arXiv:2010.15682v1 [eess.IV] 29 Oct 2020

Maximum A Posteriori Signal Recovery for Optical Coherence Tomography Angiography Image Generation and Denoising

LENNART HUSVOGT,[1,2*] STEFAN B PLONER,[1] SIYU CHEN,[2] DANIEL STROMER,[1,2] JULIA SCHOTTENHAMML,[1] A YASIN ALIBHAI,[3] ERIC MOULT,[2] NADIA K WAHEED,[3] JAMES G FUJIMOTO,[2] AND ANDREAS MAIER[1]

[1]*Pattern Recognition Lab, Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany*
[2]*Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, USA*
[3]*New England Eye Center, Tufts School of Medicine, Boston, USA*
[*]*lennart.husvogt@fau.de*

Abstract: Optical coherence tomography angiography (OCTA) is a novel and clinically promising imaging modality to image retinal and sub-retinal vasculature. Based on repeated optical coherence tomography (OCT) scans, intensity changes are observed over time and used to compute OCTA image data. OCTA data are prone to noise and artifacts caused by variations in flow speed and patient movement. We propose a novel iterative maximum a posteriori signal recovery algorithm in order to generate OCTA volumes with reduced noise and increased image quality. This algorithm is based on previous work on probabilistic OCTA signal models and maximum likelihood estimates. Reconstruction results using total variation minimization and wavelet shrinkage for regularization were compared against an OCTA ground truth volume, merged from six co-registered single OCTA volumes. The results show a significant improvement in peak signal-to-noise ratio and structural similarity. The presented algorithm brings together OCTA image generation and Bayesian statistics and can be developed into new OCTA image generation and denoising algorithms.

© 2022 Optical Society of America under the terms of the OSA Open Access Publishing Agreement

1. Introduction

Over the last years, optical coherence tomography angiography (OCTA) has become increasingly popular in research and clinical imaging. It allows micrometer-resolution imaging of retinal vessels and the choriocapillaris, the latter which is challenging to image using other existing methods [1]. In addition to that, OCTA is non-invasive compared to other methods, such as fluorescein angiography, which requires the injection of a contrast agent and carries the risk of anaphylactic shock [2]. OCTA is in clinical use and available in commercial devices by various manufacturers.

Two commonly used methods to generate OCTA data are amplitude decorrelation (AD) and speckle variance (SV). In both cases, optical coherence tomography (OCT) B-scans of the same location are acquired in immediate succession and changes in intensity over time serve as the basis on which OCTA is computed [3,4]. To the best of our knowledge, denoising of OCTA images is usually performed by applying basic filters such as a median filter and background thresholding *to the measured data*. While these approaches already achieve significant improvement with simple methodology, their underlying assumptions do not accurately apply to the OCTA image formation process. This discrepancy leaves plenty of room for further improvement of OCTA image quality without imposing changes to the acquisition protocol or the device.

Reconstruction of medical image data has made substantial progress in recent years, which was achieved by inverse problem formulations, which map the measurements backwards through

the image formation process to the true observed specimen. Such problems can be solved via iterative reconstruction, where the result image is found by optimizing an objective function comprised of two major parts: First, a data consistency term that ensures that the reconstructed data closely represents the inversely mapped measurements. Secondly, regularization terms are used to improve image quality by enforcing certain properties. Here, the formulation as inverse problem has the crucial advantage that such constraints can be enforced not only on the measurements, but also directly in the reconstructed data. Applications of this principle range through the entire field of medical imaging: In magnetic resonance imaging (MRI), compressed sensing-based regularization is used to enforce sparseness of the reconstructed data, achieving great reductions in imaging time [5]. Closely related, in computed tomography (CT), homogeneity within tissue classes was modeled by regularizers, achieving significant dose reduction at equal image quality [6–10].

A common way for setting up the objective function is to use Bayesian statistics, which has been used in the context of OCT before to reduce speckle noise [11]. More recently, Ploner et al. introduced a Bayesian signal model for OCTA which uses maximum likelihood estimation (MLE) [12]. While this initial model provided a theoretical basis for the data consistency term, image quality was not improved, because the model still lacks the key advantage of the inverse problem formulation: the inclusion of prior information on the reconstructed signal via regularization terms. This would extend the MLE to a complete Bayesian maximum a posteriori (MAP) estimation. In this paper we present, to the best of our knowledge, the first MAP image reconstruction algorithm for OCTA: we extend the OCTA MLE model by Ploner et al. to a MAP estimate and use wavelet shrinkage and total variation minimization as regularizers.

The evaluation of an algorithm like this poses some challenges though. OCT scans suffer from a wide range of motion artifacts because scan acquisition can take several seconds and patient eyes exhibit involuntary eye motion. This causes severe difficulties in acquiring any ground truth (GT) data that exhibit few motion artifacts and which have sufficient overlap with test data. To address this problem, we generated a low-noise GT OCTA volume by co-registering and merging six individual OCTA volumes with each OCTA volume generated from scans using ten B-scan repetitions each. One of the six volumes also serves as test volume. Because of the co-registration with the GT, reconstruction results can be deformed to match it and error metrics between the two can be computed.

This paper is structured as follows: in the next section we provide a description of the MAP estimate, comprised of the MLE it is based on and the newly introduced regularizers, and the reconstruction algorithm. It is followed by the experiment section that explains how the OCTA ground truth was generated. Results and discussion sections follow and the paper ends with a conclusion.

2. Method

The reconstruction is based on a MAP estimation, which is introduced first. The OCTA MLE signal model and regularizers are described afterwards, followed by the reconstruction algorithm.

2.1. Maximum A Posteriori Probability Estimation

Image reconstruction arose from the need to reconstruct some data $\mathbf{X}$ from a limited number of measurements $\mathbf{Y}$. For instance, where a linear sampling operation $\mathbf{A} \in \mathbb{R}^{M \times N}$ maps from a finite signal $\mathbf{X} \in \mathbb{R}^{N}$ to

$$\mathbf{Y} = \mathbf{A}\mathbf{X} + \mathbf{e}\,, \tag{1}$$

with $\mathbf{Y} \in \mathbb{R}^{M}$, while $\mathbf{e} \in \mathbb{R}^{M}$ models the impact noise has on the measured signal $\mathbf{Y}$. $\mathbf{A}$ models the sampling of $\mathbf{X}$ to the observed signal $\mathbf{Y}$ and is often called the *system matrix*. The goal is

to reconstruct $\mathbf{X}$ based on the measurements $\mathbf{Y}$ which is an *inverse problem*. One approach to solve this inverse problem in the case of OCTA, is to treat it as a MAP estimate [13]. By interpreting $\mathbf{X}$ and $\mathbf{Y}$ as random variables with their own distributions, the MAP estimate has the goal to find an OCTA volume $\hat{\mathbf{X}}$ which maximizes its probability given the OCT scan $\mathbf{Y}$. This is defined as

$$\hat{\mathbf{X}} = \arg\max_{\mathbf{X}} \mathrm{P}\left(\mathbf{X}|\mathbf{Y}\right) = \arg\max_{\mathbf{X}} \frac{\mathrm{P}\left(\mathbf{Y}|\mathbf{X}\right)\mathrm{P}\left(\mathbf{X}\right)}{\mathrm{P}\left(\mathbf{Y}\right)} = \arg\max_{\mathbf{X}} \mathrm{P}\left(\mathbf{Y}|\mathbf{X}\right)\mathrm{P}\left(\mathbf{X}\right) \ . \tag{2}$$

It is possible to ignore $\mathrm{P}\left(\mathbf{Y}\right)$ because it does not depend on $\mathbf{X}$. $\mathrm{P}\left(\mathbf{Y}|\mathbf{X}\right)$ is the MLE for $\mathbf{Y}$ and $\mathrm{P}\left(\mathbf{X}\right)$ is the *prior probability* of $\mathbf{X}$. If $\mathrm{P}\left(\mathbf{Y}|\mathbf{X}\right)$ is a multivariate Gaussian distribution with mean $\mu = \mathbf{A}\mathbf{X}$ and covariance matrix $\Sigma = \mathbf{I}$, maximizing the log probability yields

$$\arg\max_{\mathbf{X}} \ln \mathrm{P}\left(\mathbf{Y}|\mathbf{X}\right)\mathrm{P}\left(\mathbf{X}\right) = \arg\max_{\mathbf{X}} \ln\left[\exp\left(-\frac{1}{2}\left(\mathbf{Y}-\mathbf{A}\,\mathbf{X}\right)^{T}\Sigma^{-1}\left(\mathbf{Y}-\mathbf{A}\,\mathbf{X}\right)\right)|\Sigma|^{-\frac{1}{2}}\mathrm{P}\left(\mathbf{X}\right)\right] \tag{3}$$

$$= \arg\max_{\mathbf{X}} -\frac{1}{2}\left(\mathbf{Y}-\mathbf{A}\,\mathbf{X}\right)^{T}\left(\mathbf{Y}-\mathbf{A}\,\mathbf{X}\right) + \ln \mathrm{P}\left(\mathbf{X}\right) \tag{4}$$

$$= \arg\min_{\mathbf{X}} \frac{1}{2}\|\mathbf{Y}-\mathbf{A}\,\mathbf{X}\|_2^2 + \lambda\,\mathrm{R}\left(\mathbf{X}\right) \ . \tag{5}$$

The MAP estimate yields a least squares problem with a regularizer $\mathrm{R}\left(\mathbf{X}\right)$ in lieu for $\ln \mathrm{P}\left(\mathbf{X}\right)$ to include prior knowledge about $\mathbf{X}$ or as a regularizer. In cases such as described by Manhart et al. and Neukirchen et al. where the system matrix $\mathbf{A}$ is ill-conditioned and the data dimensions are large, the Landweber iteration was successfully used to find a solution for $\mathbf{X}$ [14, 15].

However, because OCTA computations are non-linear and non-invertible, the above Gaussian model for the maximum log likelihood estimate does not apply. This is where the work of Ploner et al. becomes relevant, in which the authors describe maximum log likelihood estimates for AD, SV, and interframe variance (IFV) which are used as $\mathrm{P}\left(\mathbf{Y}|\mathbf{X}\right)$ [12]. The next section will summarize the OCTA signal models and MLE used for reconstruction.

2.2. Probabilistic OCTA Signal Models

As noted in the introduction, intensity-based OCTA functions are based on changes in intensity over time for a given voxel. For an OCT scan where each B-scan was repeatedly scanned N times, y_i denotes the measured amplitude for scan repetition $i \in [1, 2, \ldots, N]$. AD is the variance of the pair-wise difference amplitudes y_i and y_{i+1} normalized by

$$\frac{1}{\sqrt{y_i^2 + y_{i+1}^2}} \ . \tag{6}$$

We assume that $\mu = 0$ for the differences between y_i and y_{i+1} because both are identically distributed. The likelihood to acquire the measurements y_1 to y_N is then

$$L_{\mathrm{AD}} = \prod_{i=1}^{N-1} \mathrm{P}\left(\left.\frac{y_i - y_{i+1}}{\sqrt{y_i^2 + y_{i+1}^2}}\right| \mu = 0, \sigma^2 = x\right) = \prod_{i=1}^{N-1} \frac{1}{\sqrt{2\pi x}} \exp\left(-\frac{\frac{(y_i - y_{i+1})^2}{y_i^2 + y_{i+1}^2}}{2x}\right) \ . \tag{7}$$

Using the maximum log-likelihood estimation one obtains

$$\arg\max_{x} \log L_{\mathrm{AD}} = \arg\max_{x} (N-1) \cdot \log\left(\frac{1}{\sqrt{2\pi x}}\right) - \sum_{i=1}^{N-1} \frac{\frac{(y_i - y_i)^2}{y_i^2 + y_{i+1}^2}}{2x} \ , \tag{8}$$

with the derivative over x being

$$\frac{d \log L_{\mathrm{AD}}}{dx} = \frac{-(N-1) \cdot x + \sum_{i=1}^{N-1} \frac{(y_i - y_{i+1})^2}{y_i^2 + y_{i+1}^2}}{2x^2} . \tag{9}$$

Setting the derivative to zero and solving for x results in the definition of AD [3]

$$x = \frac{1}{N-1} \sum_{i=1}^{N-1} \frac{(y_i - y_{i+1})^2}{y_i^2 + y_{i+1}^2} = 1 - \frac{1}{N-1} \sum_{i=1}^{N-1} \frac{y_i \cdot y_{i+1}}{\frac{1}{2}\left(y_i^2 + y_{i+1}^2\right)} . \tag{10}$$

Based on SV, Ploner et al. proposed a new angiography formula, IFV [12]. Compared to SV, IFV computes the variance of the differences between voxel values over time but without the normalization used in AD

$$L_{\mathrm{IFV}} = \prod_{i=1}^{N-1} \mathrm{P}\left(y_i - y_{i+1} \middle| \mu = 0, \sigma^2 = x\right) = \prod_{i=1}^{N-1} \frac{1}{\sqrt{2\pi x}} \exp\left(-\frac{(y_i - y_{i+1})^2}{2x}\right) . \tag{11}$$

Again, as with AD, the assumption is that $\mu = 0$. After applying the log function

$$\log L_{\mathrm{IFV}} = (N-1) \cdot \log\left(\frac{1}{\sqrt{2\pi x}}\right) - \sum_{i=1}^{N-1} \frac{(y_i - y_{i+1})^2}{2x} , \tag{12}$$

and taking the first degree derivative with respect to x

$$\frac{d \log L_{\mathrm{IFV}}}{dx} = -\frac{(N-1) \cdot x + \sum_{i=1}^{N-1} (y_i - y_{i+1})^2}{2x^2} , \tag{13}$$

the maximum can then be found at

$$x = \frac{1}{N-1} \sum_{i=1}^{N-1} (y_i - y_{i+1})^2 . \tag{14}$$

The MLE for SV can be found in the supplemental document.

By including prior knowledge in form of a regularizer, the MLE can be extended into a MAP estimate. For this paper we evaluated two different regularizers: wavelet shrinkage, and total variation minimization. Wavelet shrinkage works by transforming the OCTA volume into wavelet domain and setting the wavelet coefficients below a certain threshold to zero [16]. For total variation regularization, in which the variation over the signal is minimized, we used the algorithm by Chambolle et al. [17]. These regularizers are also commonly used in compressed sensing to enforce sparsity. The relationship of the algorithm described in this paper to compressed sensing is described in the supplementary.

2.3. Reconstruction Algorithm

For the purposes of this paper, both structural and angiography OCT volumes consist of N_b B-scans. A B-scan is a 2D image composed of N_a A-scans, which in turn contain N_s samples or pixels. Furthermore, in order to generate OCTA data, every structural B-scan is scanned N_r times. A structural OCT volume is thus $\mathbf{Y} \in \mathbb{R}^{N_b \cdot N_r \cdot N_a \cdot N_s}$ while an OCTA volume is $\mathbf{X} \in \mathbb{R}^{N_b \cdot N_a \cdot N_s}$.

Even though AD and IFV are non-linear and non-invertible, it is possible to compute an OCTA reconstruction of $\mathbf{X}$ using the MAP estimate. For this we propose a novel data consistency term, based on the MLE in section 2.2 and the MAP estimate in Eq. (2). The goal is to find an OCTA

Objective: Find $\hat{\mathbf{X}} = \arg\max_{\mathbf{X}} \mathrm{P}(\mathbf{X}|\mathbf{Y})$
Input: Supply the following components and parameters:

- Step size λ
- Regularizer R and parameters
- Total number of iterations N_{iter} and number of iterations N_{reg} for the regularizer

Initialization: Set $\mathbf{X}_0 = \mathbf{A}_{\text{alg}}(y)$
Loop:

- Update the OCTA volume $\mathbf{X}_{k+1} = \mathbf{X}_k + \lambda \nabla \mathrm{L}_{\text{alg}}(\mathbf{X}_k)$.
- Every N_{reg} iterations apply the regularizer $\mathbf{X}_{k+1} = \mathrm{R}(\mathbf{X}_k)$.

End of Loop:
Result: The output $\hat{\mathbf{X}}$

Algorithm 1: Description of the OCTA MAP reconstruction algorithm. The current estimate $\mathbf{X}_k$ is iteratively updated using the element-wise derivative $\nabla \mathrm{L}_{\text{alg}}(\mathbf{X}_k)$ while a regularizer is applied every N_{reg} iterations. The result is $\hat{\mathbf{X}}$.

volume $\hat{\mathbf{X}}$ that maximizes the probabilities at each voxel position that fit the OCT observations $\mathbf{Y}$ by maximizing

$$\hat{\mathbf{X}} = \arg\max_{\mathbf{X}} \log\left[\mathrm{P}(\mathbf{Y}|\mathbf{X})\,\mathrm{P}(\mathbf{X})\right] \tag{15}$$

$$= \arg\max_{\mathbf{X}} \log \mathrm{P}(\mathbf{Y}|\mathbf{X}) + \mathrm{R}(\mathbf{X}) \ . \tag{16}$$

Using the observations in $\mathbf{Y}$ we can find the OCTA volume $\hat{\mathbf{X}}$ that maximizes the posterior probability of $\mathbf{Y}$. Assuming that $\mathbf{X}$ is uniformly distributed, $\mathrm{P}(\mathbf{X})$ becomes constant, and can be disregarded. A regularizer is added in its place that fulfills a similar function, enforcing prior knowledge about $\mathbf{X}$. Finally, we search $\hat{\mathbf{X}} = \arg\max_{\mathbf{X}} \mathrm{P}(\mathbf{X}|\mathbf{Y})$ using the Landweber iteration

$$\mathbf{X}_{k+1} = \mathbf{X}_k + \lambda \nabla \mathrm{L}_{\text{alg}}(\mathbf{X}_k) \ , \tag{17}$$

with λ being the relaxation factor or step size and $\nabla \mathrm{L}_{\text{alg}}$ the element-wise derivative for $\text{alg} \in \{\text{AD}, \text{IFV}\}$ from Eq. (9) and Eq. (13) [18]. The regularizer is applied to the currently reconstructed volume $\mathbf{X}_k$ every N_{reg} iterations [14, 15]. Fig. 1 shows the complete reconstruction algorithm, including the required parameters.

3. Evaluation

This section describes the experiments, the data used, including the generation of the GT data, and patient data.

3.1. Experiments

Two experiments using the GT and patient data were conducted:

1. Comparison of reconstruction results with the GT for AD and IFV OCTA formulas using wavelet shrinkage and total variation minimization as regularizers for different numbers of repeated samples.

2. Qualitative evaluation of two scans of patient's eyes.

The wavelet shrinkage regularizer is used with a threshold of 5e-4 using Haar wavelets for a total of 1000 iterations. The total variation regularizer uses a denoising weight of 1e-4 and ten iterations for the regularizer. The reconstruction was run for 2000 iterations. For the reconstruction of AD data, a step size of $\lambda = 5e-6$ was chosen and for IFV data $\lambda = 3e-6$. The number of samples in experiment one refers to the number of repeated OCT B-scans used in computing the OCTA signal. The reconstruction and experiments were implemented in Python using NumPy and scikit-learn [19, 20].

3.2. Ground Truth Data

In order to evaluate the performance of our method, we decided to compare the reconstruction results with GT OCTA data. The GT gold standard for structural OCT is averaged B-scans [21]. However, patient motion and its associated artifacts pose a significant challenge for the acquisition of a GT. In order to generate GT OCTA data with minimal motion artifacts, we chose the motion correction algorithm by Kraus and Ploner et al. [22–25]. The motion correction algorithm requires at least two OCT volumes which were raster-scanned with orthogonal fast-scan directions. It is based on the principle of non-rigid registration, i.e. the motion correction algorithm estimates a 3D displacement field for each scan such that similarity between the displaced inputs is maximized, while using both OCT and OCTA for registration. After deforming the volumes to match each other, the motion correction algorithm is able to produce a merged and interpolated result by performing a weighted average per A-scan. This allows to correct for gaps and overlaps in the input volumes as long as the missing data is contained in another volume. The more volumes are used as inputs, the higher the signal-to-noise ratio (SNR) of the motion-corrected and merged result, which includes both OCT and OCTA channels [24, 25]. It is important to note that the merging step of the motion correction algorithm is only used for generating the GT. Since the chosen test volume is also co-registered with the GT it is possible to compute comparative metrics. The merged GT OCTA volume with high SNR allows to compare the reconstructed results. Although noise is substantively reduced, the GT OCTA volume is not perfect since it still exhibits some low levels of noise. Because of this it is more of a silver than a gold standard, but still useful as a low-noise, motion artifact-free GT to compare reconstruction results against. The overall process of obtaining the data necessary for a GT is very time consuming. A sufficient number of OCT scans needs to be acquired, which have to be in focus, suffer from a minimal amount of motion artifacts, have similar alignment, and are scanned using orthogonal scan patterns. Only five to six scans could be acquired in a single sitting due to eye fatigue of the volunteer. Of 32 scans that were acquired over the course of a week, six were suitable for the generation of a GT.

The data were acquired from a healthy 24 year old male volunteer. The OCT system used is a ultrahigh speed swept source OCT research prototype developed at the Massachusetts Institute of Technology and used by the New England Eye Center at Tufts Medical Center in Boston. The scanned field size is 3×3 mm with 400 A-scans by 400 B-scans. The system runs at a scan rate of 600 kHz and scanned each B-scan ten times with an inter-scan time of ~0.83 µs. Power on cornea is ~4.5 mW. Axial and lateral resolution are ~8 µm and ~20 µm respectively. Six scans were acquired with alternating scan patterns. Three scans were acquired using an X-fast scan pattern while the other three were acquired using a Y-fast scan pattern. Having both X and Y-fast scans ensures the necessary scan orthogonality for the motion correction algorithm. During each scan, every B-scan was repeatedly scanned ten times. This allows to compute OCTA data using the angiography formulas. Fig. 1 illustrates the data used and the relationship between the GT and test data. The six scans are depicted in the center with the scan used for testing highlighted blue. All six were co-registered and merged to obtain the OCTA GT visible on the right. Because the test scan, like the other five scans, was scanned with ten B-scan repeats,

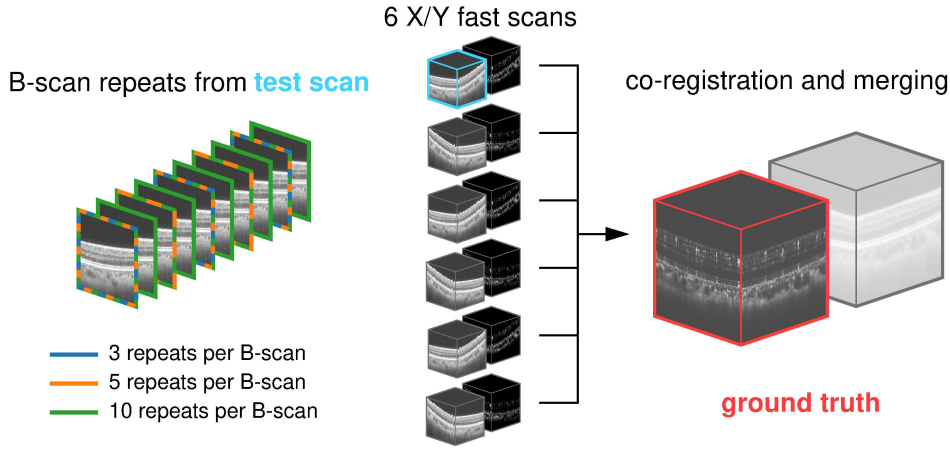

Fig. 1. Use of test data during the evaluation. Six OCT scans were acquired for testing (center). OCTA volumes were computed from the test scans and structural/angiographic volume pairs were co-registered and merged. The merged OCTA volume serves as ground truth (red, right). Every one of the six test scans was acquired using ten repeats per B-scan. One of these scans is used as test scan (blue). Of the ten repeats, reconstruction was evaluated using three, five, and all ten repeats per B-scan (left).

we decided to test reconstruction using three, five and all ten B-scans, which is illustrated on the left. In the case of three repeats, the first, fifth, and ninth repeat were used. For five repeats, every second repeat was used.

To compare the reconstructed OCTA data with the GT, en face projections of the respective volumes were generated. OCTA data are commonly viewed as en face projections, because it allows clinicians to see the retinal vasculature or the choriocapillaris. In order to obtain clean OCTA en face projections, the volumes were segmented using segmentation software developed by Schottenhamml et al. [26]. The space between the outer nerve fiber layer and the inner plexiform layer was projected using a 98^{th} percentile projection. This corresponds to the superficial capillary plexus in the retina. In the case of AD, a background thresholding is performed to avoid excess noise in the en face image. The en face projections of the reconstructed OCTA volume can then be compared to the en face projection of the GT and peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) is computed. We chose SSIM in addition to PSNR because it is not based on the absolute error, but is based on the human visual perception of changes in structure [27].

3.3. Patient Data

In order to not only test on the scans of one eye, we decided to apply the reconstruction algorithm to existing scans of patient eyes. These data were not acquired with the generation of a GT in mind, which is very time-consuming and labor-intensive, thus the reconstruction results can only be evaluated qualitatively. But it allows to observe reconstruction results on more OCTA scans where there is also pathology present.

Patient 1 is male and was 41 years old at the time of imaging and exhibits severe non-proliferative diabetic retinopathy (NPDR) in connection with diabetic macular edema. Patient 2 is female and was 88 years old at the time of imaging and shows dry age-related macular degeneration and severe NPDR in both eyes. The system used was a 400 kHz Swept Source OCT device. Field size is 3×3mm with 500 A-scans by 500 B-scans. Every B-scan was repeatedly scanned five times with an inter-scan of ~1.5 ms. Power on cornea was ~1.8 mW [1, 28].

	Repeats	Initial OCTA		Wavelet Shrinkage		Total Variation		Median Filter	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
AD	3	16.92	0.36	27.63	0.78	**29.43**	**0.84**	27.18	0.73
	5	20.96	0.55	28.64	0.81	**30.29**	**0.86**	27.25	0.73
	10	25.70	0.73	29.30	0.83	**30.50**	**0.86**	26.95	0.72
IFV	3	23.78	0.54	34.55	0.91	35.56	**0.93**	**36.03**	**0.93**
	5	27.33	0.72	35.81	0.93	**36.76**	**0.95**	36.55	0.93
	10	31.78	0.88	37.12	0.95	**37.76**	**0.96**	36.71	0.92

Table 1. Reconstruction results using the GT for comparison. The top rows show the results for AD and the bottom rows for IFV. Results for three, five, and ten scan repeats are sorted by row. PSNR and SSIM measures are provided for the initial OCTA volume, wavelet shrinkage and total variation reconstruction results and the median filtered initial volume. The final reported value corresponds to the final iteration result of the reconstruction. The best PSNR and SSIM values in each row are marked in bold.

4. Results

Fig. 2 shows PSNR and SSIM metrics for the reconstruction tests that compare against the GT. The colors indicate how many B-scan repeats were used during reconstruction. The two top rows show AD results and the bottom row IFV. The plots in the left column show results using wavelet shrinkage regularization (WS) and the plots in the right column show total variation minimization (TV) results. For AD and IFV, PSNR results are shown first, followed by SSIM results directly below. Reconstruction using ten scan repeats consistently performs best, although the largest improvements can be observed when using three repeats. IFV shows better results using both PSNR and SSIM. Fig. 3 shows 98^{th} percentile en face projections based on AD and IFV. The first row shows GT projections for AD and IFV. Enlarged sections of each en face image are shown to the right of the respective en face image. Rows are grouped by AD and IFV first, then by the initial OCTA volume, then WS and TV reconstruction results and, for comparison, median filtered results of the initial OCTA volumes using a 3D median filter using a $3 \times 3 \times 3$ kernel. Columns are grouped by the number of scan repeats used. Furthermore, Tab. 1 shows PSNR and SSIM results in more detail, including the 3D median filter results. Results for SV are included in the supplementary document.

5. Discussion

Fig. 2 shows PSNR and SSIM over reconstruction iterations while Tab. 1 shows the highest PSNR and SSIM values during reconstruction. A slight dipping of the curves during reconstruction might indicate over-regularization, but, as can be seen in Fig. 2, a decrease in PSNR does not necessarily correspond with a decrease in SSIM.

It is apparent that using more B-scan repeats, or samples, generally leads to better PSNR and SSIM while the initial OCTA and early reconstruction show larger discrepancies than the final results. The raw, unreconstructed en face projections for AD have PSNRs of ~16.92 dB, ~20.96 dB, and ~25.70 dB for three, five, and ten repeats respectively. Reconstruction results using WS regularization show PSNRs of ~27.63 dB, ~28.64 dB, and ~29.30 dB, while TV regularized PSNRs are ~29.43 dB, ~30.29 dB, and ~30.50 dB. To put it differently, using more repeats generally leads to better results but with diminishing returns during reconstruction. This is an important consideration, since more B-scan repeats give more information to compute OCTA

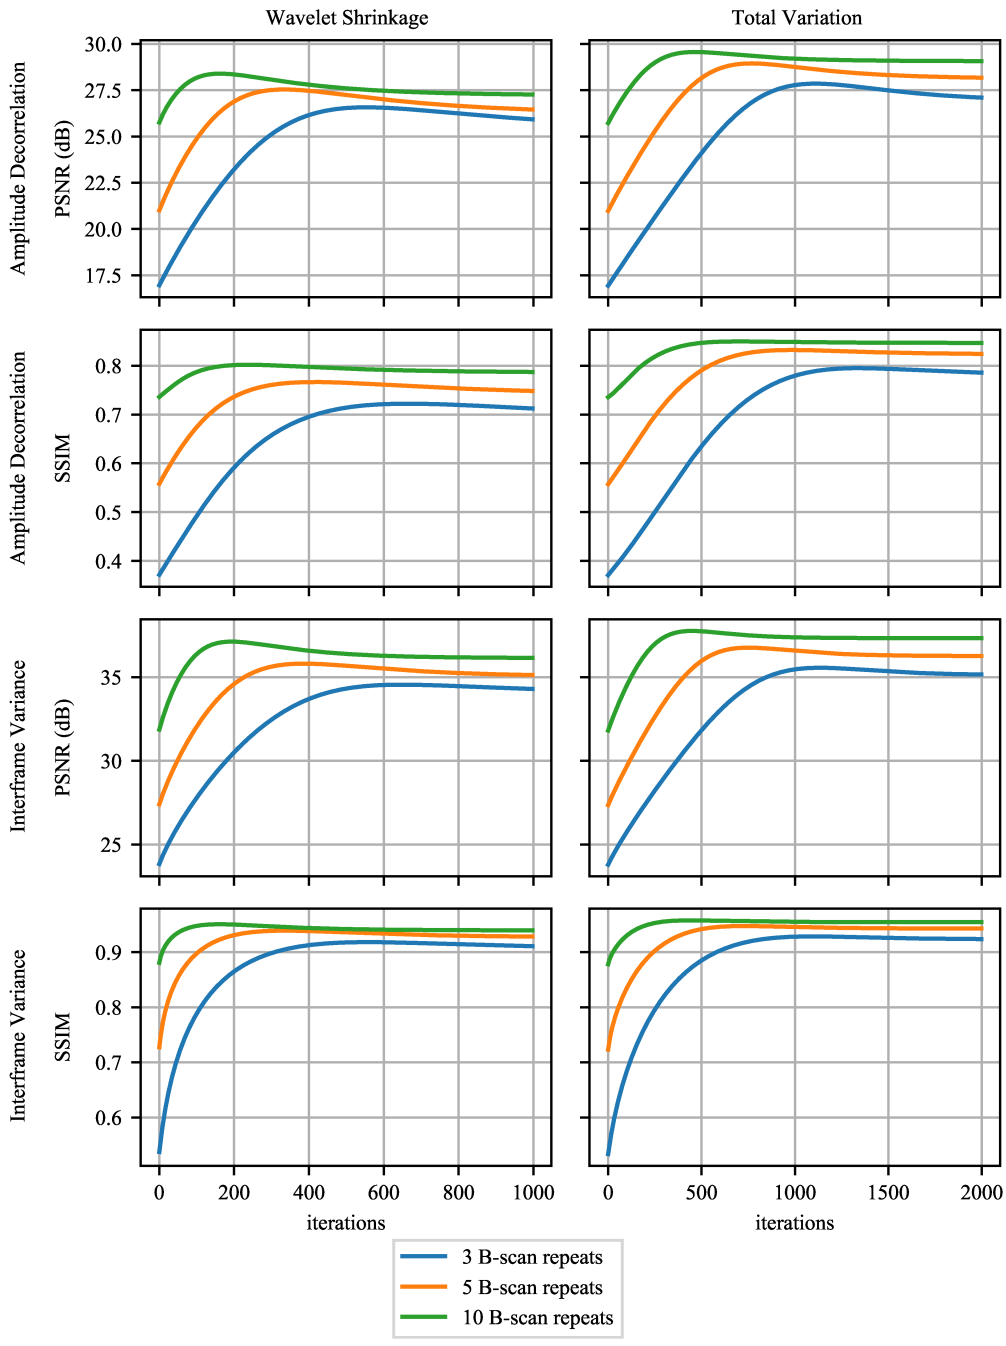

Fig. 2. Peak signal-to-noise ratio (PSNR) and structural similarity (SSIM) during reconstruction. The color indicates whether three, five, or ten repeated B-scans were used for reconstruction. The rows show PSNR for amplitude decorrelation and interframe variance, the columns indicate the regularizer.

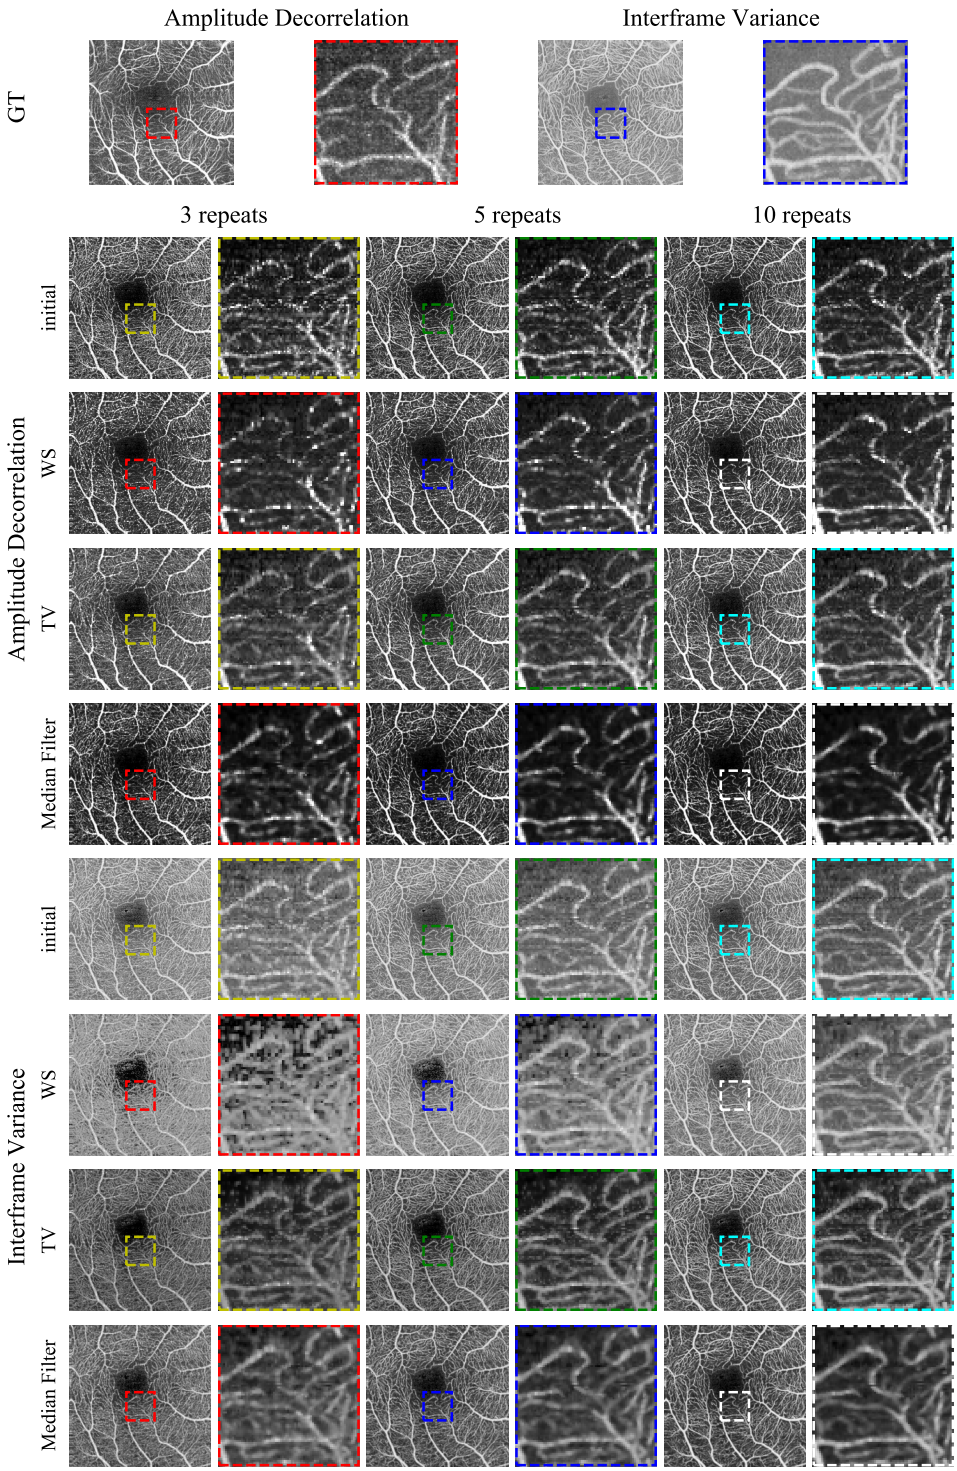

Fig. 3. En face projections of the GT and reconstruction results. More detailed areas are shown in enlargements next to the en face images. Rows are grouped by AD and IF, and subdivided by initial OCTA, wavelet shrinkage and total variation reconstructions and median filter for comparison. Columns are grouped by the number of scan repeats used.

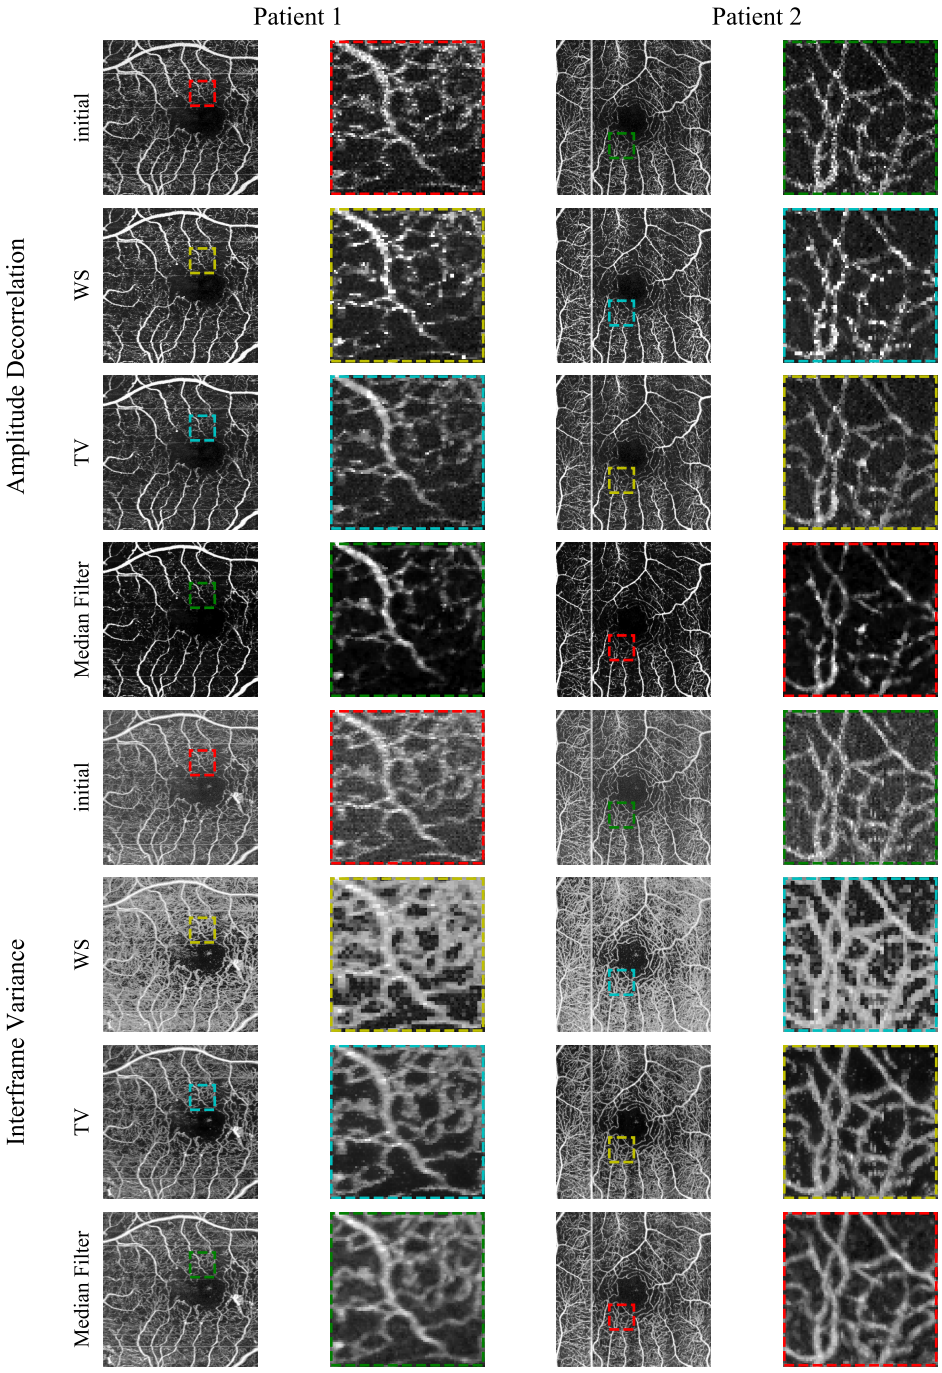

Fig. 4. En face projections of reconstructed patient data. Rows are grouped by AD and IFV, then by the initial OCTA volume, followed by wavelet shrinkage and total variation reconstructions. Columns are grouped by patient, enlarged areas are directly to the right of their projections.

from, while longer scan times make the entire scan more susceptible to motion artifacts. This is visually supported by the en face images in figure Fig. 3 where there is hardly a discernible difference between the projections for 5 and 10 repeats, while the difference to 3 repeats is easily apparent in the form of more noise. Commercial systems use two or three B-scan repeats where the benefits of reconstruction is highest, but testing reconstruction using more repeats shows that the algorithm is working correctly. One observation of interest is that SSIM of the median filtered result slightly decreases from five to ten repeats.

IFV results achieve both higher PSNR in range of ~34 dB to ~37.7 dB and SSIM of more than 0.9 when compared to AD with ~27.6 dB to ~30.50 dB and ~0.8 respectively. Direct comparisons between IFV and AD results are of limited use however. Because of the normalization which is part of the AD computation, its results tend to be very noisy. Background thresholding can compensate for this, but can erase detail such as small capillaries. Comparing AD and IFV GT in Fig. 3 shows, despite background thresholding for AD, that more noise is visible in intercapillary areas and less capillary details. For this reason, IFV en face images start out at better PSNR and SSIM to begin with than AD images.

Comparing WS to TV regularization is more meaningful though. The measurements in Tab. 1 consistently show better results when using TV regularization. In the case of AD, PSNR and SSIM results range from ~27.63 dB to ~30.50 dB and ~0.78 to ~0.86. For IFV, these results are much closer with ~37.12 dB to ~37.76 and ~0.95 to ~0.96. When looking at the images in Fig. 3 the qualitative differences between WS and TV regularization can be easily seen. The median filtered PSNR and SSIM results, however, do not show an advantage when using IFV. Median filter PSNR is higher than 36 dB for all numbers of repeats. In the case of three repeats, both PSNR and SSIM indicate that the median filter performs better than TV or WS regularized reconstruction. When computing AD data though, both WS and TV regularization have better PSNR and SSIM values than the median filter comparison.

Fig. 3 shows en face projections of the GT and reconstruction results. GT projections for both AD and IFV show that those are not completely noise-free. Co-registering and merging six volumes does not result in a perfect gold standard GT. When comparing to the reconstruction results below, it becomes evident that the GT displays significantly less noise and can still serve as a silver-standard GT. First, we take a look at the AD en face projections in Fig. 3. The initial AD projections show significant amounts of noise which continually decrease the more repeats are used despite background thresholding. WS regularization seems to lead to less noisy intercapillary areas, although blood vessels are more faint, when compared to TV regularization. WS regularization also leads to isolated bright pixels in blood vessels. TV regularization my preserver more noise in the intercapillary areas compared to WS, but fine capillaries remain more visible. This is consistent with the results in Tab. 1. While, in the case of ten repeats, it can be argued that WS regularization provides improvements over the initial projection, the same is less obvious for the TV regularization. Although the median filtered results seem to provide more contrast between vessels and intercapillary areas at first glance, vessels are blurred and fine capillaries are harder to discern and less continuous.

Fig. 4 shows the en face results for the two patient scans. Again, WS regularization used in conjunction with AD leads to brightened sections in some blood vessels. TV regularization makes the best impression, while the median-filtered comparison is exceedingly blurry with the loss of finer vessels. With IFV data, WS regularization increases contrast but leads to irregular, jagged borders along the vessels. TV regularization looks best, with enhanced contrast and a better delineation between vessels and the areas in-between, while the median-filtered comparison is more blurry.

We do not provide a formal proof for the convergence of our algorithm, but the plots in Fig. 2 indicate that the algorithm converges after a certain number of iterations depending on the regularizer and parameters. When used in practice, there are no GT data available

during reconstruction of OCTA volumes. One way of determining whether convergence has been reached, is to compare the current reconstruction result with the initial OCTA data and stop reconstruction once an error metric such as mean squared error between the two stopped changing.

6. Conclusion and Outlook

In this paper, we present a novel MAP estimation based data consistency term for the reconstruction of OCTA volumes. It extends the MLE of Ploner at al. with regularization terms, thereby harnessing the inverse problem formulation's potential. We present OCTA reconstruction results for AD and IFV data and compare them to GT OCTA data merged from six co-registered OCT scans. Generally, TV regularization outperforms WS both quantitatively and qualitatively. Contrast between vessels and areas in-between seems significantly increased, with only very minor blurring. While median-filtering also leads to increases in contrast, those en face images are significantly more blurry. TV reconstruction achieves PSNR of more ~29.43 dB for AD and more than ~35.56 dB for IFV. Except in the case of three scan repeats using IFV, TV regularized reconstruction outperforms 3D median filtering. Qualitative results show a clear reduction in noise in the FAZ and intercapillary areas when compared to the state-of-the-art. Although this is a simple iterative algorithm, it shows significant improvement for OCTA image data and opens OCTA up to more sophisticated optimization techniques, such as alternating direction method of multipliers (ADMM). The method presented in this paper is, to the best of our knowledge, the first application of Bayesian statistics for the reconstruction of OCTA images.

Funding

Deutsche Forschungsgemeinschaft (DFG) (MA 4898/12-1); National Institutes of Health (NIH) (5-R01-EY011289-31); Champalimaud Vision Award; Beckman-Argyros Award in Vision Research; Macula Vision Research Foundation (MVRF); Massachusetts Lions Clubs

Disclosures

NKW: Optovue (C), Carl Zeiss Meditec (F), Heidelberg (F), Nidek (F). JGF: Optovue (I, P), Topcon (F)

See Supplement 1 for supporting content.

References

1. W. Choi, E. M. Moult, N. K. Waheed, M. Adhi, B. Lee, C. D. Lu, T. E. de Carlo, V. Jayaraman, P. J. Rosenfeld, J. S. Duker, and J. G. Fujimoto, "Ultrahigh-Speed, Swept-Source Optical Coherence Tomography Angiography in Nonexudative Age-Related Macular Degeneration with Geographic Atrophy," Ophthalmology **122**, 2532–2544 (2015).
2. E. A. Novais, M. Adhi, E. M. Moult, R. N. Louzada, E. D. Cole, L. Husvogt, B. Lee, S. Dang, C. V. Regatieri, A. J. Witkin, C. R. Baumal, J. Hornegger, V. Jayaraman, J. G. Fujimoto, J. S. Duker, and N. K. Waheed, "Choroidal Neovascularization Analyzed on Ultrahigh-Speed Swept-Source Optical Coherence Tomography Angiography Compared to Spectral-Domain Optical Coherence Tomography Angiography," Am. J. Ophthalmol. **164**, 80–88 (2016).
3. Y. Jia, O. Tan, J. Tokayer, B. Potsaid, Y. Wang, J. J. Liu, M. F. Kraus, H. Subhash, J. G. Fujimoto, J. Hornegger, and D. Huang, "Split-spectrum amplitude-decorrelation angiography with optical coherence tomography," Opt. Express **20**, 4710 (2012).
4. A. Mariampillai, B. A. Standish, E. H. Moriyama, M. Khurana, N. R. Munce, M. K. K. Leung, J. Jiang, A. E. Cable, B. C. Wilson, I. A. Vitkin, and V. X. D. Yang, "Speckle variance detection of microvasculature using swept-source optical coherence tomography," Opt. Lett. **33**, 1530–1532 (2008).
5. M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," Magn. Reson. Medicine (2007).
6. E. Y. Sidky and X. Pan, "Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization," Phys. Medicine Biol. **53**, 4777–4807 (2008).

7. H. Wu, A. Maier, R. Fahrig, and J. Hornegger, "Spatial-temporal total variation regularization (STTVR) for 4D-CT reconstruction," in *Medical Imaging 2012: Physics of Medical Imaging,* vol. 8313 N. J. Pelc, R. M. Nishikawa, and B. R. Whiting, eds. (International Society for Optics and Photonics, 2012).
8. M. Amrehn, A. Maier, F. Dennerlein, and J. Hornegger, "Portability of TV-Regularized Reconstruction Parameters to Varying Data Sets," in *Bildverarbeitung für die Medizin 2015,* H. Handels, T. M. Deserno, H.-P. Meinzer, and T. Tolxdorff, eds. (Springer Vieweg, Berlin, Heidelberg, Berlin, Heidelberg, 2015), pp. 131–136.
9. D. Stromer, M. Amrehn, Y. Huang, P. Kugler, S. Bauer, G. Lauritsch, and A. Maier, "Comparison of SART and ETV reconstruction for increased C-arm CT volume coverage by proper detector rotation in liver imaging," in *2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI),* (IEEE, 2016), pp. 589–592.
10. Y. Huang, G. Lauritsch, M. Amrehn, O. Taubmann, V. Haase, D. Stromer, X. Huang, and A. Maier, "Image Quality Analysis of Limited Angle Tomography Using the Shift-Variant Data Loss Model," in *Bildverarbeitung für die Medizin 2016,* T. Tolxdorff, T. M. Deserno, H. Handels, and H.-P. Meinzer, eds. (Springer Vieweg, Berlin, Heidelberg, Berlin, Heidelberg, 2016), pp. 277–282.
11. 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**, 8338 (2010).
12. S. B. Ploner, C. Riess, J. Schottenhamml, E. M. Moult, N. K. Waheed, J. G. Fujimoto, and A. Maier, "A Joint Probabilistic Model for Speckle Variance, Amplitude Decorrelation and Interframe Variance (IFV) Optical Coherence Tomography Angiography," in *Bildverarbeitung für die Medizin 2018,* A. Maier, T. M. Deserno, H. Handels, K. H. Maier-Hein, C. Palm, and T. Tolxdorff, eds. (Springer Berlin Heidelberg, Berlin, Heidelberg, 2018), 211279, pp. 98–102.
13. G. L. Zeng, *Medical Image Reconstruction* (Springer Berlin Heidelberg, Berlin, Heidelberg, 2010), 1st ed.
14. M. T. Manhart, M. Kowarschik, A. Fieselmann, Y. Deuerling-Zheng, K. Royalty, A. Maier, and J. Hornegger, "Dynamic Iterative Reconstruction for Interventional 4-D C-Arm CT Perfusion Imaging," IEEE Transactions on Med. Imaging **32**, 1336–1348 (2013).
15. C. Neukirchen, M. Giordano, and S. Wiesner, "An iterative method for tomographic x-ray perfusion estimation in a decomposition model-based approach," Med. Phys. **37**, 6125–6141 (2010).
16. S. G. Chang, B. Yu, and M. Vetterli, "Adaptive wavelet thresholding for image denoising and compression," IEEE Transactions on Image Process. **9**, 1532–1546 (2000).
17. A. Chambolle, "An Algorithm for Total Variation Minimization and Applications," J. Math. Imaging Vis. **20**, 89–97 (2004).
18. L. Landweber, "An Iteration Formula for Fredholm Integral Equations of the First Kind," Am. J. Math. **73**, 615–624 (1951).
19. S. Van Der Walt, S. C. Colbert, and G. Varoquaux, "The NumPy array: A structure for efficient numerical computation," Comput. Sci. Eng. **13**, 22–30 (2011).
20. S. Van Der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, and T. Yu, "Scikit-image: Image processing in python," PeerJ **2014**, e453 (2014).
21. M. a. Mayer, A. Borsdorf, M. Wagner, J. Hornegger, C. Y. Mardin, and R. P. Tornow, "Wavelet denoising of multiframe optical coherence tomography data," Biomed. Opt. Express **3**, 572 (2012).
22. M. F. Kraus, B. Potsaid, M. a. Mayer, R. Bock, B. Baumann, J. J. Liu, J. Hornegger, and J. G. Fujimoto, "Motion correction in optical coherence tomography volumes on a per A-scan basis using orthogonal scan patterns," Biomed. Opt. Express **3**, 1182 (2012).
23. M. F. Kraus, J. J. Liu, J. Schottenhamml, C.-L. Chen, A. Budai, L. Branchini, T. Ko, H. Ishikawa, G. Wollstein, J. S. Schuman, J. S. Duker, J. G. Fujimoto, and J. Hornegger, "Quantitative 3D-OCT motion correction with tilt and illumination correction, robust similarity measure and regularization," Biomed. Opt. Express **5**, 2591 (2014).
24. S. B. Ploner, M. F. Kraus, L. Husvogt, E. Moult, A. Y. Alibhai, J. Schottenhamml, T. Geimer, C. Rebhun, B. Lee, C. R. Baumal, N. K. Waheed, J. S. Duker, J. G. Fujimoto, and A. Maier, "3-D OCT Motion Correction Efficiently Enhanced with OCT Angiography," in *Investigative Ophthalmology & Visual Science,* vol. 59 (The Association for Research in Vision and Ophthalmology, 2018), p. 3922.
25. S. B. Ploner, M. F. Kraus, E. M. Moult, L. Husvogt, J. Schottenhamml, A. Y. Alibhai, N. K. Waheed, J. S. Duker, J. G. Fujimoto, and A. K. Maier, "Efficient and high accuracy 3-D OCT angiography motion correction in pathology," (2020). arXiv:2010.06931.
26. J. Schottenhamml, E. M. Moult, E. A. Novais, M. F. Kraus, B. K. Lee, W. Choi, S. B. Ploner, L. Husvogt, C. D. Lu, P. Yiu, P. J. Rosenfeld, J. S. Duker, A. K. Maier, N. Waheed, and J. G. Fujimoto, "OCT-OCTA Segmentation," in *Bildverarbeitung für die Medizin 2018,* A. Maier, T. M. Deserno, H. Handels, K. Maier-Hein, P. C., and T. Tolxdorff, eds. (Springer Vieweg, Berlin, Heidelberg, Berlin, Heidelberg, 2018), p. 284.
27. Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," IEEE Transactions on Image Process. **13**, 600–612 (2004).
28. W. Choi, B. Potsaid, V. Jayaraman, B. Baumann, I. Grulkowski, J. J. Liu, C. D. Lu, A. E. Cable, D. Huang, J. S. Duker, and J. G. Fujimoto, "Phase-sensitive swept-source optical coherence tomography imaging of the human retina with a vertical cavity surface-emitting laser light source," Opt. Lett. **38**, 338–340 (2013).