



Titre: Improving microvascular brain analysis with adversarial learning for OCT-TPM vascular domain translation
Title: OCT-TPM vascular domain translation

Auteurs: Nadia Badawi, Jaloliddin Rustamov, Zahiriddin Rustamov, Frédéric Lesage, Nazar Zaki, & Rafat Damseh
Authors: Lesage, Nazar Zaki, & Rafat Damseh

Date: 2025

Type: Article de revue / Article

Référence: Badawi, N., Rustamov, J., Rustamov, Z., Lesage, F., Zaki, N., & Damseh, R. (2025). Improving microvascular brain analysis with adversarial learning for OCT-TPM vascular domain translation. *Scientific Reports*, 15(1), 25182 (13 pages).
Citation: <https://doi.org/10.1038/s41598-025-07410-x>

 **Document en libre accès dans PolyPublie**
Open Access document in PolyPublie

URL de PolyPublie: <https://publications.polymtl.ca/66562/>
PolyPublie URL:

Version: Version officielle de l'éditeur / Published version
Révisé par les pairs / Refereed

Conditions d'utilisation: Creative Commons Attribution-Utilisation non commerciale-Pas d'oeuvre dérivée 4.0 International / Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND)
Terms of Use:

 **Document publié chez l'éditeur officiel**
Document issued by the official publisher

Titre de la revue: Scientific Reports (vol. 15, no. 1)
Journal Title:

Maison d'édition: Nature Portfolio
Publisher:

URL officiel: <https://doi.org/10.1038/s41598-025-07410-x>
Official URL:

Mention légale: This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.
Legal notice:



OPEN Improving microvascular brain analysis with adversarial learning for OCT–TPM vascular domain translation

Nadia Badawi¹, Jaloliddin Rustamov¹, Zahiriddin Rustamov¹, Frederic Lesage², Nazar Zaki¹ & Rafat Damseh¹✉

Modeling microscopic cerebrovascular networks is essential for understanding cerebral blood flow and oxygen transport. High-resolution imaging modalities, such as Optical Coherence Tomography (OCT) and Two-Photon Microscopy (TPM), are widely used to capture microvascular structure and topology. Although TPM angiography generally provides better localization and image quality than OCT, its use is impractical in studies involving fluorescent dye leakage. Here, we exploit generative adversarial learning to produce high-quality TPM angiographies from OCT vascular stacks. We investigate the use of 2D and 3D cycle generative adversarial networks (CycleGANs) trained on unpaired image samples. We evaluate the generated TPM vascular structures based on image similarity and signal-to-noise ratio. Additionally, we evaluated the generated vascular structures after applying vessel segmentation and extracting their 3D topological models. Our results demonstrate that the 2D adversarial learning model outperforms the 3D model in terms of image quality. However, our statistical comparisons of vascular network features show the 3D model's consistent superiority in generating vascular structures. Our work provides a complementary approach to enhance vascular analysis when only OCT imaging is available.

Keywords CycleGAN, Generative adversarial networks, Optical coherence tomography, Two-photon microscopy, Vascular imaging

Optical Coherence Tomography (OCT) and Two-Photon Microscopy (TPM) are two imaging techniques that help study biological tissue anatomy and physiology at the micro-scale level. OCT uses light waves to create cross-sectional images of tissues based on low-coherence interferometry, measuring the time delay of light reflected from various layers within the biological tissue¹. On the other hand, TPM captures the deep-tissue structure based on the near-infrared two-photon excitation principle, where a fluorophore is excited by two photons of lower energy, resulting in a very localized fluorescence emission in the tissue². In brain vascular studies, OCT and TPM imaging have been widely used to characterize the microvascular structure, blood flow, and oxygen consumption. As a result, it provides invaluable insights when studying neurovascular diseases such as stroke and cerebral small vessel disease^{3,4}.

Researchers usually choose between OCT and TPM imaging modalities based on their specific research goals and the tissue layers they need to study. OCT provides non-invasive spatially- and temporally-resolved structural images of brain vascular structure, making it suitable for studying changes in blood vessel morphology and monitoring blood flow patterns⁵. However, their limited depth penetration restricts their ability to capture detailed microvascular structures in deeper brain regions. Its resolution trade-off implies that lower-resolution settings are necessary to improve depth penetration. Conversely, TPM can effectively provide deeper penetration within brain tissues, allowing the ability to visualize cellular and molecular details. TPM is essential for studying the interactions between neurons, glial cells, and blood vessels in the context of neurovascular coupling and other brain-related processes⁶. Some works exploit both modalities to provide a comprehensive understanding of brain vascular physiology and pathology⁴. Experimental studies on microstrokes involve invasive protocols that disrupt blood micro-circulation and are associated with micro-bleeds and tissue infarction⁷. TPM imaging in these scenarios is impractical since fluorescence emission from affected regions will be noisy and can not

¹Department of Computer Science and Software Engineering, United Arab Emirates University, Al Ain, UAE. ²Biomedical Engineering Institute, Ecole Polytechnique de Montreal, Montreal, QC, Canada. ✉email: rdamseh@uaeu.ac.ae

be processed to study shape and geometric changes in microvascular structures. Thus, there is a demand for imaging capabilities that combine OCT's non-invasive imaging attributes with TPM's depth and penetration abilities.

Our study investigates enhancing OCT brain vascular imaging analysis through unsupervised image translation using generative adversarial networks (GANs), specifically by employing the 2D registration generative adversarial network (RegGAN)⁸ and the 3D cycle generative adversarial network (CycleGAN)⁹ with unpaired OCT and TPM image stacks from various mice. It is to be noted that the use of unpaired OCT and TPM volumes in the training phase limits voxel-by-voxel validation between the two modalities. Nevertheless, our proposed method can support future studies focusing on improving OCT-based vascular network analysis, which is essential for understanding hemodynamic response and development, rather than providing voxel-based diagnostic measures.

This work is an extension of our previous work in Badawi et al.¹⁰ that focused only on image-based evaluation. To our knowledge, this is the first work to investigate such a problem in the context of OCT–TPM domain translation with image-based and graph-based validation of the generated vascular stacks. Our work improves image quality and depth resolution by transforming OCT images into TPM-like images, which enables more accurate representations of vascular geometry to help in further applications, e.g., studying blood flow dynamics. To enable proper validation, we have applied deep image segmentation and 3D vascular graphing on the generated TPM angiographies to carry out detailed vascular structure analysis, besides image-to-image similarity assessments. The structure of this paper is as follows: the Related Work section reviews relevant literature; the Methods and Materials section describes the methodology and data used; the Experiments and Results section presents the experimental setup and findings; and the Conclusion section summarizes the key outcomes.

Related work

GAN is a deep learning architecture comprising two competing neural networks: a generator and a discriminator. The generator creates new data instances that resemble the training data, while the discriminator tries to distinguish between real and generated data. Through this adversarial process, the generator learns to produce increasingly realistic data, and the discriminator becomes better at identifying fake data. GANs have found various applications in medical image computing, particularly in tasks involving unpaired images. Unpaired images refer to scenarios with no direct one-to-one correspondence between images in different domains, such as MRI and CT scans of the same patient. GANs can learn the mapping between these domains, enabling tasks like image-to-image translation, cross-modality synthesis, and data augmentation.

Image-to-image translation

Nie et al.¹¹ proposed a GAN-based approach for medical image synthesis using unpaired images, introducing CycleGAN to map between two domains without paired data. This method demonstrated effectiveness across various medical imaging modalities, including MRI, CT, and PET. Wolterink et al.⁹ leveraged GANs to generate synthetic CT images from unpaired MRI data, enabling realistic CT image generation without ionizing radiation exposure. These studies underscored the potential of GANs for realistic image synthesis in medical imaging.

Cross-modality synthesis

Zhang et al.¹² developed a dual-cycle GAN for cross-modality medical image synthesis, learning bidirectional mappings between domains like MRI and PET, and enhancing brain tumor segmentation performance. Huang et al.¹³ employed a multi-modal GAN to synthesize missing PET images from MRI data using a 3D U-Net generator and 3D PatchGAN discriminator, improving brain tumor segmentation tasks. These works illustrate the efficacy of GANs in cross-modality synthesis, providing complementary imaging data to enhance diagnostic accuracy.

Data augmentation

Dar et al.¹⁴ utilized CycleGAN for data augmentation in medical imaging, generating synthetic MRI images from unpaired CT scans to bolster training data for brain tumor segmentation tasks. Shin et al.¹⁵ employed GANs for data augmentation in chest X-ray images, generating synthetic images with different abnormalities to enhance deep-learning models for disease classification. These studies demonstrate the potential of GANs to generate diverse training datasets, improving the robustness of machine-learning models in medical imaging.

Unsupervised domain adaptation

Yi et al.¹⁶ proposed a GAN-based method for unsupervised cross-domain adaptation, using CycleGAN to translate images from source domains (e.g., MRI) to target domains (e.g., CT) while preserving semantic content. Chen et al.¹⁷ developed a GAN-based framework for unsupervised domain adaptation in medical image segmentation, translating images between domains and learning domain-invariant features. These approaches have shown improved segmentation performance compared to traditional methods, highlighting the adaptability of GANs in various medical imaging applications. While these studies demonstrate the potential of GANs in addressing various challenges in medical image analysis, to our knowledge, no prior work has investigated OCT–TPM domain translation using 3D and 2D adversarial networks. Our study aims to bridge this gap by exploiting generative adversarial learning to produce high-quality TPM angiographies from OCT vascular stacks, thereby providing a complementary approach to enhance vascular analysis when only OCT imaging is involved.

Methods and materials

This section details the methodology employed in our work, beginning with the initial stages of data collection and preparation, the techniques used for image generation and analysis, and the procedures we followed for methodology validation.

Data collection and preparation

To train our 2D and 3D generative models, we utilized five OCT volumes previously collected in¹⁸ from wild-type C57BL/6 mice (3 male; 2 female; twelve-week-old). The acquisitions were made using a custom spectral-domain OCT system integrated with optical-intrinsic-signal imaging (OISI): a broadband 1310 nm super-luminescent diode (SLD; LS2000C, Thorlabs), a 10× long-working-distance objective, and a Cobra-1300 spectrometer. Original OCT volumes were of size $512 \times 512 \times 150$ at voxel resolution of $\approx 2.3 \times 2.3 \times 8.3 \mu\text{m}$. We extracted 10 non-overlapping sub-volumes of size $256 \times 256 \times 150$ from the original OCT stacks for model training. Our TPM samples were previously used in^{19–21}. After tail vein injection of dextranconjugated fluorescein (FITC) with a concentration of 500 nM, the TPM stacks were acquired from six male C57BL/6 mice (25–30 g) using a custom-built two-photon microscope with a 20X Olympus objective (NA = 0.95), capturing images at voxel dimensions of $\approx 1.2 \times 1.2 \times 2.0 \mu\text{m}$. To match the number and size of OCT training stacks, we extracted 10 non-overlapping subvolumes of TPM from the original TPM data, which had image sizes ranging from $525 - 604 \times 529 - 604 \times 301 - 341$ voxels. To train the 2D RegGAN and 3D CycleGAN networks, we split our volumetric data into $\approx 0.7 - 0.3$ ratio for training and testing, with training patch sizes of 256×256 and $128 \times 128 \times 64$, respectively. The testing sub-volumes were obtained from one animal, which was excluded from the network training process. To ensure that training patches have the same spatial resolution, a third-order B-spline interpolation with Gaussian anti-aliasing was used to upsample OCT patches to match that of the TPM counterpart. In addition to the raw TPM samples that were used to train our generative AI models, 3D TPM vascular labels made in²² were utilized to train our segmentation model, as explained in the next section. There was no reported inter-rater reliability because the labels were created by a single annotator. All experiments were performed in accordance with relevant guidelines and regulations. All procedures and handling were also approved by the Animal Research Ethics Committee of the Montreal Heart Institute.

Methodological pipeline

Our methodological pipeline for improving OCT-based microvascular analysis, as shown in Fig. 1, comprises three key stages: generating images through OCT–TPM domain translation using 2D and 3D generative models, i.e., RegGAN and CycleGAN, applying deep segmentation models to extract binary maps of vascular vessels, and then using computational geometry methods for creating vascular graph representations for further analysis.

Image generation

To convert the preprocessed OCT stacks into TPM stacks, we employed 2D RegGAN and 3D CycleGAN as generative deep learning architectures for unpaired image-to-image translation, allowing the conversion of OCT images to TPM-like images without the need for paired training data.

RegGAN⁸ is designed with a registration network added as a noise layer to GANs to correct the output results, improve stability and performance in image-to-image translation, and preserve anatomical properties in

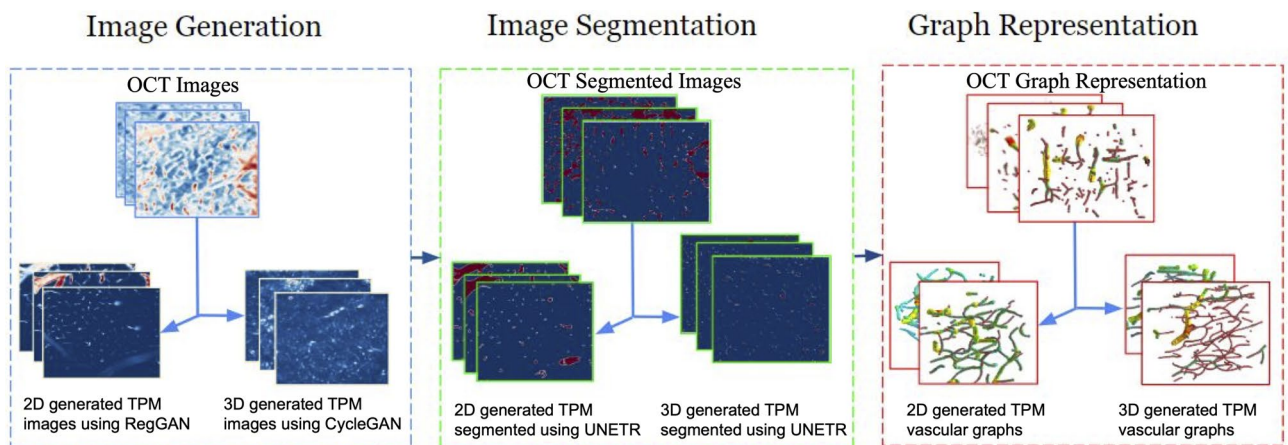


Fig. 1. Our methodological pipeline for OCT–TPM domain translation and vascular analysis. The pipeline begins with the acquisition of OCT volume stacks from biological tissue. The preprocessed OCT volumes are then input into two generative adversarial network models: a 2D RegGAN for slice-based image translation and a 3D CycleGAN for volumetric translation. The models generate corresponding TPM-like angiographies from the OCT images. After generating the TPM images, a deep segmentation model (i.e., UNETR) is applied to extract vascular structures from the translated images, yielding binary segmentation maps. These maps are further processed to construct vascular graphs, representing the microvascular network’s topology and geometry for advanced structural and functional analysis.

the generation process. Hence, an additional loss term was introduced to the GAN objective to guide the training process. The used RegGAN model architecture is shown in Fig. 2. The model, which includes a generator and a discriminator with the Leaky ReLU activation function, learns to map input images to output images while rectifying spatial discrepancies. The model accepts a source domain image with distribution X (OCT), and a target domain image with distribution Y (TPM), into a generator G that maps 2D OCT images to 2D TPM images while compensating for training data misalignment and learning to adjust for spatial inconsistencies. The following equations present the formulation of the GAN losses in the used model RegGAN, Eq. (1) expresses the correlation loss (L_{Corr}) where R is the registration layer after the generator G .

$$\min_{G,R} L_{\text{Corr}}(G, R) = \mathbb{E}_{x,y_e} [\|y_e - G(x) \circ R(G(x), y_e)\|_2] \quad (1)$$



Fig. 2. Architecture of the 2D RegGAN model for OCT-TPM image translation. The 2D RegGAN model consists of a generator that transforms OCT images into TPM-like images and a discriminator that refines image quality through adversarial learning. A registration network, R , aligns the output images with anatomical consistency by correcting spatial misalignments. The generator uses a U-Net structure with convolutional layers, while the registration network introduces a deformation field to ensure high-quality, spatially aligned TPM images from OCT slices. X : Source OCT domain. Y : Target TPM domain. G : Generator mapping OCT \rightarrow TPM. D_Y : Discriminator for the TPM domain. R : Registration network for spatial alignment.

where $R(G(x), y_e)$ is the deformation field and \circ represents the resampling operation. The registration network is based on U-Net²³. A smoothness loss (L_{Smooth})²⁴ is defined in Eq. (2), to evaluate the smoothness of the deformation field and minimize the gradient of the deformation field.

$$\min_R L_{Smooth}(R) = \mathbb{E}_{x, y_e} [\|\nabla R(G(x), y_e)\|_2] \tag{2}$$

in RegGAN, the adversarial loss (L_{adv}) is added between the generator and the discriminator, and the total loss (L_{Total}) is expressed in Eq. (3).

$$\min_{G, R} \max_D L_{Total}(G, R, D) = L_{Corr} + L_{Smooth} + L_{Adv} \tag{3}$$

On the other hand, Fig. 3 presents the architecture of CycleGAN designed for unsupervised image-to-image translation between two unpaired domains, X and Y. It consists of two generators ($G : X \rightarrow Y$ and $F : Y \rightarrow X$) and two discriminators (D_X and D_Y). The key idea is to learn mappings G and F such that $G(X) \approx Y$ and $F(Y) \approx X$. CycleGAN aims to combine the adversarial losses for both mappings with a cycle consistency loss to ensure that an image can be translated from one domain to the other and back again with minimal loss of content. The generator 'G' is the component trained to transform images from the X domain into the Y domain, and the generator 'F' does the opposite, translating images from domain Y back into domain X. CycleGAN can forge these domain translations without the need for paired data, using cycle consistency loss. The loss function acts like a quality check, ensuring that despite training on unpaired data, the model can still retain the essential attributes of the input images when it performs the domain transformations, creating a reliable and coherent translation between the domains for unpaired datasets. The following equations represent the formulation of the model. The adversarial losses for generators G and F are defined as follows²⁵:

$$\mathcal{L}_{GAN}(G, D_Y, X, Y) = \mathbb{E}_{y \sim p_{data}(y)} [\log D_Y(y)] + \mathbb{E}_{x \sim p_{data}(x)} [\log(1 - D_Y(G(x)))] \tag{4}$$

$$\mathcal{L}_{GAN}(F, D_X, Y, X) = \mathbb{E}_{x \sim p_{data}(x)} [\log D_X(x)] + \mathbb{E}_{y \sim p_{data}(y)} [\log(1 - D_X(F(y)))] \tag{5}$$

$$\mathcal{L}_{cycle}(G, F) = \mathbb{E}_{x \sim p_{data}(x)} [\|F(G(x)) - x\|_1] + \mathbb{E}_{y \sim p_{data}(y)} [\|G(F(y)) - y\|_1] \tag{6}$$

$$\mathcal{L}(G, F, D_X, D_Y) = \mathcal{L}_{GAN}(G, D_Y, X, Y) + \mathcal{L}_{GAN}(F, D_X, Y, X) + \lambda \mathcal{L}_{cycle}(G, F) \tag{7}$$

where λ is a weighting factor for the regularization term, which would depend on the goals of the regularization (e.g., improving image quality, ensuring diversity in the generated images, or stabilizing the training process).

Vascular segmentation

To represent the vascular structures in the generated TPM images and study them further, we employed UNETR²⁶, a transformer model based on the UNet architecture designed for the segmentation of medical images. The distinctive choice of UNETR is attributed to its proficiency in delineating intricate 3D patterns and features found in vascular medical images²⁷. The UNETR model integrates the Transformer's self-attention

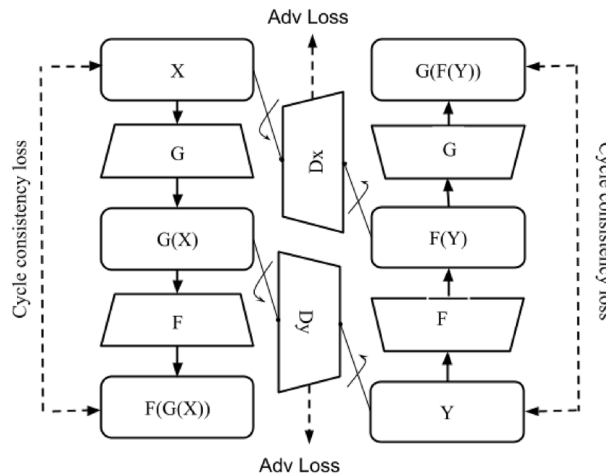


Fig. 3. Architecture of the 3D CycleGAN model for OCT-TPM image translation. The 3D CycleGAN model facilitates unsupervised volumetric image translation between OCT and TPM domains. It includes two generators, each translating between the two image domains, and two discriminators that refine image quality by distinguishing real from generated images. The model incorporates cycle consistency loss to ensure the OCT-TPM translation and its inverse retain key structural features. This architecture allows for efficient 3D image-to-image translation without requiring paired datasets. X: Source OCT domain. Y: Target TPM domain. G: Generator mapping $X \rightarrow Y$. F: Generator mapping $Y \rightarrow X$. D_X : Discriminator for OCT domain. D_Y : Discriminator for TPM domain. λ : Cycle-consistency weight in Eq. (7).

mechanisms with the UNet's structured framework to enhance medical image segmentation. It uniquely weights different parts of the input data through the self-attention equation²⁸:

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V \quad (8)$$

Q, K, and V denote the dimensions of the keys. The model's architecture progresses from an encoder utilizing Transformer blocks for extracting high-level features through a bottleneck for feature refinement to a decoder that enriches feature maps for precise segmentation. It employs skip connections to enhance feature propagation and combines them, ensuring the recovery of spatial details through concatenated features. The approach is further solidified by employing a dual-loss function to address both overlap and classification accuracy.

Vascular graphing

We utilized VascGraph¹⁹ to represent the generated and segmented TPM stacks as geometry-based vascular graphs. To generation of such a computational graph begins with the construction of an initial geometric graph within the boundaries defined by a binary vascular mask. This initial graph is composed of nodes, which are placed at every voxel within the vascular structure that has a value of '1' (indicating the presence of vascular tissue). The edges between nodes are established based on their immediate spatial relationships, typically using a 6-connectivity scheme, which considers neighboring voxels that share a face. Laplacian optimization is then applied with proper affinity weights across graph edges that influence the contraction of the graph towards the vascular centerlines during an iterative process, by minimizing the distance between the nodes and the centerline. This process effectively refines the graph to more closely follow the actual vascular paths. Each iteration of the process adjusts the positions of the nodes based on a calculated Laplacian matrix until a convergence criterion is met, indicating that the graph adequately represents the underlying vascular structure. Finally, a decimation step reduces the graph complexity to resemble skeleton-like structures while preserving critical topological and geometric features.

Evaluation procedure of the proposed pipeline

The evaluation of the proposed pipeline was carried out in two phases: image quality assessment and vascular structure analysis. For image quality, we used metrics such as Mean Absolute Error (MAE), SSIM, and Peak Signal-to-Noise Ratio (PSNR) to quantify the similarity between the generated TPM images and the original TPM samples. MAE measures the average magnitude of errors between the generated and reference images, with lower values indicating better similarity. The SSIM is a perception-based metric that considers changes in structural information when comparing the quality of the real and generated images. Values range from -1 to 1, with 1 indicating perfect similarity. PSNR, which is usually used to measure the quality of reconstruction of lossy compression, was also employed in our analysis. Lower MAE and higher SSIM and PSNR scores demonstrated the success of the models in generating accurate and high-quality TPM-like angiographies from OCT images. These metrics provided essential insights into the structural fidelity and noise reduction in the translated images. To further evaluate the pipeline, we performed vascular structure analysis by segmenting the generated TPM images and extracting key morphometric properties, such as vessel length, radius, volume, and tortuosity. We then represented the segmented vessels as graphs to study the topology of the vascular networks, allowing for a detailed comparison between the original OCT data and the generated TPM data. Paired Wilcoxon signed-rank tests were performed for each morphometric parameter (tortuosity, length, area and volume), comparing the original TPM vascular graphs to those extracted from OCT, 2D RegGAN, and 3D CycleGAN outputs.

Experiments and results

This section presents the experimental setup, training process, and evaluation results for the 2D RegGAN and the 3D CycleGAN models used in our study. All the experiments were performed on an HP workstation with the following specifications: Intel(R) Xeon (R) Gold 5220R CPU @ 2.20 GHz, RTX A5500 GPU with 24 GB of memory and 128GB of DDR4 memory.

Image generation training process

To train the 2D RegGAN model, we used OCT slices of size 256×256 and employed an Adam optimizer with a 0.0001 learning rate, decreasing linearly post the 20th epoch until epoch 200. The loss function combined correction loss (with a weight of $\lambda_{\text{corr}} = 20$) for image alignment, smoothness loss (weight of $\lambda_{\text{smooth}} = 10$) for deformation field continuity, and adversarial loss (weight of $\lambda_{\text{adv}} = 1$). These wight values were empirically chosen as done in⁸. For the translation process, the generator provides outputs based on the ReLU activation function, and the discriminator D differentiates between real and generated data based on the Leaky ReLU activation function. The discriminator provides critical feedback to the generator and propels the training process forward. The registration network R aligns the 256×256 pixel images, ensuring anatomical alignments. Network enhancements feature a spatial transformer in the final layer for output normalization and a dropout layer with a 50% rate before the last \tanh activation function for standardized outputs⁸. To train the 3D CycleGAN model, we utilized the OCT volume sized $256 \times 256 \times 150$ and preprocessed our training sample using SimpleITK for transformation to 8-bit images, edge enhancement, and contrast adjustment. The value of λ was set to 10 as suggested by previous works that used CycleGAN^{9,29}, given our 256-grayscale input voxel values. Furthermore, we performed augmentation by applying random noise, rotation, and affine deformations on the samples. Optimization was executed using Adam optimizer with initial learning rates of 0.0002, halved over 400 epochs. The generators were tasked with preserving structural integrity across modalities, leveraging cycle-

consistency loss and attention mechanisms for detailed medical imaging. The 3D CycleGAN model consists of a generator and discriminator duo, each tailored for 3D data processing. The generator, constructed on a nine-block ResNet framework, integrates Conv3d and BatchNorm3d for its convolutional and normalization operations, specifically for volumetric data. The architecture is modular, with sequential phases of down sampling and up sampling, enriched with skip connections that ensure the fidelity of intricate details from input to output. A series of ResNet blocks forms the backbone of the generator, which is essential for enabling the training of deeper networks by preventing vanishing gradients. These blocks are adept at learning identity mappings, which are necessary for adding incremental information and avoiding the degradation often associated with training deep networks. The generator utilizes ReLU activation functions to introduce non-linearity without affecting the sparse structure of the data. On the other hand, the discriminators employ a sigmoid activation function, which is valuable for classifying the originality of the generated images, providing a probabilistic output that discriminates real from generated images.

Segmentation training process and graph generation

To study the improvement of the TPM-generated images over the OCT input images, we employed the UNETR model²⁶ to perform segmentation and trained it using our 10 TPM stacks with their vascular labels (70% of the data was used for training and the remaining 30% was used for testing). The model incorporates input channels set to 1 and output channels set to 1 and specifies an image size of $64 \times 64 \times 64$. The architecture utilizes a feature size of 16, with a hidden size of 768, an MLP dimension of 3072, and 32 attention heads for processing. It employs instance normalization, includes a residual block, and incorporates a dropout rate of 0.2. This configuration enabled the model to segment vascular structures accurately by learning from complex spatial patterns in volumetric images. To further refine the segmentation output, we applied a post-processing sequence consisting of a sigmoid activation function and a discrete conversion of the sigmoid output based on a threshold value of 0.5, ensuring binary classification of pixels into vascular or non-vascular regions. The UNETR model is optimized using Dice Loss and the Adam optimizer. Training is conducted over 300 epochs with a batch size of 5, employing data loaders for efficient data handling. Validation is performed at regular intervals using the Dice metric to assess model performance, where the best-performing model update is saved for future use. After training, we evaluated our UNETR segmentation model on the real TPM volumes using standard metrics, i.e., Dice coefficient and an Intersection-over-Union (IoU). On the held-out testing set, our model, trained with Dice loss and post-processed with a 0.5 threshold, achieved a mean Dice score of 0.64 ± 0.03 and a mean IoU of 0.52 ± 0.04 . Following the segmentation, VascGraph algorithm, described in the Methods section, was applied to convert the segmented vascular structures into geometric graphs that represent the connectivity patterns of the microvessels. The resulting vascular graphs were used to calculate several parameters used in our analysis, including vessel length, radius, volume, and tortuosity. These parameters were statistically compared with the original OCT data, generated TPM samples and original TPM data. The resulting geometric graphs were also visualized to assess and analyze the vascular network's structure.

Analysis of generated images

Figure 4 presents the evaluation of the performance of RegGAN on generating 2D slices and CycleGAN on generating 3D TPM volumes. The boxplots show that the 2D model has a significantly lower MAE, 0.2938, compared to the 3D model, 0.6338. Again, the 2D model has a higher SSIM, 0.76208, compared to the 3D model, 0.5196, which implies that the 2D model preserves structural information more effectively. In terms of PSNR, the 2D model achieved a higher PSNR, 23.78, than the 3D model, 19.26, indicating that the 2D model yields images with higher quality and less noise. The 2D model's MAE has a much lower standard deviation, 0.15595, than the 3D model, 4.14532, showing that the 2D model's performance is more consistent across different images. Comparable standard deviations of the PSNR, 1.78772 and 1.81812, respectively, were observed. For the SSIM measures, the 2D model has a higher standard deviation of 0.07361 compared to 0.04279 for the 3D model, suggesting that the 3D model has more consistent structural similarity. Figure 5 shows a side-by-side qualitative analysis from one of the testing OCT volumes and their segmentation, the slices at different depth levels of 0,

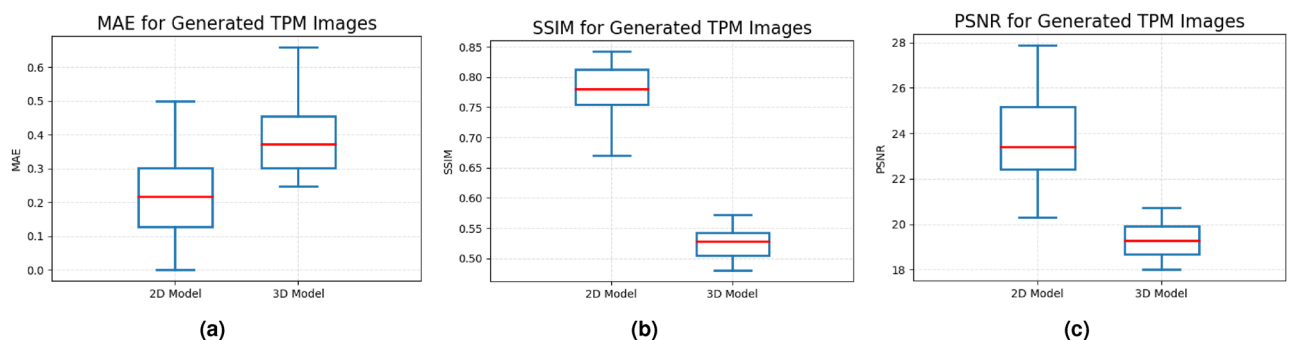


Fig. 4. Quantitative assessment of the used 2D RegGAN and 3D CycleGAN networks on generating TPM vascular stack from acquired OCT volumes. We observed that 2D RegGAN generated TPM slices with lower MAE and higher SSIM and PSNR values compared to that obtained from the 3D CycleGAN. Reprinted from Badawi et al.¹⁰, with permission from ©IEEE.




Fig. 5. The first row shows random slices taken from an OCT stack as our input data. The second row shows the corresponding segmentations obtained.




Fig. 6. The first row shows random slices taken from a TPM volume generated using the 3D CycleGAN model. The second row shows the corresponding segmentations of the generated TPM slices.




Fig. 7. The first row shows random slices taken from a TPM volume generated using the 2D RegGAN model. The second row shows the corresponding segmentations of the generated TPM slices.

50, 100, 145, and 200, downward. Figure 6 shows the 2D-based generated TPM stacks and their segmentation, whereas Fig. 7 shows the 3D-based generated volumes. As noticed from Figs. 6 and 7, the 2D model outperforms the 3D model regarding the quality and level of detail in the generated TPM volumes.

Assessment of segmentation and graph representation

Our study investigated vascular graph representations by examining various parameters, including vascular segments' lengths, average radius, volumes, cross-sectional areas, and tortuosities. Such analysis is conducted to understand the differences in vascular networks' structural characteristics when comparing original OCT and

generated TPM volumes. To compute these parameters, we first identified segments within the vascular network based on junction nodes, points where three or more vessels intersect. We calculated each vessel's length, average radius, volume, and cross-sectional area by considering the positions and radii of the segment's nodes. The volume and area of each segment were estimated by treating each segment as a cylindrical structure, allowing for the approximation of these properties based on the radii and the distance between consecutive nodes. Tortuosity, a measure of the deviation of a path from a straight line, was calculated for each segment as the ratio of the path length to the straight-line distance between the segment's endpoints within the vessel. Statistical analysis compared morphological parameters between the vascular networks in the original OCT volumes and the generated TPM versions. We explored the distribution of each parameter by calculating the Probability Density Function (PDF) and Cumulative Distribution Function (CDF). The PDF provides a visual representation of the probability distribution of a continuous random variable, in this case, each vascular parameter, allowing us to observe the likelihood of different values within the original and the generated samples. The CDF, conversely, represents the probability that a random variable takes on a value less than or equal to a specific value, offering insights into the distribution's cumulative properties. These distributions were generated using the SciPy library's stats.norm module, which assumes a normal distribution of the data and allows for estimating the PDF and CDF based on each parameter's mean and standard deviation.

Morphometric comparisons

In our study, we sought to determine whether there were statistically significant differences between the original OCT stack and the generated stacks using CycleGAN and RegGAN after applying the segmentation on both of them across various geometric properties. Additionally, we used as a reference the average TPM, which is the average of those parameters derived from original TPM stacks. Figure 8 shows boxplot-based statistical analyses comparing the geometric parameters calculated from the original OCT, 2D-based generated TPM, 3D-based generated TPM, and the average TPM. The PDF and CDF plots in Fig. 9 show the distribution of each parameter. From the boxplots of vascular lengths, cross-sectional areas, and volumes in Fig. 8, the use of generated TPM samples shows a closer vascular geometry to the average TPM examples than using the original OCT, with the mean values being closer to the average TPM. The fact that the original OCT appears to have a wider range and larger interquartile range suggests that using it for vascular assessments would produce an inconsistent vascular structure compared to using the generated TPM samples. For the case of assessing vascular tortuosity, the boxplot in Fig. 8 suggests that the use of generated TPM samples has a tighter interquartile range closer to that of the Average TPM. However, it produces more tortuous segments. The PDF and CDF distributions in Fig. 9 confirm our observation. Moreover, they highlight the efficiency of using our approach for vascular analysis based on generated TPM samples instead of original OCT versions, especially in capturing better vascular segment lengths. Table 1 reports a Wilcoxon signed-rank comparison (Bonferroni-corrected; $\alpha = 0.0167$) of the average TPM baseline against the three other counterparts, i.e., OCT, 2D RegGAN, 3D CycleGAN, based on their calculated morphometric features. For vessel tortuosity, OCT exhibited the highest adjusted p -value ($p_{\text{adj}} = 0.915$), showing better preservation of vessel curvature in the raw input. The 2D RegGAN produced the highest deviation ($p_{\text{adj}} = 0.0111$), possibly reflecting oversimplification or distortion of vessel geometry in its 2D-based transformation process. For vessel length, the original OCT input differed

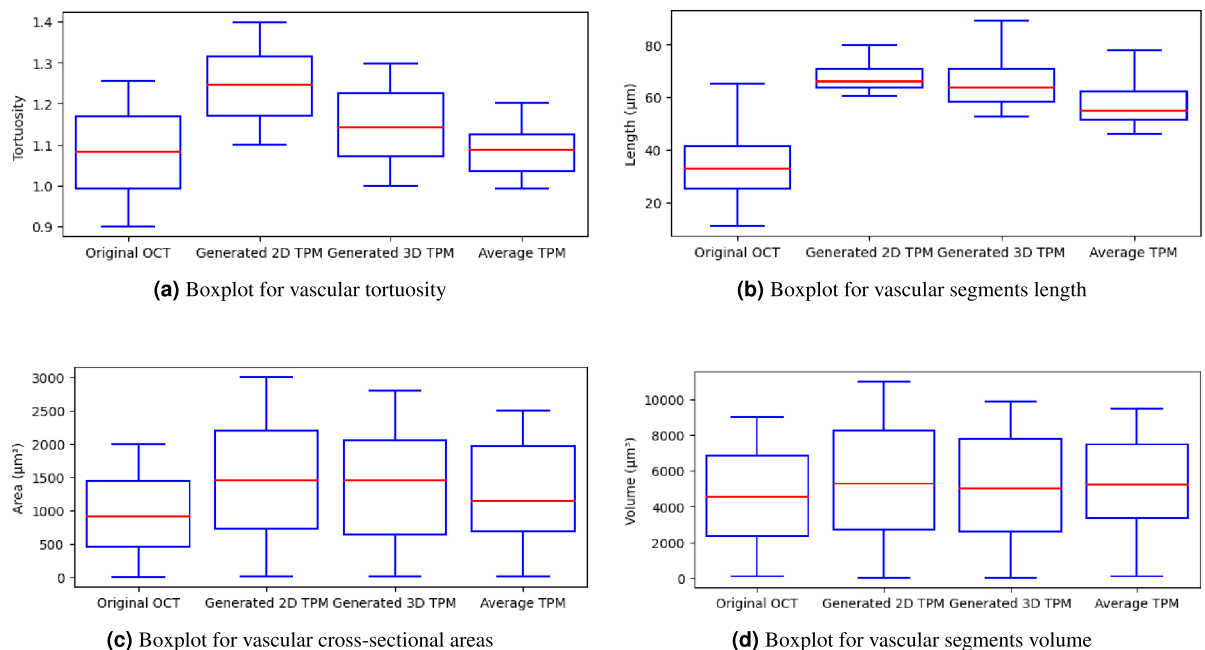


Fig. 8. Boxplots comparing the morphometric properties of the vessels network of an OCT and its TPM generated versions using CycleGAN and RegGAN models.

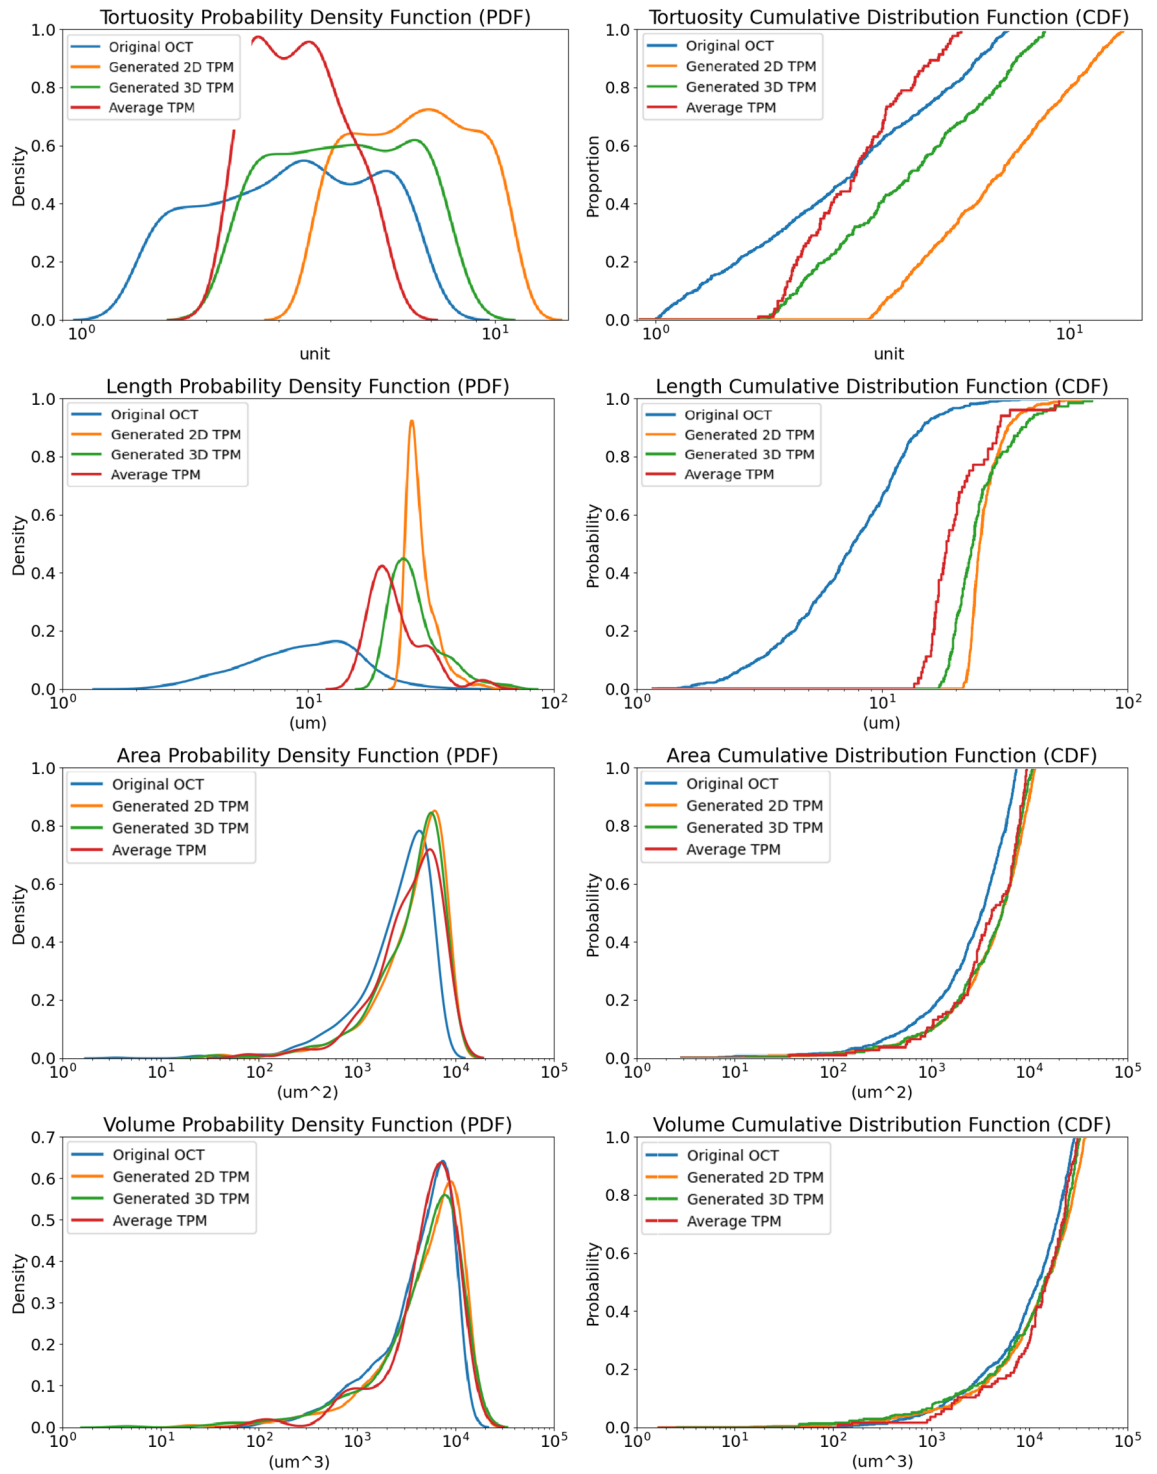


Fig. 9. Comparison of vascular graph morphometric properties obtained from original OCT stack and the two TPM generated stacks using 2D RegGAN and 3D CycleGAN models. The vascular graphs extracted from generated TPM stacks matched the average microvascular properties more closely than those obtained from processing raw OCT.

significantly from the baseline ($p_{adj} = 5.87 \times 10^{-4}$). On the other hand, generated TPM samples showed less deviation, with 3D CycleGAN vascular networks achieving the closest match ($p_{adj} = 0.880$). In terms of vessel cross-sectional area and volume, none of the models deviated significantly from the average TPM reference. Nevertheless, 2D RegGAN outputs showed the most consistent results for both features (area: $p_{adj} = 0.627$; volume: $p_{adj} = 0.978$). Figure 10 provides a visual inspection of the quality of vascular structures produced from 2D-based and 3D-based generated TPM stacks compared to that obtained from the original OCT counterparts.

Parameter	Comparison	Raw p -value	P_{adj}
Tortuosity	Avg. TPM vs OCT	3.05×10^{-1}	9.15×10^{-1}
	Avg. TPM vs 2D RegGAN	3.71×10^{-3}	<u>1.11×10^{-2}</u>
	Avg. TPM vs 3D CycleGAN	5.92×10^{-2}	1.77×10^{-1}
Length	Avg. TPM vs OCT	1.96×10^{-4}	<u>5.87×10^{-4}</u>
	Avg. TPM vs 2D RegGAN	1.13×10^{-2}	3.40×10^{-2}
	Avg. TPM vs 3D CycleGAN	2.89×10^{-2}	8.80×10^{-2}
Area	Avg. TPM vs OCT	5.07×10^{-2}	1.521×10^{-1}
	Avg. TPM vs 2D RegGAN	2.09×10^{-1}	6.27×10^{-1}
	Avg. TPM vs 3D CycleGAN	4.630×10^{-2}	1.39×10^{-1}
Volume	Avg. TPM vs OCT	5.70×10^{-2}	1.71×10^{-1}
	Avg. TPM vs 2D RegGAN	3.26×10^{-1}	9.78×10^{-1}
	Avg. TPM vs 3D CycleGAN	1.12×10^{-1}	3.34×10^{-1}

Table 1. Paired Wilcoxon signed-rank tests between the Average-TPM distribution and each alternative (OCT, 2D RegGAN, 3D CycleGAN). Raw and Bonferroni-adjusted p -values are reported. The adjusted significance level is $\alpha_{adj} = 0.0167$. The largest p_{adj} (closest match) is in **bold**; values below α_{adj} are underlined.




Fig. 10. Examples of vascular graph representations across three different regions of interest (i, ii, iii). The first row shows the vascular graph derived from the original OCT data. The second row represents the TPM volumes generated using the 2D RegGAN model, while the third row illustrates the TPM volumes generated using the 3D CycleGAN model.

It is seen that the 3D-generated TPM volumes provide more consistent outputs. This could be attributed to the fact that the 3D CycleGAN processes volumetric inputs that can bring consistent vascular connectivity compared to the 2D RegGAN model that generates TPM samples slice by slice. On the other hand, the original OCT stacks are shown to be not suitable for further vascular processing and assessments.

Conclusion

This paper investigated the effectiveness of 2D and 3D adversarial networks in generating high-quality vascular TPM volumes from OCT inputs. We used the 2D RegGAN model for slice-based image translation, whereas the 3D CycleGAN model was employed for volumetric image translation. Our qualitative and quantitative assessments were performed to examine the generated TPM images and the associated vascular geometry features compared to those of the original TPM volumes. Based on various image-quality metrics, i.e., MAE, SSIM, and PSNR, the 2D RegGAN produced better outputs compared to those generated with the 3D CycleGAN model. Nevertheless, the use of the 3D CycleGAN model provided more consistent vascular geometry in terms of vascular segment lengths. Our work suggests that generative adversarial networks can help as a complementary

vascular analysis tool when only the OCT modality is involved. Such integration can enhance the extraction of more accurate vascular parameters, e.g., cerebral blood volume and vascular density. One potential avenue to explore is integrating a 3D registration block into the existing 3D CycleGAN architecture. However, careful consideration must be given to the design of this block to ensure computational efficiency and avoid excessive overhead. Additionally, further research could focus on optimizing data collection and curation processes to enhance the quality and diversity of the training datasets to include higher vascular variabilities, which is important to improve model generalizability.

Data availability

The datasets generated and/or analyzed during the current study are not publicly available due to ongoing research efforts, but are available from the corresponding author on reasonable request.

Received: 8 September 2024; Accepted: 16 June 2025

Published online: 12 July 2025

References

- Fujimoto, J., Pitris, C., Boppart, S. & Brezinski, M. Optical coherence tomography: An emerging technology for biomedical imaging and optical biopsy. *Neoplasia* **2**, 9–25. <https://doi.org/10.1038/sj.neo.7900071> (2000).
- Helmchen, F. & Denk, W. Deep tissue two-photon microscopy. *Nat. Methods* **2**, 932–940 (2005).
- Walek, K. W. et al. Near-lifespan longitudinal tracking of brain microvascular morphology, topology, and flow in male mice. *Nat. Commun.* **14**, 2982 (2023).
- Marchand, P. J., Lu, X., Zhang, C. & Lesage, F. Characterization of RBC flux and velocity measurements in oct with concurrent two-photon microscopy. In *Optical Coherence Tomography and Coherence Domain Optical Methods in Biomedicine XXV*, vol. 11630, 28–35 (SPIE, 2021).
- Marchand, P. J., Lu, X., Zhang, C. & Lesage, F. Validation of red blood cell flux and velocity estimations based on optical coherence tomography intensity fluctuations. *Sci. Rep.* **10**, 19584 (2020).
- Shih, A. Y. et al. Two-photon microscopy as a tool to study blood flow and neurovascular coupling in the rodent brain. *J. Cereb. Blood Flow Metab.* **32**, 1277–1309 (2012).
- Damseh, R. et al. A simulation study investigating potential diffusion-based MRI signatures of microstrokes. *Sci. Rep.* **11**, 14229 (2021).
- Kong, L. et al. Breaking the dilemma of medical image-to-image translation. *Adv. Neural Inf. Process. Syst.* **34**, 1964–1978 (2021).
- Wolterink, J. M. et al. Deep MR to CT synthesis using unpaired data. In *Simulation and Synthesis in Medical Imaging: Second International Workshop, SASHIMI 2017, Held in Conjunction with MICCAI 2017, Québec City, QC, Canada, September 10, 2017, Proceedings 2*, 14–23 (Springer, 2017).
- Badawi, N. et al. Generative adversarial learning for OCT-TPM vascular domain translation. In *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*, 1–5 (IEEE, 2024).
- Nie, D. et al. Medical image synthesis with deep convolutional adversarial networks. *IEEE Trans. Biomed. Eng.* **65**, 2720–2730 (2018).
- Zhang, Z., Yang, L. & Zheng, Y. Translating and segmenting multimodal medical volumes with cycle-and shape-consistency generative adversarial network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 9242–9251 (2018).
- Huang, Y., Shao, L. & Frangi, A. F. Simultaneous super-resolution and cross-modality synthesis of 3d medical images using weakly-supervised joint convolutional sparse coding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6070–6079 (2017).
- Dar, S. U. et al. Image synthesis in multi-contrast MRI with conditional generative adversarial networks. *IEEE Trans. Med. Imaging* **38**, 2375–2388 (2019).
- Shin, H.-C. et al. Medical image synthesis for data augmentation and anonymization using generative adversarial networks. In *Simulation and Synthesis in Medical Imaging: Third International Workshop, SASHIMI 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Proceedings 3*, 1–11 (Springer, 2018).
- Yi, X., Walia, E. & Babyn, P. Unsupervised and semi-supervised learning with categorical generative adversarial networks assisted by wasserstein distance for dermoscopy image classification. Preprint at [arXiv:1804.03700](https://arxiv.org/abs/1804.03700) (2018).
- Chen, C., Dou, Q., Chen, H. & Heng, P.-A. Semantic-aware generative adversarial nets for unsupervised domain adaptation in chest x-ray segmentation. In *Machine Learning in Medical Imaging: 9th International Workshop, MLMI 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Proceedings 9*, 143–151 (Springer, 2018).
- Zhang, C., Jamshidi, M., Delafontaine-Martel, P., Linninger, A. A. & Lesage, F. Evaluation of cerebral microcirculation in a mouse model of systemic inflammation. *Neurophotonics* **11**, 035003–035003 (2024).
- Damseh, R., Delafontaine-Martel, P., Pouliot, P., Cherié, F. & Lesage, F. Laplacian flow dynamics on geometric graphs for anatomical modeling of cerebrovascular networks. *IEEE Trans. Med. Imaging* **40**, 381–394. <https://doi.org/10.1109/TMI.2020.3027500> (2021).
- Tahir, W. et al. Anatomical modeling of brain vasculature in two-photon microscopy by generalizable deep learning. *BME Front.* **2021**, 8620932. <https://doi.org/10.34133/2021/8620932> (2021).
- Gagnon, L. et al. Quantifying the microvascular origin of bold-fMRI from first principles with two-photon microscopy and an oxygen-sensitive nanoprobe. *J. Neurosci.* **35**, 3663–3675 (2015).
- Damseh, R. et al. Automatic graph-based modeling of brain microvessels captured with two-photon microscopy. *IEEE J. Biomed. Health Inform.* **23**, 2551–2562. <https://doi.org/10.1109/JBHI.2018.2884678> (2019).
- Ronneberger, O., Fischer, P. & Brox, T. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 234–241 (Springer, 2015).
- Balakrishnan, G., Zhao, A., Sabuncu, M. R., Guttag, J. V. & Dalca, A. V. Voxelmorph: A learning framework for deformable medical image registration. *IEEE Trans. Med. Imaging* **38**, 1788–1800 (2019).
- Almahairi, A., Rajeshwar, S., Sordani, A., Bachman, P. & Courville, A. Augmented cyclegan: Learning many-to-many mappings from unpaired data. In *International conference on machine learning*, 195–204 (PMLR, 2018).
- Hatamizadeh, A. et al. Unetr: Transformers for 3d medical image segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 574–584 (2022).
- Damseh, R. et al. Automated analysis of brain microvasculature: From segmentation to anatomical modeling. In *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 1907–1910 (IEEE, 2020).
- Hatamizadeh, A. et al. Unetr: Transformers for 3d medical image segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 574–584 (2022).
- Zhu, J.-Y., Park, T., Isola, P. & Efros, A. A. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, 2223–2232 (2017).

Acknowledgements

This work was supported by internal funding from the United Arab Emirates University under Grant Numbers 12T037 and 12R239.

Author contributions

R.D. conceived the study. N.B. and J.R. conducted the experiments, while F.L. assisted with data collection. N.Z. and R.D. contributed to the computational analysis. N.B. and J.R. wrote the initial draft of the manuscript. All authors reviewed and provided feedback on the final manuscript.

Declarations

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to R.D.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025