

Learnable real-time inference of molecular composition from diffuse spectroscopy of brain tissue

Ivan Ezhov,^{a,*}† Kevin Scibilia^{a,†} Luca Giannoni^{b,c} Florian Kofler^d Ivan Iliash,^a
Felix Hsieh^a Suprosanna Shit,^a Charly Caredda^e Frédéric Lange^f
Bruno Montcel^e Ilias Tachtsidis^f and Daniel Rueckert^{a,g}

^aTechnical University of Munich, Department of Computer Science, Munich, Germany

^bUniversity of Florence, Department of Physics and Astronomy, Florence, Italy

^cEuropean Laboratory for Non-Linear Spectroscopy, Florence, Italy

^dHelmholtz AI, Helmholtz Munich, Neuherberg, Germany

^eUniv Lyon, INSA-Lyon, Université Claude Bernard Lyon 1, UJM-Saint Etienne, CNRS, Inserm, CREATIS UMR, Lyon, France

^fUniversity College London, Department of Medical Physics and Biomedical Engineering, London, United Kingdom

^gImperial College London, Department of Computing, London, United Kingdom

ABSTRACT. **Significance:** Diffuse optical modalities such as broadband near-infrared spectroscopy (bNIRS) and hyperspectral imaging (HSI) represent a promising alternative for low-cost, non-invasive, and fast monitoring of living tissue. Particularly, the possibility of extracting the molecular composition of the tissue from the optical spectra deems the spectroscopy techniques as a unique diagnostic tool.

Aim: No established method exists to streamline the inference of the biochemical composition from the optical spectrum for real-time applications such as surgical monitoring. We analyze a machine learning technique for inference of changes in the molecular composition of brain tissue.

Approach: We propose modifications to the existing learnable methodology based on the Beer–Lambert law. We evaluate the method’s applicability to linear and non-linear formulations of this physical law. The approach is tested on data obtained from the bNIRS- and HSI-based monitoring of brain tissue.

Results: The results demonstrate that the proposed method enables real-time molecular composition inference while maintaining the accuracy of traditional methods. Preliminary findings show that Beer–Lambert law-based spectral unmixing allows contrasting brain anatomy semantics such as the vessel tree and tumor area.

Conclusion: We present a data-driven technique for inferring molecular composition change from diffuse spectroscopy of brain tissue, potentially enabling intra-operative monitoring.

© The Authors. Published by SPIE under a Creative Commons Attribution 4.0 International License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.JBO.29.9.093509](https://doi.org/10.1117/1.JBO.29.9.093509)]

Keywords: broadband near-infrared spectroscopy; hyperspectral; machine learning; brain imaging; Beer–Lambert law; spectral unmixing

Paper 240152SSRR received May 29, 2024; revised Sep. 4, 2024; accepted Sep. 5, 2024; published Sep. 24, 2024.

*Address all correspondence to Ivan Ezhov, ivan.ezhov@tum.de

†These authors contributed equally to this work.

1 Introduction

Various biomedical applications such as histopathology or neurosurgery require access to rapid monitoring of intrinsic tissue properties. In particular, neuronavigation would benefit by having structural and functional information on brain tissue in real time.^{1–3} Spatially resolved maps of the tissue characteristics would allow bypassing invasive disease diagnostics, e.g., biopsy, which halts the operation. Instead, a surgical decision could be made during the operation, reducing its time and preserving a healthy brain.

Diffuse optical modalities, such as near-infrared spectroscopy and hyperspectral imaging (HSI), emerge as promising technologies to address these clinical needs.^{4,5} These techniques can probe biological matter utilizing non-ionizing electromagnetic radiation within the visible and near-infrared ranges. The spectral instrumentation is inexpensive (compared with other modalities such as MRI), allows for continuous tissue monitoring, and can be easily adapted for the complex context of the operating room.⁶

The overarching principle behind spectroscopy-based molecular characterization is to relate the reflection spectrum obtained upon illumination of the tissue surface with its optical properties. The molecules constituting the tissue have unique absorption dependency on the optical radiation frequency, and thus, the reflection should exhibit molecular absorption signatures in its frequency dependency.

However, several other physical factors contribute to shaping the measured reflectance spectra. These can include light scattering on the surface and within tissue volume, autofluorescence, tissue inhomogeneity, and background illumination.⁷ Disentangling these phenomena from a reflectance spectrum is often an ambiguous, ill-posed problem, yet it is crucial for deducing the relation between the reflection and molecular composition. Another complication is a scarcity of available studies in which optical monitoring and quantitative biochemical composition analysis are performed simultaneously.

Analytical and statistical approaches exist to unmix optical spectra into the physical phenomena defining the spectra profile under a limited data regime.^{8–14} A large number of methods mitigate the data scarcity by introducing a physical prior to establishing the spectrochemical link. Typically, the modified Beer–Lambert law^{15,16} is used to provide such a link by describing the incoming light's energy dissipation as an exponentially decaying function:

$$\log[I_R(\lambda)/I_0(\lambda)] = -\left[\sum_i c_i \mu_a^i(\lambda) + s\mu_s(\lambda)\right]l + U. \quad (1)$$

Here, $I_R(\lambda)$ and $I_0(\lambda)$ are the intensities of the reflected and the incoming light; μ_a and μ_s are the absorption and scattering coefficients; the index i denotes the molecule constituting the tissue such as water, fat, hemoglobin, and cytochromes; c_i denotes the corresponding concentration (e.g., as volume fraction); and s is the weight of scattering in the total light energy dissipation. The remaining quantities are λ , which is the light wavelength, l is the light pathlength (even though several works demonstrate the importance of wavelength-dependent definition of the pathlength,¹⁷ in what follows, for simplicity, we opted for constant pathlength independent of the wavelength), and U describes other physical factors contributing to the energy dissipation of the incoming light or other sources of the optical signal captured by the light detectors or cameras.

Typical molecules whose changes in concentration are inferred include oxyhemoglobin, deoxyhemoglobin, and cytochrome-c-oxidase (CCO).⁵ Measurement of the former two chromophores can reveal, e.g., the oxygenation status of the brain, and can help determine hypoxic or hyperoxic conditions. CCO is a fundamental metabolic molecule correlated to ATP production during cellular respiration, which has previously complemented obtained hemodynamic information in various applications.^{5,18}

Now, in the case of changes in molecular composition over the course of optical monitoring, one can assume that the effects contributing to U either stay constant (e.g., which is a fair assumption for ambient illumination) or change notably less than the total absorption (in certain scenarios, even changes in scattering are not expected since it is a rather bulk effect dependent on the density of the probed matter rather than a molecule-specific one). Under this assumption, the subtraction of two reflection spectra, $\log I_R^2 - \log I_R^1$ (traditionally, these have been two different

points in time, but this could also be true for different points in space), would cancel out or make negligible the term $\delta U = U^2 - U^1$ in the following equation:

$$\log[I_R^2(\lambda)/I_R^1(\lambda)] = - \left[\sum_i \delta c_i \mu_a^i(\lambda) + \delta(s\mu_s(\lambda)) \right] l + \delta U. \quad (2)$$

In such a differential form, the modified Beer–Lambert law can now be used to identify molecular composition. For this, standard least-square optimization algorithms (or non-negative matrix factorization^{10,19}) can be employed to minimize the difference between the real spectra and the spectra obtained from the modified Beer–Lambert law. As a result of the minimization, the optimal values of the set of concentration changes $\{\delta c_i = c_i^2 - c_i^1\}$ are obtained (alongside the scattering parameters).

The overarching drawback of this approach is the computational time it takes to infer the biochemical composition. For example, the optimization methods take a subsecond time to infer the composition of a single spectrum containing a number of wavelengths typical for broadband near-infrared spectroscopy (bNIRS) and HSI (a few hundred). However, for real-time applications particularly in the case of HSI modality, one needs to solve the optimization task in a subsecond time for as many spectra as there are spatial pixels, as every pixel contains its own spectrum. The number of pixels on a hyperspectral image can be easily in the order of to $10^5 - 10^6$. Providing subsecond timings for simultaneous inference on such an number of spectra poses a challenge for traditional methods.

1.1 Contributions

There are numerous studies analyzing the application of machine learning methods to achieve fast inversion of the physical models based on Beer–Lambert law or Monte-Carlo simulations.^{20–26} Predominantly, they imply training a machine learning model on synthetic data generated by following the chosen physical formalism and then evaluating the trained model on real spectra. While proven to work for the use cases mentioned in the cited works, this approach might be inferior as synthetic data generators likely underestimate the complexity of real data. To mitigate this, we tested different strategies for model training using only the “synthetic data” as in previous works or incorporating “real data” via traditional optimization in the training procedure.

Second, we test the proposed method on its ability to approximate physical models of varying complexity: “linear” (absorption only) and “nonlinear” (absorption combined with scattering). While evaluated independently in previous works, here, we also analyze our approach to explicitly compare both models in terms of spectral fit. This comparison is motivated by a desire to elucidate the conditions under which the linear model (that can be easily solved using, e.g., pseudoinverse) is appropriate for describing the light-brain matter interaction process and where inclusion of the scattering is necessary.

Third, given that the work is carried out within the HyperProbe project²⁷ aiming to achieve real-time brain tissue monitoring, the present paper evaluates the “computational timing” for the biochemical composition inference across different methods and hardware platforms. To reinforce the comparison, in contrast to previous works manually choosing the hyperparameters’ values of the machine learning methods, we used the AutoML technique²⁸ to identify the most optimal hyperparameters set.

Finally, to our knowledge, this is the first work that applies a neural-network-based approach to provide real-time inference of chromophore composition from *in vivo* “brain tissue” spectroscopy measurements.^{29–37} We evaluate and discuss the applicability of the method on broadband NIRS (transmission mode)³⁸ and hyperspectral (reflection mode)³⁹ measurements of brain tissue.

2 Method

As mentioned in the introduction, inference of absolute chromophore concentrations from an optical spectrum is a challenging task due to multiple physical effects shaping the reflection spectrum. Thus, we instead aim to predict the changes in concentrations from changes in the spectra [Eq. (2)].

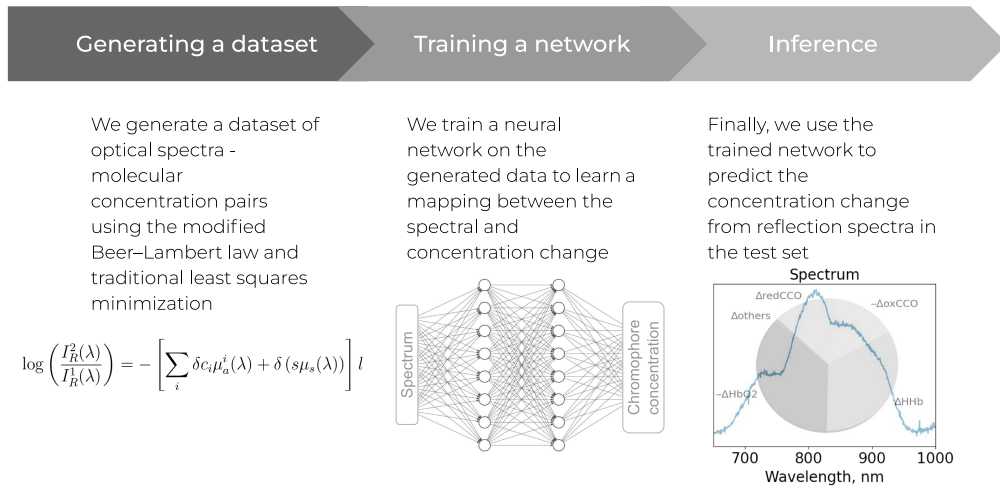


Fig. 1 General pipeline describing the learnable approach for inferring concentrations' changes of molecules such as reduced and oxidized CCO, oxy- and deoxy-hemoglobin $\{\delta c_i\} = \{\delta redCCO, \delta oxCCO, \delta HHb, \delta HbO_2, etc.\}$. The pipeline involves training on a dataset that is generated by the means of a modified Beer-Lambert model. According to the model, the light reflection I_R is shaped by the absorption μ_a and scattering μ_s phenomena.

2.1 Dataset Creation

In our method, Fig. 1, we use a supervised data-driven approach by creating a dataset of attenuation-concentration pairs to train a neural network (by attenuation, or more precisely the change of it, we imply the logarithm of the reflection: $\Delta A = \log I_R^2 - \log I_R^1$). We employ two different strategies to create the dataset.

- (a) The first strategy directly utilizes the modified Beer-Lambert law to generate the training dataset pairs $(\Delta A, \{\delta c_i\})$ with ΔA being the difference in attenuation between two spectra and $\{\delta c_i\}$ the corresponding differences in concentration of chromophores [Fig. 2(a)]. For each chromophore, we randomly sample values for changes in the molecular composition $\{\delta c_i\}$ using the uniform distribution within physiologically plausible ranges. These ranges were determined based on values typically used in the literature, and further details are provided in the [Appendix](#).

If scattering is included in Eq. (2), we may assume it to be of rational form

$$s\mu_s(\lambda) = s \left(\frac{\lambda}{500 \text{ nm}} \right)^{-b}, \quad (3)$$

with the scaling of the anisotropy $g = 0.9$ included in $s = s'/(1 - g)$.^{5,40} In the differential form, we obtain

$$\delta[s\mu_s(\lambda)] = s_2 \left(\frac{\lambda}{500 \text{ nm}} \right)^{-b_2} - s_1 \left(\frac{\lambda}{500 \text{ nm}} \right)^{-b_1}, \quad (4)$$

and therefore also uniformly sample parameters s_1, s_2, b_1, b_2 within plausible ranges.⁴⁰ Subsequently, we input the obtained values into the modified Beer-Lambert law to obtain synthetic differential attenuations $\Delta A(\lambda)$. This difference in attenuation ΔA as input and the corresponding $\{\delta c_i\}$ as output are then used for training.

- (b) Given that the distribution of the synthetic spectra obtained according to the strategy described above can be notably different from the distribution of real spectra, this can result in an unsatisfactory network prediction accuracy. Therefore, in addition to (a), we evaluate another strategy for creating a dataset trying to bridge the gap between the physical model and real data [Fig. 2(b)]. For this, we use traditional least square minimization to fit the changes in the real reflectance spectrum with the modified Beer-Lambert law. The concentrations $\{\delta c_i\}$ found upon the optimization and the corresponding ΔA constitute the training samples.

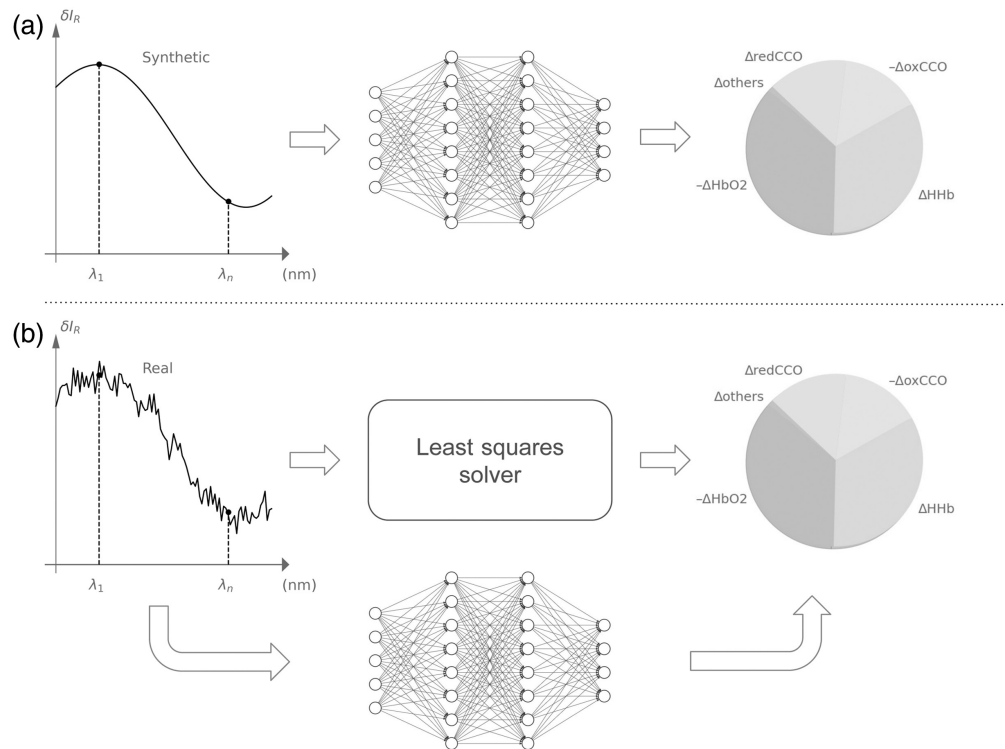


Fig. 2 Two strategies for collecting the training dataset. Strategy (a) in which we train a network on synthetic attenuation-concentration pairs generated from the modified Beer-Lambert law. Strategy (b) in which the training is performed on pairs of real attenuation and concentrations obtained through the least-squares fit to the corresponding real spectra.

2.2 Network and Optimization Details

The training was performed using a multi-layer perceptron (MLP)⁴¹ neural network for both approaches. The network takes as input a one-dimensional vector of attenuation difference and outputs molecular concentration changes.

We trained both networks with early stopping when they reached convergence. To find the optimal network architecture, we used the Ray Tune library²⁸ to validate different MLP architectures (width, number of hidden layers, and activation functions), learning rates, and batch sizes. More details regarding the networks and the training procedure can be found in the [Appendix](#).

Least squares optimization for the Beer-Lambert law, excluding the scattering effect, can be performed via multiplication of the observed attenuation with the pseudoinverse of the absorption coefficients.⁵ To perform the nonlinear least-squared optimization for the Beer-Lambert law model including scattering, we used the publicly available solver of the SciPy library.⁴² We used the least-square minimization obtained predictions as the ground truth (GT) to validate all the trained networks.

2.3 Data

For our experiments, we applied two types of Beer-Lambert law formulation, with and without scattering, to two types of spectral datasets: broadband NIRS data for which the spectra were measured in light transmission mode³⁸ and hyperspectral data, which were obtained in non-contact reflection mode.³⁹

2.4 Broadband NIRS

The first dataset is composed of broadband NIRS spectra from a study analyzing 27 piglets' brains in which a hypoxia-ischemia (HI) state was induced.³⁸ The piglets were monitored for several hours, during which the carotid arteries were surgically isolated, and a stepwise hypoxia took place for 15 to 20 min. This produced a significant hypoxic-ischaemic effect that changed

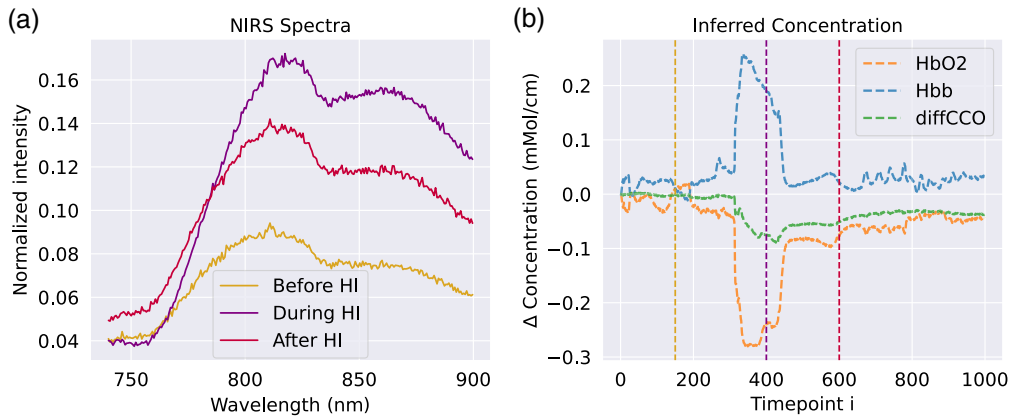


Fig. 3 (a) Optical spectra from the broadband bNIRS study in Ref. 38 before, during, and after inducing HI in the piglet's brain. (b) Predictions of the molecular concentration change $\{\delta c_{\text{HbO}_2}, \delta c_{\text{Hb}}, \delta c_{\text{diffCCO}}\}$ over the course of the optical monitoring (for the Beer–Lambert model *without* scattering). The vertical lines denote the time points corresponding to the normalized reflection spectra on the left.

the metabolic status of the brain and, in some instances, caused further brain injury. The details of the intervention protocol are described in Ref. 38. The optical device used in the study utilizes a miniature light source and a customized high-throughput miniature spectrometer, connected to high numerical aperture optical fibers. The measurements contain around eight thousand spectra per piglet. The distance between each measured time point is between 10.0 and 10.5 s. We use the first thousand measurements, i.e., we only consider the first ~ 2.5 h of measurement. For all piglets, this is sufficient to observe HI and recovery after HI. As Eq. (2) requires defining a baseline spectrum, analogously to Ref. 38, we used a spectrum at the very beginning of optical monitoring (i.e., before HI) for the baseline. We normalized the spectra with respect to dark noise and white reference. The normalized bNIRS spectra before and after the intervention-inducing hypoxia are shown in Fig. 3(a), and predictions of the concentrations change over the course of the optical monitoring are shown in Fig. 3(b). Out of the 27 piglets in the dataset, 25 had data available during HI, such that 19 were used for training, two for validation, and four for testing. For this dataset, we predict three types of molecules: oxyhemoglobin, deoxyhemoglobin, and differential CCO (as the total CCO concentration may be assumed to not change within a few hours, the oxidized-reduced difference spectrum may be used to infer changes of both oxidized and reduced CCO¹⁸), i.e., $\{\delta c_i\} = \{\delta c_{\text{HbO}_2}, \delta c_{\text{Hb}}, \delta c_{\text{diffCCO}}\}$, where $\delta c_{\text{diffCCO}} = \delta c_{\text{oxCCO}} - \delta c_{\text{redCCO}}$. We neglected the potential contribution to the spectra from water and fat due to their minimal change in concentrations during the 2.5 h of monitoring.^{38,43} Note that we assume unitary pathlength in our experiments, which results in units of (mM cm^{-1}) and (cm^{-1}) for the inferred concentrations. In the NIR range, it has been shown that the pathlength is semi-constant,³⁸ which effectively leads to a simple rescaling in our concentrations when using this assumption. This can also be observed when comparing the inferred exemplary concentrations from Fig. 3(b) with pathlength-corrected concentrations in previous work.³⁸

2.5 Hyperspectral Data

The second dataset we used consists of hyperspectral data from the HELICoiD project.³⁹ The HELICoiD dataset comprises brain HSI images obtained in surgical conditions from 22 patients diagnosed with glioma. The optical instrumentation is based on the pushbroom technique and a silicon charge-coupled device (CCD) detector array as a camera. The HSI images provide a high spectral resolution of 826 bands spread between 400 and 1000 nm and a 2D spatial resolution of a few hundred pixels in each dimension, corresponding to the maximum size of $129 \text{ mm} \times 230 \text{ mm}$.³⁹ The images were also expert-annotated into three tissue classes: normal and tumor tissues, as well as blood vessels.

A typical hyperspectral image and corresponding spectra are shown in Fig. 4. As one can see, there is a notable presence of noise at the end of the measured spectra that apparently comes from the decreasing signal-to-noise ratio of the instrumentation in these ranges. To circumvent

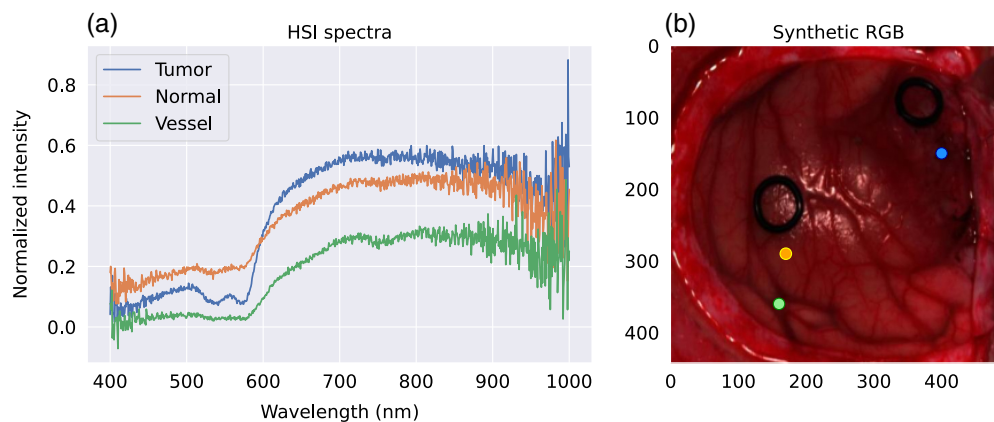


Fig. 4 (a) Optical spectra from the HSI study of patients diagnosed with glioma³⁹ for different tissue types: tumor, normal tissue, and blood vessels. (b) A typical RGB image of the brain surface, which is obtained from the HSI volume. The dots correspond to the spectra on the left image. The black circles on the RGB image are rubber rings that surgeons use to mark tumors and healthy tissues.

this problem, we used the signal from the 530 to 750 nm range, where both absorption properties of chromophores were known and a high signal-to-noise ratio (SNR) of the instrumentation was expected. Different from the bNIRS dataset, we used a spectrum of the pixel belonging to the blood vessel class as a baseline spectrum. The blood vessel was used as a reference for a couple of reasons. The blood vessel is clearly distinguishable from the other two tissue types, tumor and non-tumor tissue. These two types of tissue are highly heterogeneous, e.g., within a pixel area, they can have small capillaries, leakage of blood, agglomeration of dead cells, etc. By contrast, the blood vessel pixels, especially the ones belonging to large arteries, are less heterogeneous. Moreover, it is assumed that blood vessels do not possess cytochrome molecules. Thus, it is a better reference when one sets a goal of detecting the presence of cytochromes in the brain matter. We then subtracted the baseline spectrum from all other spectra in the same image. In other words, we performed the differential spectroscopy *not in time but in space*.

Besides predicting oxyhemoglobin δc_{HbO_2} and deoxyhemoglobin δc_{Hb} , we again infer the differential CCO concentration due to its role in capturing oxidative metabolic activity. We separately predicted oxidized CCO and reduced CCO, as the total CCO concentration may not be assumed to remain constant in space. We also predict water and fat since, for these molecules, one cannot assume minimal concentration change across different tissue types as in the case of the bNIRS spectra. For reference, the absorption spectra can be found in the [Appendix](#).

Figure 5 showcases examples of molecular inference for the HSI images from the HELICoiD dataset. Out of the nine patients with glioblastoma in the dataset, six with distinct class labeling were chosen, and three patients were used for training, one for validation, and two for testing. Note that patients might have multiple images taken, and different images from the same patient were assigned to the same training, validation, and test set to avoid set contamination. Therefore, the training set consists of five, the validation set of one, and the test set of three images.

3 Results

3.1 Scattering Versus Non-scattering

First, before discussing the learnable methods for molecular inference, we test different Beer-Lambert law formulations—with and without scattering—to elucidate the limits of applicability of both models. For the case of piglets undergoing HI, it is widely assumed that the 780 to 900 nm range is predominantly dominated by absorption, with scattering being only a minor contributor to the overall measured spectrum.³⁸ As measurements in the piglet dataset below the 780 nm threshold were available, we opted to extend the model fitting range from 740 to 900 nm. This test is motivated by our desire to assess whether a linear model (without scattering) would still be sufficient to describe the broader spectroscopy measurement of brain tissue.

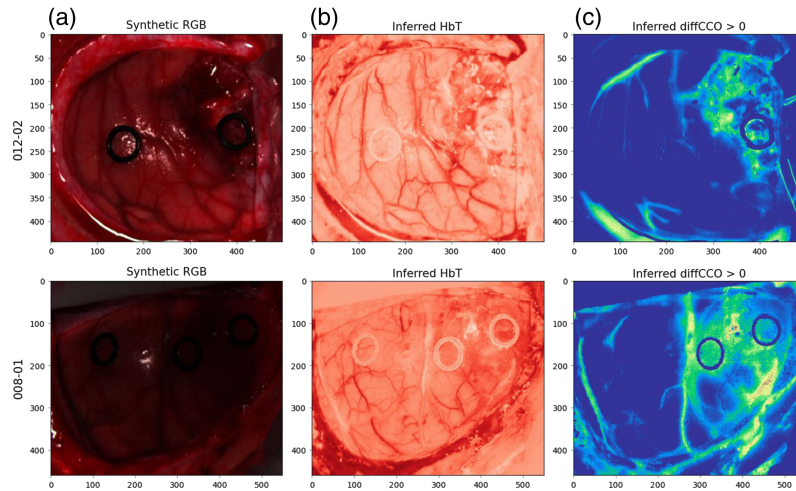


Fig. 5 (a) Examples of HSI images of two patients, shown with respective patient ID, from the HELiCoID dataset.³⁹ Each pixel in the shown 2D image possesses a spectral signature with 826 bands. From this signature, we predict the molecular concentration change for (b) hemodynamic $\delta C_{\text{HbT}} = \delta C_{\text{HbO}_2} + \delta C_{\text{Hb}}$ and (c) metabolic $\delta C_{\text{diffCCO}} = \delta C_{\text{oxCCO}} - \delta C_{\text{redCCO}}$ characterization. Here, we use the Beer–Lambert model *with* scattering, as it provides a closer fit to real spectra than the model without scattering. We observe that performing the spectral unmixing on the HSI measurement of brain tissue allows us to better contrast the vessel tree (b) and tumor area (c) than on the RGB image.

Figures 6 and 7 show spectral fits and inferred molecular concentrations using both formulations. The model with scattering provides a clearly better fit. It allows us to better describe the peak around 760 nm for the bNIRS data, whereas for the HSI data, the inclusion of scattering is often merely necessary for an accurate fit of the spectra in this wavelength range. This finding is consistent across the dataset, as shown in Table 1. Table 1 also reveals the fact that the HELiCoID data are notably noisier than the bNIRS data, which explains the models' worse fitting.

To show that the model with scattering can significantly improve model fits, especially for the higher frequency portion of the spectrum, we evaluate the relative improvement r in terms of the spectral fit of the scattering model compared with the non-scattering model. We use mean absolute error (MAE) as a measure of the fit and compute it for all piglets across different spectrum bands.

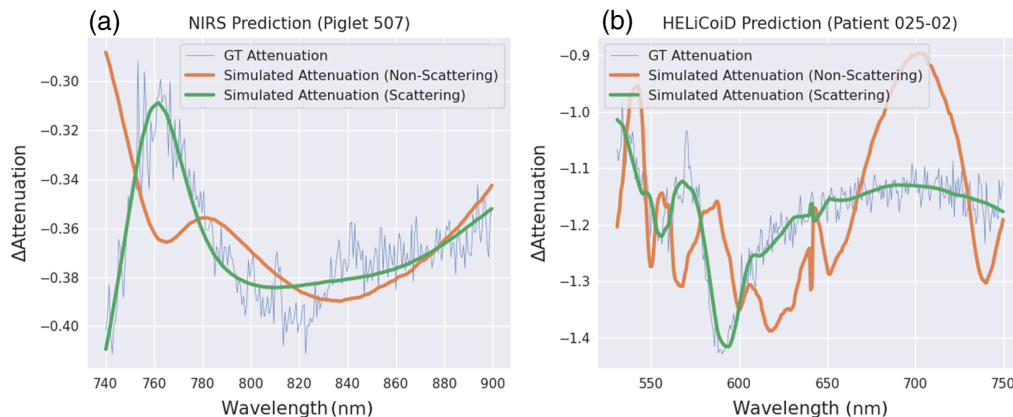


Fig. 6 Comparison between predictions using linear (no scattering) and nonlinear (with scattering) models for (a) bNIRS and (b) HSI spectra. The GT attenuation is computed from the real spectra difference. The inclusion of scattering into the formulation of the Beer–Lambert law notably improves the spectral fit to real data.

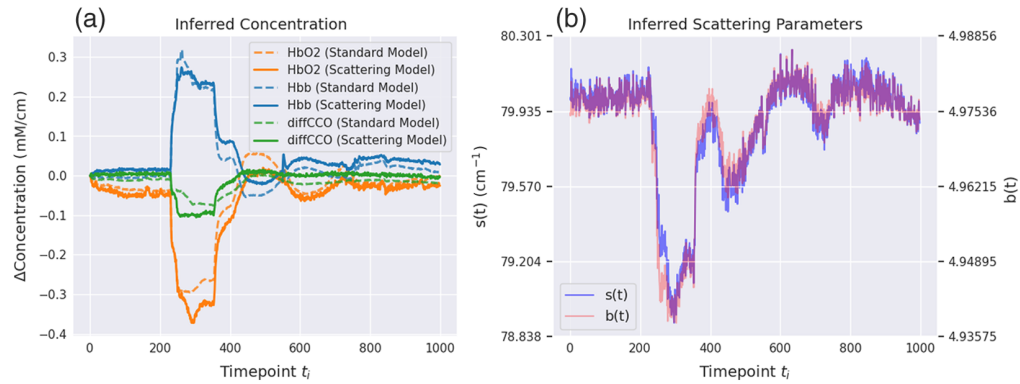


Fig. 7 (a) Comparison between chromophore predictions using the linear model without scattering (“standard model”) and nonlinear model with scattering (“scattering model”) for the bNIRS dataset. (b) The predicted coefficients s and b for the scattering term in Eq. (3). Note that the purple curve represents an overlap across the two red and blue curves.

Table 1 Quantitative performance comparison of the different Beer–Lambert models and network training strategies on the test set of the two spectral datasets.

Dataset	Individual ID	Spectral MAE		Concentration MAE	
		Non-scattering	Scattering	Strategy (a)	Strategy (b)
Broadband NIRS	507	1.23×10^{-2}	8.60×10^{-3}	1.36×10^{-2}	4.32×10^{-3}
	509	1.09×10^{-2}	9.37×10^{-3}	2.46×10^{-2}	4.81×10^{-3}
	511	8.01×10^{-3}	6.60×10^{-3}	1.16×10^{-2}	3.41×10^{-3}
	512	1.20×10^{-2}	1.09×10^{-2}	1.25×10^{-2}	4.59×10^{-3}
HELICoID	012-01	3.27×10^{-2}	2.49×10^{-2}	1.73×10^{-1}	1.64×10^{-2}
	012-02	2.24×10^{-2}	2.19×10^{-2}	1.50×10^{-1}	2.58×10^{-2}
	015-01	6.33×10^{-2}	2.54×10^{-2}	1.75×10^{-1}	1.54×10^{-2}

To compare the two (non-scattering and scattering) models, we compute the mean absolute error of the spectral fit (denoted as “Spectral MAE”) between the GT observed and predicted signals. The two network training strategies are compared by assessing the MAE of each strategy between the network and optimization-inferred concentrations (denoted as “Concentration MAE”) of all considered chromophores. In the case of the HELICoID dataset, only pixels labeled as normal, tumor, or blood were considered for these computations. The best-performing model and strategy for each individual is highlighted in bold.

The spectral fit MAE is calculated by

$$\text{MAE} = \frac{\sum_{i=1}^n |\Delta A_{\text{model}}(\lambda_i) - \Delta A_{\text{data}}(\lambda_i)|}{n}, \quad (5)$$

where $\Delta A_{\text{model}}(\lambda_i)$ represents the model-inferred attenuation at wavelength λ_i and ΔA_{data} the real measured attenuation, respectively. The relative mean improvement is computed by

$$r = \frac{\text{MAE}_{\text{scatter}} - \text{MAE}_{\text{linear}}}{\text{MAE}_{\text{linear}}},$$

comparing the improvement in mean error between the scattering and linear models, computed across all timepoints.

The results of such computation for all piglets are shown in Fig. 8, where we observed a mean relative improvement of 15.7% over the full fitting range. Such improvement is especially noticeable in the 740 to 780 nm range, where the mean improvement of the distribution almost doubled at 30.8%. The spectral fit does not improve significantly in the 780 to 900 nm range, with the mean relative improvement of the distribution being at merely 5.6%. For one of the

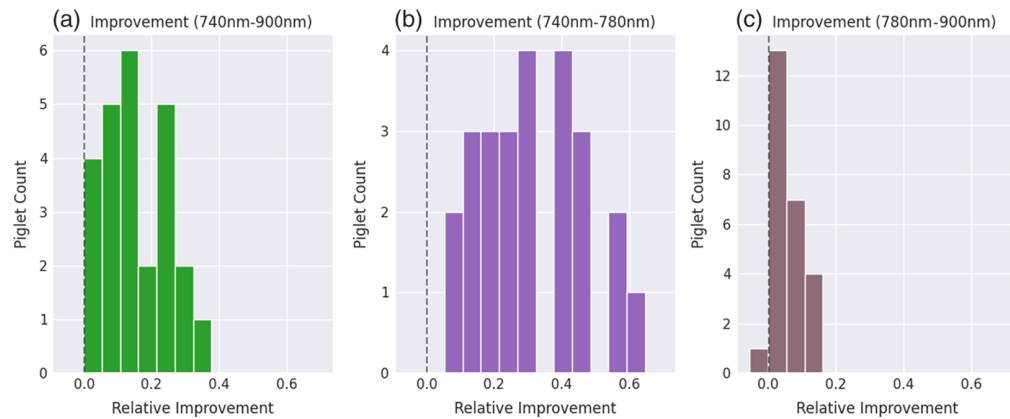


Fig. 8 Histograms showing mean relative improvement of the spectral fit MAE for the presented scattering model, compared with the linear model, across all piglets in the dataset for (a) all wavelengths, (b) in the range 740 to 780 nm, and (c) in the range 780 to 900 nm in the broadband NIRS dataset. The x-axis represents the relative improvement r between the two models, and the y-axis shows the number of piglets that achieved the corresponding mean relative improvement r . The dashed line signifies improvements below zero, i.e., cases where the spectral fit worsened.

25 piglets, we observed that the spectral fit slightly worsened in the 780 to 900 nm range through the nonlinear model. However, the fitting MAE is only worse by 0.6%, and the spectral fit was still better for the overall range and in the 740 to 780 nm range. We therefore can still confidently conclude that the presented model is able to fit the piglets' measured spectra more closely, especially for presumed scattering-dominated bands.

The necessity of the scattering consideration in the Beer–Lambert model for the HSI data can be explained by the more pronounced contribution of the scattering process. For the HSI data, we infer the difference in molecular composition between different spatial locations on an image, i.e., between different tissue types. The scattering property across brain tissues can significantly vary, and thus, the scattering shapes markedly the differential spectra. By contrast, for the bNIRS data, we perform the differential spectroscopy analysis not in space but in time (comparing two spectra for the same location taken at different time points), meaning that the molecular inference is performed for the same tissue type.

In conclusion, we find that the nonlinear model is especially helpful in describing scattering-dominated bands. However, the linear model may still be used when absorption is the prevalent physical effect.

3.2 Evaluating Different Training Strategies

Next, we evaluate the proposed machine learning approach in its ability to substitute both the linear absorption and the nonlinear scattering model.

Figure 9 demonstrates the results of the experiment in which we test the network trained on synthetic data collected according to strategy (a) and on real data according to strategy (b), for both linear and nonlinear models. For the linear case, both strategies are able to correctly infer the concentrations. The solution to the linear model can merely be found by a matrix multiplication, i.e., the pseudoinverse, which is why both strategies are able to very accurately predict the optimization-inferred concentrations.

For the nonlinear case, strategy (b) provides qualitatively closer fits. We also tested this model for the HELICoiD dataset, where we found highly matching results by the use of strategy (b), as seen in Fig. 10.

These findings are also quantitatively supported by the results in Table 1.

3.3 Computational Time

Importantly, the proposed network-based optimization comes with a significant speed-up in computational time. In Fig. 11, we show a comparative analysis for performing chromophore composition inference using standard least-square solvers (based on gradient update or

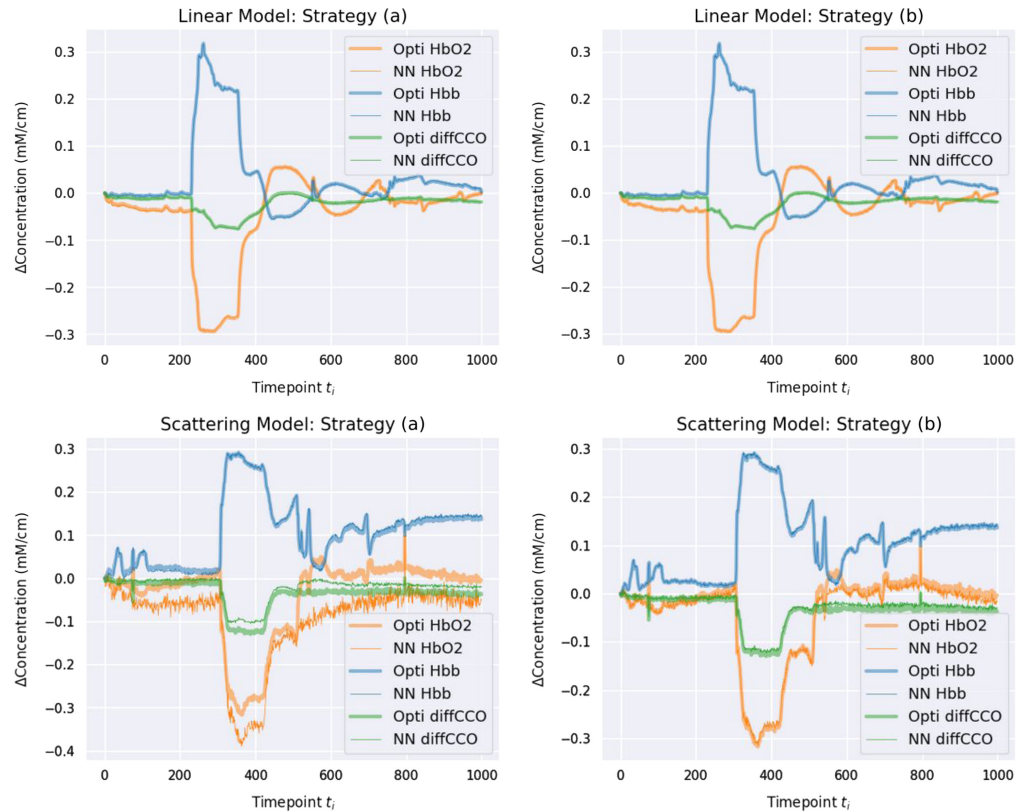


Fig. 9 Comparison between inference of the molecular composition using the standard optimization methods and proposed network-based inference for training strategies (a) and (b) on the bNIRS dataset. The top row compares both strategies when using the linear model, where highly accurate neural network predictions are visible in both cases. The bottom row compares the strategies when using the nonlinear model, with strategy (b) delivering noticeably more accurate predictions.

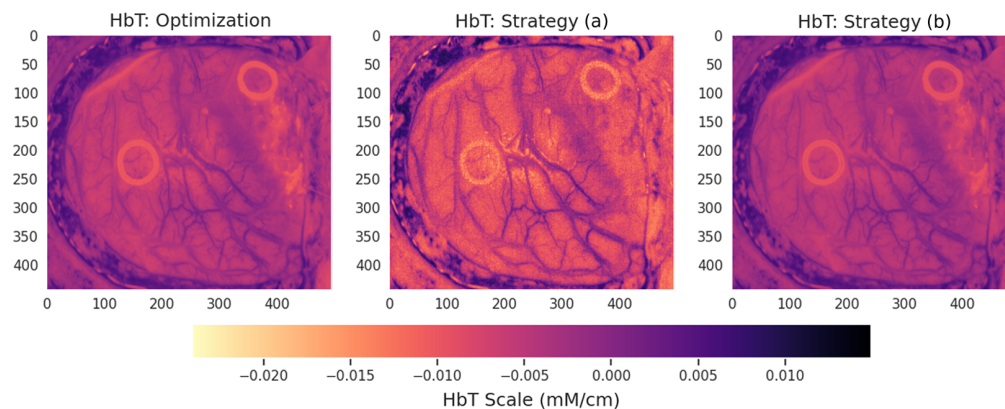


Fig. 10 Comparison between inference of the molecular composition using the standard optimization methods (Opti) and proposed network-based inference (NN) for training strategies (a) and (b) on the HELICoiD dataset. The top row compares strategy (a), left, with strategy (b), right, for the bNIRS dataset. The figure compares the inference of the hemodynamic signal of the optimization-based result, left, with strategy (a), middle, and strategy (b), right. Strategy (b) significantly improves results when using the nonlinear model.

pseudoinverse) and our proposed approach. The used spectra for this comparative analysis are taken from the broadband NIRS dataset assuming the scattering model, i.e., they are in the 740 to 900 nm range, with a total of 244 measured wavelengths per spectrum, and the underlying chromophores are oxyhemoglobin, deoxyhemoglobin, and differential CCO.

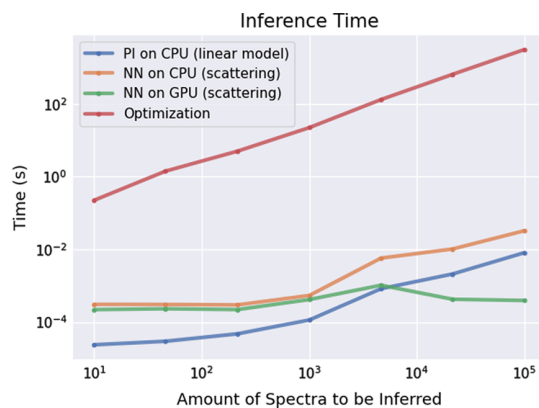


Fig. 11 Comparison between inference time for various optimization approaches for varying number of spectra (from 10 to 10^5): including the pseudoinverse for the linear model (blue) and optimization-based (red) for the nonlinear scattering model (both running on CPU), as well as network-based approach for scattering model running on CPU (orange) and GPU (green).

As solving the linear system using the pseudoinverse requires the least amount of matrix multiplication, this method provides the fastest computation. However, with the growth of the number of spectra for which we solve the optimization task, the matrix size for the inversion increases, and thus, the computational time increases. Starting from ca. 10^4 number of spectra, the proposed network having a fixed number of computational units becomes superior in terms of optimization time. Such runtime will remain approximately constant with a further increasing number of spectra, assuming sufficient GPU memory is available. More importantly, for non-linear systems, which are here represented as a Beer–Lambert model with the inclusion of scattering, one cannot utilize the pseudoinverse and has to resort to nonlinear solvers such as the ones based on gradient update. Such solvers are two to three orders of magnitude slower than the neural network approach, which has fixed compute time for linear and nonlinear systems. As Fig. 11 shows, it takes ca. 0.4 ms for the network to infer biochemical composition for 10^5 spectra on NVIDIA GeForce MX450 with 2048 MiB. Overall, on our hardware, it takes between 2.5 and 3.1 s to run the neural network for one image from the HELICoiD dataset (the largest among all tested data), from opening the normalized HSI image and loading the neural network into the GPU, to displaying the inferred concentrations.

4 Discussion and Conclusion

First, in this paper, we wanted to address the limitation of the existing machine learning approaches to infer molecular composition using a physical model. Predominantly, the training within these approaches is performed on synthetic data produced by the physical model, analogously to the strategy (a) above. However, real spectroscopy measurements include a few factors not considered by the modeling, such as instrumentation and physiological noise or other nonlinear optical phenomena. To close the gap between synthetic and real data, several works propose incorporating various kinds of noise into the synthetic data.²⁶ But the realism of the used noise formulations and their sufficiency to close the gap can still be questioned. Instead, the learning scheme proposed in this paper results in training and testing carried out on the same type of data obtained by processing real spectra with traditional least-squares optimization. Thus, no noise model is required to achieve accurate predictions.

However, we must admit that the predictions by the existing training approach are, nevertheless, close to the GT. Intuitively, this behavior is not within expectations as, again, the spectra of synthetic data on which the network was trained and real data on which it was evaluated notably vary. We attribute such behavior to the fact that the main spectral feature that networks learn to focus on is the global functional shape of the spectra. The local behavior of the spectral function is less informative, and thus, the local perturbations, i.e., random noise, do not affect the network performance significantly. This reasoning is further reinforced when comparing the network performance between bNIRS and HSI data. For the HSI data, the noise has a pronounced

wavelength dependency, in contrast to bNIRS data (see Fig. 4), which results in the network performance (trained according to the existing approach) degrading more notably compared with bNIRS data.

Another explanation for this behavior can be that here, we try to learn with a neural network a solution to a linear system (or quasi-linear system, in the case of the model with scattering, where scattering contribution to the reflection spectra is minor compared with the absorption, as for bNIRS data). We hypothesize that upon network training on the dataset representing solutions to linear systems, the network weights are learned to minimize the contribution from nonlinear network units (e.g., activation functions), as such units are unnecessary to establish a linear mapping. Generally speaking, the nonlinear units are both a strength and a limitation of the networks. The strength is that they enable us to learn an arbitrary mapping, and the limitation is the very sequential application of nonlinear units that can cause the predictions to become highly unstable for input data unseen during training. However, if the contribution from nonlinear units of the network is minimized, we would not expect such instability in the network's predictions.

Second, we want to understand the limits of applicability of the linear model (with no scattering) for describing brain tissue spectroscopy. Note that under the assumption of negligible contribution from scattering in the modified Beer–Lambert law, Eq. (2) results in a linear system. Such a system can be efficiently solved using the pseudoinverse, achieving close to real-time computation. However, in general, when including nonlinear terms in the Beer–Lambert law formulation (scattering, nonlinear absorption effects, etc.), the system of equations cannot be solved anymore using pseudoinverse. Our analysis reveals that including scattering can often be necessary to describe the spectroscopy measurements, especially for HSI data. But also for the bNIRS data, we have qualitatively observed that the nonlinear model seems to fit the wavelength range between 740 and 780 nm particularly well (Fig. 8). This finding was within expectations since the range between 780 and 900 nm is not regarded to be dominated by changes in scattering and has therefore been widely used in broadband NIRS technologies with concentrations inferred through the Beer–Lambert law solely considering absorption.³⁸

Importantly, as a by-product of our analysis, when we test the spectral unmixing on the brain tissue HSI dataset, we observe that the molecular inference driven by the physical model leads to better contrasting the blood vessel tree compared with the RGB image, Fig. 5(b), and capturing metabolic activity (right). Image semantics related to both hemodynamic and metabolic tissue properties could be highly valuable in assisting neurosurgeons during the process of tumor removal. Knowledge of the topology of the blood vessel tree would minimize the amount of undesirable resection of the vessels leading to bleeding. In turn, understanding the metabolic activity across the 2D field of view would allow for better separation of pathological tissue from the healthy parenchyma, reducing the time spent on intra- and postoperative biopsy.

To conclude, in the paper, we present a data-driven concept for inferring molecular composition change from diffuse spectroscopy of brain tissue. We test the approach on various datasets (bNIRS and HSI) and physical models of different complexity (with and without scattering, i.e., linear and nonlinear systems). Importantly, we evaluate different training strategies for neural-network-based molecular prediction. The proposed strategy provides predictions that are nearly identical to the traditional least-square-fit method, making the learnable solver an accurate alternative. Finally, the method achieves subsecond time for simultaneous inference of molecular composition across a large number of spectra, allowing for real-time tissue characterization using bNIRS and HSI imaging modalities.

5 Appendix

5.1 Absorption Spectra of Chromophores

For the broadband NIRS dataset, we used oxyhemoglobin, deoxyhemoglobin, and differential CCO as absorbing chromophores. These have been used extensively in literature in the context of NIRS imaging.³⁸ For HSI brain tumor imaging, there is no standardized chromophore set in the literature that could be used for fitting observed attenuations. We therefore resorted to the assumptions explained in the main text and reported absorption spectra of these fitted chromophores in Fig. 12.

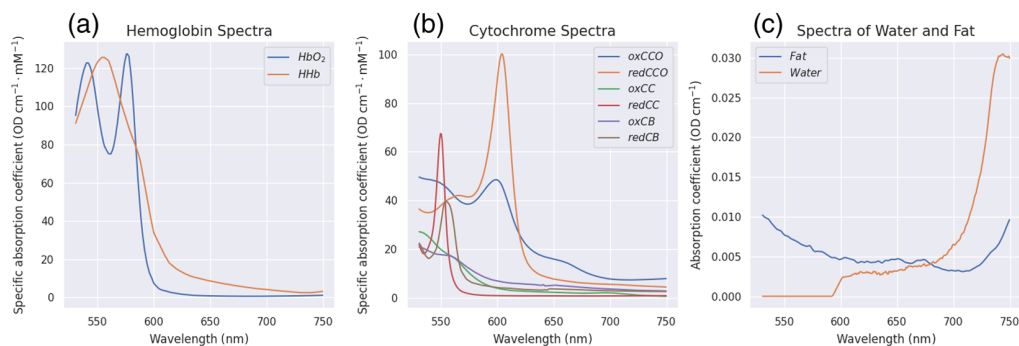


Fig. 12 Absorption coefficients of the fitted chromophores^{44–47} used for the HELICoID dataset. The first plot shows the absorption coefficients of oxyhemoglobin and deoxyhemoglobin. The second plot shows the absorption coefficients of CCO, cytochrome-c, and cytochrome-b in oxidized and reduced form, respectively. Units of these absorption coefficients are per cm per millimole, as they represent concentrations. The third plot shows absorption coefficients of fat and water in the form of volumetric content, with units per cm.

5.2 Dataset Generation Details

To generate the synthetic datasets necessary to train the neural networks with strategy (a), we uniformly sample parameters from certain ranges. For the broadband NIRS dataset, the selected ranges are shown in Table 2. We also report the physiologically used ranges for the scattering parameters a_1 , b_1 , used in both datasets, and the corresponding reference upon which they were based. For the HELICoID dataset, we do not make any direct physiological assumptions as we instead use minima and maxima of the parameters found during optimization.

Table 2 Ranges for model parameters chosen based on physiological assumption, with the corresponding references.

Parameter type	Parameter	Minimum	Maximum	Reference
Concentration (broadband NIRS)	HbO ₂ (mM/cm)	−0.5	0.5	38
	HHb (mM/cm)	−0.5	0.5	38
	diffCCO (mM/cm)	−0.25	0.25	38
Scattering	a_1 (cm ^{−1})	0	100	40
	b_1	0	5	40

5.3 Neural Network Training

All networks were trained with the Adam optimizer. The Ray Tune framework²⁸ automatically finds neural network hyperparameters that would otherwise be difficult to manually tune and find. It selected the optimal number of hidden layers H ; the network width W ; the learning rate λ ; the activation function f ; and the batch size B in 200 trials of random search from the following ranges: $H \in \{0, 1, 2, 3\}$ or $H \in \{0, 1, 2, 3, 4\}$ (depending on whether training was being performed on the broadband NIRS or HELICoID dataset), $W \in [1, 64]$, $\lambda \in [10^{-4}, 10^{-1}]$, $f \in \{\text{ELU}, \text{Hardshrink}, \text{LeakyReLU}\}$, and $B \in \{32, 64, 128, 256\}$ or $B \in \{32, 64, 128, 256, 512, 1024, 2048\}$ (depending on the dataset). Slightly different ranges were used for the HELICoID dataset to account for a possibly larger needed computational complexity, due to the inherently more complex dataset. The found optimal parameters are reported in Table 3.

Table 3 Parameters listed in the text, found by the RayTune library for each dataset (Broadband NIRS and HELICoiD), model (linear absorption and nonlinear scattering), and neural network training strategies (a) and (b), respectively.

Dataset and model	Neural network	Hidden layers	Network width	Learning rate	Activation function	Batch size
Broadband NIRS	Strategy (a)	2	35	5.28832×10^{-3}	Leaky ReLU	128
Linear	Strategy (b)	2	16	1.7495724×10^{-3}	Leaky ReLU	32
Broadband NIRS	Strategy (a)	1	62	7.86158×10^{-4}	Leaky ReLU	32
Nonlinear	Strategy (b)	2	53	5.94006×10^{-4}	ELU	32
HELICoiD	Strategy (a)	1	31	6.64275×10^{-3}	ELU	512
Nonlinear	Strategy (b)	2	46	3.52619×10^{-4}	ELU	32

Disclosures

The authors declare no financial conflicts of interest.

Code and Data Availability

In support of open science, all our code is available on GitHub at <https://github.com/HyperProbe/SpectraFit>. The HSI data are publicly available by the HELICoiD project. The bNIRS data may be obtained upon request. Both datasets received ethical approval as described in the corresponding papers.^{38,39}

Acknowledgments

The HyperProbe consortium has received funding from the European Union's Horizon Europe Research (Grant No. 101071040). UCL is supported by UKRI (Grant No. 10048387).

References

- O. Ganslandt et al., "Neuronavigation: concept, techniques and applications," *Neurol. India* **50**(3), 244–255 (2002).
- D. A. Orringer, A. Golby, and F. Jolesz, "Neuronavigation in the surgical management of brain tumors: current and future trends," *Expert Rev. Med. Devices* **9**(5), 491–500 (2012).
- I. J. Gerard et al., "Brain shift in neuronavigation of brain tumors: an updated review of intra-operative ultrasound applications," *Front. Oncol.* **10**, 618837 (2021).
- A. N. Sen, S. P. Gopinath, and C. S. Robertson, "Clinical application of near-infrared spectroscopy in patients with traumatic brain injury: a review of the progress of the field," *Neurophotonics* **3**(3), 031409 (2016).
- L. Giannoni, F. Lange, and I. Tachtsidis, "Hyperspectral imaging solutions for brain tissue metabolic and hemodynamic monitoring: past, current and future developments," *J. Opt.* **20**(4), 044009 (2018).
- D. Viderman and Y. G. Abdildin, "Near-infrared spectroscopy in neurocritical care: a review of recent updates," *World Neurosurg.* **151**, 23–28 (2021).
- G. Lu and B. Fei, "Medical hyperspectral imaging: a review," *J. Biomed. Opt.* **19**(1), 010901 (2014).
- S. Tzoumas and V. Ntzachristos, "Spectral unmixing techniques for optoacoustic imaging of tissue pathophysiology," *Philos. Trans. R. Soc., A* **375**(2107), 20170262 (2017).
- N. Dobigeon et al., "Linear and nonlinear unmixing in hyperspectral imaging," in *Data Handling in Science and Technology*, C. Ruckebusch, Ed., Vol. **30**, pp. 185–224, Elsevier (2016).
- G. Lu et al., "Estimation of tissue optical parameters with hyperspectral imaging and spectral unmixing," *Proc. SPIE* **9417**, 94170Q (2015).
- A. Bateson and B. Curtiss, "A method for manual endmember selection and spectral unmixing," *Remote Sens. Environ.* **55**(3), 229–243 (1996).
- I. Ezhov et al., "Identifying chromophore fingerprints of brain tumor tissue on hyperspectral imaging using principal component analysis," in *Eur. Conf. Biomed. Opt.*, p. 1262826 (2023).
- K. Yokoyama, M. Watanabe, and E. Okada, "Estimation of the optical path length factor for functional imaging of an exposed cortex by principal component analysis," in *Eur. Conf. Biomed. Opt.*, Optica Publishing Group, p. 5138_168 (2003).

14. A. O. Gerstner et al., “Hyperspectral imaging of mucosal surfaces in patients,” *J. Biophotonics* **5**(3), 255–262 (2012).
15. A. A. Stratonnikov and V. B. Loschenov, “Evaluation of blood oxygen saturation in vivo from diffuse reflectance spectra,” *J. Biomed. Opt.* **6**(4), 457–467 (2001).
16. I. Oshina and J. Spigulis, “Beer–Lambert law for optical tissue diagnostics: current state of the art and the main limitations,” *J. Biomed. Opt.* **26**(10), 100901 (2021).
17. A. Bahl et al., “A comparative study of analytical models of diffuse reflectance in homogeneous biological tissues: gelatin based phantoms and Monte Carlo experiments,” arXiv:2312.12935 (2023).
18. G. Bale, C. E. Elwell, and I. Tachtsidis, “From Jöbsis to the present day: a review of clinical near-infrared spectroscopy measurements of cerebral cytochrome-c-oxidase,” *J. Biomed. Opt.* **21**(9), 091307 (2016).
19. C. Caredda et al., “A priori free spectral unmixing with periodic absorbance changes: application for auto-calibrated intraoperative functional brain mapping,” *Biomed. Opt. Express* **15**(1), 387–412 (2024).
20. Y.-W. Chen and S.-H. Tseng, “Efficient construction of robust artificial neural networks for accurate determination of superficial sample optical properties,” *Biomed. Opt. Express* **6**(3), 747–760 (2015).
21. S. J. Wirkert et al., “Robust near real-time estimation of physiological parameters from megapixel multi-spectral images with inverse Monte Carlo and random forest regression,” *Int. J. Comput. Assist. Radiol. Surg.* **11**, 909–917 (2016).
22. S. Panigrahi and S. Gioux, “Machine learning approach for rapid and accurate estimation of optical properties using spatial frequency domain imaging,” *J. Biomed. Opt.* **24**(7), 071606 (2019).
23. I. Fredriksson, M. Larsson, and T. Strömberg, “Machine learning for direct oxygen saturation and hemoglobin concentration assessment using diffuse reflectance spectroscopy,” *J. Biomed. Opt.* **25**(11), 112905 (2020).
24. B. H. Hokr and J. N. Bixler, “Machine learning estimation of tissue optical properties,” *Sci. Rep.* **11**(1), 6561 (2021).
25. T. Manojlović et al., “Rapid extraction of skin physiological parameters from hyperspectral images using machine learning,” *Appl. Intell.* **53**(13), 16519–16539 (2023).
26. A. Scarbrough, K. Chen, and B. Yu, “Designing a use-error robust machine learning model for quantitative analysis of diffuse reflectance spectra,” *J. Biomed. Opt.* **29**(1), 015001 (2024).
27. L. Giannoni et al., “Hyperprobe consortium: innovate tumour neurosurgery with innovative photonic solutions,” *Proc. SPIE* **12628**, 126281C (2023).
28. R. Liaw et al., “Tune: a research platform for distributed model selection and training,” arXiv:1807.05118 (2018).
29. M. Mamouei et al., “An empirical investigation of deviations from the Beer–Lambert law in optical estimation of lactate,” *Sci. Rep.* **11**(1), 13734 (2021).
30. P.-H. Chou et al., “Deep neural network to differentiate brain activity between patients with first-episode schizophrenia and healthy individuals: a multi-channel near infrared spectroscopy study,” *Front. Psychiatry* **12**, 655292 (2021).
31. X. Y. Yap, K. S. Chia, and N. A. S. Suarin, “Adaptive artificial neural network in near infrared spectroscopy for standard-free calibration transfer,” *Chemom. Intell. Lab. Syst.* **230**, 104674 (2022).
32. N. K. Qureshi et al., “Comparison of classification performance for fNIRS-BCI system,” in *2nd Int. Conf. Robot. and Artif. Intell. (ICRAI)*, IEEE, pp. 54–57 (2016).
33. Y. Altmann et al., “Nonlinear unmixing of hyperspectral images using radial basis functions and orthogonal least squares,” in *IEEE Int. Geosci. and Remote Sens. Symp.*, IEEE, pp. 1151–1154 (2011).
34. Y. Karamavuş and M. Özkan, “Newborn jaundice determination by reflectance spectroscopy using multiple polynomial regression, neural network, and support vector regression,” *Biomed. Signal Process. Control* **51**, 253–263 (2019).
35. H. Bian et al., “Multiple kinds of pesticides detection based on back-propagation neural network analysis of fluorescence spectra,” *IEEE Photonics J.* **12**(2), 1–9 (2020).
36. M. Ewerlöf et al., “Multispectral snapshot imaging of skin microcirculatory hemoglobin oxygen saturation using artificial neural networks trained on in vivo data,” *J. Biomed. Opt.* **27**(3), 036004 (2022).
37. Z. Zhou et al., “Real-time in-situ optical detection of fluid viscosity based on the Beer–Lambert law and machine learning,” *Opt. Express* **30**(23), 41389–41398 (2022).
38. P. Kaynezhad et al., “Quantification of the severity of hypoxic-ischemic brain injury in a neonatal preclinical model using measurements of cytochrome-c-oxidase from a miniature broadband-near-infrared spectroscopy system,” *Neurophotonics* **6**(4), 045009 (2019).
39. H. Fabelo et al., “In-vivo hyperspectral human brain image database for brain cancer detection,” *IEEE Access* **7**, 39098–39116 (2019).
40. S. L. Jacques, “Optical properties of biological tissues: a review,” *Phys. Med. Biol.* **58**(11), R37 (2013).
41. D. E. Rumelhart, G. E. Hinton, and R. J. Williams, “Learning representations by back-propagating errors,” *Nature* **323**(6088), 533–536 (1986).

42. P. Virtanen et al., “SciPy 1.0: fundamental algorithms for scientific computing in Python,” *Nat. Methods* **17**(3), 261–272 (2020).
43. L. Giannoni, F. Lange, and I. Tachtsidis, “Investigation of the quantification of hemoglobin and cytochrome-c-oxidase in the exposed cortex with near-infrared hyperspectral imaging: a simulation study,” *J. Biomed. Opt.* **25**(4), 046001 (2020).
44. I. Tachtsidis, “Specific extinction spectra of tissue chromophores,” GitHub Repository, <https://github.com/multimodalspectroscopy/UCL-NIR-spectra> (2022).
45. S. J. Matcher, M. Cope, and D. T. Delpy, “Use of the water absorption spectrum to quantify tissue chromophore concentration changes in near-infrared spectroscopy,” *Phys. Med. Biol.* **39**, 177–196 (1994).
46. W. G. Zijlstra, A. Buursma, and O. W. Van Assendelft, *Visible and Near Infrared Absorption Spectra of Human and Animal Haemoglobin Determination and Application*, CRC Press, London (2021).
47. U. Heinrich, “Untersuchungen zur qualitativen photometrischen Analyse der Redox-Zustände der Atmungskette in vitro und in vivo am Beispiel des Gehirns,” PhD Thesis, Abteilung für Biologie, Ruhr-Universität Bochum, Bochum (1981).

Biographies of the authors are not available.