Effect of wavelength selection on the accuracy of blood oxygen saturation estimates obtained from photoacoustic images

Roman Hochuli, Paul C. Beard and Ben Cox

Department of Medical Physics & Biomedical Engineering, University College London, UK.

ABSTRACT

In photoacoustic tomography (PAT) the image contrast is due to optical absorption, and because of this PAT images are sensitive to changes in blood oxygen saturation (sO_2). However, this is not a linear relationship due to the presence of a non-uniform light fluence distribution. In this paper we systematically evaluate the conditions in which an approximate linear inversion scheme – which assumes the internal fluence distribution is unchanged when the absorption coefficient changes – can give accurate estimates of sO_2. A numerical phantom of highly vascularised tissue is used a test case for this assumption. It is shown that using multiple wavelengths over a broad range of the near-infrared spectrum yields inaccurate estimates of oxygenation, while a careful selection of wavelengths in the 620-920nm range is likely to yield more accurate oxygenation values. (We demonstrate that a 1D fluence correction, obtained from the average decay rate in the image, can further improve the estimates.) However, opting to use these longer wavelengths involves sacrificing signal-to-noise ratio in the image, as the absorption of blood is low in this range. This results in an inherent trade-off between uncertainty in the sO_2 estimates due to fluence variation and error due to noise. This study shows that the depth to which sO_2 can be estimated accurately using a linear approximation is limited *in vivo*, even with idealised measurements, to at most 3mm. In practice, there will be even greater uncertainties affecting the estimates, eg. due to bandlimited or partial-view acoustic detection.

Keywords: Oxygenation, oxygen saturation, Monte Carlo

1. INTRODUCTION

Blood oxygen saturation, defined as the ratio of oxygenated to total haemoglobin concentration, is a fundamental measure of physiological activity and is a key indicator of tissue function and pathology. Oxygen partial pressure was measured in the past using invasive catheterisation and gas extraction methods, but today systemic blood oxygenation can be measured noninvasively via pulse oximetry.[1] Obtaining localised measures of oxygenation has been more challenging; brain activation can be imaged using Blood Oxygenation Level Dependent Magnetic Resonance Imaging,[2] which responds to changes in both blood volume and/or venous deoxy-haemoglobin concentration but not oxygen saturation, while oxygenation can be estimated with low spatial resolution using optical techniques, such as near-infrared spectroscopy (NIRS)[3] and diffuse optical tomography (DOT).[4]

Optical methods are attractive due to their relative low cost compared with MR techniques and their high specificity; the absorption spectrum of blood in the 'near infrared window', where water and blood absorption are low so the light can penetrate the tissue, is strongly dependent on the level of oxygenation allowing highly specific imaging. However, purely optical approaches afford limited spatial resolution due to the highly diffusive nature of biological tissue. Photoacoustic imaging, which involves the deposition of optical energy in order to generate MHz-frequency acoustic waves, does not suffer from this limitation due to the fact that ultrasonic waves undergo relatively little scattering in soft tissue.

Photoacoustic imaging can provide micron-resolution 3D images of soft tissue *in vivo*[5] and it is for this reason that it is particularly well suited to imaging oxygenation in applications such as cancer staging.[6,7] However, photoacoustic images are not proportional to absorption, or indeed blood oxygenation; a photoacoustic image,

Send correspondence to: r.hochuli@ucl.ac.uk

$p_0(\mathbf{x}, \lambda)$ the initial acoustic pressure, is in fact proportional to the product of the local light fluence, $\Phi(\mathbf{x}, \lambda)$, as well as the absorption coefficient at that position, $\mu_a(\mathbf{x}, \lambda)$:

$$p_0(\mathbf{x}, \lambda) = \Gamma\Phi(\mathbf{x}, \lambda; \mu_a, \mu_s, g)\mu_a(\mathbf{x}, \lambda), \tag{1}$$

where Γ is the Grüneisen parameter, λ is the optical wavelength, μ_s is the scattering coefficient and g is the anisotropy factor. The absorption coefficient can be expanded in terms of a sum over products of α_k, the molar absorption coefficient for the k^{th} chromophore and c_k, the concentration of this chromophore:

$$p_0(\mathbf{x}, \lambda) = \Gamma\Phi(\mathbf{x}, \lambda; c_k, \mu_s, g) \sum_{k=1}^{K} \alpha_k(\lambda)c_k(\mathbf{x}) \tag{2}$$

Eq. (2) describes the propagation of light, absorption of optical energy by tissue chromophores and generation of the initial acoustic pressure. In the case where blood is assumed to be the only chromophore present in the medium, Eq. (2) can be succinctly written in matrix form for a given position $\mathbf{x}$ in the tissue

$$\begin{bmatrix} p_0(\lambda_1) \\ \vdots \\ p_0(\lambda_N) \end{bmatrix} = \Gamma \begin{bmatrix} \Phi(\lambda_1) & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \Phi(\lambda_N) \end{bmatrix} \begin{bmatrix} \alpha_{HbO2}(\lambda_1) & \alpha_{Hb}(\lambda_1) \\ \vdots & \vdots \\ \alpha_{HbO2}(\lambda_N) & \alpha_{Hb}(\lambda_N) \end{bmatrix} \begin{bmatrix} c_{HbO2} \\ c_{Hb} \end{bmatrix}, \tag{3}$$

where c_{HbO2} and c_{Hb} are the concentrations of oxygenated- and deoxygenated-haemoglobin, respectively.

As we define oxygen saturation as,

$$sO_2 = \frac{c_{HbO2}}{c_{HbO2} + c_{Hb}}, \tag{4}$$

oxygenation can be estimated by inverting Eq. (3) for the haemoglobin concentrations and substituting into Eq. (4). However, the matrix containing $\Phi(\lambda)$ along the diagonal is typically unknown, given that the fluence is itself wavelength-dependent and depends on the medium properties throughout the tissue $\mu_a(\mathbf{x})$, $\mu_s(\mathbf{x})$ and g. One approach is to model the fluence explicitly using a model of light transport, such as a Monte Carlo (MC) model of light transport, the diffusion approximation (DA) or the radiative transfer equation (RTE), in order to divide it out of the photoacoustic image.[8–12] However, as the absorption and scattering are not known *a priori*, these models must be used to iteratively estimate the fluence within an optimisation scheme at significant computational expense.

An alternative, often exploited in the literature, is to assume that the effect of the fluence on estimating chromophore concentrations is negligible,[13–16] i.e. $\Phi \propto \mathcal{I}$, the identity matrix, such that the concentrations of oxy- and deoxy-haemoglobin can be expressed as

$$\begin{bmatrix} c_{HbO2} \\ c_{Hb} \end{bmatrix} \approx \frac{1}{\Gamma} \begin{bmatrix} \alpha_{HbO2}(\lambda_1) & \alpha_{Hb}(\lambda_1) \\ \vdots & \vdots \\ \alpha_{HbO2}(\lambda_N) & \alpha_{Hb}(\lambda_N) \end{bmatrix}^{\dagger} \begin{bmatrix} p_0(\lambda_1) \\ \vdots \\ p_0(\lambda_N) \end{bmatrix}, \tag{5}$$

and $\dagger$ is used to indicate the pseudoinverse.

This approximation is not true in general and depends on the illumination wavelengths as well as tissue geometry and composition. With collimated, one-sided illumination the fluence decays much more strongly in the illumination direction, z, compared with the transverse directions; thus, by modelling the fluence as $\Phi_{\text{est}}(\lambda) = \exp(-\mu_a^{\text{fit}}(\lambda)z) \; \forall \; x, y$ this depth dependence can be reduced by dividing the entire image through by this correction factor. The method used to obtain $\mu_a^{\text{fit}}(\lambda)z$ involved integrating each image over x and y, taking

the log, and fitting a straight line to the resultant data as a function of z. Note that other approximate analytic models for the fluence exist, e.g. radial functions for the fluence in a cylindrical structure.[17] Inversions of this type do not aim to model the fluence explicitly and are approximate.

In this paper, we present a rigorous numerical examination of the accuracy of the assumptions that the fluence is wavelength-independent or it follows 1D propagation with depth using a realistic model of light propagation in tissue. We use a highly vascularised, realistic tissue phantom to demonstrate that through careful selection of the optical wavelengths, oxygenation estimates in the phantom can be improved without explicitly modelling the fluence; however, this leads to a fundamental and problematic dilemma. The wavelengths found to be optimal for the estimation of blood oxygenation are those where the absorption coefficient of blood is low and changes little between illumination wavelengths; therefore estimates made at these wavelengths contain significant uncertainties due to low signal-to-noise ratio, while estimates made using wavelengths where blood absorption is high are inaccurate due to the fluence spectrum becoming distorted.

2. GENERATION OF DIGITAL PHANTOM

Several studies in the literature aim to quantify either relative oxygenation changes or absolute oxygenation in the brain using PAI.[15, 16, 18, 19] This reason, combined with the fact that optically the brain is representative of many tissues of interest for PAI, means this anatomy is an optimal choice for this investigation.

The phantom was developed from an *ex vivo* μCT dataset of a mouse brain vascular cast, which offers high contrast of the vasculature with $2.5\mu m$ isotropic voxel size. This was down-sampled to $25\mu m$ isotropic voxel size, followed by a two-step segmentation process. An initial thresholding step was used to remove noise from the background; connected-component analysis was subsequently applied in order to identify the connected portions of the vasculature. A vessel of interest (VOI) was selected by choosing the largest component following connected-component analysis, while smaller components were determined to form the background vasculature.

In order to make this phantom more realistic, skin and skull layers had to be added in via post-processing as these were not present in the μCT image. This was achieved by dilating the segmented image volume (using Minkowski addition with a isotropic cubic $500\mu m$ kernel) by $500\mu m$ in all directions, twice sequentially, producing two layers, variable in thickness, that were approximately $500\mu m$ thick on average. These layers were used to simulate skin and skull with realistic thicknesses for mice.[20, 21]
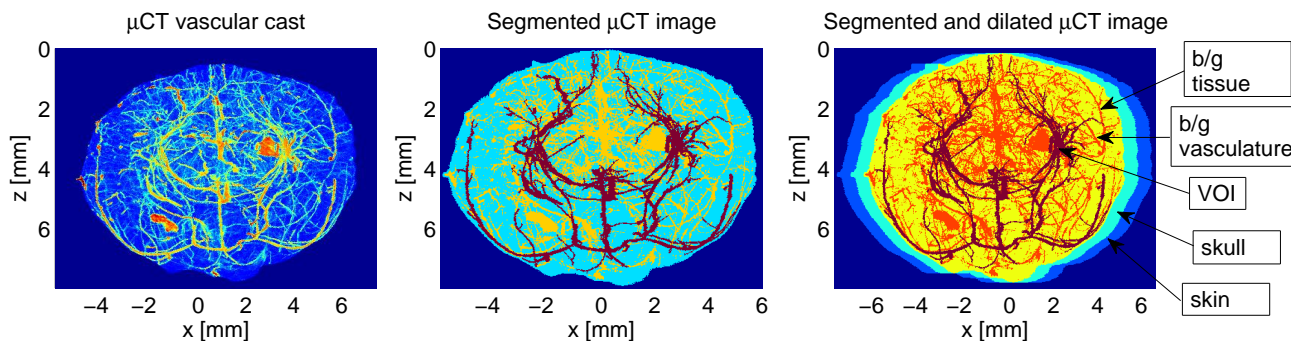

Figure 1. MIPs (coronal view) of μCT vascular cast (left), segmented μCT vascular cast (centre) and segmented and dilated (via Minkowski addition) μCT vascular cast (right). Background tissue and vasculature, skin and skull, and vessel of interest (VOI) denoted.

Vasculature in the brain had total haemoglobin concentration of $150\mathrm{gl}^{-1}$. The background tissue was assumed to consist of a volume fraction of water of 55% and 45% haemoglobin with total haemoglobin concentration $c_{HbT}^{bg} = 5.63\mathrm{gl}^{-1}$ and $sO_2 = 60.7\%$ oxygenation – values taken from rat brain cortical tissue.[22]

As the original dataset did not include skin and skull layers surrounding the brain, these had to be introduced via post-processing. The skin layer had $c_{HbT}^{skin} = 0.48\mathrm{gl}^{-1}$ at 98.5% oxygenation, 0.87% melanosome volume fraction with the absorption coefficient following the trend $\mu_{a,\mathrm{mel}} = 0.452 \left(\frac{\lambda}{500nm}\right)^{-3}$, and 21.4% water volume fraction.[23]

The skull had wavelength-dependent absorption and scattering coefficients between 0.025mm^{-1} and 0.036mm^{-1}, for μ_a, and 25mm^{-1} and 45mm^{-1}, for μ_s, where the scattering coefficient was obtained by fitting to scattering data of cortical bone.[24]

The scattering coefficient of the background tissue and skin regions varied as a function of wavelength according to $\mu'_s = A \left(\frac{\lambda}{500nm}\right)^{-B}$.[23] The brain and blood regions of the phantom had $A = 2.14$mm^{-1} and $B = 1.20$, the skin region had $A_{\mathrm{skin}} = 2.97$mm^{-1} and $B_{\mathrm{skin}} = -0.705$. The anisotropy factor was heterogeneously distributed with $g = 0.9$ for the tissue background and skin and $g = 0.9945$ for blood vessels.

3. MONTE CARLO MODELLING OF LIGHT TRANSPORT

There are many ways of modelling light transport in tissue, with each method having its own benefits and limitations. The RTE is the most commonly used equation for describing light propagation in scattering and non-scattering media, but finite element implementations of this model are excessively computationally intensive due to the need to discretise the medium in both space and angle. The diffusion approximation simplifies this model by only modelling near-isotropic propagation and reduces computational demands as no angular discretisation is required. However, the DA is only accurate where the light field is diffusive, i.e. far from index-matched boundaries and the source, and in turbid media. Monte Carlo models of light transport, considered the 'gold standard',[25] simulate the propagation of packets of energy, or 'photons', through the medium in order to accumulate the fluence. The photon packets are non-interacting and accurately model the RTE, meaning this model is both accurate in all situations which are of interest to us and can be massively parallelized.

For these reasons, MCX[26] was used to simulate the fluence in the phantom, which was illuminated using a collimated circular beam, 11.7mm in diameter, with top-hat profile at $z = 0$mm in the $+z$ direction in order to mimic illumination from a fibre tip. MC models always have some inherent spatially-varying variance in their fluence estimates due to an insufficient number of photons being simulated; as the number of photons in the simulation increases, the variance decreases and the fluence estimate converges to the 'true' solution predicted by the RTE. A series of experiments using an increasing number of photons allowed us to determine that simulating 10^{14} photons yields a *maximum* uncertainty in sO_2 estimates made using the model of 1.56% due to MC noise, meaning that estimates made on the vessel of interest show very little error due to MC noise.

4. ESTIMATING OXYGENATION IN VESSEL OF INTEREST

The phantom was illuminated as described in Section 3 with wavelengths in the 500-1000nm range at 20nm intervals in order to calculate the fluence. PA images were then formed by taking the product of the fluence with $\mu_a(\lambda)$ and with $\Gamma = 1$. It was assumed that $p_0(\lambda)$ was perfectly reconstructed, without any partial-view or bandlimited-acquisition artefacts.

4.1 Estimates using 26 wavelengths

A naive approach is to use a broad range of wavelengths across the NIR-visual range to perform the inversion with the aim of obtaining a better fit to the data. This inversion was performed as described by Eq. (5) with $N = 26$, using wavelengths in the 500-1000nm range at 20nm intervals.

Fig. 2 shows the maximum intensity projection (MIP) in the y-direction of the oxygenation estimate made in every voxel, overlaid on the anatomy of the image in greyscale. The oxygenation estimates in the VOI are of very limited accuracy and only a few voxels near the surface of the brain, around 1mm in depth, contain accurate oxygenation values to within 20% (absolute error calculated as $|sO_2^{\mathrm{true}} - sO_2|$). The impact of the 1D fluence correction is very small and only marginally improves the accuracy of sO_2 estimates made using 26 wavelengths.

4.2 Estimates using two or more wavelengths

It was demonstrated above (Fig. 2) that using many wavelengths across the 500-1000nm range, in an inversion that assumes unchanged fluence between wavelengths, yields inaccurate sO_2 estimates beyond depths of 1-2mm, even when a simple depth-dependent and wavelength-dependent fluence correction is applied to the images. In order to illustrate why the assumption that the fluence is unchanged with wavelength is unlikely to be accurate
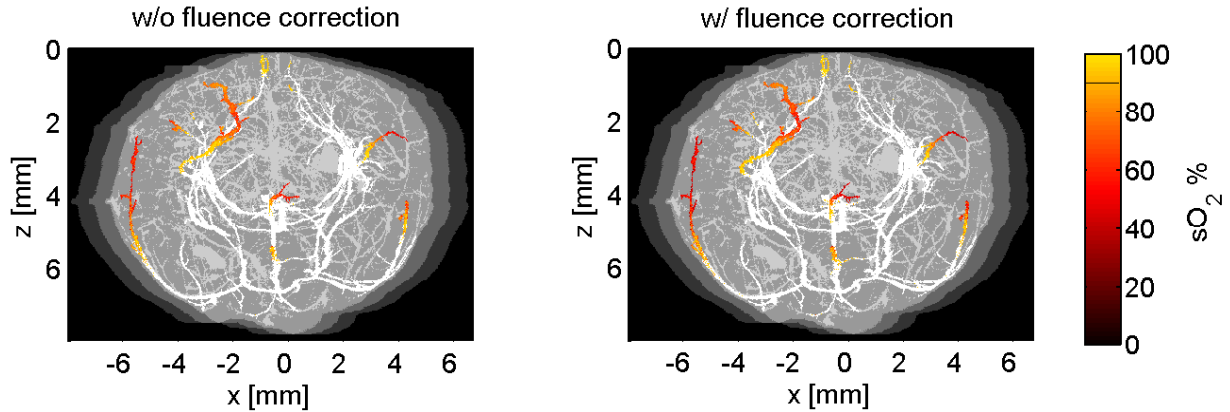
Figure 2. MIPs of oxygenation estimates in VOI made using 26-wavelength linear inversion with wavelengths in range 500-1000nm without fluence correction (left) and with fluence correction (right), both overlayed on greyscale phantom. True oxygenation 90% indicated by line on colour bar. White indicates sO_2 estimates outside valid range of 0-100%.

over a wide range of the spectrum, recall that a photoacoustic image is a function of both the absorption coefficient and the fluence; thus the change in photoacoustic amplitude with wavelength will in part be due to a change in the absorption coefficient as well as a change in fluence. In order for a linear inversion to be accurate, we require that the change in photoacoustic amplitude is predominantly due to a change in absorption, rather than fluence. This is formalised elsewhere, but it can be shown that we require $\frac{\delta\mu_a}{\mu_a(\lambda_1)} \gg \frac{\delta\Phi}{\Phi(\lambda_1)}$.

Thus, wavelengths (or indeed oxygenation states) that satisfy this condition are those that will yield the most accurate oxygenation estimates. Although not presented here, this was demonstrated using a MC experiment on a blood-filled tube submerged in an absorbing-scattering background; it was revealed that the most accurate oxygenation values are obtained when $\frac{\delta\mu_a}{\mu_a(\lambda_1)} \gg \frac{\delta\Phi}{\Phi(\lambda_1)}$. So, in a practical setting, can we determine a set of wavelengths that ensure this is the case?

Using the 26 PA images, inversions were run for *all possible* combinations of these 26 images $(\sum_{r=2}^{26} {}^{26}C_r = 6.71 \times 10^7)$ and sO2 in the VOI was averaged over all voxels. It was found that there were no particular wavelengths that consistently yielded accurate oxygenation estimates on average over the vessel; however, the wavelengths that appeared in the most accurate estimates (within 10% of the true value of 90%) were those in the 620-1000nm range, where the absorption coefficient is low and undergoes a small changes from λ1 ↔ λ2, which is what was anticipated given that the absorption coefficient undergoes small changes between different wavelengths in this range.

As choosing the set of N > 2 wavelengths in order to obtain the most accurate sO₂ estimates in an experiment configuation may be completel non-trivial, we can consider the average oxygenation estimates for the 325 combinations of N = 2 wavelengths, as these may be instructive of which regions of the spectrum are best-suited to obtaining accurate values of sO₂.

It was found that the wavelengths yielding the most accurate oxygenation estimates on average, to within 30% error at most, are those between 820nm and 960nm. This is consistent with the fact that when using more than two wavelengths, the inversion yielded accurate estimates when the wavelengths were in the 620-1000nm range and indicates that oxygenation is indeed more accurately estimated when the absorption coefficient undergoes small changes from one wavelength to another. It was again observed with two-wavelength inversions in experiments on a blood-filled tube submerged in an absorbing-scattering background that these wavelengths satisfy the relationship that the relative change in absorption must exceed that of the fluence.

It was found that with two-wavelength inversions the fluence correction has an important impact, shifting the optimal wavlengths from 820-960nm, without the fluence correction, to 600-920nm with the fluence correction applied. Without the application of the 1D fluence correction, the fluence is effectively 'coloured' as it passes through the medium, having its spectrum distorted by absorbing chromophores in the blood and skin. The minimal change in absorption in the 820-960nm range minimizes this effect. However, the 1D fluence correction

ameliorates this spectral colouring and therefore improves the accuracy of estimates made where the absorption of blood is higher (6$\tilde{0}$0nm) and therefore undergoes a more significant change between illumination wavelengths.

Averaging over many voxels at a range of depths in the domain results in a loss of spatial information, which may obscure the impact of the fluence correction spatially within the image. MIPs of sO_2 in the VOI, overlayed on a greyscale MIP of the phantom, are shown in Fig. 3, where the left-hand image was generated using data without the 1D fluence correction and the right-hand image used data with the fluence correction. The wavelengths used in these inversions were 640nm and 840nm as these were found to yield the most accurate oxygenation estimates on average with the fluence correction applied to the PA images.
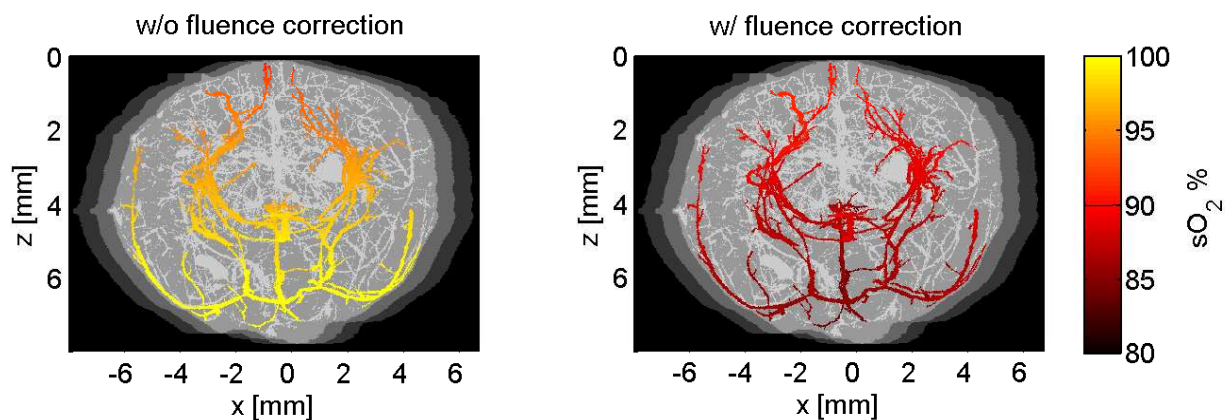

Figure 3. MIPs of oxygenation estimates in VOI made using two-wavelength linear inversion at 640/840nm without fluence correction (left) and with fluence correction (right), both overlayed on greyscale phantom. True oxygenation 90% indicated by line on colour bar. White indicates sO_2 estimates outside valid range of 0-100%.

It can be seen that without the fluence correction the oxygenation values are predominantly above 95%, while when the image is divided through by a depth-dependent fluence correction, the values are centered around the true value of 90% plus-or-minus a few % to a depth of about 8mm. Moreover, the sO_2 estimates show less variation with depth. Nevertheless, such accurate estimates were only possible by selecting wavelengths that yielded oxygenation estimates that were closest to that of the *known* true value of 90%. In practice, as the true oxygenation is not known, one would have to select two wavelengths in the 600-920nm range and apply the fluence correction as described in Section 1 and rely on the resultant oxygenation estimates' accuracy. An additional problem is that in an experimental setting, the presence of noise in the data will also corrupt estimates.

The impact of limited signal-to-noise in practice is that oxygenation estimates will be have some inherent uncertainty due to noise in the PA images. Optimg to use wavelengths in the 600-920nm range, where the absorption of blood is almost two orders of magnitude lower than in the 500-680nm range, involves sacrificing PA signal in favour of accurate sO_2 estimates. However, the presence of noise will corrupt estimates thus reducing their accuracy. This Catch-22 is fundamental to obataining accurate oxygenation estimates in PA imaging.

5. SUMMARY AND CONCLUSIONS

In this paper we investigate the accuracy of inversion schemes that assume fluence is unchanged when the absorption coefficient changes with the aim of systematically evaluating under what conditions this assumption is valid in estimating blood oxygenation. These linear inversions were applied to a highly vascularised mouse brain phantom, surrounded by skin and skull, with a vessel-of-interest selected.

We initially performed an inversion with 26 wavelengths across the NIR range (500-1000nm). This yielded highly inaccurate estimates beyond depths of 1-2mm, even with the application of a depth- and wavelength-dependent fluence correction. This led to us asking the question of whether this approach is sensible, given that the absorption coefficient of 90% oxygenated blood changes by two orders of magnitude over this range. By running multiwavelength linear inversions for between 2 and 26 distinct wavelengths, it was demonstrated

that using wavelengths where the absorption coefficient of blood is significantly different yields inaccurate sO_2 estimates due to spectral colouring of the fluence being more marked.

As hand-picking more than two wavelengths within the 600-1000nm range deemed optimal for reconstructing oxygenation is non-trivial in an experimental setting without explicit knowledge of the internal fluence distribution, the size of the problem was reduced by performing a linear inversion for 325 possible pairs of 26 wavelengths. It was found that without correcting the images using a 1D fluence correction, the most accurate oxygenation values on average over the VOI were obtained in the 820-960nm range. However, the application of such a correction shifted the range where sO_2 is estimated most accurately to 600-920nm, where the absorption coefficient can undergo more significant changes between wavelengths and potentially provide higher SNR.

Nevertheless, that the wavelength pairs that produced the most accurate oxygenation estimates were consistently in the 600-960nm range revealed a fundamental trade-off: using wavelengths where the absorption of blood is high (<600nm) yields high SNR but produces significant spectral colouring of the fluence, thus negatively impacting our ability to estimate oxygenation accurately; however, by using longer wavelengths (600-1000nm), where the fluence is less coloured throughout the domain, means sacrificing SNR which also produces inaccurate sO_2 estimates. With these two sources of uncertainty, it is clear that great care must be taken when using simple inversions to estimate sO_2 to ensure accuracy. This represents a severe limitation in using linear inversions to estimate oxygenation from PA images.

REFERENCES

[1] J. W. Severinghaus and Y. Honda, "History of Blood Gas Analysis," *Journal of Clinical Monitoring* **3**(2), pp. 135–138, 1987.

[2] S. Ogawa and T. Lee, "Brain magnetic resonance imaging with contrast dependent on blood oxygenation," *Proceedings of the National Academy of Sciences of the United States of America* **87**, pp. 9868–9872, 1990.

[3] A. Villringer, J. Planck, C. Hock, L. Schleinkofer, and U. Dirnagl, "Near infrared spectroscopy (NIRS): A new tool to study hemodynamic changes during activation of brain function in human adults," *Neuroscience Letters* **154**, pp. 101–104, May 1993.

[4] A. P. Gibson, J. C. Hebden, and S. R. Arridge, "Recent advances in diffuse optical imaging," *Physics in Medicine and Biology* **50**, pp. R1–R43, Feb. 2005.

[5] E. Z. Zhang, J. G. Laufer, R. B. Pedley, and P. C. Beard, "In vivo high-resolution 3D photoacoustic imaging of superficial vascular anatomy," *Physics in Medicine and Biology* **54**, pp. 1035–46, Feb. 2009.

[6] J. Laufer, P. Johnson, E. Zhang, B. Treeby, B. Cox, B. Pedley, and P. Beard, "In vivo preclinical photoacoustic imaging of tumor vasculature development and therapy," *Journal of Biomedical Optics* **17**, pp. 056016–1 – 056016–8, May 2012.

[7] J. Laufer, A. Jathoul, P. Johnson, E. Zhang, M. Lythgoe, R. B. Pedley, M. Pule, and P. Beard, "In vivo photoacoustic imaging of tyrosinase expressing tumours in mice," in *Proc. of SPIE, Photons Plus Ultrasound: Imaging and Sensing*, A. A. Oraevsky and L. V. Wang, eds., **8223**, pp. 82230M–1 – 82230M–5, Feb. 2012.

[8] B. T. Cox, S. R. Arridge, and P. C. Beard, "Estimating chromophore distributions from multiwavelength photoacoustic images.," *Journal of the Optical Society of America. A, Optics, image science, and vision* **26**, pp. 443–455, Feb. 2009.

[9] T. Saratoon, T. Tarvainen, B. T. Cox, and S. R. Arridge, "A gradient-based method for quantitative photoacoustic tomography using the radiative transfer equation," *Inverse Problems* **29**, p. 075006, July 2013.

[10] J. Laufer, B. Cox, E. Zhang, and P. Beard, "Quantitative determination of chromophore concentrations from 2D photoacoustic images using a nonlinear model-based inversion scheme.," *Applied Optics* **49**, pp. 1219–33, Mar. 2010.

[11] A. Pulkkinen, B. T. Cox, S. R. Arridge, J. P. Kaipio, and T. Tarvainen, "A Bayesian approach to spectral quantitative photoacoustic tomography," *Inverse Problems* **30**, p. 065012, June 2014.

[12] A. V. Mamonov and K. Ren, "Quantitative photoacoustic imaging in the radiative transport regime," *Communications in Mathematical Sciences* **12**(2), pp. 201–234, 2014.

[13] J. Laufer, C. Elwell, D. Delpy, and P. Beard, "In vitro measurements of absolute blood oxygen saturation using pulsed near-infrared photoacoustic spectroscopy: accuracy and resolution.," *Physics in Medicine and Biology* **50**, pp. 4409–4428, Sept. 2005.

[14] R. Fainchtein, "Local determination of hemoglobin concentration and degree of oxygenation in tissue by pulsed photoacoustic spectroscopy," *Proceedings of SPIE* **3916**, pp. 19–33, 2000.

[15] X. Wang, Y. Pang, G. Ku, X. Xie, G. Stoica, and L. V. Wang, "Noninvasive laser-induced photoacoustic tomography for structural and functional in vivo imaging of the brain.," *Nature Biotechnology* **21**, pp. 803–806, July 2003.

[16] K. M. Stantz, "Photoacoustic spectroscopic imaging of intra-tumor heterogeneity and molecular identification," *Proceedings of SPIE* **6086**, pp. 608605–608605–12, 2006.

[17] Z. Guo, C. Favazza, A. Garcia-Uribe, and L. V. Wang, "Quantitative photoacoustic microscopy of optical absorption coefficients from acoustic spectra in the optical diffusive regime," *Journal of Biomedical Optics* **17**(6), p. 066011, 2012.

[18] X. Wang, X. Xie, G. Ku, L. V. Wang, and G. Stoica, "Noninvasive imaging of hemoglobin concentration and oxygenation in the rat brain using high-resolution photoacoustic tomography.," *Journal of Biomedical Optics* **11**(2), p. 024015, 2006.

[19] I. Y. Petrov, Y. Petrov, D. S. Prough, D. J. Deyo, I. Cicenaite, and R. O. Esenaliev, "Optoacoustic monitoring of cerebral venous blood oxygenation through extracerebral blood.," *Biomedical optics express* **3**, pp. 125–36, Jan. 2012.

[20] L. S. Hansen, J. E. Coggle, J. Wells, and M. W. Charles, "The influence of the hair cycle on the thickness of mouse skin.," *The Anatomical Record* **210**, pp. 569–73, Dec. 1984.

[21] A. J. Lin, M. a. Koike, K. N. Green, J. G. Kim, A. Mazhar, T. B. Rice, F. M. LaFerla, and B. J. Tromberg, "Spatial frequency domain imaging of intrinsic optical property contrast in a mouse model of Alzheimer's disease.," *Annals of Biomedical Engineering* **39**, pp. 1349–57, Apr. 2011.

[22] D. Abookasis, C. C. Lay, M. S. Mathews, M. E. Linskey, R. D. Frostig, and B. J. Tromberg, "Imaging cortical absorption, scattering, and hemodynamic response during ischemic stroke using spatially modulated near-infrared illumination.," *Journal of Biomedical Optics* **14**(2), p. 024033, 2010.

[23] S. L. Jacques, "Optical properties of biological tissues: a review," *Physics in Medicine and Biology* **58**, pp. 5007–5008, July 2013.

[24] M. Firbank, M. Hiraoka, M. Essenpreis, and D. T. Delpy, "Measurement of the optical properties of the skull in the wavelength range 650-950 nm.," *Physics in Medicine & Biology* **38**, pp. 503–10, Apr. 1993.

[25] Q. Fang, "Mesh-based Monte Carlo method using fast ray-tracing in Plücker coordinates," *Biomedical Optics Express* **1**, pp. 165–75, Aug. 2010.

[26] Q. Fang and D. a. Boas, "Monte Carlo simulation of photon migration in 3D turbid media accelerated by graphics processing units," *Optics express* **17**, pp. 20178–90, Oct. 2009.