



Volumetric non-local-means based speckle reduction for optical coherence tomography

CARLOS CUARTAS-VÉLEZ,^{1,*} RENÉ RESTREPO,¹ BRETT E. BOUMA,^{2,3} AND NÉSTOR URIBE-PATARROYO²

¹Applied Optics Group, Universidad EAFIT, Carrera 49 # 7 Sur-50, Medellín, Colombia

²Wellman Center for Photomedicine, Harvard Medical School and Massachusetts General Hospital, 40 Blossom Street, Boston, MA 02114, USA

³Institute for Medical Engineering and Science, Massachusetts Institute of Technology, Cambridge, MA 02142, USA

*ccuarta1@eafit.edu.co

Abstract: We present a novel tomographic non-local-means based despeckling technique, TNode, for optical coherence tomography. TNode is built upon a weighting similarity criterion derived for speckle in a three-dimensional similarity window. We present an implementation using a two-dimensional search window, enabling the despeckling of volumes in the presence of motion artifacts, and an implementation using a three-dimensional window with improved performance in motion-free volumes. We show that our technique provides effective speckle reduction, comparable with B-scan compounding or out-of-plane averaging, while preserving isotropic resolution, even to the level of speckle-sized structures. We demonstrate its superior despeckling performance in a phantom data set, and in an ophthalmic data set we show that small, speckle-sized retinal vessels are clearly preserved in intensity images *en-face* and in two orthogonal, cross-sectional views. TNode does not rely on dictionaries or segmentation and therefore can readily be applied to arbitrary optical coherence tomography volumes. We show that despeckled esophageal volumes exhibit improved image quality and detail, even in the presence of significant motion artifacts.

© 2018 Optical Society of America under the terms of the [OSA Open Access Publishing Agreement](#)

OCIS codes: (110.4500) Optical coherence tomography; (100.6950) Tomographic image processing; (100.0100) Image processing; (110.4280) Noise in imaging systems; (170.3880) Medical and biological imaging.

References and links

1. D. Huang, E. A. Swanson, C. P. Lin, J. S. Schuman, W. G. Stinson, W. Chang, M. R. Hee, T. Flotte, K. Gregory, C. A. Puliafito, and J. G. Fujimoto, "Optical coherence tomography," *Science* **254**(5035), 1178–1181 (1991).
2. J. M. Schmitt, S. H. 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. Express* **25**(8), 545–547 (2000).
4. B. Karamata, K. Hassler, M. Laubscher, and T. Lasser, "Speckle statistics in optical coherence tomography," *J. Opt. Soc. Am. A* **22**(4), 593–596 (2005).
5. 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).
6. A. E. Desjardins, B. J. Vakoc, W. Y. Oh, S. M. R. Motaghiannezam, G. J. Tearney, and B. E. Bouma, "Angle-resolved optical coherence tomography with sequential angular selectivity for speckle reduction," *Opt. Express* **15**(10), 6200–6209 (2007).
7. B. F. Kennedy, T. R. Hillman, A. Curatolo, and D. D. Sampson, "Speckle reduction in optical coherence tomography by strain compounding," *Opt. Lett.* **35**(14), 2445–2447 (2010).
8. 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).
9. M. Pircher, E. Gotzinger, R. Leitgeb, A. F. Fercher, and C. K. Hitzenberger, "Speckle reduction in optical coherence tomography by frequency compounding," *J. Biomed. Opt.* **8**, 565–569 (2003).
10. M. Szkulmowski, I. Gorczynska, D. Sznajda, M. Sylwestrzak, A. Kowalczyk, and M. Wojtkowski, "Efficient reduction of speckle noise in optical coherence tomography," *Opt. Express* **20**(2), 1337–1359 (2012).

11. M. Szkulmowski, and M. Wojtkoski, "Averaging techniques for OCT imaging," *Opt. Express* **21**(8), 9757–9773 (2013).
12. L. Fang, S. Li, Q. Nie, J. A. Izatt, C. A. Tith, and S. Farsiu, "Sparsity based denoising of spectral domain optical coherence tomography images," *Biomed. Opt. Express* **3**(5), 927–942 (2012).
13. A. Ozcan, A. Bilenca, A. E. Desjardins, B. E. Bouma, and G. J. Tearney, "Speckle reduction in optical coherence tomography images using digital filtering," *J. Opt. Soc. Am.* **24**(7), 1901–1910 (2007).
14. H. Yu, J. Gao, and A. Li, "Probability-based non-local means filter for speckle noise suppression in optical coherence tomography images," *Opt. Express* **41**(5), 994–997 (2016).
15. 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).
16. Z. Jian, Z. Yu, L. Yu, B. Rao, Z. Chen, and B. J. Tromberg, "Speckle attenuation in optical coherence tomography by curvelet shrinkage," *Opt. Express* **34**(10), 1516–1518 (2009).
17. Z. Jian, L. Yu, B. Rao, B. J. Tromberg, and Z. Chen, "Three-dimensional speckle suppression in optical coherence tomography based on the curvelet transform," *Opt. Express* **18**(2), 1024–1032 (2010).
18. L. Fang, S. Li, R. P. McNabb, Q. Nie, A. N. Kuo, C. A. Toth, J. A. Izatt, and S. Farsiu, "Fast acquisition and reconstruction of optical coherence tomography images via sparse representation," *IEEE T. Med. Imaging* **32**(11), 2034–2049 (2013).
19. L. Fang, S. Li, D. Cunefare, and S. Farsiu, "Segmentation based sparse reconstruction of optical coherence tomography images," *IEEE T. Med. Imaging* **16**(2), 407–421 (2017).
20. J. Cheng, D. Tao, Y. Quan, D. W. Kee Wong, G. C. Ming Cheung, M. Akiba, and J. Liu, "Speckle reduction in 3D optical coherence tomography of retina by A-scan reconstruction," *IEEE T. Med. Imaging* **35**(10), 2270–2279 (2016).
21. M. Gargasha, M. W. Jenkins, A. M. Rollins, and D. L. Wilson, "Denoising and 4D visualization of OCT images," *Opt. Express* **16**(16), 641–655 (2012).
22. A. Wong, A. Mishra, K. Bizheva, and D. A. Clausi, "General Bayesian estimation for speckle noise reduction in optical coherence tomography retinal imagery," *Opt. Express* **18**(8), 8338–8352 (2010).
23. 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).
24. L. Wang, Z. Meng, X. S. Yao, T. Liu, Y. Su, and M. Qin, "Adaptive speckle reduction in OCT volume data based on block-matching and 3-D filtering," *IEEE Photonic. Tech. L.* **24**(20), 1802–1804 (2012).
25. 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).
26. P. P. Srinivasan, L. A. Kim, P. S. Mettu, S. W. Cousins, G. M. Comer, J. A. Izatt, and S. Farsiu, "Fully automated detection of diabetic macular edema and dry age-related macular degeneration from optical coherence tomography images," *Biomed. Opt. Express* **5**(10), 3568–3577 (2014).
27. Y. Gu, and X. Zhang, "Spiking cortical model based non-local means method for despeckling multiframe optical coherence tomography data," *Laser Phys. Lett.* **14**, 056201 (2017).
28. X. Zhang, G. Hou, J. Ma, W. Yang, B. Lin, Y. Xu, W. Chen, Y. Feng, "Denoising MR images using non-local means filter with combined patch and pixel similarity," *PLOS ONE* **9**(6), e100240 (2014).
29. K. Lu, N. He, and L. Li, "Nonlocal means-based denoising for medical images," *Hindawi Publishing Corporation* **2012**, 438617 (2012).
30. J. V. Manjón, P. Coupé, L. Martí-Bonmatí, D. L. Collins, and M. Robles, "Adaptive non-local means denoising of MR images with spatially varying noise levels," *J. Magn. Reson. Imaging* **31**, 192–203 (2010).
31. C.-A. Deledalle, L. Denis, and F. Tupin, "Iterative weighted maximum likelihood denoising with probabilistic patch-based weights," *IEEE T. Image Process.* **18**(12), 2661–2672 (2009).
32. C.-A. Deledalle, L. Denis, F. Tupin, A. Reigber, and M. Jäger, "NL-SAR: a unified nonlocal framework for resolution-preserving (Pol)(In)SAR denoising," *IEEE Geosci. Remote. S.* **53**(4), 2021–2038 (2015).
33. F. Argenti, A. Lapini, T. Bianchi, and L. Alparone, "A tutorial on speckle reduction in synthetic aperture radar images," *IEEE Geosci. Remote. S.* **1**(3), 6–35 (2013).
34. J. Yu, J. Tan, and Y. Wang, "Ultrasound speckle reduction by a SUSAN-controlled anisotropic diffusion method," *Pattern Recognition* **43**(9), 3083–3092 (2010).
35. P. V. Sudeep, P. Palanisamy, J. Rajan, H. Baradaran, L. Saba, A. Gupta, and J. S. Suri, "Speckle reduction in medical ultrasound images using an unbiased non-local means method," *Biomedical Signal Processing and Control* **28**, 1–8 (2016).
36. J. W. Goodman, *Speckle phenomena in optics* (W. H. Freeman, 2010).
37. K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3D transform-domain collaborative filtering," *IEEE T. Image Process.* **16**(8), 2080–2095 (2007).
38. C.-A. Deledalle, L. Denis, and F. Tupin, "How to compare noisy patches? Patch similarity beyond Gaussian noise," *International Journal of Computer Vision* **99**(1), 86–102 (2012).
39. A. Buades, B. Coll, and J. M. Morel, "A review of image denoising algorithms, with a new one," *Multiscale Model. Simul.* **4**(2), 490–530 (2005).
40. Y. Jia, J. C. Morrison, J. Tokayer, O. Tan, L. Lombardi, B. Baumann, C. D. Lu, W. Choi, J. G. Fujimoto, and D. Huang, "Quantitative OCT angiography of optic nerve head blood flow," *Biomed. Opt. Express* **3**(12), 3127–3137 (2012).

41. T. E. de Carlo, A. Romano, N. K. Waheed, and J. S. Duker, "A review of optical coherence tomography angiography (OCTA)," *International Journal of Retina and Vitreous* **1**(5), 1:5 (2015).
42. S. Shin, U. Sharma, H. Tu, W. Jung and S. A. Boppart, "Characterization and analysis of relative intensity noise in broadband optical sources for optical coherence tomography," *IEEE Photonic. Tech. L.* **22**(14), 1057–1059 (2010).
43. Y. Chen, D. M. de Bruin, C. Kerbage, and J. F. de Boer, "Spectrally balanced detection for optical frequency domain imaging," *Opt. Express* **15**(25), 16390–16399 (2007).
44. S. Makita, F. Jaillon, I. Jahan, and Y. Yasuno, "Noise statistics of phase-resolved optical coherence tomography imaging: single- and dual-beam-scan Doppler optical coherence tomography," *Opt. Express* **22**(4), 4830–4848 (2012).
45. "NvisionVLE Registry System Registry," (2014–2018), <https://clinicaltrials.gov/ct2/show/NCT02215291>.
46. H. C. Wolfsen, P. Sharma, M. B. Wallace, C. Leggett, G. Tearney, and K. K. Wang, "Safety and feasibility of volumetric laser endomicroscopy in patients with Barrett's esophagus (with videos)," *Gastrointest. Endosc.* **82**(4), 631 (2015).
47. N. Uribe-Patarroyo, and B. E. Bouma, "Rotational distortion correction in endoscopic optical coherence tomography based on speckle decorrelation," *Opt. Lett.* **40**(23), 5518–5521 (2015).
48. F. Timischl, "The contrast-to-noise ratio for image quality evaluation in scanning electron microscopy," *Scanning* **37**, 54–62 (2015).
49. D. Adler, T. H. Ko, and J. G. Fujimoto, "Speckle reduction in optical coherence tomography images by use of a spatially adaptive wavelet filter," *Opt. Lett.* **29**(24), 2878–2880 (2004).
50. Y. Gan, and C. P. Fleming, "Extracting three-dimensional orientation and tractography of myofibers using optical coherence tomography," *Biomed. Opt. Express* **4**(10), 2150–2165 (2013).
51. C. J. Goergen, H. Radhakrishnan, S. Sakadžić, E. T. Mandeville, E. H. Lo, D. E. Sosnovik, and V. J. Srinivasan, "Optical coherence tractography using intrinsic contrast," *Opt. Lett.* **37**(18), 3882–3884 (2012).
52. C. M. Ambrosi, V. V. Fedorov, I. R. Efimov, R. B. Schuessler, and A. M. Rollins, "Quantification of fiber orientation in the canine atrial pacemaker complex using optical coherence tomography," *J. Biomed. Opt.* **17**(7), 071309 (2012).
53. M. Villiger, E. Z. Zhang, S. K. Nadkarni, W.-Y. Oh, B. J. Vakoc, and B. E. Bouma, "Spectral binning for mitigation of polarization mode dispersion artifacts in catheter-based optical frequency domain imaging," *Opt. Express* **21**(14), 16353–16369 (2013).
54. D. C. Adams, L. P. Hariri, A. J. Miller, Y. Wang, J. L. Cho, M. Villiger, J. A. Holz, M. V. Szabari, D. L. Hamilos, R. Scott Harris, J. W. Griffith, B. E. Bouma, A. D. Luster, B. D. Medoff, and M. J. Suter, "Birefringence microscopy platform for assessing airway smooth muscle structure and function in vivo," *Sci. Transl. Med.* **8**(359), 359ra131 (2016).
55. M. Villiger, K. Otsuka, A. Karanasos, P. Doradla, J. Ren, N. Lippok, M. Shishkov, J. Daemen, R. Diletti, R.-J. van Geuns, F. Zijlstra, G. van Soest, P. Libby, E. Regar, S. K. Nadkarni, and B. E. Bouma, "Coronary plaque microstructure and composition modify optical polarization: A new endogenous contrast mechanism for optical frequency domain imaging," *JACC: Cardiovasc. Imag., Dec.* (2017).
56. W. C. Y. Lo, M. Villiger, A. Golberg, G. F. Broelsch, S. Khan, C. G. Lian, W. G. Austen Jr., M. Yarmush, B. E. Bouma, "Longitudinal, 3D imaging of collagen remodeling in murine hypertrophic scars in vivo using polarization-sensitive optical frequency domain imaging," *J. Invest. Dermatol.* **136**(1), 84–92 (2016).
57. C. Cuartas-Vélez, R. Restrepo, B. E. Bouma, and N. Uribe-Patarroyo, "Volumetric non-local-means based speckle reduction for optical coherence tomography," *figshare* (2018) [retrieved 4 Apr 2018], <https://doi.org/10.6084/m9.figshare.6089861>.

1. Introduction

Optical coherence tomography (OCT) is an imaging technique which produces high-resolution cross-sectional images of biological tissues [1]. Due to the coherent nature of OCT, its tomographic images present speckle [2–4] which degrades image quality and hinders visual interpretation. Consequently, the reduction of speckle has been a perennial topic of interest in the OCT community [5–27]. Indeed, the suppression of speckle and speckle-like noise while preserving the visibility of fine-scale structure is an active area of research for most coherent imaging techniques, including magnetic resonance imaging [28–30], synthetic aperture radar (SAR) [31–33] and ultrasound imaging [34, 35]. The main goal of speckle reduction techniques is to increase image quality, thereby facilitating visual interpretation and, in the case of medical applications, improving diagnostic utility.

Speckle reduction in coherent imaging is typically achieved by performing an incoherent averaging of uncorrelated speckle realizations of the imaged object [5–9]. Incoherent averaging reduces the speckle contrast, and as the number of uncorrelated realizations increases, the averaged result approaches the equivalent incoherent image [36]. However, for static objects

where the scatterer distribution does not change in time, the incoherent addition of uncorrelated speckle realizations carries an implicit resolution loss with respect to the original information content. The search for improved speckle reduction techniques relies on obtaining uncorrelated speckle patterns reducing the impact on the spatial resolution of the image.

Speckle-reduction techniques for OCT can be classified in two general families: hardware-based and post-processing algorithms. Hardware-based approaches modify the acquisition system to produce uncorrelated speckle patterns within or between B-scans. Examples include angular [5, 6], spatial [7, 8] and frequency compounding [9]. It is easy to see the loss of information content in each implementation: in angular compounding, the raw data contains information corresponding to a higher numerical aperture (NA) than the compounded image — although to create the high-NA image coherent addition of multiple angles is required; in spatial compounding, neighboring sample locations are merged into a single pixel; in frequency compounding, the raw data contains information corresponding to a higher axial resolution. Moreover, the addition of components in commercial OCT systems and special acquisition schemes required in hardware-based techniques is not a straightforward task [8]. In Ref. [7], strain induced on the sample changes the arrangement of the scatterers inside each resolution volume. This spatial approach can, in principle, preserve the spatial resolution while decreasing the speckle contrast. However, the number of uncorrelated speckle patterns that can be realistically created without incurring resolution loss is small [7].

Perhaps the most common embodiment of speckle reduction in OCT relies on the acquisition of repeated B-scans comprising the same sample structure but uncorrelated speckle. This is a special case of spatial compounding and has long been used in commercial ophthalmic OCT systems with success. However, if two B-scans are acquired at exactly the same spatial location, the speckle patterns must necessarily be highly correlated and compounding does not reduce the speckle contrast [10, 11]. When compounding B-scans, speckle reduction comes from two sources: in perfused tissue, like the choroid, speckle decorrelation arises mainly from moving red blood cells. In avascular tissue, speckle decorrelation is due to small lateral offsets given by small errors in the eye tracking system.

Available speckle suppression post-processing algorithms for OCT can be classified in general in three families: transformation based, sparse representations and spatial domain. Although many advanced post-processing algorithms have been developed and applied to OCT, it is surprising that most approaches focus on improving individual B-scans, without regard to the volumetric data structure [12–15]. Frame-to-frame analysis after filtering is generally hindered and further volumetric processing is difficult. Additionally, most of these algorithms treat speckle as additive Gaussian noise through logarithmic conversions, ignoring its specific statistical properties [16, 17].

Among post-processing algorithms, transform-based approaches represent images in a different domain by using wavelet or curvelet transforms. In the curvelet or wavelet domain, the coefficients representing the image and speckle are *assumed* to be distinguishable. Choosing the correct threshold for the speckle coefficients, however, is difficult [13, 17, 21] and the wrong selection of coefficients causes point-like structures to disappear in the filtered images while preserving the contrast between tissue layers [17]. Sparse representation techniques in OCT adopt a special scanning approach in which few B-scans are captured with high nominal signal-to-noise ratio (SNR) and used to improve the quality of low SNR B-scans. For example, multiscale sparsity-based tomographic denoising (MSBTD) [12] utilizes the high-SNR images to train dictionaries used in the denoising of neighboring low-SNR B-scans. MSBTD was further improved by sparsity based simultaneous denoising and interpolation [18], where the dictionaries were developed from a sparse representation of previous datasets consisting of pairs of low-SNR / low-resolution and high-SNR / high-resolution images, thus not requiring the peculiar scanning system of the MSBTD. More recently, segmentation based sparse reconstruction has enhanced the filter performance by using segmentation of retinal layers to construct specific structural dictionaries,

taking advantage of similarity in the patches within segmented layers [19]. Nevertheless, the reliance on a dictionary limits applicability: previously unobserved structures do not have dictionary counterparts and are therefore poorly addressed. Other methods, such as block-matching and 3D filtering (BM3D) [23–26, 37], are combinations between transform-based and sparse representation algorithms. Although these approaches may function well in specific instances, they tend to over-smooth regions, and visible artifacts are commonly found in the final images.

Resolution-preserving spatial-domain techniques are mainly represented by non-local methods, where the evaluation of each pixel is performed within a large neighborhood based on the existing redundancy of natural images [14, 27–32, 38, 39]. Some approaches, such as the non-local means (NLM) algorithm, denoise images by comparing patches in a surrounding neighborhood and mapping the similarities between patches into weights in order to perform a weighted maximum likelihood estimation of the noise-free image. The NLM algorithm has been implemented in OCT in a probabilistic context (PNLM) [14]. However, PLNM treats speckle as additive Gaussian noise and speckle nulls are seen to corrupt the probability calculation. These incorrectly characterized speckle nulls remain after the filtering process, degrading image quality [14].

We have developed an NLM-based algorithm for OCT —Tomographic Non-local-means despeckling (TNode)— built upon the polarimetric and interferometric non-local framework developed by Deledalle *et al.* [31, 32, 38] for SAR images. In analogy with the SAR work, we treat speckle based on its physical properties and use proper speckle statistics in order to adequately estimate weights in the despeckling process. The core idea is to determine small volumetric patches in the tomogram that represent different speckle realizations of the same underlying object, such as small blood vessels. The incoherent averaging is performed non-locally from these patches. Ideally, because only uncorrelated speckle realizations of the same object are compounded, resolution is preserved and speckle contrast reduced. From speckle statistics, we employ a general context of single speckle realization and multiple speckle realizations tomograms to determine the speckle-free intensity. This general context allows TNode to be used *in-vivo* with singlelook frame analysis and easily adapted to a context where multiple speckle realizations are acquired.

As opposed to the approach used in SAR, we exploit the volumetric information available in OCT by making use of 3D similarity windows to retrieve the weights from the volumetric patch-similarity. We have developed two kinds of search windows: a 2D search window oriented to include the fast-axis of transverse scanning for imaging *in-vivo* and a 3D window for motion-artifact-free or motion-corrected volumetric tomograms. The 2D search window allows speckle reduction even in the presence of motion artifacts as the weighted estimation from the NLM algorithm searches for the most-similar neighbors [32], unaffected by relative distances as long as the motion is smaller than the search window size. The 3D search window enhances the contrast of point-like structures while preserving resolution in cross-sectional images and improving the quality of *en-face* projections.

Our 3D similarity window enhances the selection of similar patches. Many tissue structures appear one-dimensional (e.g. vessels, nerves) in traditional OCT imaging due to resolution limitations [40, 41], as shown in Fig. 1. Such structures are not generally aligned with the scanning direction and therefore appear as point-like structures in cross-sectional views. State-of-the-art filtering is incapable of correctly identifying these structures in a single B-scan, hence detail is lost and only tissue layers and other larger structures are preserved. Our 3D similarity window ensures that one- and two-dimensional structures are correctly characterized and consequently preserved upon despeckling. In essence, TNode considers the 3D geometrical properties of volumetric structures while performing the averaging. We do this voxel-per-voxel in a feature-agnostic way, avoiding the need for image segmentation or dictionaries. Our results show that TNode volumes present high-quality, cross-sectional views oriented along both slow and fast axes of transverse

scanning, as well as in *en-face* views. TNode B-scans resemble those obtained from out-of-plane averaging, without the associated out-of-plane resolution loss.




Fig. 1. One-dimensional structures in OCT have speckle-sized appearance in single cross-sections. Filtering techniques are therefore incapable of distinguish between bright contrast speckle and retinal structures. The proposed 3D similarity window considers the volumetric nature of structures with 3D similarity windows and a 2D or 3D search windows.

2. TNode theoretical basis

2.1. Speckle statistics

OCT imaging exhibits speckle because of the superposition of coherent light backscattered with random phases [36]. Tomograms with a single realization of speckle have unity speckle contrast C : the speckle fluctuations are of the order of the signal itself and the image of the underlying object is severely degraded. To quantify the deleterious effect of speckle in coherent imaging, we define the signal-to-speckle ratio, SSR, as the inverse of the speckle contrast: $SSR = 1/C = \bar{I}/\theta$, where \bar{I} is the speckle-free object intensity and θ the speckle standard deviation. We also define the signal-to-noise ratio (SNR) as $SNR = I_s/I_n$, where I_s represents the OCT noise-free signal intensity and I_n is the intensity of the noise. In this formulation, I_n encompasses all sources of noise (e.g. shot, thermal, excess photon, digitization, etc.) in the OCT system [2, 42, 43].

We use speckle statistics in order to find the speckle-free object intensity by determining the parametric distribution that is characteristic of speckle. On an intensity basis, speckle is described by an exponential probability distribution [2, 36], which is described by a single parameter. Assuming the object has an incoherent, speckle-free intensity $\theta(x)$ at pixel x , the probability of measuring an intensity $I(x)$ under coherent imaging for an arbitrary speckle realization is

$$p[I(x)|\theta(x)] = \frac{1}{\theta(x)} \exp\left\{-\frac{I(x)}{\theta(x)}\right\}, \quad (1)$$

where $\theta(x)$ is the distribution parameter. In an exponential distribution, $\theta(x)$ is the distribution mean *and* standard deviation, and as explained above, it equals the mean intensity $\bar{I}(x)$ of a large set of speckle realizations. Therefore, the correct determination of $\theta(x)$ is equivalent to finding $\bar{I}(x)$, the underlying object intensity under incoherent imaging. Speckle suppression can be

understood as an estimation of the parameter $\theta(x)$ from data with speckle by using uncorrelated speckle patterns.

In OCT systems that acquire uncorrelated speckle patterns using spatial, angular, frequency or any other type of compounding, the averaged intensity from L speckle patterns has a gamma probability distribution [36]

$$p[I(x)|\theta(x)] = \frac{L^L}{\Gamma(L)\theta(x)^L} I(x)^{L-1} \exp\left\{-L \frac{I(x)}{\theta(x)}\right\}, \quad (2)$$

where Γ is the gamma function. Equation (2) describes a distribution with SSR = \sqrt{L} , and reduces to Eq. (1) for $L = 1$.

Adopting SAR terminology, we refer to single-speckle-realization tomograms [Eq. (1)] as “singlelook tomograms”, and to intensity-averaged speckle realizations [Eq. (2)] as “multilook tomograms”.

2.2. Non-local denoising and speckle suppression

In general terms, non-local methods reduce noise by performing a weighted intensity average among a set of regions within a vicinity [31, 32, 38, 39]. These regions, called patches, are not necessarily adjacent, but are ideally from anywhere in the image. The restriction to a bounded vicinity is generally driven by performance considerations. Non-local methods base the denoising process on the probability of two given patches to be different realizations of the noise with the same underlying object. In this case, the pixel values of the two patches are described by a single probability distribution with a common set of parameters and these parameters describe the underlying noise-free object [14, 15, 28–30, 35].

In the context of non-local despeckling, the goal is to determine the probability that the pixel x being despeckled has the same underlying object—its speckle is described by the same distribution parameter—as any other pixel x' inside a neighborhood or search window. To define the probability of having a common parameter θ , the intensity $I(x)$ in the pixel x is compared with the intensity $I(x')$ in the set of neighbors. For each comparison, the pair of pixels are considered similar if their intensities match distributions with a common parameter $\theta(x, x')$, hence $I(x)$ and $I(x')$ have a high probability of being two realizations of speckle of the same object. The pair of pixels are considered dissimilar if their intensities have a high probability of being observations of two distributions with different parameters $\theta(x)$ and $\theta(x')$, and therefore different underlying objects [38]. This problem can be rephrased as a hypothesis test [38]:

$$\mathcal{H}_0 : \theta(x) = \theta(x') \equiv \theta(x, x') \quad (\text{null hypothesis}), \quad (3)$$

$$\mathcal{H}_1 : \theta(x) \neq \theta(x') \quad (\text{alternative hypothesis}). \quad (4)$$

In practice, comparing just two pixel intensities does not provide a robust measure of similarity. For this reason, non-local methods consider a group of pixels \mathbf{x} , a patch, centered at the pixel of interest x . Hereafter, we use a bold symbol to denote a vector quantity with element values assigned to each pixel inside a patch. The set of pixels \mathbf{x} are compared pairwise with the pixels \mathbf{x}' centered at x' . Similarly, $\boldsymbol{\theta}(x)$ represents a vector of distribution parameters, and $\mathbf{I}(x)$ a vector of pixel intensities, for the patch centered at x . Note that all pixels inside the patch are *not* considered to share the same parameter θ . The patch size, also called similarity window, defines the smallest *structure* that will be considered during comparisons [39]. Large patch sizes improve the robustness of the similarity metric but reduce the probability of finding similar patches because the compared structures are large. Therefore, the patch size should contain enough independent speckles to have a robust similarity metric but not too many to compromise the effectiveness of the despeckling process.

A similarity criterion makes use of the two hypotheses above to map a pair of speckle patches $(\mathbf{x}, \mathbf{x}')$ into a single real value. The process consists of determining the similarity, C_r , pixelwise and then compounding the probabilities inside a patch into a log-probability $\Delta(x, x')$:

$$\Delta(x, x') = \sum_{\substack{q \in \mathbf{x} \\ q' \in \mathbf{x}'}} \log \{C_r[I(q), I(q')]\}, \quad (5)$$

where the purpose of the logarithm is to transform the product of the probabilities into a sum for all patch pixels.

It is known that the optimal criterion, C_r , based on the hypothesis in Eq. (3), is the likelihood ratio test [38]:

$$\mathcal{L}[I(x), I(x')] = \frac{p[I(x); \theta(x, x')]p[I(x'); \theta(x, x')]}{p[I(x); \theta(x)]p[I(x'); \theta(x')]}, \quad (6)$$

where the numerator is the probability p that intensities in \mathbf{x} and \mathbf{x}' have a common parameter $\theta(x, x')$ and the denominator is the probability that the intensities have different parameters $\theta(x)$ and $\theta(x')$. Note that there is one probability per pixel, and therefore \mathcal{L} is a vector containing a criterion value for each pixel in the patch. This vector is transformed into a single scalar for each patch using Eq. (5).

In practice, the parameters $\theta(x)$, $\theta(x')$ and $\theta(x, x')$ in Eq. (6) are unknown. To address this, we use the generalized likelihood ratio (GLR) \mathcal{L}_G , which solves this problem by replacing $\theta(x)$, $\theta(x')$ and $\theta(x, x')$ by their maximum likelihood estimates (MLEs) \hat{t}_1 , \hat{t}_2 and $\hat{t}_{1,2}$, respectively [38],

$$\mathcal{L}_G[I(x), I(x')] = \frac{p[I(x); \theta(x, x') = \hat{t}_{1,2}]p[I(x'); \theta(x, x') = \hat{t}_{1,2}]}{p[I(x); \theta(x) = \hat{t}_1]p[I(x'); \theta(x') = \hat{t}_2]}. \quad (7)$$

A detailed derivation of the GLR is shown in Ref. [38]. The essence of Eq. (7) is that if the intensities in compared pixels are likely to have the same parameter from a unique distribution, the numerator increases and the GLR value tends to a maximum. However, if the intensities of the pixels are more likely to be given by different individual distributions, the denominator decreases the GLR value. In the case of identical patches, $\mathcal{L}_G = 1$. As the MLE for the gamma distribution in Eq. (2) is the mean of the observed intensities (i. e. $\hat{t}_1 = I(x)$, $\hat{t}_2 = I(x')$ and $\hat{t}_{1,2} = [I(x) + I(x')]/2$), it is possible to arrive at a closed form for the GLR [38]:

$$\mathcal{L}_G[I(x), I(x')] = \left(\frac{I(x)I(x')}{\left\{\frac{1}{2}[I(x) + I(x')]\right\}^2} \right)^L. \quad (8)$$

In TNode, we use GLR in order to map the similarities between patches into weights. Now, referring to locations in the tomogram, the pixels x and x' become vectors containing the three coordinates $\vec{x} = (x, y, z)$ and $\vec{x}' = (x', y', z')$, which indicate the locations in the volume. Throughout this work, z corresponds to the depth dimension, x to the lateral fast scanning axis and y the lateral slow scanning axis. In these terms, Eq. (5) becomes:

$$\Delta(\vec{x}, \vec{x}') = \sum_{\vec{\tau} \in \vec{p}} \log (\mathcal{L}_G[I(\vec{x} + \vec{\tau}), I(\vec{x}' + \vec{\tau})]), \quad (9)$$

where $\vec{\tau}$ is a 3D shift vector indicating locations in each patch such that $\vec{x} + \vec{\tau}$ and $\vec{x}' + \vec{\tau}$ span the pixels in the neighborhood of x and x' . The components of \vec{p} correspond to the patch half-size, with total elements $P = (2p_x + 1) \times (2p_y + 1) \times (2p_z + 1)$. The windows indicated by $\vec{x} + \vec{\tau}$ and $\vec{x}' + \vec{\tau}$ represent the similarity windows. We define a 3D similarity window that accounts for the volumetric nature of OCT tomograms. This 3D similarity window considers one-dimensional structures in the volumetric data that appear as point-like structures in cross-sectional views

and ensures their preservation upon despeckling since they are correctly characterized in the volumetric patch similarity metric.

After computing the log-probability using Eq. (9), the weights are obtained by recovering the compound probability modified by the parameter h [38]:

$$w(\vec{x}, \vec{x}') = \exp \left[\frac{\Delta(\vec{x}, \vec{x}')}{h} \right], \quad (10)$$

where $h > 0$ controls the overall distribution of weight values. Assuming the pixel x is being despeckled, the weights $w(\vec{x}, \vec{x}')$ are calculated for all pixels within a vicinity and then used in the averaging process to retrieve the speckle-free intensity at \vec{x} . Note that $w(\vec{x}, \vec{x}')$ is a scalar. NLM performs a weighted average over all pixels $\vec{x}' + \vec{\tau}$ in an extended vicinity \vec{v} , called a search window, centered at \vec{x} and with total size $V = (2v_x + 1) \times (2v_y + 1) \times (2v_z + 1)$ [32, 39]. The search window defines a subvolume where the central pixel \vec{x} of the similarity window can be located, and therefore the number of similarity windows contained in the search window is always V regardless of the similarity window size. TNode is based on a weighted maximum likelihood estimation of the speckle-free intensity $\hat{I}(\vec{x})$ in terms of a weighted average from the intensity with speckle $I(\vec{x}')$ as

$$\hat{I}(\vec{x}) = \frac{\sum_{\vec{x}' \in \vec{v}} w(\vec{x}, \vec{x}') I(\vec{x}')}{\sum_{\vec{x}' \in \vec{v}} w(\vec{x}, \vec{x}')}. \quad (11)$$

In the event of having an object in the similarity window with no similar neighbors within the search window, weights will be small and thus no despeckling is performed. Now, in order to better explain the effect of the h parameter, assume, for simplicity, two one-dimensional patches consisting of only two pixels: patch A (x_{A1}, x_{A2}) and patch B (x_{B1}, x_{B2}). Explicitly, Eq. (10) states

$$w(x_A, x_B) = \exp \left\{ \frac{1}{h} \left[\log \left(\mathcal{L}_G[I(x_{A1}), I(x_{B1})] \right) + \log \left(\mathcal{L}_G[I(x_{A2}), I(x_{B2})] \right) \right] \right\} = \left(\mathcal{L}_G[I(x_{A1}), I(x_{B1})] \right)^{1/h} \left(\mathcal{L}_G[I(x_{A2}), I(x_{B2})] \right)^{1/h}. \quad (12)$$

Considering that $0 \leq \mathcal{L}_G \leq 1$: if h is near unity then the weights remain unchanged; a larger h increases the probability ratio for each pixel; and, all probability ratios approach 1 when h approaches ∞ . This implies that when h is large, all weights increase and their contributions are equalized, which increases averaging. When h is very large the behavior approaches that of local averaging with a kernel size equal to the search window. Finally, the weight definition in Eq. (10) is unaffected by the relative distance between patches; it is influenced only by the patch-similarity, making TNode completely non local.

In general, the weight for the self similarity is significantly larger than any other weight $w(\vec{x}, \vec{x}) \gg w(\vec{x}, \vec{x}')_{\{x' | x' \neq x\}}$ and therefore, for effective despeckling, we have found that h needs to take a value around $50 \lesssim h \lesssim 100$. In most cases, it is more effective to replace the self similarity by a value given by all the other values in the search window, such as a multiple of the maximum [28]. This produces a more consistent despeckling, reduces the value of h needed and homogenizes the despeckling behavior for different samples. In this work, we replace the self similarity by $w(\vec{x}, \vec{x}) = \max [w(\vec{x}, \vec{x}')_{\{x' | x' \neq x\}}]$.

Figure 2 describes the idea behind TNode. As presented in Eq. (2), the exponential distribution followed by speckle in singlelook intensities ($L = 1$) turns into a Rayleigh distribution if an average between two images is produced ($L = 2$); it transforms into a Gaussian-like distribution when more intensities are averaged [36]. In the ideal case, after acquiring a large number of realizations ($L \rightarrow \infty$), the mean intensity corresponds to the parameter of the distribution, as shown in Fig. 2(a). However, this requires the acquisition of multiple speckle realizations. The

NLM algorithm uses patches inside the image [Fig. 2(b)] in order to produce a mean intensity similar to that obtained after averaging many speckle realizations [32]. The similarity windows depicted in Fig. 2(b) are shown for illustrative purposes. Ideally, the search window should be comprised of the entire data set; in practice, the search window is limited to a given subvolume by memory usage and maximum allowable computation time. The circular patches, 1 to 3, and the square patches, 4 to 7, in Fig. 2(b) illustrate how different patches in the image look alike, and averaging their pixels can produce a despeckled image.

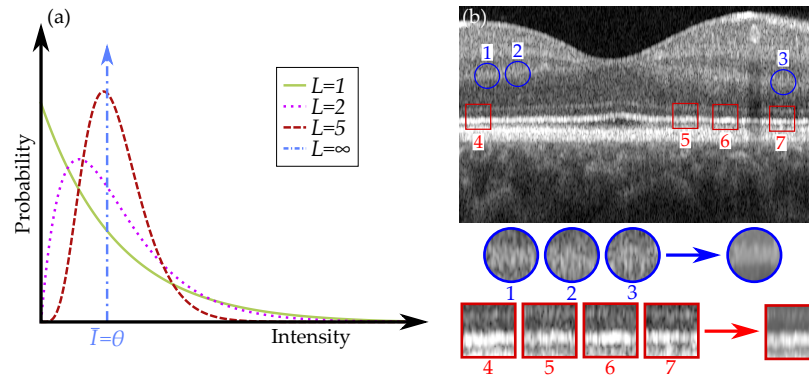


Fig. 2. Representation of the gamma probability density function in Eq. (2) obtained when averaging L realizations of speckle (a). Patch comparison performed in NLM algorithms (b): multiple patches are compared and the probability of having a common underlying object is measured with the GLR to produce the averaged intensity. Patches and their distances are for illustrative purposes; actual patches are 3D and generally comprise only a few pixels per dimension.

Unlike in radar imaging, OCT has a wide SNR range, reaching 50dB or more. In regions with low SNR, the similarity criterion in Eq. (8) may produce artificially high values, as the noise in the tomogram arises from a unique distribution parameterized by the average noise intensity with the same statistical properties as speckle—a circular complex Gaussian random variable. To account for this, we modify the filtering parameter, h , to achieve the same level of despeckling across a wide range of SNRs. Guided by the functional relationship between the correlation coefficient of OCT signals and their SNR [44], we propose the following form for h :

$$h = h_0 + \frac{h_1}{1 + \frac{1}{SNR}}, \quad (13)$$

where h_0 is the base despeckling parameter and h_1 is an SNR-dependent parameter. With this criterion, in regions with low-SNR h approaches h_0 when $SNR \rightarrow 0$, increasing the selectivity in the averaging process. In high-SNR regions, h increases ($h = h_0 + h_1$ for $SNR \rightarrow \infty$), relaxing the similarity criterion and therefore equalizing the weights of similar neighboring patches.

3. Experimental results

3.1. Ophthalmic OCT imaging

We tested the performance of TNode in an eye scan of a healthy volunteer acquired with a Spectralis OCT2 Imaging System (Heidelberg Engineering, Germany). This spectral-domain OCT system uses a super luminescent diode with 870 nm central wavelength, 50 nm bandwidth and an A-line acquisition rate of 85 kHz. We selected a region of interest (ROI) consisting of 257 B-scans with 512 A-lines per B-scan and 640 depth samples. We produced a multilook tomogram by recording four speckle realizations ($L = 4$). The axial pixel size of the reconstructed

tomograms is $3.9 \mu\text{m}$ in air, the A-line spacing is $14 \mu\text{m}$ and the B-scan spacing is $28 \mu\text{m}$. We used a similarity window with $\tau = 7 \times 7 \times 7$ pixels, equivalent to $\approx 4 \times 7 \times 2$ speckles. We despeckled the data set using a two-dimensional search window (TNode 2D with $v = 41 \times 1 \times 21$ pixels, 861 locations containing ≈ 180 distinct speckles) oriented in the fast transverse scanning plane *prior* to correction of motion artifacts between B-scans. We also performed despeckling after motion correction with a three-dimensional search window (TNode 3D with $v = 15 \times 5 \times 11$ pixels, 825 locations containing ≈ 170 distinct speckles). We set the base filtering parameter $h_0 = 0$ and the SNR-dependent parameter $h_1 = 40$ for singlelook and $h_1 = 100$ for multilook tomograms.

Figure 3 presents an individual B-scan and the results of despeckled B-scans with the following algorithms: out-of-plane averaging (OOPA), block matching and 3D filtering (BM3D) [37], probabilistic non-local means (PNLM) [14] and our TNode with both 2D and 3D search windows. We used default parameters in all algorithms and produced OOPA with seven adjacent B-scans after correction of motion artifacts. After motion correction, neighboring B-scans exhibited similar structures and the spatial compounding produced by OOPA reduced speckle successfully. However, since 1D structures may be located in different positions across B-scans, there is potential for loss in spatial resolution. This effect can be seen in the insets (red box). BM3D, as a processing algorithm designed for white noise suppression, requires transformation of speckle into additive noise through a logarithmic conversion. This may result in artifacts and speckle remains visible in the filtered B-scans. PNLM relies on characterizing speckle nulls by measuring the probability of having signal or speckle in each pixel; undetected speckle nulls remain visible in the despeckled B-scans. TNode-filtered B-scans resemble the visual appearance of OOPA but without the associated degradation of resolution. Similar results are obtained with the 2D search window with motion artifacts and the 3D window after correction of motion. However, TNode in 3D enhances layer homogeneity after despeckling, as evidenced by the photoreceptor layer (orange box). For a comparison of the tomograms after despeckling see [Visualization 1](#).

In order to show the resolution-preserving and volumetric nature of TNode, we focus now on blood vessels in the ganglion cell layer of the eye scan after despeckling the tomogram. Generally, OCT angiography would be used for visualizing blood vessels. Here, we take advantage of their one-dimensional structure to evaluate the resolution after despeckling. To facilitate visualization of the volumetric data, we flattened the tomogram, referenced to the retinal pigment epithelium layer, and performed correction of motion artifacts. Figure 4 presents a comparison of *en-face* and cross-sectional views of the eye scan, showing the original volume and results of spatial compounding with OOPA and after despeckling with TNode 3D. As OOPA takes advantage of the similarity between neighboring B-scans, retinal layers are well-defined after despeckling. The orange and green arrows show large blood vessels which are difficult to identify in single cross-sections but easier to identify in *en-face* projections. After OOPA, these large vessels become more visible due to their high scattering compared to avascular tissue. However, because their location shifts among B-scans, OOPA produces significant blurring after the averaging process. In contrast, TNode prevents this blurring as the similarity is evaluated non-locally to produce the weighted average.

The red box in Fig. 4 highlights an ROI containing two blood vessels, closely spaced along the x axis. In the original tomogram, the vessels appear as two speckles, slightly brighter than the surrounding tissue. There is simply not enough information in a single B-scan to infer the true identity of these two speckles. In the OOPA case, the resolution loss associated with averaging multiple B-scans is evident: the two vessels appear as a single blurred bright area. In contrast, as TNode exploits volumetric structures when despeckling, the geometrical distribution of the vessels is considered in the averaging. This ensures that bright pixels containing information belonging to a one-dimensional structure are preserved in both axes while surrounding speckle is suppressed. The blue box highlights an ROI where a blood vessel appears as a bright speckle. After using OOPA, the resolution loss and the low contrast produces a blurred bright area. The

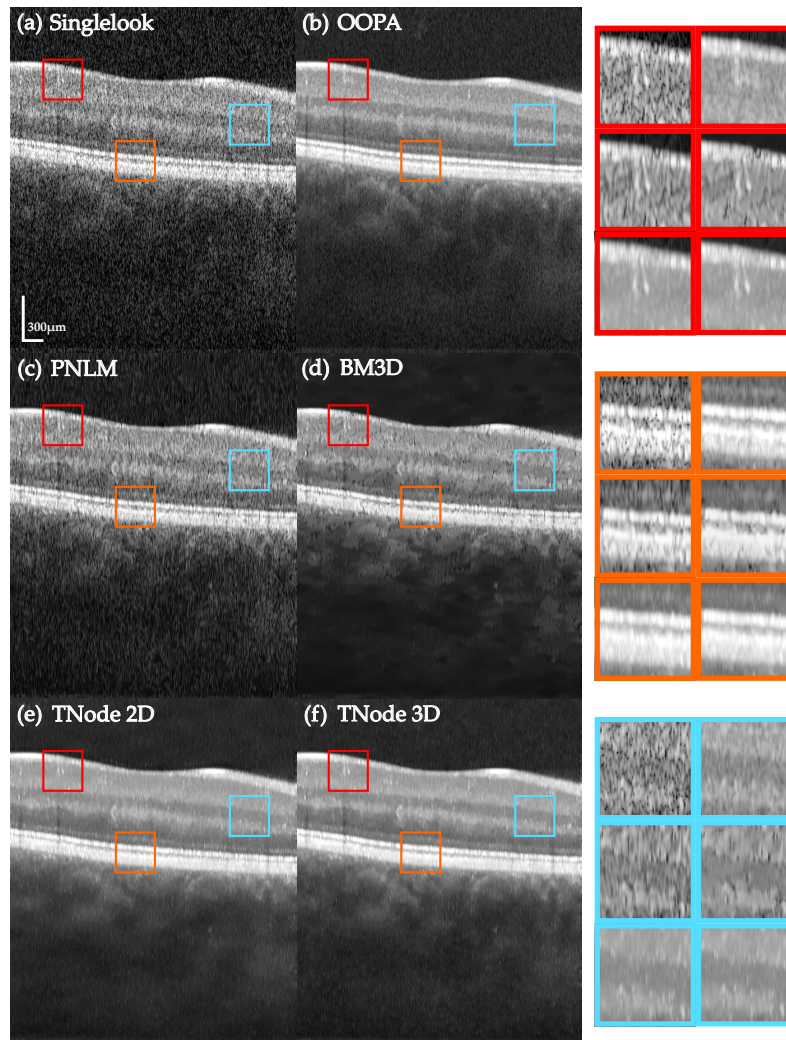


Fig. 3. Native single-look cross-sectional view of human retina *in vivo* (a) and corresponding images despeckled with (b) OOPA, (c) PNLM, (d) BM3D, (e) TNode 2D and (f) TNode 3D. The insets at right are organized in the same order. See [Visualization 1](#) for a comparison of the full tomogram.

same vessel is clearly visible after using TNode. Note that even with anisotropic sampling along axes, the *en-face* view and longitudinal cross-section zy exhibit a homogeneous speckle reduction without evident resolution degradation.

As TNode can be used in tomograms captured with multiple realizations of speckle and previously averaged tomograms, we tested its performance after averaging the four realizations acquired with the eye scan. Since perfused tissue varies speckle dynamically during acquisition, most of the improvement involved in the multiple realizations and multilook process is achieved near vascular tissue, while static avascular tissue remains largely unaffected. Figure 5 is a comparison between single and multilook *en-face* views in the outer plexiform layer of the same eye scan after flattening. We have adjusted the dynamic range of the images in the multilook case, where we have found an increase in the mean intensity after producing the compound intensities. The mean intensity of the multilook tomogram is 2.7 dB higher compared to the singlelook case,




Fig. 4. Orthogonal views of tomograms before (a) and after despeckling using OOPA (b) and TNode 3D (c). TNode preserves the two speckle-sized vessels shown in the red box as two independent structures, while OOPA merges them into a single structure. Resolution in the yz projection (small blood vessel in blue box) is mostly preserved in TNode. Scale bars correspond to $300\mu\text{m}$ in each direction.

therefore we shifted the intensity range to match this increase. This modification has been applied to the multilook tomogram after using TNode as well.

Figure 5 clearly shows how perfused tissue in the multilook images has improved contrast with respect to the singlelook images. Moreover, after despeckling with TNode, vessels are preserved. Structures strongly corrupted by the presence of speckle are clearly noticeable after despeckling, as indicated in the boxes in Fig. 5. See [Visualization 1](#) for a comparison of the *en-face* views of the tomogram.

3.2. Gastrointestinal OCT

The voxel-per-voxel approach of TNode, without the need for dictionaries, enables its use in any imaging application of OCT. To demonstrate this, we used TNode to despeckle a gastrointestinal OCT (GI-OCT) clinical data set, acquired in a recently completed clinical registry [45], with an NvisionVLE Imaging System (NinePoint Medical, Inc., Bedford, Massachusetts) [46]. This is a frequency-domain OCT system with a wavelength-swept laser centered at 1310 nm and with a 90 nm sweep range and a 50 kHz A-line repetition rate. We selected an ROI of 300 B-scans with 1024 A-lines per B-scan (out of the 4096 A-lines per B-scan in the full volume) and 900 pixels in depth. The data set was processed with OOPA using 7 adjacent B-scans, the PNLM filter with default parameters and TNode 3D, in all cases, non-uniform rotational distortion (NURD)




Fig. 5. *En-face* view comparison of a single layer in the outer plexiform layer with single and multilook tomograms: singlelook (a) and multilook (b) tomograms with speckle, and corresponding singlelook (c) and multilook (d) after using TNode 3D. Boxes highlight structures corrupted by speckle that appear clearer after TNode. See [Visualization 2](#) for a comparison of the tomogram.

correction [47] was previously performed. The axial pixel size of the reconstructed tomograms is $6 \mu\text{m}$ in air. The OCT system has a transverse e^{-2} beam radius of $33 \mu\text{m}$, the A-line spacing is $\approx 15 \mu\text{m}$ and the B-scan spacing is $\approx 50 \mu\text{m}$. Because of the highly anisotropic lateral sampling, we used a similarity window with $\tau = 11 \times 7 \times 7$ pixels equivalent to $\approx 7 \times 7 \times 2$ speckles, and a search window with $\nu = 13 \times 5 \times 9$ pixels (585 locations containing ≈ 120 distinct speckles). The parameters were set to $h_0 = 0$ and $h_1 = 45$.

Figure 6 presents the comparison of the ROI cross-sectional view with the obtained results. As motion artifacts, even after NURD correction, are routine in endoscopic OCT because of patient involuntary motion, OOPA-processed B-scans have very poor quality, with structures, such as glands, and tissue morphology appearing significantly blurred and with reduced contrast. The PNLM filtered B-scan shows speckle nulls which were unrecognized during the calculation of the corruption probability. Additionally, because of the high SNR range in GI-OCT data, which is significantly higher than in ophthalmic data, some high-SNR regions still present significant speckle, while low-SNR regions have an over-smoothed appearance. Our results indicate that TNode 3D produces the highest homogeneity in the B-scan while preserving most of the morphological information, even with evident patient motion. The TNode 3D B-scan shows a significant reduction of speckle, preserving contrast and maintaining the structural and morphological characteristics of tissue. A homogeneous reduction of speckle is evident across the whole B-scan. For a comparison of the 300 B-scans see [Visualization 3](#).

3.3. Phantom data and quantitative quality metrics

Quantitative image metrics, such as the contrast-to-noise ratio or the peak signal-to-noise ratio, rely on having *known* homogeneous regions or access to the underlying speckle-free intensity. When applied to natural images, these metrics are prone to misinterpretation because over-smoothed images with lack of detail get highly rated. In order to calculate meaningful image quality metrics, we performed an experiment with a three-dimensional structured phantom made from cellulose. The detailed cellulose structure was obtained in an OCT scan by immersing the phantom in water. This allowed us to define regions with signal produced exclusively by the phantom and void regions defined as the background. Next, we added intralipid at 20% volume




Fig. 6. GI-OCT cross-section comparison of data with speckle (a) and after using OOPA (b), PNLM (c) and TNode 3D (d). TNode exhibits a superior despeckling performance at all SNRs. See [Visualization 3](#) for a comparison of the 300 B-scans.

concentration to the water bath to create scattering in the cellulose voids. We acquired 2 volumes at different intralipid concentrations, corresponding to 1% and 0.3%.

The volumes were acquired with a custom-built wavelength-swept source OCT system, with 1310 nm central wavelength, 130 nm bandwidth and an A-line repetition rate of 54 KHz. We selected an ROI having 256 B-scans with 256 A-lines per B-scan and 256 depth samples, acquired with an objective lens with $35 \mu\text{m } e^{-2}$ beam diameter. The reconstructed tomograms had an A-line spacing of $10 \mu\text{m}$, a B-scan spacing of $10 \mu\text{m}$ and an axial pixel size of $6 \mu\text{m}$ in air. In order to quantitatively compare the performance of TNode with other speckle-reduction algorithms, we despeckled the phantom tomograms with OOPA, BM3D and PNLM. For TNode, we used a $\tau = 7 \times 7 \times 7$ pixels similarity window, equivalent to $\approx 4 \times 7 \times 2$ speckles, a $v = 41 \times 1 \times 21$ pixels (861 locations containing ≈ 180 distinct speckles) two-dimensional search window for TNode 2D and a $v = 15 \times 5 \times 11$ pixels three-dimensional search window for TNode 3D (825 locations containing ≈ 170 distinct speckles). Although the objective lens has the same lateral resolution in the x and y axes, the Brownian motion of the lipid particles produced decorrelated speckle patterns across B-scans regardless of the B-scan spacing; this effect was considered in the calculations above. We set the base filtering parameter $h_0 = 0$ and the SNR-dependent parameter $h_1 = 38$. We produced OOPA with seven adjacent B-scans and set default parameters for PNLM and BM3D. For BM3D and PNLM tomograms images were first transformed into logarithmic scale, while in OOPA and TNode tomograms were kept on a linear scale.

With clearly defined homogeneous areas (corresponding to cellulose voids filled with intralipid), we defined three two-dimensional ROIs corresponding to small cellulose structures surrounded by intralipid for all possible projections: A is an xy (*en-face*) ROI, B is an yz ROI and C an xz ROI [see Fig. 7(a)]. The phantom in water was used to define a mask for the pixels containing the

cellulose structure, which when immersed in intralipid had an intensity $I_S(\vec{x})$, and the inverse mask for the background containing intralipid with intensity $I_{IL}(\vec{x})$. Logarithmic counterparts to these intensities are denoted by L [$L_X(\vec{x}) \equiv 10 \log_{10} I_X(\vec{x})$]. Based on these, we calculated the following quality metrics:

1. The average contrast between cellulose and intralipid, defined as \bar{I}_S/\bar{I}_{IL} , where $\bar{\cdot}$ denotes the average among all the pixels \vec{x} in a given ROI. This average implicitly removes the speckle fluctuations and therefore, when calculated in a singlelook tomogram, defines a speckle-free ideal contrast. Despeckling algorithms should preserve this contrast as much as possible.
2. The contrast-to-noise ratio (CNR) [48, 49] defined as $(\bar{L}_S - \bar{L}_{IL})/\sqrt{\sigma(L_S)^2 + \sigma(L_{IL})^2}$, where $\sigma(L)$ corresponds to the standard deviation of the intensity L in an ROI. CNR describes the contrast between two regions of different intensities in presence of noise, and uses logarithmic intensity to characterize this metric in a similar manner to when tomograms are examined visually. Despeckling algorithms should heavily reduce σ_{IL} , increasing the CNR. σ_S is expected to decrease. However, its value will be dominated by changes in the intensity of the structure itself. Nevertheless, a reduction in the overall contrast in the image as a result of despeckling will contribute to a reduction of the CNR.
3. The equivalent number of looks (ENL) [49] is defined as $ENL = SSR^2 = \bar{I}_{IL}^2/\sigma_{IL}^2$, and provides an estimation of the number of uncorrelated speckle patterns used to achieve a given image quality (see Eq. (2)). As such, this is a metric related solely to the quality of the despeckling. To guarantee enough data for proper statistics, a larger dedicated background ROI was used to calculate the ENL.

Figure 7(a) presents three orthogonal views of the phantom immersed in water. The three ROIs with cellulose structure (A, B, C) are indicated in solid-line boxes and three homogeneous (background) ROIs (A_{bg} , B_{bg} , C_{bg}) are indicated in dashed-line boxes. The contrast and CNR were evaluated inside (A, B, C), while the ENL was calculated in their corresponding (A_{bg} , B_{bg} , C_{bg}). Figures 7(b)–7(e) show the yz cross-sectional view containing ROIs B and B_{bg} after despeckling at 0.3% intralipid concentration. The image processed with OOPA [Fig. 7(c)] exhibits the classic deterioration of out-of-plane resolution. BM3D [Fig. 7(d)] and PNLM [Fig. 7(e)] exhibit remnant speckle fluctuations throughout the image, while TNode 2D and TNode 3D show a similar preservation of the cellulose structure with a superior degree of speckle reduction.

Figure 8(a) shows the despeckling results for ROI B at two intralipid concentrations. An increase in the intralipid concentration reduces the contrast with the cellulose. In order to assess the resolution loss after despeckling, we measured the full width at half-maximum (FWHM) of the cellulose structure across the lines indicated in Fig. 8(a) and (b) at 0.3% intralipid concentration. The FWHM for the singlelook tomograms was measured at a location away from a speckle null. The FWHM in water corresponds to the FWHM expected for the PSF of the objective lens, once the confocality is taken into account ($35\mu\text{m}/\sqrt{2}/\sqrt{2\ln 2} = 21\mu\text{m}$ FWHM). Our results indicate that the FWHM increases from $\approx 20\mu\text{m}$ to $\approx 30\mu\text{m}$ after immersion in intralipid. However, it does not further increase after despeckling with BM3D, PNLM and TNode. The increase in FWHM between water and intralipid could be attributed to the change in the liquid-air interface geometry given by surface tension, which would deform the cellulose structure between the two measurements. In the case of OOPA, as expected, the FWHM of the structure is increased by a factor of ≈ 2 compared with other despeckling algorithms. This is not a definitive conclusion, because assessing resolution in scattering media is a complex task. Alternative experiments, such as evaluating the ability to separate two sources inside a scattering medium, implies the use of phantoms with extremely large homogeneous regions, which facilitate despeckling unrealistically.




Fig. 7. Tomogram of the phantom: (a) cellulose structure immersed in water, (b) single cross-section of the cellulose structure immersed in 0.3% intralipid, and results after despeckling with (c) OOPA, (d) BM3D, (e) PNLM, (f) TNode 2D and (g) TNode 3D. Scale bars are 300 μm for all axes.

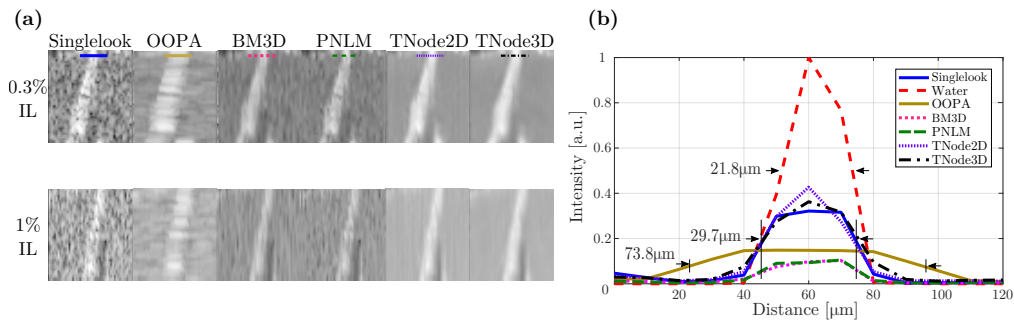


Fig. 8. Region B in Fig. 7 despeckled with different concentrations of intralipid (IL) (a), and intensity line profiles (in linear scale) at 0.3% intralipid concentration (b).

We measured the contrast, CNR and ENL for the two intralipid concentrations. The numerical results are tabulated in Table 1. Overall, the contrast drops less than 1 dB for most algorithms and therefore this performance metric is comparable for all of them. For the 0.3% intralipid, the visually strong performance of TNode is numerically supported by the results of the CNR and ENL. TNode 3D had consistently the highest CNR among all algorithms, meaning that details in the image are preserved and are visually more evident after despeckling. Interestingly, despite the drop in contrast and CNR of the singlelook tomogram going from 0.3% to 1% intralipid, TNode 3D is able to generate a despeckled ROI with a similar CNR for both concentrations. In terms of pure speckle suppression, TNode significantly outperforms all other techniques as quantified by the ENL. We believe our phantom experiments—as well as the ophthalmic data—support the conclusion that the possible resolution loss of TNode in a realistic sample, with intertwined regions of fine structures and homogeneous scattering, is very small. Together with its excellent despeckling performance as quantified by the ENL, its improvement of CNR and its preservation of contrast, TNode outperforms all other state-of-the-art techniques for despeckling OCT scans, improving the interpretation of OCT volumes of scattering media.

Table 1. Comparison of the performance of despeckling techniques in terms of contrast, CNR and ENL for the three concentrations.

Region Projection	Contrast [dB]			CNR			ENL		
	A <i>xy</i>	B <i>yz</i>	C <i>xz</i>	A <i>xy</i>	B <i>yz</i>	C <i>xz</i>	A _{bg} <i>xy</i>	B _{bg} <i>yz</i>	C _{bg} <i>xz</i>
0.3% IL									
<i>Singlelook</i>	<i>10.1</i>	<i>12.8</i>	<i>10.4</i>	<i>1.4</i>	<i>1.5</i>	<i>1.3</i>	<i>1.2</i>	<i>1.1</i>	<i>1.3</i>
OOPA	7.2	13.8	7.4	1.6	2.2	1.9	8.1	8.1	8.6
BM3D	9.2	11.3	9.4	1.8	2.1	1.9	8.3	5.9	8.4
PNLM	9.2	11.5	9.4	1.7	1.8	1.7	12	8.3	10
TNode 2D	9.5	12.7	10.0	2.1	2.2	2.1	20	72	61
TNode 3D	9.5	12.7	10.3	2.4	2.5	2.1	12	99	102
1% IL									
<i>Singlelook</i>	<i>6.6</i>	<i>8.1</i>	<i>7.6</i>	<i>1.0</i>	<i>1.1</i>	<i>1.2</i>	<i>1.2</i>	<i>1.1</i>	<i>0.9</i>
OOPA	5.4	8.6	5.6	1.5	1.4	2.2	7.9	7.7	5.1
BM3D	5.8	7.0	6.6	1.7	1.7	1.8	7.7	7.1	5.0
PNLM	5.3	6.5	6.0	1.4	1.4	1.8	6.5	6.9	4.7
TNode 2D	5.4	7.1	6.6	1.8	1.6	2.6	43	85	44
TNode 3D	5.5	6.9	6.8	2.4	1.8	2.9	30	100	65

Singlelook parameters are indicated in *italics*. The best parameter value for each region and quality metric among all despeckling techniques is highlighted in purple. Regions A, B, C and backgrounds A_{bg}, B_{bg} and C_{bg} are those defined in Fig. 7.

4. Conclusion

In conclusion, we have presented a novel despeckling technique for OCT, TNode, which is based on the non-local means algorithm with proper speckle statistics. TNode uses a 3D similarity window that leverages the volumetric nature of OCT data and accounts for the full three-dimensional geometry of structures normally found in OCT imaging. By using the generalized likelihood ratio to retrieve the volumetric patch similarity, we used a weighted maximum likelihood estimation to compute the speckle-free intensity tomogram, while permitting the use of single and compound tomograms with multiple speckle realizations. By accounting for the high dynamic range of OCT tomograms, we achieved homogeneous reduction of speckle at all depths.

In the retinal data set, visual comparison of the results obtained with different despeckling techniques demonstrated the strong performance of TNode in terms of resolution and speckle reduction. Cross-sectional and *en-face* views after despeckling showed the preservation of resolution in all dimensions, while conserving even speckle-sized structures. In singlelook and multilook tomograms the contrast improvement accomplished with TNode was evident. Even when applied to OCT data having significant motion artifacts, such as images acquired endoscopically in the gastrointestinal tract, TNode showed excellent speckle reduction while preserving resolution. With the phantom data, we demonstrate the potential of TNode for resolution and structure preservation and reduction of speckle, by using different metrics to quantitatively evaluate the quality of filtered images.

Since speckle statistical properties are common to all coherent imaging techniques, TNode could be used in other volumetric techniques, such as ultrasound imaging. However, we believe its strongest impact could come with its use as a pre-processing step in advanced OCT signal processing. For example, OCT intensity-based tractography [50–52] could greatly benefit from a resolution-preserving despeckling technique to improve the tracking process. Polarization-sensitive OCT in the Stokes formalism could also greatly benefit from TNode [53–56], where it

could mitigate the negative effect of speckle nulls in retardation and optical axis calculations and reduce the associated lateral resolution degradation.

5. Supporting material

We have made a Matlab implementation of TNode, [Code 1](#), available at Ref. [57] as well as on <http://octresearch.org>, including the full source code of TNode 2D and TNode 3D, as well as exemplary raw datasets for testing purposes.

Funding

National Institutes of Health (NIH) (P41 EB015903 and K25 EB024595); Universidad EAFIT (767-000063); COLCIENCIAS (761).

Acknowledgments

The authors would like to thank Heidelberg Engineering for access to the Spectralis OCT2 system and assistance with imaging, and NinePoint Medical Inc. for access to the clinical registry database.

Disclosures

The authors declare that there are no conflicts of interest related to this article.