







Probabilistic volumetric speckle suppression in OCT using deep learning

BHASKARA RAO CHINTADA,^{1,2,*}  SEBASTIÁN RUIZ-LOPERA,^{1,3}
RENÉ RESTREPO,⁴  BRETT E. BOUMA,^{1,2,5} MARTIN
VILLIGER,^{1,2}  AND NÉSTOR URIBE-PATARROYO^{1,2} 

¹Wellman Center for Photomedicine, Massachusetts General Hospital, Boston, MA 02114, USA

²Harvard Medical School, Boston, MA 02115, USA

³Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02142, USA

⁴Applied Optics Group, Universidad EAFIT, Carrera 49 # 7 Sur-50, Medellín, Colombia

⁵Institute of Medical Engineering and Science, Massachusetts Institute of Technology, Cambridge, MA 02142, USA

*bchintada@mgh.harvard.edu

Abstract: We present a deep learning framework for volumetric speckle reduction in optical coherence tomography (OCT) based on a conditional generative adversarial network (cGAN) that leverages the volumetric nature of OCT data. In order to utilize the volumetric nature of OCT data, our network takes partial OCT volumes as input, resulting in artifact-free despeckled volumes that exhibit excellent speckle reduction and resolution preservation in all three dimensions. Furthermore, we address the ongoing challenge of generating ground truth data for supervised speckle suppression deep learning frameworks by using volumetric non-local means despeckling–TNode– to generate training data. We show that, while TNode processing is computationally demanding, it serves as a convenient, accessible gold-standard source for training data; our cGAN replicates efficient suppression of speckle while preserving tissue structures with dimensions approaching the system resolution of non-local means despeckling while being two orders of magnitude faster than TNode. We demonstrate fast, effective, and high-quality despeckling of the proposed network in different tissue types that are not part of the training. This was achieved with training data composed of just three OCT volumes and demonstrated in three different OCT systems. The open-source nature of our work facilitates re-training and deployment in any OCT system with an all-software implementation, working around the challenge of generating high-quality, speckle-free training data.

© 2024 Optica Publishing Group under the terms of the [Optica Open Access Publishing Agreement](#)

1. Introduction

Optical coherence tomography (OCT) is a cross-sectional optical imaging technique that provides high-resolution images of biological tissue [1] and has become a well-established clinical diagnostic imaging tool in ophthalmology. Due to the coherent nature of OCT, tomograms contain speckle, which degrades the image quality and hinders visual interpretation [2–5]. Speckle reduction has been an active topic of interest in OCT community and a plethora of techniques have been developed in the literature, which can be broadly classified as hardware-based [6–10] and signal-processing methods [11–24]. Hardware-based methods have the potential to produce higher-quality images. However, hardware modifications and different data acquisition strategies make them too complex for broad adoption and incompatible with imaging *in vivo*. For instance, angular compounding requires long acquisition times and the sample must be static during the entire acquisition; modulation of the point spread function (PSF) of the illumination beam carries a signal-to-noise ratio (SNR) penalty as well as increased acquisition times [6–10].

Cuartas-Vélez, et al, have discussed in more detail the merits and limitations of the hardware-based and signal-processing methods for speckle reduction in our previous work [24]. Among signal-processing methods, the non-local probabilistic despeckling method TNode [24] exploits volumetric information in OCT tomograms to estimate the incoherent intensity value at each voxel. TNode efficiently suppresses speckle contrast while preserving tissue structures with dimensions approaching the system resolution. It is, however, computationally very expensive, a common problem with exhaustive search methods: processing a typical retinal OCT volume takes a few hours.

Deep learning methods have been explored for speckle mitigation by posing the task as an image-to-image translation problem [25–45]. The main goal of an image-to-image translation network is to learn the mapping between an input image and an output image [46,47] either through a supervised modality using training pairs of input and target images, or through an unsupervised modality using independent sets of input images and target images when paired examples are scarce or not easy to generate. Supervised methods are known to yield better results as the model has access to the ground truth information during training. Previous despeckling efforts using deep learning include the use of variants of generative adversarial networks (GANs) [25–31], variants of convolutional neural networks (CNN) based methods [32–34,36–41] and fusion networks [42–45]. However, with few exceptions [48], *these methods have been based on the compounding of multiple B-scans*, acquired at the same sample location, to generate ground truth despeckled tomograms. This presents significant limitations: B-scan compounding only reduces speckle contrast if the component images contain multiple speckle realizations by capturing a variation of microstructural organization within the whole sample. This condition is typically only satisfied inside blood vessels. We note that inaccuracies in scanning or motion artifacts may also be leveraged to provide data for B-scan compounding, but these approaches induce a direct penalty to spatial resolution [10]. In addition, all these methods have focused on two-dimensional speckle suppression on a B-scan per B-scan basis; their performance on slices or projections that include the out-of-plane dimension of volumetric tomograms has not been evaluated. We expect that any three-dimensional manipulation of processed tomograms (i.e., *en face* projections) will contain artifacts due to the aforementioned B-scan-wise processing, disrupting the continuity of tissue structures along the slow-scan axis. In addition, we argue that speckle suppression based on two-dimensional data cannot provide the neural network with complete information on volumetric structures in the training data, and thus it is expected to perform poorly on structures that have a small cross section in a given B-scan.

To overcome the limitations in the state of the art, we present a workflow to utilize volumetric information present in OCT data for near-real-time speckle suppression and enable high-quality deep-learning based volumetric speckle suppression in OCT. We exploit ground-truth training data generated using the tomographic non-local-means despeckling (TNode) [24] and our neural network uses a new cGAN that receives OCT partial volumes as inputs, exploiting three-dimensional structural information for speckle mitigation. Our hybrid deep-learning–TNode-3D (DL-TNode-3D) enables easy training and implementation in a multitude of OCT systems without relying on specialty-acquired training data.

2. Methods

2.1. Data

In this study, we demonstrated our network performance in three different custom-built frequency domain OCT systems; one system used a wavelength-swept light source (Axsun Technologies, Inc., MA, USA) having a spectral bandwidth of 91 nm centered at 1040 nm and a sweep repetition rate of 100 kHz. This system was integrated with the ophthalmic interface of a Spectralis OCT device (Heidelberg Engineering, Germany), which provided eye tracking and fixation capabilities and a transverse resolution (defined hereafter as e^{-2} focal spot diameter after considering the

confocal effect) of 18 μm [49]. The second system operated with a vertical-cavity surface-emitting laser (VCSEL) that provided a spectral bandwidth of 90 nm centered at 1300 nm wavelength and 100 kHz sweep repetition rate [50]. This system had a transverse resolution of 12 μm . The third system was based on a custom-made wavelength-swept laser that utilized a semiconductor optical amplifier (Covega Corp., BOA-4379) as the gain medium and a polygon mirror scanner (Lincoln Laser Co.) as the tunable filter to rapidly sweep the wavelength at a rate of 54 kHz. This laser had a center wavelength of 1300 nm and bandwidth of 110 nm [51] and the system provided a transverse resolution of 16.5 μm . Herein, we refer to these systems as ophthalmic, VCSEL, and polygon systems, respectively. All the datasets used in this study were acquired using Nyquist sampling in the fast and slow scan axes. In the present case, we define Nyquist sampling as the A-line spacing equal to half e^{-2} diameter of the resolution volume. All systems were provided with polarization-diverse balanced detection; thus, all A-lines were recorded for two detection polarization channels.

We selected a volume of interest (VOI) in each dataset consisting of a varying number of B-scans, A-lines per B-scan, and depth samples per A-line. We trained our network for each system separately. For the ophthalmic system, the network was trained and tested using different regions of the retina from two healthy subjects. The smallest features we can see in the retina are capillaries at our ophthalmic system resolution. Generally, OCT angiography would be used for visualizing blood vessels. However, in order to evaluate the resolution-preserving and volumetric nature of our network, we focus now on small blood vessels in the ganglion cell layer after despeckling the tomograms. We trained the network with different tissue types *ex vivo* and *in vivo*, and tested using an additional tissue type not used in training for VCSEL and polygon systems in order to test the generalizability of our network for the tissue types that are not part of the training dataset.

The tissue types used for training and testing for ophthalmic, VCSEL and polygon systems are tabulated in Table 1. In this study, we have used 5 OCT volumes for training and testing for each system ($3 \times 5 = 15$ volumes in total).

Table 1. Overview of datasets used for training and testing for different OCT systems in this study.

System	Training	Testing
Ophthalmic [49]	Retina ($\times 3$)	Retina ($\times 2$)
VCSEL [50]	Chicken heart ($\times 1$) and leg ($\times 1$), nail bed ($\times 1$)	Ventral ($\times 1$) and dorsal ($\times 1$) finger skins
Polygon [51]	Chicken leg ($\times 2$), dorsal finger skin ($\times 1$)	Nail bed ($\times 2$)

2.2. Tomogram pre-processing

Using MATLAB (MathWorks, USA), the acquired OCT fringes were mapped to a linear wavenumber space, numerically compensated for dispersion, apodized with a Hanning window, zero-padded to the next power of two, and Fourier transformed to reconstruct raw complex-valued tomograms with a final pixel size in the axial direction of 4.8 μm , 5.3 μm and 5.8 μm for the ophthalmic, VCSEL, and polygon systems, respectively, assuming unity index of refraction. Data from the VCSEL system was acquired using the k -clock from the light source, therefore, it did not require linearization in wavenumber space.

Data acquired *in vivo* further required inter-B-scan bulk-motion correction to preserve the continuity of tissue along the slow axis direction and enable volumetric despeckling. The reconstructed complex-valued tomograms were phase-stabilized along the fast scan axis and then low-pass filtered with the optimum filter [52]. Phase-stabilization and filtering, applied to each polarization detection channel, enabled the use of efficient sub-pixel image registration [53], which we used to determine and correct for axial and lateral sub-pixel shifts between adjacent

B-scans. After these steps, we calculated the tomogram intensity as the sum of the intensity (squared absolute value) of each polarization detection channel. We then saved the tomogram intensity in logarithmic scale in single precision format for further processing with TensorFlow library in Python.

2.3. Non-local-means despeckling (TNode)

Tomographic non-local-means despeckling (TNode) [24] exploits volumetric information in OCT by making use of 3D similarity windows to retrieve the weights from the volumetric patch-similarity. This method uses a 3D search window, which consists of depth, z , and both fast-axis, x , and slow-axis, y , information. For tomograms acquired *in vivo*, motion correction was applied to guarantee continuity of tissue structures along the slow axis (see Sec. 2.2). TNode efficiently suppresses speckle while preserving tissue structures with dimensions approaching the system resolution. Despite the merits of this method, it is very computationally demanding; even after making it more computationally efficient by an order of magnitude compared to our original implementation (see Supplement 1 and Code 1, Ref. [54]).

In this study, a search window of size $2 \times 8 + 1$ was chosen along the slow-scan axis for TNode speckle reduction, and thus the network was defined to accept partial volumes with $n = 8$; the network is easily modified to accept the partial volumes with any desired n . We set the base filtering parameter, $h_0 = 80 \times 10^{-3}$ and the SNR-dependent parameter, $h_1 = 40 \times 10^{-3}$ to process the OCT volumes.

2.4. cGANs for volumetric speckle suppression

To suppress the speckle in a cross-section x_i , several conventional signal-processing methods in the literature use information from its $2n$ neighborhood cross-sections from x_{i-n} to x_{i+n} in an OCT volume V . We denote the cross sections from x_{i-n} to x_{i+n} as partial volume v_i , and y_i as the speckle-suppressed cross-section corresponding to x_i . Mapping x_i to y_i can be treated as an image-to-image translation problem. cGANs are known to perform extremely well in these problems compared to other CNN-based methods [46]. cGANs can be adapted to learn a speckle suppression mechanism from partial volume v_i and a random noise vector z_i to a target speckle-suppressed cross-section y_i ; $G : \{v_i, z_i\} \rightarrow y_i$. G is the generator, that is trained to produce speckle-suppressed cross-sections that cannot be differentiated from y_i by an adversarial discriminator D .

The objective function of our network is defined as

$$\mathcal{L}_{\text{cGAN}}(G, D) = \mathbb{E}_{v_i, y_i} [\log D(v_i, y_i)] + \mathbb{E}_{v_i, z_i} [\log \{1 - D(v_i, G(v_i, z_i))\}], \quad (1)$$

where \mathbb{E} is an ensemble average and G tries to minimize this objective function against an adversarial D that tries to maximize it, i.e.,

$$G^* = \arg \min_G \left\{ \max_D [\mathcal{L}_{\text{cGAN}}(G, D) + \lambda \mathcal{L}_{L1}(G)] \right\}, \quad (2)$$

where

$$\mathcal{L}_{L1}(G) = \mathbb{E}_{v_i, y_i, z_i} [\|y_i - G(v_i, z_i)\|_1]. \quad (3)$$

Our cGAN network is a modified Pix2Pix network [46], which consists of a generator and a discriminator, as illustrated in Fig. 1. The generator we have used in this study is a U-Net with skip connections [55]. The discriminator is a convolutional patchGAN classifier that penalizes at the scale of patches, as illustrated in Fig. 2.

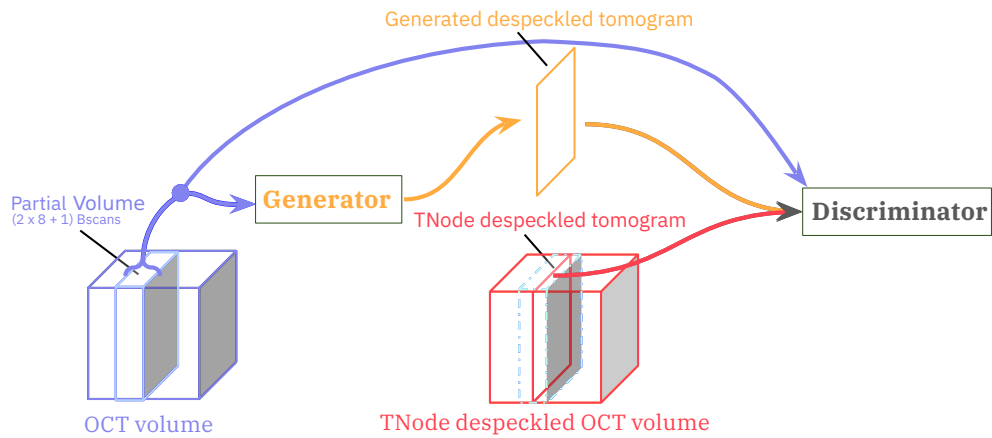


Fig. 1. In DL-TNode-3D, a partial OCT volume is given as input to the generator to learn to despeckle tomograms with the help of a discriminator, making use of the given ground-truth speckle suppressed tomogram produced by TNode.

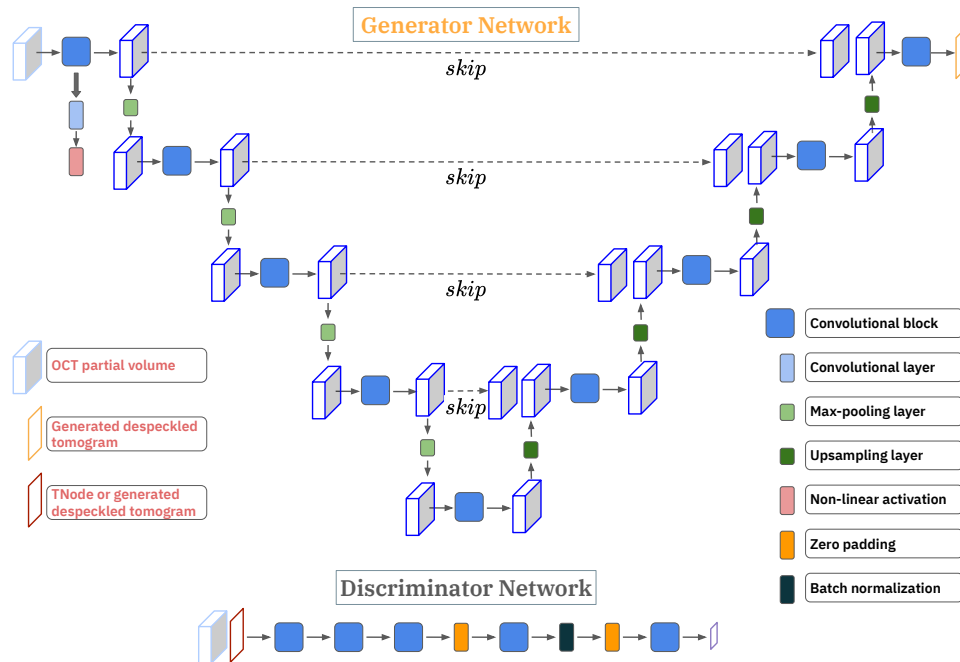


Fig. 2. cGAN architecture of our volumetric speckle suppression network; U-Net as a Generator, convolutional patchGAN classifier as a Discriminator.

2.4.1. Generator

Our U-Net takes a partial OCT volume (v_i) as an input and consists of a series of convolutions and pooling layers, as an encoder, followed by a series of deconvolution and upsampling layers, as a decoder. In the encoder path, the input arrays are gradually downsampled in z_x by half (number of pixels in each dimension) in each layer, while at the same time doubling the number of filter banks. In contrast, in the decoder path, the input feature maps are upsampled in z_x by a factor of two in each layer, while decreasing the number of filters in half. The U-Net architecture also consists of ‘skip’ connections, which allows the network to access the information from earlier layers, which otherwise might be lost due to the vanishing gradient problem [55]. The generator loss penalizes the network if the generated speckle-suppressed cross-section is different from the targeted speckle-suppressed cross-section. The generator loss was defined by a combination of entropy loss and the L_1 loss as in the work by Isola et al. [46]. L_1 loss allows the generated speckle-suppressed image, \hat{y}_i , to be structurally similar to the targeted speckle-suppressed image y_i . As in the work by Isola et al. [46], we provide the noise vector z_i as an input in the form of dropout, applied on several layers of our U-Net generator at both training and testing time.

2.4.2. Discriminator

Our discriminator, patchGAN, enables penalization at a patch level to quantify granular details in the generated images. It consists of a series of blocks tapering down to a desired patch level classification layer, each block consists of a convolution layer, batch normalization, and leaky ReLU as in [46]. We compute discriminator loss in each training step; discriminator loss is a combination of real and fake loss. The real loss was a binary cross-entropy loss of patchGAN output for a provided ground-truth speckle-suppressed tomogram y_i and a matrix of ones, while the fake loss was a binary cross-entropy loss of patchGAN output for a generated speckle-suppressed tomogram \hat{y}_i and a matrix of zeros. Hence, if the generator produced a speckle-suppressed tomogram that matches the provided ground-truth speckle-suppressed tomogram, the real loss will be equal to the fake loss.

2.5. Training data preparation

To create a training pair, a random B-scan was first selected from the TNode-processed volume as a ground truth target image, y_i , together with 16 adjacent B-scans (x_{i-8} to x_{i+8}), and the same B-scan from the raw OCT volume as an input partial volume v_i (17 B-scans in total). We trained our network using 300 partial volumes that were randomly selected from 3 datasets (100 partial volumes from each dataset). The logarithmic intensity values in single precision were loaded into Python and normalized to the uint16 range with defined limits described below. Before being fed into the neural network, the data was converted into TensorFlow-32 values (19-bit precision) to leverage the improved performance of modern GPUs compared to traditional 32-bit single-precision data. We used conventional data augmentation strategies to increase the diversity in our training dataset. In each training step, we changed the contrast dynamically by varying the lower and upper limits of the uint16 representation. The lower limit of the contrast range was set to the noise floor of the dataset, and a random value drawn from a uniform distribution ([0, 10] dB) was added to the lower limit of the contrast range. Similarly, the upper limit of the contrast range was set to an average value computed from the volume of interest ($11 \times 11 \times 3 \text{ px}^3$) centered around the maximum value in the input partial volume, and a random value drawn from a uniform distribution ([-15, 1] dB) was added to the upper limit of the contrast range. In the next step, intensity values lesser than or equal to the lower limit of the contrast range were set to 0, and the intensity values equal to or greater than the upper limit of the contrast range were set to 65535; then, all the remaining intensity values were linearly scaled to the interval of [0 65535]. In addition to this data augmentation step, we randomly flipped, rotated, and performed

random crop and resize on the training pairs to simulate different geometric orientations. We normalized the training pairs from -1 to 1 before we gave them as inputs to our network.

For a given input partial volume, v_i , our generator, G , generated an estimate of the speckle-suppressed cross-section, \hat{y}_i . The discriminator received two inputs: first, the partial volume v_i and the generated speckle-suppressed cross-section \hat{y}_i ; second, the partial volume v_i and the ground-truth speckle-suppressed cross-section y_i . In the next step, we computed the generator loss and the discriminator loss and optimized them using an Adam optimizer [56] with parameters; learning rate of 2×10^{-4} , $\beta_1 = 0.5$, $\beta_2 = 0.99$ and $\epsilon = 1 \times 10^{-7}$. The network was trained on an NVIDIA RTX 5000 with 24 GB of memory. We trained both networks for 200 epochs or until the discriminator loss approached $\approx 2 \ln 2$.

2.6. Experiments

We used the random search approach to determine the best architecture for our network, we conducted experiments by changing the number of convolutional blocks in the generator and changing the number of convolutional blocks and the output size of the discriminator with different training and testing datasets. We converged to the network shown in Fig. 2, which gave us the best results in terms of time efficiency and accuracy. The number of A-lines and depth samples of OCT volumes used in this study were less than 1024. Because the search window size used for TNode processing in the out-of-plane direction was 17 px in size, we chose an input partial volume size in our network of $1024 \times 1024 \times 17$ px³. We used a U-Net with the encoder composed of 5 convolutional blocks with 256, 512, 1024, 2048, and 2048 filter banks, which downsample the input partial volume of size $1024 \times 1024 \times 17$ px³ to $32 \times 32 \times 2048$ px³. We used 256, 512, 1024, and 2048 filter banks in respective convolutional blocks in the encoder part. We padded with zeros the input partial volumes of size less than the specified input partial volume size in the fast and slow axes. Our discriminator consisted of 3 convolutional blocks with 512, 1024, and 1024 filter banks followed by a zero-padding layer and another convolution layer with 2048 filter banks followed by batch normalization, and zero-padding layers, and the last layer was a fully connected layer that outputs a patch-wise classification matrix with size 126×126 corresponding to the classification of patches of size $\approx 8 \times 8$.

To compare and contrast with the current two-dimensional approach in deep learning despeckling, where a speckle-suppressed B-scan is learned from its corresponding raw B-scan, we also trained our network using solely the central raw B-scan of each subvolume as the input, in combination with the same ground truth target image obtained using the TNode of the full subvolume. We herein refer to this approach as cGAN-2D, while DL-TNode-3D refers to our volumetric despeckling approach using partial volumes. We used the same generator and discriminator architecture and data augmentation strategies for cGAN-2D and the same datasets and tissue types for the training. Detailed comparison of DL-TNode-3D and cGAN-2D is discussed in Sec. 3.

2.7. Evaluation metrics

In this study, the following metrics were used for a quantitative analysis of our network performance:

- **Peak-signal-to-noise ratio (PSNR):** measures the quality of our framework generated speckle-suppressed OCT volume compared to ground-truth speckle-suppressed volume obtained using TNode.

$$\text{PSNR} = 10 \log_{10} \left(\frac{I_{\max}^2}{\text{MSE}} \right),$$

where $\text{MSE} = \sum_{p=1}^D (v_p - \hat{v}_p)^2$ is the mean square error between the ground truth voxels v_p and the estimated voxels \hat{v}_p , D is the total number of voxels in the volume, and I_{\max} is the

maximum value possible. MSE computes the cumulative error between the ground truth and the estimated speckle-suppressed OCT volume obtained using our framework. We compute PSNR using 16-bit OCT intensity volumes, hence I_{max} would be $2^{16} - 1$. PSNR quantifies the quality of the generated speckle-suppressed OCT volume, a higher PSNR indicates a better quality, as it means that the generated speckle-suppressed OCT volume is closer to the original TNode processed volume in terms of voxel values, therefore less distorted or noisy.

- **Contrast-to-noise ratio (CNR):** measures the contrast between two tissue types.

$$CNR = \frac{|\mu_{t_2}| - |\mu_{t_1}|}{\sqrt{\sigma_{t_2}^2 + \sigma_{t_1}^2}},$$

where μ_{t_1} and μ_{t_2} are the mean of tissue type #1 and tissue type #2 respectively. Similarly, σ_{t_1} and σ_{t_2} are the standard deviations of tissue type #1 and tissue type #2 respectively. We have computed the CNR between two tissue samples of interest on the speckle-suppressed OCT volume obtained using our framework and compared it with the CNR of the corresponding ground-truth speckle-suppressed OCT volume. The closer the CNR computed on the speckle-suppressed OCT volume obtained using our network to the CNR computed on the ground-truth speckle-suppressed OCT volumes indicates that our framework preserved the contrast between different tissue types as in the ground-truth speckle-suppressed OCT volume. In our CNR computations, we utilized subvolumes of interest with dimensions of $20 \times 40 \times 10$ px³ pixels. These subvolumes of interests were selected from two tissue types and were all centered along the y (slow)-axis.

- **Structural similarity index (SSIM) and multi-scale-SSIM (MS-SSIM):** are image quality assessment methods that assess the similarity between corresponding patches of two images [57,58]. The SSIM is computed using three components, namely the luminance (I), the contrast (C) and the structural (S), and they are defined as

$$I(y, \hat{y}) = \frac{2\mu_y\mu_{\hat{y}} + C_1}{\mu_y^2 + \mu_{\hat{y}}^2 + C_1}, \quad C(y, \hat{y}) = \frac{2\sigma_y\sigma_{\hat{y}} + C_2}{\sigma_y^2 + \sigma_{\hat{y}}^2 + C_2} \quad \text{and} \quad S(y, \hat{y}) = \frac{\sigma_{y\hat{y}} + C_3}{\sigma_y\sigma_{\hat{y}} + C_3},$$

where $\mu_{\hat{y}_i}$ and $\sigma_{\hat{y}_i}$ are the mean and the standard deviation of the generated speckle-suppressed cross-section \hat{y}_i , respectively; μ_{y_i} and σ_{y_i} are the mean and the standard deviation of the label image y_i , respectively; $\sigma_{y_i\hat{y}_i}$ denotes the cross-covariance between \hat{y}_i and y_i ; C_1 , C_2 and C_3 are small positive values used to avoid numerical instability. $SSIM(y_i, \hat{y}_i)$ is the product of these three components,

$$SSIM(y, \hat{y}) = I(y, \hat{y})^\alpha C(y, \hat{y})^\beta S(y, \hat{y})^\gamma,$$

where α , γ and β are exponent weights for the luminance, contrast and structural components. SSIM is measured using a fixed patch size, which does not capture complex variations between the input two images to assess the similarity. MS-SSIM considers the input patches that are iteratively downsampled by a factor of two with low-pass filtering, with scale j denoting the original images downsampled by a factor of $2^j - 1$ and given as

$$MS-SSIM(y, \hat{y}) = I_M(y, \hat{y})^{\alpha M} \prod_{j=1}^M C_j(y, \hat{y})^{\beta_j} S_j(y, \hat{y})^{\gamma_j}.$$

Note that these equations are valid regardless of the dimensionality of the images being compared: for use in volumes, we compute $I(y, \hat{y})$, $C(y, \hat{y})$, and $S(y, \hat{y})$ over 3D-patches instead of 2D-patches.

We used MATLAB functions *ssim* and *multissim* with default values for patch size, exponent weights α , γ and β and constants C_1 , C_2 and C_3 to compute both volume SSIM and MS-SSIM respectively in this study. MATLAB *ssim* function also returns the local SSIM value for each voxel. Using these local SSIM maps, we evaluated the statistical significance between DL-TNode-3D and cGAN-2D using the Student's t-test.

3. Results and discussion

Figure 3 shows representative results for a retinal volume using the ophthalmic OCT system, when trained with three different OCT volumes consisting of distinct fields of view of the retina of two individual subjects. Our network produced OCT volumes that closely resembled TNode volumes; however, the cGAN processed the entire volume of size ($448 \times 818 \times 808 \text{ px}^3$) in just 2 minutes (40 ms per B-scan), as opposed to the 4 hours required for TNode on an NVIDIA RTX 5000 with 24 GB of GPU memory, which is an *improvement of two orders of magnitude in processing time*. We quantified the similarity between the ground truth and our network-produced tomograms using the volume SSIM and MS-SSIM; for this example, the respective values were 0.988 and 0.996. To the best of our knowledge, our network produced speckle-suppressed tomograms more similar to the ground truth—as measured by SSIM and MS-SSIM metrics, and regardless of the method used for ground-truth generation—than any other method in the literature. Our network enhanced the contrast between the layers, similar to ground-truth TNode speckle-suppressed cross-sections while preserving small structures. For instance, the small capillaries, marked with yellow arrows in Fig. 3, became much clearer and easier to identify. Similarly, our network led to improved visualization and differentiation the nerve fiber bundles as shown in Fig. S2 in the Supplement 1.

Figure 4 shows a comparison of speckle reduction with cGAN-2D, the case in which we trained our network using the single raw B-scan of interest as the input—instead of a partial volume—and its corresponding TNode processed B-scan as the target image. Limiting the learning and inference process to 2D processing with cGAN-2D produced a decrease in the quality of the results (similarity metrics were 0.947 and 0.976 for volume SSIM and MS-SSIM, respectively). It is clear from the results in Fig. 4 that 2D processing produces high-frequency artifacts along the slow-scan axis direction, and the quality of despeckling is qualitatively and quantitatively inferior to 3D processing. This demonstrates that the excellent performance of DL-TNode-3D is due to the use of partial volumes for network training and inference, which contain more structural information than individual B-scans. We observed similar DL-TNode-3D performance in tomograms acquired using the two other OCT systems, for which the network had been trained independently.

Figure 5 shows representative results for a ventral finger skin volume (tissue type not part of training data) using the VCSEL-based OCT system when trained with three OCT volumes consisting of the nail bed, chicken heart, and chicken leg tissues. The similarity between the ground truth and our network-produced tomograms using the volume SSIM and MS-SSIM for this example were 0.968 and 0.993, respectively. In contrast, testing with cGAN-2D resulted in a decrease in the quality of the results with volume SSIM and MS-SSIM 0.960 and 0.986, respectively. Furthermore, cGAN-2D processing produced high-frequency artifacts along the slow-axis scan direction, and the quality of speckle suppression is qualitatively inferior to DL-TNode-3D. This network gave similar results when we tested on dorsal finger skin volume with volume SSIM and MS-SSIM 0.988 and 0.996, respectively. This experiment shows that our network is generalizable for tissue types that are not part of the training data. Speckle reduction using our network enhanced the contrast between different ventral and dorsal finger skin layers. Sweat ducts can be identified easily on our network speckle-suppressed cross-sections.

Similarly, Fig. 6 shows representative results for a nail bed volume using the polygon-based OCT system when trained with three OCT volumes consisting of chicken leg (2 \times) and dorsal




Fig. 3. Orthogonal views of tomograms before and after despeckling using TNode (i.e. Ground Truth) and DL-TNode-3D. DL-TNode-3D produces OCT volumes close to the ground truth without any visible artifacts along the out-of-plane axis, y . z is the depth and $x(y)$ is the fast- (slow-) scan axis direction. Yellow arrows indicate small capillaries that are preserved after despeckling with both TNode and DL-TNode-3D. Scale bars = 0.5 mm.

finger skin ($1\times$), respectively. The similarity between the ground truth and our network-produced tomograms using the volume SSIM and MS-SSIM for this example were 0.989 and 0.996, respectively. In contrast, testing with cGAN-2D resulted in a decrease in the quality of the results with volume SSIM and MS-SSIM 0.943 and 0.983, respectively. It is clear from the results in Fig. 6 that 2D processing produces high-frequency artifacts along the slow-axis scan direction, and the quality of despeckling is qualitatively inferior to 3D processing. The quality metrics for the ophthalmic, VCSEL, and polygon systems are summarized in Table 2. It is evident that considering the out-of-plane information for speckle suppression using our DL-TNode-3D framework produced volumes most similar in terms of volume PSNR, SSIM and MS-SSIM to the ground truth volumes compared to using only B-scans in cGAN-2D. The p-value is $<10^{-6}$ between the SSIM of DL-TNode-3D and cGAN-2D in all the cases. cGAN-2D produced volumes with CNR close to the ground truth CNR for models trained on data from the ophthalmic and VCSEL systems but not from the polygon-based system. However, they are corrupted with high-frequency artifacts in the out-of-plane direction in all cases.

Typical loss curves obtained during DL-TNode-3D training are provided in Fig. S3 of Supplement 1. Notably, all three models trained on data from different OCT systems exhibited

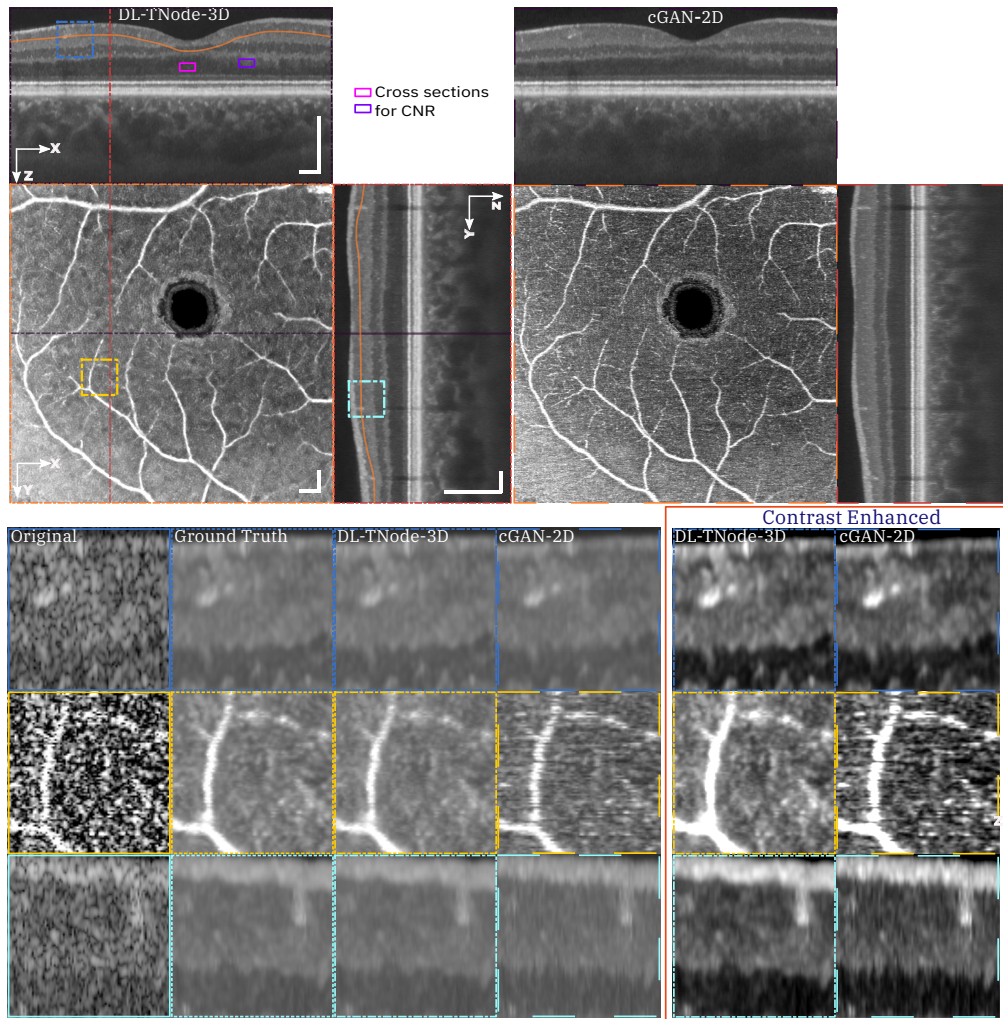


Fig. 4. Comparison of speckle reduction using DL-TNode-3D and cGAN-2D for the retinal volume in Fig. 3. Orthogonal views where z is depth and $x(y)$ is the fast- (slow-) scan axis direction. DL-TNode-3D produces OCT volumes close to the ground truth without any visible artifacts along the out-of-plane axis, y . Contrast-enhanced boxes show superior speckle suppression ability of DL-TNode-3D compared to cGAN-2D, which exhibits high-frequency artifacts along the slow-scan axis. Magenta and violet rectangles represent the two tissue types used for calculating CNR. Scale bars = 0.5 mm.




Fig. 5. Comparison of speckle reduction using DL-TNode-3D and cGAN-2D for the ventral finger skin volume acquired using VCSEL system. Orthogonal views where z is depth and $x(y)$ is the fast (slow) axis direction. DL-TNode-3D produces OCT volumes that match the ground truth without any visible artifacts along the out-of-plane axis, y . Purple arrow in the en face view indicates a motion artifact. Contrast-enhanced boxes show superior speckle suppression ability of DL-TNode-3D compared to cGAN-2D, which exhibits high-frequency artifacts along the slow-scan axis. Magenta and violet rectangles represent the two tissue types used for calculating CNR. Scale bars = 0.5 mm.

Table 2. Quantitative evaluation of DL-TNode-3D and cGAN-2D for 3 OCT systems. DL-TNode-3D is our method, which uses the volumetric information for speckle suppression whereas cGAN-2D uses only 2D (depth-fast axes) information for speckle suppression. CNR values that are nearest to ground truth volume CNR are highlighted in bold. The highest values of PSNR, SSIM, and MS-SSIM are highlighted in bold. CNR: contrast-to-noise ratio, PSNR: peak-signal-to-noise ratio; SSIM: structural similarity index; MS-SSIM: multi-scale structural similarity index.

OCT System	Trained Model	CNR	PSNR (dB) ↑	SSIM ↑	MS-SSIM ↑
Ophthalmic	Ground truth	1.193	-	-	-
	cGAN-2D	1.188	34.726	0.943	0.976
	DL-TNode-3D	1.302	38.076	0.988	0.996
VCSEL	Ground truth	1.623	-	-	-
	cGAN-2D	1.612	37.175	0.954	0.982
	DL-TNode-3D	1.675	41.095	0.978	0.994
Polygon	Ground truth	1.394	-	-	-
	cGAN-2D	1.525	36.876	0.949	0.985
	DL-TNode-3D	1.505	40.654	0.988	0.996

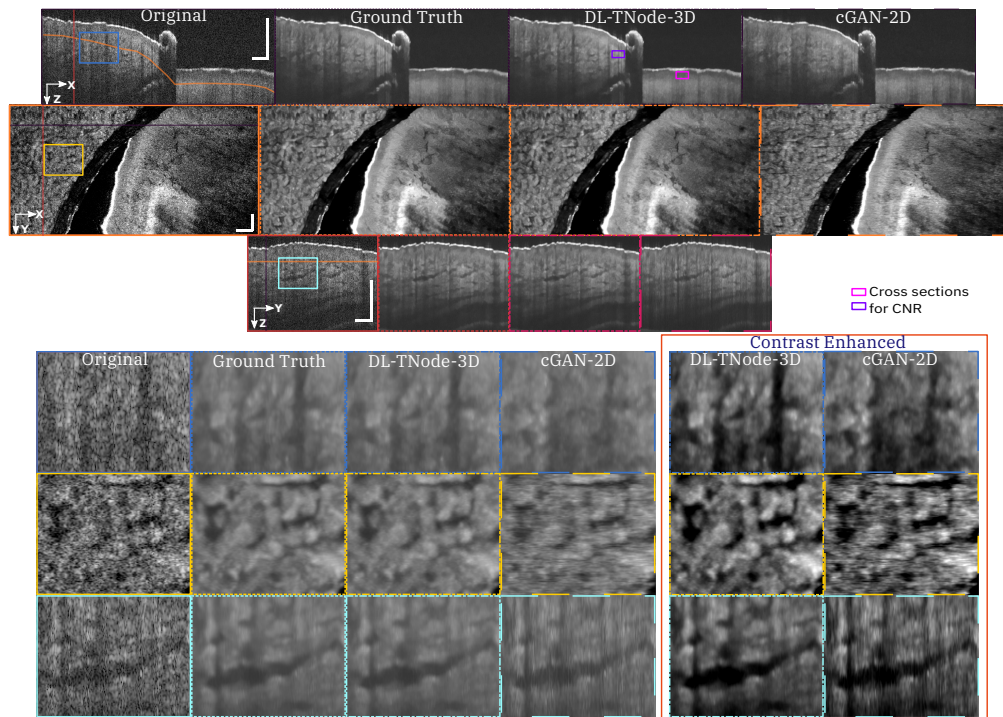


Fig. 6. Comparison of speckle reduction using DL-TNode-3D and cGAN-2D for the ventral nail bed volume acquired using the polygon-based system. Orthogonal views where z is depth and $x(y)$ is the fast (slow) axis direction. DL-TNode-3D produces OCT volumes close to the ground truth without any visible artifacts along the out-of-plane axis, y . Contrast-enhanced boxes show superior speckle suppression ability of DL-TNode-3D compared to cGAN-2D, which exhibits high-frequency artifacts along the slow-scan axis. Magenta and violet rectangles represent the two tissue types used for calculating CNR. Scale bars = 0.5 mm.

similar stable learning curves. This suggests DL-TNode-3D's capability for good generalizability. With our publicly available source code, training DL-TNode-3D for any specific OCT system should be straightforward. DL-TNode-3D process utilizes overlapping subvolumes of 17 B-scans for volume processing. This inevitably leads to the exclusion of the first 7 and last 7 B-scans from the final despeckled output. To ensure a fair comparison in the figures above, we correspondingly trimmed the same number of B-scans from the original OCT and ground truth TNode volumes. However, to retain these excluded B-scans, one potential solution is to utilize repeat padding, which is the padding used in TNode for the generation of the ground-truth data. This padding technique would allow the model to process the entire volume without discarding any B-scans.

Given the excellent performance of our model on the disparate data that were not seen during the training, we anticipate that DL-TNode-3D, trained on retinal volumes obtained with our Ophthalmic system, will demonstrate generalizability to retinal data from patients with various ocular diseases acquired using the same system. Figure 4 illustrates the potential use of our network for glaucoma monitoring as it increases the contrast between retinal layers (nerve fiber layer, ganglion cell layer and inner plexiform layer) for which longitudinal monitoring of thickness is important [59,60]. The value of despeckled OCT data goes beyond potential improvement in layer segmentation: we believe its impact to be higher in improving image interpretation in untrained readers in clinical applications in which small tissue features are of diagnostic importance, such as anterior segment, dermatology, intravascular and endoscopic imaging. We

demonstrated improved visibility of small tissue features with TNode in anterior segment imaging with moderate resolution [52] (see Figs. 4(c) vs 4(e) and their insets in [52]). In cellular-resolution anterior segment imaging there is an even stronger need for despeckling [61]. In the case of retinal imaging, the spatial resolution in commercial ophthalmic OCT systems does not enable the visualization of individual cells in the retina. However, there is increased interest in this capability, as demonstrated by the many recent advances in adaptive-optics OCT [62,63], including the development of a specialized deep learning algorithm for retinal pigment epithelium cell counting in presence of speckle [64]. We believe that DL-TNode-3D will find ever increasing applicability in this kind of high-resolution ophthalmic imaging as well. Even for deep-learning segmentation routines, we expect that the use of despeckled data to generate more accurate training data will be beneficial [65,66]. TNode assumes speckle follows an exponential distribution, which is a reasonable approximation for single-look tomograms. Functional extensions of OCT, such as OCTA or polarization-sensitive OCT, produce noisy parametric images (for parameters such as decorrelation or birefringence) with very different statistics. For this reason, TNode is not intended to be used on the outputs of these modalities.

Our results show that our network trained with only three OCT volumes of readily available tissue can produce despeckled volumes that replicate the efficient suppression of speckle and preservation of tissue structures with dimensions approaching the system resolution known from TNode, while being two orders of magnitude faster. Because DL-TNode-3D relies on an all-software approach for training, it can be easily re-trained and deployed in virtually any OCT system. The updated TNode code for generating the training data and the source code for our neural network are available in [Code 1](#), Ref. [54]. This [Dataset 1](#), Ref. [67] contains a retinal OCT intensity volume as a demo dataset to generate volumetric speckle-suppressed training data using our TNode script and four OCT intensity volumes and their corresponding TNode-processed intensity volumes of different tissue samples acquired using VCSEL OCT system to train and test our deep learning framework.

4. Conclusion

In this work, we presented a despeckling cGAN framework to utilize volumetric information for speckle suppression in OCT. Our framework was trained using partial OCT volumes as input and TNode speckle-suppressed tomogram as targeted output using U-Net as a generator and patchGAN as the discriminator. We trained and tested our framework using 300 partial volumes randomly drawn from three OCT volumes (100 partial volumes from each OCT volume) from three OCT systems separately. Our network produced volumes that approximate the ground-truth despeckled volumes with unprecedented fidelity, for models trained on all three OCT systems. Additionally, our framework is two orders of magnitude faster than TNode and reaches near real-time performance at ~ 20 fps.

Funding. National Institutes of Health (P41EB015903, R01EB033306, K25EB024595); Universidad EAFIT (11100252021).

Acknowledgments. We gratefully acknowledge Dr. Neerav Karani, Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA for discussions on cGAN.

Disclosures. We have no conflicts of interest to disclose.

Data availability. The latest TNode code for generating training data and the source code for DL-TNode-3D are available on [Code 1](#), Ref. [54]. The dataset used for training and testing DL-TNode-3D on VCSEL OCT system is available in [Dataset 1](#), Ref. [67]. The dataset used for training and testing DL-TNode-3D on our Ophthalmic and Polygon systems are not publicly available at this time but may be obtained from the authors upon reasonable request.

Supplemental document. See [Supplement 1](#) for supporting content.

References

1. D. Huang, E. A. Swanson, C. P. Lin, *et al.*, "Optical coherence tomography," *Science* **254**(5035), 1178–1181 (1991).

2. J. M. Schmitt, S. Xiang, and K. M. Yung, "Speckle in optical coherence tomography," *J. Biomed. Opt.* **4**(1), 95–105 (1999).
3. M. Bashkansky and J. Reintjes, "Statistics and reduction of speckle in optical coherence tomography," *Opt. Lett.* **25**(8), 545–547 (2000).
4. B. Karamata, K. Hassler, M. Laubscher, *et al.*, "Speckle statistics in optical coherence tomography," *J. Opt. Soc. Am. A* **22**(4), 593–596 (2005).
5. J. W. Goodman, *Speckle Phenomena in Optics: Theory and Applications* (Roberts and Company Publishers, 2007).
6. N. Iftimia, B. E. Bouma, and G. J. Tearney, "Speckle reduction in optical coherence tomography by 'path length encoded' angular compounding," *J. Biomed. Opt.* **8**(2), 260–263 (2003).
7. M. Pircher, E. Götzinger, R. A. Leitgeb, *et al.*, "Speckle reduction in optical coherence tomography by frequency compounding," *J. Biomed. Opt.* **8**(3), 565–569 (2003).
8. A. Desjardins, B. Vakoc, W.-Y. Oh, *et al.*, "Angle-resolved optical coherence tomography with sequential angular selectivity for speckle reduction," *Opt. Express* **15**(10), 6200–6209 (2007).
9. D. Alonso-Caneiro, S. A. Read, and M. J. Collins, "Speckle reduction in optical coherence tomography imaging by affine-motion image registration," *J. Biomed. Opt.* **16**(11), 116027 (2011).
10. B. F. Kennedy, T. R. Hillman, A. Curatolo, *et al.*, "Speckle reduction in optical coherence tomography by strain compounding," *Opt. Lett.* **35**(14), 2445–2447 (2010).
11. A. Ozcan, A. Bilenca, A. E. Desjardins, *et al.*, "Speckle reduction in optical coherence tomography images using digital filtering," *J. Opt. Soc. Am. A* **24**(7), 1901–1910 (2007).
12. M. Gargsha, M. W. Jenkins, A. M. Rollins, *et al.*, "Denoising and 4d visualization of OCT images," *Opt. Express* **16**(16), 12313–12333 (2008).
13. Z. Jian, Z. Yu, L. Yu, *et al.*, "Speckle attenuation in optical coherence tomography by curvelet shrinkage," *Opt. Lett.* **34**(10), 1516–1518 (2009).
14. A. Wong, A. Mishra, K. Bizheva, *et al.*, "General bayesian estimation for speckle noise reduction in optical coherence tomography retinal imagery," *Opt. Express* **18**(8), 8338–8352 (2010).
15. Z. Jian, L. Yu, B. Rao, *et al.*, "Three-dimensional speckle suppression in optical coherence tomography based on the curvelet transform," *Opt. Express* **18**(2), 1024–1032 (2010).
16. L. Fang, S. Li, Q. Nie, *et al.*, "Sparsity based denoising of spectral domain optical coherence tomography images," *Biomed. Opt. Express* **3**(5), 927–942 (2012).
17. M. Szkulmowski, I. Gorczynska, D. Sznajda, *et al.*, "Efficient reduction of speckle noise in optical coherence tomography," *Opt. Express* **20**(2), 1337–1359 (2012).
18. L. Wang, Z. Meng, X. S. Yao, *et al.*, "Adaptive speckle reduction in oct volume data based on block-matching and 3-d filtering," *IEEE Photonics Technol. Lett.* **24**(20), 1802–1804 (2012).
19. M. Szkulmowski and M. Wojtkowski, "Averaging techniques for oct imaging," *Opt. Express* **21**(8), 9757–9773 (2013).
20. D. Yin, Y. Gu, and P. Xue, "Speckle-constrained variational methods for image restoration in optical coherence tomography," *J. Opt. Soc. Am. A* **30**(5), 878–885 (2013).
21. B. Chong and Y.-K. Zhu, "Speckle reduction in optical coherence tomography images of human finger skin by wavelet modified BM3D filter," *Opt. Commun.* **291**, 461–469 (2013).
22. J. Aum, J.-h. Kim, and J. Jeong, "Effective speckle noise suppression in optical coherence tomography images using nonlocal means denoising filter with double gaussian anisotropic kernels," *Appl. Opt.* **54**(13), D43–D50 (2015).
23. J. Cheng, D. Tao, Y. Quan, *et al.*, "Speckle reduction in 3d optical coherence tomography of retina by a-scan reconstruction," *IEEE Trans. Med. Imaging* **35**(10), 2270–2279 (2016).
24. C. Cuatrecasas-Vélez, R. Restrepo, B. E. Bouma, *et al.*, "Volumetric non-local-means based speckle reduction for optical coherence tomography," *Biomed. Opt. Express* **9**(7), 3354–3372 (2018).
25. Y. Ma, X. Chen, W. Zhu, *et al.*, "Speckle noise reduction in optical coherence tomography images based on edge-sensitive CGAN," *Biomed. Opt. Express* **9**(11), 5129–5146 (2018).
26. K. J. Halupka, B. J. Antony, M. H. Lee, *et al.*, "Retinal optical coherence tomography image enhancement via deep learning," *Biomed. Opt. Express* **9**(12), 6205–6221 (2018).
27. Z. Dong, G. Liu, G. Ni, *et al.*, "Optical coherence tomography image denoising using a generative adversarial network with speckle modulation," *J. Biophotonics* **13**(4), e201960135 (2020).
28. A. Guo, L. Fang, M. Qi, *et al.*, "Unsupervised denoising of optical coherence tomography images with nonlocal-generative adversarial network," *IEEE Trans. Instrum. Meas.* **70**, 1 (2020).
29. N. A. Kande, R. Dakhane, A. Dukkupati, *et al.*, "Siamesegcn: a generative model for denoising of spectral domain optical coherence tomography images," *IEEE Trans. Med. Imaging* **40**(1), 180–192 (2020).
30. Y. Zhou, K. Yu, M. Wang, *et al.*, "Speckle noise reduction for oct images based on image style transfer and conditional gan," *IEEE J. Biomed. Health Inform.* **26**(1), 139–150 (2021).
31. M. Wang, W. Zhu, K. Yu, *et al.*, "Semi-supervised capsule cgan for speckle noise reduction in retinal oct images," *IEEE Trans. Med. Imaging* **40**(4), 1168–1183 (2021).
32. F. Shi, N. Cai, Y. Gu, *et al.*, "Despnet: a cnn-based method for speckle reduction in retinal optical coherence tomography images," *Phys. Med. Biol.* **64**(17), 175010 (2019).
33. A. Abbasi, A. Monadjemi, L. Fang, *et al.*, "Three-dimensional optical coherence tomography image denoising through multi-input fully-convolutional networks," *Comput. Biol. Med.* **108**, 1–8 (2019).

34. S. K. Devalla, G. Subramanian, T. H. Pham, *et al.*, “A deep learning approach to denoise optical coherence tomography images of the optic nerve head,” *Sci. Rep.* **9**(1), 14454 (2019).
35. T. L. Bobrow, F. Mahmood, M. Inserni, *et al.*, “DeepIsr: a deep learning approach for laser speckle reduction,” *Biomed. Opt. Express* **10**(6), 2869–2882 (2019).
36. B. Qiu, Z. Huang, X. Liu, *et al.*, “Noise reduction in optical coherence tomography images using a deep neural network with perceptually-sensitive loss function,” *Biomed. Opt. Express* **11**(2), 817–830 (2020).
37. S. N. Menon, V. Vineeth Reddy, A. Yeshwanth, *et al.*, “A novel deep learning approach for the removal of speckle noise from optical coherence tomography images using gated convolution–deconvolution structure,” in *Proceedings of 3rd International Conference on Computer Vision and Image Processing* (Springer, 2020), pp. 115–126.
38. N. Gour and P. Khanna, “Speckle denoising in optical coherence tomography images using residual deep convolutional neural network,” *Multimed. Tools Appl.* **79**(21–22), 15679–15695 (2020).
39. G. Gisbert, N. Dey, H. Ishikawa, *et al.*, “Self-supervised denoising via diffeomorphic template estimation: application to optical coherence tomography,” in *International Workshop on Ophthalmic Medical Image Analysis* (Springer, 2020), pp. 72–82.
40. S. Apostolopoulos, J. Salas, J. L. Ordóñez, *et al.*, “Automatically enhanced oct scans of the retina: a proof of concept study,” *Sci. Rep.* **10**(1), 7819 (2020).
41. D. Varadarajan, C. Magnain, M. Fogarty, *et al.*, “A novel algorithm for multiplicative speckle noise reduction in ex vivo human brain oct images,” *NeuroImage* **257**, 119304 (2022).
42. D. Hu, J. D. Malone, Y. Atay, *et al.*, “Retinal oct denoising with pseudo-multimodal fusion network,” in *International Workshop on Ophthalmic Medical Image Analysis* (Springer, 2020), pp. 125–135.
43. I. Oguz, J. D. Malone, Y. Atay, *et al.*, “Self-fusion for oct noise reduction,” in *Medical Imaging 2020: Image Processing*, vol. 11313 (SPIE, 2020), pp. 45–50.
44. J. J. Rico-Jimenez, D. Hu, E. M. Tang, *et al.*, “Real-time oct image denoising using a self-fusion neural network,” *Biomed. Opt. Express* **13**(3), 1398–1409 (2022).
45. G. Ni, R. Wu, J. Zhong, *et al.*, “Hybrid-structure network and network comparative study for deep-learning-based speckle-modulating optical coherence tomography,” *Opt. Express* **30**(11), 18919–18938 (2022).
46. P. Isola, J.-Y. Zhu, T. Zhou, *et al.*, “Image-to-image translation with conditional adversarial networks,” in *Proceedings of the IEEE Conference On Computer Vision And Pattern Recognition* (2017), pp. 1125–1134.
47. Y. Pang, J. Lin, T. Qin, *et al.*, “Image-to-image translation: methods and applications,” *IEEE Trans. Multimedia* **24**, 3859–3881 (2021).
48. G. Ni, Y. Chen, R. Wu, *et al.*, “SM-NET OCT: a deep-learning-based speckle-modulating optical coherence tomography,” *Opt. Express* **29**(16), 25511–25523 (2021).
49. B. Braaf, S. Donner, A. S. Nam, *et al.*, “Complex differential variance angiography with noise-bias correction for optical coherence tomography of the retina,” *Biomed. Opt. Express* **9**(2), 486 (2018).
50. T. M. Cannon, B. E. Bouma, and N. Uribe-Patarroyo, “Layer-based, depth-resolved computation of attenuation coefficients and backscattering fractions in tissue using optical coherence tomography,” *Biomed. Opt. Express* **12**(8), 5037 (2021).
51. J. Ren, H. Choi, K. Chung, *et al.*, “Label-free volumetric optical imaging of intact murine brains,” *Sci. Rep.* **7**, 1–8 (2017).
52. S. Ruiz-Lopera, R. Restrepo, C. Cuartas-Vélez, *et al.*, “Computational adaptive optics in phase-unstable optical coherence tomography,” *Opt. Lett.* **45**(21), 5982–5985 (2020).
53. M. Guizar-Sicairos, S. T. Thurman, and J. R. Fienup, “Efficient subpixel image registration algorithms,” *Opt. Lett.* **33**(2), 156–158 (2008).
54. B. R. Chintada, S. Ruiz-Lopera, R. Restrepo, *et al.*, “Probabilistic volumetric speckle suppression in oct using deep learning: Code,” GitHub, (2023). <https://github.com/bhaskarachintada/DLTNode>.
55. O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18* (Springer, 2015), pp. 234–241.
56. D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv*, arXiv:1412.6980 (2014).
57. Z. Wang, A. C. Bovik, H. R. Sheikh, *et al.*, “Image quality assessment: from error visibility to structural similarity,” *IEEE Trans. on Image Process.* **13**(4), 600–612 (2004).
58. Z. Wang, E. P. Simoncelli, and A. C. Bovik, “Multiscale structural similarity for image quality assessment,” in *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, vol. 2 (Ieee, 2003), pp. 1398–1402.
59. K. Banister, C. Boachie, R. Bourne, *et al.*, “Can automated imaging for optic disc and retinal nerve fiber layer analysis aid glaucoma detection?” *Ophthalmology* **123**(5), 930–938 (2016).
60. J.-C. Mwanza, D. L. Budenz, J. L. Warren, *et al.*, “Retinal nerve fibre layer thickness floor and corresponding functional loss in glaucoma,” *Br. J. Ophthalmol.* **99**(6), 732–737 (2015).
61. S. Khan, K. Neuhaus, O. Thaware, *et al.*, “Corneal imaging with blue-light optical coherence microscopy,” *Biomed. Opt. Express* **13**(9), 5004–5014 (2022).
62. Y. Zhang, J. Rha, R. S. Jonnal, *et al.*, “Adaptive optics parallel spectral domain optical coherence tomography for imaging the living retina,” *Opt. Express* **13**(12), 4792–4811 (2005).
63. D. T. Miller and K. Kurokawa, “Cellular-scale imaging of transparent retinal structures and processes using adaptive optics optical coherence tomography,” *Annu. Rev. Vis. Sci.* **6**(1), 115–148 (2020).

64. V. Das, F. Zhang, A. J. Bower, *et al.*, “Revealing speckle obscured living human retinal cells with artificial intelligence assisted adaptive optics optical coherence tomography,” *Commun. Med.* **4**(1), 68 (2024).
65. C. Ge, X. Yu, M. Yuan, *et al.*, “Self-supervised self2self denoising strategy for OCT speckle reduction with a single noisy image,” *Biomed. Opt. Express* **15**(2), 1233–1252 (2024).
66. B. Yao, L. Jin, J. Hu, *et al.*, “PSCAT: a lightweight transformer for simultaneous denoising and super-resolution of OCT images,” *Biomed. Opt. Express* **15**(5), 2958–2976 (2024).
67. B. R. Chintada, S. Ruiz-Lopera, R. Restrepo, *et al.*, “Probabilistic volumetric speckle suppression in OCT using deep learning: Dataset,” Zenodo, (2023) <https://zenodo.org/records/10258100>.