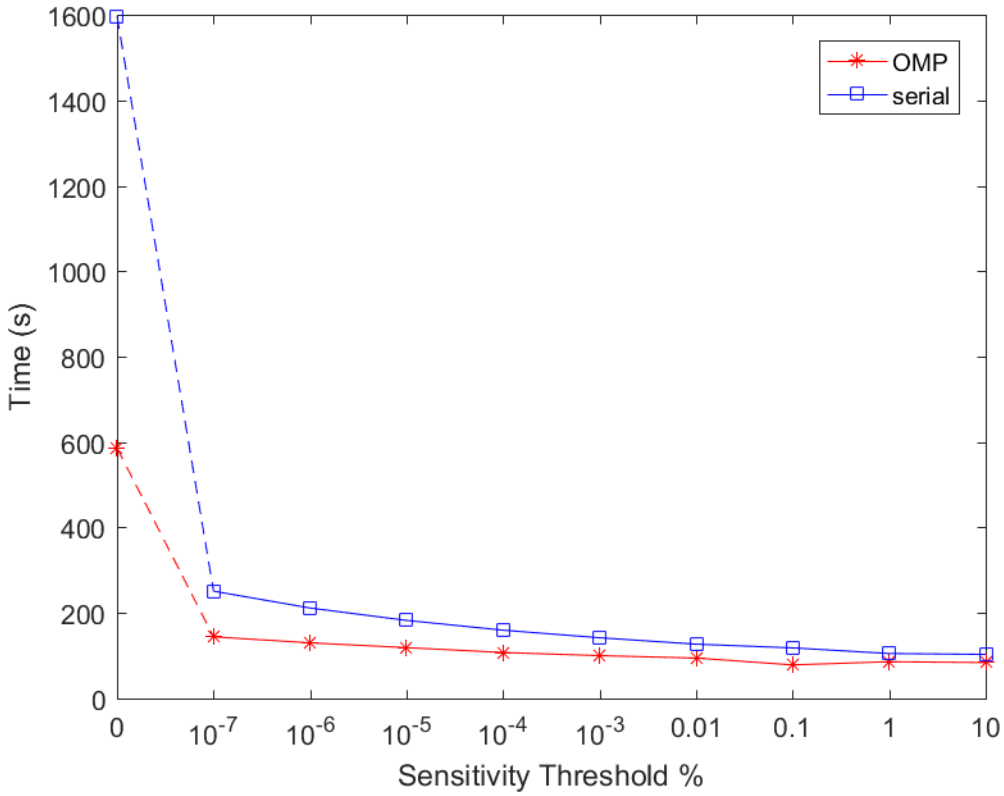

Figure 3.14 Computational time over different sensitivity thresholds. Estimated a for a 300,000 nodes mesh with 3478 measurements. Serial implementation in blue, parallel with OpenMP in red.

The sparsity of the Jacobian is the leading factor for the computational speed improvements. However, a parallel implementation using OpenMP was developed in order to further accelerate the construction of the sensitivity matrix. The maximum improvement of the parallel implementation is observed when calculating the full Jacobian, where the procedure is accelerated by 2.7 times for a 300,000 nodes mesh with 3478 measurements (Figure 3.14). The computational speed is improved by 34% using the parallel implementation over the serial, for a sensitivity threshold of 0.1%, however, this improvement is larger for meshes with a larger number of nodes (Figure 3.15).

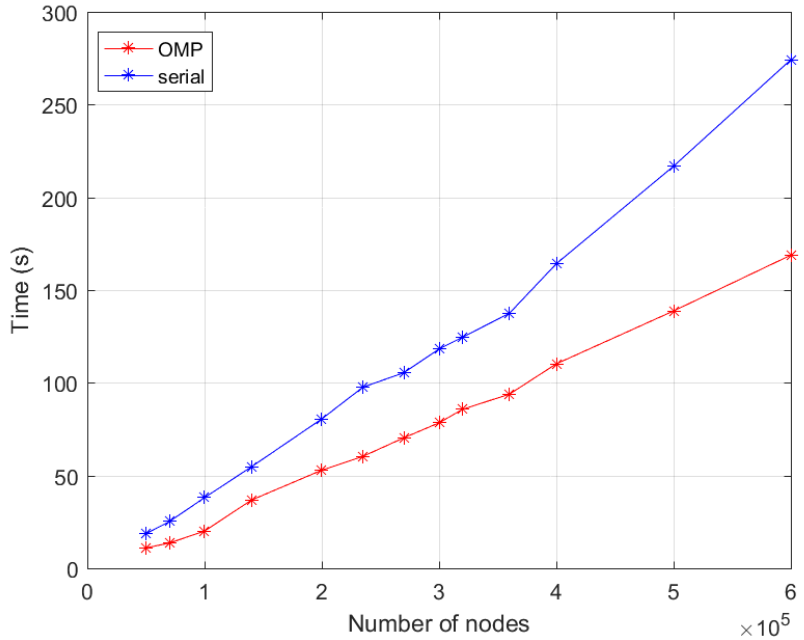

Figure 3.15 Computational time over different mesh resolutions. For 3478 measurements and 0.1% sensitivity threshold. Serial implementation in blue, parallel with OpenMP in red.

3.3.2 Qualitative Comparison

To qualitatively evaluate the proposed framework, recoveries of local activations with difference imaging were performed in 24 subject-specific models, rigidly registered onto a head atlas[117], to ensure the activations were always at the same anatomical location. An activation on the brain cortex, with a radius of 7mm, was simulated by perturbating the absorption values around the brain functional centres in Broca, motori and visual area (Figure 3.16). Boundary data were generated and noise was added for each measurement neighbourhood: 0.12%, 0.15%, 0.41%, and 1.42% for first to fourth neighbourhood respectively according to[118]. Also, the inverse problem was considered linear, due to the relatively small changes between the temporal measurements, and the employed inversion strategy was the generalised Moore–Penrose pseudoinverse, with Tikhonov and spatial variant regularizations as described in[118]. Finally, reconstruction was performed for 10 different noise instances for each subject, allowing the production of Z-maps of statistical significance.

The creation of a sparse Jacobian does not negatively affect the reconstruction quality. Using difference imaging (see section 4.1.2 for more details), the z-maps recovered with the full Jacobian, the sparse 0.1%, and 1% sensitivity threshold Jacobians have visually no difference and the spatial overlap between their recoveries for z>0.01 is on average 96% with a standard deviation of ±2.3%.

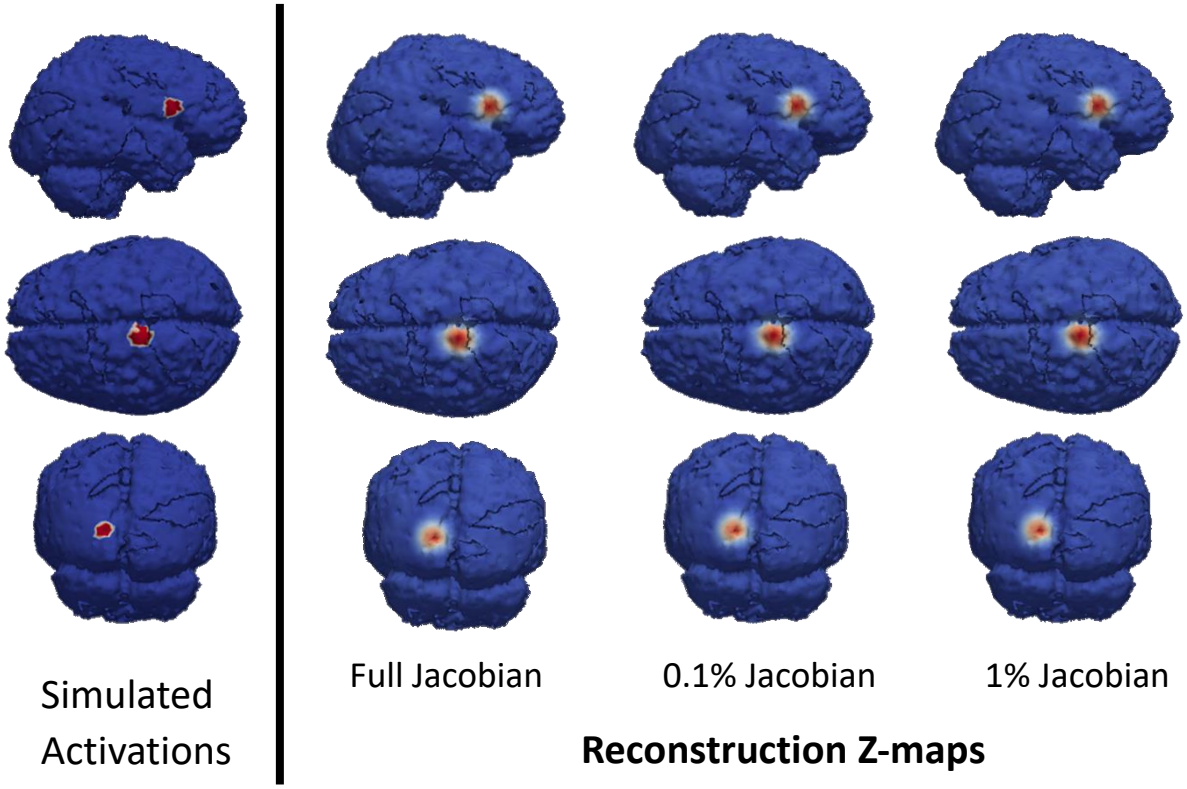

Figure 3.16 Simulated brain activations (left) and Z-maps of recoveries with full Jacobian, 0.1% and 1% sensitivity threshold Jacobian.

As discussed later in chapter 5, measurements of phase have larger depth penetration, while the algorithm presented here does not consider phase data. Therefore, considering the maximum depth reached for each measurement (Figure 3.10) the 0.1% sensitivity threshold was chosen for use in the rest of this work as the computational efficiency is sufficient to produce fast reconstructions, whilst maintaining image quality. This

threshold should be re-evaluated in the future, as data from more FD-HD-DOT studies become available.

3.4 Discussion and Conclusions

Prior knowledge of the underlying volume geometry and of optical properties distribution greatly affects DOT modelling. When this prior knowledge is available, for example from a subject's MRI scan, it is possible to create subject specific FEM meshes to solve the diffusion approximation. Higher mesh resolutions minimise tissue localisation uncertainty and also reduce numerical inaccuracy induced by the FEM. Therefore, if the aim is to improve imaging quality (in terms of resolution, contrast and localisation), high density meshes should be used. However, this approach causes computational tractability issues, arising from high memory requirements and low computational speed. These issues are shown to be minimised by employing iterative linear solvers and parallel computing architectures. The developed solvers can be applied to any numerical model of light propagation in tissue, such as the Radiative Transport Equation, which may provide more accuracy in situations where the Diffusion Approximation is not valid, such as short source/detector separations and presence of purely non-scattering regions. Although throughout this work it is assumed that scattering will always dominate absorption, making the Diffusion Approximation a valid model of light propagation, the utilisation of the Radiative Transport Equation should also be further investigated to determine further improvement in imaging accuracy.

In chapter 3.1.3.1, it is shown that numerical errors introduced by iterative solvers are spatially located away from the excitation source; however, the distance of the numerical errors from the excitation source is related to the termination criterion, indicating that

choosing large termination criterion to accelerate the modelling procedure, could negatively affect the quality of the reconstruction, depended on the application. Nevertheless, if the application allows, the computational time of any iterative solver can be greatly reduced by increasing the termination criterion. For example, for the models examined in this work, increasing the termination to 10^{-6} from 10^{-12} will reduce the computational time by half, but will increase the modelling error above 1% for the farthest SD separation. Therefore, the underlying physics, and the modelling and reconstruction procedure must be considered before attempting to solve with higher termination criterion. However, it is now computationally feasible to select lower termination criterion for the iterative solvers, practically eliminating any error induced by the approximate solving or the complexity of the volume, as the GPU parallelised approach has overly significantly lower computational time. Furthermore, the proposed approaches can be very efficient for systems with a large number of sources and detectors since the computational time is not greatly affected by solving for multiple sources; and in addition, can be employed in frequency domain simulations. Based on the performed experiments, the fastest approach is to parallelise the matrix to vector operations involved in iterative solvers in GPU architectures.

The inverse problem of DOT, where the creation and the inversion of the Jacobian are currently the most computationally expensive parts of the algorithm, was also optimised and parallelised. In functional brain imaging, creating sparse Jacobians enables to solve the linear inverse problem in real-time speed for each temporal set of measurements. Furthermore, creating a sparse Jacobian massively reduces memory requirements, allowing reconstructions in high-resolution mesh basis, whilst employing

numerous frequency domain measurements in multiple wavelengths. The selected sensitivity threshold was 0.1%, as it allows nodes that reach at least the brain cortex depth to contribute towards reconstructions for all measurements.

4 Frequency-Domain High-Density Diffuse Optical Tomography: A Numerical Study

Measurements of dynamic near-infrared (NIR) light attenuation across the human head together with model-based image reconstruction algorithms allow the recovery of three-dimensional spatial brain activation maps. Previous studies using high-density diffuse optical tomography (HD-DOT) systems have reported improved image quality over sparse arrays. These HD-DOT systems incorporated multi-distance overlapping continuous wave measurements that only recover differential intensity attenuation. We investigate the potential improvement in reconstructed image quality due to the additional incorporation of phase shift measurements, which reflect the time-of-flight of the measured NIR light, within the tomographic reconstruction from high-density measurements. To evaluate image reconstruction with and without the additional phase information, we simulated point spread functions across a whole-scalp field of view in 24 subject-specific anatomical models using an experimentally derived noise model. The addition of phase information improves the image quality by reducing localization error by up to 59% and effective resolution by up to 21% as compared to using the intensity attenuation measurements alone. Furthermore, we demonstrate that the phase data enable images to be resolved at deeper brain regions where intensity data fail. The computational tools described in chapter 3 have been employed for the calculation of the forward model and the formulation of the inverse problem.

4.1 Methodology

4.1.1 Forward problem

To model the light propagation in tissue, commonly known as the forward model, a numerical model based on the finite element method (FEM) has been employed to solve

the diffusion approximation (Eq. 2.1) through the NIRFAST[58] modelling and image reconstruction software package.

High-resolution meshes (mean mesh element volume < 1 mm^3, leading to ~650,000 nodes) have been used throughout this work, which improves the accuracy of the forward model at a significant increase in the computational cost. To overcome the computational cost, GPU parallelised linear solvers have been deployed that facilitate the efficient calculation of the forward problem in FD simulations; when using high resolution meshes with ~450,000 nodes, with 158 sources and 160 detectors the forward model calculation time is less than 3 minutes[119].

4.1.2 Inverse Problem

The sensitivity matrix, $\mathbf{J}$, also known as the Jacobian or weight matrix is calculated using the Adjoint method[76] that relates changes in boundary fluence measurements, $\partial\mathbf{y}$, with respect to small changes in underlying tissue optical parameters, $\partial\mathbf{x}$:

$$\partial \mathrm{y} = \mathrm{J}\partial \mathrm{x} \qquad \text{Eq. 4.1}$$

From the diffusion approximation equation (Eq. 2.1), in the frequency domain, the measured data are complex numbers and can be decomposed in two parts: amplitude $\mathbf{y_I}$, and phase $\mathbf{y_\Theta}$. Using the Rytov approximation, differential measurements with respect to a baseline $\mathbf{y_0}$ are described as a logarithmic ratio for amplitude and a difference in phase:

$$\ln\left(\frac{\mathrm{y}}{\mathrm{y}_0}\right) = \ln\left(\frac{\mathrm{y_I e^{iy_\Theta}}}{\mathrm{y_{I_0} e^{iy_{\Theta_0}}}}\right) = \ln\left(\frac{\mathrm{y_I}}{\mathrm{y_{I_0}}}\right) + \mathrm{i}\left(\mathrm{y_\Theta} - \mathrm{y_{\Theta_0}}\right) \qquad \text{Eq. 4.2}$$

Therefore, the frequency domain Jacobian describes changes in logarithmic intensity, $\partial \ln \mathbf{y_I}$, and linear differences in phase $\partial \mathbf{y_\theta}$, with respect to changes in the absorption, $\partial \boldsymbol{\mu_a}$, and the diffusion, $\partial \boldsymbol{\kappa}$ coefficients:

$$\begin{bmatrix} \partial ln\mathbf{y}_\mathrm{I} \\ \partial \mathbf{y}_\Theta \end{bmatrix} = \begin{bmatrix} \dfrac{\partial ln\mathbf{y}_\mathrm{I}}{\partial \mu_\mathrm{a}} & \dfrac{\partial ln\mathbf{y}_\mathrm{I}}{\partial \kappa} \\ \dfrac{\partial \mathbf{y}_\Theta}{\partial \mu_\mathrm{a}} & \dfrac{\partial \mathbf{y}_\Theta}{\partial \kappa} \end{bmatrix} \begin{bmatrix} \partial \mu_\mathrm{a} \\ \partial \kappa \end{bmatrix} \qquad \text{Eq. 4.3}$$

The existing literature on fNIRS and fDOT largely report on vascular response when utilised in functional brain imaging, namely detecting changes in haemoglobin concentration only. This is depicted as a measured change in absorption only at different wavelengths, bilateral of the isotropic point between oxygenated and non-oxygenated haemoglobin. A change in tissue absorption will also affect the measured phase, as the average photon path changes with absorption changes. The same is in theory true for changes in scatter, however, there are no reported scattering characteristics that are related to vascular changes. When correlation is shown between phase information and detected brain activation in fMRI, comparing the BOLD effect with the synchronous phase signal[41], it is still studying the same underlying haemodynamic vascular effect, that is recovering the absorption related changes uses the change in phase measurement. When phase is linked directly to brain activity, as in the Event Related Optical Signal (EROS) studies, minuscule differences in the structural organization of the neurons are hypothesised to affect the scattering coefficient which are then observed through a very small phase changes – which is not synchronous with the hemodynamic but with the electrical activity of the brain[40]. Therefore, assuming that only changes in the absorption

coefficient are associated with vascular haemodynamic changes within the brain, this work will focus on modelling and reconstructing changes in absorption, $\partial \boldsymbol{\mu}_{\boldsymbol{a}}$, alone, reducing the previous equation to:

$$\begin{bmatrix} \partial ln\mathbf{y}_{\mathrm{I}} \\ \partial \mathbf{y}_{\Theta} \end{bmatrix} = \begin{bmatrix} \dfrac{\partial ln\mathbf{y}_{\mathrm{I}}}{\partial \mu_{\mathrm{a}}} \\ \dfrac{\partial \mathbf{y}_{\Theta}}{\partial \mu_{\mathrm{a}}} \end{bmatrix} [\partial \mu_{\mathrm{a}}] \qquad \text{Eq. 4.4}$$

and using measurements at two wavelengths of 690 nm and 830 nm, these wavelengths were employed because the available FD system (ISS Imagent™, Champaign, Illinois) used for the experiments operated with these wavelengths. The dynamic changes of oxy (HbO_2) and deoxy (HbR) haemoglobin concentration can be calculated using:

$$\begin{bmatrix} \partial HbO_2 \\ \partial HbR \end{bmatrix} = \begin{bmatrix} \varepsilon_{HbO2,690} & \varepsilon_{HbR,690} \\ \varepsilon_{HbO2,830} & \varepsilon_{HbR,830} \end{bmatrix} \begin{bmatrix} \partial \mu_{a,690} \\ \partial \mu_{a,830} \end{bmatrix} \qquad \text{Eq. 4.5}$$

where $\varepsilon_{c,\lambda}$ is the extinction coefficient for chromophore c (HbO_2/HbR) at wavelength λ. Conversely, it is possible to map the changes in oxy (HbO_2) and deoxy (HbR) haemoglobin concentration directly to the measured dynamic data, by considering a Spectral Jacobian:

$$\begin{bmatrix} \partial ln\mathrm{y}_{\mathrm{I}_{690}} \\ \partial \mathrm{y}_{\Theta_{690}} \\ \partial ln\mathrm{y}_{\mathrm{I}_{830}} \\ \partial \mathrm{y}_{\Theta_{830}} \end{bmatrix} = \begin{bmatrix} \dfrac{\partial ln\mathrm{y}_{\mathrm{I}_{690}}}{\partial \mu_{\mathrm{a},690}} \cdot \varepsilon_{\mathrm{HbO2},690} & \dfrac{\partial ln\mathrm{y}_{\mathrm{I}_{690}}}{\partial \mu_{\mathrm{a},690}} \cdot \varepsilon_{\mathrm{HbR},690} \\ \dfrac{\partial \mathrm{y}_{\Theta_{690}}}{\partial \mu_{\mathrm{a},690}} \cdot \varepsilon_{\mathrm{HbO2},690} & \dfrac{\partial \mathrm{y}_{\Theta_{690}}}{\partial \mu_{\mathrm{a},690}} \cdot \varepsilon_{\mathrm{HbR},690} \\ \dfrac{\partial ln\mathrm{y}_{\mathrm{I}_{830}}}{\partial \mu_{\mathrm{a},830}} \cdot \varepsilon_{\mathrm{HbO2},830} & \dfrac{\partial ln\mathrm{y}_{\mathrm{I}_{830}}}{\partial \mu_{\mathrm{a},830}} \cdot \varepsilon_{\mathrm{HbR},830} \\ \dfrac{\partial \mathrm{y}_{\Theta_{830}}}{\partial \mu_{\mathrm{a},830}} \cdot \varepsilon_{\mathrm{HbO2},830} & \dfrac{\partial \mathrm{y}_{\Theta_{830}}}{\partial \mu_{\mathrm{a},830}} \cdot \varepsilon_{\mathrm{HbR},830} \end{bmatrix} \begin{bmatrix} \partial \mathrm{HbO} \\ \partial \mathrm{HbR} \end{bmatrix} \qquad \text{Eq. 4.6}$$

This formula is used in conjunction with a noise model to produce realistic boundary data by perturbating haemoglobin concentration values.

To demonstrate the difference in the sensitivity distribution and magnitude of both log-amplitude and phase, as a function of source and detector distance, the Jacobian for a 3-layered model, at a single wavelength of 830 nm, at a modulation frequency of 140 MHz is shown in Figure 4.1. Here, the FD Jacobian has been calculated for 4 source and detector separations of 10, 20, 30 and 40 mm, for a 3-layered 2D model, with a skin/scalp layer of 5 mm, skull layer of 8 mm and brain layer (combining grey and white matter) with optical properties shown in Table 3.2.

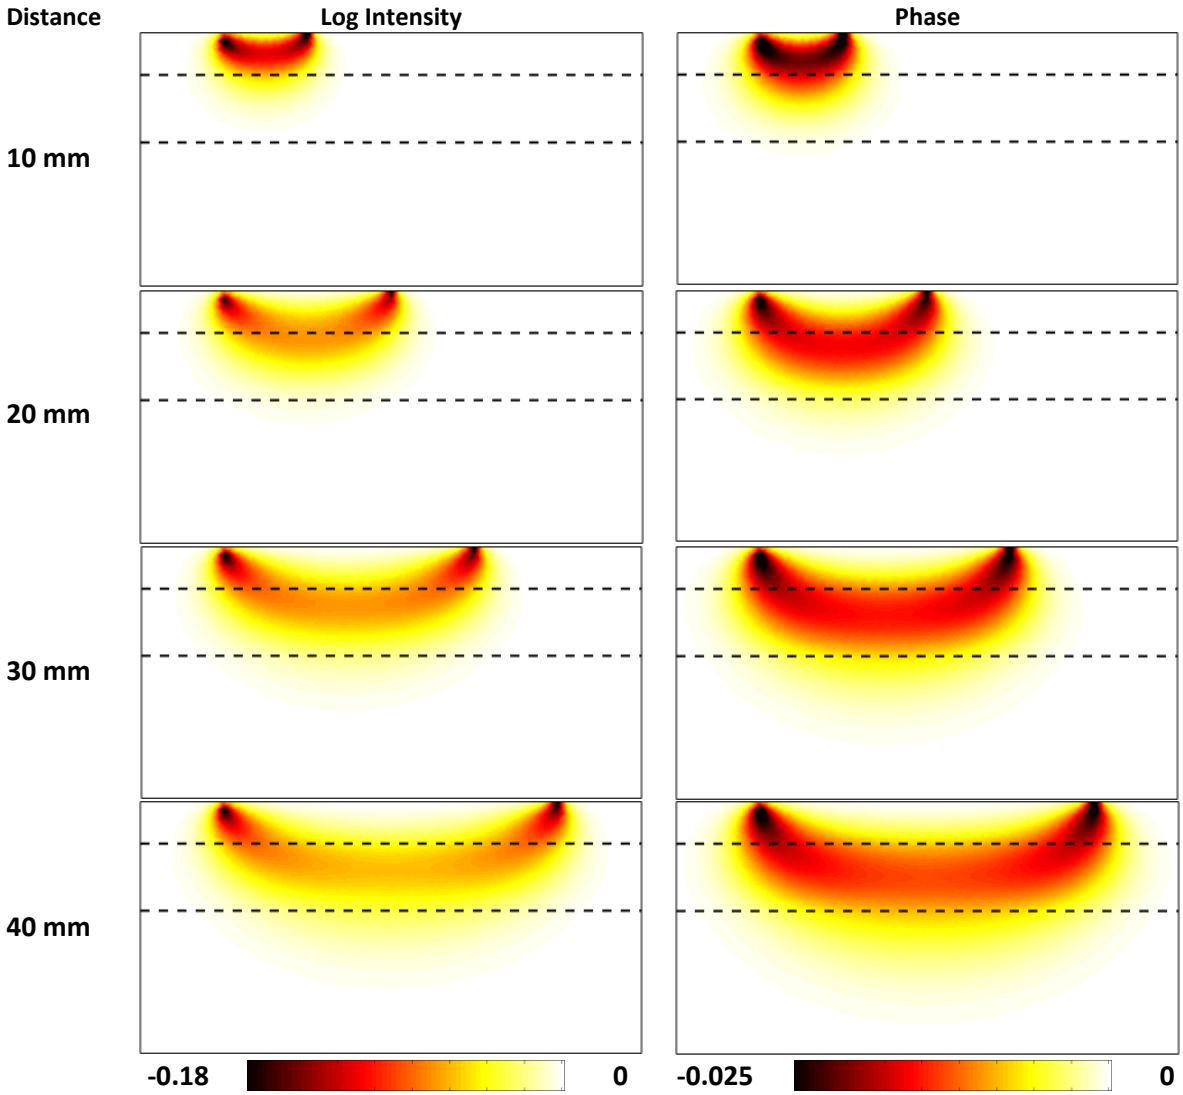

Figure 4.1 Frequency domain Jacobian for Log Intensity (Joules/mm^2/mM) and Phase (radians/mM) for HbO2. Each row corresponds to a different source and detector distances of 10 – 40 mm. The dashed line represents the boundary of each layer, skin/scalp (top), bone (middle) and brain (bottom) tissue, each having different optical properties, as depicted in Table 4.1.

Several features are evident from Figure 4.1: (1) The Jacobian for Phase is showing sensitivity weighting at deeper regions as compared to log Intensity, highlighting the potential for deeper tissue sampling, (2) the magnitude of change for log Intensity measurements is larger than that as compared to Phase but (3) the distribution of the sensitivity for Phase appears more homogeneous across all tissues as compared to log Intensity reducing the 'hyper-sensitivity' to superficial signal contamination. These

same features can also be demonstrated for HbR and is also true for a wavelength of 690 nm.

Table 4.1 Baseline haemoglobin concentrations and optical properties in the head model[114].

Region	Oxy-haemoglobin (mM)	Deoxy-haemoglobin (mM)	Scatter Power	Scatter Amplitude
Skin	0.057	0.031	1.16	0.53
Skull	0.044	0.019	0.89	0.73
CSF	0.011	0.008	0	0.3
Grey Matter	0.056	0.035	1.74	0.51
White Matter	0.068	0.027	1.31	0.82

Since the vascular dynamic changes during functional activations are small, the inverse model as used for parameter recovery can be considered as a linear problem, where a Moore-Penrose pseudoinverse with Tikhonov regularization Eq. 2.11 is used to perform a linear recovery (Eq. 2.12).

For the FD Jacobian, Eq. 4.4, the regularization is typically applied separately for log intensity and phase parts of the Jacobian[58]. In this work, the absorption coefficient is recovered for each wavelength then spectrally unmixed to recover HbO_2 and HbR as described in equation 4.6. In previous work, spatially constrained reconstructions [83] or spatially-variant regularization schemes[32,120] have been proposed that facilitate reconstructions of perturbations at deeper locations. To allow a direct evaluation of the potential increase in imaging performance afforded by the additional phase-based measurements, no spatially-variant regularization methods have been considered in this work.

4.1.3 Subject-specific models

The subject-specific models are based on 14 female and 10 male individual subjects with a mean age of 26 (+/-4), with T1-weighted MPRAGE [echo time (TE)= 3.13 ms, repetition time (TR) = 2400 ms, flip angle = 8 deg, 1 × 1 × 1 mm3 isotropic voxels] scans acquired for each subject. All subjects passed MR screening to ensure their safe participation. Informed consent was obtained, and the research was approved by the Human Research Protection Office at Washington University School of Medicine.

High-resolution subject-specific Finite Element Models (FEM), consisting of five tissue layers were created, with mean element volume of ~1 mm^3, resulting in ~450,000 nodes for each mesh using NIRFAST[97]. A sparse sensitivity matrix, J, (sensitivity threshold set to 0.1%[121]), was then created for each of the 24 subject-specific models with realistic optical properties for a 690 nm and 830 nm wavelengths, shown in Table 4.1, where the scattering is expressed in scatter power and amplitude, coefficients that when considering Mie scattering[122], and elastic spherical particles, relate to particle size and density. The reduced scattering coefficient was derived from the scatter amplitude and power, using $\mu_s' = s_a\lambda^{-s_p}$, where s_a is the scatter amplitude, s_p is the scatter power and λ is the wavelength in micrometres. The haemoglobin concentrations and scattering properties as noted in Table 4.1 are a convention used that will allow the calculation of each tissue's specific absorption and reduced scattering coefficient at any NIR wavelength. Specifically using known spectrally varying extinction coefficient of Oxy/Deoxy-haemoglobin it is possible to calculate the absorption coefficient using Beer's law, whereas the spectrally varying scattering parameters are calculated using Mie's theory[97]. It is worth stating that the effective haemoglobin concentration and

scattering properties for CSF are derived by utilizing published values of absorption and reduced scattering (at 750 and 850 nm)[114] with Beer's law and Mie's theory.

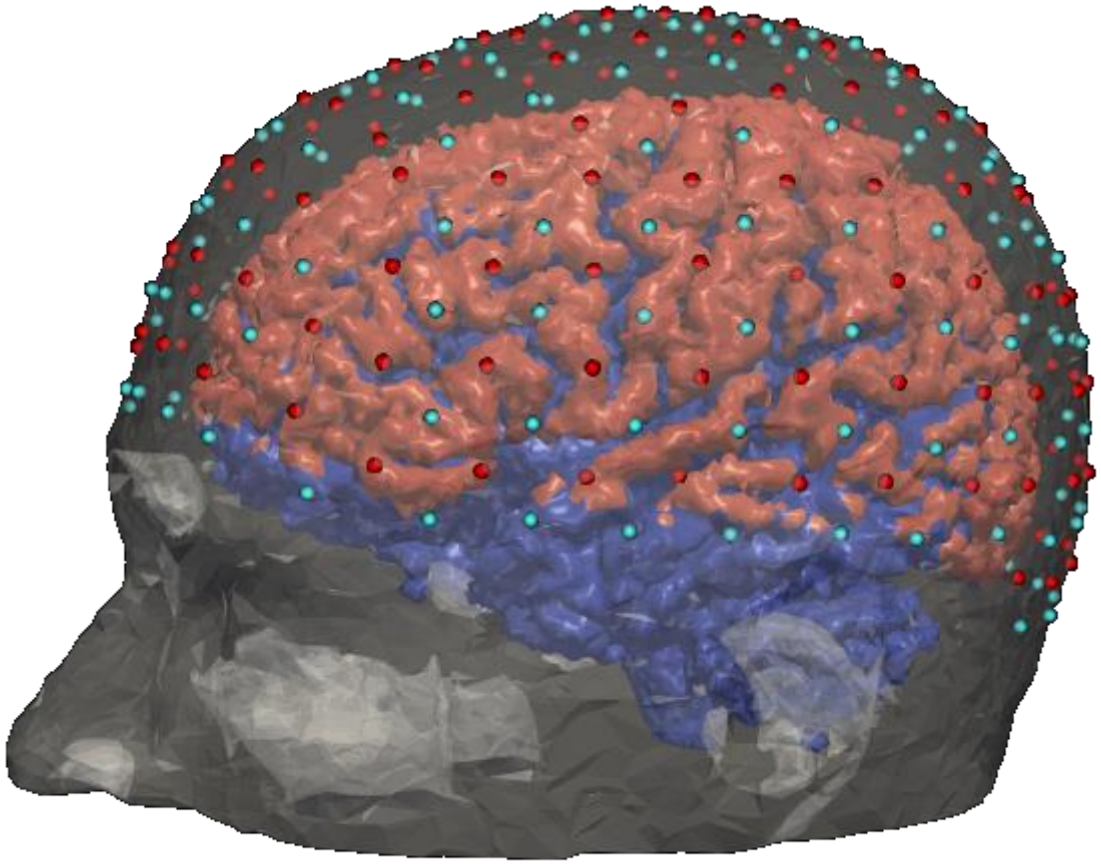

Figure 4.2 The modelled HD-DOT system: 158 sources (red dots) and 166 detectors (cyan dots), the surface of brain is in orange and blue colour. The orange region denotes cortex locations in the field of view of the optical sensors.

In order to generate the boundary data, the nodes on the surface of the brain for each subject laying directly under the imaging array were selected to simulate point perturbations for HbO2 and HbR, representing focal activations; this region is the orange part of the brain surface in Figure 4.2. This data was then used for the calculation of a point spread function (PSF) to allow an analysis of the quality of the image reconstruction within a field of view directly under the imaging cap. For each subject, the selected region depth from the surface of the brain varies, as the brain surface changes

along the sulci and gyri morphology. To better demonstrate the depth variation of the surface of the brain, the histograms of brain surface depth across all 24 subjects are shown in Figure 4.3(a) whereas 4.3(b) shows the brain surface depth probability distribution for each subject-specific model.

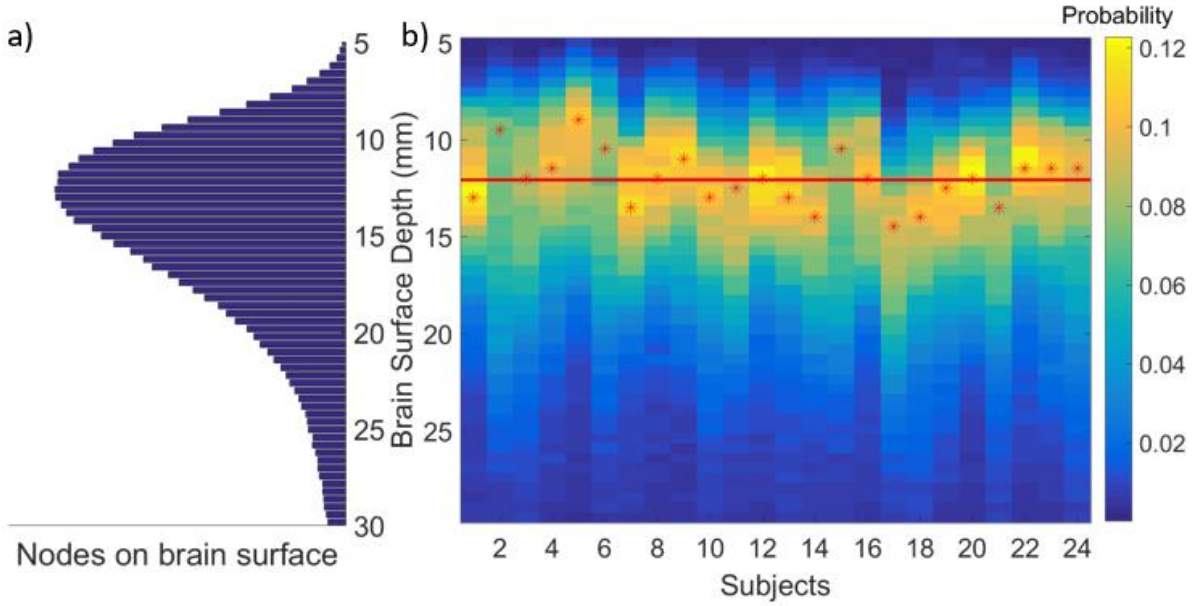

Figure 4.3 Brain surface depth probability. a) Histogram of brain surface depth across 24 subjects. b) Brain surface depth probability distribution for each subject-specific model. The red stars mark the maximum probability for each subject while the red line is the most probable depth across all subjects (12.5 mm).

This further highlights the importance and the need for improving the depth resolution of HD-DOT and demonstrates the need for modelling individual surface activation points for the presented analysis. The brain surface depth varies across the subjects, ranging from 5 to 30mm and brain surface points presented in this study lie on the cortical surface with a median depth of 12.5mm, Figure 4.3 (a) and (b). This brain cortex depth analysis informed the decisions of the Jacobian threshold, allowing the determination of the best parameters for computational optimisation without loss in imaging quality. Specifically, the 0.1% threshold was chosen to ensure that nodes that

reach at least the surface of the white matter (and hence cortical surface) are considered for reconstruction.

4.1.4 Realistic system noise model

The modelled system array adopted in this work is an HD-DOT system, with source-detector (SD) separation distances ranging from 13 to 47 mm as described previously[15]. The SD pairs are extended over the whole of the head, in order to allow coverage of a large area of the optically accessible cortical surface of the brain that has previously been used as a benchmark for fDOT simulation studies[117]. The modelled HD-DOT array contains a total of 158 light sources at two wavelengths of 690 and 830 nm, and 166 detectors, resulting in 3,500 associated SD pairs within 50 mm separation at each wavelength. The SD pairs are grouped into sets of nearest neighbours (NN) measurements (i.e. NN1, NN2, NN3, NN4, for first through fourth nearest neighbour, respectively) with mean distances of 13, 29, 39, and 47 mm. In this work, the labelling for reconstruction with respect to a set of nearest neighbour measurements used is the maximum nearest neighbour; for example, for NN4 includes NN1, NN2, NN3, and NN4 measurements.

Simulating PSFs without noise can provide insights into the theoretical aspects of image reconstruction capabilities, and it is a useful tool towards optical array spatial design and development. However, in practice the reconstruction quality is affected by the noise present in the experimental system. Moving towards realistic simulations, the magnitude of simulated haemoglobin perturbations should match the appropriate physiological values, while the noise should match this of a real system. Since the main source for the presence of noise is the reduction of measured photons due to absorption,

source-detection distance is a major factor for noise modelling, thus a noise model can be generalised to be expressed as a function of source-detector separation.

To generate a realistic noise model to be used with simulations, data were recorded using a frequency domain system console (ISS Imagent™, Champaign, Illinois) employing six detectors at distances of 18 mm to 58 mm from the source, placed over the optical cortex. The source was modulated at 140 MHz and data were sampled at 39.74 Hz. Data were recorded for 120 seconds while the subject sat still and quietly fixated on a blank screen. Data were then processed following methods that have been previously described in[80]. Briefly, the data were band-pass filtered to 0.01-0.1 Hz to remove slow drifts, systemic physiology (heartbeat and respiration), down sampled to 1 Hz, and block averaged in ten blocks of 12 seconds each. The same processing was applied to the measured amplitude and phase signals. The noise was then estimated as the standard deviation of the log mean intensity signals, therefore it is a ratio-metric difference for the intensity and a linear phase difference (in degrees) for phase and is shown in Figure 4.4. As the system utilised Photomultiplier tubes for signal detection, each detector was at a different voltage bias to increase signal gain and throughout the experiment, all data was ensured to be above the expected noise floor.

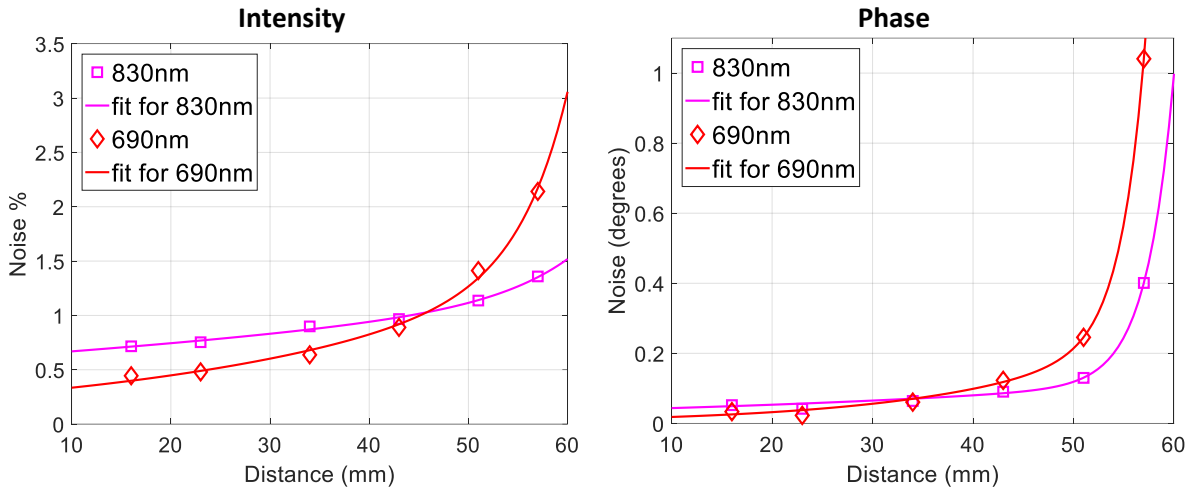

Figure 4.4 In-vivo measured noise and a two-term exponential fit of noise versus distance from the source. For intensity (left) and phase (right) measurements, on 690 nm (red diamonds) and 830 nm (pink squares) wavelengths.

The noise model n was calculated as a function of source-detector distance, r, via a two-term exponential of the form $\mathrm{n}(r) = \mathrm{a}e^{\mathrm{br}} + \mathrm{c}e^{\mathrm{dr}}$, and was found to provide the best fit with coefficients shown in Table 4.2. In previous work, using the same HD array configuration and similar signal processing, but acquiring measurements with a CW system with LEDs at 830 nm, the amplitude noise was reported to be 0.12% for NN1, 0.15% for NN2, 0.41% for NN3 and 1.42% for NN4 [118]. The experimentally derived noise model generated in this work corresponds to a noise level estimate for each neighbourhood as shown in Table 4.3.

Table 4.2 The estimated coefficients for the two-term exponential fit of the noise model.

	Amplitude noise model		Phase noise model	
Coefficients	830 nm	690 nm	830 nm	690 nm
a	0.6019	0.2502	1.917e-10	3.933e-11
b	0.01052	0.02913	0.3708	0.4161
c	9.685e-05	4.625e-06	0.03573	0.0105
d	0.1382	0.2128	0.02002	0.05585

To simulate realistic measured data, perturbations of HbO_2 and HbR values were induced employing Eq. 4.6, using perturbation values scaled to minimise the difference between simulated and measured data. The perturbation for HbO_2 was scaled to be 3.8

μM while HbR was -1.8 μM. Evaluations were carried out using noise-free simulations and simulations that included the empirically derived noise model. Stochastic realistic noise was linearly added to each measurement based on sampling from a Gaussian distribution with zero mean and a standard deviation that was a function of source-detector distance with parameters given in Table 4.3.

Table 4.3 The estimated noise levels for each measurement neighbourhood.

	Amplitude noise (%)		Phase noise (degrees)	
NN	830 nm	690 nm	830 nm	690 nm
NN1	0.69	0.36	0.044	0.021
NN2	0.81	0.56	0.062	0.05
NN3	0.91	0.77	0.076	0.088
NN4	1	0.95	0.088	0.126

4.1.5 Evaluation metrics

To evaluate the image quality of each PSF, for each modelled focal activation, metrics as previously described in similar work have been utilised [79]:

(1) Full width at half maximum (FWHM) defined as the maximum separation between all pairs of points in the activation above half of the maximum recovery. Since the activation is simulated on a single node, representing a very small volume in space, smaller FWHM represents a better recovery.

(2) Localization error defined as the distance between the centroid of the recovery and the known location of the perturbed node. A perfectly located recovery would have localisation error zero, therefore smaller localization error values are better.

(3) Effective resolution defined as the diameter of a sphere centred at the known perturbed node that can enclose all recovered nodes. This is a combination of FWHM and localisation error, therefore once more, the smaller the effective resolution value the better the recovery.

It is likely that the recovered volume above half of the maximum will not create a single focal recovery, resulting in multiple scattered activations[84]. This is more evident in the cases with added noise, as the reconstruction often has multiple recovered activations. Therefore, to automate the quantification process, such activations are isolated by a connected component algorithm, using the FEM elements as the connectivity index and considering only the nodes that are associated with each other through at least one connected element. The recovered activations are then selected based on the maximum accumulated recovery value for evaluation without the use of any prior regional or locational knowledge. Therefore, a recovery near the perturbation will not necessarily be selected for evaluation, unless it is the strongest integrated recovery in the volume. Any recovery with localization error greater than 8 mm is considered a product of noise and therefore not included in the statistical analysis according to previous literature[84]. The Wilcoxon signed-rank test was used to assess the statistical difference between metrics of CW and FD reconstructions.

4.2 Results

In this section when referring to “qualitative” evaluation the reconstruction quality is assessed based on visual cues, and “quantitative” when assessing based on the values of the employed metrics. If the employed metrics are appropriate, these two assessments should correlate. In section 4.2.1, the qualitative results are paired with quantitative to only visually point out that there is an agreement between visuals and metrics, therefore the chosen metrics are appropriate. Then, the image reconstruction quality is measured only with metrics in section 4.2.2.

4.2.1 Qualitative Evaluation

In this work, the brain surface under the imaging array was selected in 24 subject-specific models and point activations modelled as perturbations were induced for each node in the selected area, allowing the recovery of PSFs both with and without noise. To provide an overview of the spatial distribution characteristics of the recovered PSFs, consider three points on the primary motor cortex within Brodmann Area 4, along the precentral gyrus, as shown in Figure 4.5. These three focal activations are modelled at depths of 13.3 mm, 17.2 mm and 20.5 mm (Figure 4.5). PSF of each individual modelled focal activation for both HbO_2 and HbR have been reconstructed and as both demonstrated similar results, for conciseness, only those of HbO_2 are shown here.

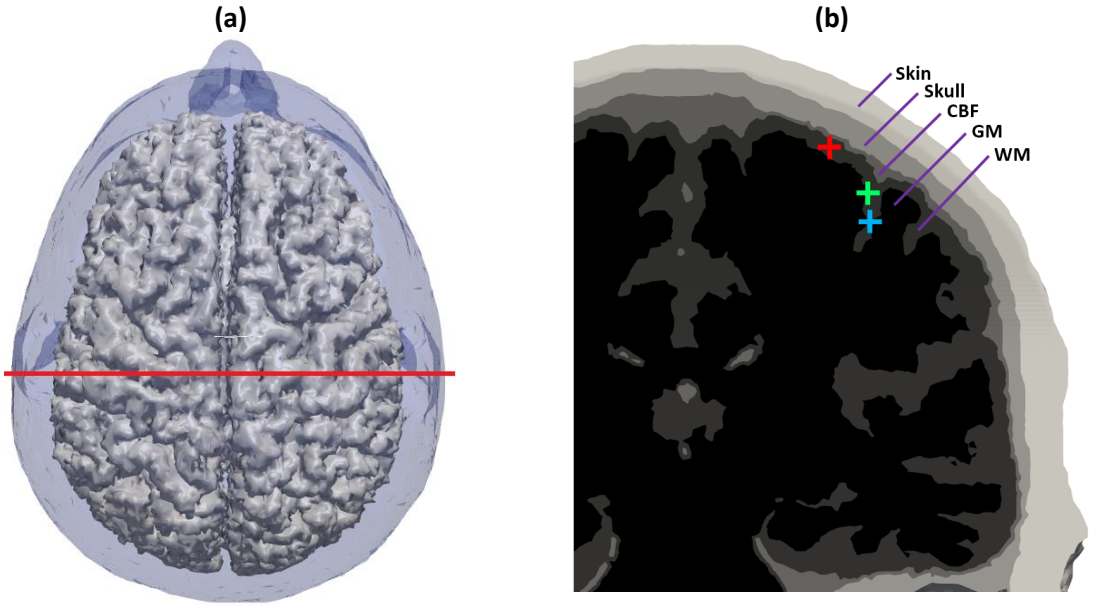

Figure 4.5 Specific brain areas activations. a) Axial view of the brain, the red line denotes the coronal slice selected to view the primary motor cortex, Brodmann Area 4, along the precentral gyrus. b) The 5 layers of the model: skin, skull, cerebral fluids, white matter, grey matter. Selected cortex points at different depths: 13.3 mm (left/red), 17.2 mm (middle/green), and 20.5 mm (bottom/blue).

For the focal point at 13.3 mm depth, Figure 4.6, top row, both CW and FD provide qualitatively similar reconstructions for all SD pairs, but the FD demonstrated a visually better reconstruction in each case particularly for shorter NN pairs (NN1 and NN2)

measurements. Specifically, for the noise-added recoveries, calculated metrics are summarised in Table 4.4 and 4.5 demonstrating that for larger NN, the image recovery of the focal activation is improving for both CW and FD, and FD is consistently outperforming CW.

For a perturbation at 17.2 mm depth, Figure 4.6, middle row, the noise added CW data fails to reconstruct for NN2 and NN3, with calculated metrics summarised in Table 4.4 and 4.5. This case demonstrates that for some depth range the CW reconstructions may fail if the dynamic range of the detectors cannot provide measurements with a good signal to noise ratio (SNR), while the FD recoveries can perform reasonably well even when using only NN2. As highlighted in Ch. 1.3, intensity measurements are dominated by the early arriving photons, which travel a more superficial route, since intensity attenuates exponentially with distance. In contrast phase changes linearly with distance, therefore phase measurements can capture perturbations that take place in deeper regions.

Table 4.4 Localisation accuracy for individual activation using noise-added data.

Depth (mm)	CW Localisation error (mm)			FD Localisation error (mm)		
	NN2	NN3	NN4	NN2	NN3	NN4
13.3	6.5	2.3	1	2.2	1.6	0.85
17.2	-	-	5.2	6.2	3.8	2.2
20.5	-	-	-	-	-	6.1

Table 4.5 Effective resolution for individual activation using noise-added data.

Depth (mm)	CW Effective resolution (mm)			FD Effective resolution (mm)		
	NN2	NN3	NN4	NN2	NN3	NN4
13.3	22.9	21.4	11	17.7	16.2	9.8
17.2	-	-	17.9	21.42	20.58	17.9
20.5	-	-	-	-	-	24.8

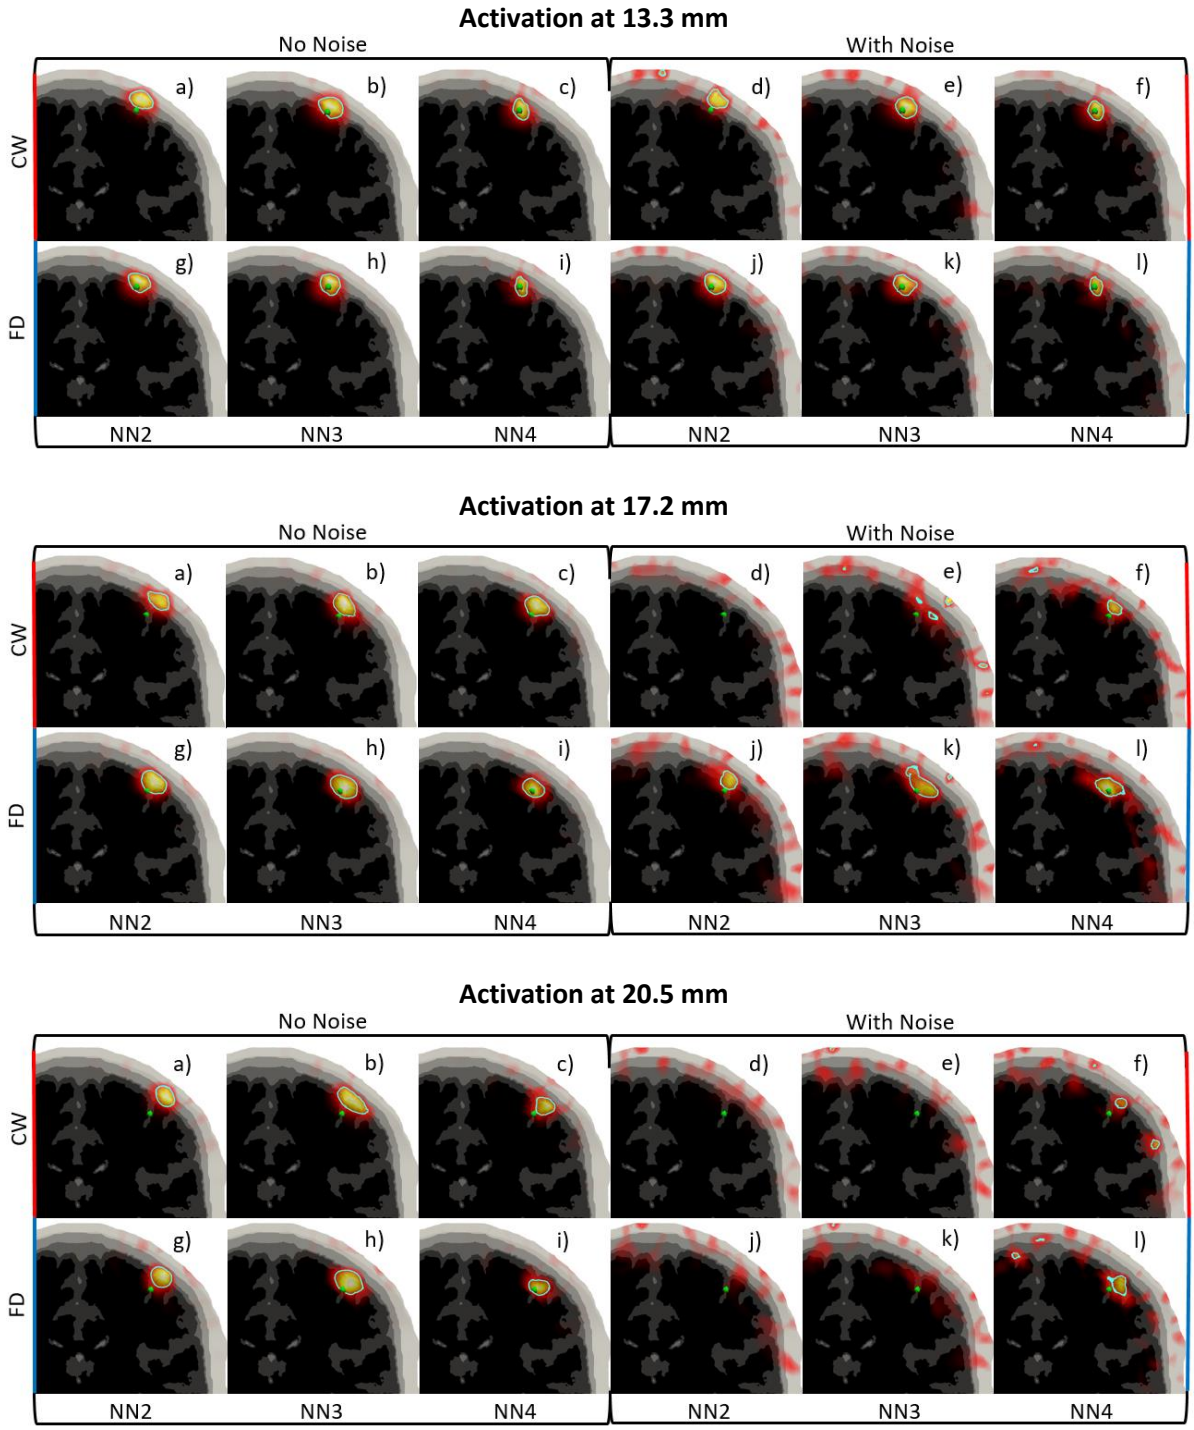

Figure 4.6 Recovered point spread functions with different nearest-neighbour measurements (NN2, NN3, NN4), for activations at different depths. The perturbation location is denoted with the green dot. The FWHM contour (when present) is noted in cyan. Recoveries with CW without noise (a,b,c), CW with noise (d,e,f), FD without noise (g,h,i) and FD with noise (j,k,l).

When reconstructing for simulated data without noise, an activation positioned 20.5 mm deep from the surface of the head may be recovered reasonably well using any number of measurement neighbourhoods, using CW only or FD data (Figure 4.6) though

the FD based data provide reconstructions with better image quality for all NN cases. However, in the presence of noise, recovery based on the NN2 and NN3 completely fails in both CW and FD. The recovery for CW using NN3 appears to provide a recovery near the perturbated point, but this is not the strongest recovery in the volume, therefore it is considered as noise artefact and as such is not selected for evaluation. However, when using NN4 the FD data recovery demonstrates a better recovery, achieving a 6.1 mm localization error with 24.8 mm effective resolution. In the presence of noise recoveries for such deep points often fail, which leads to large increases in the localization and resolution metrics.

4.2.2 Quantitative Evaluation

To provide a comprehensive evaluation of the benefits of FD data over CW, images were reconstructed for a set of NN1-NN4 SD pairs, using both noise-free and noise added data. In the presence of noise, the recoveries for deep activation points fail, which leads to significant increases in the calculated error metrics. The most obvious change is seen in the localization error, as beyond certain depths the recovered focal activations are not successful. In accordance with existing literature[84], any recovery with localization error greater than 8 mm is considered a product of noise. To identify the percentage of recoveries within the imposed localization limit, a success rate metric is introduced (Figure 4.7), where a 50% success rate limit was set to examine the depths at which different approaches perform adequately. Recoveries exceeding this limit are not included in any further results or statistical analysis.

Each recovered focal activation is assessed using metrics of the full-width at half maximum (FWHM) (Figure 4.9), localization error (Figure 4.8), and effective resolution

(Figure 4.10). These results, quantised according to depth from the surface of the head, are summarised in Table 4.6 and Table 4.7 and compared in Table 4.8. The noise-free evaluation is also presented to provide a benchmark of limitations imposed due to the ill-posed inverse problem, the employed regularization, and to provide a lower-bound on metric quality given further advancements in instrumentation that will lead to hardware with lower noise characteristics.

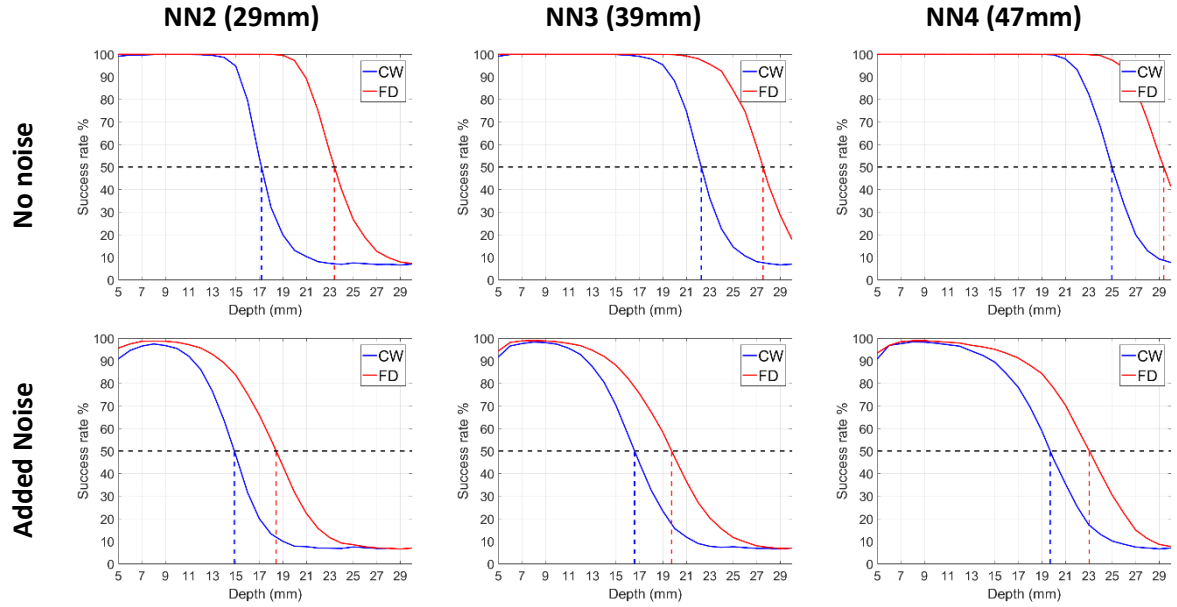

Figure 4.7 HbO_2 success rate for 8mm localization error limit. The dashed blue and red line denotes the depth where the recoveries reach 50% success rate for CW and FD respectively.

For the noise-free cases, the success rate maintains near 100% until it steadily declines at a characteristic distance (Figure 4.7). The FD approach consistently performs better than the CW case in deeper regions. For CW recoveries, the 50% success rate is crossed at 17.18 mm for NN2, 22.27 mm for NN3, and 24.97 mm for NN4; while the FD 50% success rate points are at 23.04 mm, 27.53 mm and 29.38 mm respectively. For the noise added cases the success rate drops under 50% at shallower depths. Specifically, for the CW case, the 50% success rate is crossed at 14.87 mm for NN2, 16.59 mm for NN3, and 19.71 mm for NN4, while the FD 50% success rate points are at 18.43 mm, 19.75 mm

and 23.04 mm respectively. The increase in depth sensitivity while using depth data appear to be 4-5 mm. Examining the brain cortex depth probability (Figure 4.3) we can assume that NN2 and higher measurements can perform adequately in the 10-15 mm depth range – where most of the brain cortex volume lies.

The FD data significantly improves the localization throughout the imaging domain for noiseless simulations (Figure 4.8). At each depth, the use of FD data demonstrates a statistically significant improvement ($p<0.01$) in localization accuracy. For the noise added case, the localization error is larger for activations deep within the brain, while the improvement of the FD reconstructions is again statistically significant ($p<0.01$) at each depth interval. Specifically, when comparing a depth range of 10-15 mm, the localization error for FD presents a 58.9% improvement for NN2 recoveries, 47.3% for NN3 recoveries and 21.1% for NN4 (Table 4.8). These improvements are statistically significant for all modelled focal activations ($p<0.01$) using the Wilcoxon signed-rank test. It is also worth noting that the improvement is mainly in the depth direction with the directional error between FD and CW being 10-fold larger in depth than the lateral direction.

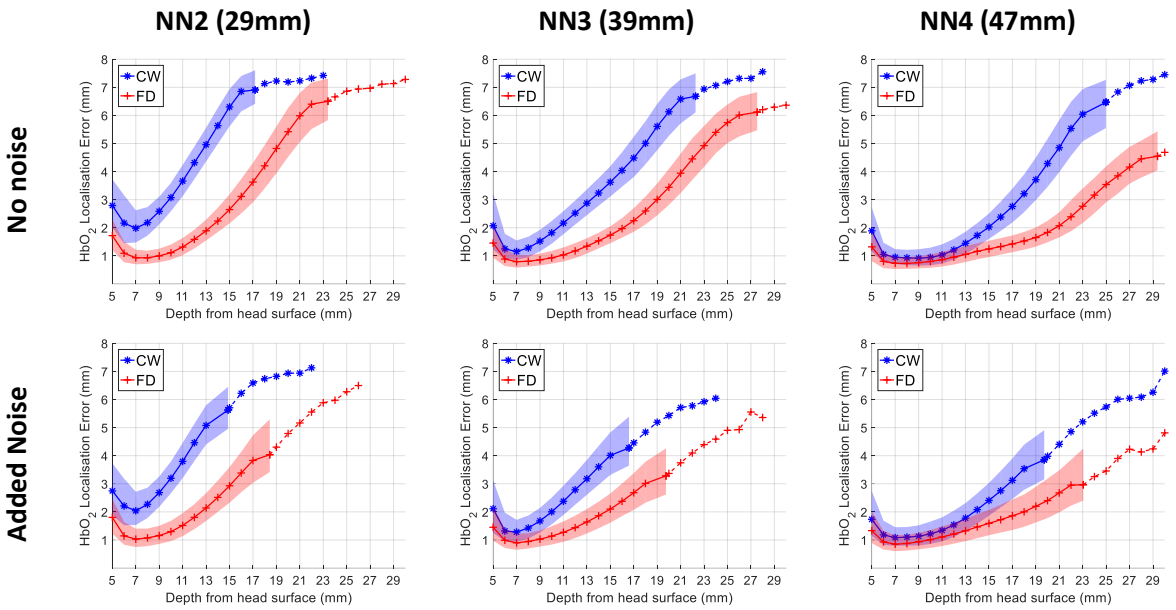

Figure 4.8 HbO_2 localization error. Median for CW (blue asterisk line) and FD (red cross line); the coloured bands represent the 25th to 75th percentile range. The percentile bands stop at the point where the 50% success rate is reached, while the median line continues until the 10% success rate is reached.

The FWHM for the noise-free cases increases with depth, but the FD reconstructions do not always significantly improve the FWHM (e.g. in the NN2 case), however employing the FD approach results in slightly improved performance and smaller variations within each depth interval (Figure 4.9). For the noise added case, the FWHM still presents slight improvements when employing FD recoveries, however, the improvement is not always statistically significant (Table 4.8), particularly for NN2 at activation depths of greater than 15 mm.

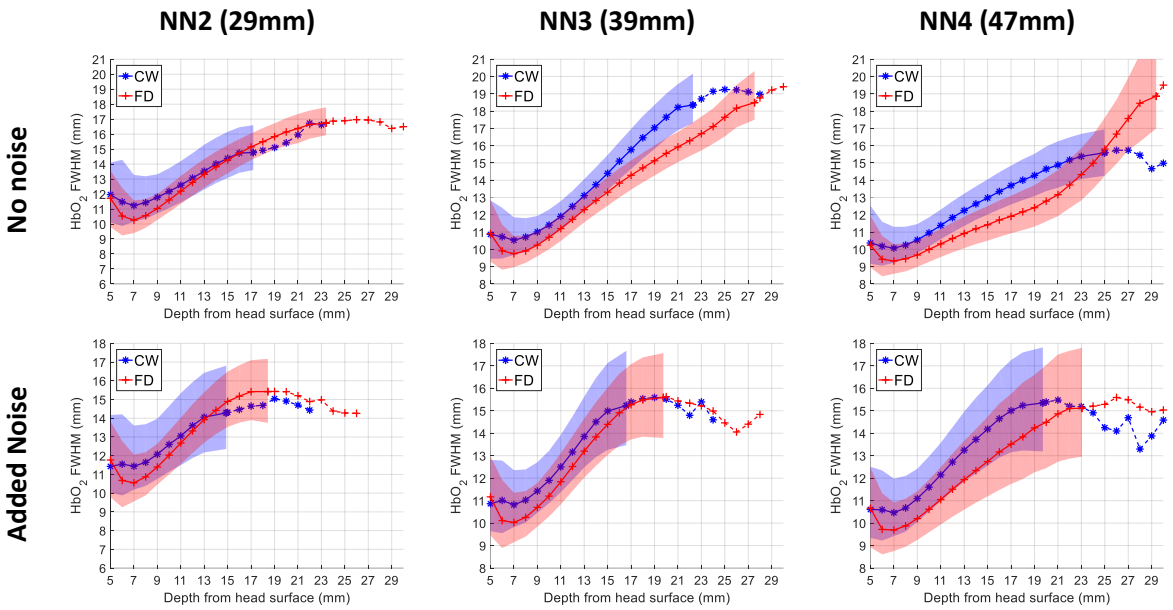

Figure 4.9 HbO_2 full width at half maximum (FWHM). Median for CW (blue asterisk line) and FD (red cross line); the coloured bands represent the 25th to 75th percentile range. The percentile bands stop at the point where the 50% success rate is reached, while the median line continues until the 10% success rate is reached.

Finally, the effective resolution with respect to depth from the head surface is shown in Figure 4.10. As expected, in the noise-free cases the effective resolution increases with depth. As the effective resolution is a function of both FWHM and localization, in the noise-added cases, the improvements observed are larger than those of the FWHM but smaller than the localization improvements. However, the observed improvement for FD data is statistically significant for each depth interval. Specifically, when comparing for a depth range of 10-15 mm, the effective resolution for FD presents a 20.6% improvement for NN2 recoveries, 13.7% for NN3 recoveries and 10.5% for NN4 (Table 4.8). These improvements are all statistically significant for all modelled focal activations ($p<0.01$) using the Wilcoxon signed-rank test.

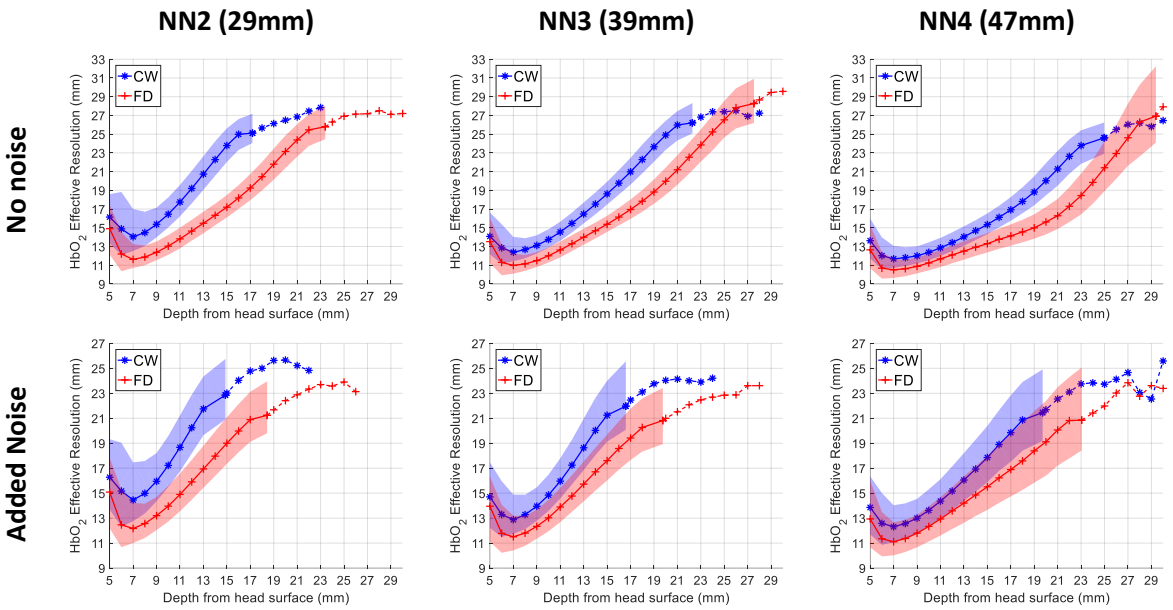

Figure 4.10 HbO_2 effective resolution. Median for CW (blue asterisk line) and FD (red cross line); the coloured bands represent the 25th to 75th percentile range. The percentile bands stop at the point where the 50% success rate is reached, while the median line continues until the 10% success rate is reached.

The results are summarised in Table 4.6 and Table 4.7 where the reconstruction metrics for recoveries with noise-added data have been binned and analysed in 5 mm depth intervals and the median, upper and lower quartile for each interval is reported. Finally, in Table 4.8 the percentage relative change for reconstruction using FD data as compared to CW is shown.

Table 4.6 Image quality metrics medians and quartiles for recoveries with CW noise-added model.

Depth(mm)	FWHM (mm)			Localization error (mm)			Effective resolution (mm)		
	NN2	NN3	NN4	NN2	NN3	NN4	NN2	NN3	NN4
5-10	$11.6^{+13.7}_{-10.3}$	$11^{+12.4}_{-9.9}$	$10.6^{+12.1}_{-9.5}$	$2.2^{+2.9}_{-1.7}$	$1.4^{+1.8}_{-1}$	$1.1^{+1.5}_{-0.8}$	$14.9^{+17.7}_{-13.2}$	13.2^{+15}_{-12}	$12.5^{+14.2}_{-11.2}$
10-15	$13.3^{+15.7}_{-11.7}$	$12.9^{+14.5}_{-11.6}$	$12.5^{+14.1}_{-11.2}$	$4.2^{+5}_{-3.3}$	$2.6^{+3.2}_{-2}$	1.5^{+2}_{-1}	$19.6^{+22.5}_{-17.3}$	$16.8^{+19.1}_{-15}$	$14.9^{+17}_{-13.3}$
15-20	$14.6^{+17.2}_{-12.6}$	$15.2^{+17.5}_{-13.4}$	$14.6^{+16.8}_{-12.8}$	$6.4^{+7.2}_{-5.6}$	$4.4^{+5.4}_{-3.6}$	$2.8^{+3.6}_{-2.2}$	$24.5^{+27.1}_{-22.3}$	$22^{+25.3}_{-19.8}$	$19^{+22}_{-16.8}$
20-25	$14.5^{+17.1}_{-12.7}$	$15.2^{+17.8}_{-12.8}$	$15.3^{+17.8}_{-13}$	$7.1^{+7.6}_{-6.1}$	$5.8^{+6.7}_{-4.7}$	$4.7^{+5.8}_{-3.7}$	$25.1^{+28}_{-23.2}$	$24^{+27.2}_{-21.6}$	$23^{+26.1}_{-20.3}$

Table 4.7 Image quality metrics medians and quartiles for recoveries with FD noise-added model.

Depth (mm)	FWHM (mm)			Localization error (mm)			Effective resolution (mm)		
	NN2	NN3	NN4	NN2	NN3	NN4	NN2	NN3	NN4
5-10	$10.8^{+12.2}_{-9.8}$	$10.2^{+11.5}_{-9.3}$	$9.8^{+11.1}_{-8.9}$	$1.1^{+1.4}_{-0.78}$	$0.94^{+1.2}_{-0.68}$	$0.87^{+1.2}_{-0.62}$	$12.5^{+14.2}_{-11.3}$	$11.7^{+13.3}_{-10.6}$	$11.3^{+12.9}_{-10.2}$
10-15	$13.1^{+14.6}_{-11.8}$	$12.3^{+13.7}_{-11.1}$	$11.3^{+12.9}_{-10.1}$	$1.7^{+2.2}_{-1.2}$	$1.4^{+1.8}_{-1}$	$1.2^{+1.6}_{-0.83}$	$15.6^{+17.5}_{-14}$	$14.5^{+16.3}_{-13}$	$13.3^{+15.3}_{-11.9}$
15-20	$15.1^{+16.8}_{-13.7}$	$14.9^{+16.6}_{-13.4}$	$13.2^{+15.4}_{-11.5}$	$3.4^{+4.2}_{-2.6}$	$2.4^{+3.1}_{-1.8}$	$1.7^{+2.3}_{-1.3}$	$20.1^{+22.3}_{-18.2}$	$18.7^{+21}_{-16.9}$	$16.3^{+19.1}_{-14.3}$
20-25	$15^{+17}_{-12.9}$	$15.3^{+17.4}_{-13.4}$	$14.8^{+17.4}_{-12.7}$	$5.4^{+6.5}_{-4.3}$	$4^{+5}_{-3.1}$	$2.7^{+3.5}_{-2}$	$23.1^{+25.4}_{-21.2}$	$21.9^{+24.5}_{-19.6}$	$20^{+23.4}_{-17.1}$

Table 4.8 FD reconstruction relative difference. The shaded cells are not statistically significant ($p>0.01$).

	FWHM (%)			Localization error (%)			Effective resolution (%)		
Depth (mm)	NN2	NN3	NN4	NN2	NN3	NN4	NN2	NN3	NN4
5-10	7	7	7.5	51.5	32.2	20.7	16.1	11	9.5
10-15	1.9	4.8	9.3	58.9	47.3	21.1	20.6	13.7	10.5
15-20	-4	2.4	9.6	47.2	44.8	37.6	17.9	15.7	14
20-25	-4	-1.3	3.5	23.5	30.5	43.3	7.9	9.1	13

4.3 Discussion and conclusions

To date, most HD-DOT imaging studies have concentrated on the utilization of CW data. Although some studies have reported the use of FD NIR systems for spectroscopic mapping of brain functions, these have been mainly limited to utilization of phase and amplitude data for the separation of optical scattering components from absorption. It has been recently demonstrated that the information content of phase data from a retinotopic experiment contains rich information, potentially arriving from deeper regions within the brain and therefore minimizing the superficial tissue contamination [123]. This chapter presented a detailed study of the potential improvements to image quality gained through incorporating measurements of phase along with intensity in the reconstruction process.

Utilizing a simple layered slab model, the Jacobian (i.e. the sensitivity of both log amplitude and absolute phase to absorption changes) has been calculated and shown in Figure 4.1. Several features are evident: (1) both log-amplitude and phase data show a decrease in signal due to a change to absorption; (2) while the sensitivity for both log-amplitude and phase data increases in depth of sampling with source-detector separation, the phase data demonstrates deeper sensitivity for a given source-detector separation; and (3) the hyper-sensitivity in superficial tissues observed with log amplitude data is less pronounced for phase data. This supports our earlier findings that phase data is able

to detect changes deeper within the brain and are less prone to be contaminated to changes within the superficial tissue that may mainly be attributed to systemic cardiac and respiratory signal.

Utilizing a set of 24 subject-specific models, highlights that the surface of the cortex can be from 5-30 mm deep with respect to the surface of the head (Figure 4.3), further highlighting the need for the utilization of data-types that will allow deeper sampling of the brain tissue. Although this can, in theory, be achieved through the use of larger source-detector separations, in reality, this is not possible due to system design and low SNR at measurement distances above 45 mm [44,80]. To demonstrate that FD data-types (log intensity and phase) can provide a better sampling of deeper tissue, the subject-specific models have been utilised to simulate a set of data corresponding to focal activations at all cortex surface nodes lying underneath the imaging array, highlighted as orange regions in Figure 4.2. In order to ensure that these simulated data are realistic, the noise characteristics of a commercially available FD system, ISS Imagent has been measured, as a function of source-detector nearest neighbour measurements (Figure 4.4) and were added to the simulated data (equation 4.6). To investigate the benefits of using differing NN (NN2-NN4) measurements for activations at different depths three sets of representative focal activation recovery maps are shown in Figure 4.6. In all cases as the number NN measurements increases, the depth recovery improves. In each case, the presence of noise is shown to introduce artefacts within the recovered maps. Importantly, the utilization of FD data is shown to improve all metrics of image quality relative to using CW. In almost all cases, the localization accuracy of utilizing FD data is shown to

be much better than those of CW, even including an activation at 20.5 mm where the CW data failed to recover a meaningful map.

To provide an overview of these benefits, images of focal activations for all noise-added simulated cortical activations have been reconstructed and analysed to provide group average metrics of localization accuracy (Figure 4.8) full-width-half-max (Figure 4.9) and effective resolution (Figure 4.10) with data summarised in Table 4.6, Table 4.7 and Table 4.8. For all metrics in the noise-free case, the accuracy decreases as a function of activation depth, with FD data out-performing CW in all cases. To achieve meaningful statistical comparisons a localization error limit of 8 mm has been imposed as the upper-bound for statistical analysis, excluding all recoveries that are solely products of noise[84]. For all NN measurements, the FD data provides a significant improvement in almost all recovered activations for all reported metrics (summary in Table 4.8). For example, assuming a set of NN3 measurements, it is shown that as compared to CW the FD data shows up to 59% improvement in localizing accuracy and up to 21% improvement in effective resolution. Additionally, the FD data using NN4 is observed to successfully recover activations as deep as 20-25mm, beyond the range accessible to CW.

The presented results suggest that for the same configuration of source-detector, the FD will outperform the CW reconstruction. Increasing the density of a CW imaging array will considerably increase the imaging quality: increasing from NN2 to NN3 will reduce the localisation error from 4.2 mm to 2.6 mm, while increasing to NN4 will decrease it further, to 1.5 mm, in the 10-15 mm depth range. Switching from CW to FD for a NN2 imaging array will decrease the localisation error from 4.2 mm to 1.7 mm, in the same depth range. The benefit of FD measurements seems to be greater than the

benefit from more dense measurements using CW alone, nevertheless this improvement appears to decrease when adding additional measurement neighbourhoods. However, there is a limit in the spatial distribution of optodes that can be laid out in regular-distanced measurement pairs. The described HD-DOT configuration is already challenging to implement in practice due to space constraints on the imaging pad and frequency/time multiplexing options. For this reason, denser arrays have not been examined in the literature, the most recent development being larger arrays for full head coverage. Considering this, if the maximum improvement in terms of localisation error, depth sampling and FWHM is required for a HD-DOT system, switching to FD seems to be a logical step.

It is important to highlight that utilizing a time-resolved system for NIRS, utilization of mean-time of photon travel has been previously reported[124], which is analogous to phase for an FD system, however in this work, the utilization of both log intensity and phase for tomographic reconstructions have been highlighted to demonstrate benefits. Although the sensitivity of the log intensity is shown to be more weighted towards the superficial tissues, it is this variation of sensitivity between log intensity and phase data which is providing the additional accuracy as demonstrated.

5 Frequency-Domain High-Density Diffuse Optical Tomography: *In Vivo* Retinotopy Experiments

5.1 Introduction

Retinotopy is the mapping of neural activations on the visual cortex occurring during visual stimuli. The retinotopic maps are in the primary sensory region of the brain, therefore is easier to understand the correlation between stimuli and brain activity than in the case of higher mental functions. The paradigm of retinotopy has been used with PET[125] and fMRI[126] and has been used for neuroimaging quality validation in a simultaneous fMRI and fDOT study, which revealed bilateral visual cortex activation pattern for both approaches[127]. This result was improved using HD-DOT, which allowed recoveries at four distinct quadrants[44]. Finally, a high correlation between fMRI and CW-HD-fDOT results was observed in studies employing subject-specific[114] and atlas-based[128] models for the DOT reconstruction. In this chapter, the computational tools presented in chapter 3 and the FD-HD-fDOT methodology presented in chapter 4 will be employed to perform all the reconstructions.

5.2 Experimental Setup

5.2.1 System set-up

A frequency-domain NIRS (FD NIRS) device (IMAGENT™, ISS Inc. Illinois USA) was used to obtain the NIRS-based data. The system consists of 32 source and 30 detector fibres with each source coupled to laser diodes emitting at 690 nm and 830 nm with the detectors being compact and fast photomultiplier tubes (PMT). The light sources were modulated at different frequencies, each offset a few Kilohertz from the central modulation frequency of 140 MHz. The signal was collected by the detectors modulated at a fifth frequency (cross-correlation frequency). Using this strategy, two parameters were recorded at each detector, for each wavelength, the AC amplitude and

phase-shift of the signal, as depicted in Figure 1.2. In contrast to a continuous wave system (CW), where the sources are at a constant amplitude (I_{in}) and only the attenuated signal (I_{out}) is detected (red lines in Figure 1.2), FD allows the measurement of both the AC amplitude (analogous to CW intensity) as well as the associated phase-shift ($\Delta\theta$) for each source/detector pair. 15 subjects were recruited from the University of Birmingham community, and written informed consent was obtained, none of these subjects were the same as the 25 subjects used for mesh generation earlier in this work. Subjects were seated facing an adjustable screen while the imaging pad was attached over the occipital cortex with hook and loop strapping. Scans were in line with a previously reported CW high-density DOT imaging system setup[44]. Briefly, the imaging cap consisted of 24 source positions (with two near-infrared wavelengths—690 nm and 830 nm—at each position) and 28 detectors interleaved in a high-density array. First- (13 mm, NN1), through fourth nearest neighbour (48 mm, NN4) optode pairs were utilised (Figure 5.1), giving a total of 348 total possible measurements, at a frame rate of ~40 Hz.

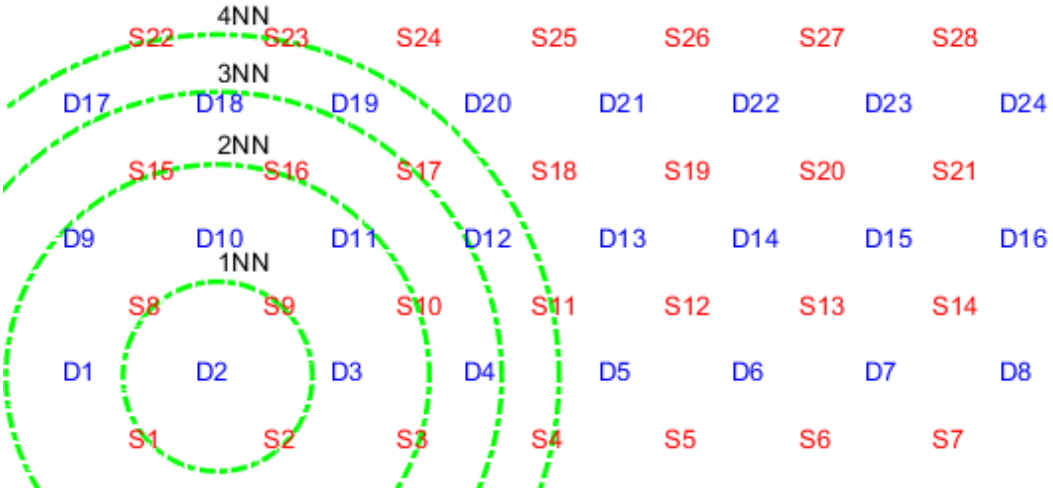

Figure 5.1 The optical cortex pad, with 28 sources(red) and 24 detectors (blue), featuring the 1st to 4th measurement areas for detector D2.

5.2.2 Retinotopy experiment

The human visual system is organised such that neighbouring neurons correspond to nearby locations of the visual field, in a continuous manner. In order to examine this mapping between the visual field and the optical cortex functional areas, retinotopy experiments have been designed, relying on carefully designed periodic visual stimuli.

In this study visual stimuli consisted of rotating logarithmic checkerboards, reversing intensity at 10 Hz, on a 50% intensity grey background. Each rotation completed in 36 seconds, corresponding to an event frequency of ~ 0.0277 Hz, and 10 rotations occurred in each run, resulting in 7 minutes total scanning time per run, including 30 second baseline periods at the beginning and the end. Typically, the subjects were comfortable to go through 3 runs (around 21 minutes) before taking a break.

5.2.3 Eye blinking experiment

This experiment was designed in order to confirm that the measured intensity fluctuations along the visual pad were indeed brain activations and not a product of screen intensity. Therefore, the screen was set on 50% intensity grey background, with no fluctuations, but the subjects were asked to open or close their eyes following a provided audio stimuli. Each event was completed in 34 seconds, 10 seconds of eyes open followed by 24 seconds with eyes closed, corresponding to an event frequency of ~ 0.0294 Hz, resulting in 6 minutes and 40 seconds total scanning time per run, including 30 second baseline periods at the beginning and the end. Similarly to the retinotopy experiment, the subjects were comfortable to repeat 3 times (around 20 minutes) before having a break.

5.3 *In vivo* fNIRS Data

5.3.1 Visualisation

The raw signals were visualised in vivo using NeuroDot[129], an extendable and user-friendly environment for analysis of optical data from raw light measurements. NeuroDot provides excellent quality control visualization options: the raw intensity signals, the signal strength, and noise across the imaging pad, the logarithmic intensity over the measurement distance and the frequency spectrum of the measured data (Figure 5.2).

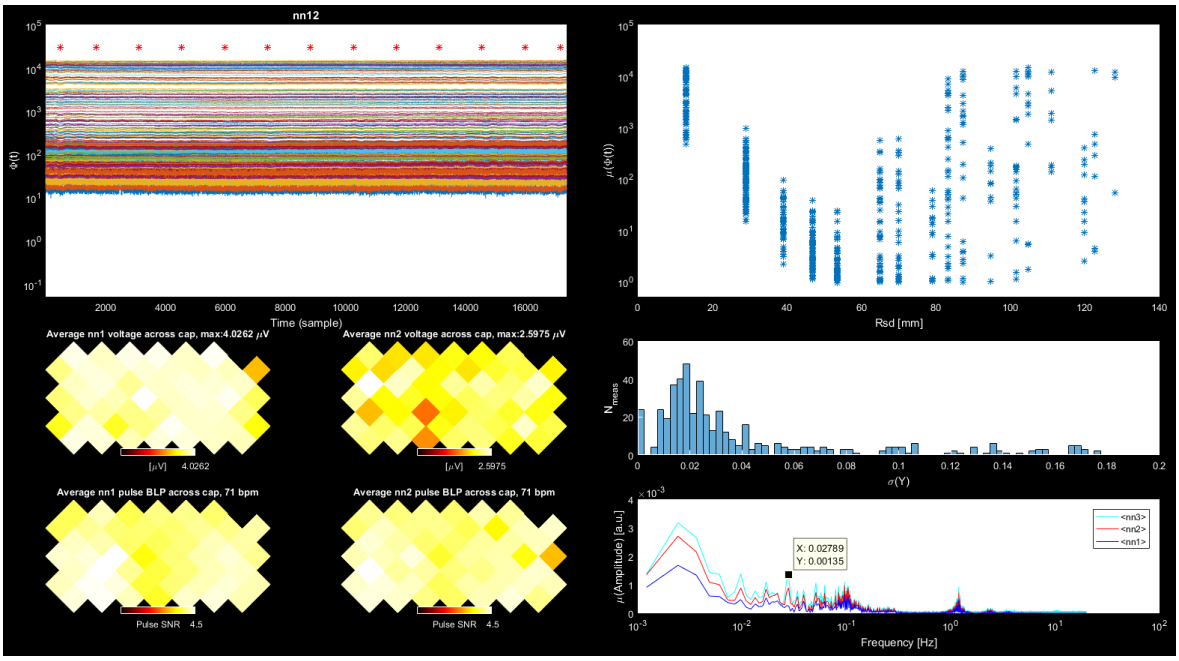

Figure 5.2 Screenshot of the NeuroDot analysis tool. In the top-left-hand, the raw intensity data is presented, with average signal intensity across the measurements pad for all source/detector pairs at the bottom left. The top right hand is the log-intensity as a function of source-detector separation, and the variance of the signal and power spectrum shown in the bottom right.

NeuroDOT was modified to provide similar visualisations for the measured phase data (

Figure 5.3). These visualisation tools were used to provide quality control during the data recording sessions, allowing for fast fitting of the optical array and ensuring high-quality data are collected.

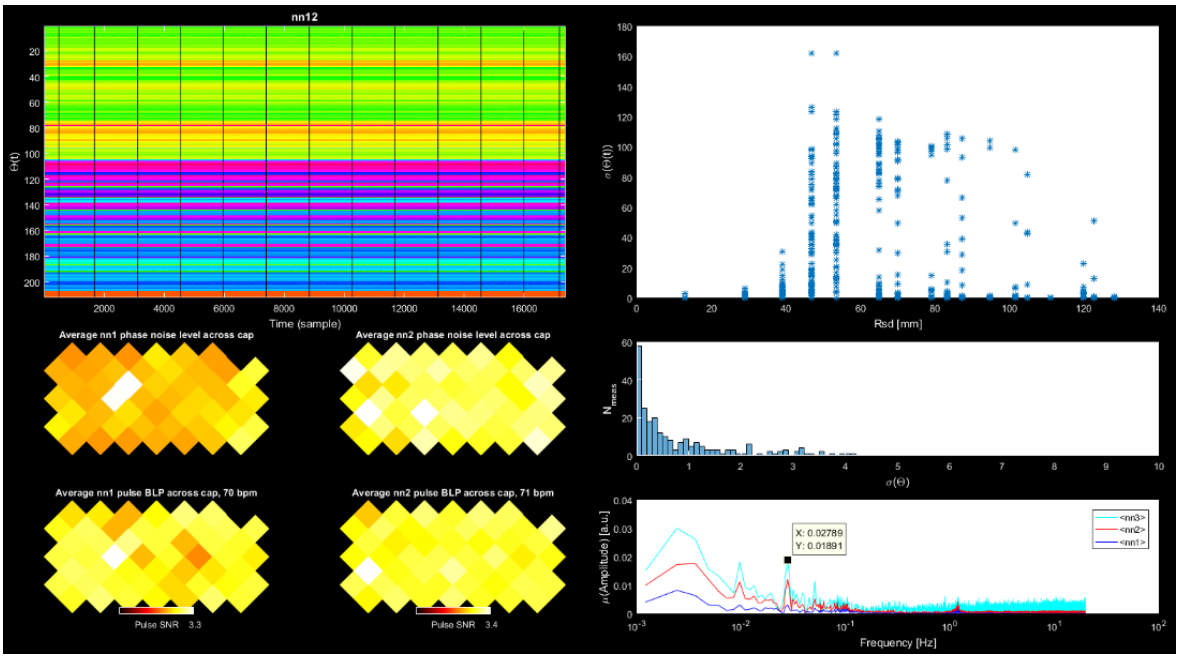

Figure 5.3 Data quality display for phase data. Top left: Each row represents the time course of the phase of a given source-detector measurement pair within the NN1 and NN2 sets. The actual mean phase value is arbitrary. Lower left: Spatial distribution of mean phase noise for NN1 (left) and NN2 (right) for each source and detector on the imaging array. Top right: Temporal standard deviation of the phase measurement as a function of source-detector distance. Middle right: Histogram of the temporal standard deviation of phase measurements. Lower right: Average power spectra for the NN1, NN2, and NN3 measurement sets with the frequency of the stimulus cycle highlighted.

In both visualisations, the peak near the stimuli frequency is clearly visible indicating the brain activity caused by the stimuli, which for the shown examples was the rotating pattern, with an event frequency of 0.278 Hz. A clear peak on the stimuli frequency was observed in 12 of the 15 participants.

5.3.2 Frequency Domain Spectrum

The frequency spectrum of the measured data can provide useful initial information regarding the quality of the data. Haemodynamic phenomena of regularly repeatable nature, for example heartbeat and breath, produce a peak in the frequency spectrum.

Prior to data analysis, measurements with high noise (amplitude standard deviation > 7.5%[114] or phase standard deviation >5 degrees) were disregarded, resulting in ~300 measurements passing the quality control for each scan. Each remaining measurement was subjected to a Fourier transform to acquire its frequency power spectrum. Power spectra for each measurement neighbourhood (e.g. NN1) were averaged separately for the amplitude and phase measurements.

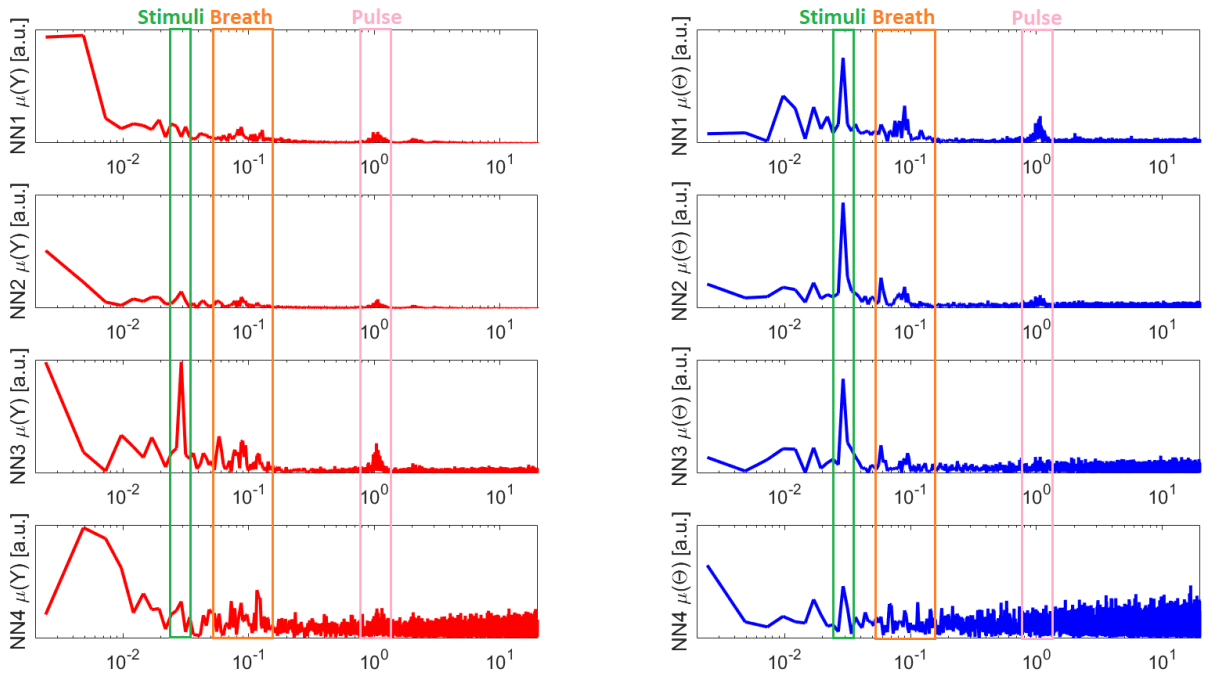

Figure 5.4 Fourier transforms of the time traces obtained using the FD system at 830 nm. The y-axis showing power (arbitrary units) x-axis depicting frequency (Hz). Each row corresponds to NN1-NN4 optode pairs. Left column (red plots) correspond to AC amplitude data, whereas the right column (blue plots) corresponds to Phase data. Note the strong peak at the rotation frequency of the stimulus wedge (~0.028Hz), as well as those from Breathing (~0.1 Hz) and Cardiac pulse (~1Hz).

Several features are evident from the power spectrum of the signals shown in Figure 5.4. Using AC amplitude data (intensity) the signal becomes weaker as the source/detector separation increases. The peak due to the stimuli is only evident for source/detector separations of 30mm or larger, as the lower separations are only

sampling the superficial tissues. The peaks due to breathing and the cardiac pulse are evident in all source-detector separations.

Using phase data, it is evident that the data contains more noise at the largest source/detector separation of 48 mm, which is primarily due to the strong signal attenuations at this longer distance. But more crucially, the peak due to the visual stimuli is evident in all source/detector measurements, even those at 1st NN of 13 mm, demonstrating the deeper sampling of brain tissue obtained by this phase data. The most substantial finding is that in the phase data, the peaks due to breathing and cardiac pulse are much smaller in all cases, almost diminishing for 2nd and 3rd NN measurements, demonstrating that the phase data is less sensitive to signal contamination from the superficial tissues.

5.3.3 Discussion

While a CW system measures only the attenuated light intensity, FD and Time-Resolved (TR) systems can measure the pathlength of the light as it has travelled through the tissue which allows for a more accurate quantitative assessment of its optical properties. The use of mean-time-of-flight of photons from TR systems has been reported and used for recovery of functional activations[124], but the information content of phase data has not been previously studied to demonstrate depth selective information content. It is shown that the intensity data for short optode distances do not contain information about brain activations but are corrupted by systematic noise (breathing and cardiac) from the superficial layers. However, the measured phase data demonstrate that brain activation signals are available at these short source/detector separations, while at longer distances, the contamination from superficial regions is minimised. This

demonstrates the benefit of FD fNIRS for quantitative and accurate detection of brain activations.

5.4 *In vivo* reconstructions

The raw signals were processed according to previously proposed pipeline[80], with a few modifications to accommodate for the phase data. For each channel the passed the noise threshold, the amplitude data was log-meaned, $\ln\left(\frac{y_I}{y_{I_0}}\right)$, and the phase data was mean abstracted $y_\Theta - y_{\Theta_0}$, to create differential data according to the Rytov approximation (Eq. 4.2), where y_{I_0} and y_{Θ_0} are the mean values for amplitude and phase for this channel respectively. Initially, signal quality was assured using a previously reported[80] noise threshold of 7.5% standard deviation for amplitude. Since this was the first time FD-HD-fDOT tomographic reconstructions were performed, a 5 degrees standard deviation threshold for phase was selected through trial and error, aiming to reduce noise in reconstructed images.

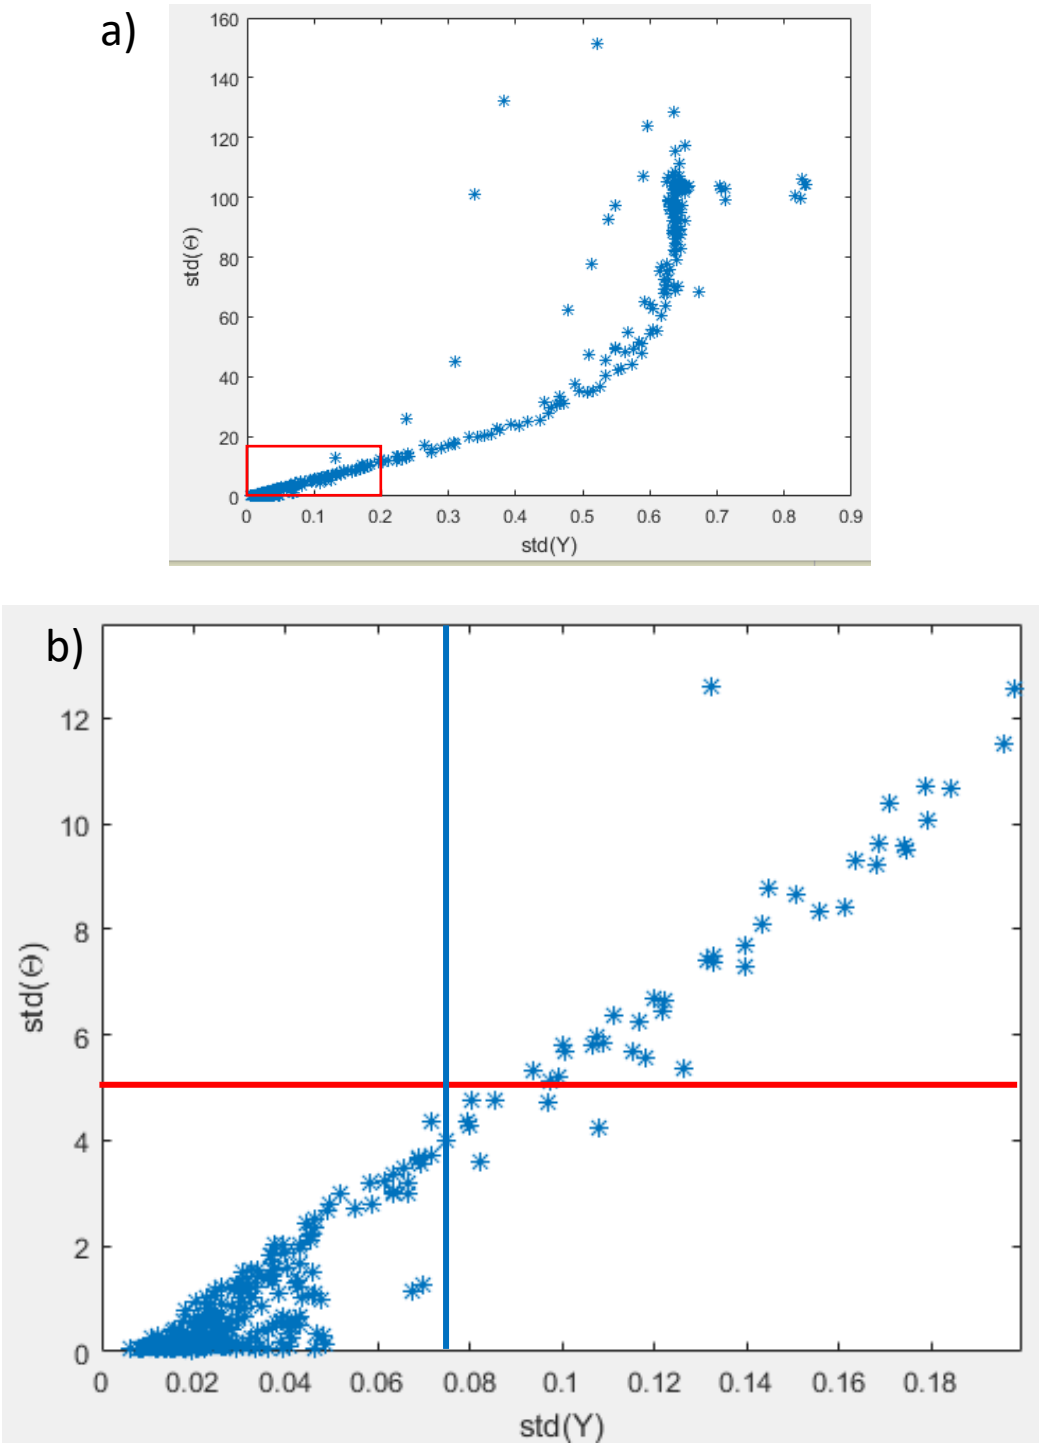

Figure 5.5 Standard deviation of amplitude versus standard deviation for phase for each measurement. Part b is a zoomed-in at the red-boxed area of a. In part b the red line denotes the 5 degrees and the blue line the 7.5% noise thresholds.

Only the channels that score noise under 5 degrees for phase and, at the same time, under 7.5% for amplitude were considered for further processing, these were the measurements lying within the area defined by the red and blue lines in Figure 5.5 (b).

The frequency spectrum of the differential data, averaged for each measurement neighbourhood, is displayed in Figure 5.6. The number of channels below the noise thresholds, N, and the total number of channels for each neighbourhood is noted on top of each graph. The top row has NN1 on the left and NN2 on the right, while the bottom row has NN3 and NN4 respectively. The noise, for both amplitude and phase, is increasing as source-detector separation increases, therefore fewer measurements are passing the quality thresholds at more distant measurement neighbourhoods (Figure 5.6). This is also reflected in the noise model estimated for section 4.1.4.

Amplitude Data

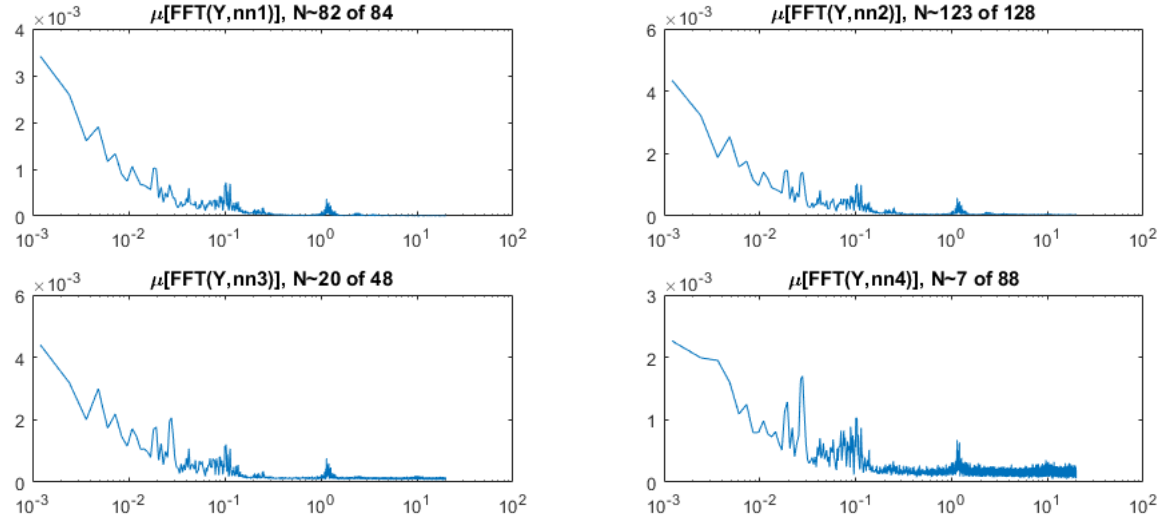

Phase Data

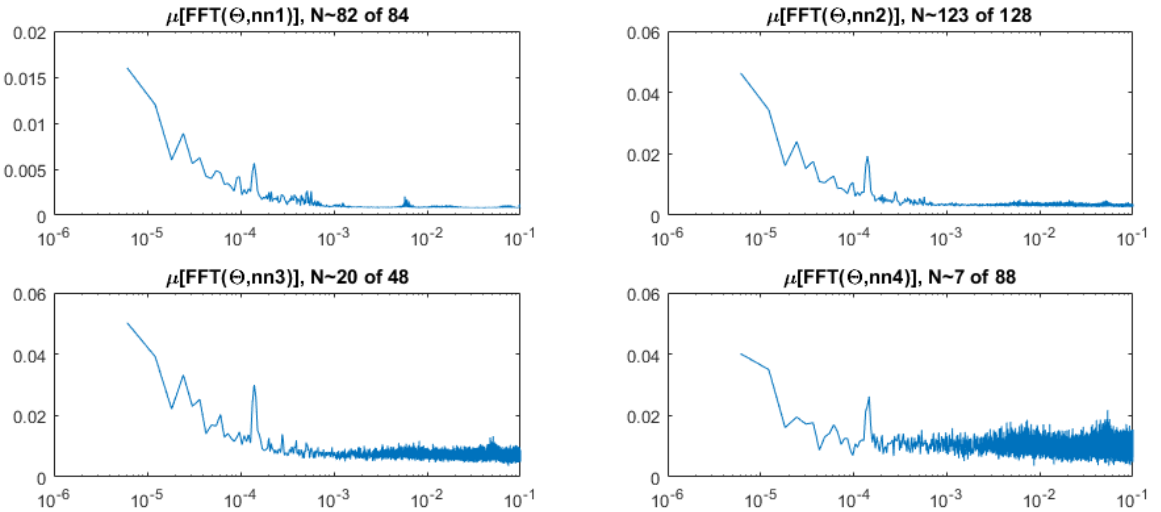

Figure 5.6 Differential amplitude and phase data frequency spectrum, averaged for each measurement neighbourhood. The number of channels passing the noise thresholds, N, and the total number of channels for each neighbourhood is noted on top of each graph. The top row has NN1 on the left and NN2 on the right, while the bottom row has NN3 and NN4 respectively.

Then the data was band-pass filtered to 0.01-0.1 Hz to remove slow drifts and systemic physiology noise as heartbeat and respiration. To implement this filtering, first a zero-phase digital high-pass filter at 0.01 Hz was applied (Figure 5.7), then a zero-phase digital lowpass filter at 0.1 Hz was applied on the outcome (Figure 5.8).

Amplitude data

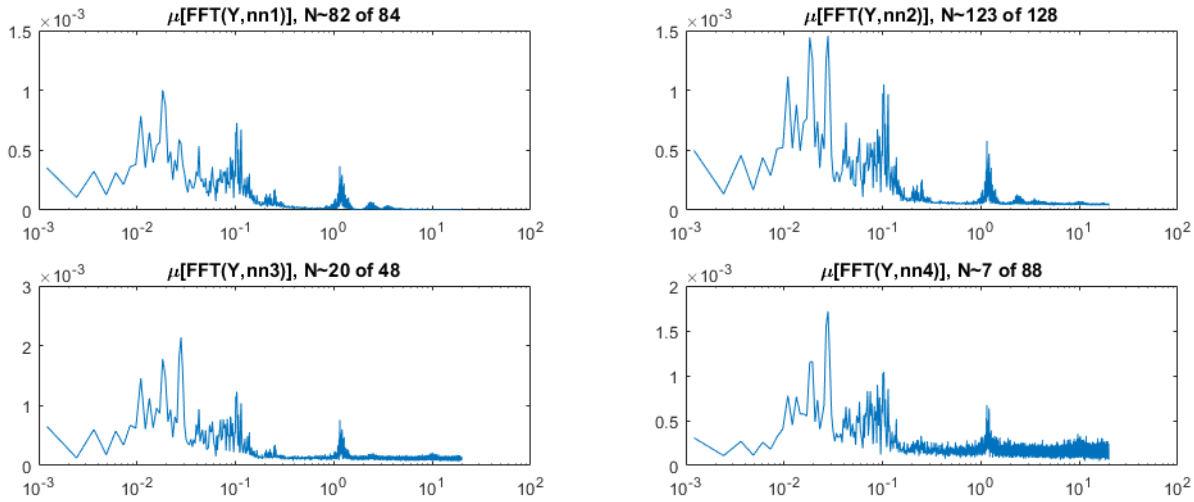

Phase data

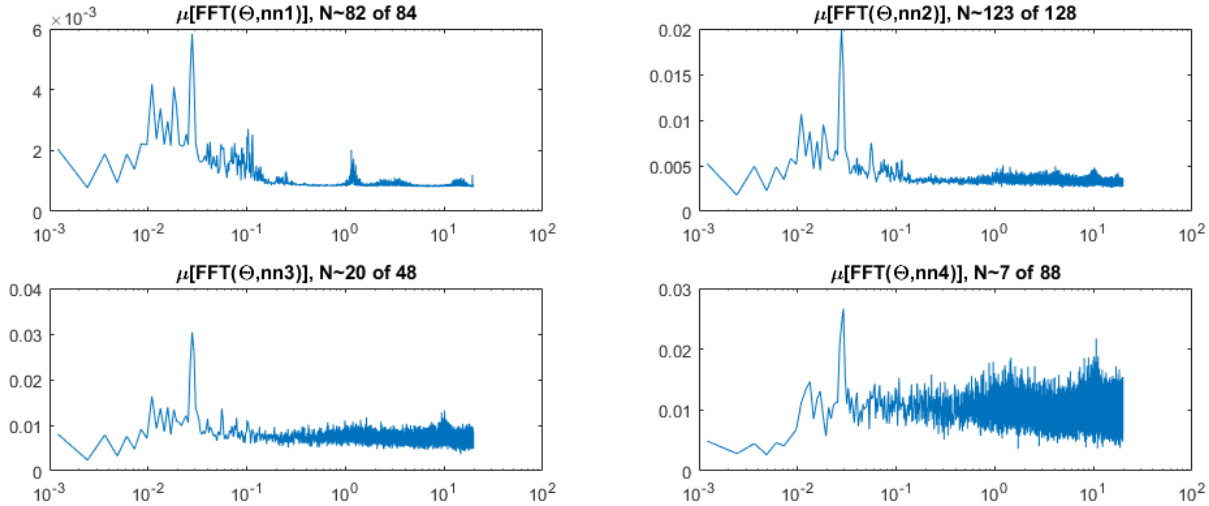

Figure 5.7 High-passed differential amplitude and phase data frequency spectrum, averaged for each measurement neighbourhood.

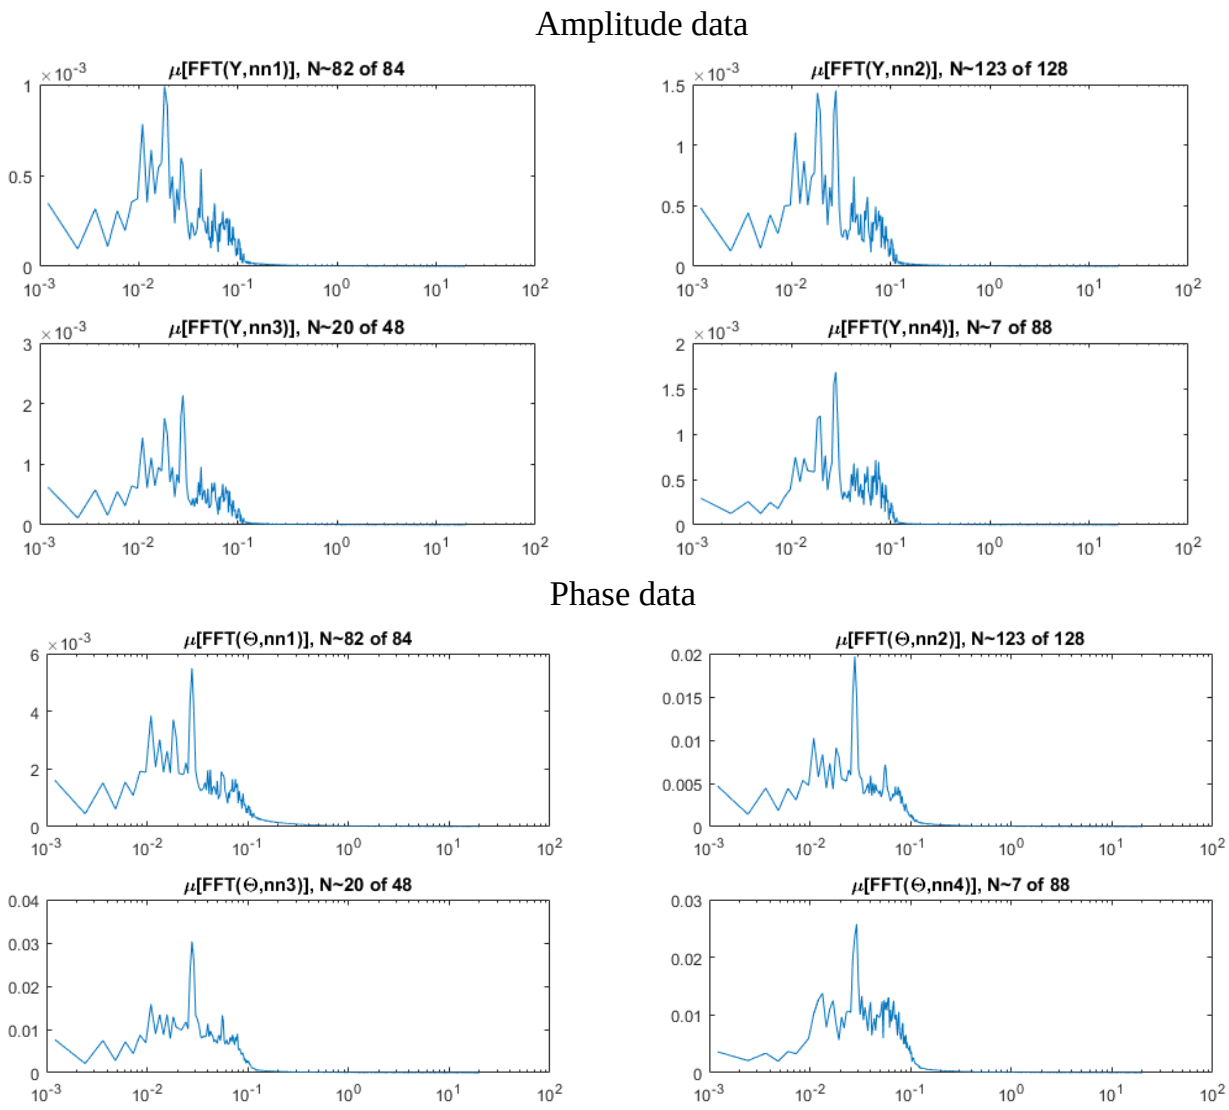

Figure 5.8 Low-passed differential amplitude and phase data frequency spectrum, averaged for each measurement neighbourhood.

Finally, the data was down-sampled to 1 Hz and block averaged in ten blocks, according to the ten repetition cycles of the visual stimuli, of 36 seconds each.

5.4.1 Retinotopy

A single snapshot of a recovered focal activation for a single subject is shown in Figure 5.9. Here the cross-sectional sagittal and transverse planes of a single focal activation within the left visual cortex are shown as a change in recovered HbO_2. The CW reconstruction is shown to recover an activation at a much more superficial depth, mainly appearing in the area of the superficial tissue; whereas the FD based

reconstruction is able to recover the activation overlapping with the cortex of the brain. It is important to highlight that the same regularisation parameter for the inverse model has been used, namely ensuring that the improvement in image recovery and accuracy is primarily due to the inclusion of phase data within the FD based reconstruction.

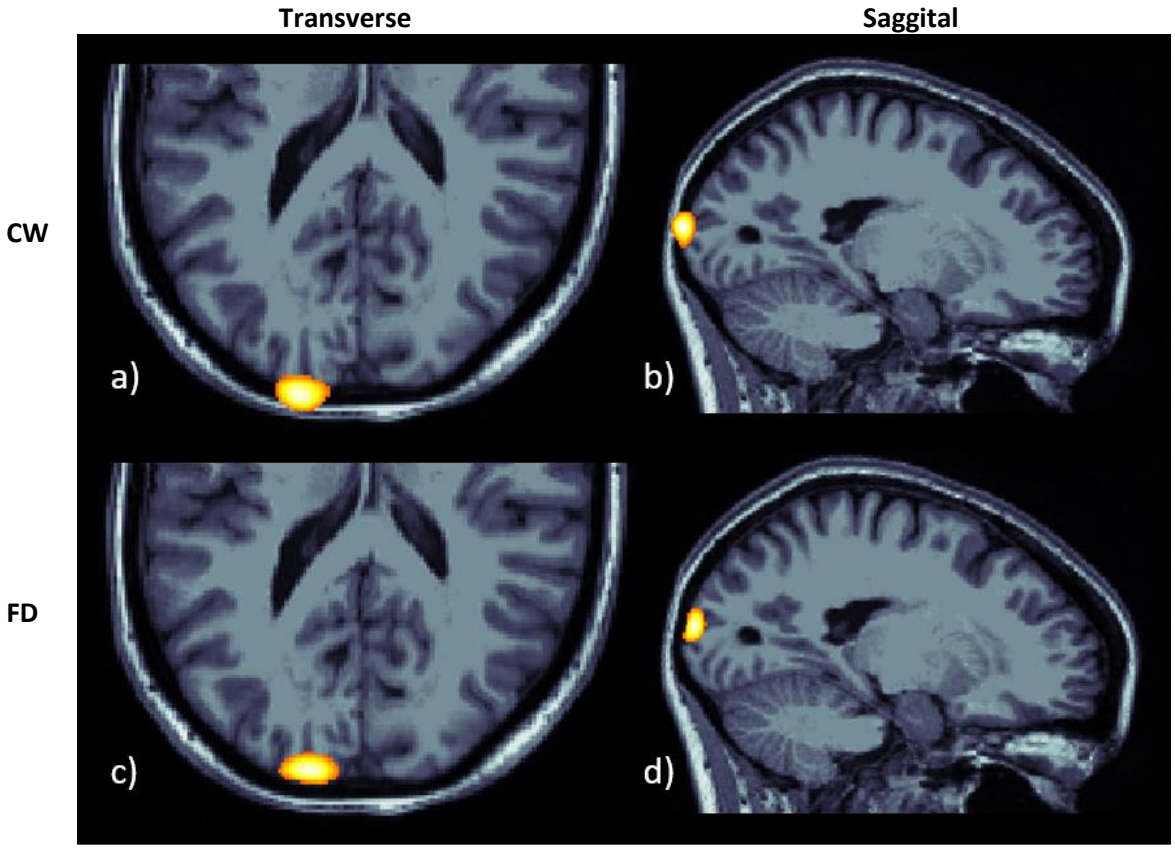

Figure 5.9 In-vivo visual cortex activation (50% threshold, ΔHbO_2 [a.u.]) using CW data (a,b) and FD data (c,d).

5.4.2 Phase encoded retinotopy

Phase encoded retinotopy is a technique whereby tomographic reconstruction is performed and processed such that only nodes that have maximum recovered peak at a given stimuli frequency are selected. Therefore, the reconstruction is confined spatially only for nodes that have been affected by the visual stimuli and produce strong variance over time for the given stimuli frequency (Figure 5.10): this simultaneously ignores all

information from other frequency bins, dramatically reducing noise and hence imaging artifact.

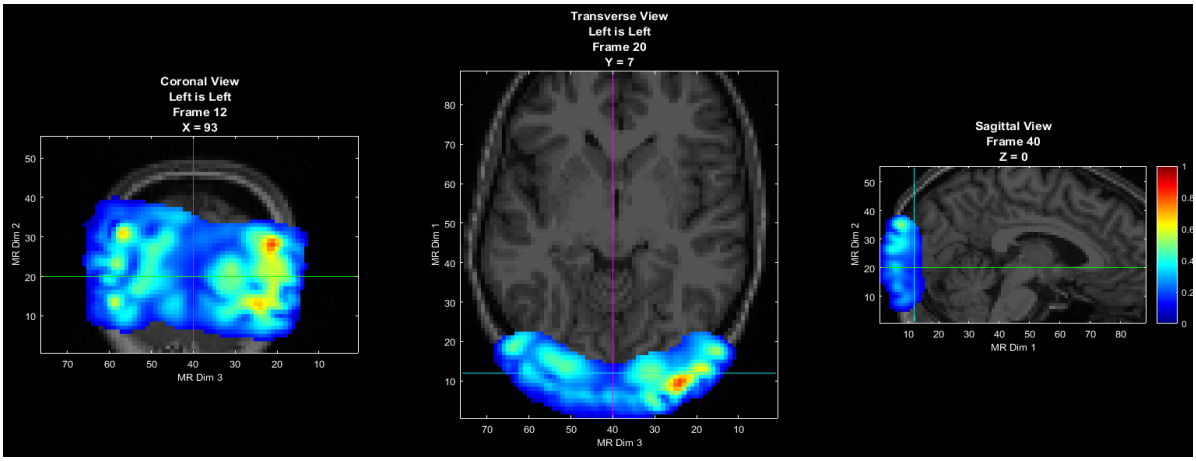

Figure 5.10 HbO_2 recovery (normalised) variance over time only for the stimuli frequency.

Specifically, the recovered HbO_2 values for each node are considered through the temporal dimension, and the frequency spectrum of this time trace is obtained. If the node presents a significant peak at the stimuli frequency, the phase of this frequency bin, which is directly related to the time offset the HbO_2 variance peak, is examined. For visualization, in order to represent the angles in which the stimuli wedge is rotating through, each angle is color-coded (Figure 5.11). Ideally, each of the colours in the stimuli wheel represent a different angle of the checkerboard wedge as it rotates, typically resulting in 36 colours (in steps of 10 degrees), until it completes a full circle. The nodes on the surface of the brain should activate with a specific time offset that will also classify them into similarly arranged colours. This is because spatially adjacent visual stimuli should activate spatially adjacent neurons. In the example as shown in Figure 5.11 the rotational bins have been quantised only into 16 bins to make the effect visually clearer.

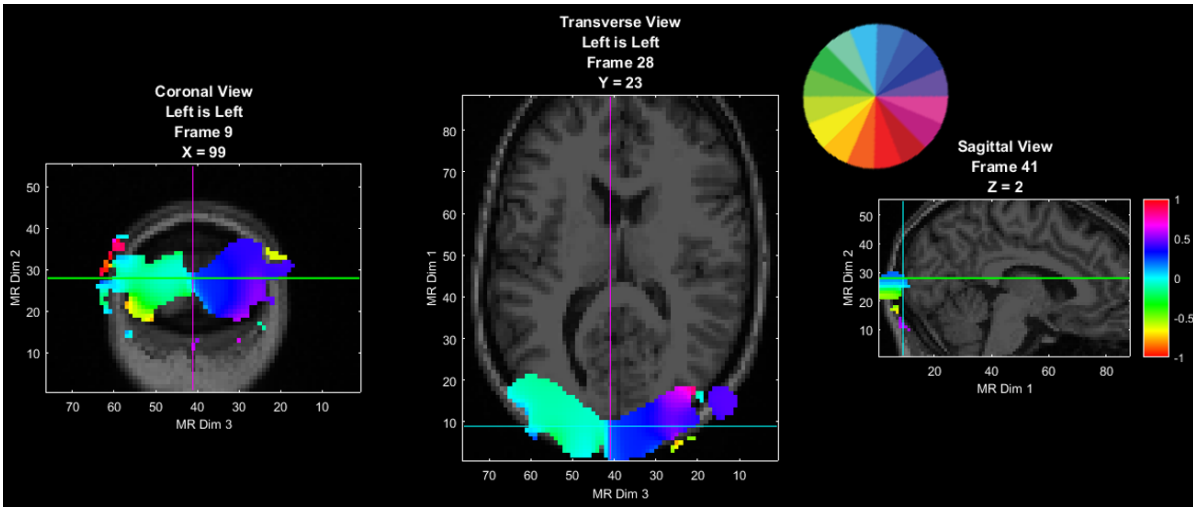

Figure 5.11 Phase encoded retinotopy. Each colour represents a different angle on the rotation of the wedge stimuli.

We were able to achieve good quality phase-encoded retinotopy for CW and FD based reconstructions. The FD reconstructions achieved more realistic localization of the activations, however, they were noisier. This is primarily due to the fact that phase appears to be a more noisy measurement, with noise levels drastically increasing when source-detector distance increases. For the PER analysis, in all of the subjects where a peak was detected on the stimuli frequency (12 subjects), there was a clear distinction between left and right hemisphere activations, in most subjects (9 subjects) quadrants could be resolved. However only 4 subjects had PER that captured a fuller range of activations, similar to Figure 5.11. The failure cases are probably due to poor imaging array placement, in terms of location over the optical lobe and optode contact to the head surface.

5.4.3 Discussion and conclusions

The inclusion of phase data improves the reconstruction quality by localising the activations more appropriately on the cortical surface rather than within superficial layers, as shown in Figure 5.9. As the purpose of the work was the proof of concept that FD-HD-DOT outperforms CW-HD-DOT in reconstructed image quality, there was no

fine-tuning of parameters as data filtering and regularisation; furthermore, no superficial signal regression and no depth normalisation technique was employed, to ensure that all the improvements are due to the enhanced nature of measurements. The results indicate that there is much room for improvement of HD-DOT, only by switching to FD measurements, a possible even further improvements if the reconstruction parameters are fine-tuned to accommodate phase data.

6 Conclusion and Future Work

DOT is a promising imaging modality, steadily gaining ground amongst the established imaging techniques. The harmless and patient-friendly procedure enables use in applications where other techniques are inadequate. This work presents a framework for frequency domain high density functional diffuse optical tomography, whilst providing computational tools to enable fast and computationally efficient reconstructions.

The main motivation for the development of new computational tools was to allow researchers to perform HD-DOT reconstructions, with CW and FD measurements, whilst solving the forward and inverse problem in very high-resolution FEM meshes, while using a standard desktop/laptop computer. In this work, the accuracy and computational speed of iterative solvers is evaluated on anatomically realistic head models, with high-resolution meshes and HD-DOT system, the solver's termination criterion that was found to be optimal for this application is 10^{-12}. The produced tools require only 10% of the time in comparison to the tools available up to now; additionally, these tools can scale to take advantage of high-end computing facilities, where available. Additionally, a fast and memory efficient formulation of the inverse problem is demonstrated. Creating a sparse Jacobian massively reduces memory requirements, allowing reconstructions in high-resolution mesh basis, whilst employing numerous frequency domain measurements in multiple wavelengths. The selected sensitivity threshold of 0.1% allows nodes that reach at least the brain cortex depth to contribute towards reconstructions for all measurements.

Furthermore, the memory-efficient handling of the forward and inverse models enable high-quality reconstructions in scenarios where it was impossible up to now due to lack of computational power, for example as a bedside monitoring tool or as a piece of ambulance or operating theatre equipment.

The inclusion of phase data, alongside the amplitude measurements, enhances the reconstruction quality as while the sensitivity for both log-amplitude and phase data increases in depth of sampling with source-detector separation, the phase data demonstrates deeper sensitivity for a given source-detector separation; and the hyper-sensitivity in superficial tissues observed with log amplitude data is less pronounced for phase data.

A framework for FD-HD-fDOT was proposed and evaluated in simulations, revealing a significant improvement in imaging quality, achieving a reduction of localisation error by up to 59% and an improvement of effective resolution by up to 21%, as compared to using the intensity attenuation measurements alone. Also, the simulations display sampling at larger depth when phase data is employed, allowing DOT to resolve functional events in deeper brain regions.

The benefit of FD measurements seems to be greater than the benefit from more dense measurements using CW alone, however, this improvement appears to decrease when adding additional measurement neighbourhoods. There is a limit in the spatial distribution of optodes that can be laid out in regular-distanced measurement pairs. HD-DOT configuration is already challenging to implement in practice due to space constraints on the imaging pad and frequency/time multiplexing options, therefore denser arrays have not been examined in the literature. The configuration of a FD

imaging array if of the same complexity as an CW array. Considering this, if the maximum imaging quality is required for a HD-DOT system, the transition to FD measurements should be pursued.

Finally, *in vivo* retinotopic experiments were performed that verified that phase can provide additional useful information in fNIRS, and subsequently improve fDOT imaging quality. Specifically the FD reconstructions appear to located the brain activation in a more appropriate position on the brain cortex, while containing information about brain activity even in very short distance measurements, due to increase depth penetration. However, in practice, phase is a more noisy measurement, therefore, whilst improving localisation, often added noise to the reconstruction. This is a current drawback of the FD approach, but as is mainly a hardware limitation, it could be overcome with finetuning the source/detector gain and multiplexing settings, and the optodes placement and contact to the head surface.

This work demonstrates that phase data is a valuable measurement for fDOT, however it is only an initial study that proves the concept and provide the foundations for further research. For example, in this work the same regularisation factor was used for amplitude and phase recoveries, no regional sensitivity weighting, and no depth normalization was used in any FD reconstruction. Fine-tuning such parameters could further increase the imaging quality of HD-DOT.

In addition, some unique characteristics of the phase data, can lead to development of new approaches. For example, while phase data capture brain-related activity, is not as sensitive as the amplitude data to systemic noise, like breathing and heartbeat. This

discrepancy could be exploited to suppress those noise sources in the whole of the reconstruction and further increase functional imaging quality.

Whilst this work focused only on the recovery of relative absorption changes, that mainly reflect oxy and deoxy haemoglobin concentrations, the availability of phase measurements can provide information on scattering related changes and potentially provide absolute values for the recovered optical properties, which could be a useful tool in many clinical applications.

Although it is accepted that CW systems are simpler and cheaper to develop as compared to FD systems, the information content as outlined above may provide the motivation for future applications and system development. CW systems rely on cheaper components and electronics such as laser diodes and avalanche photodiode and system calibration may be considered easier, particularly when concerned with difference imaging, when detecting changes in intensity of the signal due to any systematic offset will cancel. FD systems however rely on more expensive and bulky components such as amplitude modulated laser diodes, signal generators, lock-in amplifiers, and photomultiplier tubes (PMT). Although the cost and size of these components have reduced over the years, system calibration and characterisation becomes more challenging, particularly with regards to the achievable dynamic range of PMTs, trade-off between signal detectability and saturation, as well as potentially reduced phase SNR for large source/detector separation. These factors should be taken into consideration, together with the potential gains as outlined in this work. However, what has been demonstrated here, is that the availability of this additional phase data is no longer a bottleneck when considering model-based parameter recovery whilst it provides

detectable functional contrast at better depth within the brain as compared to current CW based fNIRS systems.

Bibliography

1. Sanfey AG. The Neural Basis of Economic Decision-Making in the Ultimatum Game. *Science (80-)*. 2003;1755. doi:10.1126/science.1082976

2. Walter H, Abler B, Ciaramidaro A, Erk S. Motivating forces of human actions Neuroimaging reward and social interaction. *Brain Res Bull*. 2005;67:368-381. doi:10.1016/j.brainresbull.2005.06.016

3. Singer T. The neuronal basis of empathy. *Empathy Fairness Novartis Found Symp*. 2006;278:20-40.

4. Johnston S, Linden DEJ, Healy D. Upregulation of emotion areas through neurofeedback with a focus on positive mood. *Cogn Affect Behav Neurosci*. 2011:44-51. doi:10.3758/s13415-010-0010-1

5. Sitaram R, Lee S, Ruiz S, Rana M, Veit R, Birbaumer N. Real-time support vector classi fi cation and feedback of multiple emotional brain states. *Neuroimage*. 2011;56(2):753-765. doi:10.1016/j.neuroimage.2010.08.007

6. Brigadoi S, Cutini S, Scarpa F, Scatturin P, Dell'Acqua R. Exploring the role of primary and supplementary motor areas in simple motor tasks with fNIRS. *Cogn Process*. 2012;13(1 SUPPL):97-101. doi:10.1007/s10339-012-0446-z

7. Decety J, Perani D, Jeannerod M, et al. Mapping motor representations with positron emission tomography. *Nature*. 1994;371(6498):600-602. doi:10.1038/371600a0

8. Craig AD. How do you feel? Interoception: the sense of the physiological condition of the body. *Nat Rev | Neurosci*. 2002;3(August).

9. Sitaram R, Veit R, Stevens B, et al. Acquired Control of Ventral Premotor Cortex Activity by Feedback Training: An Exploratory Real-Time fMRI and TMS Study. *Neurorehabil Neural Repair*. 2012. doi:10.1177/1545968311418345

10. Ruiz S, Lee S, Soekadar SR, et al. Acquired Self-control of Insula Cortex Modulates Emotion Recognition and Brain Network Connectivity in Schizophrenia. *Wiley Period Inc Hum Brain Mapp*. 2011;000(July). doi:10.1002/hbm.21427

11. Tangermann M, Dornhege G, Krauledat M, Curio G, Blankertz B. Machine learning for real-time single-trial EEG-analysis : From brain – computer interfacing to mental state monitoring. *J Neurosci Methods*. 2008;167:82-90. doi:10.1016/j.jneumeth.2007.09.022

12. Coyle SM, Ward E, Markham CM. Brain – computer interface using a simplified functional near-infrared spectroscopy system. *J Neural Eng*. 2007;219. doi:10.1088/1741-2560/4/3/007

13. Mason SG, Bohringer R, Borisoff JF, Birch GE. Real-Time Control of a Video Game With a Direct Brain – Computer Interface. *J Clin Neurophysiol*. 2004;21(6):404-408.

14. Weiskopf N, Mathiak K, Bock SW, et al. Principles of a Brain-Computer Interface (BCI) Based on Real-Time Functional Magnetic Resonance Imaging (fMRI). *IEEE Trans Biomed Eng*. 2004;51(6):966-970.

15. Wheelock MD, Eggebrecht AT, Culver JP. High-density diffuse optical tomography for imaging human brain function. *Rev Sci Instrum*. 2019;051101(May). doi:10.1063/1.5086809

16. Crone NE, Miglioretti DL, Gordon B, Lesser RP. Functional mapping of human sensorimotor

cortex with electrocorticographic spectral analysis. II. Event-related synchronization in the gamma band. *Brain*. 1998;121(12):2301-2315. doi:10.1093/brain/121.12.2301

17. Waldert S, Pistohl T, Braun C, Ball T, Aertsen A, Mehring C. A review on directional information in neural signals for brain-machine interfaces. *J Physiol Paris*. 2009;103(3-5):244-254. doi:10.1016/j.jphysparis.2009.08.007

18. Neil Cuffin B, Cohen D. Comparison of the magnetoencephalogram and electroencephalogram. *Electroencephalogr Clin Neurophysiol*. 1979;47(2):132-146. doi:10.1016/0013-4694(79)90215-3

19. Lux HD. Neurophysiological Basis of the EEG and DC Potentials. *Electroencephalogr Basic Princ Clin Appl Relat Fields*. 1981:24-28.

20. Sharon D, Hämäläinen MS, Tootell RBH, Halgren E, Belliveau JW. The advantage of combining MEG and EEG: Comparison to fMRI in focally stimulated visual cortex. *Neuroimage*. 2007;36(4):1225-1235. doi:10.1016/j.neuroimage.2007.03.066

21. Strangman G, Culver JP, Thompson JH, Boas DA. A quantitative comparison of simultaneous BOLD fMRI and NIRS recordings during functional brain activation. *Neuroimage*. 2002;17(2):719-731. doi:10.1016/S1053-8119(02)91227-9

22. MacIntosh BJ, Klassen LM, Menon RS. Transient hemodynamics during a breath hold challenge in a two part functional imaging study with simultaneous near-infrared spectroscopy in adult humans. *Neuroimage*. 2003;20(2):1246-1252. doi:10.1016/S1053-8119(03)00417-8

23. Hoshi Y, Kobayashi N, Tamura M. Interpretation of near-infrared spectroscopy signals: a study with a newly developed perfused rat brain model. *J Appl Physiol*. 2001;90(5):1657-1662. http://www.ncbi.nlm.nih.gov/pubmed/11299252.

24. Siegel AM, Culver JP, Mandeville JB, Boas D a. Temporal comparison of functional brain imaging with diffuse optical tomography and fMRI during rat forepaw stimulation. *Phys Med Biol*. 2003;48(10):1391-1403. doi:10.1088/0031-9155/48/10/311

25. Boss A, Bisdas S, Kolb A, et al. Hybrid PET/MRI of intracranial masses: initial experiences and comparison to PET/CT. *J Nucl Med*. 2010;51(8):1198-1205. doi:10.2967/jnumed.110.074773

26. Cooper RJ, Magee E, Everdell N, et al. MONSTIR II: A 32-channel, multispectral, time-resolved optical tomography system for neonatal brain imaging. *Rev Sci Instrum*. 2014;85(5):053105. doi:10.1063/1.4875593

27. Leary MAO, Boas DA, Chance B, Yodh AG. Experimental images of heterogeneous turbid media by frequency-domain diffusing-photon tomography. 1995;20(5):426-428.

28. Srinivasan S, Ph D, Pogue BW, et al. Near-Infrared Characterization of Breast Tumors In Vivo using Spectrally-Constrained Reconstruction. 2005;4(5).

29. Hielscher AH, Kim HK, Montejo LD, et al. Frequency-domain optical tomographic imaging of arthritic finger joints. *IEEE Trans Med Imaging*. 2011;30(10):1725-1736. doi:10.1109/TMI.2011.2135374

30. Bortfeld H, Diamond SG, Boas DA, Arvin K, Grant PE, Hospital MG. Assessment of Infant Brain Development With Frequency-Domain Near-Infrared Spectroscopy. 2007;61(5):546-551. doi:10.1203/pdr.0b013e318045be99

31. Ghosh A, Tachtsidis I, Kolyva C, Highton D, Elwell C, Smith M. Normobaric Hyperoxia Does Not Change Optical Scattering or Pathlength but Does Increase Oxidised Cytochrome c Oxidase

Concentration in Patients with Brain Injury. In: Springer, New York, NY; 2013:67-72. doi:10.1007/978-1-4614-4989-8_10

32. Culver JP, Durduran T, Furuya D, Cheung C, Greenberg JH, Yodh AG. Diffuse optical tomography of cerebral blood flow, oxygenation, and metabolism in rat during focal ischemia. *J Cereb Blood Flow Metab*. 2003;23(8):911-924. doi:10.1097/01.WCB.0000076703.71231.BB

33. Power SD, Falk TH, Chau T. Classification of prefrontal activity due to mental arithmetic and music imagery using hidden Markov models and frequency domain near-infrared spectroscopy. *J Neural Eng*. 2010;7(2):026002. doi:10.1088/1741-2560/7/2/026002

34. Zimmermann R, Marchal-Crespo L, Edelmann J, et al. Detection of motor execution using a hybrid fNIRS-biosignal BCI: a feasibility study. *J Neuroeng Rehabil*. 2013;10(1):4. doi:10.1186/1743-0003-10-4

35. Scholkmann F, Klein SD, Gerber U, Wolf M, Wolf U. Cerebral hemodynamic and oxygenation changes induced by inner and heard speech: a study combining functional near-infrared spectroscopy and capnography. *J Biomed Opt*. 2014;19(1):017002. doi:10.1117/1.JBO.19.1.017002

36. Choi J, Wolf M, Toronov V, et al. Noninvasive determination of the optical properties of adult brain: near-infrared spectroscopy approach. *J Biomed Opt*. 2004;9(1):221. doi:10.1117/1.1628242

37. Powell S, Dempsey LA, Cooper RJ, Gibson AP, Hebden JC, Arridge SR. Real-time dynamic image reconstruction in time-domain diffuse optical tomography. In: *Biomedical Optics 2016*. Washington, D.C.: OSA; 2016:OM4C.5. doi:10.1364/OTS.2016.OM4C.5

38. Dempsey LA, Brigadoi S, Cooper RJ, et al. Whole-head functional brain imaging of neonates at cot-side using time-resolved diffuse optical tomography. *Diffus Opt Imaging V*. 2015;9538(0):953818. doi:10.1117/12.2190420

39. Gratton G, Corballis PM, Cho E, Fabiani M, Hood DC. Shades of gray matter: noninvasive optical images of human brain responses during visual stimulation. *Psychophysiology*. 1995;32(5):505-509. doi:10.1111/j.1469-8986.1995.tb02102.x

40. Gratton G, Fabiani M. The event-related optical signal (EROS) in visual cortex : Replicability , consistency , localization , and resolution. 2003;40:561-571.

41. Toronov V, Webb a, Choi JH, et al. Study of local cerebral hemodynamics by frequency-domain near-infrared spectroscopy and correlation with simultaneously acquired functional magnetic resonance imaging. *Opt Express*. 2001;9(8):417-427. doi:65396 [pii]

42. Martelli F, Binzoni T, Pifferi A, Spinelli L, Farina A, Torricelli A. There's plenty of light at the bottom: Statistics of photon penetration depth in random media. *Sci Rep*. 2016;6:1-14. doi:10.1038/srep27057

43. Binzoni T, Sassaroli A, Torricelli A, et al. Depth sensitivity of frequency domain optical measurements in diffusive media. *Biomed Opt Express*. 2017;8(6):2990. doi:10.1364/BOE.8.002990

44. Zeff BW, White BR, Dehghani H, Schlaggar BL, Culver JP. Retinotopic mapping of adult human visual cortex with high-density diffuse optical tomography. *Proc Natl Acad Sci U S A*. 2007;104(29):12169-12174. doi:10.1073/pnas.0611266104

45. Boas D a., Chen K, Grebert D, Franceschini MA. Improving the diffuse optical imaging spatial

resolution of the cerebral hemodynamic response to brain activation in humans. *Opt Lett*. 2004;29(13):1506-1508. doi:10.1364/OL.29.001506

46. Leung TS, Tachtsidis I, Smith M, Delpy DT, Elwell CE. Measurement of the absolute optical properties and cerebral blood volume of the adult human head with hybrid differential and spatially resolved spectroscopy. *Phys Med Biol*. 2006;51(3):703-717. doi:10.1088/0031-9155/51/3/015

47. Xu H, Springett R, Dehghani H, Pogue BW, Paulsen KD, Dunn JF. Magnetic-resonance-imaging-coupled broadband near-infrared tomography system for small animal brain studies. *Appl Opt*. 2005;44(11):2177-2188. doi:10.1364/AO.44.002177

48. Hueber DM, Franceschini M a, Ma HY, et al. Non-invasive and quantitative near-infrared haemoglobin spectrometry in the piglet brain during hypoxic stress, using a frequency-domain multidistance instrument. *Phys Med Biol*. 2001;46(1):41-62. doi:10.1088/0031-9155/46/1/304

49. Austin T, Gibson AP, Branco G, et al. Three dimensional optical imaging of blood volume and oxygenation in the neonatal brain. *Neuroimage*. 2006;31(4):1426-1433. doi:10.1016/j.neuroimage.2006.02.038

50. Liao SM, Ferradal SL, White BR, Gregg N, Inder TE, Culver JP. High-density diffuse optical tomography of term infant visual cortex in the nursery. *J Biomed Opt*. 2012;17(8):081414. doi:10.1117/1.JBO.17.8.081414

51. Dehghani H, Pogue BW, Poplack SP, Paulsen KD. Multiwavelength three-dimensional near-infrared tomography of the breast: initial simulation, phantom, and clinical results. *Appl Opt*. 2003;42(1):135-145. doi:10.1364/AO.42.000135

52. Hebden J, Veenstra H, Dehghani H, et al. Three-dimensional time-resolved optical tomography of a conical breast phantom. *Appl Opt*. 2001;40(19):3278-3287. doi:10.1364/AO.40.003278

53. Choe R, Corlu A, Lee K, et al. Diffuse optical tomography of breast cancer during neoadjuvant chemotherapy: a case study with comparison to MRI. *Med Phys*. 2005;32(4):1128-1139. doi:10.1118/1.1869612

54. Brooksby B, Jiang S, Dehghani H, et al. Combining near-infrared tomography and magnetic resonance imaging to study in vivo breast tissue: implementation of a Laplacian-type regularization to incorporate magnetic resonance structure. *J Biomed Opt*. 2013;10(5):051504. doi:10.1117/1.2098627

55. Pogue BWW, Jiang S, Dehghani H, et al. Characterization of hemoglobin, water, and NIR scattering in breast tissue: analysis of intersubject variability and menstrual cycle changes. *J Biomed Opt*. 2004;9(3):541-552. doi:10.1117/1.1691028

56. Hielscher AH, Klose AD, Scheel AK, et al. Sagittal laser optical tomography for imaging of rheumatoid finger joints. *Phys Med Biol*. 2004;49(7):1147-1163. doi:10.1088/0031-9155/49/7/005

57. Musgrove C, Bunting CF, Dehghani H, Pogue BW, Piao D. Computational aspects of endoscopic (trans-rectal) near-infrared optical tomography: initial investigations. 2007;6434:643409-643409-643410. doi:10.1117/12.700507

58. Dehghani H, Eames ME, Yalavarthy PK, et al. Near infrared optical tomography using NIRFAST: Algorithm for numerical model and image reconstruction. *Commun Numer Methods Eng*. 2008;25(6):711-732.

http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2826796&tool=pmcentrez&rendertype=abstract. Accessed March 24, 2016.

59. Arridge SR, Schweiger M. Direct calculation of the moments of the distribution of photon time of flight in tissue with a finite-element method. *Appl Opt*. 1995;34(15):2683-2687. doi:10.1364/AO.34.002683

60. Ripoll J, Culver JP, Pattanayak DN, Yodh AG, Nieto-vesperinas M. Recovery of optical parameters in multiple-layered diffusive media: theory and experiments. *J Opt Soc Am A Opt Image Sci Vis*. 2001;18(4):821-830. doi:10.1364/JOSAA.18.000821

61. Schotland JC. Continuous-wave diffusion imaging. *J Opt Soc Am A*. 1997;14(1):275. doi:10.1364/JOSAA.14.000275

62. Dehghani H, Pogue BW, Jiang S, Poplack S, Paulsen KD. Optical images from Pathophysiological signals within breast tissue using three-dimensional near-infrared light. *Cancer*. 2007;85(603):9-11. doi:10.1088/1742-6596/85/1/012028

63. Hielscher AH, Bartel S. Use of penalty terms in gradient-based iterative reconstruction schemes for optical tomography. *J Biomed Opt*. 2001;6(2):183-192. doi:10.1117/1.1352753

64. Ripoll J. Derivation of the scalar radiative transfer equation from energy conservation of Maxwell's equations in the far field. *J Opt Soc Am A*. 2011;28(8):1765. doi:10.1364/josaa.28.001765

65. Arridge SR. Optical tomography in medical imaging. *Inverse Probl*. 1999;15:41-93.

66. Carpenter CM, Pogue BW, Jiang S, et al. Image-guided optical spectroscopy provides molecular-specific information in vivo: MRI-guided spectroscopy of breast cancer hemoglobin, water, and scatterer size. *Opt Lett*. 2007;32(8):933-935. doi:10.1364/OL.32.000933

67. Brooksby B, Jiang S, Dehghani H, et al. Magnetic resonance-guided near-infrared tomography of the breast. *Rev Sci Instrum*. 2004;75(12):5262-5270. doi:10.1063/1.1819634

68. Mazziotta J, Toga A, Evans A, et al. A probabilistic atlas and reference system for the human brain: International Consortium for Brain Mapping (ICBM). *Philos Trans R Soc Lond B Biol Sci*. 2001;356(1412):1293-1322. doi:10.1098/rstb.2001.0915

69. Fang Q. Mesh-based Monte Carlo method using fast ray-tracing in Plücker coordinates. *Biomed Opt Express*. 2010;1(1):165-175. doi:10.1364/BOE.1.000165

70. Yao R, Intes X, Fang Q. Generalized mesh-based Monte Carlo for wide-field illumination and detection via mesh retessellation. *Biomed Opt Express*. 2015;7(1):171. doi:10.1364/BOE.7.000171

71. Schweiger M, Arridge SR, Hiraoka M, Delpy DT. The finite element method for the propagation of light in scattering media: boundary and source conditions. *Med Phys*. 1995;22(11 Pt 1):1779-1792. doi:10.1118/1.597634

72. Dehghani H, Brooksby B, Vishwanath K, Pogue BW, Paulsen KD. The effects of internal refractive index variation in near- infrared optical tomography: a finite element modelling approach. *Phys Med Biol*. 2003;48(16):2713-2727.

73. Paulsen KD, Jiang H. Spatially varying optical property reconstruction using a finite element diffusion equation approximation. *Med Phys*. 1995;22(6):691-701. doi:10.1118/1.597488

74. Arridge SR, Schweiger M, Hiraoka M, Delpy DT. A finite element approach for modeling

photon transport in tissue. *Med Phys*. 1993;20(2 Pt 1):299-309. doi:10.1118/1.597069

75. Arridge SR, Schotland JC. Optical tomography: Forward and inverse problems. *Inverse Probl*. 2009;25(12). doi:10.1088/0266-5611/25/12/123010

76. Arridge SR, Schweiger M. Photon-measurement density functions. Part 2: Finite-element-method calculations. *Appl Opt*. 1995;34(34):8026-8037. doi:10.1364/AO.34.008026

77. Dehghani H, Srinivasan S, Pogue BW, Gibson A. Numerical modelling and image reconstruction in diffuse optical tomography. *Philos Trans A Math Phys Eng Sci*. 2009;367(1900):3073-3093. doi:10.1098/rsta.2009.0090

78. Culver JP, Ntziachristos V, Holboke MJ, Yodh AG. Optimization of optode arrangements for diffuse optical tomography: A singular-value analysis. *Opt Lett*. 2001;26(10):701. doi:10.1364/OL.26.000701

79. White BR, Culver JP. Quantitative evaluation of high-density diffuse optical tomography: in vivo resolution and mapping performance. *J Biomed Opt*. 2010;15(2):026006. doi:10.1117/1.3368999

80. Eggebrecht AT, Ferradal SL, Robichaux-Viehoever A, et al. Mapping distributed brain function and networks with diffuse optical tomography. *Nat Photonics*. 2014;8(6):448-454. doi:10.1038/nphoton.2014.107

81. Paulsen KD, Jiang H. Enhanced frequency-domain optical image reconstruction in tissues through total-variation minimization. *Appl Opt*. 1996;35(19):3447-3458. doi:10.1364/AO.35.003447

82. Schweige M, Arridge SR. Optical tomographic reconstruction in a complex head model using a priori region boundary information. *Phys Med Biol*. 1999;44:2703-2721.

83. Boas DA, Dale AM. Simulation study of magnetic resonance imaging – guided cortically constrained diffuse optical tomography of human brain function. *Appl Opt*. 2005;44(10):1957-1968. doi:10.1364/AO.44.001957

84. Zhan Y, Eggebrecht AT, Culver JP, Dehghani H. Image quality analysis of high-density diffuse optical tomography incorporating a subject-specific head model. *Front Neuroenergetics*. 2012;(MAY). doi:10.3389/fnene.2012.00006

85. Wu X, Eggebrecht AT, Ferradal SL, Culver JP, Dehghani H. Fast and efficient image reconstruction for high density diffuse optical imaging of the human brain. *Biomed Opt Express*. 2015;6(11):4567-4584. doi:10.1364/BOE.6.004567

86. Scherl H, Keck B, Kowarschik M, Hornegger J. Fast GPU-based CT reconstruction using the Common Unified Device Architecture (CUDA). *IEEE Nucl Sci Symp Conf Rec*. 2007;6:4464-4466. doi:10.1109/NSSMIC.2007.4437102

87. Yoo S-H, Park J-H, Jeong C-S. Accelerating Multi-scale Image Fusion Algorithms Using CUDA. *Soft Comput Pattern Recognition, 2009 SOCPAR '09 Int Conf*. 2009:278-282. doi:10.1109/SoCPaR.2009.63

88. Zhang Z, Miao Q, Wang Y. CUDA-based Jacobi's iterative method. *IFCSTA 2009 Proc - 2009 Int Forum Comput Sci Appl*. 2009;1:259-262. doi:10.1109/IFCSTA.2009.68

89. Ren N, Liang J, Qu X, Li J, Lu B, Tian J. GPU-based Monte Carlo simulation for light propagation in complex heterogeneous tissues. *Opt Express*. 2010;18(7):6811-6823. doi:10.1364/OE.18.006811

90. Fang Q, Boas D a. Monte Carlo simulation of photon migration in 3D turbid media accelerated by graphics processing units. *Opt Express*. 2009;17(22):20178-20190. doi:10.1364/OE.17.020178

91. Alerstam E, Svensson T, Andersson-Engels S. Parallel computing with graphics processing units for high-speed Monte Carlo simulation of photon migration. *J Biomed Opt*. 2008;13(6):060504. doi:10.1117/1.3041496

92. Prakash J, Chandrasekharan V, Upendra V, Yalavarthy PK. Accelerating frequency-domain diffuse optical tomographic image reconstruction using graphics processing units. *J Biomed Opt*. 2010;15(6):066009. doi:10.1117/1.3506216

93. Zhang T, Zhou J, Carney PR, Jiang H. Towards real-time detection of seizures in awake rats with GPU-accelerated diffuse optical tomography. *J Neurosci Methods*. 2015;240:28-36. doi:10.1016/j.jneumeth.2014.10.018

94. Saikia MJ, Kanhirodan R. High Performance Single and Multi-GPU Acceleration for Diffuse Optical Tomography. In: *International Conference on Contemporary Computing and Informatics, IC3I 2014*. ; 2014:1320-1323.

95. Yi X, Wang X, Chen W, Wan W, Zhao H, Gao F. Full domain-decomposition scheme for diffuse optical tomography of large-sized tissues with a combined CPU and GPU parallelization. *Appl Opt*. 2014;53(13):2754-2765. doi:10.1364/AO.53.002754

96. Schweiger M. GPU-accelerated finite element method for modelling light transport in diffuse optical tomography. *Int J Biomed Imaging*. 2011;2011. doi:10.1155/2011/403892

97. Jermyn M, Ghadyani H, Mastanduno MA, et al. Fast segmentation and high-quality three-dimensional volume mesh creation from medical images for diffuse optical tomography. *J Biomed Opt*. 2013;18(8):086007. doi:10.1117/1.JBO.18.8.086007

98. Mathworks. spparams. https://uk.mathworks.com/help/matlab/ref/spparms.html. Published 2017. Accessed October 5, 2017.

99. Saad Y. Parallel iterative methods for sparse linear systems. *Stud Comput Math*. 2001;8(C):423-440. doi:10.1016/S1570-579X(01)80025-2

100. Nachtigal M, Reddy SC, Trefethen LN. How fast are nonsymmetric matrix iterations?* NOeL. *SIAM J Matrix Anal Appl*. 1992;13(3):778-795.

101. Faberf V, Manteuffel T. Necessary and sufficient conditions for the existance of a conjugate gradient method. *SIAM J Numer Anal*. 1984;21(2):352–362. doi:10.1080/03081087.2013.851198

102. Axelsson O, Farouq S, Neytcheva M. *Comparison of Preconditioned Krylov Subspace Iteration Methods for PDE-Constrained Optimization*. Numerical Algorithms; 2016. doi:10.1007/s11075-016-0111-1

103. Benzi M. Preconditioning Techniques for Large Linear Systems: A Survey. *J Comput Phys*. 2002;182(2):418-477. doi:http://dx.doi.org/10.1006/jcph.2002.7176

104. Elman H, Howle VE, Shadid J, Shuttleworth R, Tuminaro R. Block preconditioners based on approximate commutators. *SIAM J Sci Comput*. 2006;27(5):1651-1668. doi:10.1137/040608817

105. Saad Y, Zhang J. Enhanced multi-level block ILU preconditioning strategies for general sparse linear systems. *J Comput Appl Math*. 2001;130(1):99-118. doi:10.1016/S0377-0427(99)00388-X

106. Kolotilinat LY, Yeremin AY. Factorized Sparse Approximate Inverse Preconditionings I. Theory*. *Siam J Matrix Anal Appl*. 1993;14(1):45-58. doi:10.1137/0614004

107. Dagum L, Menon R. OpenMP: an industry standard API for shared-memory programming. *IEEE Comput Sci Eng*. 1998;5(1):46-55. doi:10.1109/99.660313

108. Nickolls J, Buck I, Garland M, Skadron K. Scalable parallel programming with CUDA. *AMC Queue*. 2008;6(April):40-53. doi:10.1145/1365490.1365500

109. Labs P. PARALUTION - v1.0.0. https://www.paralution.com/. Published 2015.

110. Tillet P, Rupp K. Towards Performance-Portable, Scalable, and Convenient Linear Algebra. *HotPar'13*. 2013. http://0b4af6cdc2f0c5998459-c0245c5c937c5dedcca3f1764ecc9b2f.r43.cf2.rackcdn.com/11423-hotpar13-tillet.pdf.

111. Day D, Heroux MA. Solvig complex-valued linear systems via equevelant real formulation. 2001;23(2):480-498.

112. Humphrey JR, Price DK, Spagnoli KE, Paolini AL, Kelmelis EJ. CULA: hybrid GPU accelerated linear algebra routines. *Defense*. 2010;7705(1):770502-770502-770507. doi:10.1117/12.850538

113. Klingner M. Improving tetrahedral meshes. 2008.

114. Eggebrecht AT, White BR, Ferradal SL, et al. A quantitative spatial comparison of high-density diffuse optical tomography and fMRI cortical mapping. *Neuroimage*. 2012;61(4):1120-1128. doi:10.1016/j.neuroimage.2012.01.124

115. Guven M, Yazici B, Kwon K, Giladi E, Intes X. Effect of discretization error and adaptive mesh generation in diffuse optical absorption imaging: I. *Inverse Probl*. 2007;23(3):1115-1133. doi:10.1088/0266-5611/23/3/017

116. Eames ME, Pogue BW, Yalavarthy PK, Dehghani H. An efficient Jacobian reduction method for diffuse optical image reconstruction Abstract. *Opt Express*. 2007;15(24):15908-15919.

117. Wu X, Eggebrecht AT, Ferradal SL, Culver JP, Dehghani H. Evaluation of rigid registration methods for whole head imaging in diffuse optical tomography. *Neurophotonics*. 2015;2(3):035002. doi:10.1117/1.NPh.2.3.035002

118. Dehghani H, White BR, Zeff BW, Tizzard A, Culver JP. Depth sensitivity and image reconstruction analysis of dense imaging arrays for mapping brain function with diffuse optical tomography. *Appl Opt*. 2009;48(10):D137-43. doi:10.1364/AO.48.00D137

119. Doulgerakis M, Eggebrecht A, Wojtkiewicz S, Culver J, Dehghani H. Toward real-time diffuse optical tomography: accelerating light propagation modeling employing parallel computing on GPU and CPU. *J Biomed Opt*. 2017;22(12):1. doi:10.1117/1.jbo.22.12.125001

120. Pogue BW, McBride TO, Prewitt J, Osterberg UL, Paulsen KD. Spatially variant regularization improves diffuse optical tomography. *Appl Opt*. 1999;38(13):2950-2961. doi:10.1364/AO.38.002950

121. Doulgerakis M, Eggebrecht A, Culver J, Dehghani H. Towards real-time functional human brain imaging with diffuse optical tomography. In: Dehghani H, Wabnitz H, eds. *European Conference on Biomedical Optics, Diffuse Optical Spectroscopy and Imaging VI*. Vol 10412. ; 2017:1041209. doi:10.1117/12.2285983

122. Wang X, Pogue BW, Jiang S, et al. Image reconstruction of effective Mie scattering parameters of breast tissue in vivo with near-infrared tomography. *J Biomed Opt*.

2006;11(4):041106. doi:10.1117/1.2342747

123. Doulgerakis-Kontoudis M, Eggebrecht AT, Dehghani H. Information rich phase content of frequency domain functional Near Infrared Spectroscopy. In: Luo Q, Ding J, Fu L, eds. *Neural Imaging and Sensing 2019*. Vol 1. SPIE; 2019:13. doi:10.1117/12.2505874

124. Diop M, St. Lawrence K. Improving the depth sensitivity of time-resolved measurements by extracting the distribution of times-of-flight. *Biomed Opt Express*. 2013;4(3):447. doi:10.1364/boe.4.000447

125. Fox PT, Miezin FM, Allman JM, Van Essen DC, Raichle ME. Retinotopic organization of human visual cortex mapped with positron-emission tomography. *J Neurosci*. 1987;7(3):913-922. doi:10.1523/jneurosci.07-03-00913.1987

126. Engel SA, Rumelhart DE, Wandell BA, et al. FMRI of human visual cortex. *Nature*. 1994;369(6481):525. doi:10.1038/369525a0

127. Zhang X, Toronov VY, Webb AG. Simultaneous integrated diffuse optical tomography and functional magnetic resonance imaging of the human brain. *Opt Express*. 2005;13(14):5513. doi:10.1364/opex.13.005513

128. Ferradal SL, Eggebrecht AT, Hassanpour M, Snyder AZ, Culver JP. Atlas-based head modeling and spatial normalization for high-density diffuse optical tomography: In vivo validation against fMRI. *Neuroimage*. 2014;85:117-126. doi:10.1016/j.neuroimage.2013.03.069

129. Muccigrosso D, Eggebrecht A. NeuroDOT : A New Neuroimaging Toolbox for DOT. 2018;2018:6-7.