A sparse deep learning approach for automatic segmentation of human vasculature in multispectral optoacoustic tomography

Nikolaos-Kosmas Chlis [a,b,g,1], Angelos Karlas [b,c,d,e,1], Nikolina-Alexia Fasoula [b,c], Michael Kallmayer [e], Hans-Henning Eckstein [e], Fabian J. Theis [a,f], Vasilis Ntziachristos [b,c,d,1], Carsten Marr [a,*,1]

[a] Institute of Computational Biology, Helmholtz Center Munich, Neuherberg, Germany
[b] Institute of Biological and Medical Imaging, Helmholtz Center Munich, Neuherberg, Germany
[c] Chair of Biological Imaging and Center for Translational Cancer Research (TranslaTUM), Munich, Germany
[d] DZHK (German Centre for Cardiovascular Research), Partner Site Munich Heart Alliance, Munich, Germany
[e] Clinic for Vascular and Endovascular Surgery, Rechts Der Isar Hospital, Munich, Germany
[f] Department of Mathematics, Technical University of Munich, Munich, Germany
[g] Roche Pharma Research and Early Development, Large Molecule Research, Roche Innovation Center Munich, Penzberg 82377, Germany

ARTICLE INFO

ABSTRACT

Multispectral Optoacoustic Tomography (MSOT) resolves oxy- (HbO_2) and deoxy-hemoglobin (Hb) to perform vascular imaging. MSOT suffers from gradual signal attenuation with depth due to light-tissue interactions: an effect that hinders the precise manual segmentation of vessels. Furthermore, vascular assessment requires functional tests, which last several minutes and result in recording thousands of images. Here, we introduce a deep learning approach with a sparse-UNET (S-UNET) for automatic vascular segmentation in MSOT images to avoid the rigorous and time-consuming manual segmentation. We evaluated the S-UNET on a test-set of 33 images, achieving a median DICE score of 0.88. Apart from high segmentation performance, our method based its decision on two wavelengths with physical meaning for the task-at-hand: 850 nm (peak absorption of oxy-hemoglobin) and 810 nm (isosbestic point of oxy-and deoxy-hemoglobin). Thus, our approach achieves precise data-driven vascular segmentation for automated vascular assessment and may boost MSOT further towards its clinical translation.

1. Introduction

The abundant presence of hemoglobin in the blood renders multispectral optoacoustic tomography (MSOT) an ideal technique for imaging vasculature [1–3]. By illuminating tissue at multiple different light wavelengths at the near infrared range (~680−980 nm), MSOT is capable of resolving several tissue chromophores, in particular oxy-(HbO_2) and deoxy-hemoglobin (Hb), with a wide range of clinical applications, such as Crohn's disease, systemic sclerosis, breast cancer, brown adipose tissue imaging and thyroid disease [4–8]. MSOT can provide precise structural visualizations of arteries and veins by recording multispectral data and resolving the different oxygenation states of human hemoglobin molecule. Moreover, the dynamic nature of the vascular system requires the acquisition not only of structural but also of functional data over multiple seconds or minutes to observe, for example, the vascular wall kinetics during the cardiac cycle or the arterial responses to stimuli such as the transient arterial occlusion or hyperthermia, which are valid descriptors of cardiovascular risk [9,10]. The need to record multispectral data in order to extract molecular information and to perform longitudinal measurements over several minutes radically increases the number of recorded images and the data volume.

Both structural and functional vascular imaging require the precise segmentation of the vascular lumen in several applications, such as the quantification of an atheromatous arterial stenosis, the detection of a venous thrombosis or the tracking of the arterial diameter over a 5-minute arterial occlusion challenge to quantify the degree of endothelial dysfunction. The segmentation of the vascular lumen is usually

---

* Corresponding author.
  E-mail address: carsten.marr@helmholtz-muenchen.de (C. Marr).
[1] Equal contribution.

performed by expert physicians who manually draw the regions of interest (ROIs) on the recorded MSOT images. However, manual segmentation is a time-consuming process, in particular in the case of longitudinal recordings of several minutes and thus of hundreds or thousands of frames. Furthermore, because of the gradual light attenuation due to scattering and absorption when propagating in living tissue, the vascular lumen shows an inhomogeneous and fainting intensity profile with increasing depth, making its manual delineation a challenging process. But even in routine imaging diagnostics, a reliable automated segmentation method can be beneficial by aiding the clinician in performing the same task much faster. Deep learning has been recently shown to be very effective in computer vision tasks [11,12] and segmentation in particular [13–15]. As such, deep learning has been successfully applied to clinical diagnostics [16–18], with medical image segmentation applications including prostate [19], retinal disease [20], brain [21,22] and cervical cell segmentation [23]. Surveys of deep learning applications for medical imaging can be found in [24,25].

We present herein a pilot study to achieve automated vascular segmentation in clinical raw MSOT images via a deep learning approach, based on an extension of the UNET architecture originally introduced in [15] that is specifically tailored for multispectral optoacoustic data. The proposed Sparse UNET (S-UNET) allows for automated segmentation of vascular ROIs in clinical MSOT images, while simultaneously identifying which of the employed illumination wavelengths are relevant to the specific task. This way we aim at radically reducing the time needed for vascular segmentation in longitudinal scans as well as the number of illumination wavelengths for future task-specific scans, facilitating this way the data analysis, increasing the time resolution and reducing the data volume.

2. Methods

2.1. Network architecture

The proposed Sparse-UNET (S-UNET) is based on the fully convolutional architecture of the original UNET [15], with the added capability of sparse wavelength selection. The goal of S-UNET is to transform each input image with dimensions $400 \times 400 \times 28$ (Height x Width x Wavelengths) into a 400×400 probability map p that corresponds to a ground truth segmentation mask, while simultaneously assigning a weight w_c (wavelength importance) to each of the 28 illumination wavelengths (from 700 to 970 nm at steps of 10 nm), which correspond to the 28 channels of the input image. The ground truth segmentation mask y is a binary image (each pixel is either 0 or 1), extracted from the recorded MSOT image in consensus between two clinical MSOT experts. To arrive at a predicted segmentation mask, the resulting S-UNET probability map p is discretized by thresholding at 0.5: pixels with probabilities less than 0.5 are set to 0, while the rest are set to 1.

In order to perform wavelength selection, the first layer of the S-UNET corresponds to a 1×1 2D convolution of a single filter and no bias. Given each $400 \times 400 \times 28$ input image stack, the first layer essentially performs a linear combination of the 28 wavelengths, resulting in a $400 \times 400 \times 1$ image that is forward-propagated to the rest of the network. In this manner, each wavelength is assigned a unique scalar weight. Moreover, to ensure sparsity of wavelength selection we add L1 regularization [26] on the wavelength weights. L1 regularization refers to adding a term $\lambda|\beta|$, where λ is a scalar hyperparameter and β refers to the first convolutional layer's trainable parameters, in this case, the parameters of the first 1×1 convolutional layer. Here, we employed a regularization parameter $\lambda = 0.01$. Regularization does not necessarily result in an interpretable model. To ensure interpretability of wavelength selection we force the weights of the first layer to be non-negative. As such, there is no possibility to have irrelevant wavelengths of similar wavelengths cancelling each other out with weights of similar, potentially high, magnitude and opposite signs.

Taken together, the two constraints of L1 regularization and non-negative weights ensure that only few relevant wavelengths will be assigned with positive weights, while all other non-relevant wavelengths will be set to zero and will effectively be excluded from the model. After wavelength selection, we add a batch normalization [27] layer between every convolution layer and its respective activation function. The S-UNET architecture employed is visualized in Fig. 1.

2.2. Training and data augmentation

The original dataset of 164 raw MSOT images was randomly split into training, validation and test sets of 98, 33 and 33 images, respectively.

Each raw MSOT image corresponds to spatial dimensions of 400×400 pixels (which corresponds to 4×4 cm) and 28 wavelengths. Each wavelength is normalized to values between 0 and 1 separately, as part of pre-processing. Normalization of the input image is a common practice in deep learning applications [28]. We train the model on a training subset of the data using Adam [29] while evaluating model performance on a validation set. The model is trained for a maximum number of 200 epochs, or until model performance on the validation set has not improved for 20 consecutive epochs (early stopping). The instance of the model that achieved the best performance on the validation set is saved as the final model. We keep a separate test set that is hidden from the model during training.

The model is trained using a batch size of 4 images and data augmentation is performed on-the-fly on each image in every batch to increase model performance. Data augmentation includes flipping the x axis and rotating the image in a random angle from 5 to 15 degrees. Each of the two augmentation schemes has a 50 % probability of being performed on any given image. According to our experiments more aggressive augmentation hinders model performance on the given task. The last layer of the model corresponds to a pixel-wise binary classification problem of computing a probability map of the predicted segmentation mask. The model's loss corresponds to the loss of the 400×400 binary classification tasks. Thus, the total binary cross entropy loss function L, is used to train the model:

$$L = -\sum_{h=1}^{H}\sum_{w=1}^{W} \left(y_{hw} \cdot ln(p_{hw}) + (1 - y_{hw}) \cdot ln(1 - p_{hw}) \right).$$

Here, H and W correspond to the image height and width in pixels (each being 400), $y_{hw} \in \{0, 1\}$ corresponds to the ground truth segmentation class, $p_{hw} \in [0, 1]$ corresponds to the predicted class probability for the corresponding pixel in position (h, w) and ln is the natural logarithm.

3. Experiments

3.1. Data acquisition

In this pilot study we scanned six (n = 6) healthy volunteers (3 men, 3 women, age 30 ± 5.44 years). All healthy volunteers consented to participate in this study in full accordance with the work safety regulations of the Helmholtz Center Munich (Neuherberg, Germany). The radial artery, the brachial artery, the dorsal artery of the foot, as well as the cephalic vein, the radial veins and the dorsal vein of the foot were scanned by means of a clinical hand-held MSOT/Ultrasound system (iThera Medical GmbH, Munich, Germany). All subjects were asked to consume no food or caffeine for 8 h before the examination, which was conducted in a quiet dark room with normal temperature of 25 °C. Each scan lasted for 5–10 seconds. The system used was equipped with a near-infrared laser for achieving optimal penetration depth in tissue (3−4 cm) even with low illumination energy (~15 mJ per pulse). For multispectral data recording we used 28 wavelengths (700:10:980 nm). Tissue was illuminated by short light pulses (~10 ns) at a frame rate of 25fps. The ultrasound detection was performed by 256 ultrasound
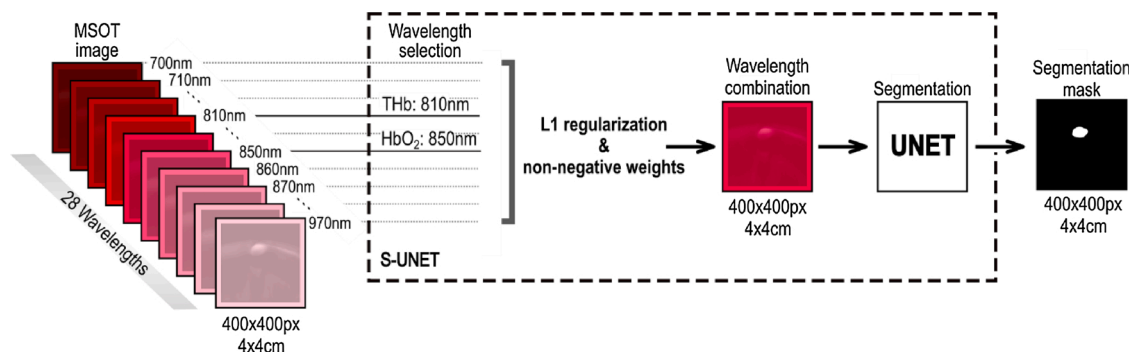
Fig. 1. The S-UNET identifies important illumination wavelengths in MSOT images while learning to predict segmentation masks of human blood vessels. Each wavelength is weighted by a corresponding non-negative weight and all weighted wavelengths are combined before being inserted as input into a UNET architecture. Sparsity of wavelength selection is enforced by L1 regularization on the non-negative wavelength weights and the weights themselves are learned through standard back-propagation, along with the rest of the UNET parameters.

sensors with a central frequency of 4 MHz which covered an angle of 145° and was mounted on the hand-held scanning probe. Acquired ultrasound signals for each illumination pulse were reconstructed into a tomographic image using a model-based reconstruction algorithm [30]. For each MSOT image a co-registered ultrasound image was recorded. The segmentation of the scanned arteries and veins was manually performed on the appropriate MSOT frame by simultaneous view of the co-registered ultrasound image. We decided to segment the blood vessels directly on the MSOT frames because of better contrast, compared to ultrasound, provided by the high light absorption of hemoglobin at the near-infrared illumination range. The appropriate frame for vein segmentation was the frame corresponding to the 750 nm illumination wavelength were the absorption of Hb is clearly higher than that of HbO_2. The appropriate frame for artery segmentation was the frame corresponding to the 850 nm illumination wavelength were the absorption of HbO_2 is clearly higher than that of Hb. Manual segmentation was conducted in consensus of two clinicians with experience in MSOT and clinical ultrasound imaging.

3.2. Model comparison

We compared the performance of four segmentation methods on the recorded MSOT dataset: the proposed S-UNET, UNET++ [31], and two differently-sized variants of a standard UNET on the segmentation task. The S-UNET architecture is described in section 2.1 above. The wavelength selection layer is followed by a downsized UNET where every convolutional layer corresponds to 1/8 of filters compared to the architecture in [15]. Two variants of the UNET were applied: one with the same number of filters as in [15] ('original') and a variant with 1/8 of filters ('downsized'). Additionally, a batch normalization layer was inserted between every convolutional layer and its corresponding activation function in both UNET variants. Training was performed as described in the previous section. The results of all segmentation methods were compared to the binary ground truth segmentation mask, which was manually generated from expert clinicians on the recorded MSOT images under co-registered ultrasound guidance (see Methods). Model comparison is based on the Dice coefficient [22] defined as:

$$Dice = 2TP/(2TP + FP + FN)$$

where TP, FP and FN correspond to true positive, false positive and false negative classified pixels: A TP pixel is a correctly classified foreground pixel, a FP pixel is a background pixel falsely classified as foreground, and a FN pixel corresponds to a foreground pixel that was incorrectly classified as background by the model. The Dice coefficient is well-suited to tackle the class imbalance inherent to the segmentation task [32], where more than 99 % of the pixels in our dataset are background pixels. As such, it is preferred for model assessment compared to the standard

cross entropy used to train the model. The Dice coefficient lies between 0 and 1, with higher values being better since they correspond to larger overlap between the ground truth and predicted segmentation masks.

3.3. Wavelength selection

Wavelength selection was performed by the first layer of the S-UNET (see Methods). However, since feature selection is an inherently noisy process [33,34] it is good practice to average a number of models [35] in order to obtain a smoothed version of wavelength importance. We thus train 100 different instances of the S-UNET and aggregate their results for the tasks of segmentation, as well as for wavelength selection. In the case of segmentation, we average the probability maps of all models before discretizing in order to obtain the binary segmentation mask.

3.4. Results

The performance of all four segmentation models (Dice coefficient) is reported in Table 1. The original UNET with over 30 million parameters is potentially slightly overfitting the training dataset while the downsized UNET, as well as the S-UNET achieve very similar segmentation results with roughly half a million parameters. The downsized UNET achieves slightly higher Dice scores on average (0.90 ± 0.08) than the S-UNET ensemble (0.86 ± 0.11), but the difference is not statistically significant given the test set size of 33 images (p-value = 0.37, two-sample Wilcoxon rank-sum test). This similarity in performance is to be expected since both methods correspond to a similar number of parameters. However, this also suggests that the added sparsity of wavelength selection does not affect, at least not significantly, the quality of the generated segmentation masks in the case of S-UNET. Finally, UNET++ performs worse than all other deep learning methods, achieving a Dice coefficient of 0.61 ± 0.26.

The advantage of the proposed S-UNET approach over the other UNET approaches is clearly its interpretability of results due to the embedded wavelength selection. As visualized in Fig. 2, out of the 28 input wavelengths the model has identified two as being the most important in a purely data-driven manner. These wavelengths correspond to the maximum of the absorption spectra of total blood volume (HbO_2 and Hb) at 810 nm and HbO_2 at 850 nm. Both of these identified

Table 1
Model Performance. Dice results correspond to mean ± std.

| Model	Test Set Dice	Parameters	Wavelength Selection
UNET (original)	0.75 ± 0.28	31,416,897	No
UNET (downsized)	0.90 ± 0.08	495,881	No
UNET++	0.61 ± 0.26	9,049,377	No
S-UNET	0.86 ± 0.11	493,965	Yes

Fig. 2. The S-UNET identifies wavelengths relevant to the segmentation task. Each boxplot (the box's edges correspond to quartiles 1 and 3 while whiskers extend to ± 1.5 times the interquartile range) corresponds to the weights assigned by the ensemble of 100 S-UNET instances to each wavelength. Averaging results is necessary since feature selection is an inherently noisy process. For every S-UNET instance, each of the 28 wavelengths of the input image is multiplied by its corresponding weight and all 28 weighted single-wavelength images are added in a pixel-wise manner. This step results in a single-channel image being passed on to the following layers of the network. According to the median weight of each wavelength, the two most important wavelengths are 850 nm and 810 nm, corresponding to the maximum absorption of HbO_2 and total hemoglobin (THb), respectively.

wavelengths are thus meaningful since they mark the presence of blood in the detected image regions. To further highlight the importance of wavelength selection we re-evaluated the S-UNET performance using only the two most important wavelengths (810 and 850 nm). The 2-wavelength S-UNET achieved a mean Dice score of 0.73 on the test set. This corresponds to a 13 % decrease in performance compared to the 28-wavelength S-UNET (0.73 Dice instead of 0.86), while it simultaneously reduces the required data volume by 93 % (2 wavelengths instead of 28). Additionally, when we train the same model using only two alternative wavelengths (750 and 870 nm), it achieves a lower mean Dice score of 0.70. This suggests that while segmentation is possible using other wavelengths, the two wavelengths identified by S-UNET have a higher predictive value. The segmentation results of S-UNET on an exemplary set of images are visualized in Fig. 3. Interestingly, our approach is able to discriminate blood vessels from similar objects probably by exploiting the wavelength information.

Blood vessel segmentation via deep learning is more accurate than classical thresholding methods, such as Sauvola's adaptive [36] and Otsu's global thresholding [37]. Both methods are available in scikit-image, a Python package for image processing [38] and require a single grayscale image as input. For this reason, a grayscale image was produced by calculating the average value of each pixel across the 28 image channels. On the one hand, local thresholding (Sauvola's method) achieved poor results (mean Dice of $0.02 + 0.01$) since the region of interest (vascular lumen) corresponds to a single region of maximum intensity values. On the other hand, Otsu's global thresholding performs better than local thresholding (mean Dice of $0.24 + 0.23$), but considerably worse than deep learning approaches. When using only the two most important wavelengths (instead of all 28) to compute the grayscale input image, the performance of Sauvola thresholding remains practically identical while Otsu's method achieves better results on average (mean Dice of $0.41 + 0.33$).

4. Discussion

In this work we applied a deep learning approach based on an

Fig. 3. The S-UNET successfully segments human vasculature from MSOT images. Each row corresponds to a different image of the test set. The first column (images a, e, i, m) shows the 850 nm channel of the MSOT image. The second (images b, f, j, n) and third columns (c, g, k, o) show the ground truth (true mask, blue) and predicted segmentation masks (red), respectively, visualized on top of the input image. The true segmentation mask is identified by expert physicians, while the S-UNET predicted segmentation mask corresponds to the output of the S-UNET ensemble. The fourth column (images d, h, l, p) corresponds to the absolute difference between the true and predicted binary segmentation masks and is equivalent to the logical operation of XOR (exclusive or). The predicted masks almost completely overlap with the ground truth segmentation. The S-UNET is successful even in the last two cases (rows) where the mask is relatively small and located in an area where similar bright spots are present. The white dashed line represents the skin surface. The white arrows point to the blood vessel of interest. The scale bar is 5 mm. The gray color bar ranges from 0 to 1 and corresponds to the normalized intensity of each image (columns 1-3) or the difference of the true and predicted segmentation masks (column 4).

adapted S-UNET to perform automated vascular segmentation in clinical MSOT images. Our model successfully segments blood vessels (arteries and veins) and its performance is comparable to a standard UNET of similar model size. Furthermore, our model is capable of selecting the illumination wavelengths that are most important for the segmentation task at hand in a purely data driven manner. Our results show that among the 28 illumination wavelengths used for data acquisition, two wavelengths are associated with the light absorption of hemoglobin at the near-infrared range of illumination (700–970 nm). These correspond to 810 nm, which is the isosbestic point of HbO_2 and Hb and reflects the absorption of total hemoglobin or else the total blood within the vasculature and 850 nm, which is the point where HbO_2 absorbs significantly more than Hb and reflects the arterial blood.

Our approach achieves accurate automated segmentation of both arteries and veins on raw clinical MSOT data. Apart from facilitating the segmentation process, which is time-costly for longitudinal scans of several minutes during functional vascular testing, it may help tackling a significant limitation of optical and optoacoustic imaging: the attenuation of light due to scattering and absorption when propagating in living tissue. This effect causes a gradual attenuation of the signal intensity in the vascular lumen with increasing depth. Thus, the accurate visualization and segmentation of the lumen constitute real challenges even for clinicians with extensive MSOT experience. In the current study, we scanned blood vessels where this effect was apparent (e.g. Fig. 3a) but not to an extent that would jeopardize the accurate manual segmentation of the vascular lumen directly on the MSOT images under ultrasound guidance. Thus, future studies are required to further investigate the efficacy of deep learning approaches in automatically detecting and segmenting vessels with clinical interest (e.g. the carotid artery) deep in tissue in clinical MSOT data.

In this study, we preferred to work on raw MSOT data. However, the discrete spectral difference of HbO2 and Hb at the near-infrared range as well as the strong presence of HbO2 in arteries and Hb in veins would allow for the direct spectral unmixing of HbO2 and Hb in the MSOT data and thus for direct vascular segmentation. Nevertheless, spectrally unmixed data suffer from errors related to imaging depth and motion, either exogenous (e.g. operator's hand and random patient movement) or endogenous (e.g. arterial pulsation or breathing).

Regarding motion-related errors, the dynamic character of the vascular system introduces significant inaccuracy when it comes to spectral unmixing results, especially when illuminating at multiple different wavelengths (e.g. 28) to achieve high spectral quality. For example, the recording of a multispectral stack of 28 wavelengths at a frame rate of 25 Hz takes more than one second. Considering that the cardiac cycle of a normal individual with a heart rate of $70-75$ Hz is approximately 0.8 s, the use of 28 wavelengths renders the spectrally unmixed data vulnerable to errors due to arterial wall motion, especially in the periphery of the vascular lumen, potentially degrading the precision of vascular segmentation when performed by means of direct spectral unmixing.

Moreover, multispectral optoacoustic imaging at increased tissue depths (> 1 cm), where normally the blood vessels lie, renders the spectral unmixing output vulnerable to the spectral coloring effect: the random absorption and scattering of each illumination wavelength before reaching the HbO2 or the Hb of the vascular lumen according to the optical properties of the set of tissues covering them (e.g. skin, subcutaneous fat, muscle). Thus, usual linear spectral unmixing methods fail to unmix the absorbers of interest (e.g. HbO2 and Hb) at increasing depths since the measured spectra have been colored and thus deflected from the known absorption spectra, as measured in the lab. For the above mentioned reasons, we decided to work on the recorded raw MSOT data.

Our model showed that the decision for segmenting the vasculature was mainly based on two near-infrared wavelengths: the 810 nm where HbO_2 and Hb absorb light to the same extent and the 850 nm where the light absorption of HbO_2 is significantly higher than that of Hb. Our results provide evidence for effective and task-specific wavelength selection via the suggested deep learning model for accurate segmentation of blood vessels in clinical MSOT data. Apart from increasing the time resolution by skipping a number of unnecessary illumination wavelengths and decreasing the data volume, the effective wavelength selection may be used for indirect spectral characterization of more complex tissues or even homogeneous tissues at high depths by identifying the wavelengths critical for achieving their segmentation. This approach may help overcoming the limitations introduced by the spectral coloring effect and thus providing a blind or data-driven spectral unmixing with great implications for clinical MSOT imaging. Our method may be used for segmenting and characterizing tissues with clinical relevance (e.g. the subcutaneous fat or the atherosclerotic plaques which contain lipids, the skeletal muscle which contains water) or even the detection and distribution mapping of injected contrast agents targeting specific molecules involved in the pathophysiology of a disease.

To the best of our knowledge, while deep learning has been used before in the context of optoacoustic imaging data [39,40], this is the first time where a deep learning method is applied to clinical MSOT data. Our approach has significant implications for future MSOT applications with clinical relevance, such as the automated segmentation of more complex soft tissues (e.g. muscle, fat, atherosclerotic plaques) and foreseeable for more accurate diagnosis of vascular disease.

Code availability

The keras implementation of the S-UNET is freely available online at https://github.com/nchlis/sunet.

Declaration of Competing Interest

The authors declare that there are no conflicts of interest.

Acknowledgements

References

[1] A. Karlas, N.-A. Fasoula, K. Paul-Yuan, J. Reber, M. Kallmayer, D. Bozhko, M. Seeger, H.-H. Eckstein, M. Wildgruber, V. Ntziachristos, Cardiovascular optoacoustics: from mice to men – a review, Photoacoustics 14 (2019) 19–30, https://doi.org/10.1016/j.pacs.2019.03.001.

[2] A. Karlas, J. Reber, G. Diot, D. Bozhko, M. Anastasopoulou, T. Ibrahim, M. Schwaiger, F. Hyafil, V. Ntziachristos, Flow-mediated dilatation test using optoacoustic imaging: a proof-of-concept, Biomed. Opt. Express, BOE 8 (2017) 3395–3403, https://doi.org/10.1364/BOE.8.003395.

[3] M. Masthoff, A. Helfen, J. Claussen, A. Karlas, N.A. Markwardt, V. Ntziachristos, M. Eisenblätter, M. Wildgruber, Use of multispectral optoacoustic tomography to diagnose vascular malformations, JAMA Dermatol. 154 (2018) 1457–1462, https://doi.org/10.1001/jamadermatol.2018.3269.

[4] W. Roll, N.A. Markwardt, M. Masthoff, A. Helfen, J. Claussen, M. Eisenblätter, A. Hasenbach, S. Hermann, A. Karlas, M. Wildgruber, V. Ntziachristos, M. Schäfers, Multispectral optoacoustic tomography of benign and malignant thyroid disorders – a pilot study, J. Nucl. Med. (2019), https://doi.org/10.2967/jnumed.118.222174.

[5] J. Reber, M. Willershäuser, A. Karlas, K. Paul-Yuan, G. Diot, D. Franz, T. Fromme, S.V. Ovsepian, N. Bézière, E. Dubikovskaya, D.C. Karampinos, C. Holzapfel, H. Hauner, M. Klingenspor, V. Ntziachristos, Non-invasive measurement of brown fat metabolism based on optoacoustic imaging of hemoglobin gradients, Cell Metab. 27 (2018) 689–701, https://doi.org/10.1016/j.cmet.2018.02.002, e4.

[6] M. Masthoff, A. Helfen, J. Claussen, W. Roll, A. Karlas, H. Becker, G. Gabriëls, J. Riess, W. Heindel, M. Schäfers, V. Ntziachristos, M. Eisenblätter, U. Gerth, M. Wildgruber, Multispectral optoacoustic tomography of systemic sclerosis, J. Biophotonics 11 (2018) e201800155, https://doi.org/10.1002/jbio.201800155.

[7] G. Diot, S. Metz, A. Noske, E. Liapis, B. Schroeder, S.V. Ovsepian, R. Meier, E. Rummeny, V. Ntziachristos, Multispectral optoacoustic tomography (MSOT) of human breast cancer, Clin. Cancer Res. 23 (2017) 6912–6922, https://doi.org/10.1158/1078-0432.CCR-16-3200.

[8] F. Knieling, C. Neufert, A. Hartmann, J. Claussen, A. Urich, C. Egger, M. Vetter, S. Fischer, L. Pfeifer, A. Hagel, C. Kielisch, R.S. Görtz, D. Wildner, M. Engel, J. Röther, W. Uter, J. Siebler, R. Atreya, W. Rascher, D. Strobel, M.F. Neurath, M. J. Waldner, Multispectral optoacoustic tomography for assessment of Crohn's disease activity, N. Engl. J. Med. 376 (2017) 1292–1294, https://doi.org/10.1056/NEJMc1612455.

[9] J. Green Daniel, Jones Helen, Thijssen Dick, N.T. Cable, Atkinson Greg, Flow-mediated dilation and cardiovascular event prediction, Hypertension 57 (2011) 363–369, https://doi.org/10.1161/HYPERTENSIONAHA.110.167015.

[10] S.C. Agarwal, J. Allen, A. Murray, I.F. Purcell, Comparative reproducibility of dermal microvascular blood flow changes in response to acetylcholine iontophoresis, hyperthermia and reactive hyperaemia, Physiol. Meas. 31 (2009) 1–11, https://doi.org/10.1088/0967-3334/31/1/001.

[11] A. Krizhevsky, I. Sutskever, G.E. Hinton, ImageNet classification with deep convolutional neural networks, in: F. Pereira, C.J.C. Burges, L. Bottou, K. Q. Weinberger (Eds.), Advances in Neural Information Processing Systems 25, Curran Associates, Inc., 2012, pp. 1097–1105 (Accessed 17 December 2018), http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf.

[12] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016) 770–778, https://doi.org/10.1109/CVPR.2016.90.

[13] V. Badrinarayanan, A. Kendall, R. Cipolla, SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation, IEEE Trans. Pattern Anal. Mach. Intell. 39 (2017) 2481–2495, https://doi.org/10.1109/TPAMI.2016.2644615.

[14] J. Long, E. Shelhamer, T. Darrell, Fully convolutional networks for semantic segmentation, 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2015) 3431–3440, https://doi.org/10.1109/CVPR.2015.7298965.

[15] O. Ronneberger, P. Fischer, T. Brox, U-net: convolutional networks for biomedical image segmentation, in: N. Navab, J. Hornegger, W.M. Wells, A.F. Frangi (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI, Springer International Publishing, 2015, pp. 234–241. https://link.springer.com/chapter/10.1007/978-3-319-24574-4_28.

[16] U.R. Acharya, S.L. Oh, Y. Hagiwara, J.H. Tan, M. Adam, A. Gertych, R.S. Tan, A deep convolutional neural network model to classify heartbeats, Comput. Biol. Med. 89 (2017) 389–396, https://doi.org/10.1016/j.compbiomed.2017.08.022.

[17] A. Esteva, B. Kuprel, R.A. Novoa, J. Ko, S.M. Swetter, H.M. Blau, S. Thrun, Dermatologist-level classification of skin cancer with deep neural networks, Nature 542 (2017) 115–118, https://doi.org/10.1038/nature21056.

[18] P. Rajpurkar, J. Irvin, R.L. Ball, K. Zhu, B. Yang, H. Mehta, T. Duan, D. Ding, A. Bagul, C.P. Langlotz, B.N. Patel, K.W. Yeom, K. Shpanskaya, F.G. Blankenberg, J. Seekins, T.J. Amrhein, D.A. Mong, S.S. Halabi, E.J. Zucker, A.Y. Ng, M. P. Lungren, Deep learning for chest radiograph diagnosis: a retrospective comparison of the CheXNeXt algorithm to practicing radiologists, PLoS Med. 15 (2018) e1002686, https://doi.org/10.1371/journal.pmed.1002686.

[19] T. Brosch, J. Peters, A. Groth, T. Stehle, J. Weese, Deep learning-based boundary detection for model-based segmentation with application to MR prostate segmentation, in: A.F. Frangi, J.A. Schnabel, C. Davatzikos, C. Alberola-López, G. Fichtinger (Eds.), Medical Image Computing and Computer Assisted Intervention – MICCAI 2018, Springer International Publishing, 2018, pp. 515–522.

[20] J.D. Fauw, J.R. Ledsam, B. Romera-Paredes, S. Nikolov, N. Tomasev, S. Blackwell, H. Askham, X. Glorot, B. O'Donoghue, D. Visentin, G. van den Driessche, B. Lakshminarayanan, C. Meyer, F. Mackinder, S. Bouton, K. Ayoub, R. Chopra, D. King, A. Karthikesalingam, C.O. Hughes, R. Raine, J. Hughes, D.A. Sim, C. Egan, A. Tufail, H. Montgomery, D. Hassabis, G. Rees, T. Back, P.T. Khaw, M. Suleyman, J. Cornebise, P.A. Keane, O. Ronneberger, Clinically applicable deep learning for diagnosis and referral in retinal disease, Nat. Med. 24 (2018) 1342, https://doi.org/10.1038/s41591-018-0107-6.

[21] K. Kamnitsas, C. Ledig, V.F.J. Newcombe, J.P. Simpson, A.D. Kane, D.K. Menon, D. Rueckert, B. Glocker, Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation, Med. Image Anal. 36 (2017) 61–78, https://doi.org/10.1016/j.media.2016.10.004.

[22] S. Pereira, A. Pinto, V. Alves, C.A. Silva, Brain tumor segmentation using convolutional neural networks in MRI images, IEEE Trans. Med. Imaging 35 (2016) 1240–1251, https://doi.org/10.1109/TMI.2016.2538465.

[23] Y. Song, E. Tan, X. Jiang, J. Cheng, D. Ni, S. Chen, B. Lei, T. Wang, Accurate cervical cell segmentation from overlapping clumps in pap smear images, IEEE Trans. Med. Imaging 36 (2017) 288–300, https://doi.org/10.1109/TMI.2016.2606380.

[24] J. Ker, L. Wang, J. Rao, T. Lim, Deep learning applications in medical image analysis, IEEE Access 6 (2018) 9375–9389, https://doi.org/10.1109/ACCESS.2017.2788044.

[25] G. Litjens, T. Kooi, B.E. Bejnordi, A.A.A. Setio, F. Ciompi, M. Ghafoorian, J.A.W. M. van der Laak, B. van Ginneken, C.I. Sánchez, A survey on deep learning in medical image analysis, Med. Image Anal. 42 (2017) 60–88, https://doi.org/10.1016/j.media.2017.07.005.

[26] R. Tibshirani, Regression shrinkage and selection via the lasso, J. R. Stat. Soc. Ser. B 58 (1996) 267–288.

[27] S. Ioffe, C. Szegedy, Batch normalization: accelerating deep network training by reducing internal covariate shift, International Conference on Machine Learning (2015) 448–456 (Accessed 17 December 2018), http://proceedings.mlr.press/v37/ioffe15.html.

[28] A. Galea, L. Capelo, Applied Deep Learning with Python: Use Scikit-learn, TensorFlow, and Keras to Create Intelligent Systems and Machine Learning Solutions, Packt Publishing, 2018.

[29] D.P. Kingma, J. Ba, Adam: A Method for Stochastic Optimization, ArXiv:1412.6980 [Cs], 2014 (Accessed 17 December 2018), http://arxiv.org/abs/1412.6980.

[30] A. Rosenthal, V. Ntziachristos, D. Razansky, Model-based optoacoustic inversion with arbitrary-shape detectors, Med. Phys. 38 (2011) 4285–4295, https://doi.org/10.1118/1.3589141.

[31] Z. Zhou, M.M. Rahman Siddiquee, N. Tajbakhsh, J. Liang, UNet++: a nested U-net architecture for medical image segmentation, in: D. Stoyanov, Z. Taylor, G. Carneiro, T. Syeda-Mahmood, A. Martel, L. Maier-Hein, J.M.R.S. Tavares, A. Bradley, J.P. Papa, V. Belagiannis, J.C. Nascimento, Z. Lu, S. Conjeti, M. Moradi, H. Greenspan, A. Madabhushi (Eds.), Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support, Springer International Publishing, Cham, 2018, pp. 3–11, https://doi.org/10.1007/978-3-030-00889-5_1.

[32] H. Kervadec, J. Bouchtiba, C. Desrosiers, E. Granger, J. Dolz, I.B. Ayed, Boundary loss for highly unbalanced segmentation, International Conference on Medical Imaging with Deep Learning (2019) 285–296 (Accessed 17 October 2019), http://proceedings.mlr.press/v102/kervadec19a.html.

[33] A. Kalousis, J. Prados, M. Hilario, Stability of feature selection algorithms, Fifth IEEE International Conference on Data Mining (ICDM'05) (2005), https://doi.org/10.1109/ICDM.2005.135, p. 8.

[34] S. Nogueira, K. Sechidis, G. Brown, On the stability of feature selection algorithms, J. Mach. Learn. Res. 18 (2018) 1–54.

[35] N. Chlis, E.S. Bei, M. Zervakis, Introducing a stable bootstrap validation framework for reliable genomic signature extraction, IEEE/ACM Trans. Comput. Biol. Bioinform. 15 (2018) 181–190, https://doi.org/10.1109/TCBB.2016.2633267.

[36] J. Sauvola, M. Pietikäinen, Adaptive document image binarization, Pattern Recognit. 33 (2000) 225–236, https://doi.org/10.1016/S0031-3203(99)00055-2.

[37] N. Otsu, A threshold selection method from gray-level histograms, IEEE Trans. Syst. Man Cybern. 9 (1979) 62–66, https://doi.org/10.1109/TSMC.1979.4310076.

[38] S. van der Walt, J.L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J.D. Warner, N. Yager, E. Gouillart, T. Yu, scikit-image: image processing in Python, PeerJ. 2 (2014) e453, https://doi.org/10.7717/peerj.453.

[39] D. Allman, A. Reiter, M.A.L. Bell, Photoacoustic source detection and reflection artifact removal enabled by deep learning, IEEE Trans. Med. Imaging 37 (2018) 1464–1477, https://doi.org/10.1109/TMI.2018.2829662.

[40] S. Antholzer, M. Haltmeier, J. Schwab, Deep learning for photoacoustic tomography from sparse data, Inverse Probl. Sci. Eng. 0 (2018) 1–19, https://doi.org/10.1080/17415977.2018.1518444.

Angelos Karlas studied Medicine (M.D.) and Electrical and Computer Engineering (Dipl.-Ing.) at the Aristotle University of Thessaloniki, Greece. He holds a Master of Science in Medical Informatics (M.Sc.) from the same university and a Master of Research (M.Res.) in Medical Robotics and Image-Guided Intervention from Imperial College London, UK. He currently works as clinical resident at the Department for Vascular and Endovascular Surgery of the Rechts der Isar University Hospital in Munich, Germany. He holds the position of the Clinical Translation Manager at the Institute for Biological and Medical Imaging of the Helmholtz Center Munich and the Technical University of Munich, Germany, while pursuing his Ph.D. (Dr. rer.nat.) in Experimental Medicine at the same university. He also serves as the Group Leader of the Clinical Bioengineering Group at the Institute for Biological and Medical Imaging of the Technical University of Munich, Germany. His research interests are in the areas of innovative vascular imaging and image-guided vascular interventions.

Vasilis Ntziachristos received his PhD in electrical engineering from the University of Pennsylvania, USA, followed by a postdoctoral position at the Center for Molecular Imaging Research at Harvard Medical School. Afterwards, he became an Instructor and following an Assistant Professor and Director at the Laboratory for Bio-Optics and Molecular Imaging at Harvard University and Massachusetts General Hospital, Boston, USA. Currently, he is the Director of the Institute for Biological and Medical Imaging at the Helmholtz Zentrum in Munich, Germany, as well as a Professor of Electrical Engineering, Professor of Medicine and Chair for Biological Imaging at the Technical University Munich. His work focuses on novel innovative optical and optoacoustic imaging modalities for studying biological processes and dis-eases as well as the translation of these findings into the clinic.