

Clinical Validation of Deep Learning for Real-Time Tissue Oxygenation Estimation Using Spectral Imaging

Jens De Winne^{1,2}[0000–0003–2506–7351], Siri Willems¹[0000–0002–4269–1976], Siri Luthman¹[0000–0002–8539–826X], Danilo Babin²[0000–0002–2881–6760], Hiep Luong²[0000–0002–6246–5538], and Wim Ceelen^{3,4}[0000–0001–7692–4419]

¹ Imec, Kapeldreef 75, 3001 Leuven, Belgium

² Department of Telecommunications and Information Processing, Ghent University/imec, Sint-Pietersnieuwestraat 41, 9000 Ghent, Belgium

³ Department of Gastrointestinal Surgery, Ghent University Hospital, 9000 Ghent, Belgium

⁴ Department of Human Structure and Repair - Experimental Surgery Lab, Ghent University, 9000 Ghent, Belgium

Abstract. Accurate, real-time monitoring of tissue ischemia is crucial to understand tissue health and guide surgery. Spectral imaging shows great potential for contactless and intraoperative monitoring of tissue oxygenation. Due to the difficulty of obtaining direct reference oxygenation values, conventional methods are based on linear unmixing techniques. These are prone to assumptions and these linear relations may not always hold in practice. In this work, we present deep learning approaches for real-time tissue oxygenation estimation using Monte-Carlo simulated spectra. We train a fully connected neural network (FCN) and a convolutional neural network (CNN) for this task and propose a domain-adversarial training approach to bridge the gap between simulated and real clinical spectral data. Results demonstrate that these deep learning models achieve a higher correlation with capillary lactate measurements, a well-known marker of hypoxia, obtained during spectral imaging in surgery, compared to traditional linear unmixing. Notably, domain-adversarial training effectively reduces the domain gap, optimizing performance in real clinical settings.

Keywords: multispectral imaging · tissue oxygenation · domain adaptation · deep learning · simulations · real-time imaging

1 Data

1.1 Simulated data

Diffuse reflectance spectra are simulated in silico using Monte-Carlo simulations of light propagation in an artificial tissue. As tissues's oxygenation is defined within the simulation, reflectance spectra with a ground-truth oxygenation label

are obtained and used to train the deep learning models. This simulation framework follows a similar approach as Ayala et al. [2, 8]. An overview is illustrated in Fig. 1.

Soft tissue typically consists of three distinct layers, from outermost to innermost: the serosa, muscularis and submucosa. Hence, we implement a three-dimensional 3-layered tissue model, see Fig. 1a. For each of the layers in the tissue model, physiological tissue parameters are uniformly sampled from a prior distribution. Additionally, the absorption coefficient μ_a and scattering coefficient μ_s are calculated based on the sampled values and used as input for the simulation of the light propagation. Prior distributions and μ_a and μ_s calculation are identical to [2]. The size of the tissue comprised $20 \times 20 \times D$ voxels, with a voxel size of 0.01 mm. The depth D of the tissue depends on the sampled thickness value for each layer.

Simulations were performed using a GPU-accelerated Python implementation of the MCX software [5]. A uniform planar light source was defined above the tissue and the number of photons at launch was set to $1 \cdot 10^5$. A total of $6.4 \cdot 10^5$ reflectance spectra were simulated at wavelengths 440 nm to 640 nm in steps of 4 nm, see Fig. 1b. The diffuse reflectance obtained at the surface was averaged resulting in a reflectance spectra $r(\lambda)$ of 51 wavelengths.

Labeling the oxygenation value from a multi-layered tissue is performed using the method proposed by Ayala et al. [1]. During every simulation, the penetration depth $p(\lambda)$ is calculated as the depth inside the tissue at which the intensity of incident light at wavelength λ drops to a ratio of $\frac{1}{e}$. A weighted average oxygenation is then calculated based on $p(\lambda)$ and the sampled thickness for every layer.

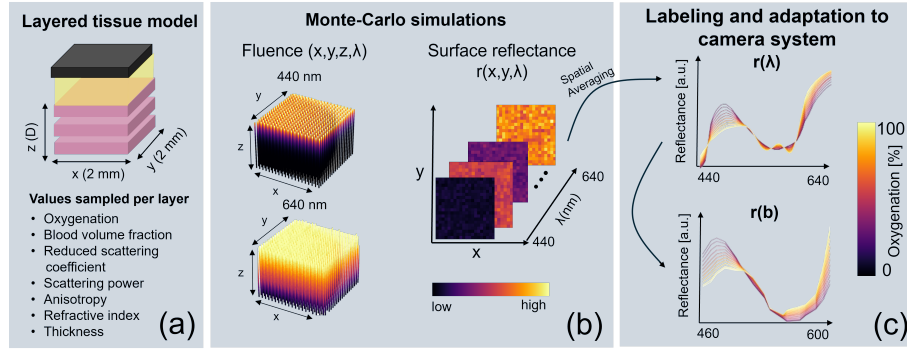


Fig. 1: Overview of the approach to simulate oxygenation labeled reflectance spectra.

The obtained reflectance spectra $r(\lambda)$ from the simulations are independent of the imaging system. These are adapted to resemble measurements using a multispectral camera (Fig. 1c):

$$r(b) = \frac{w(b) + \int_{440nm}^{640nm} r(\lambda) \cdot L(\lambda) \cdot T(\lambda) \cdot s(b, \lambda) d\lambda}{\int_{440nm}^{640nm} L(\lambda) \cdot T(\lambda) \cdot s(b, \lambda) d\lambda} \quad (1)$$

with $r(b)$ denoting the reflectance at band b of the multispectral camera, $w(b)$ the camera noise for band b , $L(\lambda)$ the relative irradiance of the light source, $T(\lambda)$ the transmission profile of optical components, and $s(b, \lambda)$ the spectral response curve for band b . Finally, the spectra were normalized by dividing their area under the curve to compensate for changes in imaging distance and angles between the camera and tissue.

1.2 Clinical data

Spectral data was acquired for 17 patients scheduled for elective laparoscopic esophagectomy. Immediately before resecting the stomach graft, as part of the procedure to connect it to the esophagus by means of a anastomosis, the surgeon visually identified and marked three regions of interest (ROIs). One corresponding to the chosen anastomotic site. The second being a well-perfused region located 2-3 *cm* proximal from the pylorus. In addition, a small piece of tissue excised from the proximal stomach acted as an ischemic control. In 11 patients, capillary lactate was measured at every ROI using a handheld lactate analyzer (EDGE Analyzer, Apex Biotechnology, Taiwan) to further assess hypoxia.

Video-rate MSI was performed immediately before the lactate measurements. The XIMEA SNAPSHOT VIS (Imec, Belgium), a multispectral snapshot camera capturing images at 16 different wavelengths in the 460 - 600 *nm* spectral range was used. The camera was mounted to a standard laparoscope using a C-mount coupler (C-mount zoom adapter, RVA Synergies, UK), and connected to a Xenon light source (CLV-S190, Olympus, Japan). During imaging, all three ROIs were within the camera's field of view. The study was approved by the ethical committee of the Ghent University Hospital (B670201836427) and informed consent was obtained from all patients. The trial is registered at Clinicaltrials.gov with registration number NCT03587532.

A radiometric calibration method was employed [7]. In short, the raw images underwent dark subtraction and demosaicking was performed using an algorithm developed by Muszyński and Luong [6]. Only in the case of linear unmixing, each pixel spectrum was multiplied with the correction matrix supplied by the manufacturer. Values were further divided by the spectra of the light source, obtaining reflectance data. The spectra were further normalized by dividing by their area under the curve. To obtain unlabeled real spectra later used for training the domain-adversarial models, a single frame was selected for every patient and the stomach graft was manually segmented.

2 Methods

2.1 Deep learning models

Four deep learning models are trained. Model architectures were kept small to ensure real-time oxygenation estimation. First, a fully connected (FCN) and convolutional neural network (CNN) were trained without a domain-adversarial approach. The FCN consists of three hidden layers containing 64, 128 and 256 nodes. In the CNN architecture, 2 convolutional layers with sizes 16 and 32 are introduced, with a kernel size of 2 performing convolutions in the spectral domain. Each layer is followed with a ReLU activation function, batch normalization and a dropout rate of 20%. The output layer is a single linear layer estimating an oxygenation value.

To construct the domain adversarial networks, the same 2 model structures are used as before with the addition of a domain discriminator: a domain-adversarial fully connected network (DA-FCN) and a domain-adversarial convolutional network (DA-CNN). These models produce deep features in their hidden layers acting as a generator. A linear output layer, the discriminator, is added to the network classifying whether the produced features from the generator originate from a simulated or real spectrum.

2.2 Experiments

In total, $6.4 \cdot 10^5$ labeled spectra were simulated and adapted to the used MSI system. They were split into a training (80%) and validation (20%) set, and stratified based on the oxygenation. Before every epoch, the training data was augmented by simulating camera noise by means of additive Gaussian noise. The noise was added to every band b by means of the $w(b)$ parameter in Eq. 1: $w(b) \sim \eta(0, \sigma^2)$, where σ is adjusted to obtain an SNR of 40 dB. The validation data was similarly augmented, but only once at the start of the training. Training of the FCN and CNN was performed using the Mean Squared Error (MSE) loss. The learning rate was set to 10^{-3} , and scheduled to reduce by a factor of 10 if the validation loss did not decrease for 10 epochs by as much as 1%.

In the case of DA-FCN and DA-CNN, unlabeled real spectra were added to train both the generator and discriminator. To fully separate the training and test set, 6 out of the 17 patients without a corresponding lactate measurement were used for training. From these, a total of $1.2 \cdot 10^5$ spectra were derived from their manually annotated stomach graft and split into a training (80%) and validation set (20%). For testing, a total of 7328 spectra were derived from 20x20 pixel ROIs of the 11 other patients that did have a corresponding lactate measurement. During training, a balanced sampler is used to ensure equal representations of simulated and real spectra. A domain label of 0 was given to simulated spectra and 1 to real spectra. A combined loss function was used incorporating the MSE for the regression task (L_r) and Binary Cross Entropy (BCE) loss as adversarial loss for the domain classification task (L_a and L_D):

$$L = L_G + L_D = (L_r + \lambda_a L_a) + L_D \quad (2)$$

with L_G the loss used for training the generator, L_D the loss for training the discriminator, and λ_a a factor to control the contribution of the adversarial loss set to 0.25. Before computing BCE losses, the output was transformed using a sigmoid transformation. The learning rate for the generator was set to 10^{-3} and that of the discriminator was to 10^{-6} , with the same schedulers as before. Every model was trained with an Adam optimizer with a weight decay of 10^{-6} , a batch size of 512, and for 100 epochs.

2.3 Clinical validation

Retrospective clinical validation is performed on the test set containing 11 patients with a corresponding capillary lactate measurement. The oxygenation quantification methods are evaluated by analyzing the relationship between oxygenation and capillary lactate, a clinical marker of hypoxia. The relationship between lactate and oxygenation is commonly described using an exponential function as follows [3]:

$$\text{Capillary lactate} = A \cdot e^{B \cdot O_2} \quad (3)$$

where O_2 corresponds to the oxygenation value and A and B are unknown fitting parameters of the exponential function.

The parameters A and B are determined using an ordinary least squares fit between the capillary lactate values and corresponding oxygenation estimates (averaged over a region of 20x20 pixels around the lactate sampling point). The mean absolute error (MAE) and R^2 value of the fit, and the correlation coefficient, are selected to assess the performance of the different oxygenation quantification methods. The higher the correlation with lactate and the better the fit with the exponential function, the better the relation between predicted oxygenation and the hypoxia marker capillary lactate. In addition, these values are compared to those obtained by predicting oxygenation using a linear mixing approach outlined in previous work [4].

3 Results

The performance of different models for oxygenation quantification is summarized in Table 1, containing the training and validation regression losses, along with the evaluation metrics from the clinical validation. In general, the deep learning approach shows an improved correlation with lactate measurements compared to the linear unmixing technique. Moreover, adversarial training leads to the highest correlation resulting in clinically improved oxygenation estimations. The best performing model on the clinical test set is the DA-FCN. Fig. 2 shows the exponential fit of the capillary lactate with the oxygenation predictions of this model compared to linear unmixing. Fig. 3 shows oxygenation maps

of the stomach graft from a single multispectral image from the clinical test data. All models successfully estimate a high oxygenation for the well-perfused tissue and critically low oxygenation for the ischemic tissue. Furthermore, the anastomosis site seems to lie at a borderline, as a transition into lower oxygenated tissue can be seen. The deep learning models seem to estimate less extreme, and more realistic, values compared to linear unmixing.

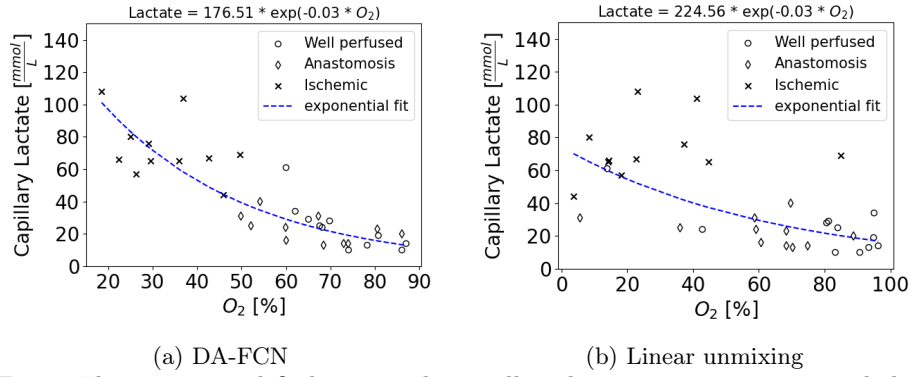


Fig. 2: The exponential fit between the capillary lactate measurements and the predicted oxygenation of the best model (a) and linear unmixing (b).




Fig. 3: (a) shows a synthetic RGB image of a patient included in the clinical test set, along with the selected ROIs of the well-perfused, anastomosis and ischemic tissue, (b) Oxygenation maps produced by linear unmixing and machine learning methods. The scale ranges from 0% (dark purple) to 100% (bright yellow).

	Train loss ($\cdot 10^{-3}$)	Validation loss ($\cdot 10^{-3}$)	MAE [mmol/L]	R^2 [a.u.]	Correlation [a.u.]	Inference time [ms]
FCN	6.47	9.59	10.81 ± 10.73	0.74	-0.86	16.64
CNN	8.80	9.86	13.33 ± 11.98	0.71	-0.84	25.28
DA-FCN	8.04	11.78	9.83 ± 10.05	0.76	-0.87	16.80
DA-CNN	11.04	12.06	12.26 ± 10.54	0.74	-0.86	25.88
Linear unmixing			15.32 ± 15.94	0.46	-0.68	94.68

Table 1: Model training and validation regression losses, and their performance on the clinical test data with corresponding lactate measurements. In addition, the average inference times over 1000 iterations on a multispectral image are shown using an NVIDIA GeForce RTX 4060 Laptop GPU.

4 Discussion

Deep learning models improve oxygenation estimation compared to traditional linear unmixing, showing a stronger correlation with the hypoxia marker capillary lactate. Despite higher training loss, domain-adversarial training further enhances performance by reducing the domain gap between simulated and real data. This occurs because the generator learns deep features optimized for oxygenation estimation while maintaining domain invariance through adversarial training. As a result, these models demonstrate greater robustness, leading to improved performance on the clinical test set.

Both deep learning and linear unmixing detect large oxygenation differences, with linear unmixing showing more extreme values, particularly between well-perfused and ischemic tissue. However, their estimates diverge in regions with subtle differences, such as the anastomosis. As shown in Fig. 2b, linear unmixing tends to underestimate oxygenation in the anastomosis and overestimate it in ischemic tissue, leading to outliers in the fit. Deep learning mitigates these errors, resulting in a more accurate correlation.

Additionally, deep learning models achieve oxygenation mapping at 50–62 frames per second on a low-end laptop GPU, which is 5–6 times faster than linear unmixing. This makes them a promising tool for intraoperative guidance based on real-time oxygenation estimation.

As the study is limited by the number of patients and tissue types included, future work will enhance statistical significance and robustness by addressing this. Additionally, we are exploring the transformation of 1D simulated spectra into 2D data to incorporate spatial information.

5 Conclusion

Deep learning methods, particularly domain-adversarial networks, improve oxygenation estimation compared to traditional linear unmixing. Our results demonstrate the potential of deep learning to translate spectral imaging into a clinically

viable tool for real-time oxygenation assessment, showing strong correlations with a hypoxia marker.

Acknowledgments. This work has received funding from the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” programme.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Ayala, L.: Translational Functional Imaging in Surgery Enabled by Deep Learning. Ph.D. thesis, Dissertation, Heidelberg, Universität Heidelberg, 2023 (2023). <https://doi.org/10.11588/heidok.00033281>, <http://www.ub.uni-heidelberg.de/archiv/33281>
2. Ayala, L., Adler, T.J., Seidlitz, S., Wirkert, S., Engels, C., Seitel, A., Sellner, J., Aksenov, A., Bodenbach, M., Bader, P., Baron, S., Vemuri, A., Wiesenfarth, M., Schreck, N., Mindroc, D., Tizabi, M., Pirmann, S., Everitt, B., Kopp-Schneider, A., Teber, D., Maier-Hein, L.: Spectral imaging enables contrast agent-free real-time ischemia monitoring in laparoscopic surgery. *Science Advances* **9**(10), eadd6778 (2023). <https://doi.org/10.1126/sciadv.add6778>, <https://www.science.org/doi/abs/10.1126/sciadv.add6778>
3. Barberio, M., Longo, F., Fiorillo, C., Seeliger, B., Mascagni, P., Agnus, V., Lindner, V., Geny, B., Charles, A.L., Gockel, I., Worreth, M., Saadi, A., Marescaux, J., Diana, M.: Hyperspectral enhanced reality (hyper): a physiology-based surgical guidance tool. *Surgical Endoscopy* **34**, 1736–1744 (2020). <https://doi.org/10.1007/s00464-019-06959-9>
4. De Winne, J., Strumane, A., Babin, D., Luthman, S., Luong, H., Philips, W.: Multispectral indices for real-time and non-invasive tissue ischemia monitoring using snapshot cameras. *Biomed. Opt. Express* **15**(2), 641–655 (Feb 2024). <https://doi.org/10.1364/BOE.506084>, <https://opg.optica.org/boe/abstract.cfm?URI=boe-15-2-641>
5. Fang, Q., Boas, D.A.: Monte carlo simulation of photon migration in 3d turbid media accelerated by graphics processing units. *Opt. Express* **17**(22), 20178–20190 (Oct 2009). <https://doi.org/10.1364/OE.17.020178>, <https://opg.optica.org/oe/abstract.cfm?URI=oe-17-22-20178>
6. Muszyński, R., Luong, H.: Cube it: Training hyperspectral demosaicing models using synthetic datasets. In: 2024 14th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS). pp. 1–5 (2024). <https://doi.org/10.1109/WHISPERS65427.2024.10876469>
7. Vunckx, K., Charle, W.: Accurate video-rate multi-spectral imaging using imec snapshot sensors. In: 2021 11th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS). pp. 1–7 (2021). <https://doi.org/10.1109/WHISPERS52202.2021.9483975>
8. Wirkert, S.J.: Multispectral image analysis in laparoscopy – A machine learning approach to live perfusion monitoring. Ph.D. thesis, Karlsruher Institut für Technologie (KIT) (2018). <https://doi.org/10.5445/IR/1000086188>