Department of Precision and Microsystems Engineering

A Simulation Framework for Diffuse Correlation Spectroscopy using the Modified Born Series

Marco Loddo

Report no : 2025.066
Coach : Dr. Nandini Bhattacharya, Binbin Zhang
Professor : Dr. Nandini Bhattacharya
Specialisation : Optics for Technology (OPT)
Type of report : Master Thesis
Date : 29 August 2025

A Simulation Framework for Diffuse Correlation Spectroscopy using the Modified Born Series

A Computational Study using COMSOL and the Modified Born Series

MSc Thesis

Marco Loddo (4902580)

Delft University of Technology

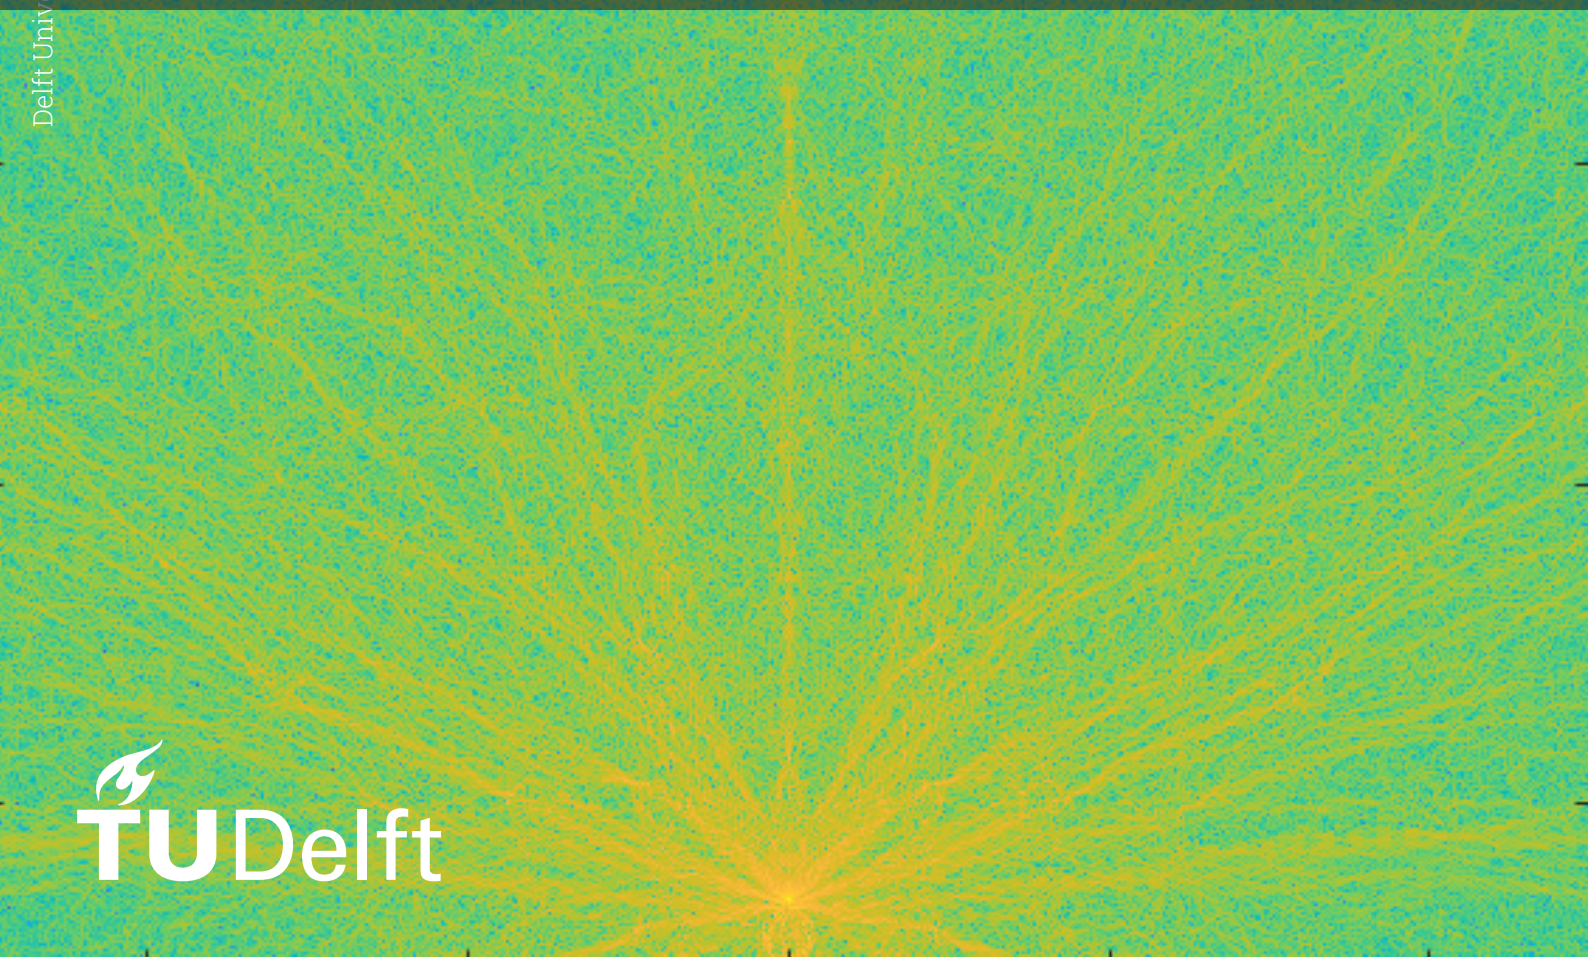

TU Delft

A Simulation Framework for Diffuse Correlation Spectroscopy using the Modified Born Series

A Computational Study using COMSOL and the Modified Born Series

by

Marco Loddo (4902580)

Instructor:	Dr. N. Bhattacharya
Teaching Assistant:	B. Zhang
Project Duration:	Sep, 2024 - Aug, 2025
Faculty:	Faculty of Mechanical Engineering, Delft

Cover:	Scattered light simulated in Matlab
Style:	TU Delft Report Style, with modifications by Daan Zwaneveld

TU Delft

Abstract

Accurate measurement of blood flow in deep tissue is essential for diagnosing and monitoring a wide range of clinical conditions, yet current imaging methods are often invasive, limited to snapshots, or impractical for continuous bedside use. Diffuse correlation spectroscopy (DCS) offers a non-invasive optical alternative, measuring blood flow dynamics through temporal fluctuations in scattered light. However, modelling DCS signals remains computationally challenging, particularly when accounting for structured flow and tissue complexity.

This thesis investigates a coupled simulation framework that integrates particle dynamics from COMSOL with optical propagation computed via a modified Born series solver (WaveSim). The framework was used to test three validation cases: temperature-dependent Brownian motion, laminar flow, and parameter extraction through exponential fitting of the autocorrelation function.

The results demonstrate that the framework captures the qualitative behaviour of DCS. Brownian motion produced linear mean squared displacement (MSD) curves, with faster optical decorrelation at higher temperatures, while laminar flow introduced quadratic MSD growth and accelerated decay. However, quantitative agreement with theoretical scaling was not achieved. A key limitation was identified as a pixel-quantisation effect: particle displacements per frame were typically far smaller than the 0.2 µm WaveSim pixel size, rendering most motion invisible to the solver. This suppressed decorrelation, distorted temperature and flow scaling, and prevented unbiased parameter extraction. Additional constraints included limited statistical averaging, absence of phase tracking, reliance on decay rates without correlation diffusion equation (CDE) modelling, and high computational costs.

The findings confirm that the COMSOL–Born framework provides proof of concept for simulating DCS but requires significant refinement for quantitative accuracy. Future work should focus on subpixel rendering, improved temporal sampling, integration of CDE solutions, and domain decomposition. With these improvements, the framework has the potential to advance real-time, non-invasive monitoring of blood flow in clinical settings.

Contents

1

Introduction

Accurate blood flow measurement is essential for understanding and monitoring various physiological and pathological states, particularly within complex biological tissues such as the brain and muscle. Real-time blood flow monitoring is critical for diagnosing medical conditions such as stroke, neurological disorders, cardio-cerebral diseases, brain health, wound healing, and skeletal diseases. Currently, available tools for blood flow measurements are mostly Doppler ultrasound-based, requiring a highly skilled operator to operate and monitor. Medical imaging scanners such as xenon-enhanced computed tomography (XeCT), dynamic susceptibility contrast magnetic resonance imaging (DSC-MRI), arterial spin labelling MRI (ASL-MRI), positron emission tomography (PET) and single photon emission computed tomography (SPECT) can map cerebral perfusion. However, these scanners are limited, as they can only provide a snapshot observation and require moving patients to imaging suites. These scanners are impractical and inappropriate for continuous monitoring. Additionally, techniques such as MRI, PET and CT require a supine scan and PET, SPECT and CT have a risk of radiation exposure. Another perfusion technique is Laser Doppler Flowmetry, however, this technique is limited to only superficial tissue blood flow [1]. A technique needs to be developed that is free from these limitations and can non-invasively monitor microvascular bloodflow in deep tissue.

Diffuse Correlation Spectroscopy (DCS) is a non-invasive optical technique that measures blood flow in deep tissues by capturing fluctuations in light intensity as red blood cells (RBCs) scatter light, producing a characteristic autocorrelation decay signal that reflects blood flow dynamics [2]. The real-time nature of DCS and its portability makes it an appealing alternative to traditional imaging techniques like MRI and CT., especially for bedside or resource-limited settings [3]. Despite its promise, DCS faces challenges in computational modelling, particularly when dealing with complex, structured flow regimes such as laminar flow in large blood vessels, where current models fall short in both accuracy and efficiency [4].

In DCS modelling, Monte Carlo (MC) simulations have long been the gold standard for accurately capturing photon transport within biological tissue due to their stochastic nature, which closely simulates the random scattering of photons by RBCs [5]. MC's strength lies in its ability to handle complex tissue geometries, making it suitable for layered or heterogeneous tissue where photon scattering varies significantly [2]. However, the main limitation of MC in DCS applications is its computational intensity. Running MC simulation in real-time is challenging due to the high computational demand, especially as tissue complexity increases [6](Yuan and Fang, 2020). Despite advances in GPU acceleration and hybrid MC models that combine mesh- and voxel-based methods to improve computational efficiency, MC methods remain largely impractical for real-time clinical monitoring [7].

Similarly, the Finite Element Method (FEM) has been applied to DCS modelling for its ability to solve partial differential equations governing photon diffusion, making it particularly useful for simulating light propagation in media with defined boundaries, such as layered tissue [8]. FEM offers flexibility and can handle boundary conditions well, which is advantageous in heterogeneous media where struc-

tural differences impact light transport. However, FEM struggles with convergence and accuracy in large-scale simulations, particularly in highly scattering or absorbing media like biological tissue. Additionally, FEM's computational demands still hinder real-time application in DCS, limiting its feasibility for dynamic blood flow measurement in clinical settings. In response to these limitations, hybrid methods that combine elements of MC and FEM have been developed, aiming to leverage the strength of each while mitigating their weaknesses [6]. These methods offer improved efficiency in photon transport simulations; however, they still lack specific adaptions for structured flow environments, like laminar flow, which are critical for accurate DCS signal modelling in certain clinical applications. The modified Born series, as introduced by Osnabrugge et al (2016) [9], offers a promising alternative to MC and FEM for light propagation in complex, inhomogeneous media. By building on the foundational Born series approximation for scattering, the modified version incorporated additional convergence factors that enhance computational efficiency without sacrificing accuracy. Mathematically, the modified Born series introduced a correction term to the initial Born approximation, which accelerates convergence in media with high scattering coefficients by selectively refining the scattering events that contribute most to the overall signal. This modification allows the Born series to account for variations in scattering intensity across large, layered tissue structures, significantly improving convergence where MC and FEM might diverge or require excessive computations. For DCS applications, the modified Born series holds particular relevance, as it allows for accurate blood flow modelling in homogeneous tissue with computationally efficient convergence even in high-scattering environments. This improvement is critical for achieving real-time DCS flow modelling, as the modified Born series reduces the computational resources required, enabling DCS to become more practical in clinical environments.

The aim of this thesis is to evaluate whether a coupled COMSOL–Born series framework can reproduce the characteristic autocorrelation dynamics of Diffuse Correlation Spectroscopy and, in particular, whether it can separate Brownian and flow contributions to blood-flow–related motion. To address this, the thesis proceeds as follows. Chapter 2 reviews the theoretical background of DCS, including light transport in tissue, analytical models, and computational approaches such as Monte Carlo, FEM, and the Born series. Chapter 3 describes the methodology of the coupled COMSOL–Born framework, including the simulation of particle trajectories, wave propagation, and autocorrelation fitting. Chapter 4 presents the results of Brownian motion and laminar flow tests, comparing them with theoretical predictions. Chapter 5 discusses the key findings, limitations, and implications for DCS modelling. Finally, Chapter 6 concludes the thesis by summarising the main contributions and outlining directions for future research.

2

Theory

This chapter reviews the theoretical foundations and recent advancements in Diffuse Correlation Spectroscopy (DCS) and related optical modelling techniques. The aim is to establish the current state of knowledge in diffuse optics, highlight methodological developments in modelling light transport and particle dynamics, and identify the limitations that remain unresolved. By systematically reviewing analytical models, numerical methods such as Monte Carlo and finite element approaches, and the recently introduced modified Born series, the chapter provides the context needed to assess existing challenges in accurately simulating DCS signals. This review forms the basis for a gap analysis, which in turn motivates the research question addressed in this thesis.

2.1. Diffuse Optics

Diffuse optics is a promising field that utilises light to measure tissue properties non-invasively. One technique for spatially mapping tissue characteristics is Near-Infrared Spectroscopy (NIRS), which measures optical properties such as scattering and absorption. Scattering occurs when a photon's energy is insufficient to excite an electron to a higher energy state. Instead, the atom's electron cloud oscillates and re-radiates the photon in a random direction. Scattering can be classified into three main types: Rayleigh scattering (for particles smaller than the wavelength), Mie scattering (for particles larger than the wavelength), and inelastic scattering, where the photon's wavelength changes due to energy transfer. NIRS typically uses coherent light in the near-infrared range because this wavelength region exhibits the lowest absorption in biological tissues. This allows for deeper light penetration, enabling the measurement of tissue properties. An example of tissue absorption spectra is shown in Figure 2.1. A complementary technique to NIRS is Diffuse Correlation Spectroscopy (DCS), which provides additional insights into tissue dynamics, such as blood flow. Together, these methods form a powerful toolbox for investigating tissue characteristics and functionality. While powerful in conjunction, they require separate instrumentation, adding difficulty in bedside monitoring and portability [5].

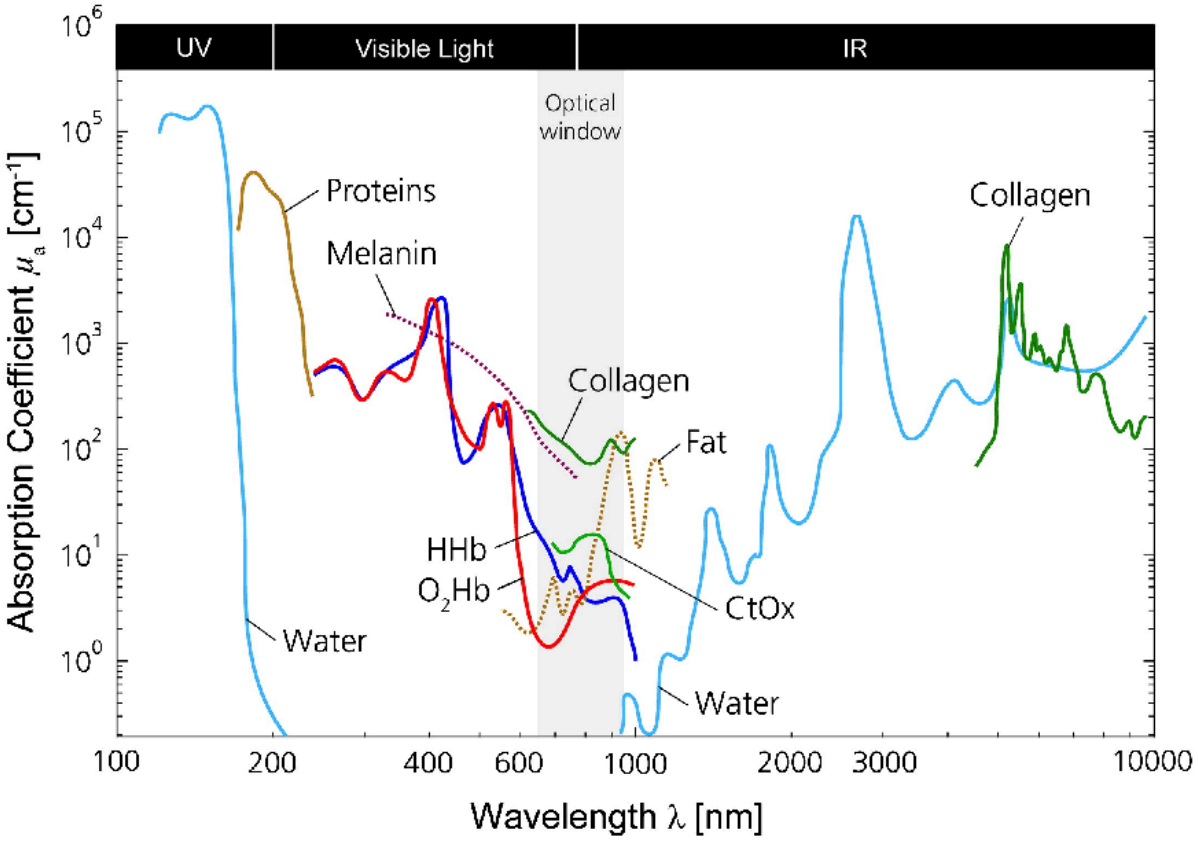

Figure 2.1: Absorption spectra of various biological chromophores at concentrations typically seen in tissue. [10]

2.2. Diffusion Correlation Spectroscopy

Diffuse Correlation Spectroscopy (DCS) is a non-invasive optical technique used to measure blood flow in biological tissues. It relies on the interaction of coherent near-infrared (NIR) light with moving red blood cells (RBCs) within a highly scattering medium. When NIR light is introduced into a tissue, it undergoes multiple scattering events before being detected. The motion of RBCs causes temporal fluctuations in the scattered light intensity, which can be quantified using an autocorrelation decay function. A schematic of this process can be seen in Figure 2.2. By analysing the decay rate of this autocorrelation, DCS can estimate the blood flow index, providing valuable information on microvascular blood flow in real-time [2].

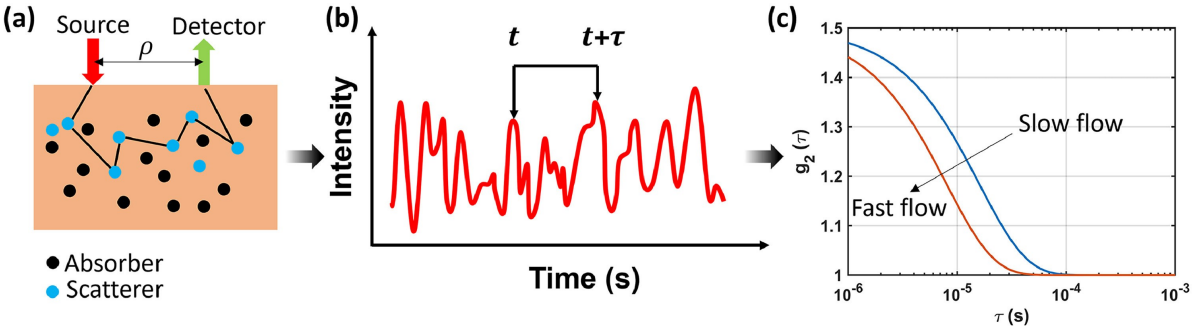

Figure 2.2: (a) The schematic of DCS measurements for semi-infinite geometry. A highly coherent laser is used as a source at a separation distance ρ from the detector. The detector measures the backscattered light; (b) Light intensity fluctuation measured over time due to photons having a longer time-of-flight before arriving at the detector caused by the moving scatterers; (c) The autocorrelation diagram showing the different decay rates between a fast and slow blood flow.[1]

In 2001, Yodh's research team was the first to formally introduce the term "Diffuse Correlation Spectroscopy" to describe a technique for analysing light scattering in biological tissues [11]. Since then, the term DCS has become widely used because it offers a mathematical framework that explains how light interacts with moving particles, like red blood cells, within the tissue. This framework relies on a simplified version of the radiative transfer equation known as the diffuse approximation. The diffusion approximation is a popular method in biomedical optics because it effectively models light transport in highly scattering media, making it suitable for estimating blood flow in tissues through DCS measurements. The development of the diffuse correlation theory by Boas and Yodh was an important milestone in DCS theory. This theory provided the first comprehensive model to predict particle movement in highly scattered media by analysing the behaviour of diffuse speckle fields [12] [13]. This theory became the foundation for the DCS model that is now used to estimate the blood flow index (BFi), a key parameter to measure in vivo microvascular blood flow [14]. Over the last two decades, DCS technology has seen significant advancement and validation. Today, it is used to perform non-invasive blood flow (BF) measurements in deep tissues. The technology has been used in areas such as the skin, muscle [15], breast tumour [16], and brain [17] [18] [19]. A schematic for the approaches is shown in Figure 2.3.

The deepest recorded measurements were 2 centimetres below the surface with a source-detector separation of 4 cm [20]. This capability has led to the widespread application of DCS in clinical and research settings, including cerebral blood flow monitoring, muscle perfusion, and tumour evaluation [21] [22] [23]. Furthermore, the integration of DCS with Near-Infrared Spectroscopy (NIRS) and Diffuse Optical Spectroscopy (DOS) has expanded its potential by enabling simultaneous monitoring of blood flow and tissue oxygenation. Initially demonstrated in rat brains [11], this combined approach has since been applied to adult humans [24], further enhancing the versatility of DCS.

2.2.1. Analytical models

The development of analytical models in Diffuse Correlation Spectroscopy has evolved significantly, enabling more precise and adaptable blood flow measurements within biological tissues. DCS originated with models based on the assumption of a semi-infinite homogeneous medium, which treated biological tissues as boundless structures with uniform optical properties. These early models utilised the diffusion approximation to the radiative transfer equation to describe photon propagation in tissue. This approximation simplifies photon transport by assuming isotropic scattering and a direct, uniform decay in photon density with distance, allowing for basic blood flow quantification [12]. However, while efficient, this approach was limited in its ability to capture the complexities of biological tissues, which often exhibit layered structures and boundaries that significantly impact photon scattering [25].

Researchers developed layered tissue models that account for distinct optical properties across tissue layers to address these limitations, enhancing DCS's sensitivity to blood flow in structured tissues. Two-layer models were introduced for applications where tissue structures, such as scalp and brain layers, could be differentiated, particularly in neonatal brain monitoring. Subsequently, three-layer models were created to accommodate additional structures, such as the skull, thereby improving the accuracy of cerebral blood flow measurements in adults. These layered models allow DCS to produce more accurate blood flow measurements by incorporating more realistic representations of tissue structure [17] [26]. A representation of layered models is shown in Figure 2.4, and the resulting temporal field autocorrelation plot is shown in Figure 2.5 using Continuous Wave-DCS.

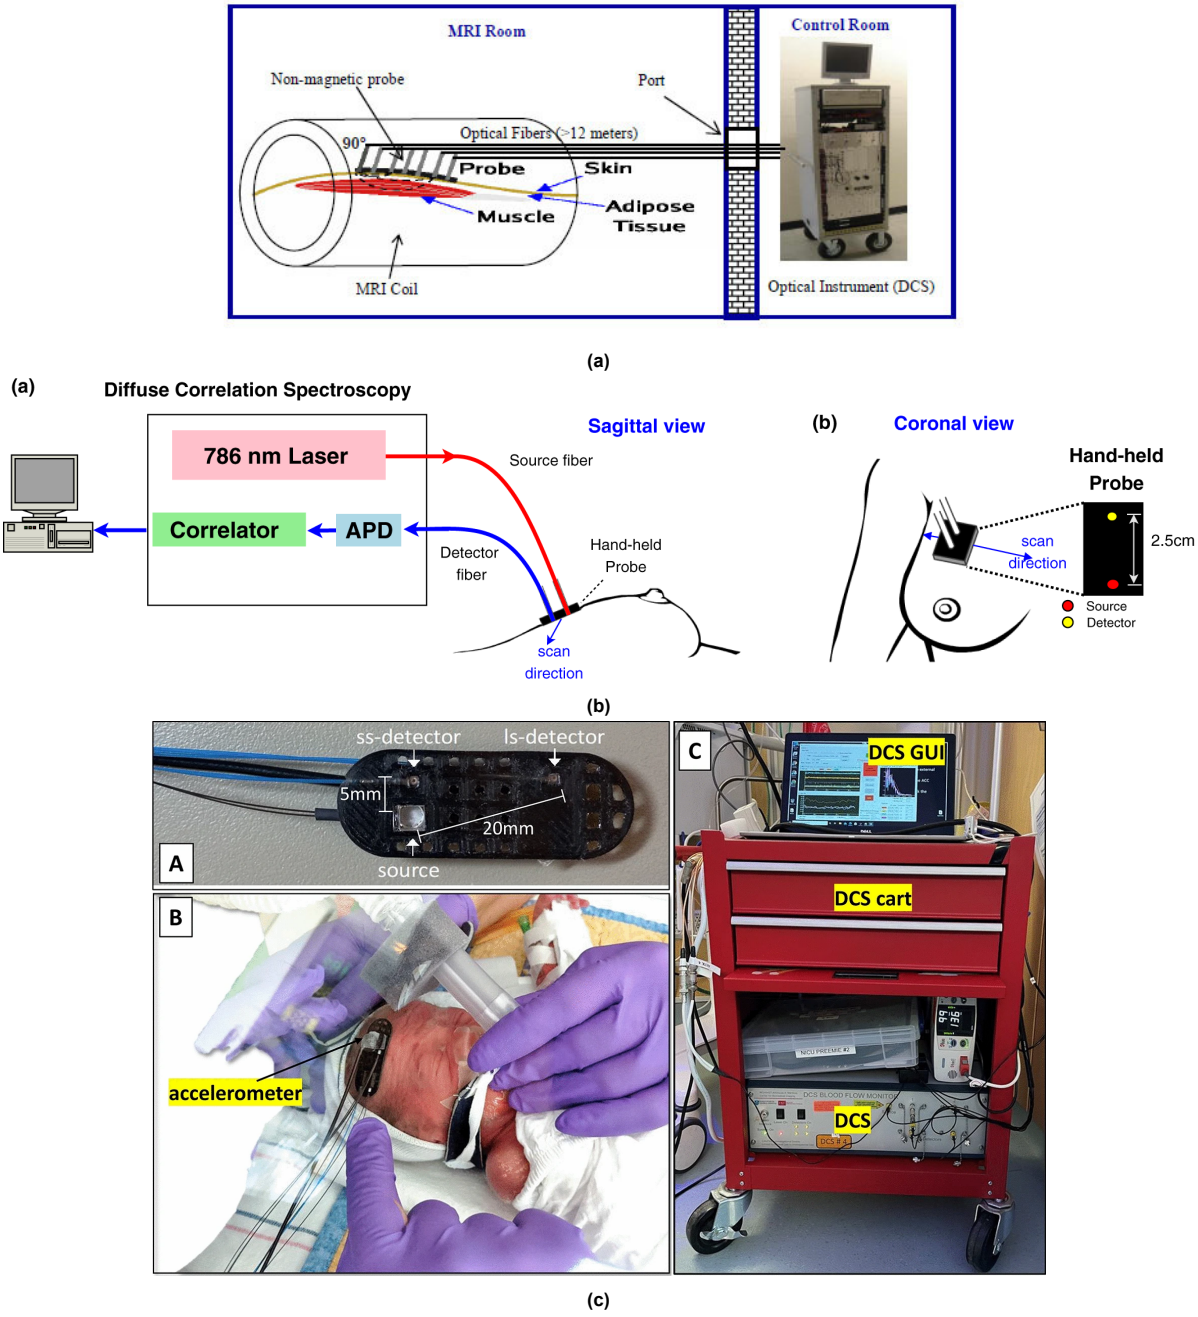

Figure 2.3: (a) Schematic of an optical-MRI measurement. The calf (with the optical probe) was then placed into the MRI knee coil. The optical probe in the MRI room was connected to the DCS instrument in the control room by 12-meter-length fibres through a port in a magnetic-field-shielded wall [15].; (b) Diagram of DCS working on a breast [16]; (c) DCS sensor was attached to the infant's head for blood flow monitoring [18].

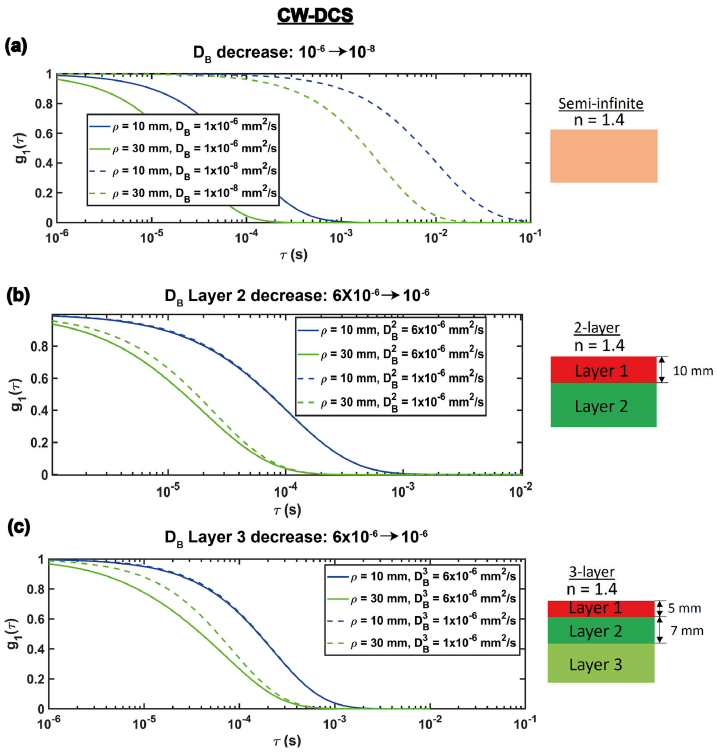

Figure 2.5: Representation of temporal field autocorrelation $g_1(\tau)$ for a medium with source-detector separation of $\rho = 10mm$ and $\rho = 30mm$ and a source of $\lambda = 785nm$. (a) Semi-infinite medium, where the D_b^1 is varied between 1E-6 to 1E-8 mm^2/s; (b) 2-layer medium, with D_b^1 = 1E-6 mm^2/s and varying D_b^2; (c) 3-three layer model, with D_b^1 and D_b^2 = 1E-6 mm^2/s and varying D_b^3 between 1E-6 to 1E-8 mm^2/s. [1]

Building upon these static models, DCS techniques were advanced to include time-dependent or time-domain models. Time-domain DCS (TD-DCS) captures the arrival times of photons scattered through tissue, providing additional depth sensitivity by reducing contamination from superficial layers and isolating signals from deeper tissues. This is achieved by selecting photons with longer path lengths, which are more likely to penetrate the tissue deeper, thus enhancing the measurement accuracy [25]. In their work, Sutin et al. introduced this novel method by utilising the time domain to simultaneously obtain the temporal point spread function to quantify tissue properties and the autocorrelation function to determine BFi. The arrival times of photons are used to differentiate between short and long photon paths, to be able to determine BFi at different depths. The research derived a path length-dependent normalised field temporal autocorrelation function, showing that the autocorrelation decay rate increases linearly with photon path length. The authors showed that narrowing the time gate width increases the correlation amplitude, improving the signal-to-noise ratio (SNR) of the blood flow measurement. They also confirmed the linear relationship between the decay rate of the autocorrelation function and photon path length, validating the theoretical predictions. The TD-DCS system demonstrated the ability to distinguish between blood flow in the scalp and brain by analysing early and late-arriving photons. During hypercapnic challenges, the system detected a 55% increase in cerebral blood flow, consistent with expected physiological responses. While this method is promising, TD-DCS requires advanced instrumentation, such as single-photon avalanche diodes (SPADs) and time-correlated single-photon counting (TCSPC) systems, making it a technically demanding and more expensive approach, often limiting its clinical accessibility. The recent introduction of Frequency-Domain DCS (FD-DCS) by Moka et al. (2022) represents a further evolution in DCS technology [26]. FD-DCS utilises light sources modulated at specific frequencies, allowing simultaneous assessment of both dynamic (blood flow) and static (tissue absorption and scattering) optical properties. By extending and generalising conventional DCS into the frequency domain, this novel approach measures a frequency-dependent intensity autocorrelation function, which is then fitted to a frequency-domain solution of the correlation diffusion equation to estimate tissue optical parameters. To verify the functionality of FD-DCS, measurements were performed on two tissue-simulating phantoms; the results are shown in Figure 2.6. The results

demonstrated that FD-DCS successfully retrieved the static and dynamic optical properties with errors comparable to or lower than those of Fast DCS and FD-DOS. Specifically, for the first phantom, the estimated static optical properties deviated by less than 9%, while the dynamic flow index was found to be on the order of 1e-8 cm^2/s, consistent with expected values. For the second phantom, FD-DCS retrieved the static optical properties with a higher error of up to 29%, which was attributed to volumetric uncertainties in phantom preparation rather than limitations of the method itself. This approach provides a comprehensive picture of tissue hemodynamics by capturing blood flow indices alongside information about tissue composition. FD-DCS is especially valuable for applications requiring continuous monitoring, as it enhances DCS's ability to differentiate between static and dynamic scatterers within the tissue. Despite these advantages, FD-DCS introduces additional complexity to DCS systems, as it requires high-speed detectors and carefully modulated light sources, which can pose challenges for portable and cost-effective deployment.

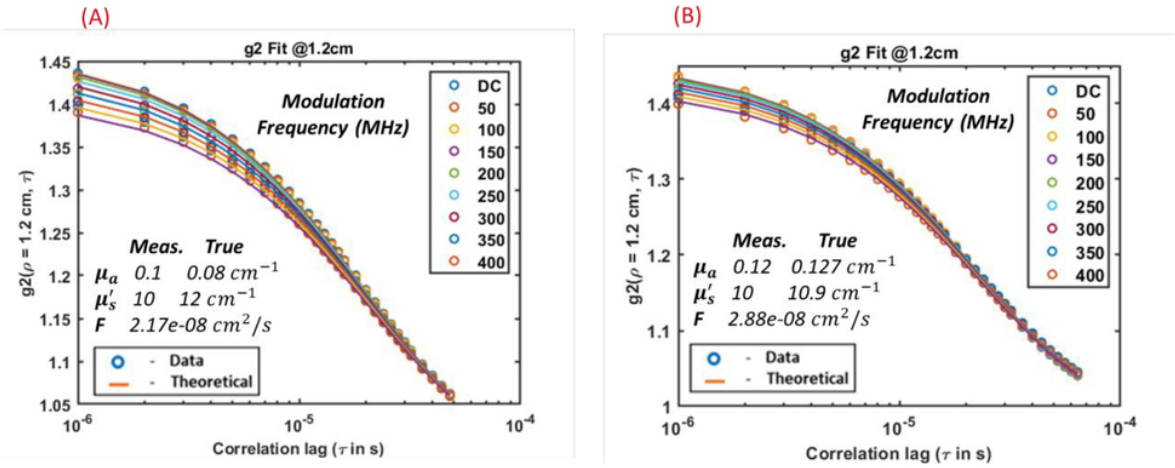

Figure 2.6: Experimental results: Frequency dependent DCS intensity auto-correlation functions (circles) fit to a Frequency dependent DCS model at multiple modulation frequencies. Optical properties are retrieved with $\sim$ 20% error [26]

Other developments in DCS have involved combining it with Near-Infrared Spectroscopy and Diffuse Optical Spectroscopy, creating hybrid systems capable of simultaneously monitoring tissue blood flow and oxygenation. This combination was first applied in animal studies [11] and later adapted for human brain monitoring [24], enabling more comprehensive monitoring of tissue perfusion and oxygenation in clinical settings like neurosurgery and intensive care. Additionally, deep-learning algorithms have been explored to enhance DCS data interpretation, aiming to improve the accuracy and speed of blood flow measurements in real time. These computational methods are particularly promising for adapting DCS to more complex flow environments, making DCS a more versatile and accessible tool in clinical settings [26].

2.2.2. Blood flow

A primary application of diffuse correlation spectroscopy lies in measuring blood flow, achieved specifically by estimating the blood flow index (BFi). Derived from DCS, the BFi is a critical parameter that reflects microvascular blood flow within biological tissues, calculated by analysing the temporal fluctuations in scattered light caused by the movement of red blood cells (RBCs) within tissue. This approach builds upon diffusing-wave spectroscopy principles, where light scattering patterns reveal the dynamics of scatterers [3]. By quantifying these dynamics, the BFi enables DCS to assess blood flow in real-time, a capability of significant value in both research and clinical settings [14].

The study by Boas et al is a seminal contribution to the field, providing both theoretical and practical insights into DCS measurements. This work establishes a critical link between the DCS signal and blood flow, providing a robust theoretical foundation and practical tools for advancing noninvasive tissue perfusion measurement [14]. By elucidating the dominance of diffusive RBC dynamics and validating the linearity of BFi with absolute blood flow, the study paves the way for more accurate and reliable clinical

applications of DCSs for RBCs, accounting for both convective motion and shear-induced diffusion. The study captures the complex interplay of photon migration and RBC dynamics by simulating scattering events across varying vessel densities and geometries. The authors derive and validate explicit relationships between the RBC diffusion coefficient and absolute tissue blood flow, linking these to the widely used correlation diffusion equation (CDE). They also test the impact of factors such as vessel diameter, hematocrit, and scattering properties on the BFi. The findings reveal that shear-induced diffusion of RBCs overwhelmingly dominates the autocorrelation decay observed in typical DCS measurements, contrary to earlier assumptions favouring convective dynamics. This insight has significant implications for interpreting BFi in clinical applications. The study demonstrates that BFi is linearly proportional to absolute blood flow under standard experimental conditions. However, the proportionality is modulated by factors such as vessel diameter and hematocrit, underscoring the importance of these variables in accurate flow estimation.

Recent studies, such as that by Zhu et al. (2020) [27], have further validated the complexity of RBC motion by conducting phantom experiments simulating microvascular blood flow environments. Their work explored the influence of both Brownian motion and random ballistic flow on the DCS signal, establishing a hybrid model that incorporates these two types of RBC movement. The results can be seen in Figure 2.7. The results highlighted that without applied pressure, the Brownian motion fits the curve perfectly, and the random ballistic flow model underestimates the decay. The results also showed that, under increased hydraulic pressure, RBC motion becomes more complex, with both flow types underestimating the decay rate. The proposed hybrid model fits the correlation curve perfectly, having both flow types contributing to the DCS signal. The proposed hybrid model offers a more realistic framework for interpreting DCS data under dynamic conditions, emphasising the importance of adapting DCS models to various flow regimes.

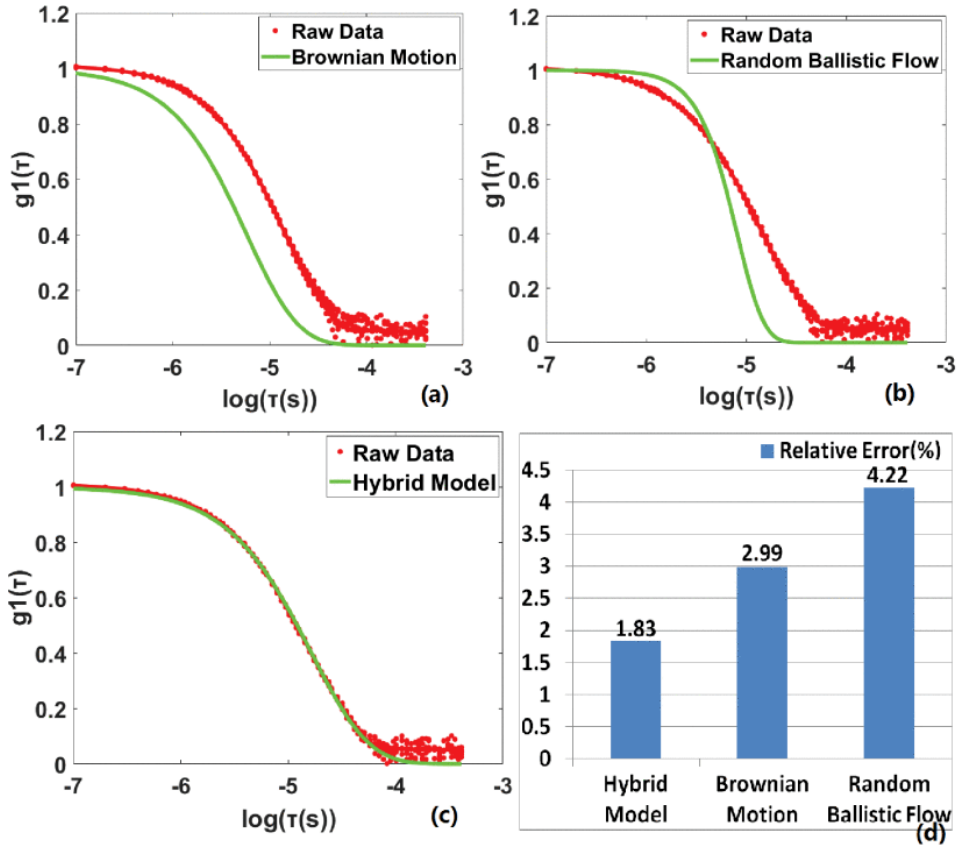

Figure 2.7: At 200 ml/h pump speed, the g1 curves of the raw data and the curve fitting by (a) the single Brownian motion model, (b) the single random ballistic model, and (c) the hybrid model; (d) The relative error (RE) between the raw data of g1 curve and that fitted by the three flow models. [27]

The theoretical model proposed by Sakadžić, Boas, and Carp (2017) refines Diffuse Correlation Spectroscopy (DCS) by incorporating both diffusive and advective red blood cell (RBC) motions, offering a more precise interpretation of blood flow measurements [3]. Using a combination of Monte Carlo (MC) simulations and theoretical derivations, the study demonstrates that shear-induced RBC diffusion is the dominant contributor to the DCS autocorrelation decay, particularly in small and medium-sized vessels where random motion exceeds the influence of structured, advective flow. The authors derived a correlation diffusion equation (CDE) and a correlation transfer equation (CTE) to model light transport through heterogeneous vascular networks, accounting for variations in vessel size, flow velocity, and RBC displacement. Their simulations confirmed that, for typical source-detector separations (1–3 cm), the observed DCS signal decay is primarily driven by diffusion, aligning with previous experimental findings. The advective motion was found to play a minor role in capillaries and smaller arterioles, but its influence increased in larger vessels, where laminar flow effects became more pronounced. However, as Zhu et al. demonstrated, introducing hydraulic pressure into the microvasculature can significantly increase the contribution of ballistic flow, necessitating hybrid modelling approaches to represent real-world conditions accurately. Moreover, traditional assumptions in DCS modelling, such as the cumulant approximation, often fall short when attempting to capture the full complexity of RBC motion in biological tissues. The cumulant approximation simplifies the DCS autocorrelation function by assuming RBC motion is purely Brownian, thus failing to account adequately for structured, advective flows typical in larger vessels and under varied physiological conditions. Du Le and Srinivasan (2020) argue that the interference between the first and second cumulant terms can complicate interpretations, especially in dynamic or heterogeneous flow environments, leading to inaccuracies in BFi estimations [28]. This highlights the need for hybrid motion models capable of accurately representing both random and directed flow dynamics within tissues, thereby offering a more robust framework for a wide range of clinical applications. Alongside these limitations in motion assumptions, Irwin et al. (2011) emphasise the significant impact of tissue optical properties on BFi accuracy [29]. Through experiments using liquid phantoms, they demonstrated that variations in the reduced scattering coefficient (μs') and absorption coefficient (μa) can lead to notable errors in BFi calculations, particularly in tissues with high optical variability, such as tumours or layered structures like the brain. Specifically, inaccuracies in μs' have been shown to result in BFi errors as high as 280% in some clinical settings, underscoring the need to account for optical properties when interpreting DCS signals. These findings suggest that incorporating concurrent measurements or making precise assumptions regarding optical properties is crucial for achieving reliable blood flow assessments across different tissue types and experimental conditions.

2.3. Modelling Methods

This section will explore modelling methods for light propagation with a focus on DCS. In this section, two main methods will be reviewed, Monte Carlo and Finite Element methods, which are well-known and used. A third potential method will be explored that has been used for modelling light propagation, called the Born series.

2.3.1. Monte Carlo

Monte Carlo (MC) simulations are regarded as the gold standard in photon transport modelling due to their stochastic nature, which enables them to simulate light scattering with high accuracy. The core principle of MC simulations is probabilistic modelling, simulating the random paths taken by photons as they scatter and are absorbed within a medium. This method models each photon as it moves through the tissue, interacting with particles and following a randomised trajectory based on predefined probabilities for scattering and absorption events. The motion of the photons is governed by the mean free path (l) and the scattering angle (θ), which are expressed as a probability distribution [30]. The probability distribution is expressed in Equation 2.1 and Equation 2.2

$$P(s_1) = \frac{P(s < s_1)}{ds_1} = \mu_t \exp(-\mu_t s_1) \longrightarrow \mu_t = \frac{1}{l} \tag{2.1}$$

$$P(cos(\theta)) = \frac{1 - g^2}{2(1 + g^2 - 2g\cos(\theta)^{3/2})} \tag{2.2}$$

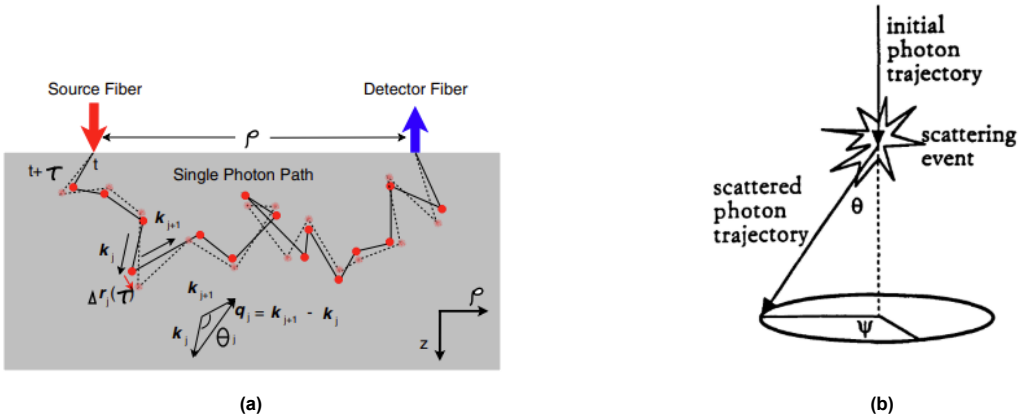

Figure 2.8: Illustration of Monte Carlo simulation. (a) Photon propagation through a medium. [31] (b) Deflection of a photon by a scattering event. The angle of deflection, θ and the azimuthal angle, ψ [30]

In MC simulations, photons undergo various interactions such as reflection, refraction, absorption and scattering, all of which are dictated by the optical properties of the medium. These properties include the refractive index (n), absorption coefficient (μ_a), scattering coefficient (μ_s), and scattering anisotropy (g). The absorption coefficient quantifies the likelihood of photon absorption per unit path length, while the scattering coefficient measures the probability of photon scattering. Scattering anisotropy, defined as the mean cosine of the scattering angle, indicates the directionality of scattering, with biological tissues typically exhibiting g=0.9, signifying strongly forward-directed scattering [32] [31]. The relation between free path, scattering and absorption coefficient is expressed in Equation 2.3. By tracing a large number of photons, MC simulations produce statistical representations of light transport, which can then be analysed to generate autocorrelation functions for DCS blood flow measurements. The generated autocorrelation function is fitted into a theoretical model, obtained by solving the correlation diffusion equation for the given case.

$$\mu_t = \mu_s + \mu_a \tag{2.3}$$

Over the past decade, significant efforts have focused on addressing the primary limitations of MC simulations: low computational efficiency and difficulties in modelling complex anatomical structures. Traditional MC simulations, which require tracing millions of photons, have historically demanded extensive computational time, often taking hours on single-core central processing units (CPUs). To overcome this, researchers have adopted massively parallel computing architectures, including graphic processing units (GPUs) and field-programmable gate arrays (FPGAs), dramatically reducing simulation times to minutes or even seconds. In 2009 a GPU-based MC framework was introduced that leverages the parallel processing capabilities of graphics hardware [33]. This approach allowed for large amounts of photons to be simulated simultaneously, significantly improving the computational efficiency compared to traditional CPU-based Monte Carlo. Their model was validated against analytical results from diffusion theory for a homogeneous semi-infinite medium, showing good agreement. The model was also tested with an MRI scan of a human head, and it was found to be capable of handling arbitrary complex media.

Modelling complex anatomical structures has been another issue for the MC method. Significant progress has been made in recent years, permitting simulations in simple layered domains and complex heterogeneous tissues. Recent research conducted on the accuracy of multi-layered tissue models highlighted that two-layered models are limited in separating extracerebral signals, such as scalp and skull blood flow, from brain-derived signals in cerebral blood flow index (CBFi) measurements [34]. Traditional DCS techniques rely heavily on single-layer analytical models that assume homogeneous tissues, often failing to differentiate between cerebral blood flow and extracerebral contributions, such as scalp blood flow (SBF). This limitation is particularly pronounced in adult brain monitoring, where thicker extracerebral layers diminish brain sensitivity and amplify systemic interference. These constraints have hindered the clinical applicability of DCS for accurate and noninvasive CBF monitoring.

The multi-layer MC method overcomes these challenges by incorporating anatomically realistic, heterogeneous geometries into photon transport simulations. Using both slab and head-like models derived from MRI scans, the algorithm allows for tunable adjustments to layer thicknesses representing scalp, skull, cerebrospinal fluid (CSF), and brain. A representation of this can be seen in Figure 2.9. By simulating photon paths and momentum transfer through these layers, the method produces precise time-resolved autocorrelation functions, enabling differentiation between superficial and deep tissue blood flow. This approach is a marked improvement over analytical models, achieving greater accuracy in scenarios with significant SBF interference, such as during hypercapnia. The study validates the MC-based models against transcranial Doppler ultrasound (TCD), a gold standard for CBF measurement. The results shown in Figure 2.10, demonstrate that three-layer MC models, particularly those incorporating curved head geometries, closely align with TCD-derived middle cerebral artery velocity (MCAV) timecourses. Compared to analytical models, MC-based methods significantly reduce systemic contamination, accurately capturing the rapid dynamics of CBF changes during hypercapnia and achieving near-complete recovery to baseline after systemic interference subsides.

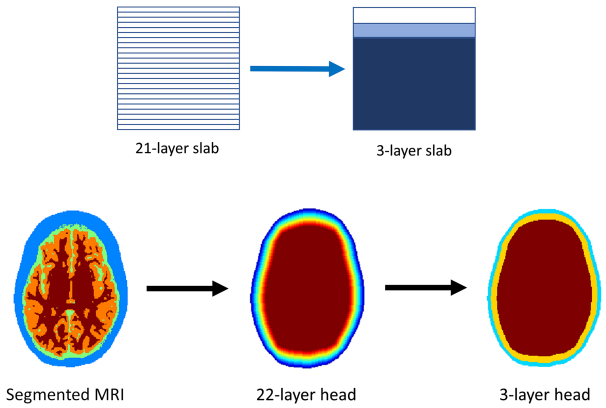

Figure 2.9: Top row: the 21-layer slab is concatenated into a 2-, 3-, or 4-layer volume in post-processing (3-layer is shown). Bottom row: the segmented MRI is iteratively image-eroded to make a 22-layer head volume and is analogously post-processed into a 2-, 3-, or 4-layer volume. [34]

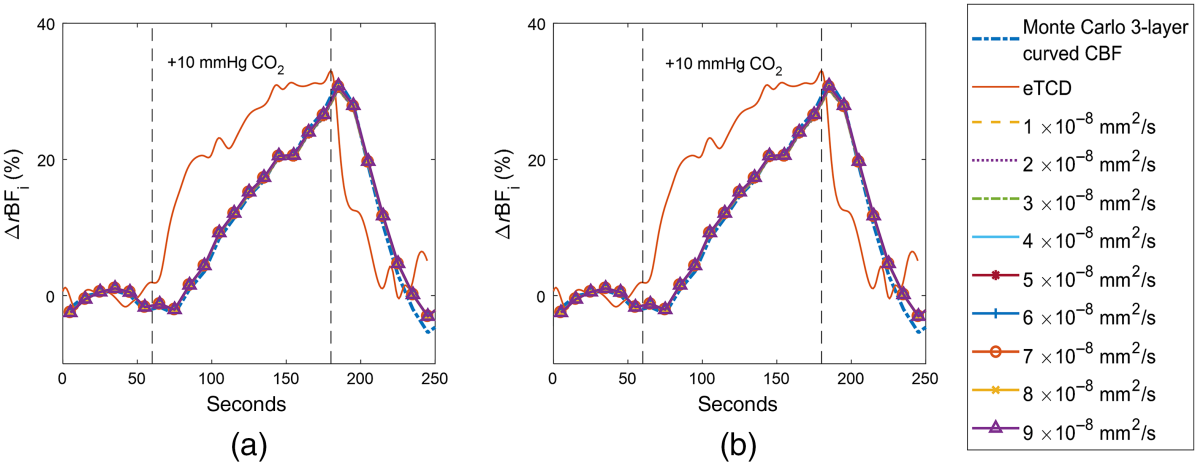

Figure 2.10: MC curved, four-layer fits with various diffusion coefficients used for the CSF layer for subject 1. A 1-mm CSF thickness was used for (a) and a 2-mm thickness was used for (b). [34]

The paper by Tran and Jacques introduces a significant improvement in Monte Carlo (MC) light transport simulations, addressing a major limitation of voxel-based MC methods: their inability to accurately model curved and oblique boundaries [8]. Traditional MC techniques use discrete Cartesian grids, where curved surfaces are approximated by a series of flat voxel facets. This simplification leads to inaccuracies in Fresnel-based calculations of reflection and transmission at boundary interfaces, which

are critical for simulating light transport in biological tissues. The methodology proposed involves the preprocessing of voxel boundaries to calculate a gradient map of refractive index mismatches, smoothing this gradient to mitigate discretisation errors, and employing trilinear interpolation to estimate gradual changes in normal vectors across curved surfaces. A representation of the algorithm can be seen in Figure 2.11. These innovations allow for significantly improved boundary representation without the need for higher voxel density, which traditionally incurs steep computational and memory costs. By focusing on efficient implementations—such as the use of Sobel filters and interpolation schemes—the authors strike an effective balance between computational efficiency and accuracy. Results from simulations involving spheres and lenses demonstrate that the new algorithm closely aligns with analytical solutions, with less than 5-degree errors in reflectance and transmittance angles. This is a substantial improvement over the traditional voxel-based facet-normal method, which can exhibit errors of up to 180 degrees under certain conditions. Moreover, the method is designed to integrate seamlessly with GPU acceleration, enabling scalability for large simulations.

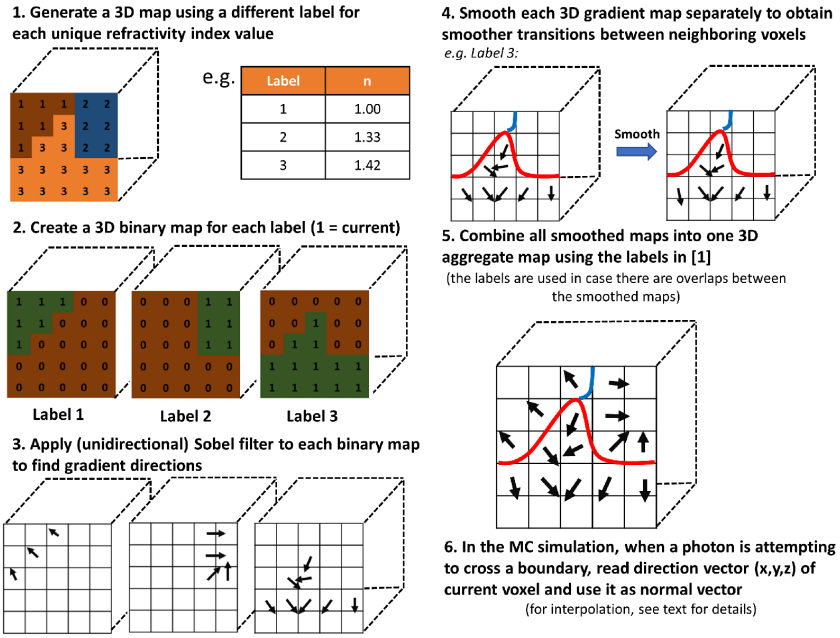

Figure 2.11: Proposed algorithm to handle curvature and oblique angles in MC light simulations. illustration of the method for specifying the boundaries and the unit vectors that are normal to that surface in each voxel containing the surface [8].

These advancements have been particularly beneficial in photon transport simulations across complex curved geometries, such as layered brain tissue and vascular networks. Recent advancements in hybrid Monte Carlo (MC) methods have addressed limitations in voxel-based and mesh-based approaches, particularly for modelling complex tissue geometries. The Split-Voxel Monte Carlo (SVMC) algorithm introduced by Yan and Fang (2020) combines the computational efficiency of voxel-based methods with the accuracy of mesh-based approaches [7]. This paper presents the Split-Voxel Monte Carlo (SVMC) algorithm, which offers a novel hybrid approach combining the benefits of voxel-based and mesh-based Monte Carlo (MC) methods for modelling light transport in complex bio-tissues. An illustration can be seen in Figure 2.12. Traditional voxel-based MC methods excel in simplicity and computational efficiency, especially on GPUs, but suffer from inaccuracies due to rasterisation errors at curved or oblique boundaries (Type I errors) and inherent surface discretisation artefacts (Type II errors). Mesh-based methods, while accurate in representing complex geometries, are resource-intensive and less efficient on memory-sensitive platforms. SVMC bridges this gap by employing a marching-cubes algorithm to preprocess voxelized domains, extracting and embedding sub-voxel boundary information directly into the voxel structure. The SVMC methodology introduces several critical innovations. First, the marching-cubes algorithm generates triangular boundary patches within boundary voxels, which are further processed to compute surface normals and positions. These refined boundaries are encoded as additional geometric information in an extended voxel data structure, allowing the simu-

lation to incorporate intra-voxel boundary interactions during photon transport. Non-boundary voxels are treated as in standard voxel-based MC, preserving computational simplicity. The hybrid approach significantly reduces Type I and Type II errors without the overhead of generating and managing complex tetrahedral meshes. Validation across multiple benchmarks demonstrates the efficacy of SVMC. In simple domains, such as concentric spheres, SVMC with Gaussian smoothing achieves near-parity accuracy with Dual-Grid Mesh-Based Monte Carlo (DMMC), significantly outperforming conventional voxel-based MC methods. Realistic domains, including a segmented brain model, further highlight SVMC's ability to handle complex geometries with reduced error margins, particularly at curved tissue-air interfaces. Cross-sectional contour plots to compare light fluence distributions between three MC algorithms for all benchmarks can be seen in Figure 2.13. In (a) and (b), it can be seen that for benchmarks B1 and B2, the SVMC with and without smoothing has improved accuracy compared to VMC. In (c), an excellent agreement between SVMC and DMMC solutions is observed for benchmark B3 along the scalp-air surface, suggesting that the Type II error was largely removed. However, the discrepancies in the CSF layer under detectors 7-9, a result of Type I error, were reduced but not eliminated. Simulation results show that SVMC is 2x to 6x faster than DMMC and requires lightweight preprocessing, with a trade-off of approximately 50% slower runtime compared to conventional voxel-based MC due to additional boundary handling.

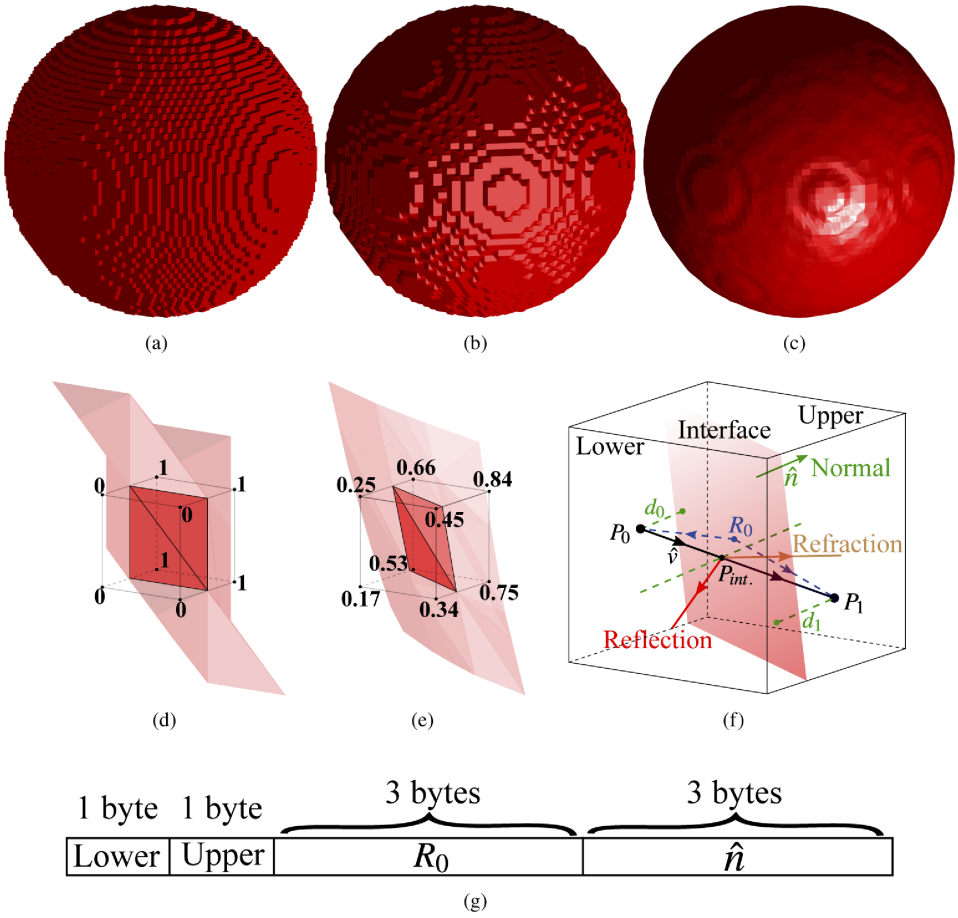

Figure 2.12: Illustrations of boundaries of a spherical domain extracted using (a) voxel representation, and marching-cubes (b) without and (c) with smoothing by a Gaussian-filter. Zoom-in views of the surfaces in (b) and (c) can be found in (d) and (e), respectively. We also show diagrams explaining (f) ray-tracing computation and (g) our extended voxel data memory layout that encodes additional shape information [7].

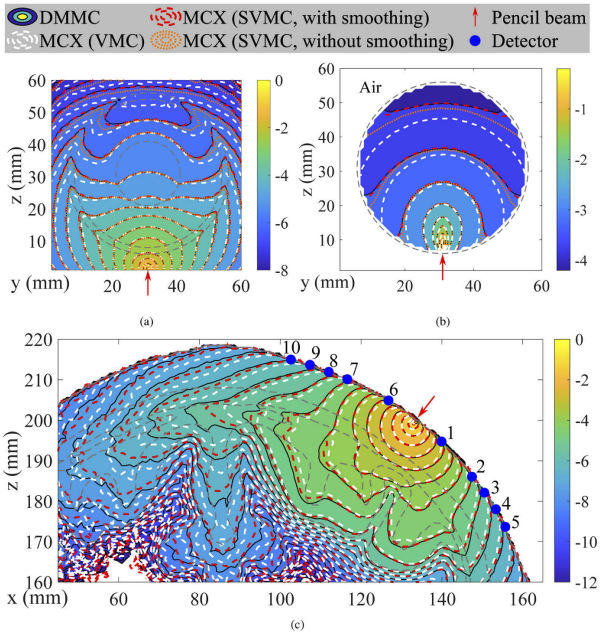

Figure 2.13: Fluence (mm^{-2}, in log-10 scale) contour plots of dual-grid MMC (DMMC), MCX (conventional VMC mode) and MCX (SVMC mode with and without volume smoothing) in a set of benchmarks: (a) B1, (b) B2 and (c) B3. The red arrow represents an inward-pointing pencil beam source. Grey-dashed lines mark media boundaries. [7]

A key development was later made by Yuan et al., introducing the implicit mesh-based Monte Carlo (iMMC) algorithm [6]. Traditional methods, such as voxel-based Monte Carlo (MC) and mesh-based MC (MMC), while foundational, suffer from critical limitations. Voxel-based methods rely on dense volumetric grids to approximate complex geometries, which leads to excessive memory requirements and computational inefficiencies. Similarly, MMC improves geometric representation by employing tetrahedral meshes but still struggles with high memory consumption and reduced scalability in anatomically intricate domains. These constraints have historically impeded progress in simulating light transport within large, detailed, and heterogeneous biological tissues, where geometric accuracy is paramount. The iMMC algorithm addresses these challenges by introducing a hybrid approach that leverages the strengths of both voxel- and mesh-based methodologies while circumventing their weaknesses. The different meshes can be seen in Figure 2.14. At its core, iMMC employs coarse skeletal tetrahedral meshes to represent key structural elements, such as vessel centrelines, and implicitly defines shapes, such as cylinders or spheres, to represent surrounding fine details. This implicit modelling paradigm is revolutionary, as it decouples the computational burden of explicitly meshing intricate features from the need to accurately simulate photon interactions. By relying on parameterised shapes instead of high-density mesh elements, iMMC achieves memory savings exceeding 100-fold compared to traditional methods. Furthermore, the methodology integrates computationally efficient intersection tests for implicit shapes, such as ray-cylinder or ray-sphere calculations. While these approximations introduce minor errors, they have been benchmarked against standard MMC and Dual-Grid MMC (DMMC) methods, showing less than 1% deviation in accuracy. In Figure 2.15, cross-sectional plots of log-scaled fluence along the source planes are shown to compare solutions derived from MMC, DMMC, and iMMC for B1, B2, and B3 benchmark domains. This figure shows an excellent match between IMMC and DMMC for all benchmarks. The MMC results show minor mismatches when compared to those of iMMC and DMMC, especially near the light source. These results show that IMMC is capable of achieving a similar level of accuracy to DMMC and MMC, despite only using a fraction of the memory for the mesh data. The algorithm's ability to maintain geometric fidelity without resorting to dense mesh discretisation represents a pivotal advancement in balancing accuracy and efficiency. This capability is particularly critical for simulating light transport in multi-scale biological tissues, where features such as microvascular networks, porous media, and thin membranes coexist and require representation at varying resolutions.

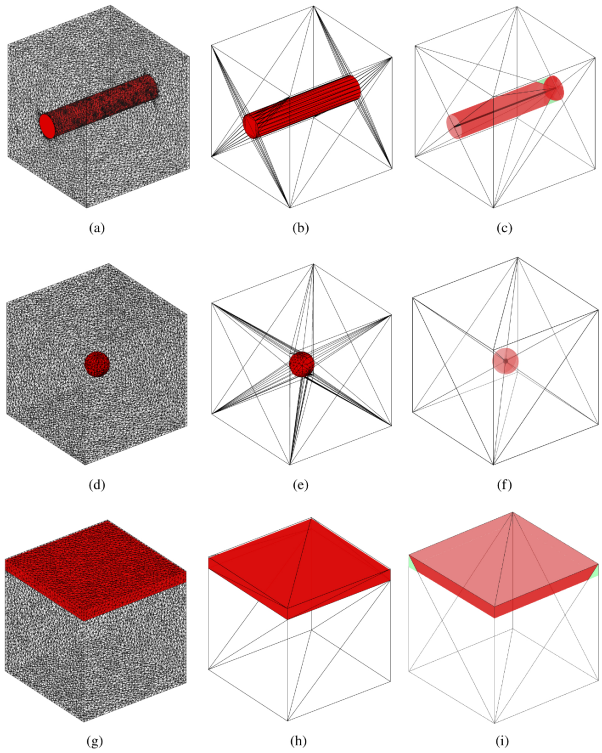

Figure 2.14: The meshes generated for three benchmarks and three Monte Carlo methods. Benchmarks (a-c) B1, (d-f) B2 and (g-i) B3 are built to validate the edge-based, node-based and face-based iMMC algorithms, respectively. Each column, from left to right, shows the meshes created for standard mesh-based Monte Carlo (MMC), dual-grid mesh-based Monte Carlo (DMMC) and iMMC, respectively. [6]

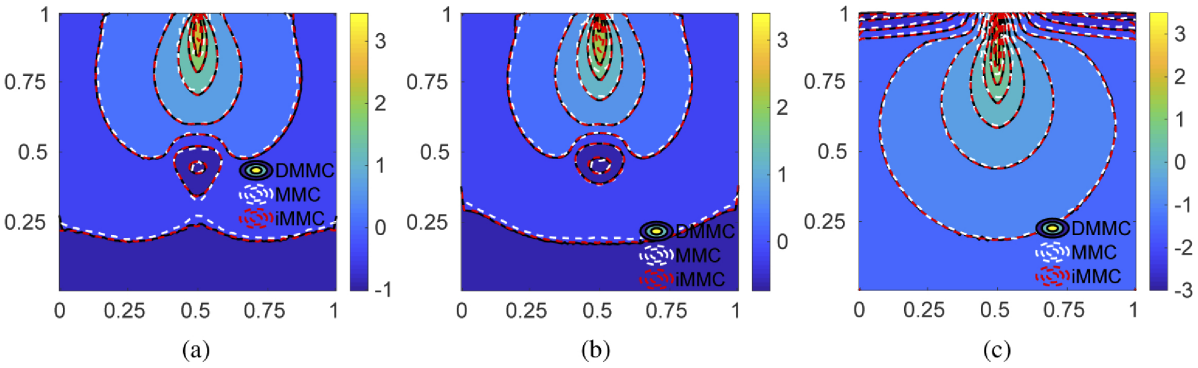

Figure 2.15: Log-scaled fluence (J · mm−2) cross-sectional plots along the source planes to compare solutions derived from MMC, DMMC, and iMMC for (a) B1, (b) B2, and (c) B3 benchmark domains [6].

2.3.2. Finite Element

The Finite Element Method (FEM) is a deterministic computational technique widely used for solving partial differential equations (PDEs) in complex geometries. In the field of photon transport modelling, FEM is typically employed to solve the diffusion approximation of the radiative transfer equation (RTE), which provides a simplified framework for understanding light propagation in highly scattering media. By discretising the computational domain into finite elements, FEM enables efficient modelling of light transport across tissues with complex structures and varying optical properties. FEM operates by dividing the domain into smaller subregions or ”elements,” typically in the form of triangles (2D) or tetrahedra (3D). Within each element, the photon density or light fluence is represented as a polynomial function. The PDE governing photon transport is then solved iteratively, with boundary conditions applied to account for reflections and refractions at tissue interfaces. Optical properties, including the absorption

coefficient, scattering coefficient, and anisotropy factor, are assigned to each element, dictating how photons interact with the medium. FEM's strength lies in its ability to adapt to irregular geometries and varying material properties, making it well-suited for modelling multi-layered or heterogeneous biological tissues. For example, FEM has been widely applied to simulate light propagation in tissues such as the brain, skin, and muscle, where layers of differing optical properties and curvatures significantly influence photon transport.

Recent advancements have focused on enhancing FEM's computational efficiency and accuracy to better meet the demands of modern optical imaging and therapeutic applications. The studies by Vasudevan and Narayanan Unni (2022, 2023) present a comprehensive framework for assessing skin blood flow using finite element methods integrated with diffuse correlation spectroscopy [35] [36]. In their 2022 work, the authors developed a 3D FEM model of a bilayer skin structure with embedded microcapillaries to simulate light-tissue interactions and blood flow. By solving the correlation diffusion equation (CDE), they derived the blood flow index and demonstrated its sensitivity to flow velocity while remaining consistent across capillary depths. In Figure 2.16, the FEM model and the results are shown. The resulting autocorrelation plot shows the sensitivity for different flow speeds, validating that a higher flow speed correlates to a faster decay rate. The results for the calculated BFi can be seen in Figure 2.17. The results show that the BFi changed proportionally to the inlet blood velocities through the capillaries for different surrounding static tissue scattering. The determined BFi does not vary significantly with capillary depth. This in-silico model provided detailed insights into microvascular dynamics but lacked direct experimental validation.

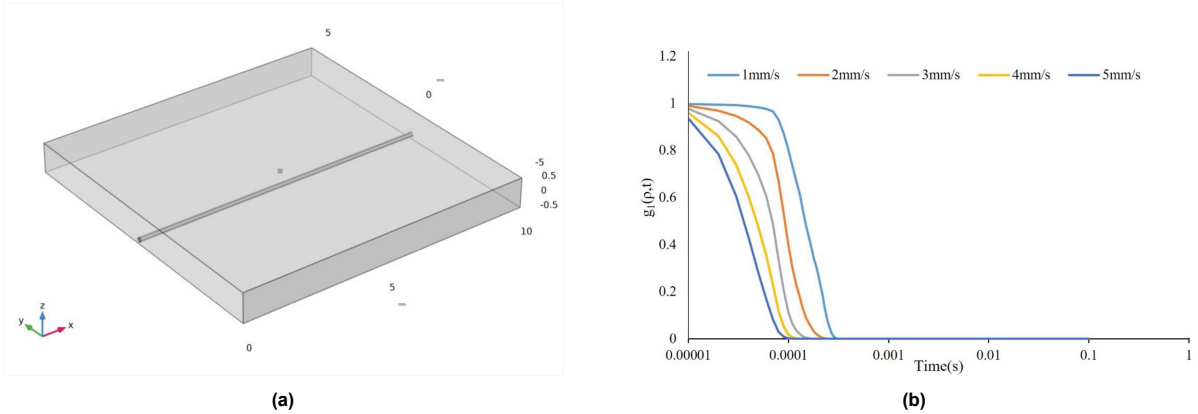

Figure 2.16: Model and simulated results (a) 3D geometry model of dermal skin tissue with embedded capillary. (b) Autocorrelation plot for varying inflow velocities through the designed capillary from FEM simulations [35].

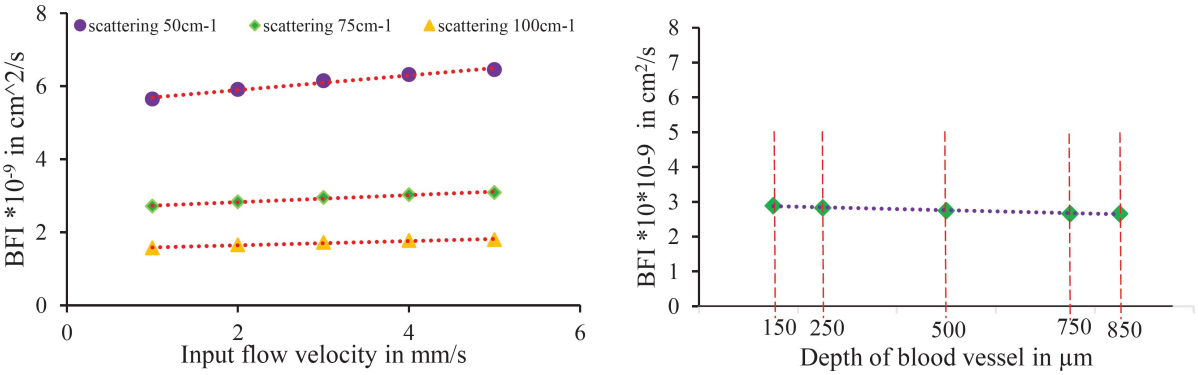

Figure 2.17: Calculated BFI from the developed algorithm for varying input velocities at varying static property conditions of surrounding tissues, and 3b illustrates the calculated BFI with varying depth of capillary vessel from the tissue surface [35]

To build upon their previous work, the 2023 study expanded the investigation by conducting in vitro experiments using capillary-mimicking microchannels embedded in skin phantoms. The experimental

setup combined a short source-detector separation (SDS) DCS system with finite element method (FEM)-based simulations to quantify blood flow dynamics. The results, presented in Figure 2.18, demonstrated a strong correlation between the simulated and experimental Blood Flow Index (BFI) values, validating the model's accuracy. The FEM simulations provided a detailed representation of light propagation and RBC dynamics, effectively capturing the underlying physics of photon scattering and diffusion. However, their high computational cost and dependence on idealised tissue properties posed challenges for broader applicability. Conversely, the in vitro experiments served as practical validation tools, offering real-world insight into flow variations and system performance. Nonetheless, they introduced potential variability due to phantom fabrication inconsistencies and flow control limitations. While this hybrid approach presents a promising direction for microvascular blood flow assessment, its validation remains limited to short-range applications. Future research is needed to adapt and optimise this methodology for deeper or more heterogeneous tissue structures, ensuring its scalability and clinical viability.

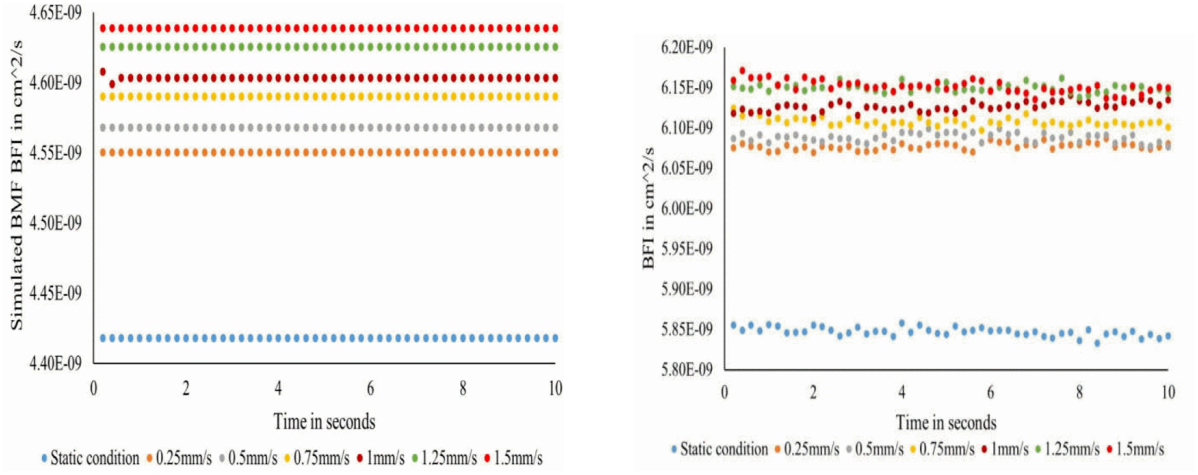

Figure 2.18: Results from short-range DCS. (a) Simulated bfi with bmf as the flow medium and its variation with changes in input flow velocity. (b) Determined bfi from the in-vitro studies with bfi as the flow medium [36]

The study by Ao Teng et al. (2020) investigates the interaction of pulsed laser light with brain tissue using a two-dimensional finite element model developed in COMSOL Multiphysics [37]. The model consists of a simplified brain structure, including a water layer, grey matter, and embedded blood vessels, irradiated by a short-pulse laser source with a wavelength of 840 nm. A representation of the model can be seen in Figure 2.19. Light propagation was modelled using the diffusion equation, while temperature changes were computed with the biological heat transfer equation. The results reveal that light energy decreases exponentially with penetration depth, with blood vessels absorbing significantly more light than surrounding grey matter, retaining approximately 85.8% of the light energy absorbed by the grey matter. The higher absorption in blood vessels results in a localised temperature increase of 0.15 K compared to the grey matter, which exhibits minimal temperature changes. This study highlights the strengths of finite element modelling in capturing the optical and thermal dynamics of laser-tissue interactions. However, its simplified two-dimensional geometry limits its applicability to the intricacies of real brain structures, and experimental validation remains necessary to confirm the model's predictive accuracy.

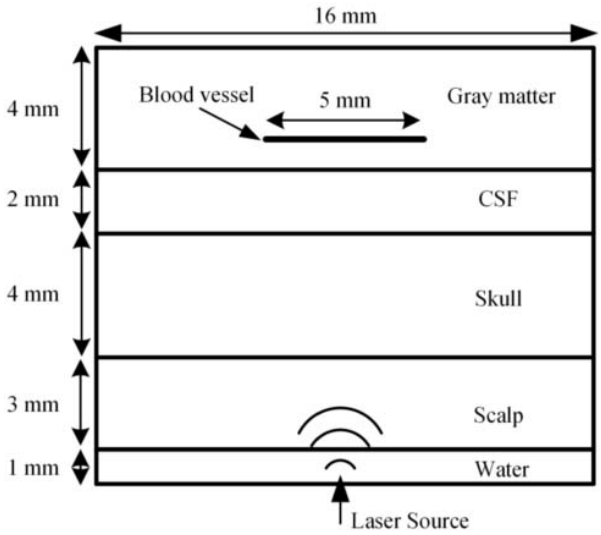

Figure 2.19: Model geometry; The diameter of the blood vessel is 60 µm and its length is 5 mm. The pulsed laser point source is placed in the middle of the water layer to irradiate the brain tissue [37].

The study by Shimaa Mahdy et al. (2022) presents a non-invasive method for detecting brain tumours, specifically gliomas and meningiomas, by numerically analysing the optical fluence rate at the scalp. Using a five-layer human head model developed in COMSOL Multiphysics, the researchers simulated light propagation through tissue using the steady-state diffusion equation modelled by the Helmholtz equation [38]. This model incorporates key tissue layers such as the scalp, skull, cerebrospinal fluid (CSF), grey matter, and white matter, each with unique optical properties. The simulation investigates the spatial fluence rate distribution for two infrared laser wavelengths (1000 and 1100 nm). An illustration for the model can be seen in Figure 2.20. Results show that the presence of tumours significantly alters the minimum fluence rate values on the scalp due to differences in absorption and scattering properties between normal and tumour tissues. Gliomas and meningiomas exhibited distinct fluence rate profiles, enabling differentiation from normal brain tissue. While this approach provides a promising theoretical framework for early tumour detection, it is constrained by its dependence on simulation-based results and requires experimental validation for clinical application. Additionally, the reliance on predefined optical properties may limit its adaptability to inter-patient variability. Nonetheless, this work offers valuable insights for improving non-invasive brain tumour diagnostic techniques.

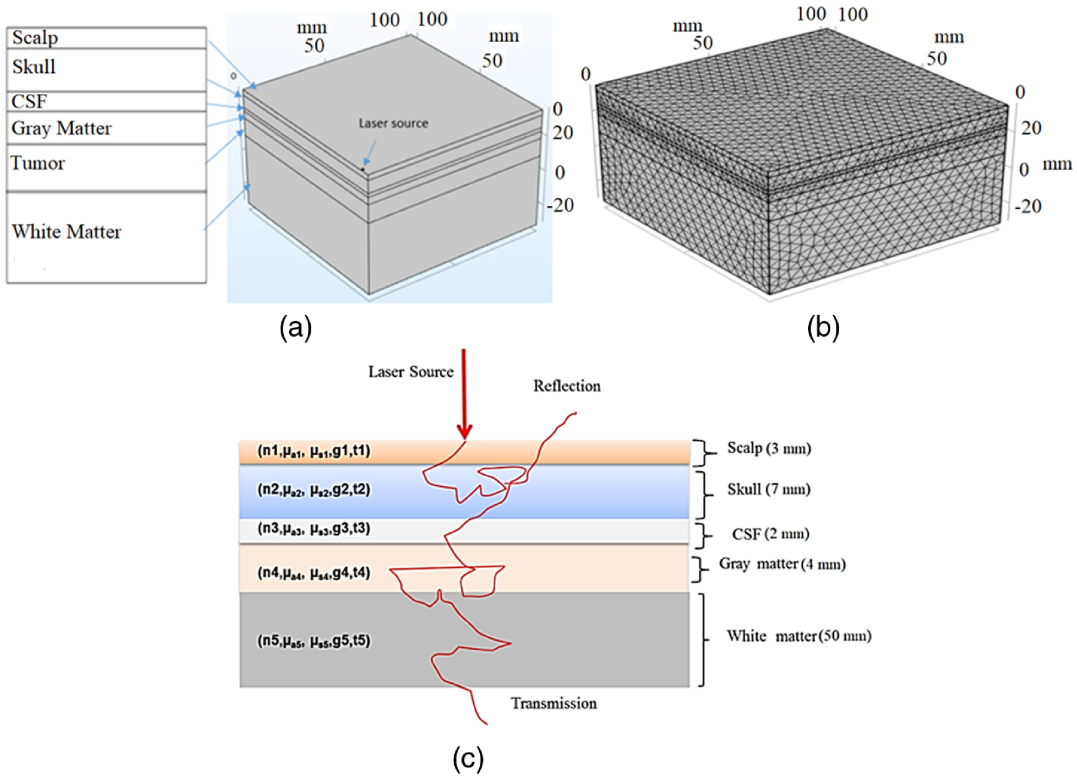

Figure 2.20: Geometry of the proposed model: (a) 3D model layers of the human head with tumour, (b) the implemented mesh, and (c) dimensions of the model layers [38].

The study by Doulgerakis et al. (2017) focuses on accelerating light propagation modelling in diffuse optical tomography (DOT) using parallel computing on both GPU and CPU platforms [39]. This work addresses the computational bottleneck posed by the forward problem in DOT, particularly when employing high-resolution finite element models with up to 600,000 nodes for realistic adult head imaging. By leveraging GPU-based parallelisation and implementing advanced iterative solvers such as BiCGStab with preconditioners, the study achieved a tenfold speed increase compared to traditional CPU methods, with computational times as low as 0.25 seconds per excitation source. This innovation enables near-real-time light propagation calculations while maintaining high numerical accuracy, making it suitable for functional brain imaging and high-density DOT systems. Despite these advances, the approach faces limitations in handling even larger-scale models due to hardware memory constraints and remains dependent on achieving an optimal balance between computational speed and termination criteria to minimise errors. The study offers a significant step toward real-time DOT applications, but further optimisation for more extensive models and diverse tissue structures remains a challenge.

2.3.3. The Born Series

The Born series offers a fundamentally different approach to modelling light propagation in scattering media compared to methods like Monte Carlo (MC) and Finite Element Method (FEM). Instead of relying on probabilistic distributions or directly solving the correlation diffusion equation, the Born series addresses the problem by solving the inhomogeneous Helmholtz equation using integral methods and Green's functions. The Helmholtz equation governs wave propagation and, when applied to electromagnetic waves, provides a scalar approximation of light behaviour in three-dimensional scattering media [40].

Early approaches to solving the Helmholtz equation in scattering media involved discretisation techniques, such as finite element [41] and finite difference methods [42]. These methods transform the continuous equation into a system of linear equations, which can then be solved using direct matrix inversion or iterative solvers. However, these approaches suffer from high computational costs, particularly in large or highly inhomogeneous domains. More advanced numerical techniques, such as

Krylov subspace methods [43] extended with multigrid solvers [44], have improved efficiency but remain computationally demanding for large-scale problems [45].

A more efficient alternative to these methods is the Born series, derived from Green's function theorem. The Born series expresses the scattered field as a series of successive perturbations, where each term represents an increasing order of scattering interaction. This framework has been extensively studied and applied to wave scattering problems, including photon propagation in biological tissues [46]. The first-order Born approximation assumes a single scattering event and neglects higher-order interactions, making it computationally efficient but limited to weakly scattering media. Early studies, such as those by Sevick-Muraca et al. (1997), criticised the use of first-order Born and Rytov approximations, particularly in optical imaging [47]. They demonstrated that these approximations underestimate the scattering and absorption properties when multiple scattering events dominate. By extending the Born series to second and higher-order perturbations, they showed significant improvements in accuracy for optical imaging. However, this approach still struggled with convergence in media with high scattering contrast. The work by Rossum et al. (1999) further explored the application of the second-order Born approximation in random media, focusing on light scattering in highly diffusive environments [40]. Their study highlighted a critical limitation: the divergence of Green's function in three-dimensional media with point scatterers. To manage this divergence, they introduced a regularisation technique, ensuring the stability of the series for small scatterers. While this approach allowed for more accurate modelling in weak scattering environments, it remained impractical for strongly scattering biological tissues.

A significant contribution to understanding the limitations of the Born series was made by Çapoğlu et al. (2009), who conducted an error analysis of the Born approximation applied to biological tissues [48]. Using a finite-difference time-domain (FDTD) simulation as a benchmark, they demonstrated that the Born approximation is valid only for weak refractive index fluctuations. For stronger fluctuations, the approximation becomes unreliable as wave localisation effects become significant, leading to a breakdown of the scattering coefficient's definition. This work underscored the need for alternative methods or modifications to the Born series to handle highly scattering media effectively.

Further exploration of the Born approximation was carried out by Joonoh Lim et al. (2011), who applied the first-order Born approximation to model light scattering by red blood cells (RBCs). Their study focused on deriving RBC parameters, such as diameter and dimple size, by solving the inhomogeneous Helmholtz equation under the Born approximation [49]. Although this method proved effective for analysing blood samples and provided a fast, cost-effective solution, it was limited to scenarios with weak scattering contrast.

In 2014, Rogers et al. introduced a versatile correlation function for the refractive index of tissues, offering a new way to model how refractive index fluctuations influence light scattering in complex biological tissues [50]. Instead of treating tissues as discrete scatterers, they modelled them as continuous random media. This approach, while more accurate for representing real biological tissues, relied on the Rayleigh-Gans-Debye (RGD) approximation, a variant of the first-order Born approximation, which assumes weak scattering. Although effective for weak scattering environments, it struggled in the presence of large refractive index fluctuations or multiple scattering effects.

The inherent limitations of the Born series in strongly scattering media were addressed by Osnabrugge et al. (2016), who introduced the Modified Born Series (MBS) [9]. By incorporating a pre-conditioner, the MBS ensured convergence even in media with high scattering potential. This innovation allowed the series to be applied to large, strongly scattering systems, making it a viable alternative to Monte Carlo and FEM-based methods for photon transport modelling. The MBS was tested against PSTD and FDTD methods and was shown to be several orders of magnitude faster in finding a steady-state solution to the inhomogeneous wave equation. In their research, the MBS was demonstrated for 1, 2 and 3-dimensional media, in Figure 2.21 the results for a 2D inhomogeneous medium with random complex potential and a 3D disordered medium. The method was extended to handle vector fields [51] and birefringent media [52], demonstrating its versatility in various optical applications. The research by Kruger et al. (2017) extended the MBS to handle the full vectorial Maxwell's equation by deriving a vectorial Green's function and adapting the preconditioning strategy. This enabled the MBS to accurately simulate the polarisation effects and the coupling between different components of the electrical field. These effects are crucial when the reflective index varies spatially or when the medium induces strong polarisation-dependent phenomena. Their work demonstrated that with the extension, the MBS

can compute solutions for complex vector fields orders of magnitude faster than conventional methods like FDTD. Further expanding on this work, Vettenburg et al (2019) generalised the MBS to handle arbitrarily complex electromagnetic media. In their work, they developed a more sophisticated preconditioning approach that ensured convergence even when the medium's response is non-Hermiation and when strong electric-magnetic coupling is present. With this extension, MBS can simulate birefringent materials as well as chiral and magnetic materials.

Building on the foundation and the extensions, Osnabrugge et al. (2021) further refined the MBS by focusing on reducing boundary-related errors and improving computational efficiency. Their work introduced ultra-thin boundary layers and a pipelined acyclic convolution scheme to minimise wrap-around errors and reflection artefacts at domain boundaries [53]. The padding-free acyclic convolution is a new convolution method introduced in this research which eliminates the need for padding, improving the computational efficiency. These enhancements resulted in up to 80% reduction in simulation size and a fivefold decrease in computational time, making MBS highly suitable for large-scale simulations in biomedical optics.

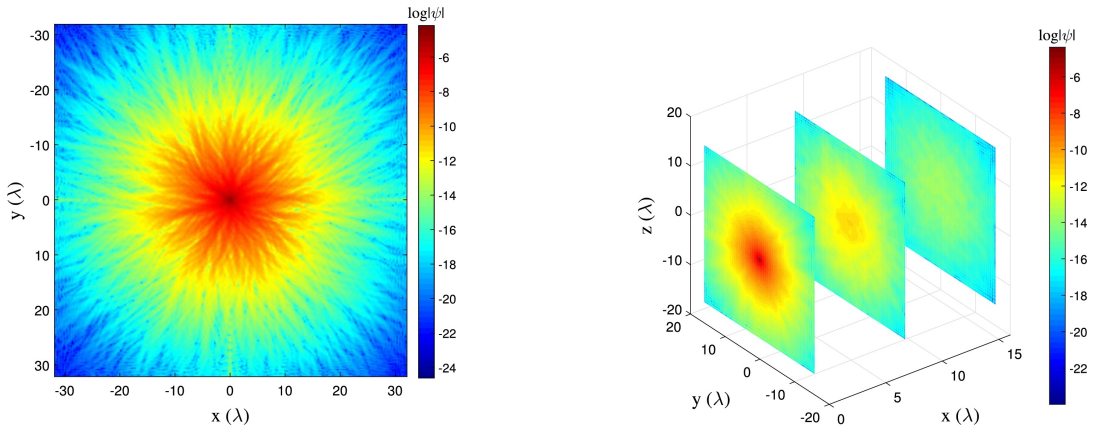

Figure 2.21: Illustrating results of the MBS. (a) Numerical experiment simulating an oscillating point source at the center of an inhomogeneous medium with a random potential distribution. The steady-state solution is found using our modified Born approach. The field amplitude is represented on a logarithmic scale. (b) Simulation results for a point source at x = y = z = 0λ in a 3-D disordered medium of 128 × 128 × 128 pixels. The 3D-solution is represented by three cross-sectional slices at x = 0λ, x = 8λ, and x = 16λ, where the field amplitude is shown on a logarithmic scale. The interference fringes at the boundaries of the slices are the result of the periodic boundary conditions. [9]

2.4. Research Question and Objectives

The literature reviewed in this chapter highlights the significant progress in modelling light transport for Diffuse Correlation Spectroscopy (DCS), particularly through Monte Carlo and Finite Element methods. While these approaches are well-validated and widely used, they are computationally expensive and limited in their ability to directly incorporate time-resolved particle dynamics. The modified Born series (MBS) has recently emerged as a highly efficient solver for the inhomogeneous Helmholtz equation, capable of handling large and strongly scattering systems. However, its potential for simulating DCS has not yet been systematically explored.

This leads to the central research question of this thesis:

> *Can a coupled COMSOL–modified Born series framework reproduce the expected Brownian and flow-induced decorrelation in DCS, and allow separate estimation of diffusion and flow contributions?*

To address this, the following objectives are defined:

1. Simulate Brownian and flow-driven particle trajectories in COMSOL under varying physical parameters.
2. Use the modified Born series to compute the resulting speckle fields from these trajectories.
3. Derive temporal autocorrelation functions of the simulated speckle fields and compare the decay characteristics with DCS theory.
4. Assess whether Brownian diffusion and flow parameters can be independently extracted from exponential fits to the autocorrelation decay.

Having established the research question and objectives, the following chapter outlines the methodology developed to address them, detailing the coupled COMSOL–Born series simulation framework and analysis pipeline.

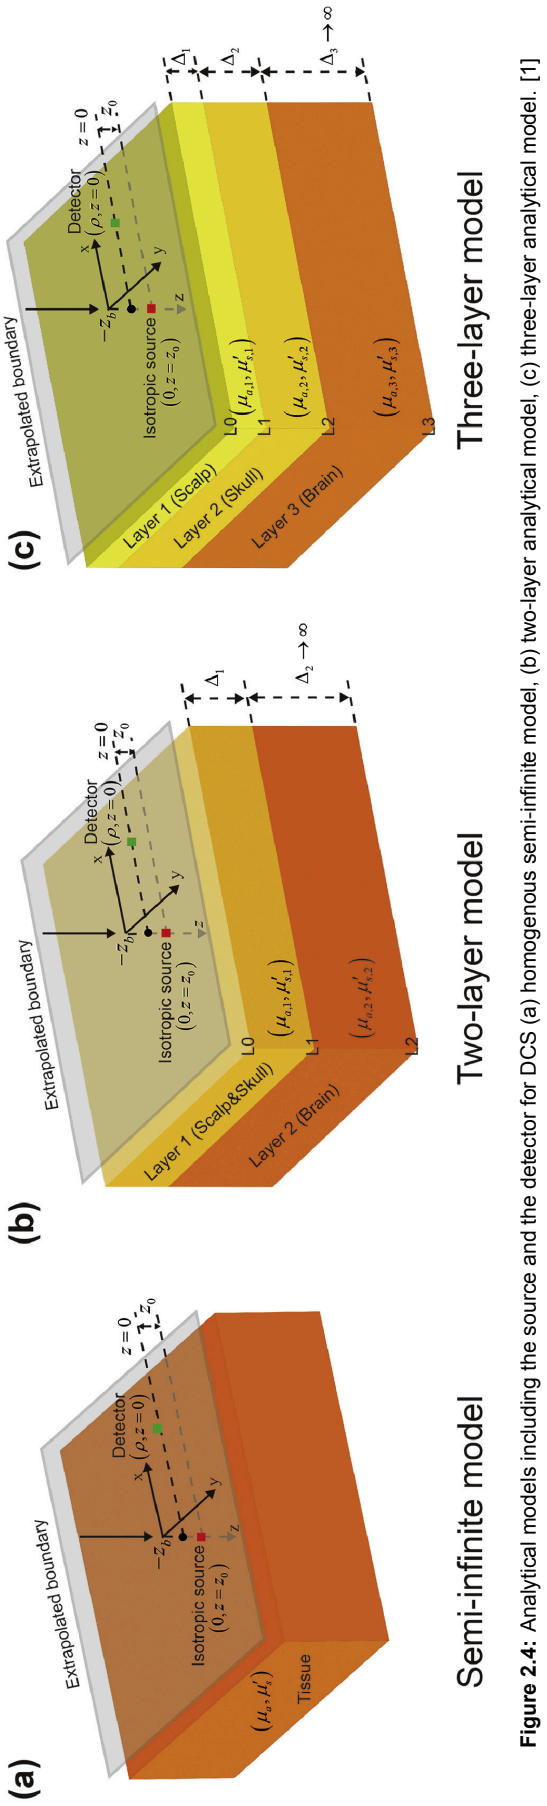

Figure 2.4: Analytical models including the source and the detector for DCS (a) homogenous semi-infinite model, (b) two-layer analytical model, (c) three-layer analytical model. [1]

3

Simulation

This chapter presents the simulation framework developed to evaluate the modified Born series as a solver for the Helmholtz equation in the context of diffuse correlation spectroscopy. The objective is to determine whether Brownian motion can be extracted from simulated optical signals and to assess the ability of the model to separate Brownian motion from laminar flow effects.

The framework combines two components: (i) Brownian motion trajectories generated in COMSOL by solving the overdamped Langevin equation, and (ii) electric field distributions solved using the modified Born series. The simulated fields are used to interpolate temporal autocorrelation functions, which are then fitted to theoretical models based on the correlation diffusion equation.

The chapter is organized as follows: the medium and source setup are introduced first, followed by the Born series solver and the integration of dynamic scatterers. The calculation of autocorrelation functions and the fitting procedure for Brownian motion and laminar flow are then described. Finally, the framework is scaled to larger domains to reach the diffusive regime and a test plan is outlined.

3.1. Medium setup

The simulated domain consists of a homogeneous tissue-like slab, representing a cross-sectional slice with dynamically moving scatterers. To capture fine wave-propagation details at a resolution of four points per wavelength (PPW) for a vacuum wavelength of $0.8\mu m$, the resolution for the discretisation of the grid was set at $\frac{\lambda}{4}$, yielding a pixel size of $0.2\mu m$. A single–pixel–thick air layer ($n = 1.0$) is appended above and below the slab. To simulate the static properties of the medium, the complex refractive index is defined as:

$$n_{slab} = n_0' + \mathrm{i}\, n_0'' \tag{3.1}$$

Here, n_0' models the background medium, while the imaginary part n_0'' ensures uniform absorption throughout the slab. In terms of the absorption coefficient μ_a, one has $\mu_a = 4\pi n_0''/\lambda_0$, which controls the exponential attenuation of intensity within the slab [54].

The source illuminating the medium is defined as a narrow Gaussian beam with a waist w_0. To mimic the behaviour of a point source autocorrelation placed at a distance equal to the mean free transport path, the beamwidth of the source needs to be much less than the thickness of the sample (L), i.e., $\delta = \frac{w_0}{L} \leq 0.1$ [4]. A figure of the simulated source profile can be seen in Figure 3.1.

The source is placed at the bottom of the grid inside the air layer and propagates in the y-direction through the medium and is detected at the top inside the air layer. This methodology simulates the transmission geometry, which gives the ability to measure particle movements much smaller than the wavelength of the source. The transmission geometry is also more accurately described within the photon diffusion approximation because all the paths that light can take through the medium are long compared to the mean free transport path. A concept of the domain and the propagation of the source through the domain can be seen in Figure 3.3.

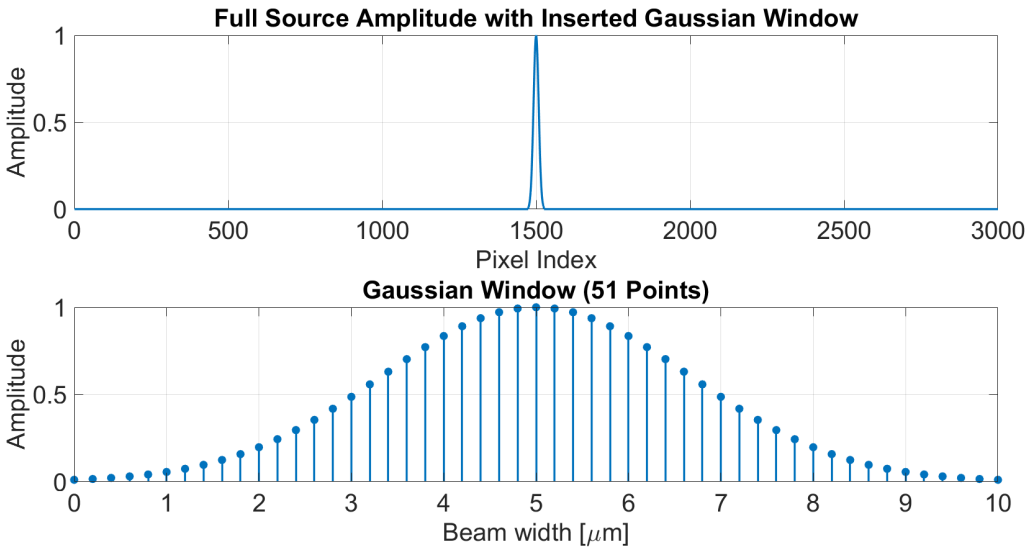

Figure 3.1: Gaussian source profile used in the WaveSim simulations. (a) Amplitude distribution along the x-axis of the medium in pixel units, showing a Gaussian window of 51 pixels centred at the grid midpoint. (b) Amplitude profile, with the 51-pixel beam width mapped to micrometres.

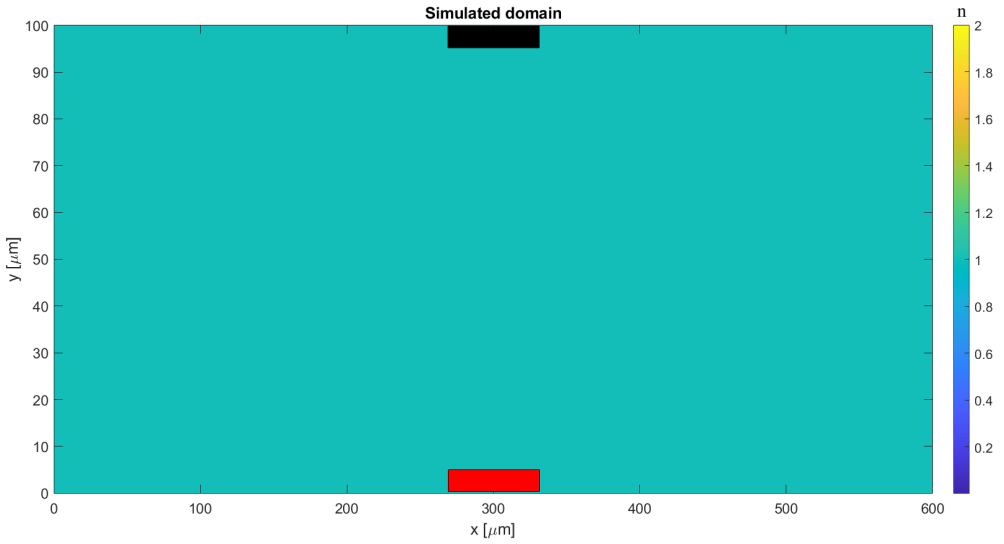

Figure 3.2: Illustration of the simulated domain without scatterers, with physical dimensions of $100\mu m \times 600\mu m$. The real part of the refractive index distribution is plotted ($n = 1.0 + i \times 0.001$). The red box indicates the source location, while the black box marks the detector position.

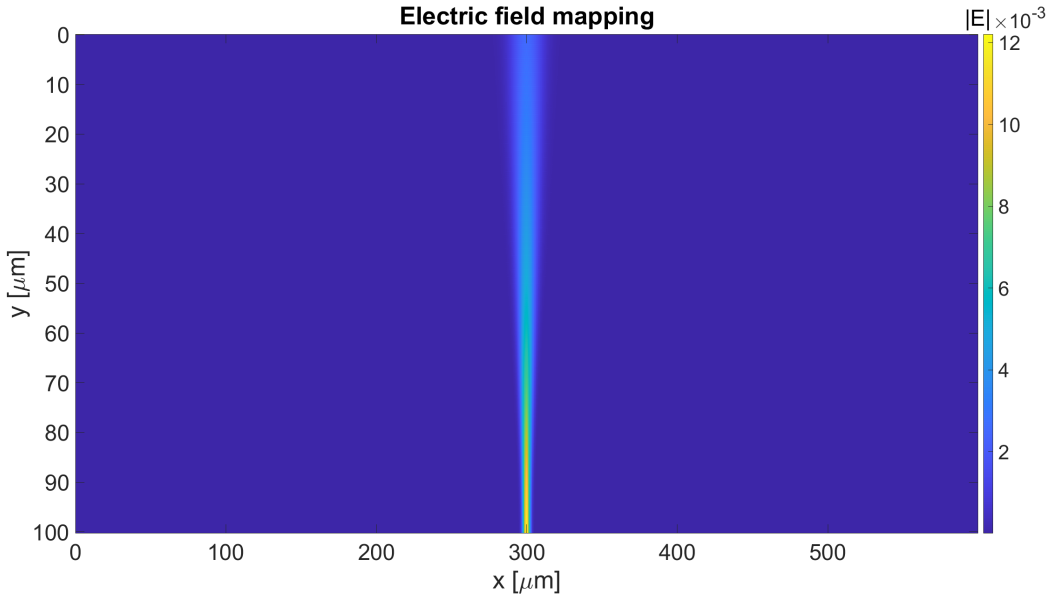

(a)

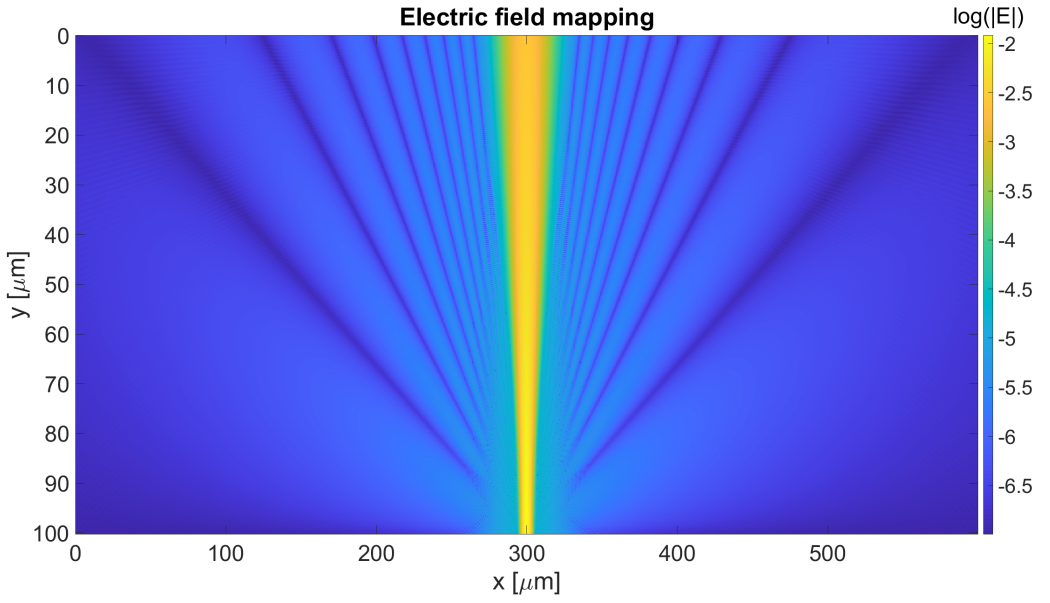

(b)

Figure 3.3: Field distribution of the Gaussian input beam in the simulated domain. (a) Absolute electric field amplitude $|E|$, showing the beam propagation profile. (b) Logarithmic scale of the absolute electric field, $\log(|E|)$, highlighting weaker field regions and the full dynamic range of the propagation.

3.1.1. Born Series Solver

The optical propagation through the medium was solved using the modified Born series method proposed by Osnabrugge *et al.* [9]. This solver provides an iterative scheme for the inhomogeneous Helmholtz equation that is unconditionally convergent due to the inclusion of a preconditioner. At each iteration, the wave field is updated using the scattering potential of the medium and the Green's function of the background. The update rule can be written as:

$$E_{k+1} = ME_k + \gamma GS, \tag{3.2}$$

where M is a preconditioned operator containing the scattering potential V(r), G is the Green's function operator, and γ is a stabilisation factor chosen such that convergence is guaranteed. Full definitions of these terms are given in [9]. Physically, each iteration of the solver corresponds to a pseudo-propagation step: the wave packet is displaced forward by a distance proportional to $2k_0/\epsilon$. While the method is time-independent, this iterative scheme allows us to sample scattered paths until the wave has effectively reached the detector. In the present implementation, independent realisations of the medium were generated for each timestep of the COMSOL Brownian motion simulation. This allowed the temporal autocorrelation function to be reconstructed from a sequence of steady-state field solutions.

3.1.2. Dynamic Scatterers

Embedded within the medium are scatters, modelled as circular disks (in 2-D) of radius $r_{rbc} = 4\mu m$. Each scatterer is assigned a refractive index to simulate static properties:

$$n_{scat} = n'_{scat,real} + i * n''_{scat} \tag{3.3}$$

The contrast between n'_{scat} and n'_0 generates scattering, while the imaginary term n''_{scat} accounts for weak absorption within the cells. These scatterers follow dynamic trajectories extracted from Brownian motion simulations. At each timestep, their 2-D positions are mapped onto the computational grid and the local portion of the dielectric map is updated accordingly. In Figure 3.4, the refractive index mapping and the resulting electric field mapping are plotted for a small medium. Comparing the electric field mapping with and without scatters, it can be seen that the scatterers have created more scattered light paths within the medium, confirming that the model behaves as expected. To obtain a temporal field autocorrelation graph, the scatterers need to move inside the medium, causing the path lengths of the scattered light that reach the detector to change over time, resulting in a fluctuating measured intensity plot over time.

In this study, the motion of each scatterer was simulated by solving the overdamped Langevin equation within COMSOL's Particle Tracing for Fluid Flow interface:

$$-6\pi\eta r_p v_i(t) + F_{B,i}(t) = 0, \tag{3.4}$$

where the drag force follows Stokes' law and the Brownian force $F_{B,i}(t)$ has zero mean and Gaussian statistics. The simulation domain was defined as a $6000\ \mu m \times 1000\ \mu m$ rectangular chamber filled with quiescent fluid (velocity set to zero), in which scatterers were released at $t = 0$ from a uniform grid of initial positions. The computational mesh employed free triangular elements with a maximum edge length of $0.05, \mu$m, ensuring accurate resolution of the $8, \mu$m-diameter particles as they experienced Brownian and drag forces. Drag was modelled according to Stokes' law,

$$F_D = -6\pi\eta r_p v, \tag{3.5}$$

while the Brownian force was generated using a reproducible random seed (ds = 3112) and satisfied the relation:

$$F_B = \frac{2k_B T}{6\pi\eta r_p} \tag{3.6}$$

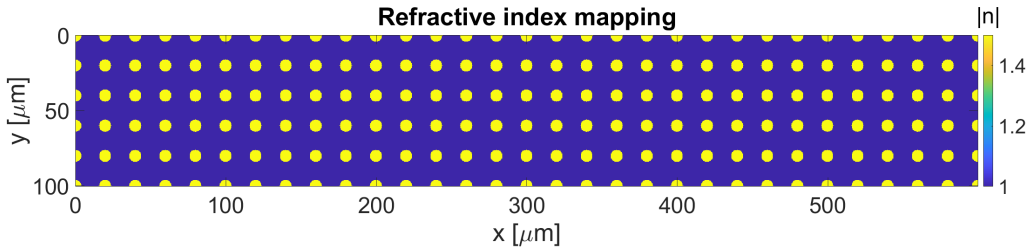

(a)

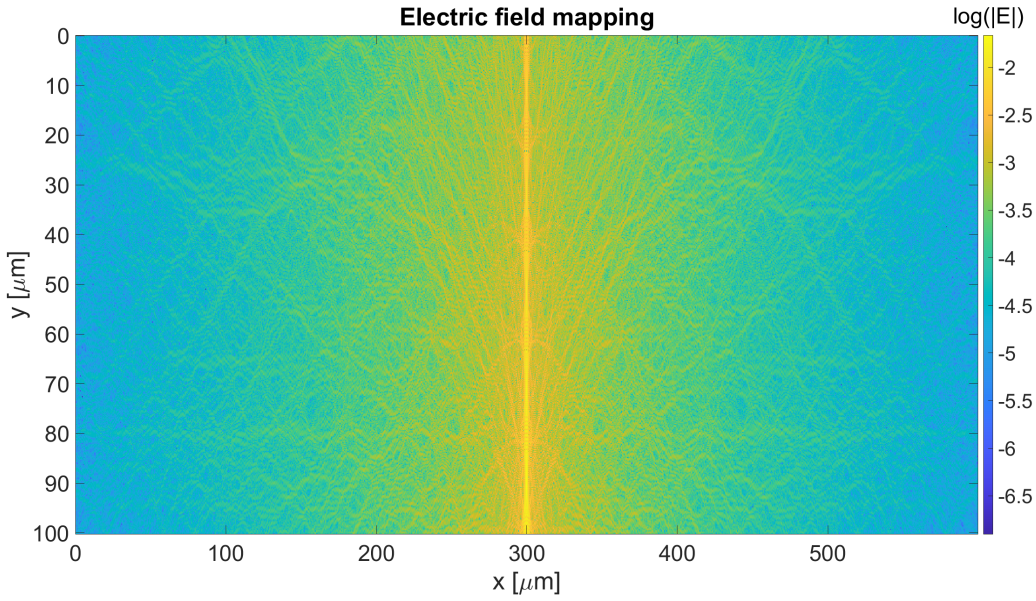

(b)

Figure 3.4: Refractive index distribution and resulting field map in the scattering domain. (a) Real part of the refractive index, showing scatterers with $n = 1.5 + i \times 0.0001$ embedded in the background medium. (b) Logarithmic mapping of the absolute electric field, $\log(|E|)$, illustrating the scattering pattern generated by the Gaussian input beam.

In these equations, k_B is the Boltzmann constant, T is the temperature, η is the fluid viscosity, r_p is the particle radius and v is the particle velocity. The boundary interactions were governed by no-slip, diffuse-scattering conditions. The transient problem was solved in uniform time-steps using the Generalised-α method.
In diffuse optics, the Brownian motion has a linear relation to the mean-squared displacement of the particles.

$$\langle r(\tau)^2 \rangle = 4D_b\tau \tag{3.7}$$

As a sanity check, the simulated positions obtained from the COMSOL simulation are checked to see if this linearity holds. The mean-squared displacement (msd) is obtained by:

$$\begin{aligned} dy(\tau) &= y(\tau+1) - y(1) \\ dx(\tau) &= x(\tau+1) - x(1) \\ dr^2(\tau) &= dx^2(\tau) + dy^2(\tau) \\ msd(\tau) &= \frac{dr^2(\tau)}{N} \end{aligned} \tag{3.8}$$

The resulting mean-squared displacement plot can be seen in Figure 3.5. In this figure, the linear relation between the msd and time can be seen. For this case, the simulated Brownian motion within COMSOL is $D_b = 6.782 \times 10^{-8} \frac{m^2}{s}$ and fitting the msd data from Equation 3.8 into Equation 3.7 gives a Brownian motion of $D_b = 6.6115 \times 10^{-8} \frac{m^2}{s}$.

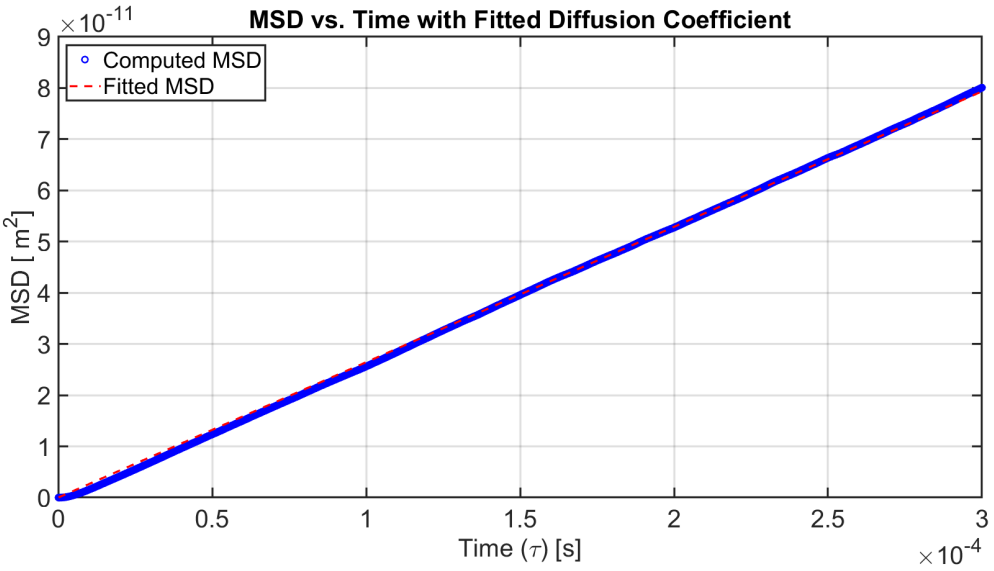

Figure 3.5: Mean squared displacement (MSD) of COMSOL-simulated particle trajectories (blue) compared with the theoretical fit $4D_B\tau$ (red). The agreement confirms Brownian motion behaviour and validates the trajectory generation.

Another sanity check is a parameter sweep, checking that increasing the temperature increases the Brownian motion. Starting with a base temperature of 25 K, and increasing the temperature to 300 K results in Figure 3.6. In this figure, the linear increase in temperature shows a linear increase in Brownian motion. This confirms that the COMSOL model performs as expected. With the COMSOL model, a simulation can be done on moving scatterers to check if the moving scatterers cause a measured intensity fluctuation. In Figure 3.7, an intensity plot can be seen for a small domain where the scatterers started at a grid position and their positions on the grid were updated for each next time-step. The resulting plot shows that the model correctly shows that moving the scatterers on the grid results in a fluctuating measured intensity.

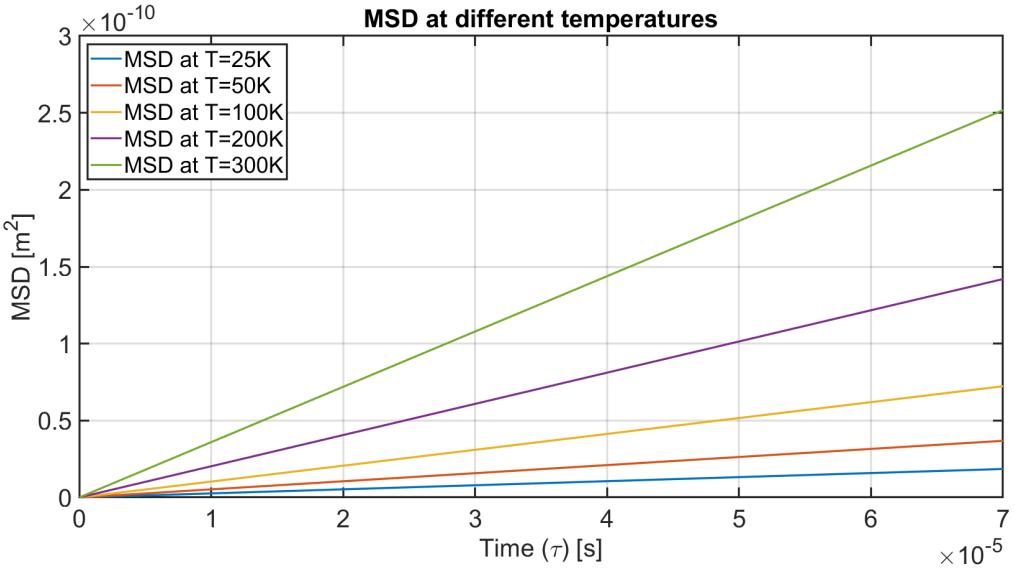

Figure 3.6: Mean squared displacement (MSD) plots computed from COMSOL particle trajectories at different temperatures. The increasing slope with temperature reflects the linear dependence of the Brownian diffusion coefficient on T, with higher temperatures producing faster particle dispersion.

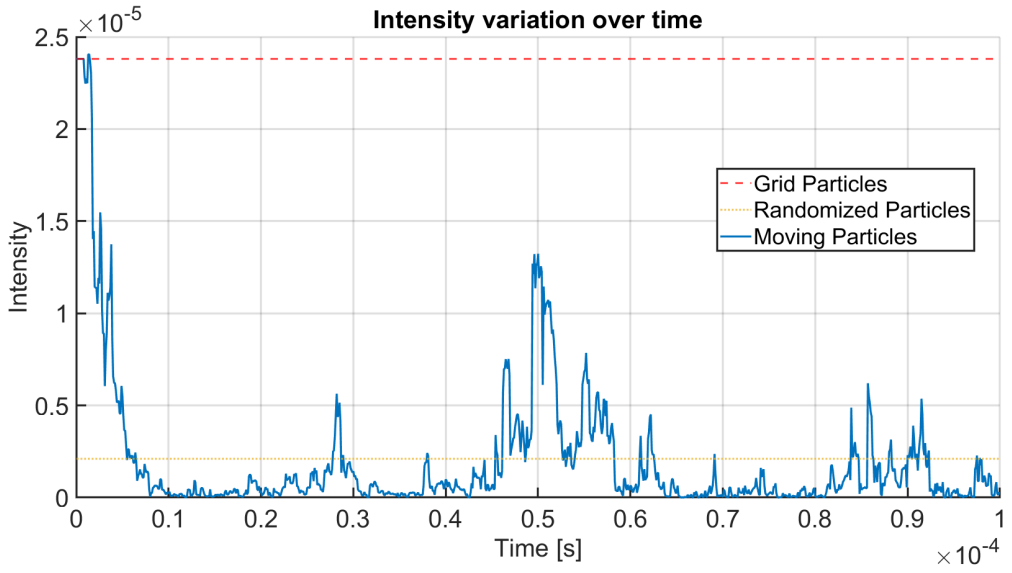

Figure 3.7: Simulated detector intensity signals for different particle configurations. The red line shows the baseline case with particles placed on a fixed grid, the yellow line corresponds to particles at randomized fixed positions, and the blue plot represents particles undergoing Brownian motion. The temporal fluctuations in the moving-particle case reflect the speckle dynamics that underpin DCS measurements.

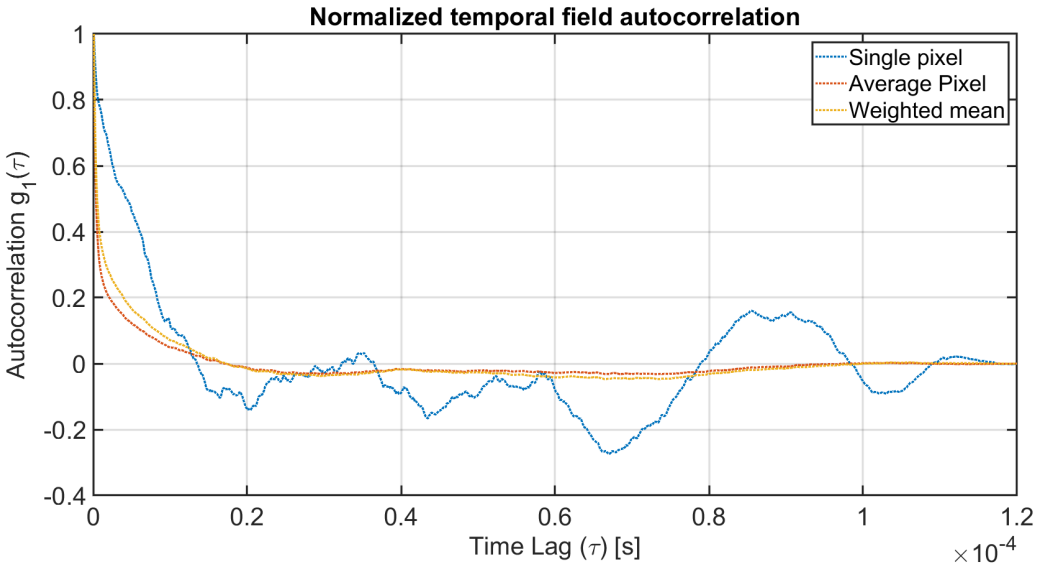

Figure 3.8: Normalized temporal autocorrelation functions computed from WaveSim outputs under different averaging schemes. The blue curve shows the autocorrelation from a single pixel, while the red curve represents the mean over 500 pixels. The yellow curve applies a weighted mean across the same 500 pixels. Averaging across pixels substantially reduces noise compared to the single-pixel trace, and the weighted mean further improves stability of the correlation decay.

3.2. Autocorrelation

The autocorrelation of the fluctuating field is needed to obtain information on the moving particles. The autocorrelation measures the degree of similarity between a time series and its lagged version. In optics, the field autocorrelation function of the complex field $E(t)$ is given by:

$$A(\tau) = \int_{-\infty}^{+\infty} E(t)E(t-\tau)dt \tag{3.9}$$

In the NIRS spectral window, the field autocorrelation is defined as the correlation between the electric field $E(t)$ and its delayed replica $E(t-\tau)$. For the case of diffuse optics, the unnormalized field autocorrelation $G_1(\tau)$ can be expressed as:

$$G_1^T(r,\hat{\Omega},\tau) = \langle E(r,\hat{\Omega},\tau) \cdot E^*(r,\hat{\Omega},\tau) \rangle \tag{3.10}$$

In this equation, $E(r,\hat{\Omega},\tau)$ is the electric field at position r and time τ propagating in the $\hat{\Omega}$ direction, inside the domain and $\langle \dots \rangle$ denotes the time average of the product of the electric field and its delayed replica [1]. In this research, the autocorrelation is obtained by measuring the complex electric field at the opposite end from the source at a separation distance of $\rho = 0$ for each time-step. As shown in Figure 3.8, the autocorrelation from a single pixel exhibits oscillations and noise. This behaviour arises because a single speckle provides only one independent observable (NIO), and the number of sampled decorrelation events is not statistically sufficient to represent the entire decorrelation process. In other words, each speckle fluctuates randomly; with only one speckle, these random fluctuations dominate, leading to a jagged autocorrelation curve. To improve the signal-to-noise ratio (SNR), autocorrelation functions from neighbouring pixels are included and averaged. Since neighbouring pixels correspond to slightly different source–detector separations, a weighted mean is applied to reduce their influence. This procedure effectively increases the number of independent observables, as the SNR of decorrelation time measurements scales with $\sqrt{\text{NIO}}$. The approach is consistent with the framework of Xu *et al.* [55], who demonstrated that temporal sampling and speckle-ensemble averaging provide equivalent descriptions of field autocorrelation in diffusing wave spectroscopy. Accordingly, the use of neighbouring pixels can be interpreted as an ensemble averaging strategy that increases NIO, averages out random intensity fluctuations, and yields a smoother and more statistically robust estimate of $g_1(\tau)$.

3.2.1. Exponential fitting

In Diffusion correlation spectroscopy, the autocorrelation data is fitted to the solution of the correlation diffusion equation. For example, in the case of a semi-infinite homogeneous medium with the backscattering geometry, the solution to the CDE would be [56]:

$$G_1(\rho, t) = \frac{3\mu_s'}{4\pi}[\frac{\exp(-k_D r_1)}{r_1} - \frac{\exp(-k_D r_2)}{r_2}] \tag{3.11}$$

where,

$$k_D^2 = 3\mu_a\mu_s' + \mu_s'^2 k_0^2 \alpha \langle \Delta r^2(\tau) \rangle \tag{3.12}$$

By fitting the correlation data onto the solution for the CDE, the mean squared displacement can be determined, which is linear to the Brownian coefficient. In this case, the optical properties, such as the absorption and scattering coefficients, are unknown. Without accurate measurements of these parameters, the Brownian motion coefficient will be inaccurate. For the transmission geometry, it is known that all autocorrelation functions decay with a characteristic decay time of approximately $(l^*/L)^2\tau$. The average decay of the autocorrelation can be described with $\exp(-2t/\tau)$ [4]. The numerical values of the autocorrelation are insensitive to the exact choice of l^*, since $l^* << L$. The choice of l^*, would only affect the first few scattering interactions of a large number of scattering. Thus, the autocorrelation for a transmission geometry with point illumination can be described with:

$$g_1(\tau) \approx \exp(-t/\tau) \tag{3.13}$$

with,

$$\tau = r(l^*/L), \quad r = (D_{\text{diff}} k_0^2)^{-1} \tag{3.14}$$

With this, the autocorrelation is going to be fitted onto an exponential function:

$$g_1(\tau) = B * \exp(-\sqrt{A\tau}) + C \tag{3.15}$$

where A describes the decay rate of the exponential decay, and B and C are for normalisation and correct placement. The decay rate from the exponential function contains the Brownian coefficient, and its relation to A can be described with:

$$A = \frac{1}{\tau} = \frac{1}{r}(\frac{L}{l^*})^2 = (L/l^*)^2 k_0^2 D_{\text{diff}} = 4(L/l^*)^2 k_0^2 D_B \tag{3.16}$$

Laminar flow

The influence of laminar flow on the autocorrelation decay has been studied before; in this study, the born series will be tested to see if it is possible to separate the autocorrelation influenced by flow from the autocorrelation signal influenced by Brownian motion. With a flow speed, the particles have an additional velocity vector that influences the MSD. In this case, the MSD is given by:

$$\langle r(\tau)^2 \rangle = 4D_b\tau + v_{scat}^2\tau^2 \tag{3.17}$$

With the additional velocity term within the msd, the theoretical autocorrelation becomes:

$$g_1(\tau) = B * \exp(-\sqrt{A\tau + D\tau^2}) + C \tag{3.18}$$

where D describes the influence of the flow velocity on the decay rate. D can be described with:

$$D = (L/l^*)^2 k_0^2 v^2 \tag{3.19}$$

3.3. Full Scale Medium

In the previous subsections, a simplified example medium was used to verify the fundamental working principles of the simulation. However, in order to apply the diffusion approximation, it is essential that the detected light lies within the diffusive regime. As shown in Figure 3.4, the light reaching the detector in the smaller domain has not undergone sufficient scattering to be considered diffused. Initially, the light is anisotropic, meaning it predominantly propagates in the forward direction. As scattering

events accumulate, the propagation direction becomes increasingly randomised. When the angular distribution of photons becomes approximately uniform, the light is considered isotropic and can be described using the diffusion approximation.

To reach this regime, photons must undergo multiple scattering events, which requires a longer propagation path within the medium. This can be achieved by increasing the size of the computational domain. In Figure 3.9, the magnitude of the electric field, $|E|$, is shown across the larger domain for cases without scatterers (a) and with scatterers (b). In the presence of scattering particles, the field exhibits a more spatially spread and smoother distribution, indicating that the light has experienced sufficient scattering and can be considered to be in the diffusive regime.

As a validation step, the corresponding autocorrelation function was computed. Figure 3.10 shows that the autocorrelation decays as expected from DCS theory.

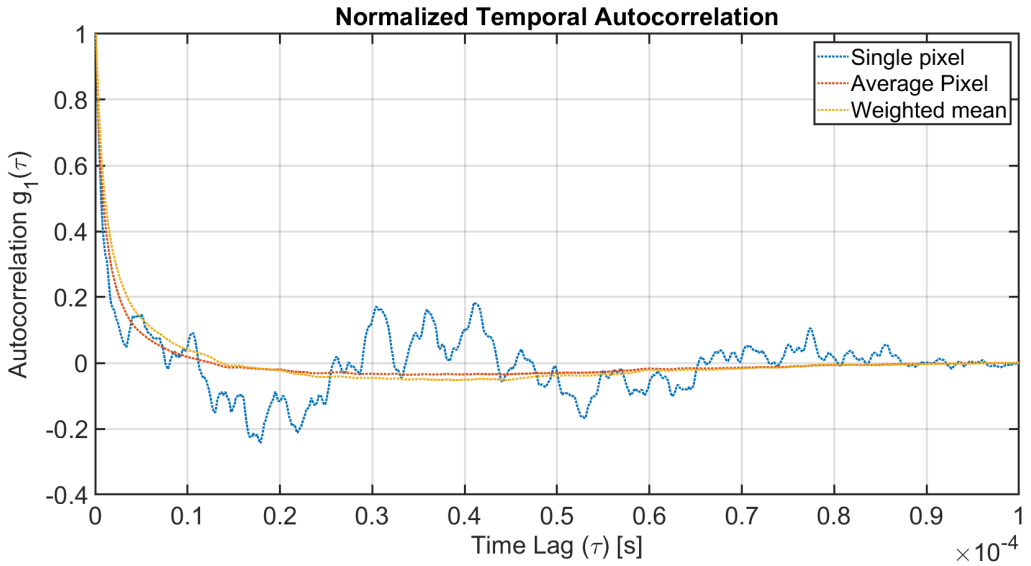

Figure 3.10: Normalized temporal autocorrelation functions from WaveSim for the full-scale simulation domain. The blue curve corresponds to the autocorrelation from a single pixel, while the red and yellow curves show the mean and weighted mean over 500 pixels, respectively. As in the smaller domain, ensemble averaging across many pixels improves the smoothness and reliability of the correlation decay.

3.4. Test plan

With the full-scale medium established in the diffusive regime, a series of tests were designed to verify the simulation framework and assess its ability to recover particle dynamics. Three categories of validation were performed:

1. Temperature dependence of Brownian motion; The first test examines whether increasing temperature leads to a corresponding increase in the autocorrelation decay rate, as predicted by the relation between mean squared displacement and the Brownian coefficient (Equation 3.7). Successful reproduction of this trend confirms that the model captures thermal diffusion.
2. Laminar flow contribution; In the second test, a uniform flow velocity is added to the particle trajectories in COMSOL. The expectation is that laminar flow introduces an additional quadratic term in the mean squared displacement (Equation 3.18), resulting in faster decorrelation. Simulations at multiple flow speeds test whether this contribution can be distinguished from purely Brownian motion.
3. Extraction of decay rates; Fitting the data from the first and the second test onto the exponential models. The expectation is that the decay rates will have a linear relation with temperature for Brownian motion and a non-linear relation for laminar flow.

Together, these tests establish whether the combined COMSOL–Born series framework can reliably

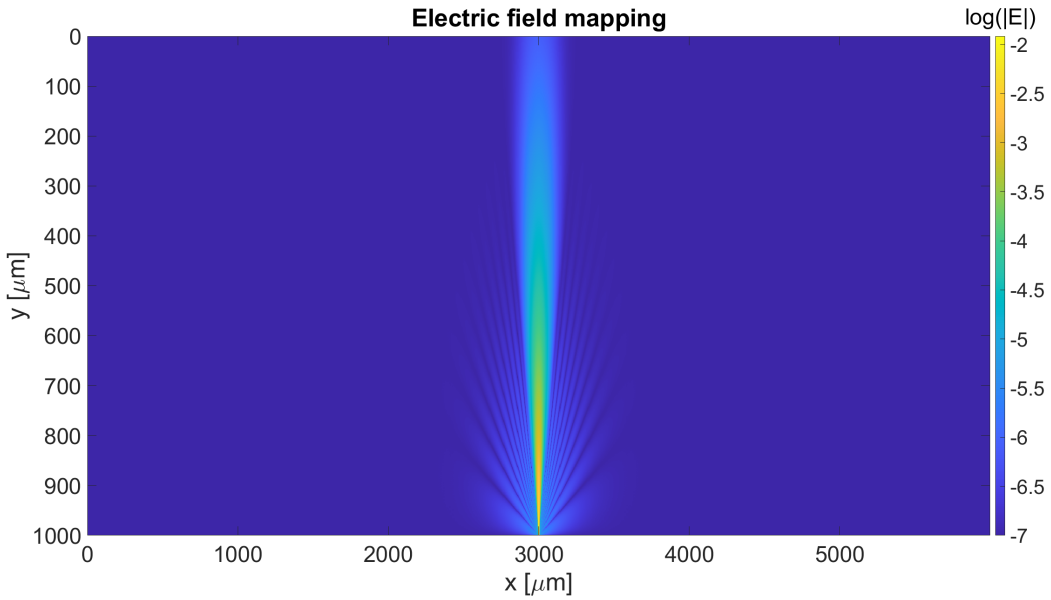

(a)

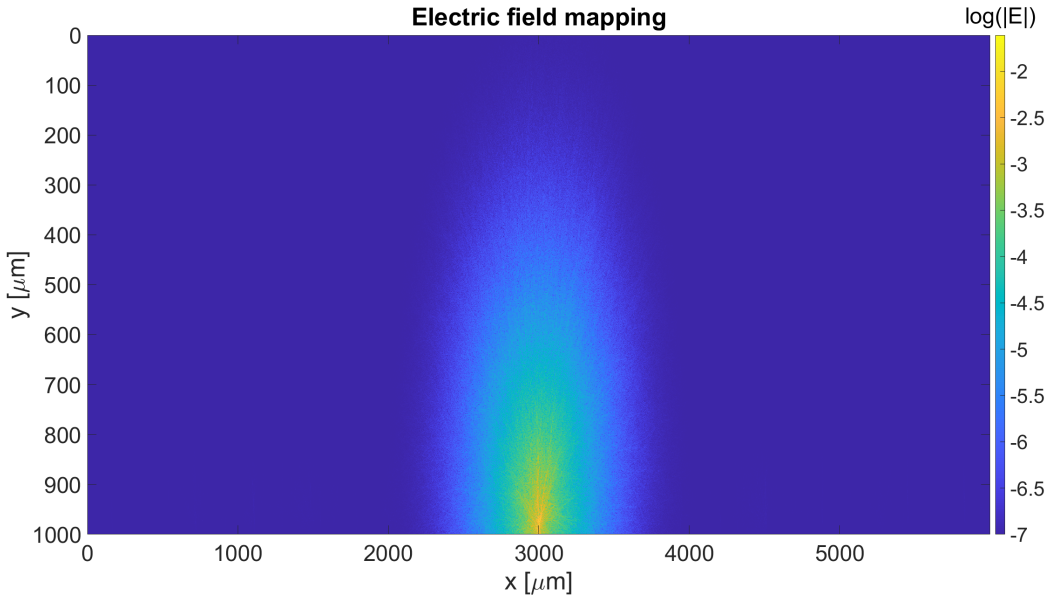

(b)

Figure 3.9: Mapping of the logarithmic absolute electric field distribution in the full-scale simulation domain. (a) Domain without scatterers, where the Gaussian source propagates smoothly through the medium. (b) Domain with scatterers

extract both Brownian and flow dynamics from simulated autocorrelation data. A schematic overview is shown in Figure 3.11.

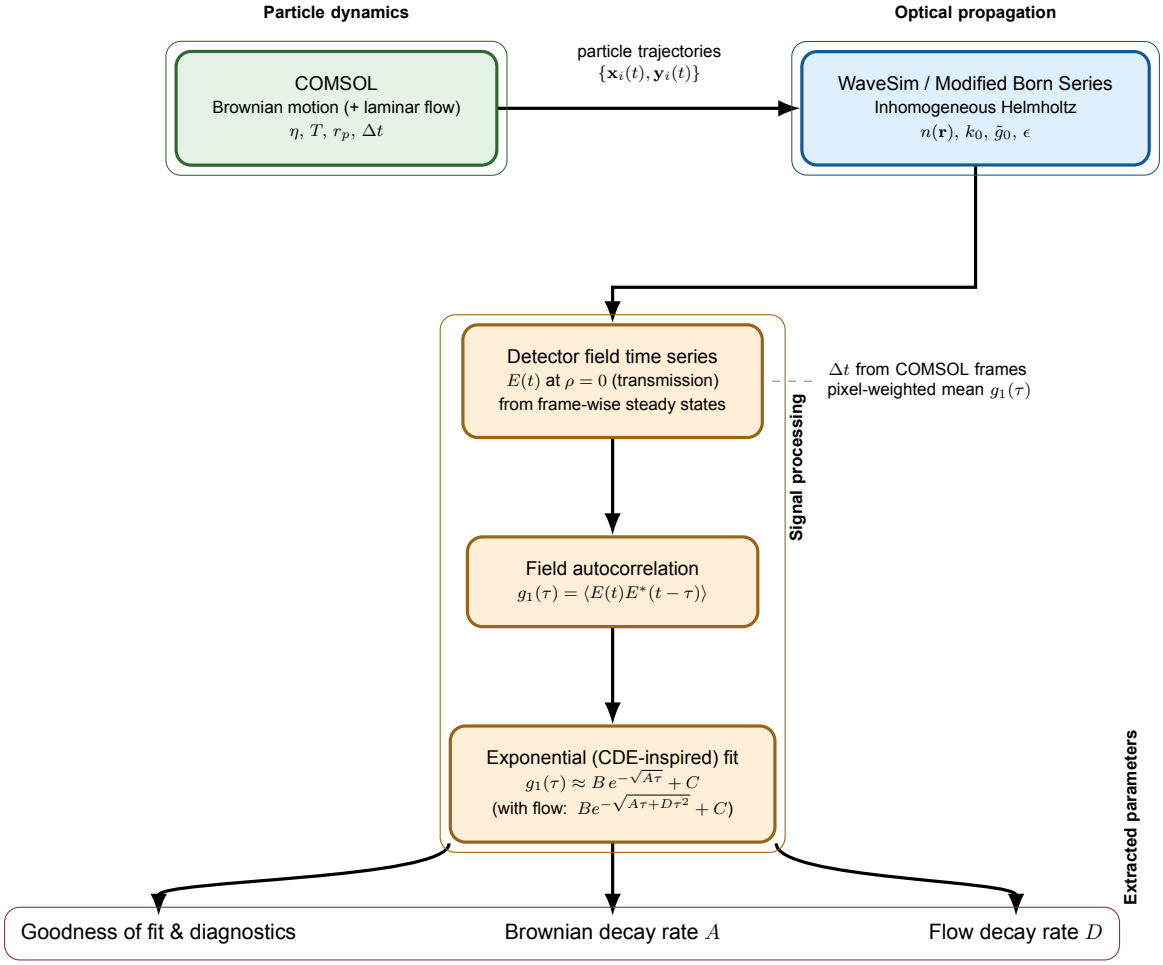

Figure 3.11: Schematic overview of the simulation workflow. Brownian motion trajectories are generated in COMSOL and mapped onto the optical grid of the modified Born series solver to compute the scattered field. Temporal fluctuations at the detector are converted into field autocorrelation functions and fitted to exponential models inspired by the correlation diffusion equation to extract Brownian and flow parameters.

4

Results

This chapter presents the results obtained from the coupled COMSOL–Born series simulation framework described in chapter 3. The aim is to evaluate whether the model reproduces the expected behaviour of diffusing correlation spectroscopy (DCS) and can separate Brownian and flow contributions to the autocorrelation decay. The results are organised according to the validation tests outlined in the methodology. Each section first presents the particle trajectories from COMSOL as mean squared displacement (MSD) plots, followed by the corresponding autocorrelation functions from WaveSim and their exponential fits.

4.1. Brownian Motion

4.1.1. Main case

The first validation case considered a viscosity of 1.0×10^{-5} Pa·s, with the temperature varied from 300K to 1800K in steps of 300 K. The particle trajectories generated in COMSOL show a linear mean squared displacement (MSD) as a function of time across all temperatures, confirming Brownian behaviour consistent with Equation 3.7. The MSD results across all temperatures are shown in Figure 4.1. The results show that the MSD slopes increase linearly with temperature as expected. These trajectories were then used, in the WaveSim model to compute the corresponding electric field for each frame and then formed the corresponding temporal autocorrelation functions. The resulting autocorrelation curves, shown in Figure 4.2, show a clear trend of faster decay at higher temperatures, reflecting shorter decorrelation times as particle motion increases. However, the increase in decay rate is not uniform across the entire temperature range.

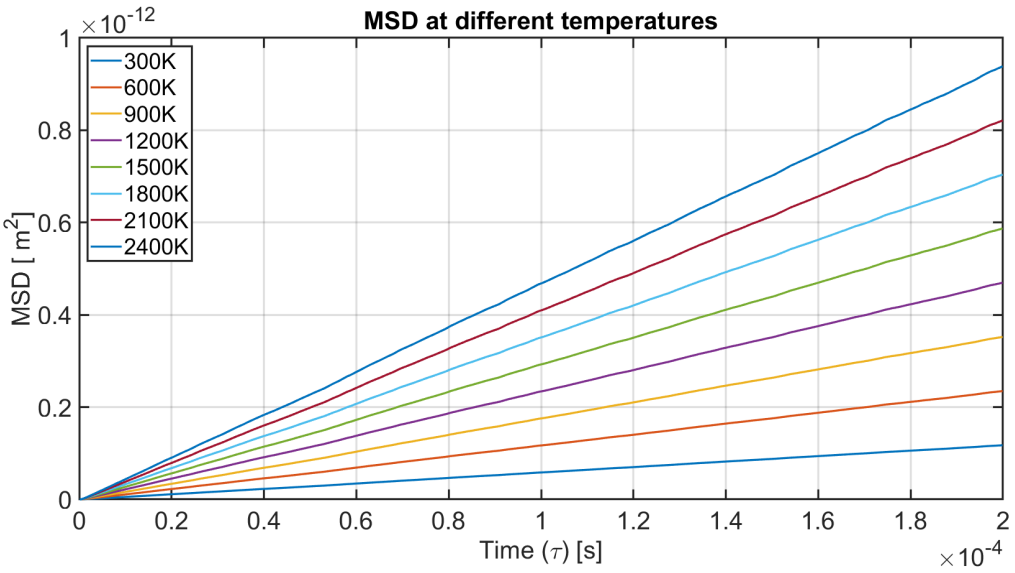

Figure 4.1: Mean squared displacement for $T = 300$–2400 K at viscosity 1.0×10^{-5}Pa·s. The MSD shows a linear relation with time across all cases.

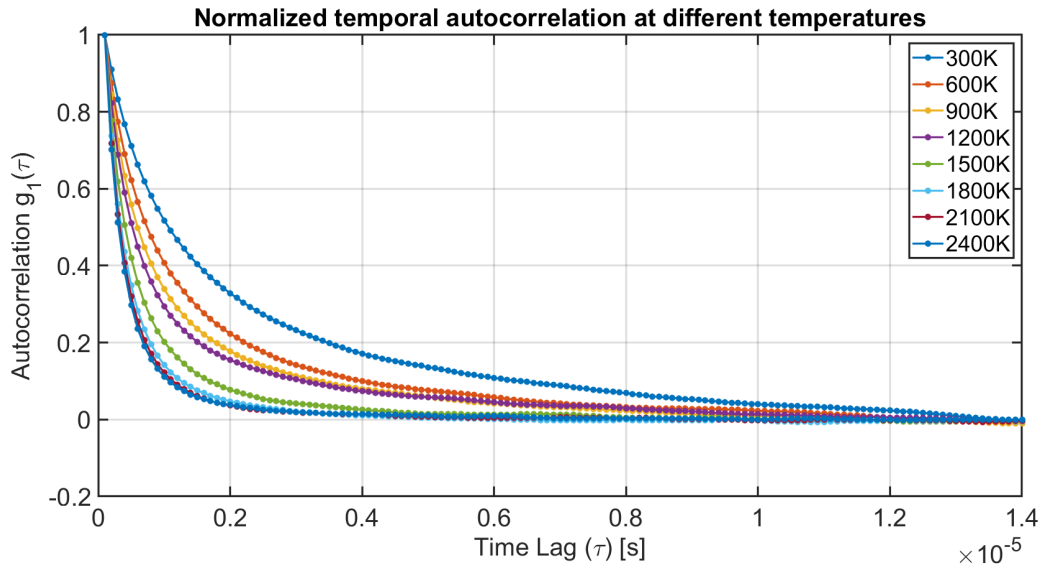

Figure 4.2: Normalized temporal autocorrelation functions using a 500-pixel weighted mean, shown for increasing temperatures between 300 K and 2400 K. Higher temperatures lead to faster decorrelation, consistent with the expected acceleration of Brownian dynamics.

To quantify this effect, the autocorrelation curves were fitted to the exponential model in Equation 3.15. The extracted decay coefficients are plotted in Figure 4.3. The results show a clear non-linear dependence on temperature, with the growth of the decay rate coefficient increasing at higher T. At low temperatures (300–600 K), the decay rate increased gradually, whereas at higher temperatures (1200–1800 K) the increments became much larger. This behaviour suggests that, although the MSD scales linearly with temperature as expected, the mapping between MSD and autocorrelation decay becomes increasingly inconsistent as temperature rises.

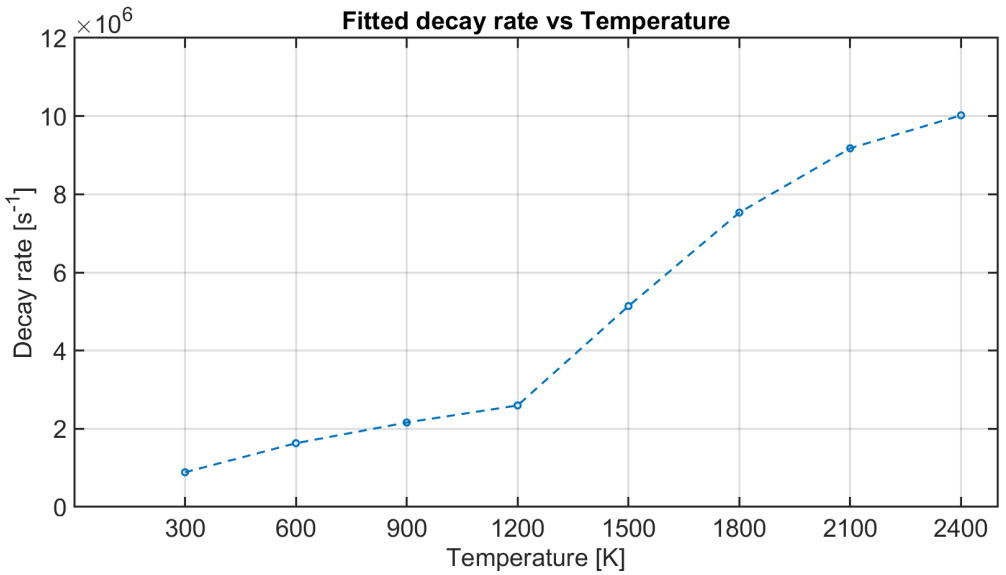

Figure 4.3: Extracted decay rates from exponential fits using Equation 3.15, plotted as a function of temperature.

4.1.2. Supporting Case 1

A second set of simulations was performed with viscosity increased to 1.0×10^{-4} Pa·s, and temperature varied between 280 K and 400 K. The MSDs remained linear with time, confirming Brownian motion. However, the corresponding autocorrelation functions (Figure 4.4) showed significant overlap between curves, with crossovers obscuring any clear separation by temperature.

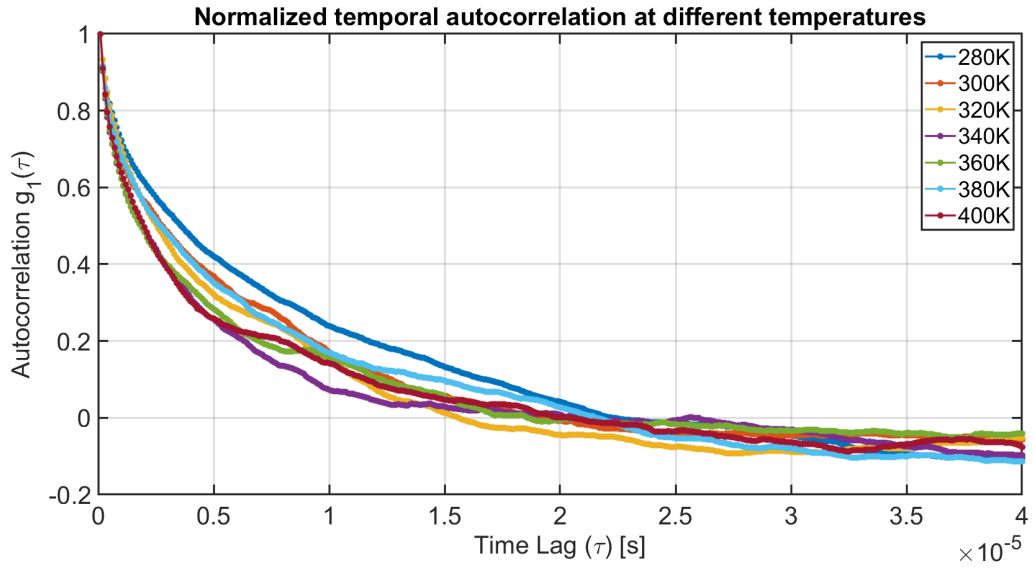

Figure 4.4: Normalized temporal autocorrelation functions using a 500-pixel weighted mean, shown for increasing temperatures between 280 K and 400 K.

When fitted to the exponential model, the extracted decay coefficients initially followed a weak linear increase between 280 K and 320 K. Beyond this range, however, the results became inconsistent: a sharp increase at 340 K, decreases at 360 K and 380 K, and another increase at 400 K. This irregular behaviour reflects the poor separation between autocorrelation curves and highlights the strong influence of noise at these parameters. As a result, no reliable linear relation with temperature could be established in this case.

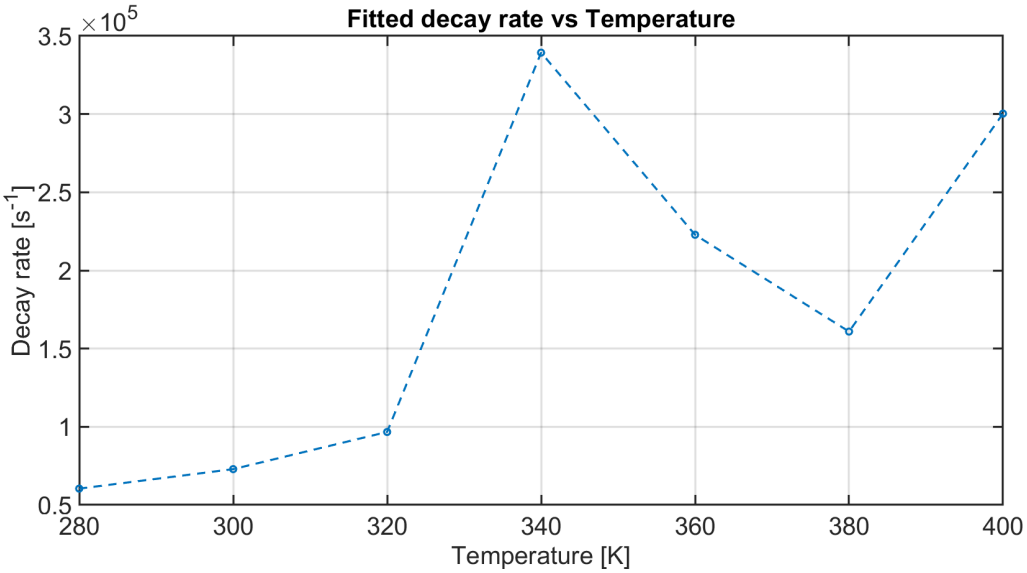

Figure 4.5: Extracted decay coefficients as a function of temperature.

4.1.3. Supporting Case 2

In the third set of simulations, the viscosity was reduced to 1.8×10^{-8} Pa·s, and temperatures between 25 K and 200 K were tested. Unlike the previous cases, the MSDs showed a non-linear increase with time, deviating from the expected linear Brownian behaviour. This nonlinearity carried through to the autocorrelation functions (Figure 4.7), which exhibited increasingly rapid decay with temperature but with a non-linear scaling pattern.

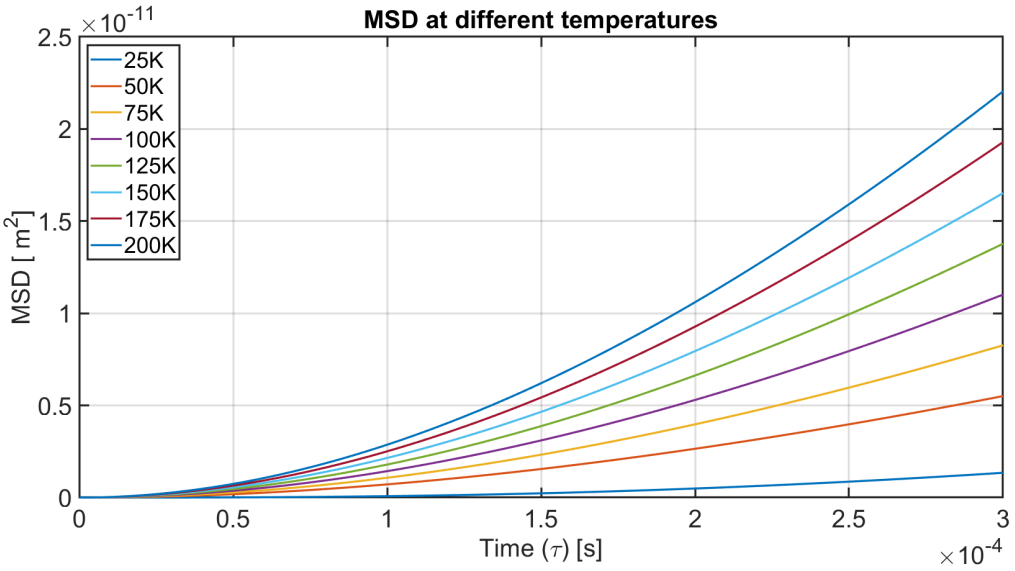

Figure 4.6: Mean squared displacement for $T = 25$–200 K at viscosity 1.8×10^{-8}Pa·s. The MSD shows a non-linear relation with time across all cases.

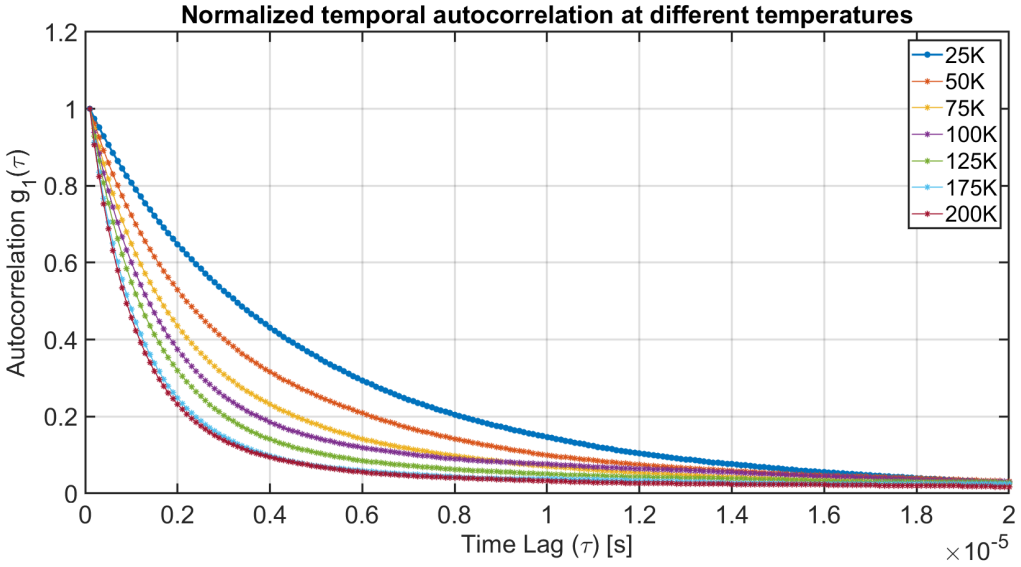

Figure 4.7: Normalized temporal autocorrelation functions using a 500-pixel weighted mean, shown for increasing temperatures between 25 K and 200 K (150 K missing due to corrupted data).

The extracted decay coefficients (Figure 4.8) confirmed this trend, displaying a non-linear dependence on temperature consistent with the underlying particle motion. This case demonstrates that when the fundamental dynamics deviate from Brownian theory, the optical signal reflects this departure directly.

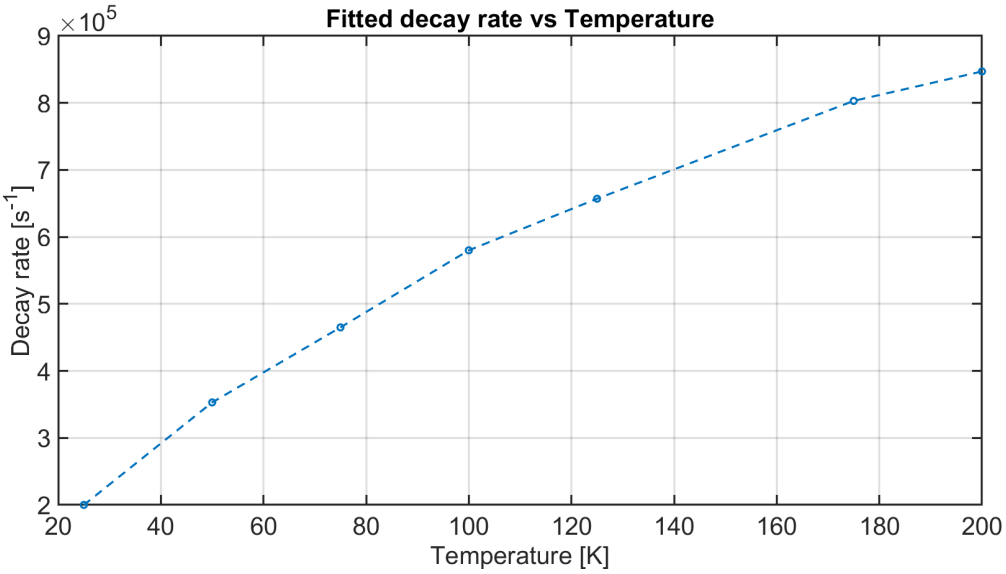

Figure 4.8: Extracted decay coefficients as a function of temperature.

4.2. Laminar Flow

To investigate whether the simulation can distinguish between Brownian motion and directed flow, a constant external force was added to each particle in COMSOL to generate a laminar flow profile. The particle parameters were kept identical to the Brownian motion case, with temperature fixed at 25 K and viscosity at $\eta = 1.8 \times 10^{-8}$. The flow speeds were simulated between 8×10^{-3} m/s to 1.8×10^{-2} m/s. The resulting MSD curves are shown in Figure 4.9. The addition of flow introduces a quadratic component that dominates at higher velocities. This behaviour is consistent with the theoretical MSD expression in Equation 3.7, extended with a flow term. At low flow speeds, the Brownian contribution

remains visible, but at higher speeds the quadratic growth becomes dominant.

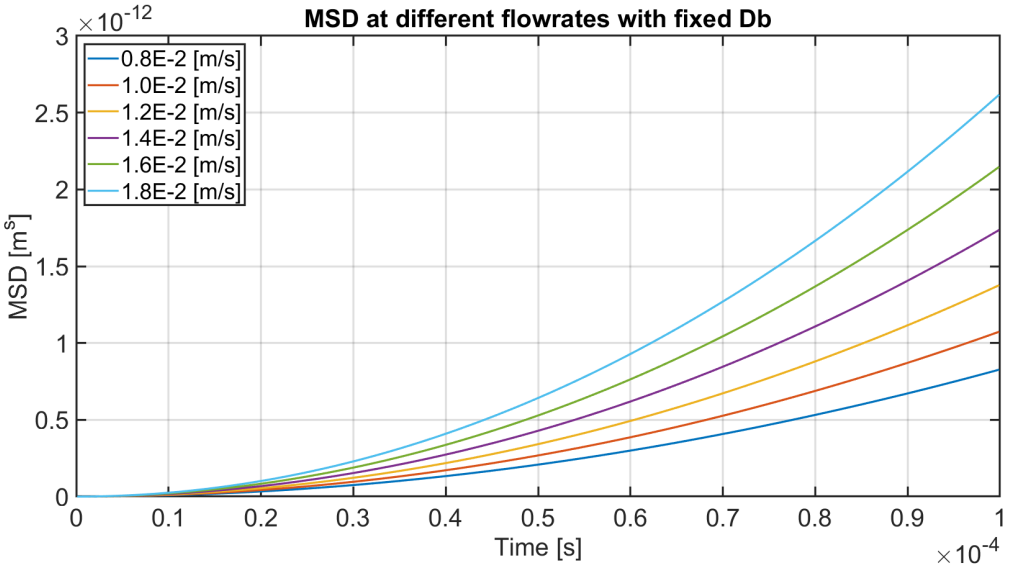

Figure 4.9: Mean squared displacement for $v = 8 \times 10^{-3}$ – 1.8×10^{-2} m/s at viscosity 1.8×10^{-8}Pa·s and $T = 25$ K. The MSD shows a non-linear relation with time across all cases.

The particle trajectories were supplied to WaveSim to calculate optical fields and temporal autocorrelation functions. The results in Figure 4.10 show that increasing flow speed produces faster autocorrelation decay, as expected. However, the increments are not uniform; higher flow speeds cause disproportionate shifts, and some curves partially overlap. This indicates that while flow influences the decay, the scaling with velocity is not showing a clear non-linear pattern. To extract quantitative flow parameters, the autocorrelation data were fitted to the extended exponential model in Equation 3.18. The extracted flow-related coefficients are shown in Figure 4.11. As expected, D increases with flow speed, confirming sensitivity to directed motion. However, the scaling was not strictly proportional to velocity, and noise in the higher-speed cases limited the precision of the fitted values. The results show a monotonic increase, confirming that higher flow speeds produce faster decorrelation. However, the scaling is not showing a clear non-linear relation: the increments were modest at intermediate speeds, followed by a large jump at the highest flow setting. This deviation from proportional scaling is consistent with the overlapping autocorrelation curves and reflects the difficulty of extracting flow parameters with quantitative precision.

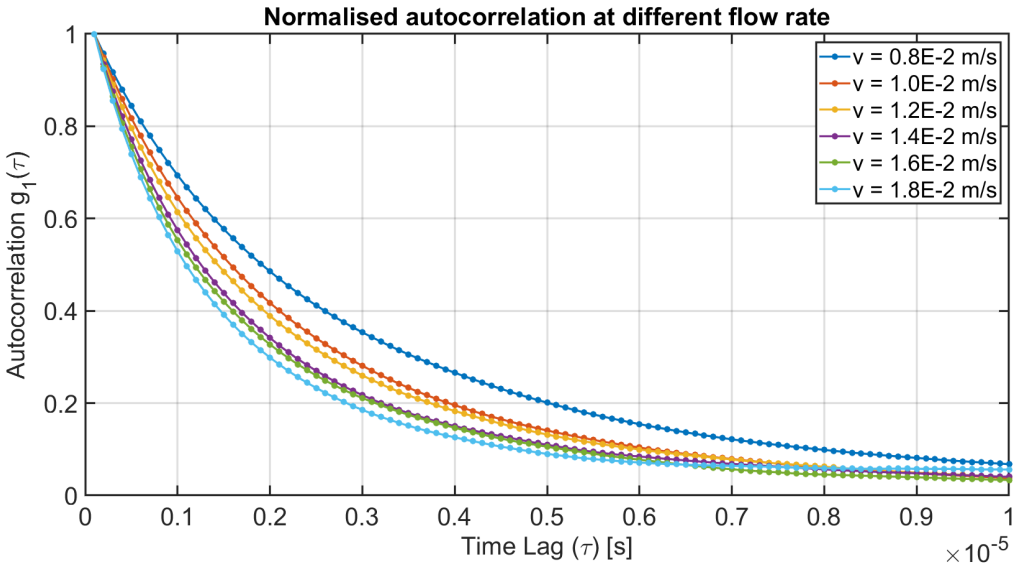

Figure 4.10: Normalized temporal autocorrelation functions using a 500-pixel weighted mean, shown for increasing flow speed

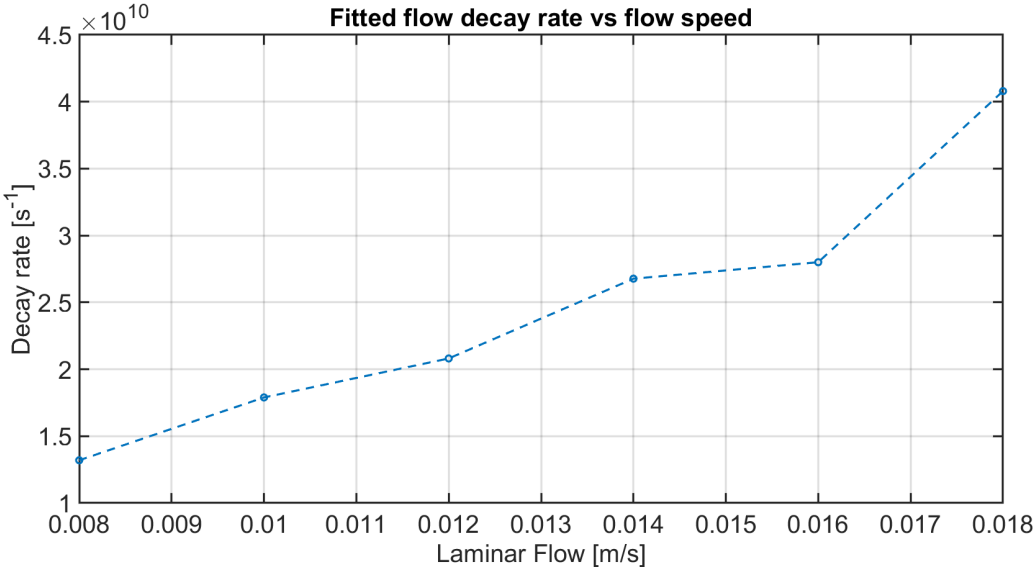

Figure 4.11: Extracted decay coefficients as a function of flow speed.

5

Discussion

The aim of this study was to evaluate whether a coupled COMSOL–Born series simulation framework can reproduce the expected behaviour of diffusing correlation spectroscopy (DCS). Specifically, the model was designed to capture particle dynamics through COMSOL-generated trajectories and to translate these into optical field fluctuations via the Born series solver, enabling direct calculation of autocorrelation functions. The central questions were whether the model reproduces Brownian motion scaling with temperature, whether flow contributions can be detected and separated, and whether exponential fitting allows quantitative extraction of diffusion and flow parameters.

The results showed that the model captures the essential qualitative features of DCS. In the Brownian motion tests, COMSOL simulations consistently produced linear mean squared displacement (MSD) curves, confirming that particle dynamics followed Brownian behaviour under appropriate viscosity and temperature conditions. The corresponding autocorrelation functions from WaveSim exhibited faster decay at higher temperatures, in agreement with the theoretical expectation that diffusion accelerates decorrelation. However, the scaling of the extracted decay rates with temperature was not strictly linear: in the main case (300–1800 K, low viscosity) the growth was non-linear, with diminishing increments at higher temperatures, while in other regimes noise or non-linear MSD behaviour limited agreement with theory.

For the laminar flow tests, adding a constant external force produced the expected quadratic growth in the MSD and resulted in accelerated autocorrelation decay. Exponential fits using an extended decay model confirmed sensitivity to both Brownian and flow terms. Nevertheless, the extracted flow coefficients did not scale proportionally with velocity, and overlapping decay curves at higher speeds limited quantitative precision.

Taken together, these findings indicate that the simulation framework qualitatively reproduces the mechanisms underlying DCS: Brownian motion leads to linear MSD growth and temperature-dependent decorrelation, while flow introduces an additional quadratic term that accelerates decay. At the same time, the deviations from linear scaling and the influence of noise demonstrate practical limitations in recovering parameters quantitatively.

The observed mismatch between COMSOL motion and WaveSim dynamics is explained by the discretisation of the dielectric map. Sub-pixel particle movements do not alter the refractive index distribution until a boundary is crossed, so most of the physical motion is invisible to the optical solver. The average per-frame particle displacement is 0.001–0.003 μm, which is far smaller than the WaveSim pixel size of 0.2 µm. As a consequence, only a small fraction of particles crossed the pixel boundaries between frames (e.g. 1/15,352 at 300 K, rising to ≈236/15,352 at 1800 K). This quantisation effect suppresses decorrelation and leads to overlapping decay curves at different temperatures. A simple resolution metric can be defined,

$$\chi = \frac{\sqrt{4D_B\Delta t}}{pixel\ size} \tag{5.1}$$

where D_B is the Brownian diffusion coefficient and Δt is the frame interval. In all cases, $\chi \ll 0.2$, explaining the underestimation of decorrelation and the irregular scaling of decay rate with temperature.

In addition to these round-off effects, the present work did not explicitly track the phase shifts driving speckle decorrelation. In DCS, the speckle pattern arises from interference of many scattered photon paths, and each particle displacement alters optical path lengths and thus the phase of the electric field. Even sub-nanometre displacements contribute to decorrelation, provided they induce a phase shift on the order of the optical wavelength. In this framework, however, most small displacements remain invisible, and the corresponding phase increments between frames are effectively quantised away. The result is an underestimation of decorrelation, since much of the phase dynamics is not represented in the current pixel-based rendering.

Another factor influencing the results is the number of independent observables (NIO) available in the simulation. In DCS, the statistical reliability of the autocorrelation function depends on averaging over many speckles and time samples. Here, the number of pixels (spatial samples) and frames (temporal samples) directly sets the effective NIO. A limited pixel count or too few frames increases the statistical variance of the autocorrelation function. This effect is especially visible in the high-temperature runs, where the expected monotonic scaling of decay rate becomes obscured by fluctuations. Xu et al. emphasise that insufficient NIO leads to “shape deviation” in the autocorrelation function, as random speckle fluctuations dominate the statistics [55]. Our results are consistent with this limitation.

A further complication arises from the use of COMSOL-generated particle trajectories as input to WaveSim. While COMSOL provides continuous trajectories with sub-nanometre resolution, WaveSim requires discretisation of these trajectories onto a pixel grid at each frame. This translation introduces an additional layer of complexity, because small but physically meaningful displacements in COMSOL may not map to a detectable pixel shift in WaveSim. As a result, the true underlying Brownian motion is partially obscured, and the apparent dynamics captured in the autocorrelation function differ from theoretical expectations. This highlights the challenge of coupling two models with very different resolution scales and numerical formalisms.

The aim of this work was to test whether the COMSOL–Born framework could reproduce DCS autocorrelation scaling and ultimately separate the Brownian and flow contributions. The present results indicate that, under the current rendering resolution, the framework cannot yet provide unbiased scaling of blood-flow-related motion. The sublinear temperature dependence and mismatched decorrelation rates highlight the need to address pixel-level quantisation before the framework can meet its intended goal. The expectation was that increasing temperature would linearly accelerate decorrelation, consistent with theory where $D_B \propto T$. This expectation was met in the MSD data from COMSOL, but not in the WaveSim autocorrelations, where the decay rate grew sublinearly and irregularly with temperature. Similarly, laminar flow simulations did show faster decay, but the results were distorted by the same quantisation limit. Thus, our theoretical expectations were only partially met.

The main implication is that parameter extraction from sub-pixel–limited simulations is biased. Fitting exponential models to autocorrelations in this regime underestimates diffusion and flow coefficients, and produces nonlinear trends even when the physics is linear. For practice, this means that simulation parameters (Δt, grid size, pixel resolution) must be chosen such that motion becomes visible above the quantisation threshold. More broadly, this work highlights a methodological limitation relevant to all pixel-based DCS models: accurate inference requires matching spatial resolution to the expected motion scale.

Methodologically, the study was constrained by the choice of a 0.2 µm pixel grid in WaveSim, which is too coarse relative to the COMSOL-predicted Brownian displacements. The approach also relied on binary occupancy maps, which do not capture subpixel shifts. These choices reduce the validity of the results and explain the mismatch between theory and simulation. Reliability is further affected by noise in the exponential fitting, which amplifies the quantisation artefact.

Another key limitation of this study is that only the autocorrelation decay rate could be extracted. We could not reliably determine absolute Brownian diffusion coefficients or flow velocities, because this would require solving the correlation diffusion equation (CDE) with knowledge of the medium’s optical properties specifically, the absorption coefficient (μ_a), scattering coefficient (μ_s), anisotropy factor (g), and transport mean free path (l^*). These parameters were not characterised in the present work,

meaning the link between decay rate and absolute motion parameters remains incomplete. As such, the current framework is best suited for studying relative scaling trends, such as temperature dependence of Brownian motion or the presence of flow contributions, rather than providing quantitative flow indices comparable to in vivo DCS measurements.

A further limitation of the present framework is computational cost. The WaveSim solver required approximately 300 seconds to compute each frame, and each autocorrelation function was generated from 1000 frames. This results in runtimes on the order of days for a single condition, which severely restricts the range of parameters that can be explored and precludes any real-time applications. The high computational demand limits scalability and makes systematic parameter sweeps impractical.

Looking ahead, several directions emerge for extending and strengthening this framework. The most immediate priority is to overcome the pixel-quantisation limit that currently suppresses decorrelation. This could be achieved either by adopting finer spatial grids in WaveSim or by matching the frame-by-frame motion within WaveSim with the time-step motion in COMSOL. In parallel, optimisation of temporal sampling is needed. Choosing frame intervals that balance physical motion with computational feasibility would ensure that displacements become visible at the lags used for autocorrelation fitting, while still keeping runtimes manageable.

Another important step is to incorporate phase dynamics explicitly. Since speckle decorrelation is driven by cumulative phase shifts from particle motion, methods that track or approximate phase evolution between frames could help capture motion on scales smaller than the current pixel size. This would bring the simulation closer to the physical basis of DCS and reduce the underestimation of decorrelation.

To enable quantitative parameter extraction, future work must also integrate a correlation diffusion equation (CDE) framework. This requires specifying the medium's optical absorption and scattering coefficients, the anisotropy factor, and the transport mean free path. With these parameters in place, the simulated decay rates could be mapped onto absolute diffusion coefficients and flow velocities, allowing direct comparison to experimental DCS data.

Finally, computational performance remains a major barrier. Each autocorrelation function currently requires several days of computation, which limits systematic testing across parameters. Future work should therefore explore domain decomposition, and Python-based solver implementations that can reduce runtimes substantially. Such improvements would open the door to larger parameter sweeps, sensitivity analyses, and eventually, real-time modelling capabilities.

6

Conclusion

The aim of this thesis was to evaluate whether a coupled COMSOL–Born series simulation framework can reproduce the expected behaviour of diffusing correlation spectroscopy (DCS). The framework was designed to link particle dynamics, modelled in COMSOL, with optical field fluctuations, calculated using a modified Born series solver, to generate autocorrelation functions characteristic of DCS. The central objective was to determine whether such a framework could qualitatively and quantitatively capture Brownian diffusion and laminar flow contributions to the autocorrelation decay.

The results demonstrated that the framework successfully reproduces the qualitative behaviour expected from DCS. In the Brownian motion simulations, COMSOL generated linear mean squared displacement (MSD) curves, and WaveSim translated these into autocorrelation functions with faster decay at higher temperatures. In the flow simulations, externally imposed velocities produced quadratic MSD growth and correspondingly faster decorrelation, consistent with theoretical expectations. These findings confirm that the framework captures the essential physical mechanisms underlying DCS.

However, the model did not reproduce the expected quantitative scaling of DCS parameters. The extracted decay rates grew sublinearly with temperature in the Brownian case and failed to scale proportionally with velocity in the flow case. Detailed analysis identified the source of this mismatch as a resolution gap between COMSOL and WaveSim: particle displacements per frame were typically two orders of magnitude smaller than the 0.2 µm pixel size, meaning most motion was invisible to the optical solver. This pixel quantisation effect suppressed decorrelation and distorted scaling, preventing reliable recovery of diffusion or flow parameters. Additional limitations include the absence of phase-resolved modelling, restricted number of independent observables due to finite pixels and frames, reliance on decay rates alone without a correlation diffusion equation (CDE) solution, and long computational runtimes.

The central conclusion of this work is that the COMSOL–Born series framework provides a promising basis for DCS simulation but, in its current form, cannot deliver unbiased parameter estimates. The qualitative agreement with DCS physics demonstrates proof of concept, but accurate inference of diffusion coefficients or flow velocities requires resolving subpixel motion, incorporating phase dynamics, and linking decay rates to optical transport parameters through the CDE.

Looking forward, methodological refinements offer a clear path toward quantitative capability. Subpixel rendering, adaptive temporal sampling, and improved statistical averaging would address the quantisation limit. Integration of optical properties and CDE solutions would enable mapping decay rates onto absolute motion parameters. Finally, advances in computational efficiency through domain decomposition could make large-scale parameter sweeps and real-time modelling feasible. With these improvements, the framework could become a powerful tool for understanding DCS signals and supporting the development of bedside, non-invasive blood flow monitoring technologies.

References

[1] Quan Wang et al. “A comprehensive overview of diffuse correlation spectroscopy: Theoretical framework, recent advances in hardware, analysis, and applications”. In: *NeuroImage* 298 (2024), p. 120793. ISSN: 1053-8119. DOI: https://doi.org/10.1016/j.neuroimage.2024.120793. URL: https://www.sciencedirect.com/science/article/pii/S1053811924002908.

[2] Turgut Durduran and Arjun G. Yodh. “Diffuse correlation spectroscopy for non-invasive, micro-vascular cerebral blood flow measurement”. In: *NeuroImage* 85 (2014). Celebrating 20 Years of Functional Near Infrared Spectroscopy (fNIRS), pp. 51–63. ISSN: 1053-8119. DOI: https://doi.org/10.1016/j.neuroimage.2013.06.017. URL: https://www.sciencedirect.com/science/article/pii/S105381191300654X.

[3] Sava Sakadžić, David A. Boas, and Stefan A. Carp. “Theoretical model of blood flow measurement by diffuse correlation spectroscopy”. In: *Journal of Biomedical Optics* 22.2 (2017), p. 027006. DOI: 10.1117/1.JBO.22.2.027006. URL: https://doi.org/10.1117/1.JBO.22.2.027006.

[4] David A Weitz and David J Pine. “Diffusing-wave spectroscopy”. In: *Dynamic Light Scattering* (Jan. 1993), pp. 652–720. DOI: 10.1093/oso/9780198539421.003.0016. URL: https://cir.nii.ac.jp/crid/1360298451714441344.

[5] Danil Tyulmankov, Danil Tyulmankov, and Danil Tyulmankov. “Time-domain diffuse correlation spectroscopy: instrument prototype, preliminary measurements, and theoretical modeling”. In: (2017). URL: http://hdl.handle.net/1721.1/113444.

[6] Yaoshen Yuan, Shijie Yan, and Qianqian Fang. “Light transport modeling in highly complex tissues using the implicit mesh-based Monte Carlo algorithm”. In: *Biomed. Opt. Express* 12.1 (Jan. 2021), pp. 147–161. DOI: 10.1364/BOE.411898. URL: https://opg.optica.org/boe/abstract.cfm?URI=boe-12-1-147.

[7] Shijie Yan and Qianqian Fang. “Hybrid mesh and voxel based Monte Carlo algorithm for accurate and efficient photon transport modeling in complex bio-tissues”. In: *Biomed. Opt. Express* 11.11 (Nov. 2020), pp. 6262–6270. DOI: 10.1364/BOE.409468. URL: https://opg.optica.org/boe/abstract.cfm?URI=boe-11-11-6262.

[8] Anh Phong Tran and Steven L. Jacques. “Modeling voxel-based Monte Carlo light transport with curved and oblique boundary surfaces”. In: *Journal of Biomedical Optics* 25.2 (2020), p. 025001. DOI: 10.1117/1.JBO.25.2.025001. URL: https://doi.org/10.1117/1.JBO.25.2.025001.

[9] Gerwin Osnabrugge, Saroch Leedumrongwatthanakun, and Ivo M. Vellekoop. “A convergent Born series for solving the inhomogeneous Helmholtz equation in arbitrarily large media”. In: *Journal of Computational Physics* 322 (2016), pp. 113–124. ISSN: 0021-9991. DOI: https://doi.org/10.1016/j.jcp.2016.06.034. URL: https://www.sciencedirect.com/science/article/pii/S0021999116302595.

[10] Felix Scholkmann and Martin Wolf. “General equation for the differential pathlength factor of the frontal human head depending on wavelength and age”. In: *Journal of Biomedical Optics* 18.10 (2013), p. 105004. DOI: 10.1117/1.JBO.18.10.105004. URL: https://doi.org/10.1117/1.JBO.18.10.105004.

[11] Cecil Cheung et al. “In vivo cerebrovascular measurement combining diffuse near-infrared absorption and correlation spectroscopies”. In: *Physics in Medicine Biology* 46.8 (July 2001), p. 2053. DOI: 10.1088/0031-9155/46/8/302. URL: https://dx.doi.org/10.1088/0031-9155/46/8/302.

[12] D. A. Boas, L. E. Campbell, and A. G. Yodh. “Scattering and Imaging with Diffusing Temporal Field Correlations”. In: *Phys. Rev. Lett.* 75 (9 Aug. 1995), pp. 1855–1858. DOI: 10.1103/PhysRevLett.75.1855. URL: https://link.aps.org/doi/10.1103/PhysRevLett.75.1855.

[13] David A. Boas et al. “Diffusion of temporal field correlation with selected applications”. In: *CIS Selected Papers: Coherence-Domain Methods in Biomedical Optics*. Ed. by Valery V. Tuchin. Vol. 2732. International Society for Optics and Photonics. SPIE, 1996, pp. 34–46. DOI: 10.1117/12.231685. URL: https://doi.org/10.1117/12.231685.

[14] David A. Boas et al. “Establishing the diffuse correlation spectroscopy signal relationship with blood flow.” In: *Neurophotonics* (2016). DOI: 10.1117/1.nph.3.3.031412.

[15] Guoqiang Yu et al. “Validation of diffuse correlation spectroscopy for muscle blood flow with concurrent arterial spin labeled perfusion MRI”. In: *Opt. Express* 15.3 (Feb. 2007), pp. 1064–1075. DOI: 10.1364/OE.15.001064. URL: https://opg.optica.org/oe/abstract.cfm?URI=oe-15-3-1064.

[16] Regine Choe et al. “Optically measured microvascular blood flow contrast of malignant breast tumors”. In: *PLOS ONE* 9 (2014), e99683. DOI: https://doi.org/10.1371/journal.pone.0099683. URL: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0099683.

[17] Turgut Durduran et al. “Transcranial Optical Monitoring of Cerebrovascular Hemodynamics in Acute Stroke Patients”. In: *Opt. Express* 17.5 (Mar. 2009), pp. 3884–3902. DOI: 10.1364/OE.17.003884. URL: https://opg.optica.org/oe/abstract.cfm?URI=oe-17-5-3884.

[18] John Sunwoo et al. “Diffuse correlation spectroscopy blood flow monitoring for intraventricular hemorrhage vulnerability in extremely low gestational age newborns”. In: *Scientific Reports* (2022). DOI: 10.1038/s41598-022-16499-3.

[19] Wenjun Zhou et al. “Functional interferometric diffusing wave spectroscopy of the human brain”. In: *Science Advances* 7.20 (2021), eabe0150. DOI: 10.1126/sciadv.abe0150. eprint: https://www.science.org/doi/pdf/10.1126/sciadv.abe0150. URL: https://www.science.org/doi/abs/10.1126/sciadv.abe0150.

[20] Francescopaolo Mattioli Della Rocca et al. “A 512×512 SPAD Laser Speckle Autocorrelation Imager in Stacked 65/40nm CMOS”. In: *2024 IEEE Symposium on VLSI Technology and Circuits (VLSI Technology and Circuits)*. 2024, pp. 1–2. DOI: 10.1109/VLSITechnologyandCir46783.2024.10631557.

[21] Turgut Durduran et al. “Diffuse optical measurement of blood flow in breast tumors”. In: *Opt. Lett.* 30.21 (Nov. 2005), pp. 2915–2917. DOI: 10.1364/OL.30.002915. URL: https://opg.optica.org/ol/abstract.cfm?URI=ol-30-21-2915.

[22] Hossein S. Yazdi et al. “Mapping breast cancer blood flow index, composition, and metabolism in a human subject using combined diffuse optical spectroscopic imaging and diffuse correlation spectroscopy”. In: *Journal of Biomedical Optics* 22.4 (2017), p. 045003. DOI: 10.1117/1.JBO.22.4.045003. URL: https://doi.org/10.1117/1.JBO.22.4.045003.

[23] Markus Belau et al. “Noninvasive observation of skeletal muscle contraction using near-infrared time-resolved reflectance and diffusing-wave spectroscopy”. In: *Journal of Biomedical Optics* 15.5 (2010), p. 057007. DOI: 10.1117/1.3503398. URL: https://doi.org/10.1117/1.3503398.

[24] Turgut Durduran et al. “Diffuse optical measurement of blood flow, blood oxygenation, and metabolism in a human brain during sensorimotor cortex activation”. In: *Opt. Lett.* 29.15 (Aug. 2004), pp. 1766–1768. DOI: 10.1364/OL.29.001766. URL: https://opg.optica.org/ol/abstract.cfm?URI=ol-29-15-1766.

[25] Jason Sutin et al. “Time-domain diffuse correlation spectroscopy”. In: *Optica* 3.9 (Sept. 2016), pp. 1006–1013. DOI: 10.1364/OPTICA.3.001006. URL: https://opg.optica.org/optica/abstract.cfm?URI=optica-3-9-1006.

[26] Sadhu Moka et al. “Frequency domain diffuse correlation spectroscopy: a new method for simultaneous estimation of static and dynamic tissue optical properties”. In: *Multiscale Imaging and Spectroscopy III*. Ed. by Paul J. Campagnola, Kristen C. Maitland, and Darren M. Roblyer. Vol. 11944. International Society for Optics and Photonics. SPIE, 2022, p. 1194409. DOI: 10.1117/12.2610115. URL: https://doi.org/10.1117/12.2610115.

[27] Yanwen Zhu et al. “Experimental Validation of Microvasculature Blood Flow Modeling by Diffuse Correlation Spectroscopy”. In: *IEEE Access* 8 (2020), pp. 15945–15951. DOI: 10.1109/ACCESS.2020.2966750.

[28] Vinh Nguyen Du Le et al. “Beyond diffuse correlations: Deciphering random flow in time-of-flight resolved light dynamics”. In: *Optics Express* (2020). DOI: 10.1364/ots.2020.sm3d.4.

[29] Daniel Irwin et al. “Influences of tissue absorption and scattering on diffuse correlation spectroscopy blood flow measurements”. In: *Biomed. Opt. Express* 2.7 (July 2011), pp. 1969–1985. DOI: 10.1364/BOE.2.001969. URL: https://opg.optica.org/boe/abstract.cfm?URI=boe-2-7-1969.

[30] Ashley Welch and Martin van Gemert. *Optical-Thermal Response of Laser-Irradiated Tissue*. Springer Science, Jan. 2011, p. 938. ISBN: 978-90-481-8830-7. DOI: 10.1007/978-90-481-8831-4.

[31] T Durduran et al. “Diffuse optics for tissue monitoring and tomography”. In: *Reports on Progress in Physics* 73.7 (June 2010), p. 076701. DOI: 10.1088/0034-4885/73/7/076701. URL: https://dx.doi.org/10.1088/0034-4885/73/7/076701.

[32] D. A. Boas et al. “Three dimensional Monte Carlo code for photon migration through complex heterogeneous media including the adult human head”. In: *Opt. Express* 10.3 (Feb. 2002), pp. 159–170. DOI: 10.1364/OE.10.000159. URL: https://opg.optica.org/oe/abstract.cfm?URI=oe-10-3-159.

[33] Qianqian Fang and David A. Boas. “Monte Carlo Simulation of Photon Migration in 3D Turbid Media Accelerated by Graphics Processing Units”. In: *Opt. Express* 17.22 (Oct. 2009), pp. 20178–20190. DOI: 10.1364/OE.17.020178. URL: https://opg.optica.org/oe/abstract.cfm?URI=oe-17-22-20178.

[34] Melissa M. Wu et al. “Improved accuracy of cerebral blood flow quantification in the presence of systemic physiology cross-talk using multi-layer Monte Carlo modeling”. In: *Neurophotonics* 8.1 (2021), p. 015001. DOI: 10.1117/1.NPh.8.1.015001. URL: https://doi.org/10.1117/1.NPh.8.1.015001.

[35] Vysakh Vasudevan and Sujatha N. “Immediate subsurface skin blood flow monitoring using diffuse correlation spectroscopy: finite element simulations”. In: *Tissue Optics and Photonics II*. Ed. by Valery V. Tuchin, Walter C. P. M. Blondel, and Zeev Zalevsky. Vol. 12147. International Society for Optics and Photonics. SPIE, 2022, p. 1214709. DOI: 10.1117/12.2621523. URL: https://doi.org/10.1117/12.2621523.

[36] Vysakh Vasudevan and Sujatha Narayanan Unni. “Short-range diffuse correlation spectroscopic system for peripheral skin tissue blood flow assessment: in-vitro studies”. In: *2023 International Electrical Engineering Congress (iEECON)*. 2023, pp. 434–437. DOI: 10.1109/iEECON56657.2023.10126753.

[37] Ao Teng et al. “Simulation study of interaction of pulse laser with brain using COMSOL”. In: *SPIE Future Sensing Technologies*. Ed. by Masafumi Kimata, Joseph A. Shaw, and Christopher R. Valenta. Vol. 11525. International Society for Optics and Photonics. SPIE, 2020, 115252J. DOI: 10.1117/12.2584025. URL: https://doi.org/10.1117/12.2584025.

[38] Shimaa Mahdy et al. “Numerical analysis of the optical fluence rate at the scalp for noninvasive brain tumor detection”. In: *J. Opt. Soc. Am. A* 39.4 (Apr. 2022), pp. 587–593. DOI: 10.1364/JOSAA.446677. URL: https://opg.optica.org/josaa/abstract.cfm?URI=josaa-39-4-587.

[39] Matthaios Doulgerakis-Kontoudis et al. “Toward real-time diffuse optical tomography: accelerating light propagation modeling employing parallel computing on GPU and CPU”. In: *Journal of Biomedical Optics* 22.12 (2017), p. 125001. DOI: 10.1117/1.JBO.22.12.125001. URL: https://doi.org/10.1117/1.JBO.22.12.125001.

[40] M. C. W. van Rossum and Th. M. Nieuwenhuizen. “Multiple scattering of classical waves: microscopy, mesoscopy, and diffusion”. In: *Rev. Mod. Phys.* 71 (1 Jan. 1999), pp. 313–371. DOI: 10.1103/RevModPhys.71.313. URL: https://link.aps.org/doi/10.1103/RevModPhys.71.313.

[41] Lonny L. Thompson. “A review of finite-element methods for time-harmonic acoustics”. In: *The Journal of the Acoustical Society of America* 119.3 (Mar. 2006), pp. 1315–1330. ISSN: 0001-4966. DOI: 10.1121/1.2164987. eprint: https://pubs.aip.org/asa/jasa/article-pdf/119/3/1315/14874278/1315_1_online.pdf. URL: https://doi.org/10.1121/1.2164987.

[42] Majid Nabavi, M.H. Kamran Siddiqui, and Javad Dargahi. “A new 9-point sixth-order accurate compact finite-difference method for the Helmholtz equation”. In: *Journal of Sound and Vibration* 307.3 (2007), pp. 972–982. ISSN: 0022-460X. DOI: https://doi.org/10.1016/j.jsv.2007.06.070. URL: https://www.sciencedirect.com/science/article/pii/S0022460X07004877.

[43] Y. A. Erlangga, C. W. Oosterlee, and C. Vuik. “A Novel Multigrid Based Preconditioner For Heterogeneous Helmholtz Problems”. In: *SIAM Journal on Scientific Computing* 27.4 (2006), pp. 1471–1492. DOI: 10.1137/040615195. eprint: https://doi.org/10.1137/040615195. URL: https://doi.org/10.1137/040615195.

[44] M. J. Gander, I. G. Graham, and E. A. Spence. “Applying GMRES to the Helmholtz equation with shifted Laplacian preconditioning: what is the largest shift for which wavenumber-independent convergence is guaranteed?” In: *Numerische Mathematik* 131.3 (Nov. 2015), pp. 567–614. ISSN: 0945-3245. DOI: 10.1007/s00211-015-0700-2. URL: https://doi.org/10.1007/s00211-015-0700-2.

[45] O. G. Ernst and M. J. Gander. “Why it is Difficult to Solve Helmholtz Problems with Classical Iterative Methods”. In: *Numerical Analysis of Multiscale Problems*. Ed. by Ivan G. Graham et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 325–363. ISBN: 978-3-642-22061-6. DOI: 10.1007/978-3-642-22061-6_10. URL: https://doi.org/10.1007/978-3-642-22061-6_10.

[46] Gregory Beylkin, Christopher Kurcz, and Lucas Monzón. “Fast convolution with the free space Helmholtz Green’s function”. In: *Journal of Computational Physics* 228.8 (2009), pp. 2770–2791. ISSN: 0021-9991. DOI: https://doi.org/10.1016/j.jcp.2008.12.027. URL: https://www.sciencedirect.com/science/article/pii/S0021999108006566.

[47] E. M. Sevick-Muraca et al. “Role of higher-order scattering in solutions to the forward and inverse optical-imaging problems in random media”. In: *Appl. Opt.* 36.34 (Dec. 1997), pp. 9058–9067. DOI: 10.1364/AO.36.009058. URL: https://opg.optica.org/ao/abstract.cfm?URI=ao-36-34-9058.

[48] İlker R. Çapoğlu et al. “Accuracy of the Born approximation in calculating the scattering coefficient of biological continuous random media”. In: *Opt. Lett.* 34.17 (Sept. 2009), pp. 2679–2681. DOI: 10.1364/OL.34.002679. URL: https://opg.optica.org/ol/abstract.cfm?URI=ol-34-17-2679.

[49] Joonoh Lim et al. “Born approximation model for light scattering by red blood cells”. In: *Biomed. Opt. Express* 2.10 (Oct. 2011), pp. 2784–2791. DOI: 10.1364/BOE.2.002784. URL: https://opg.optica.org/boe/abstract.cfm?URI=boe-2-10-2784.

[50] Jeremy D. Rogers et al. “Modeling Light Scattering in Tissue as Continuous Random Media Using a Versatile Refractive Index Correlation Function”. In: *IEEE Journal of Selected Topics in Quantum Electronics* 20.2 (2014), pp. 173–186. DOI: 10.1109/JSTQE.2013.2280999.

[51] Benjamin Krüger, Thomas Brenner, and Alwin Kienle. “Solution of the inhomogeneous Maxwell’s equations using a Born series”. In: *Opt. Express* 25.21 (Oct. 2017), pp. 25165–25182. DOI: 10.1364/OE.25.025165. URL: https://opg.optica.org/oe/abstract.cfm?URI=oe-25-21-25165.

[52] T. Vettenburg, S. A. R. Horsley, and J. Bertolotti. “Calculating coherent light-wave propagation in large heterogeneous media”. In: *Opt. Express* 27.9 (Apr. 2019), pp. 11946–11967. DOI: 10.1364/OE.27.011946. URL: https://opg.optica.org/oe/abstract.cfm?URI=oe-27-9-11946.

[53] Gerwin Osnabrugge, Maaike Benedictus, and Ivo M. Vellekoop. “Ultra-thin boundary layer for high-accuracy simulations of light propagation”. In: *Opt. Express* 29.2 (Jan. 2021), pp. 1649–1658. DOI: 10.1364/OE.412833. URL: https://opg.optica.org/oe/abstract.cfm?URI=oe-29-2-1649.

[54] Steven L. “Corrigendum: Optical properties of biological tissues: a review”. In: *Physics in Medicine Biology* 58.14 (June 2013), p. 5007. DOI: 10.1088/0031-9155/58/14/5007. URL: https://dx.doi.org/10.1088/0031-9155/58/14/5007.

[55] Jian Xu, Ali K Jahromi, and Changhuei Yang. “Diffusing wave spectroscopy: A unified treatment on temporal sampling and speckle ensemble methods”. In: *APL Photonics* 6.1 (Jan. 2021), p. 016105. ISSN: 2378-0967. DOI: 10.1063/5.0034576. URL: http://dx.doi.org/10.1063/5.0034576.

[56] Alwin Kienle and Michael S Patterson. “Determination of the optical properties of semi-infinite turbid media from frequency-domain reflectance close to the source”. In: *Physics in Medicine Biology* 42.9 (Sept. 1997), p. 1801. DOI: 10.1088/0031-9155/42/9/011. URL: https://doi.org/10.1088/0031-9155/42/9/011.

A

Source Code

A.1. COMSOL Code for Brownian motion and Laminar flow

```
%% %%%%%%%%%%%%%%%%%%%%%%%%%% COMOSL Initialization %%%%%%%%%%%%%%%%%%%%
import com.comsol.model.*
import com.comsol.model.util.*
outcham_l = 6000;
outcham_w = 1000;

model = ModelUtil.create('Model');
model.label('brownian_motion.mph');
model.param.set('rp', '1.5E-7[m]', 'Particle radius');
model.param.set('T', '2400[K]', 'Temperature'); %%%%%% T
model.param.set('eta', '1E-5[Pa*s]', 'Fluid viscosity');
model.param.set('D', '(k_B_const*T)/(6*pi*eta*rp)', 'Diffusivity');
model.param.set('ds', '3112', 'Input to random number generator');
%model.param.set('Uin', '0.6E-2[m/s]', 'Inlet velocity'); %%% flow speed
model.modelNode.create('comp1', true);

model.geom.create('geom1', 2);
model.geom('geom1').model('comp1');

model.mesh.create('mesh1', 'geom1');

model.geom('geom1').lengthUnit([native2unicode(hex2dec({'00' 'b5'}), 'unicode') 'm']); %%% um
model.geom('geom1').create('r1', 'Rectangle');
model.geom('geom1').feature('r1').set('base', 'center');
model.geom('geom1').feature('r1').set('size', [outcham_l outcham_w]);
model.geom('geom1').run;

model.variable.create('var1');
model.variable('var1').model('comp1');
model.variable('var1').set('smooth', '2E-7', 'Smoothing distance');
model.variable('var1').set('xd', 'x[1/m]', 'x coordinate');
model.variable('var1').set('yd', 'y[1/m]', 'y coordinate');
model.variable('var1').set('c0', '1', 'Peak initial concentration');
model.variable('var1').set('c_init', '2*c0*(1-flc2hs(xd^2+yd^2-smooth^2,5e-11))', 'Initial
    concentration');

model.material.create('mat1', 'Common', 'comp1');

model.physics.create('fpt', 'FluidParticleTracing', 'geom1');
model.physics('fpt').model('comp1');
model.physics('fpt').create('df1', 'DragForce', 2);
model.physics('fpt').feature('df1').selection.all;
%model.physics('fpt').feature('force1').label('Force 1');
%model.physics('fpt').feature('force1').selection.all;
%model.physics('fpt').feature('force1').set('F', { '6*pi*eta*rp*Uin'; '0' ; '0'});  %%% Flow
    force
```

```
model.physics('fpt').create('bf1', 'BrownianForce', 2);
model.physics('fpt').feature('bf1').selection.all;
model.physics('fpt').create('relg1', 'ReleaseGrid', -1);
model.physics('fpt').create('pcnt1', 'ParticleCounter', 2);
model.physics('fpt').feature('pcnt1').selection.set([1]);

model.mesh('mesh1').create('sca1', 'Scale');
model.mesh('mesh1').create('ftri1', 'FreeTri');
model.mesh('mesh1').feature('sca1').selection.geom('geom1', 0);

model.view('view1').axis.set('xmin', -158.11143493652344);
model.view('view1').axis.set('xmax', 158.11143493652344);
model.view('view1').axis.set('ymin', -137.5);
model.view('view1').axis.set('ymax', 137.5);

model.material('mat1').propertyGroup('def').set('density', '1000');

model.physics('fpt').prop('WallAccuracyOrder').set('WallAccuracyOrder', 1);
model.physics('fpt').prop('RandomNumberArgs').set('RandomNumberArgs', 'UserDefined');
model.physics('fpt').feature('wall1').set('WallCondition', 'DiffuseScattering');
model.physics('fpt').feature('wall1').set('i', 1000);
model.physics('fpt').feature('pp1').set('dp', '2*rp');
model.physics('fpt').feature('pp1').set('rhop_mat', 'userdef');
model.physics('fpt').feature('pp1').set('rhop', '1090[kg/m^3]');
model.physics('fpt').feature('df1').set('mu_mat', 'userdef');
model.physics('fpt').feature('df1').set('mu', 'eta');
model.physics('fpt').feature('bf1').set('mu_mat', 'userdef');
model.physics('fpt').feature('bf1').set('mu', 'eta');
model.physics('fpt').feature('bf1').set('i', 'ds');
model.physics('fpt').feature('bf1').set('minput_temperature', 'T');
model.physics('fpt').feature('relg1').set('x0', {'range(-3000,20,3000)'; 'range(-500,20,500)'
    }); %%%%%%release grid of particles
model.physics('fpt').feature('relg1').set('Us', 0);
model.physics('fpt').feature('relg1').set('Nvel', 5000);
model.physics('fpt').feature('pcnt1').set('ReleaseFeature', 'relg1');

model.mesh('mesh1').feature('size').set('hauto', 2);
model.mesh('mesh1').feature('sca1').set('scale', 0.05);
model.mesh('mesh1').run;
model.mesh('mesh1').feature('size').set('table', 'cfd');
model.mesh('mesh1').feature('size').set('hauto', 4);
model.mesh('mesh1').run;

model.study.create('std1');
model.study('std1').create('time', 'Transient');

model.sol.create('sol1');
model.sol('sol1').study('std1');
model.sol('sol1').attach('std1');
model.sol('sol1').create('st1', 'StudyStep');
model.sol('sol1').create('v1', 'Variables');
model.sol('sol1').create('t1', 'Time');
model.sol('sol1').feature('t1').create('fc1', 'FullyCoupled');
model.sol('sol1').feature('t1').create('i1', 'Iterative');
model.sol('sol1').feature('t1').feature('i1').create('ja1', 'Jacobi');
model.sol('sol1').feature('t1').feature.remove('fcDef');

model.result.dataset.create('part1', 'Particle');
model.result.create('pg1', 'PlotGroup2D');
model.result('pg1').set('data', 'part1');
model.result('pg1').create('traj1', 'ParticleTrajectories');
model.result('pg1').feature('traj1').create('col1', 'Color');
model.result('pg1').feature('traj1').feature('col1').set('expr', 'fpt.V');
model.result.export.create('anim1', 'Animation');

model.study('std1').feature('time').set('tunit', [native2unicode(hex2dec({'00' 'b5'}), '
    unicode') 's']); %%%% us

model.sol('sol1').attach('std1');
model.sol('sol1').feature('st1').label('Compile␣Equations:␣Time␣Dependent');
model.sol('sol1').feature('v1').label('Dependent␣Variables␣1.1');
```

```
model.sol('sol1').feature('v1').set('clist', {'range(0,0.1,60)' ['1.0000000000000002['
    native2unicode(hex2dec({'00' 'b5'}), 'unicode') 's']});
model.sol('sol1').feature('t1').label('Time-Dependent␣Solver␣1.1');
model.sol('sol1').feature('t1').set('tunit', [native2unicode(hex2dec({'00' 'b5'}), 'unicode')
     's']);
model.sol('sol1').feature('t1').set('tlist', ...
  'range(0,5,5000)␣range(5000,0.1,5200)');%%%%%%%%%%%%%%%%%%% simulation time

model.sol('sol1').feature('t1').set('rtol', 1.0E-5);
model.sol('sol1').feature('t1').set('ewtrescale', false);
model.sol('sol1').feature('t1').set('timemethod', 'genalpha');
model.sol('sol1').feature('t1').set('estrat', 'exclude');
model.sol('sol1').feature('t1').set('tstepsgenalpha', 'strict');
model.sol('sol1').feature('t1').feature('dDef').label('Direct␣1');
model.sol('sol1').feature('t1').feature('aDef').label('Advanced␣1');
model.sol('sol1').feature('t1').feature('fc1').label('Fully␣Coupled␣1.1');
model.sol('sol1').feature('t1').feature('fc1').set('ntolfact', 0.1);
model.sol('sol1').feature('t1').feature('i1').label('Iterative␣1.1');
model.sol('sol1').feature('t1').feature('i1').feature('ilDef').label('Incomplete␣LU␣1');
model.sol('sol1').feature('t1').feature('i1').feature('ja1').label('Jacobi␣1.1');
model.sol('sol1').runAll;

model.result('pg1').label('Particle␣Trajectories␣(fpt)');
model.result('pg1').feature('traj1').set('linetype', 'none');
model.result('pg1').feature('traj1').set('pointtype', 'point');
model.result('pg1').feature('traj1').set('sphereradiusscale', 349.99999999999994);
model.result('pg1').feature('traj1').set('sphereradiusscaleactive', false);
model.result.numerical.create('par1', 'Particle');
model.result.numerical('par1').set('expr', 'qx');
model.result.numerical.create('par2', 'Particle');
model.result.numerical('par2').set('expr', 'qy');
model.result.table.create('tbl1', 'Table');
model.result.table('tbl1').comments('Particle␣Evaluation␣1');
model.result.numerical('par1').set('table', 'tbl1');
model.result.numerical('par1').setResult;
model.result.table('tbl1').save('posx.csv');
model.result.table.create('tbl2', 'Table');
model.result.table('tbl2').comments('Particle␣Evaluation␣2');
model.result.numerical('par2').set('table', 'tbl2');
model.result.numerical('par2').setResult;
model.result.table('tbl2').save('posy.csv');

model.label('brownian_motion.mph');
```

A.2. MATLAB Code for WaveSim

```
clear all; close all;
addpath('..')

%% Simulation Options
PPW = 4;                            % Points per wavelength
opt.lambda = 0.8;                     % Wavelength in vacuum (um)
opt.energy_threshold = 1E-10;       % Energy threshold for convergence
opt.pixel_size = opt.lambda / PPW;  % Pixel size (um)
opt.boundary_widths = [2, 2];       % Periodic boundaries
opt.usemex = false;
if(opt.usemex)
    addpath('../MexBin');
end

%% Grid and Refractive Index Setup
S = 10;
s2p = 1/0.2;
p2s = 0.2;
N = [5000, 30000];
DS = (30000/2) - 2500;
DE = (30000/2) + 2500;
% Time steps and storage for scattering field
time_steps = 1001:1200;

```

```
%% Tissue Parameters
n0 = 1.00;
a0 = 0.001;
u_sl = 1;

%% Red Blood Cells Parameters
num_RBCs = 15352;           % Number of RBCs
rbc_radius = (8/2)*s2p;           % Radius of each RBC in pixels
n0r = 1.5;
a0r = 0.0001;

rbc_index = n0r + 1.0i*(a0r);        % Refractive index of RBCs


%% Simulation loop
n_sample = 1.0*(n0) + 1.0i*(a0)*ones(N);
air_layer = ones(1, N(2));
bottom_layer = ones(1, N(2));
n_sample_org = n_sample;

%Extracting RBC motion data
posx = readmatrix('posx.csv');
posy = readmatrix('posy.csv');

%%
% Precompute the grid and source
[x, y] = ndgrid(1:N(1), 1:N(2));  % Precompute grid coordinates
%% define pencil beam source (5 pixels wide) placed in the middle of the x-axis using
    gausswin
beam_width = 51;                            % number of pixels for the beam width
center_idx = round(N(2)/2);                % center position along the x-axis
beam_profile = gausswin(beam_width, 3);    % create a Gaussian window of length 5
source_amplitude = zeros(1, N(2));        % initialize the source array for the full x-axis
source_amplitude(center_idx-25:center_idx+25) = beam_profile;  % insert the Gaussian pencil
    beam at the center
source = Source(source_amplitude, [N(1)+2,1,1,1]); % create a normally incident source (no
    angle)


for t=time_steps
    % Update refractive index map with new RBC positions
    n_sample = n_sample_org;

    % Update only a local region around each RBC
   for i = 2:num_RBCs
    % Get the current RBC position in grid coordinates
    rbc_x = round(posy(t, i) * s2p + N(1) / 2);
    rbc_y = round(posx(t, i) * s2p + N(2) / 2);

    % Ensure the RBC positions are within the grid bounds
    rbc_x = max(min(rbc_x, N(1)), 1);
    rbc_y = max(min(rbc_y, N(2)), 1);

    % Define the bounding box around the RBC
    x_min = max(rbc_x - rbc_radius, 1);
    x_max = min(rbc_x + rbc_radius, N(1));
    y_min = max(rbc_y - rbc_radius, 1);
    y_max = min(rbc_y + rbc_radius, N(2));

    % Create a mask for the local region (bounding box) around the RBC
    x_local = x(x_min:x_max, y_min:y_max);  % Extract local x grid coordinates
    y_local = y(x_min:x_max, y_min:y_max);  % Extract local y grid coordinates
    rbc_mask = sqrt((x_local - rbc_x).^2 + (y_local - rbc_y).^2) <= rbc_radius;

    % Update the refractive index map within the bounding box
    n_sample(x_min:x_max, y_min:y_max) = rbc_index * rbc_mask + ...
        n_sample(x_min:x_max, y_min:y_max) .* ~rbc_mask;
   end


    disp('updated␣grid␣for␣new␣location')
```

```
93
94      % Run the wave simulation for the current refractive index
95      n_sample = [air_layer; n_sample; bottom_layer];
96      sim = WaveSim(n_sample, opt);
97      [E, ~] = exec(sim, source);            % Run simulation
98      disp('Wavesim␣done')
99
100     % Store the scattered field for the current time step
101     scattered_fields(1,:, t) = E(1,DS:DE);
102
103 end
104
105 % Save the first row of scattered fields
106 first_row_fields = scattered_fields;
107
108 %% Save the entire workspace
109 save('scatteredfields.mat', 'first_row_fields', '-v7.3');
```

B

Additional figures

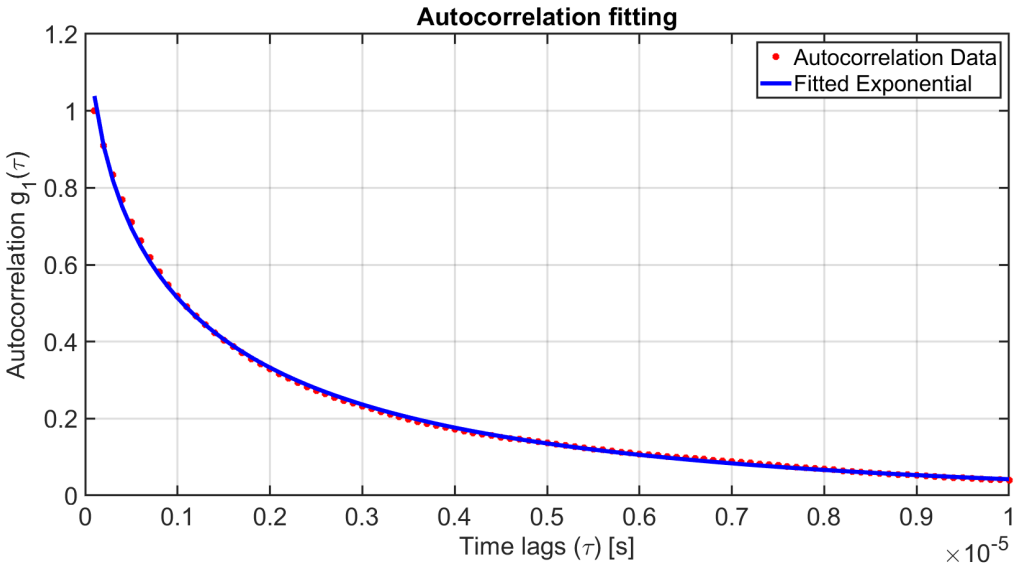

Figure B.1: Normalized temporal autocorrelation function at T = 300 K. The red dots shows the result obtained using a 500-pixel weighted mean, while the blue curve represents the fitted exponential model ($A = 1.0088 \times 10^6$, $B = 1.4523$ and $C = -0.0184$).

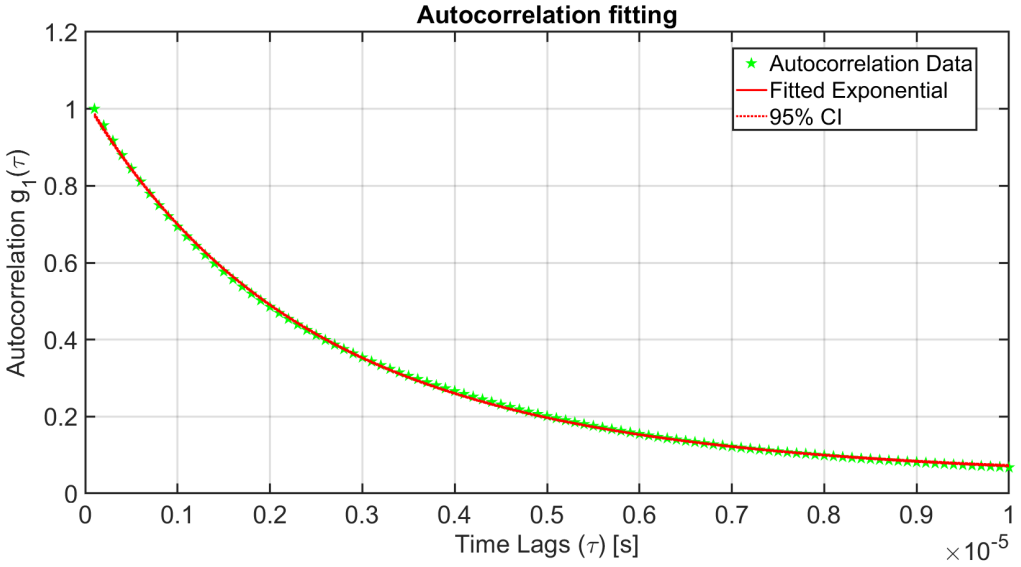

Figure B.3: Normalized temporal autocorrelation function at T=25 K with a flow speed of v=8.0×10^{-3} m/s. The green stars represent the simulated autocorrelation data, the red line shows the fitted exponential model (A = 3.49×10^{5}, B = 1.044, C = −0.0034, D = 1.32×10^{10}), and the striped red line indicates the 95% confidence interval.

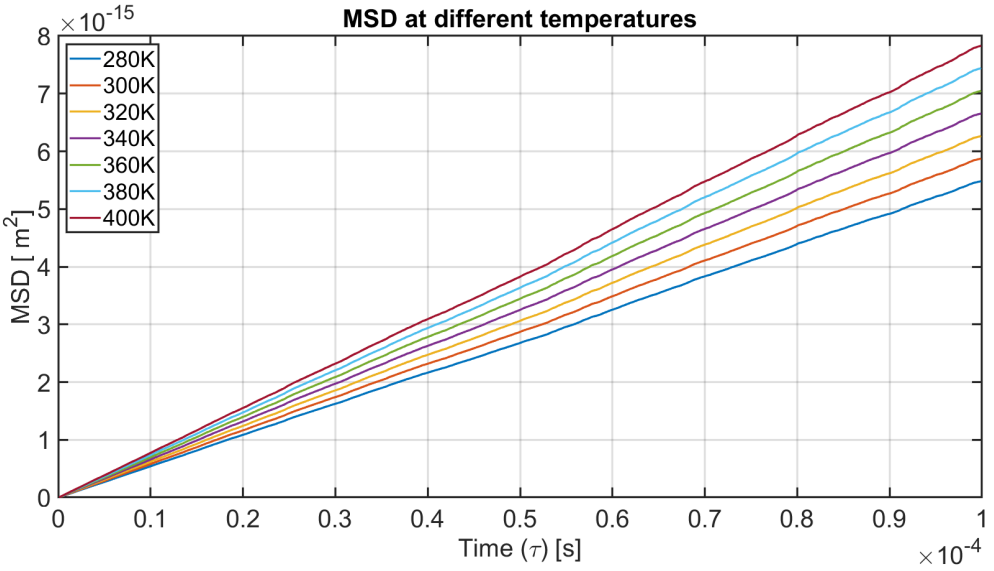

Figure B.2: Mean squared displacement for $T = 280$–400 K at viscosity 1.0×10^{-4}Pa·s. The MSD shows a linear relation with time across all cases.