Characterization of Diffuse Optical Tomography Scans using NIRFAST

Sean C. Youn

Advisors: Professor W. Cooke and Professor D. Manos

Senior Research Coordinator: Professor G. Hoatson

May 11, 2015

Abstract

Diffuse optical tomography (DOT) and other bio-optical imaging methods have recently emerged as comparable and viable alternatives to more well-established medical imaging technologies. DOT utilizes near-infrared light to create functional images based on differences in the relative scattering and absorbing properties of cellular components and molecules within the tissue in question. This study focused on the forward model portion of the DOT process, which calculates predicted values for the re-emission of near-infrared light across different discretized points within the tissue being imaged. NIRFAST, a software package developed specifically for the purpose of generating DOT scans based on experimental data, was used to simulate and conduct the forward model calculations. Initially, large variations were observed in the calculated values for the intensity and phase shift of the near-infrared light as measured by detectors in the NIRFAST forward model. The primary motivation of this study was to determine the underlying cause between the observed, unrealistic variations in intensity and phase shifts and to determine how the variations could be minimized through changes in the different input parameters involved in the forward model calculation.

1 Motivation

Over the past century, medical imaging technologies have developed into integral components of both diagnostics and research, with different imaging methods varying widely in their sets of advantages, disadvantages, and functionality. While X-rays, magnetic resonance imaging, and positron emission tomography rank among the most widely-used methods, optical imaging has recently gained traction in both clinical and experimental environments due to its versatility and the the low-risk associated with its use (no ionizing radiation or magnetic fields). Diffuse optical tomography (DOT) is a subsection of medical optical imaging that possesses the additional advantages of relatively low costs of production and use associated with its apparatus, high data-acquisition speeds, and potential portability [1,3]. Furthermore, the DOT set-up runs almost silently and allows patients the comfort of being either seated or in bed throughout the duration of the test, making it ideal for use in clinical and research settings alike [3,4]. All of these factors represent significant conveniences for both professionals and patients, and they collectively make DOT an attractive option in the treatment of patients, including pregnant women, infants, and patients with metallic or electronic implants, who would otherwise be precluded from other

imaging techniques.

The long-term goal of this project is to construct a frequency-domain DOT apparatus to scan and image zebra finch brains. Zebra finches are frequently-used model organisms in neurobiological research due to the fact that the mechanisms through which they learn to vocalize, communicate, and develop their auditory system are remarkably similar to the development of the same processes in humans [6]. Once an apparatus is constructed that successfully images finch brains, the same ideas and technology can be reproduced on a larger scale to create functional images and scans of human brains. Because DOT excels in its speed of image acquisition and apparatus portability, this technology could be used in situations, such as combat zones and emergency response settings, when medical professionals are required to quickly diagnose and treat traumatic head and brain injuries.

2 Introduction

Like other tomographic methods, DOT utilizes penetrating radiation to generate numerous projections from different angles of a particular tissue or object in order to generate a fully-reconstructed image. DOT relies on near-infrared (NIR) radiation, with wavelengths in the range of 700 nm to 900 nm, to characterize and discern the properties of tissues based on how photons propagate through the biological media in question. Within diffuse media such as medical organs and tissues, NIR photons exhibit a much stronger propensity towards scattering (rather than absorption). The probabilities with which photons are absorbed or scattered through a given length of media then become important quantities in describing the path taken by NIR light; these quantities are known as the absorption and reduced scattering coefficients respectively and possess units of mm^{-1}. The absorption coefficient, denoted (μ_a), typically varies in the range of 0.1 to 10 mm^{-1} for biological tissues [4]. Absorption at the cellular level predominantly originates from molecular chromophores including hemoglobin, melanin, water, and various lipids and is given by the Beer-Lambert relation:

$$\mu_a = [\varepsilon]C$$

where ε is the molar absorption spectra and C represents the concentration of the chromophore respectively [1,4]. Of primary interest in DOT are the relative levels of oxygenated and deoxygenated hemoglobin (HbO_2 and Hb respectively). Due to the different μ_a signatures of $Hb0_2$ and Hb, it is possible to use DOT to perceive relative changes in hemodynamics within underlying tissues and to elucidate the chemical composition and relative concentrations of both hemoglobin varieties within the area being imaged [4].

The reduced scattering coefficient is represented by μ'_s and has typical values between 50 to 200 mm^{-1} for biological tissues [4]. These values are larger than those of the absorption coefficient, which is consistent with diffuse media such as various organs and tissues in the body and allows for the path of scattered photons to be modeled as a random walk. Due to the strong tendency towards scattering, NIR photons can penetrate only up to ten centimeters into living biological tissues, with sub-cellular components (including the cell membrane, mitochondria, and nuclei) making up the chief scattering agents [1,4]. Mie theory can be used to approximate the reduced scattering coefficient by the following relation:

$$\mu'_s = a\lambda^{-b}$$

where a is the scatter amplitude, b designates the scatter power, and λ represents the wavelength [4]. Objects with a higher scattering coefficient typically have lower values for a and b, while the opposite

is true for objects with lower scattering coefficients [4]. μ'_s is also dependent on a linear scaling factor g, which is a parameter that describes the anisotropic scattering properties of different media (ranging from 0.75 to 0.99 for biological media). The differing indices of refraction and densities of components and fluids within the tissues and cells give rise to unique reduced scattering coefficients for particular tissues and biological structures [4]. Therefore, the scattering of NIR light can elucidate the nature of the physical components associated with the structure of interest, such as swelling and the presence of any abnormal growths.

Due to its usefulness in determining relative Hb and HbO_2 concentrations and physical characteristics of tissues, DOT has emerged as a useful tool in the diagnosis of cancers as angiogenesis (the production of new blood vessels from existing ones) and hypermetabolism are two hallmarks of the disease that are associated with markedly distinct hemodynamics compared to healthy tissue. Thus far, the most promising applications of DOT have been in brain imaging including detection of subcranial hematomas, functional imaging, and stroke prediction and monitoring [4]. Applications in imaging breast tissue and joints to assess the progression of rheumatoid arthritis have also been explored [4]. While these applications highlight the versatile nature of DOT, there is an inherent trade-off associated with its use as a diagnostic tool. Despite providing functional imaging capabilities at high speeds and low costs, DOT has a significantly lower spatial resolution than other imaging technologies (on the order of 20% of the imaging depth) due to the relatively sparse nature of its optical sources and detectors [1,3]. Thus, one of the main challenges in DOT implementation is in ensuring that the parameters and set-up of sources and detectors allows for sufficient differentiation of objects within close proximity in tissues.

The three principle varieties of DOT instrumentation are time-domain systems, frequency-domain systems, and steady-state domain systems. Of particular interest in this study are the frequency-domain systems due to its relatively high data-acquisition speeds, low costs, and superior signal-to-noise ratio compared to the other two methods [1]. A typical frequency-domain DOT apparatus consists of four main components: modulated light sources, a method of delivering the light to the tissue of interest, an array of detectors, and a set-up that uses a cross-correlation technique to measure observed phase shifts of light after having traveled through the medium [4]. Optical fibers are used to guide light to a array of sources and detectors arranged in various configurations around the area to be imaged. Light-emitting diodes (LED's) with a power of approximately 100 mW and sinusoidally modulated at a rate of around 100 to 1000 MHz are typically sufficient to probe biological tissues [4]. Photodiodes or photomultipler tubes (PMT's) are typically used as detectors of the intensities of outgoing light through the tissue in the DOT system [4]. The cross-correlation technique transforms the frequencies of the modulated light into the 1 kHz range while still preserving the original phase shift information. The lower frequency allows the phase shift to be more easily measured using various digital techniques [4].

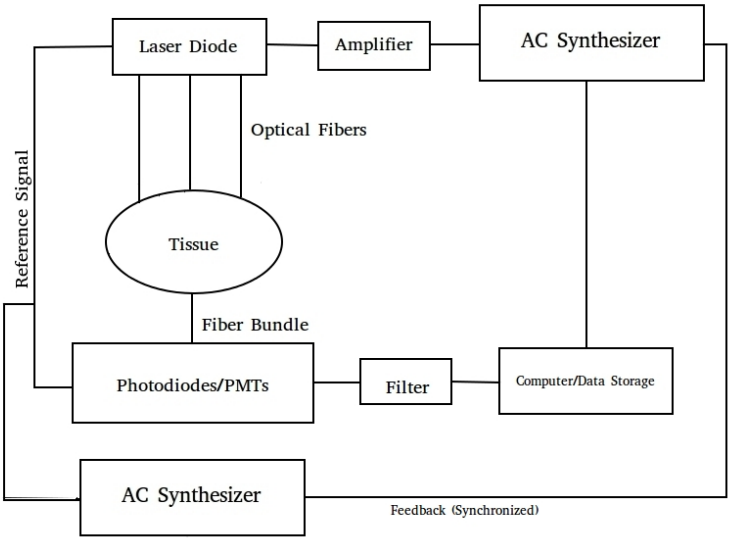

Figure 1: A typical schematic for a frequency-domain DOT apparatus

Once the outgoing photon intensities and phase shifts due to modulation have been measured, it is possible to generate cross-sectional images or 3-D volumetric scans of distributions of the optical properties of the scanned biological tissue. This process of image reconstruction is divided into two components: the forward model and reverse model. Put concisely, the forward model uses a diffusion approximation to predict the distribution of re-emitted light on the basis of presumed parameters of both the light sources and the object being imaged while the inverse model uses the results derived from the forward model to reconstruct the distributions of the optical properties of the object from a measured data set [1].

Because scattering of NIR photons in biological tissues (diffuse media) can be assumed to be isotropic in intensity and almost constant over time, the diffusion equation approximation is used to estimate the expected optical output at the detectors based on the source parameters and the properties of the tissue in question [1,2]. The diffusion approximation is given as the following:

$$q_0(r,w) = -\nabla \cdot \kappa(r)\nabla\Phi(r,w) + (\mu_a(r) + \frac{i\omega}{c_m(r)})\Phi(r,w)$$

where $q_0(r,w)$ is an isotropic source term, Φ(r,w) is the photon fluence rate at position r, ω is the modulation frequency, $\kappa = (1/3)(\mu_a + \mu'_s)$ and is the diffusion coefficient, and c_m(r) is the speed of light in a vacuum divided by the index of refraction at a given point [2]. Light propagation through biological tissue is simulated by the finite element method (FEM), a numerical technique by which a continuous medium is divided into smaller, discrete subdomains [2]. This enables the diffusion approximation to be used at specific, discrete points rather than a continuous volume. As such, solving for the photon fluence (Φ) for a given point in the diffusion equation can be reduced to a system of linear equations [2]:

$$q_0 = \Phi(K(\kappa) + C(\mu_a + \frac{i\omega}{c_m}) + F(\frac{1}{2A})$$

where K(κ), C($\mu_a + \frac{i\omega}{c_m}$), and F are matrices whose elements are given by the following:

$$K_{\mathrm{ij}} = \int_{\Omega} \kappa(r) \nabla u_{\mathrm{i}}(r) \cdot \nabla u_{\mathrm{i}}(r) d^{\mathrm{n}} r$$

$$C_{\mathrm{ij}} = \int_{\Omega} (\mu_{\mathrm{a}}(r) + \frac{i\omega}{c_{\mathrm{m}}(r)}) u_{\mathrm{i}}(r) u_{\mathrm{j}}(r) d^{\mathrm{n}} r$$

$$F_{\mathrm{ij}} = \oint_{\partial\Omega} u_{\mathrm{i}}(r) u_{\mathrm{j}}(r) d^{\mathrm{n\text{-}1}} r$$

and vector q_0 is defined as

$$q_{0\mathrm{i}} = \int_{\Omega} u_{\mathrm{i}}(r) q_0(r) d^{\mathrm{n}} r$$

While the purpose of the forward model is to solve for the photon fluence at different positions, the objective of the inverse model is to recover the optical properties (μ_{a} and μ'_{s}) at each FEM node using the measurements from the DOT scan [2,5]. This is done by minimizing a χ^2 statistic defined by the Tikhonov function [2]:

$$\chi^2 = min(\sum_{i=1}^{NM} (\Phi^{\mathrm{M}}{}_{\mathrm{i}} - \Phi^{\mathrm{C}}{}_{\mathrm{i}})^2 + \lambda \sum_{j=1}^{NN} (\mu_{\mathrm{a,k}} - \mu_0)^2)$$

where the photon fluence measured by the DOT apparatus is given by Φ^{M}, photon fluence calculated by the forward solver is Φ^{C}, NM represents the total number of measurements obtained (the number of sources times the number of detectors), NN is the number of FEM nodes, and λ is defined as the ratio of the variances of Φ^{M} and μ. μ_0 is an initial estimate of the optical properties. The minimization of the Tikhonov function is found by taking the partial derivative of the χ^2 statistic with respect to μ and setting the result to 0 (and ignoring higher order terms). The result of this first-order approximation is a Jacobian matrix which, when solved iteratively, yields the relationship between changes in the calculated photon fluence Φ^{C} resulting from small changes in the optical properties at each FEM node. The Jacobian matrix is then normalized and the variances ratio λ is systematically reduced. This process continues until the experimental data matches the calculated data to a certain tolerance level. The optical properties at each FEM node are then used in conjunction with the Beer-Lambert relation and the Mie scattering relation to determine Hb and HbO_2 relative concentrations (as they, along with water, are the chief absorbing agents in biological tissue) and scattering power/amplitude respectively.

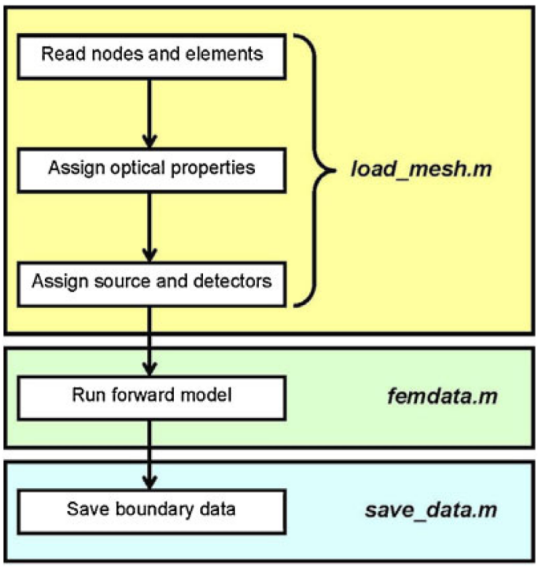

Figure 2(a): Flow-chart of DOT Forward Model (taken from Dehghani et al.)

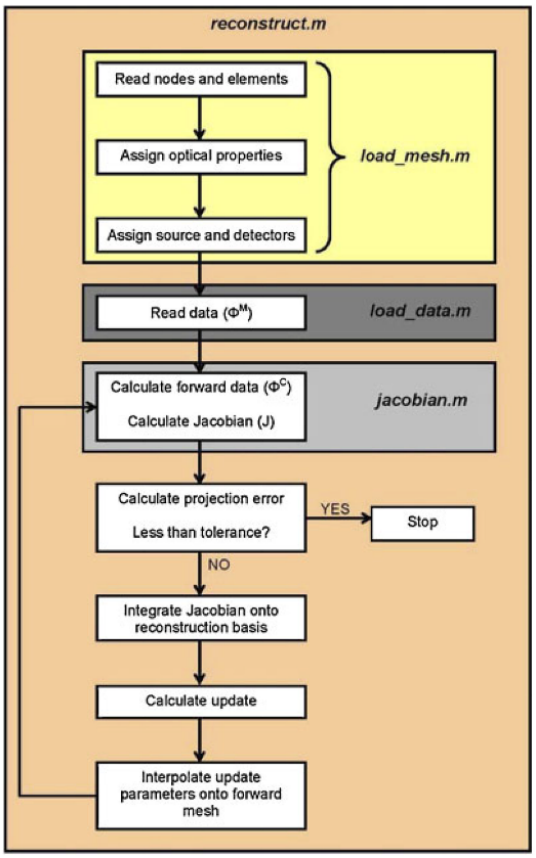

Figure 2(b): Flow-chart of DOT Reverse Model (taken from Dehghani et al.)

As of the writing of this paper, the DOT apparatus is still under construction. As a result, the scope of this study instead focused on the forward model portion of the DOT methodology. In particular, the primary aim of this paper was to conduct multiple sensitivity analyses regarding the calculated outputs for intensity and phase shifts with respect to specific initial conditions to minimize observed variability in calculated outputs from the forward model.

3 Methodology

Throughout the course of this study, a software package known as NIRFAST (Near Infrared Fluorescence and Spectral Tomography) was used to conduct the forward model simulations. NIRFAST is an open-source, MATLAB®-based software package developed by researchers at the University of Exeter and Dartmouth College specifically for the reconstruction of DOT scans based on experimental data [2].

Running a forward model simulation in NIRFAST begins by generating a volume mesh. The volume mesh is an arrangement of optical nodes (defined by the FEM implementation to discretize the continuous nature of the medium being imaged) that are interconnected using a Delaunay triangulation algorithm to approximate the random walk nature of light within a diffuse medium [2]. As a result, the surface shape and structure of the mesh closely resembles that of the object being imaged. In order to generate a mesh in NIRFAST, the distance at which each FEM node will be placed from its nearest neighbor (referred to as the node distance, abbreviated ND) must be specified. Each of these nodes are then assigned default optical parameter values of 0.01 mm^{-1} for the absorption coefficient, 0.33 mm for the diffusion coefficient, and 1.33 for the index of refraction, which are consistent with values for water [2]. No anomalous points with differing absorption coefficients, diffusion coefficients, or indices of refraction from the standard values were introduced to the mesh in this study. The coordinates at which the different sources and detectors will be placed on the surface of the mesh must then be specified before construction of the mesh is complete.

While NIRFAST provided pre-made meshes for various tissues, a simple mesh shaped as a rectangular prism (with dimensions of 50 x 30 x 10 mm) with an array of six sources and twelve detectors was used for the entire duration of this study. A visualization for this mesh can be seen in Figure 3.

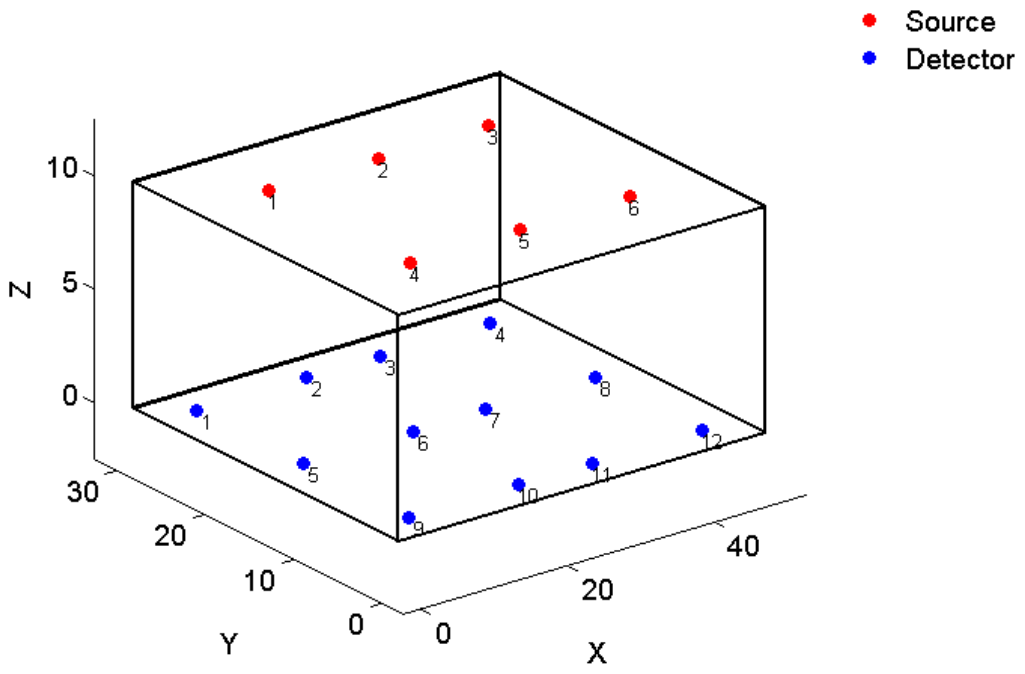

Figure 3: Arrangement of sources and detectors in the constructed mesh used in the forward models

The coordinates for the placement of sources and detectors within this mesh were as follows:

Sources		Detectors		
1: (10, 23, 10)	4: (10, 7, 10)	1: (5, 27, 0)	5: (5, 15, 0)	9: (5, 3, ,0)
2: (25, 23, 10)	5: (25, 7, 10)	2: (20, 27, 0)	6: (20, 15, 0)	10: (20, 3, 0)
3: (40, 23, 10)	6: (40, 7, 10)	3: (30, 27, 0)	7: (30, 15, 0)	11: (30, 3, 0)
		4: (45, 27, 0)	8: (45, 15, 0)	12: (45, 3, 0)

Light sources were modulated sinusoidally at a frequency of 100 MHz, which allowed for the values of intensity and phase shifts from specific source-detector pairs to be known. As a result, each forward model yielded seventy-two total measurements for both the intensity and phase shifts calculated at each detector.

4 Results

4.1 Comparison of Identical Forward Model Simulations

In a homogeneous medium with no variation in the absorption, scattering, or refraction parameters, the calculated intensity and phase shift should only depend on the geometric distance between a source and detector. However, because the path of light in a diffuse media is a random walk with absorption, a stochastic component in NIRFAST's forward calculation will produce some variation from the expected average value. In order to quantify the extent of this variation, two identical forward simulations were run and their respective intensity values for identical source-detector pairs were plotted against each other in Figure 4.

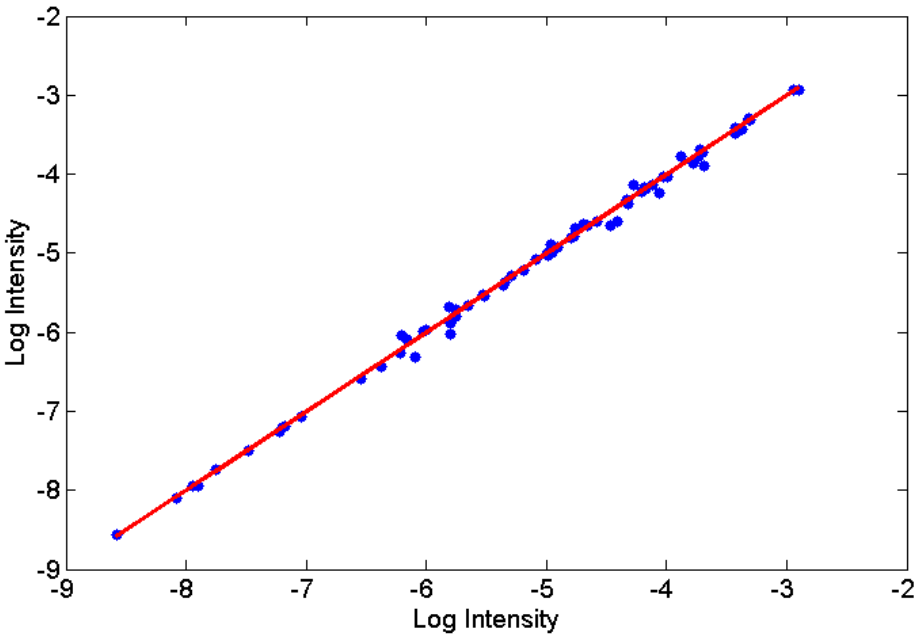

Figure 4: Comparative intensity calculations for two identical forward models. Mesh ND = 0.7 mm

If NIRFAST did not introduce any stochastic variables into its calculations, the results of the two simulations would have been identical and each of the plotted intensity values in Figure 4 would have been collinear along y = x (shown in red in Figure 4). A residual plot of Figure 4 was created in order

to determine which values yielded the largest overall differences in intensity values and to determine if any patterns existed in the magnitude of the differences in intensity between the two forward models (Figure 5). It can be seen that the variation in intensity between the two forward models is largest for intensities between 0.1% and 10% of the maximum signal size. It appears that the vast majority of the calculated intensities vary within approximately 12% of each other (with residual values between -0.05 and 0.05 on a log scale), with the exception of the aforementioned variation between points 0.1% and 10% of the maximum signal size.

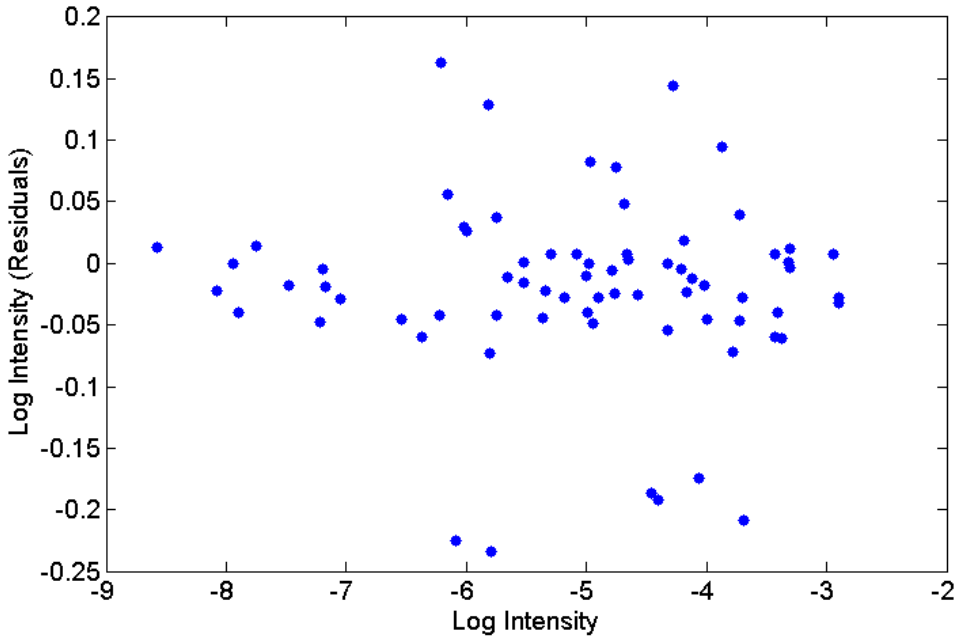

Figure 5: Residuals of intensities for two forward models

To determine if any relationships existed between the values of the residuals in intensity and the source at which they originated, the residuals were ordered from most negative (assigned a value of one) to most positive (seventy-two) and plotted against the source number (specified in Figure 3) associated with that residual (Figure 6). The process was then repeated for the specific detector numbers at which the original intensity value for each residual was calculated (Figure 7).

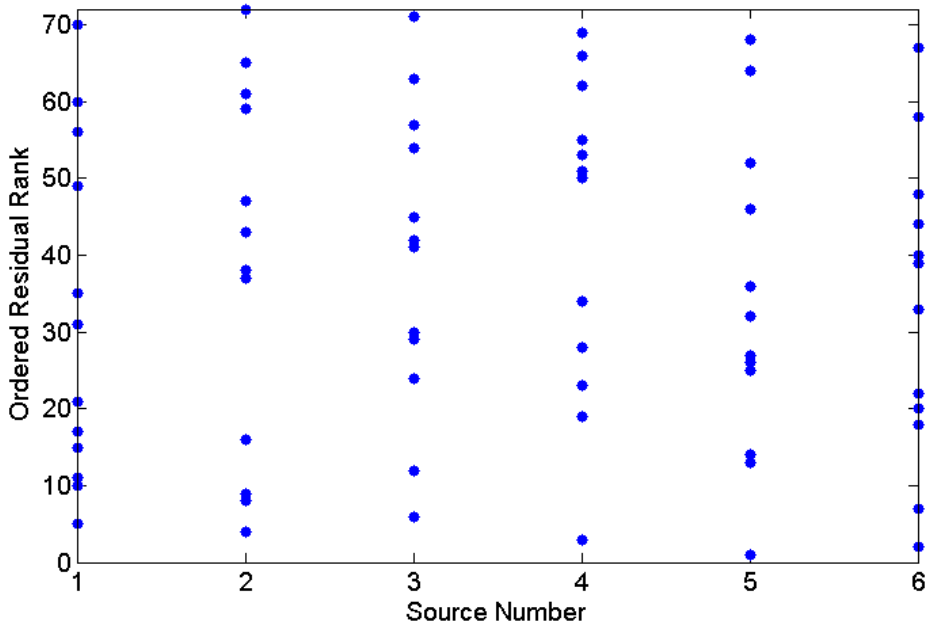

Figure 6: Distribution of residuals for the log in calculated intensity at each source. The smallest residual value (-0.234) was assigned a rank of one, and the largest value (0.163) was assigned a rank of seventy-two

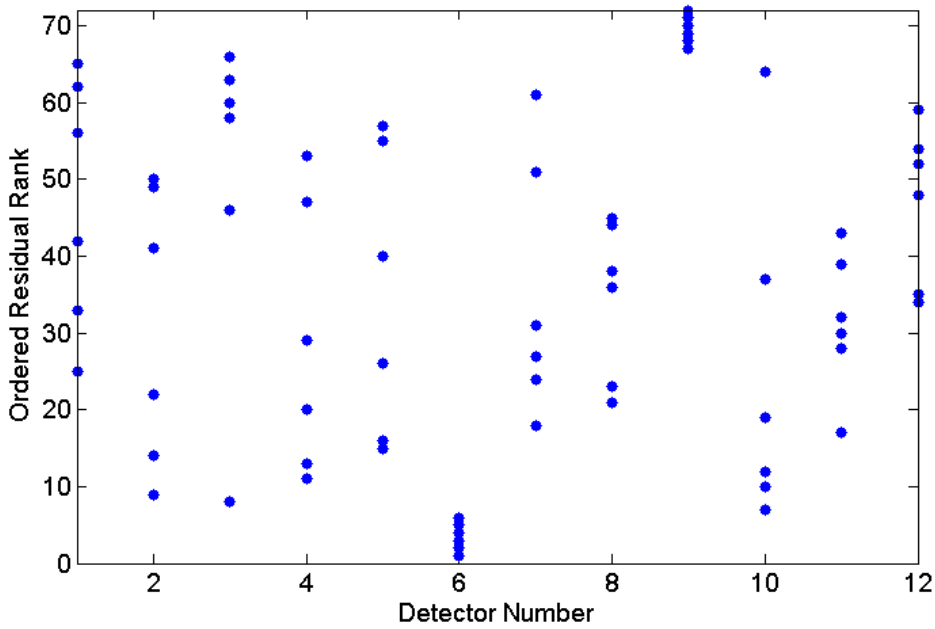

Figure 7: Distribution of residuals for the log in calculated intensity at each detector. Residuals were ranked identically to those in Figure 6

As the distribution of the ordered residuals across the sources in Figure 6 appears to be quite uniform across the domain, there does not appear to be a convincing relationship between the magnitude of the residual and the source from which the simulated photon originated. However, after comparing the residuals in the log of the intensity to the specific detector involved in their calculation, it was found that the six most positive differences in intensity were all calculated at detector nine while the six most

negative differences in intensity were all calculated at detector six. There did not appear to be any other significant patterns in the differences in intensities relative to the other detectors. If the large variations in intensities were due to photons being simulated as leaving the mesh (at which point NIRFAST removes them from the system), then the expectation would be that the most negative differences would be calculated at detectors along the edges of the mesh while the detectors in the center should remain more or less constant. However, the fact that the largest negative differences were calculated at detector six, a central detector, and that the largest positive differences were found at detector nine, an edge detector, indicates that other (possibly unrealistic) factors are inherent in NIRFAST's calculations rather than just a random walk and loss of signal at the boundaries.

The RMS for the differences in the logs of the intensity between identical forward model simulations was 0.0748, which represents the expected deviation in the log of calculated intensities from source-detector pairs in identical forward model simulations. This corresponds to an average percent change of 18.8% for identical source-detector calculations in intensity, which is much larger than what should be permitted if the only variable is a stochastic random walk component. Again, this indicates that other factors are at play in NIRFAST's calculations.

The same procedures used to analyze the magnitude of the residuals for the calculated intensities and their relationship to individual sources and detectors was repeated for the phase shifts calculated by the forward model. The phase shift values were plotted against each other (Figure 8), as were the corresponding residual values for each (Figure 9).

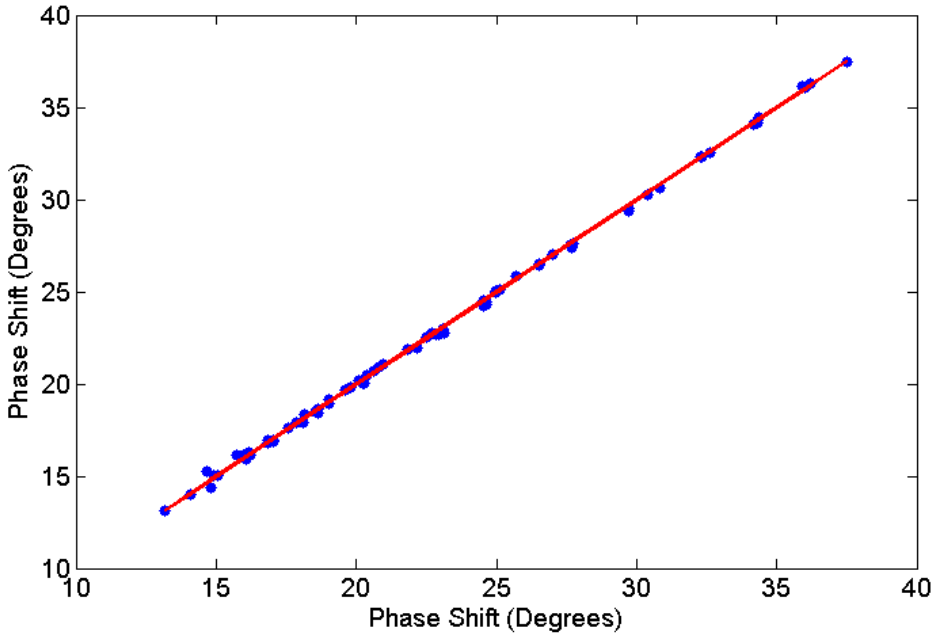

Figure 8: Comparative phase shift calculations for two identical forward models. Mesh ND = 0.7 mm

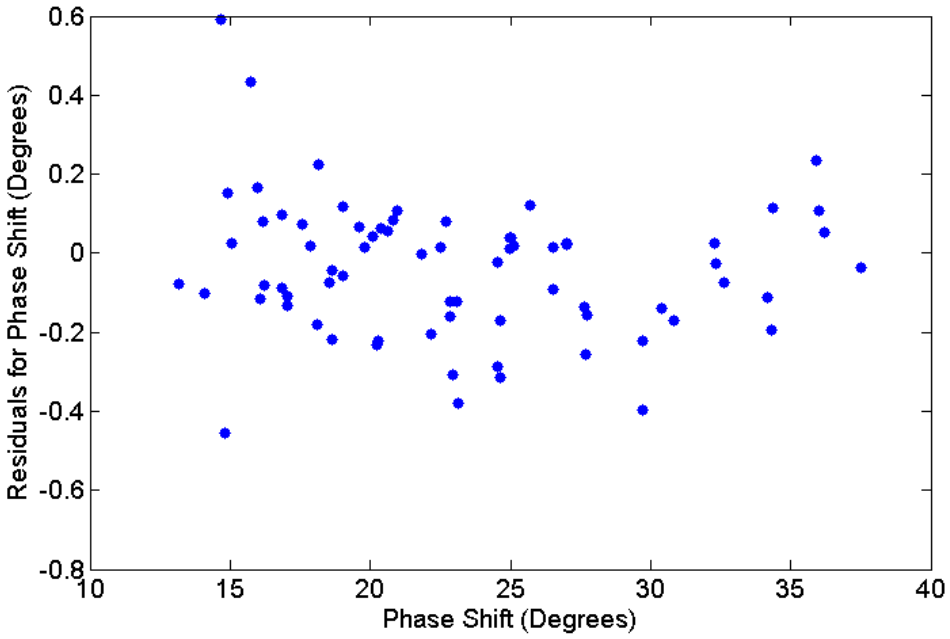

Figure 9: Residuals of phase shifts for two identical forward models

The average difference in phase shift between identical forward model simulations was 0.2049 degrees and was fairly uniform with the residuals never exceeding 0.6 degrees or falling below -0.5 degrees. Given that the average calculated phase shift across both identical forward models was 23.1627 degrees, this does not represent nearly as significant a variation as was seen in the intensity case. In a random walk process, the expectation would be that phase shift values would increase as a function of the total distance traveled (or, correspondingly, the distance from a source to a particular detector). Figure 9 would then seem to suggest that the largest variation in phase shifts are seen in the case of sources and detectors that are relatively close to each other, while phase shifts for larger source-detector distances approach a greater degree of uniformity.

In order to determine if any pattern existed between the individual sources and detectors and the variation in the phase shift values they were associated with, the ordered phase shift residuals of each source-detector pair were plotted with respect to their source number (Figure 10) and detector number (Figure 11). Unlike the intensity case, there did not appear to be any patterns for the variability in residuals in the phase shifts when stratifying for both sources and detectors.

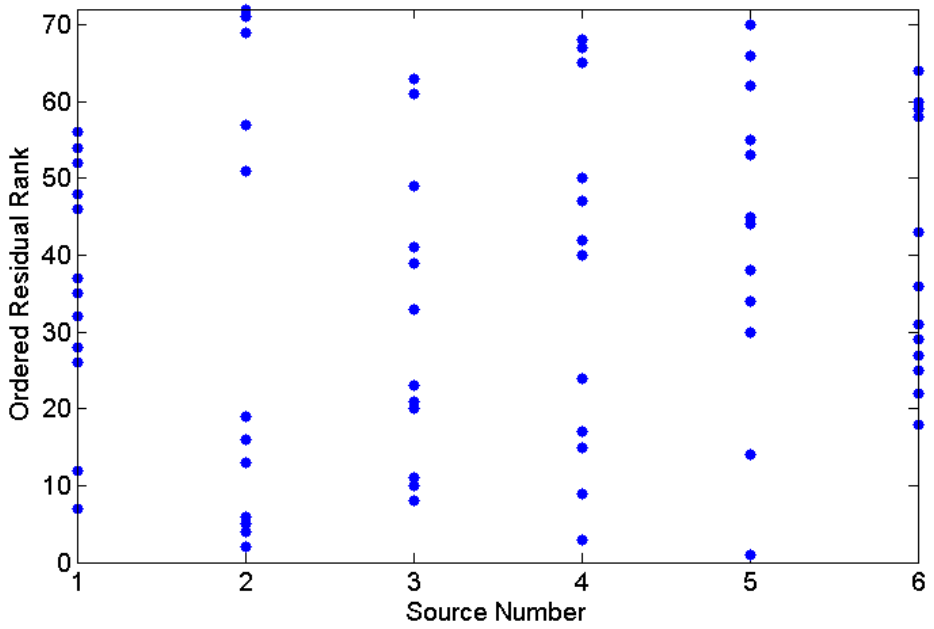

Figure 10: Distribution of residuals for the calculated phase shift at each source. The smallest residual value (-0.456 degrees) was assigned a rank of one, and the largest value (0.590 degrees) was assigned a rank of seventy-two

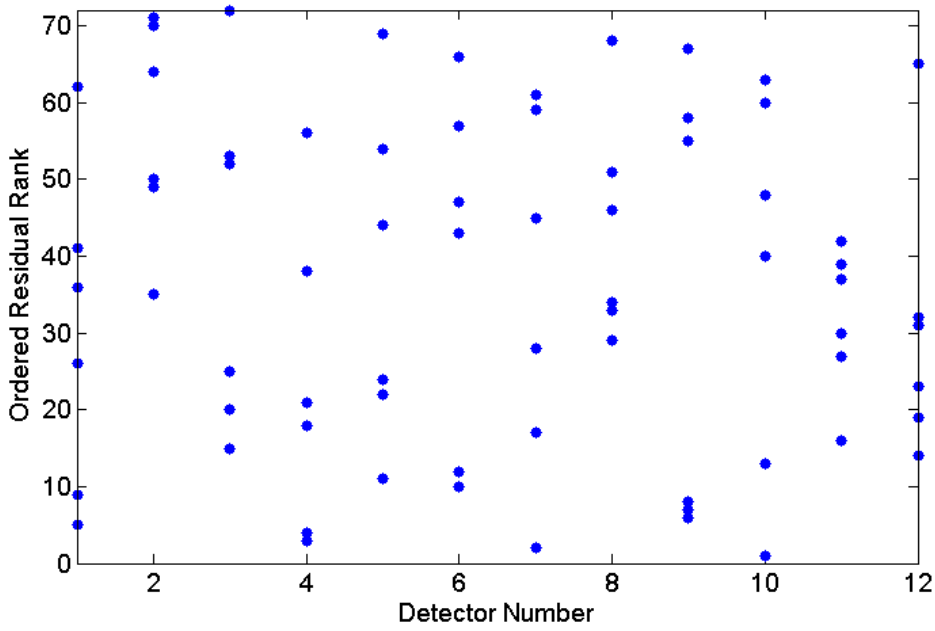

Figure 11: Distribution of residuals for the calculated phase shift at each detector. Residuals were ranked identically to those in Figure 10

4.2 Forward Model Calculations and Source-Detector Distance

As previously stated, the only factor affecting calculated intensities and phase shifts within a homogeneous medium should be the distance between the optical sources and detectors. However, stochastic variables in the absorbing random walk and additional, potentially unphysical, contributions

have been shown to have an effect on NIRFAST's forward model calculations. Therefore, in order to determine the extent to which NIRFAST adheres to the physical expectations regarding this system, a single forward model simulation was conducted, and the calculated intensities and phase shifts generated by each combination of sources and detectors were analyzed with respect to their respective source-detector distances in the mesh.

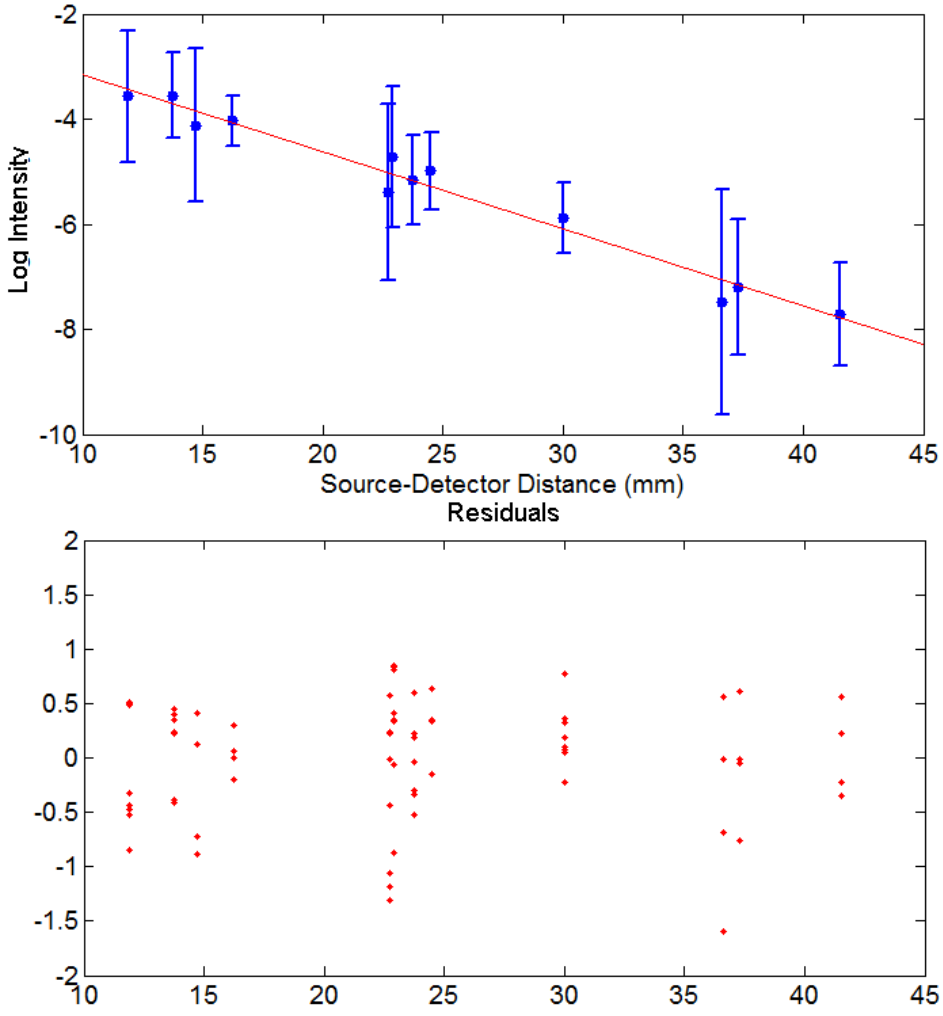

Figure 12: Average intensity as a function of source-detector distance and the residuals of each data point with respect to the linear regression line. The error bars represent the standard deviation in the log of intensity for each source detector distance. Mesh ND = 0.7 mm

Due to the symmetric and planar arrangements of sources and detectors as seen in Figure 3, there were twelve unique source-detector distances in the constructed mesh. The relationship between the average log intensity at each distance and the residual values of each individual data point with respect to a best-fit linear regression line is seen in Figure 12. As one would expect, the general trend appears to be that the intensity falls off as source-detector distance increases. The error bars represent the standard deviation in the log of intensity for a given source-detector distance. The residuals for the log of intensity in Figure 12 were extremely varied for each source-detector distance, with all but two having a spread of at least one order of magnitude.

Ordering the residuals from most-negative to most-positive and plotting each at its associated source and detector number (Figure 13 and Figure 14 respectively) illuminates whether or not any particular sources and detectors disproportionately exhibit non-uniform variability in intensity with respect to source-detector distance.

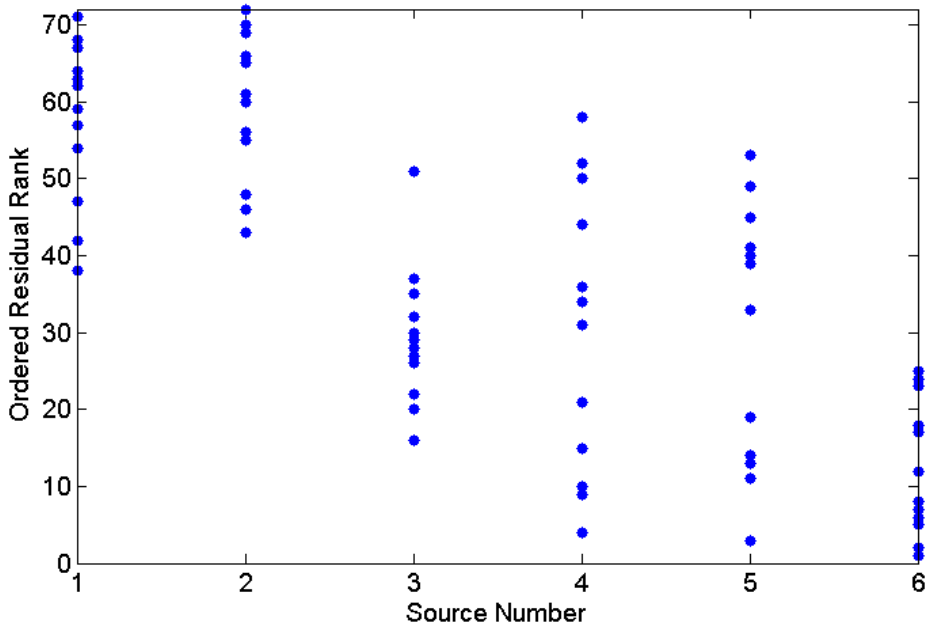

Figure 13: Distribution of residuals for the log of the calculated intensities at each source. The smallest residual value (-1.594) was assigned a rank of one, and the largest value (0.851) was assigned a rank of seventy-two

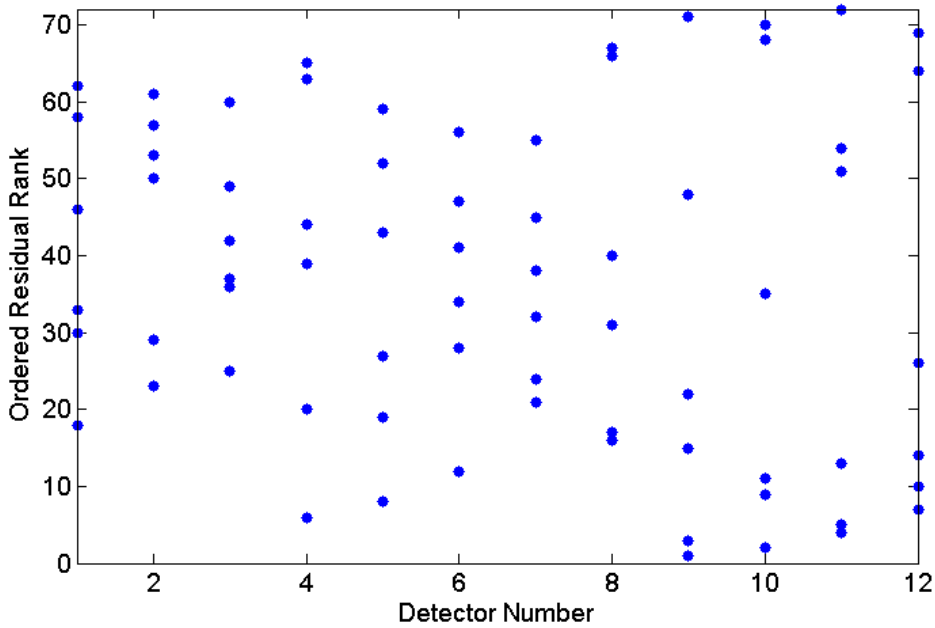

Figure 14: Distribution of residuals for the log in calculated intensity at each detector. Residuals were ranked identically to those in Figure 13

The magnitude of the residuals in Figure 12 seemed to overwhelmingly depend on the location of the source (as opposed to dependence on the detector position, which was seen in Figure 7 in the case comparing variations in identical forward model simulations). The most negative residuals originated from sources six, five, and four, while the most positive residuals originated from sources one and two (Figure 13). Additionally, detectors twelve, eleven, ten, and nine appear to calculate largest positive and negative residual values, while the others generally experienced much smaller residual values (Figure 14).

As was expected, plotting the phase shifts from the forward model versus the source-detector distance yielded a fairly strong linear, upward trend (Figure 15). The residuals for the calculated phase shifts were plotted against the source-detector distance of each. Relative variation in the residuals was found to be much smaller than the variation in intensities, as was also found to be the case in Section 4.1 of this paper.

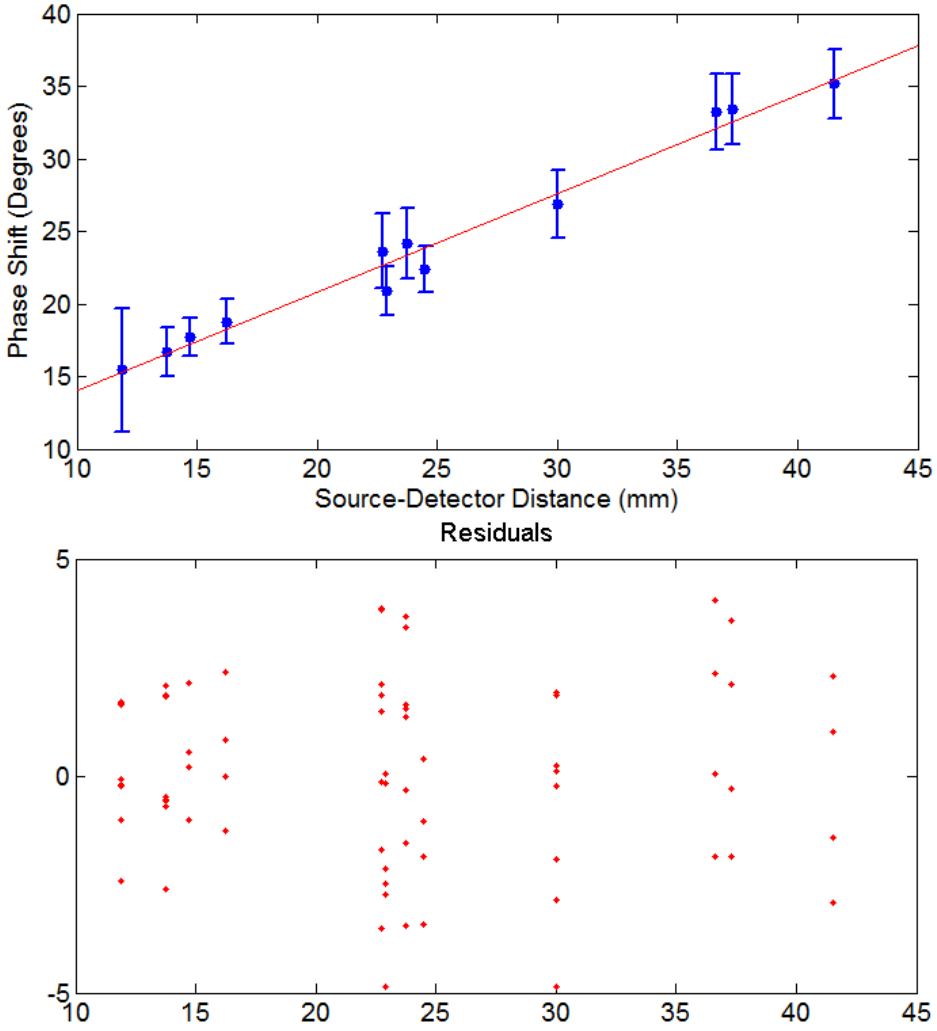

Figure 15: Average phase shift as a function of source-detector distance and the residuals of each data point with respect to the linear regression line. The error bars represent the standard deviation in the phase shift for each source detector distance. Mesh ND = 0.7 mm

Ordered residual plots of the calculated phase shift with respect to each source (Figure 16) does not appear to elucidate any sort of relationship between the two variables. However, the most negative variations in the phase shift appear to be calculated at detectors twelve, eight, and four, which all lie along the same edge of the mesh as seen in Figure 3 (Figure 17).

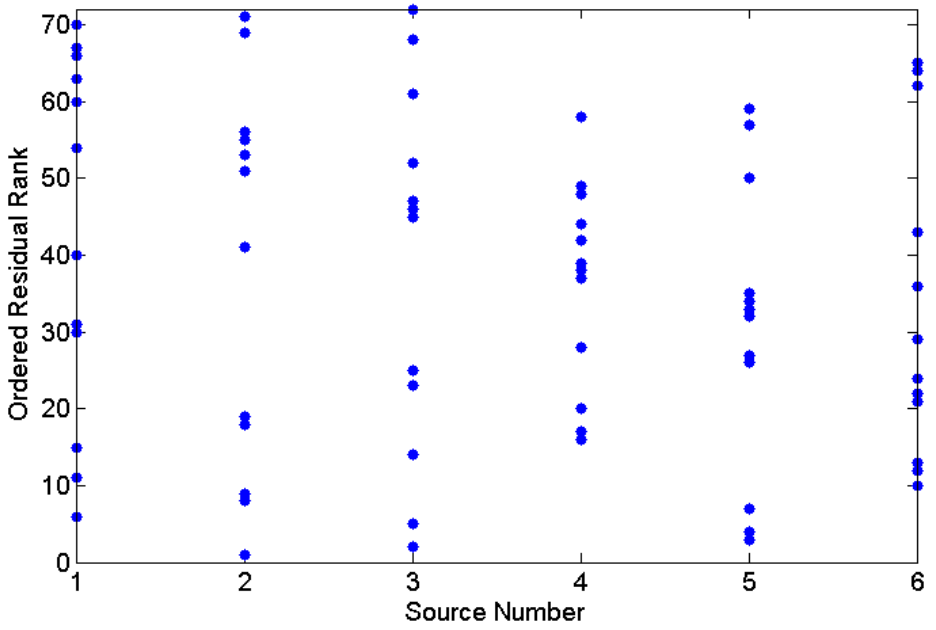

Figure 16: Distribution of residuals for calculated phase shifts at each source. The smallest residual value (-4.857 degrees) was assigned a rank of one, and the largest value (4.048 degrees) was assigned a rank of seventy-two

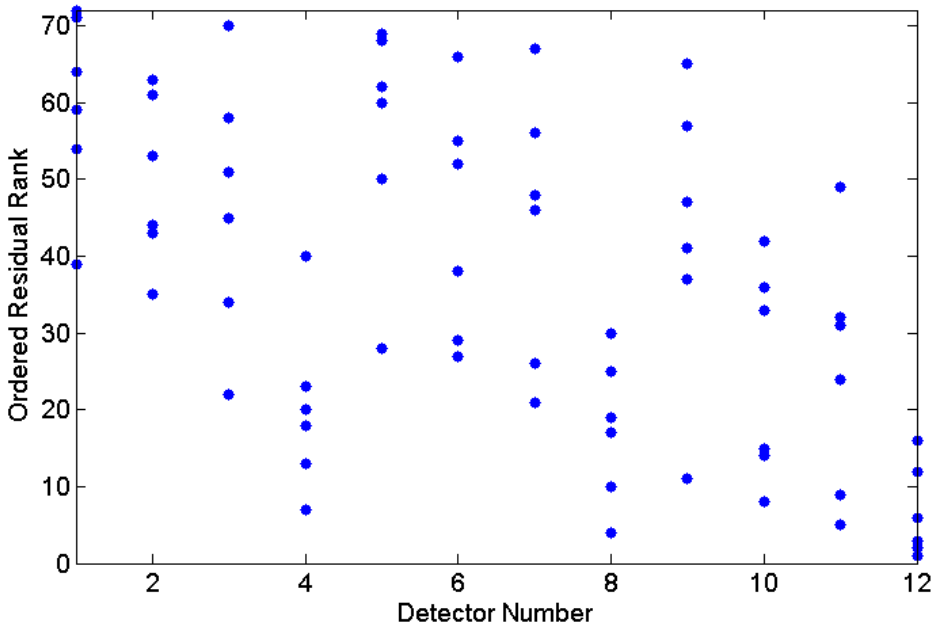

Figure 17: Distribution of residuals for the calculated phase shifts at each detector. Residuals were ranked identically to those in Figure 16

It was determined in Figures 12 and 15 that the large variations in intensity and phase shift (though the residual values for phase shift were much more reasonable compared to the values for intensity) could not be explained by differences in source-detector distances as the residuals were approximately uniform and homoscedastic when viewed solely as a function of their associated source-detector distance. However, by considering how each residual value varied with respect to specific sources and detectors, some patterns began to emerge. In regards to intensity, the magnitude of residuals were found to be stratified with respect to the y-coordinate of the sources, with the most negative residuals originating from sources six, five, and four (on edge) while the most positive residuals originated at the opposite sources. Predictably, the detectors located the shortest geometric distance from sources four, five, and six were found to have the largest calculated positive and negative residual values. While it does not appear that variations in phase shift were affected by the locations of different sources, detectors twelve, eight, and four appear to consistently calculate larger negative phase shift residual values. The lack of symmetry in these results is somewhat perplexing as there does not appear to be a good reason for why certain clusters of sources and detectors experience more drastic variations than others. While the patterns noted above may hint at photons leaving the mesh having an impact on calculated results, the effects do not appear to be dramatic or persuasive enough to fully account for the high degree of variation seen in the calculated intensities and phase shifts.

4.3 Variable Node Distances for Forward Model Calculations

Implementing the finite element method allows for the discretization of continuous volumes into smaller subdomains. The distance separating these subdomains is referred to as the node distance (ND) in NIRFAST. As the node distance decreases, the finite element approximation more-closely approximates the real, continuous case. However, this occurs at the cost of significant increases in computing time. Conversely, increasing the node distance would lead to a decrease in the accuracy of the simulation. To determine the effect that the node distance had on the simulation results, meshes with three different node distances were run in otherwise identical forward models, and their resulting intensities and phases were compared.

The first set of compared meshes had node distances of 0.65, 0.7, and 0.75 mm. To see how relatively small changes in the node distance would effect calculated intensity and phase shift values, the log of intensity values from forward models with node distances of 0.65 and 0.75 mm were plotted against the log intensity of the mesh with a node distance of 0.7 mm (Figure 18). No change in the intensity values would be represented by all points lying on the line y = x (labeled in black). There appears to be a great amount of consistency in the intensity values for all three node distances within an order of magnitude of the peak intensity. As the intensity falls off past that point, however, the variations in intensity for the three node distances increases dramatically.

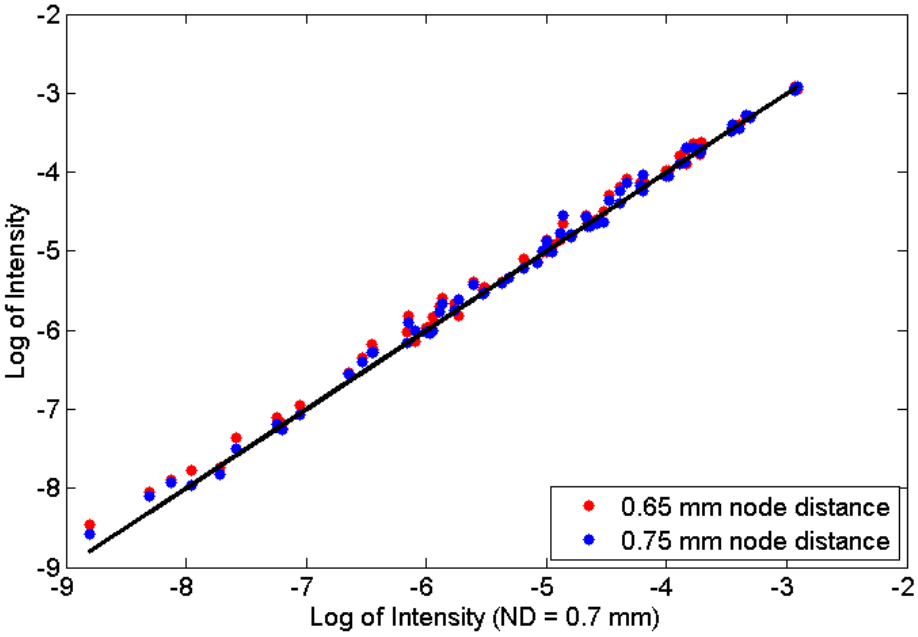

Figure 18: Relative comparison of intensities for different node distances

The aforementioned increase in variation past one order of magnitude from the peak intensity in Figure 18 can be more-clearly seen in a residual plot of the same data (Figure 19). The variation in the intensity is somewhat contained (around a 0.05 difference in the log of the intensity) up to an order of magnitude from the maximum intensity value. Past that point, however, the variation increases dramatically, but does not follow an obvious pattern or exhibit predictable behavior.

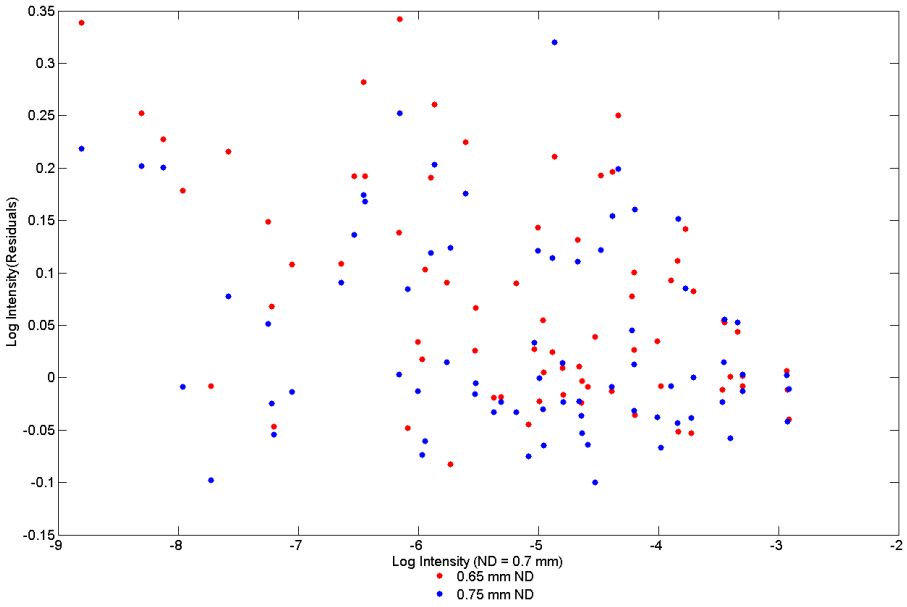

Figure 19: Residuals of intensities for different node distances

To see the extent of the effect that reducing the node distance would have on the variation in

intensity, forward models with three new meshes (with node distances about a factor of two smaller than those used in Figures 18 and 19) were calculated, and their resulting variations in intensity were compared in Figure 20.

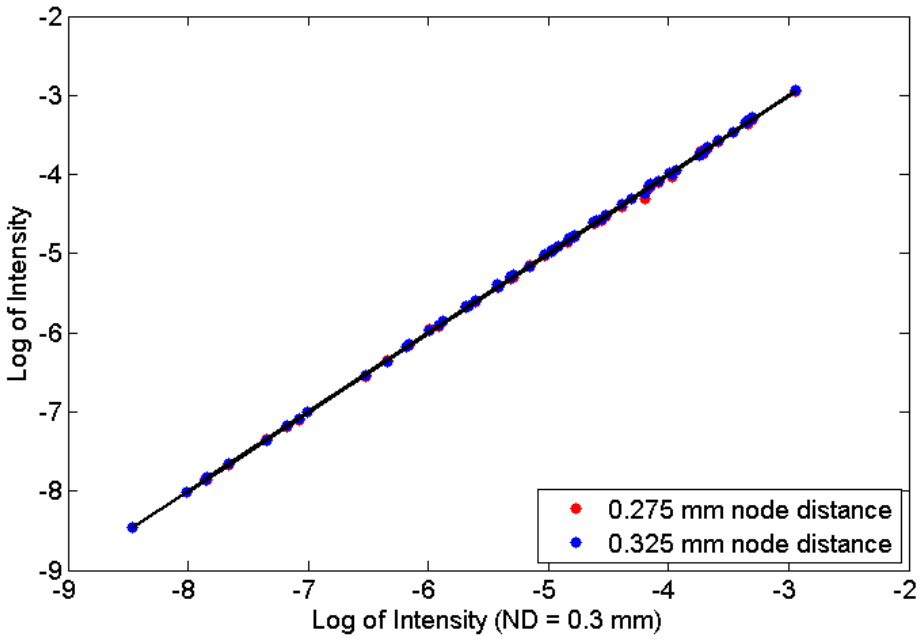

Figure 20: Relative comparison of intensities for different node distances

We can immediately see that the intensities of these three node distances are much more consistent than in Figure 9 for both the strongest and weakest intensities. Interestingly, the largest variation seems to occur between one and three orders of magnitude from the strongest signal, although almost no significant, discernible variation occurs for any other values along the domain.

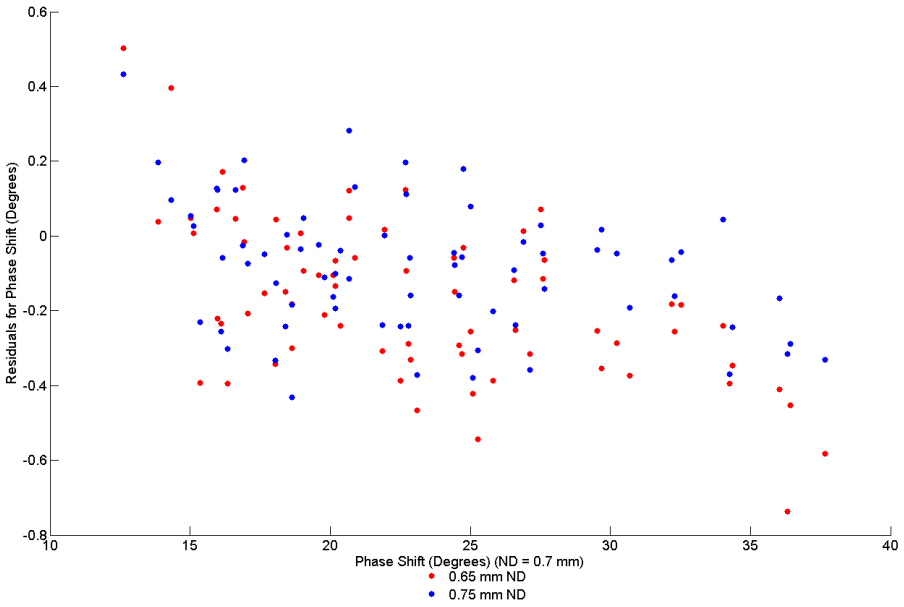

Figure 21: Residuals of intensities for different node distances

A residual plot of Figure 20 (seen in Figure 21) confirms our suspicions regarding the variations in intensity between Figure 18 and Figure 20. Compared to the varied intensity values seen in Figure 18, Figure 20 represents a dramatic improvement. The maximum residual values are decreased by roughly a factor of two (with the exception of two outlier points in the 0.275 mm node distance mesh). The distributions of the differences also appear more uniform in Figure 21 and the residuals do not increase as dramatically for values past an order of magnitude from the maximum intensity value.

The phase shifts for meshes with identical node distances to the meshes used previously in Section 4.3 were then compared. Meshes with node distances of 0.65 and 0.75 mm were compared against a mesh with node distance of 0.7mm with respect to their phase shifts (Figure 22). The variation in phase shift appeared to be quite uniform across all values of the phase shift with no extreme variation at any point. There did not appear to be a generally larger variability in phase shift for smaller phase shift values as was seen in Figure 9. A residual plot of the phase shift values in Figure 22 confirmed these observations (Figure 23). Figures 22 and 23 indicate that the residual values for the phase shifts were much smaller in relative magnitude and more uniform than the values observed when comparing the intensities for different node distances.

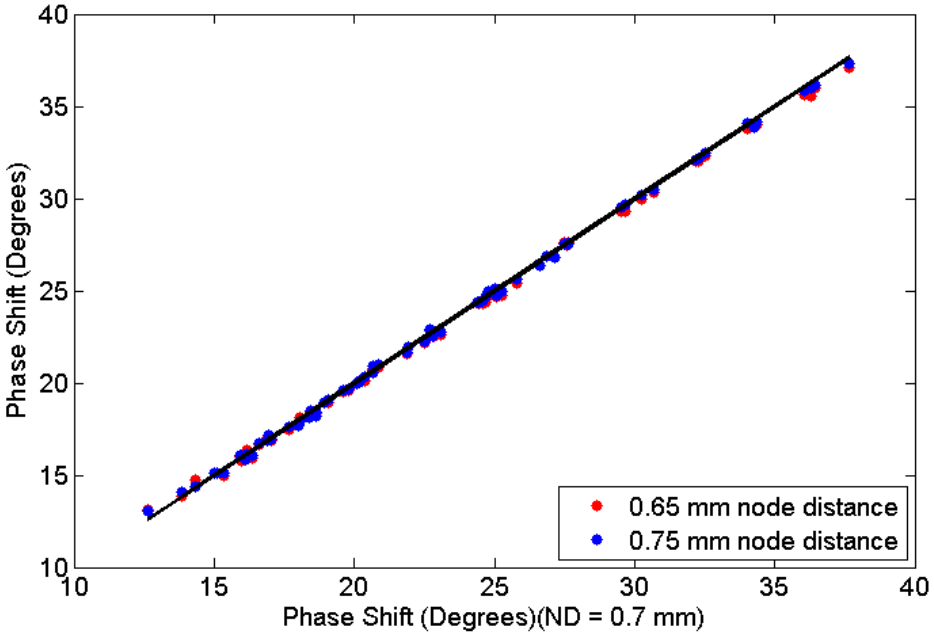

Figure 22: Relative comparison of phase shifts for different node distances

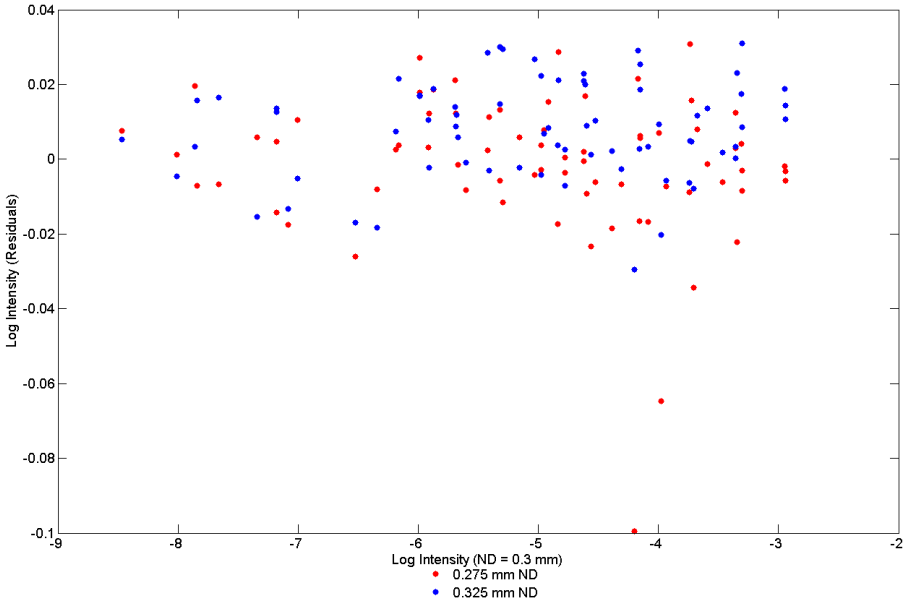

Figure 23: Residuals of phase shifts for different node distances

The phase shifts were then compared to meshes with node distances of 0.275, 0.3, and 0.325 mm (Figure 24). Even though the residuals in phase shifts seen in Figure 23 were relatively small compared to the phase shift values themselves, reducing the node distance still did much to reduce the variation. Compared to Figure 23, the residuals in phase shift for the meshes with smaller node distances were reduced by at least a factor of 1.5 for all values (Figure 25).

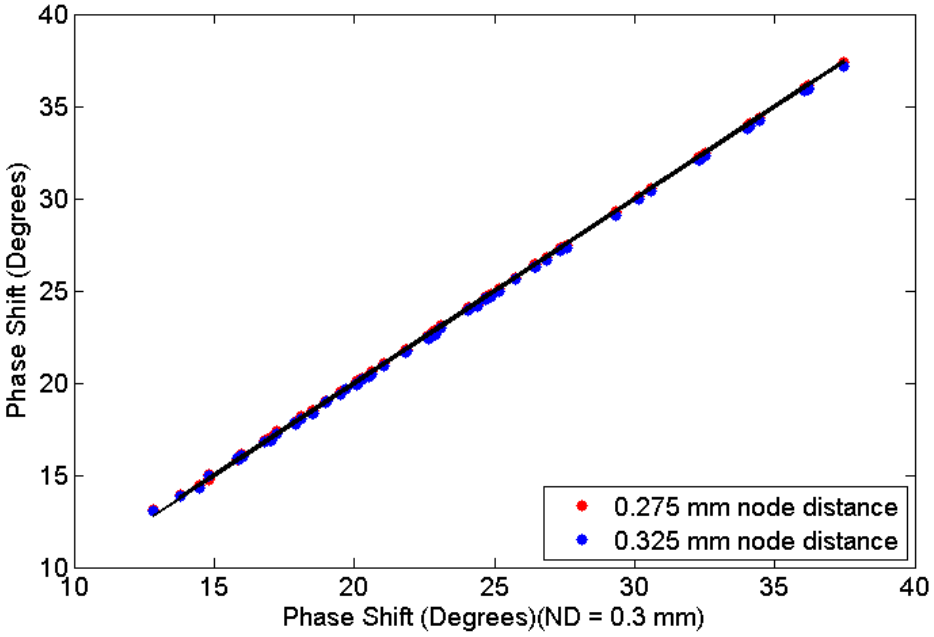

Figure 24: Relative comparison of phase shifts for different node distances

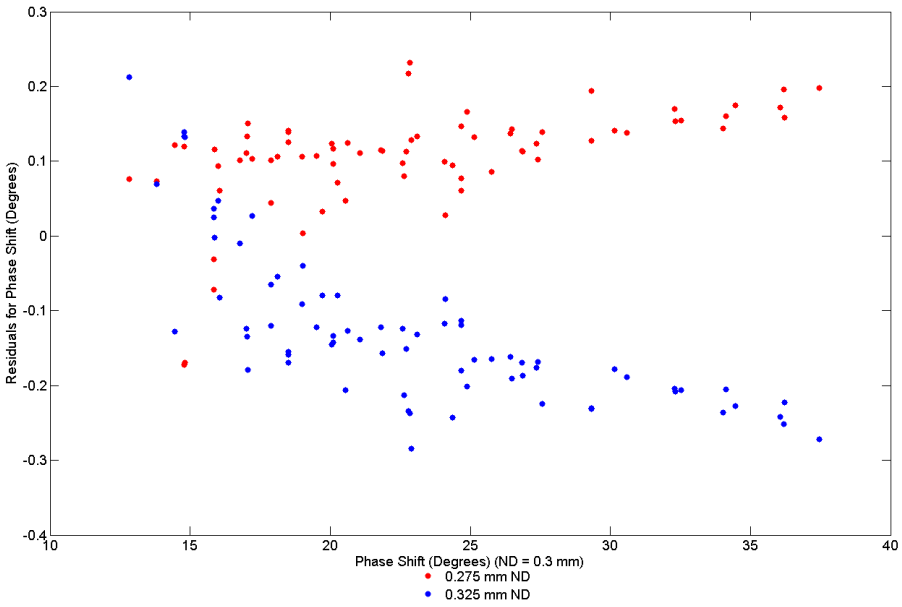

Figure 25: Residuals of phase shifts for different node distances

The results in this section indicate that a significant portion of the high degree of variation seen in values for intensity and, to a lesser extent, phase shift as a result of forward model calculations is due in large part to the node distance parameter. Node distances of 0.7 mm resulted in the simulated values from the forward model assuming unrealistic values for intensity. While the variations for phase shift values were generally less variable than the intensity case, both intensity and phase shift were found to be sensitive to the node distance, and both demonstrated significantly more uniformity in their values when the node distance was decreased.

5 Discussion

The impetus behind this analysis occurred when significant differences in intensity and phase shift values occurred for (theoretically) identical forward models conducted by two different researchers. This was particularly concerning as results from the forward simulation are then compared to the experimental data in the reverse model in order to determine the optical properties at each node in the real-life sample. The introduction of such drastic errors in the forward model could then potentially be compounded in the reverse model, which would result in inaccurate reconstructions of real DOT data.

Analysis of the forward model calculations performed in NIRFAST in this study was based on three general tests. Two forward model calculations, with identical input settings and parameters, were compared to determine what inherent variability (due to the stochastic random walk) existed in NIRFAST's calculations. It was found that variation in intensities for identical sources and detectors for the two forward models was typically within 12%, but exceeded 60% in some cases. This error range far surpassed what would have been expected due to a stochastic component in NIRFAST and could also not be convincingly explained by photons leaving the outer boundary of the mesh (as was seen when comparing the variations in intensity to the particular source and detector utilized in simulating the original intensity value). Because the mesh used could be considered optically homogeneous, the expectation would be that the intensity and phase shift calculated values should exist solely as a function of the source-detector distance. However, even when intensities from a single forward model calculation were plotted against the source-detector distance, calculated intensities were routinely found to vary by an order of magnitude or more. Again, this high level of variability could not be sufficiently explained by the individual sources, detectors, and their relative placement within the mesh. The third test involved performing a sensitivity analysis on the node distance and observing how slight variations yielded different intensity and phase shift values. While a high degree of variations in intensity were observed for meshes with node distances on the order of 0.7 mm, reducing the node distance to 0.3 mm was found to drastically reduce the previously observed variation in intensities. Although the relative variation in the phase shifts were consistently much smaller than the differences in intensity values, variation in the phase shifts were also found to decrease considerably with the smaller node distances.

6 Conclusion

The implications of these results are that NIRFAST produces physically-unsound results when the node distance is too large. In general, a node distance greater than the diffusion coefficient (standardized at 0.33 mm in NIRFAST) should not be used, and obviously smaller node distances will better-approximate real-life results, although this is done at the cost of much greater computational times. While the variations in intensity and phase shift from the forward model calculations were found to be sensitive to the node distance, it still remains unclear if increasingly smaller node distances will completely eliminate the unexpected variations in intensity, or if a case of decreasing marginal returns exists such that past a certain point, reducing the mesh size will do little in decreasing variability. In addition, it still remains to be seen just how much of an effect (and how consistent of an effect) photons leaving the outer boundary of the mesh has on variability of intensity and phase shift values. Once the variation in intensity and phase shifts are better-understood in the forward model, seeing how these variations

and errors manifest themselves in the image reconstruction process in the reverse model will be critical to increasing the overall accuracy of the DOT process.

7 Acknowledgments

I would like to express my immense gratitude to my advisors within the William and Mary Physics Department, Professor Cooke and Professor Manos, for their guidance, expertise, and patience throughout the duration of this project. I would also like to extend my sincerest thanks to our collaborators in the Biology Department: Professor Buchser and especially Lyndah Lovell. In addition, special thanks to Professor Hoatson for coordinating and guiding the physics senior research projects for the 2014-2015 academic year.

8 References

[1] Wang, L. V., Wu, H. I. (2007) Biomedical Optics (Principles and Imaging). Hoboken, New Jersey: Wiley

[2] Dehghani, H., Eames, M. E., Yalavarthy, P. K., Davis, S. C., Srinivasan, S., Carpenter, C. M., ... Paulsen, K. D. (2008). Near infrared optical tomography using NIRFAST: Algorithm for numerical model and image reconstruction. Communications in Numerical Methods in Engineering, 25(6), 711–732. doi:10.1002/cnm.1162

[3] Eggebrecht, A. T., Culver, J. P. (2014). High-Density Diffuse Optical Tomography: Imaging Distributed Function and Networks in the Human Brain. Bio-Optics World, 7(4). Retrieved from http://www.bioopticsworld.com

[4] Hielscher, A. H., Bluestone, A. Y., Abdoulaev, G. S., Klose, A. D., Lasker, J., Stewart, M., ... Beuthan, J. (2002). Near-Infrared Diffuse Optical Tomography. Disease Markers, 18(5-6), 313–337. doi:10.1155/2002/164252

[5] Zhan, Y., Eggebrecht, A. T., Culver, J. P., & Dehghani, H. (2012). Image Quality Analysis of High-Density Diffuse Optical Tomography Incorporating a Subject-Specific Head Model. Frontiers in Neuroenergetics, 4, 6. doi:10.3389/fnene.2012.00006

[6] Hall, Z. J., Bertin, M., Bailey, I. E., Meddle, S. L., & Healy, S. D. (2014). Neural correlates of nesting behavior in zebra finches (Taeniopygia guttata). Behavioural Brain Research, 264(100), 26–33. doi:10.1016/j.bbr.2014.01.043

[7] M. Jermyn, H. Ghadyani, M.A. Mastanduno, W. Turner, S.C. Davis, H. Dehghani, and B.W. Pogue, "Fast segmentation and high-quality three-dimensional volume mesh creation from medical images for diffuse optical tomography," J. Biomed. Opt. 18 (8), 086007 (August 12, 2013), doi: 10.1117/1.JBO.18.8.086007.

[8] H. Dehghani, M.E. Eames, P.K. Yalavarthy, S.C. Davis, S. Srinivasan, C.M. Carpenter, B.W. Pogue, and K.D. Paulsen, "Near infrared optical tomography using NIRFAST: Algorithm for numerical model and image reconstruction," Communications in Numerical Methods in Engineering, vol. 25, 711-732 (2009).