

Research article

Optoacoustic model-based inversion using anisotropic adaptive total-variation regularization

Shai Biton¹, Nadav Arbel¹, Gilad Drozdov, Guy Gilboa, Amir Rosenthal*

Andrew and Erna Viterbi Faculty of Electrical Engineering, Technion – Israel Institute of Technology, Technion City 32000, Haifa, Israel

ARTICLE INFO

Keywords:

Optoacoustic imaging
Total variation
Inversion algorithms
Model-based reconstruction

ABSTRACT

In optoacoustic tomography, image reconstruction is often performed with incomplete or noisy data, leading to reconstruction errors. Significant improvement in reconstruction accuracy may be achieved in such cases by using nonlinear regularization schemes, such as total-variation minimization and L_1 -based sparsity-preserving schemes. In this paper, we introduce a new framework for optoacoustic image reconstruction based on adaptive anisotropic total-variation regularization, which is more capable of preserving complex boundaries than conventional total-variation regularization. The new scheme is demonstrated in numerical simulations on blood-vessel images as well as on experimental data and is shown to be more capable than the total-variation- L_1 scheme in enhancing image contrast.

1. Introduction

Optoacoustic tomography (OAT) is a hybrid imaging modality capable of visualizing optically absorbing structures with ultrasound resolution at tissue depths in which light is fully diffused [1–5]. The excitation in OAT is most often performed via high-energy short pulses whose absorption in the tissue leads to the generation of acoustic sources via the process of thermal expansion. The acoustic signals from the sources are measured over a surface that partially or fully surrounds the imaged object and used to form an image of the acoustic sources, which generally represents the local energy absorption in the tissue [6]. Since hemoglobin is one of the strongest absorbing tissue constituents, optoacoustic images often depict blood vessels and blood-rich organs, such as the kidneys [7] and heart [8].

Numerous algorithms exist for reconstructing optoacoustic images from measured tomographic data [6]. In several imaging geometries, analytical formulae exist that may be applied directly on the measured data in either time [9,10] or frequency domain [11] to yield an exact reconstruction. The popularity of analytical formulae may be attributed to the simplicity of their implementation and low computational burden [12]. However, analytical algorithms are not exact for arbitrary detection surfaces or detector geometries and lack the possibility of regularizing the inversion in the case of noisy or incomplete data. In such cases, it is often preferable to use model-based algorithms, in which the relation between the image and measured data is represented

by a matrix whose inversion is required to reconstruct the image.

In the last decade, numerous regularization approaches have been demonstrated for model-based image reconstruction. The most basic approach is based on energy minimization and includes techniques such as Tikhonov regularization [13] and truncated singular-value decomposition [14]. In these techniques, a cost function on the image or components thereof is used to avoid divergence of the solution in the case of missing data, generally without making any assumptions on the nature of the solution. More advanced approaches to regularization exploit the specific properties of the reconstructed image. Since natural images may be sparsely represented when transformed to an alternative basis, e.g. the wavelet basis, using nonlinear cost functions that promote sparsity in such bases may be used for denoising and image reconstruction from missing data [15–21]. In images of blood vessels, in which the boundaries of the imaged structures may be of higher importance than the texture in the image, total-variation (TV) minimization has been shown to enhance image contrast and reduce artifacts [22–24].

In TV regularization, the cost function regularizer is the L_1 norm of the image gradient, which is generally lower for images with sharp, yet very localized, variations than for images in which small variations occur across the entire image. Therefore TV regularization enhances boundaries and reduces texture, where over-regularization may lead to almost piecewise constant images, which are often referred to as cartoon-like. While TV regularization is capable of accentuating the

* Corresponding author.

E-mail address: amir.r@technion.ac.il (A. Rosenthal).

¹ These authors contributed equally to this work.

boundaries of imaged objects, it does not treat all boundaries the same. In particular, boundaries with short lengths will lead to a lower TV cost function than boundaries with long lengths. Thus, complex, non-convex boundaries may be rounded by TV regularization to the closest convex form. Recently, Wang et al. have shown that if the directionality of the TV functional is adapted to the image features, TV regularization may be applied for optoacoustic reconstruction of objects with non-convex boundaries without distorting the boundaries [25].

In this paper, we demonstrate a new regularization framework for model-based optoacoustic image reconstruction that overcomes the limitations of TV regularization and is compatible with objects with complex non-convex boundaries. In our scheme, an adaptive anisotropic total variation (A²TV) cost function is used, in which the cost function is determined by the specific geometry of the imaged objects [26]. In particular, the A²TV cost function wishes to minimize the total variation in directions that are orthogonal to the boundary of the object. In contrast to [25], where the boundaries were calculated using geometrical considerations limited to 2D images, the A²TV framework developed in this work is based on eigenvalue decomposition of the image structure tensor, which may be applied in higher dimensions. The proposed formalism in the current study is based on a recent work by part of the authors which is concerned with nonlinear spectral analysis of the A²TV functional [26]. The work of Ref. [26] examines shapes which are perfectly preserved under A²TV regularization. Earlier works concerning TV have shown that only convex rounded shapes of low curvature are preserved [27]. For A²TV, however, it is shown in [26] that a parameter controlling the local extent of directionality is directly related to the degree of convexity (in the sense of [28]) and to the curvature magnitude of structures which are preserved. Thus, with appropriate parameters, long vessels of complex-nonconvex structure can be better regularized, keeping the original structure intact.

The performance of A²TV regularization was tested numerically in this work in numerical simulations for complex images of blood vessels and experimentally on 2D phantoms. The simulations were performed for the cases of noisy data and missing data and compared to unregularized reconstructions as well as to TV- L_1 reconstructions [15]. In both the numerical and experimental reconstructions, A²TV significantly increased the image contrast and was more capable than TV- L_1 in preserving non-convex structures when strong regularization was performed. In the experimental reconstructions, A²TV achieved a higher level of contrast enhancement of weak structures than the one achieved by TV- L_1 .

The rest of the paper is organized as follows: in Section 2 we give the theoretical background for OAT image reconstruction. Section 3 introduces the framework of A²TV and the A²TV algorithm for OAT image reconstruction developed in this work. The simulation results are given in Section 4, while the experimental ones are given in Section 5. We conclude the paper with a Discussion in Section 6.

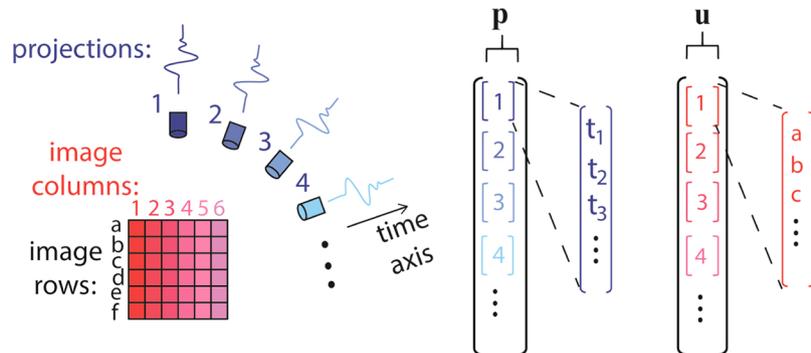


Fig. 1. An illustration of the structure of vectors \mathbf{p} and \mathbf{u} used in the matrix construction in Eq. (4).

2. Optoacoustic image reconstruction

2.1. The forward problem

The acoustic waves in OAT are commonly described by a pressure field $p(\mathbf{r}, t)$ that fulfills the following wave equation [29]:

$$\frac{\partial^2 p(\mathbf{r}, t)}{\partial t^2} - c^2 \nabla^2 p(\mathbf{r}, t) = \Gamma H_r(\mathbf{r}) \frac{\partial H_t(t)}{\partial t} \quad (1)$$

where c is the speed of sound in the medium, t is time, $\mathbf{r} = (x, y, z)$ denotes position in 3D space, $p(\mathbf{r}, t)$ is the generated pressure, Γ is the Grüneisen parameter, and $H_r(\mathbf{r})H_t(t)$ is the energy per unit volume and unit time. The spatial distribution function of energy deposited in the imaged object, $H_r(\mathbf{r})$, is referred to in the rest of the paper as the optoacoustic image.

The analysis of (1) for an optoacoustic point source at \mathbf{r}' , i.e. $H_r(\mathbf{r}) = \delta(\mathbf{r} - \mathbf{r}')$, may be performed in either time or frequency domain. In the time domain, a short-pulse excitation $H_t(t) = \delta(t)$ is used. In this case, the solution to Eq. (1) is given by [30]

$$p_\delta(\mathbf{r}, t) = \frac{\Gamma}{4\pi c} \frac{\partial}{\partial t} \frac{\delta(|\mathbf{r} - \mathbf{r}'| - ct)}{|\mathbf{r} - \mathbf{r}'|}. \quad (2)$$

For a general image $H_r(\mathbf{r})$, the solution for $p(\mathbf{r}, t)$ may be obtained by convolving $H_r(\mathbf{r})$ with the expression in Eq. (2), which yields:

$$p(\mathbf{r}, t) = \frac{\Gamma}{4\pi c} \frac{\partial}{\partial t} \int_{|\mathbf{r} - \mathbf{r}'| = ct} \frac{H_r(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|}. \quad (3)$$

Since the relation between the measured pressure signals, or projections, and the image is linear, it may be represented by a matrix relation in its discrete form:

$$\mathbf{p} = \mathbf{M}\mathbf{u} \quad (4)$$

where \mathbf{p} and \mathbf{u} are vector representations of the acoustic signals and originating image respectively, and \mathbf{M} is the model matrix that represents the operations in Eq. (3). In our work, the images is given on a two dimensional grid, and the measured pressure signals are also two dimensional, where one dimension represents the projection number, and the other time. An illustration of the image grid and projection and their respective mapping to the vectors \mathbf{u} and \mathbf{p} , is shown in Fig. 1. Briefly, the vector \mathbf{u} is divided into sub-vectors, each representing the image values for a given column, whereas the vector \mathbf{p} is divided to sub-vectors, each of which represents the time-domain pressure signal for a given location of the acoustic detector.

The i th column of the matrix \mathbf{M} represents the set of acoustic signals generated by a pixel corresponding to the location of the i th entry in the vector \mathbf{u} . Accordingly, to calculate the matrix \mathbf{M} one needs to define time domain signal for a given detector location expected for a discrete pixel. Since the operations in Eq. (3) relate to continuous, rather than

discrete images, one first needs to define the continuous representation of a single pixel, and then calculate its respective time-domain signals. For example, in [31] it was assumed that the image value was constant over each of the square pixels, leading to an image $H_r(\mathbf{r})$ that is piecewise constant. While simple to implement, a simple piece-wise uniform model for $H_r(\mathbf{r})$ includes discontinuities that lead to significant numerical errors owing to the derivative operation in Eq. (3). In the current work, we use the model of [32], in which the image $H_r(\mathbf{r})$ is represented by a linear interpolation between its grid points.

2.2. The inverse problem

While several approaches to OAT image reconstruction exist, we focus herein on image reconstruction within the discrete model-based framework described in the previous sub-section, which involves inverting the matrix relation in Eq. (4) to recover \mathbf{u} from \mathbf{p} . The most basic method to invert Eq. (4) is based on solving the following optimization problem:

$$\mathbf{u}^* = \arg \min \|\mathbf{p} - \mathbf{M}\mathbf{u}\|_2^2, \quad (5)$$

where \mathbf{u}^* is the solution and $\|\cdot\|_2$ is the L_2 norm. A unique solution to Eq. (5) exists, which is given by the Moore–Penrose inverse:

$$\mathbf{u}^* = (\mathbf{M}^T\mathbf{M})^{-1}\mathbf{M}^T\mathbf{p}. \quad (6)$$

Alternatively, Eq. (5) may be solved via iterative optimization algorithms. In particular, since the matrix \mathbf{M} is sparse, efficient inversion may be achieved by the LSQR algorithm [33].

In many cases, the measured projection data \mathbf{p} is insufficient to achieve a high-quality reconstruction of \mathbf{u} that accurately depicts its morphology. For example, when the density or coverage of the projections is too small, the matrix \mathbf{M} may become ill-conditioned, leading to significant, possibly divergent, image artifact. In other cases, \mathbf{M} may be well conditioned, but the measurement data may be too noisy to accurately recover \mathbf{u} . In both these cases, regularization may be used to improve image quality by incorporating previous knowledge on the properties image in the inversion process.

One of the simplest forms of regularization is Tikhonov regularization, in which an additional cost function is added [13]:

$$\mathbf{u}^* = \arg \min \|\mathbf{p} - \mathbf{M}\mathbf{u}\|_2^2 + \lambda \|\mathbf{L}\mathbf{u}\|_2^2 \quad (7)$$

where \mathbf{L} is weighting matrix. In the simplest form of Tikhonov regularization, \mathbf{L} is equal to the identity matrix, i.e. $\mathbf{L} = \mathbf{I}$, thus putting a penalty on the energy of the image. The value of the regularization parameter $\lambda > 0$ controls the tradeoff between fidelity and smoothness, where over-regularization may lead to the smearing of edges and texture in the image.

An alternative to the energy-minimizing cost function of Tikhonov are sparsity-maximizing cost functions. Since natural images may be sparsely represented in an alternative basis, e.g. the wavelet basis, a cost function that promotes sparsity may reduce reconstruction errors. Denoting the transformation matrix by Φ , one wishes that $\Phi\mathbf{u}$ be sparse, i.e. that most of its entries be approximately zero. In practice, sparsity is often enforced by using the L_1 norm because of its compatibility of optimization algorithms [16–21]. Accordingly, the inversion is performed by solving the following optimization problem:

$$\mathbf{u}^* = \arg \min \|\mathbf{p} - \mathbf{M}\mathbf{u}\|_2^2 + \mu \|\Phi\mathbf{u}\|_1, \quad (8)$$

where $\mu > 0$ is the regularization parameter, controlling the tradeoff between sparsity and signal fidelity. When over regularization is performed, compression artifacts may appear in the reconstructed image.

Sparsity may be enforced not only on alternative representations of the image, but also on image variations. The discrete TV cost function approximates the l_1 norm on the image gradient, and is given by

$$\|\mathbf{u}\|_{\text{TV}} = \sum_n \sqrt{|u_{x,y}^n - u_{x-1,y}^n|^2 + |u_{x,y}^n - u_{x,y-1}^n|^2}, \quad (9)$$

where $u_{x,y}^n$ is the n th entry of the vector \mathbf{u} , and the subtraction of “1” to x or y in the subscript corresponds to an entry of \mathbf{u} that is respectively shifted in relation to $u_{x,y}^n$ by one pixel in the x or y direction of the 2D image. The inversion using TV is thus given by [22–24]

$$\mathbf{u}^* = \arg \min \|\mathbf{p} - \mathbf{M}\mathbf{u}\|_2^2 + \alpha \|\mathbf{u}\|_{\text{TV}} \quad (10)$$

where $\alpha > 0$ is the regularization parameter. In the case of TV minimization, over-regularization may lead to cartoon-like images and rounding of complex boundaries into convex shapes.

In some cases, it is beneficial to promote both sparsity of the image in an alternative basis and TV minimization. In such cases, the optimization problem is as follows [15]:

$$\mathbf{u}^* = \arg \min \|\mathbf{p} - \mathbf{M}\mathbf{u}\|_2^2 + \mu \|\Phi\mathbf{u}\|_1 + \alpha \|\mathbf{u}\|_{\text{TV}}. \quad (11)$$

We will refer to the regularization described in Eq. (11), as TV- L_1 regularization.

3. Adaptive anisotropic total variation

3.1. The functional

We would like to use a regularizer that is adapted to the image in such a way that it regularizes more along edges (level-lines of the image) and less across edges (in the direction of the gradient). This idea has been introduced for nonlinear scale-space flows by Weickert [34] in the anisotropic diffusion formulation. However there is no known functional associated with anisotropic diffusion, and it is therefore not trivial to include an anisotropic-diffusion operation in our inverse-problem formulation. A more recent study of Grasmair et al. [35] uses a similar adaptive scheme within a TV-type formulation. In the study of [26], a comprehensive theoretical and numerical analysis was performed for adaptive-anisotropic TV (A^2 TV). It was shown that stable structures can be non-convex and in addition can have high curvature on the boundaries. Illustrations of the stable sets characterize the TV and A^2 TV regularizers are shown in Fig. 2. The degree of anisotropy directly controls the degree of allowed nonconvexity and the upper bound on the curvature. We adopt the formulation of [26], in which the mathematical underpinnings of A^2 TV are described in detail.

Let the A^2 TV functional be defined by

$$\mathcal{J}_{A^2\text{TV}}(\mathbf{u}) = \int_{\Omega} \|A(\mathbf{x})\nabla\mathbf{u}(\mathbf{x})\|_2 \, d\mathbf{x} = \int_{\Omega} \|\nabla_A\mathbf{u}(\mathbf{x})\|_2 \, d\mathbf{x} \quad (12)$$

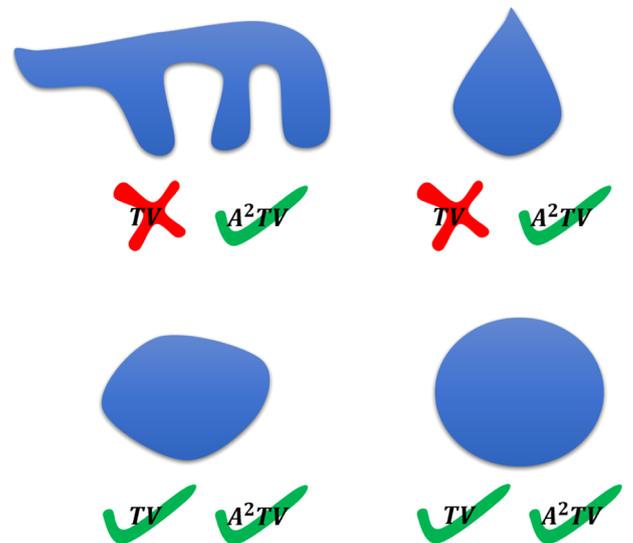


Fig. 2. An illustration of sets which are stable for TV and A^2 TV – notice A^2 TV admits non-convex and highly curved functions, including ones which resemble arteries.

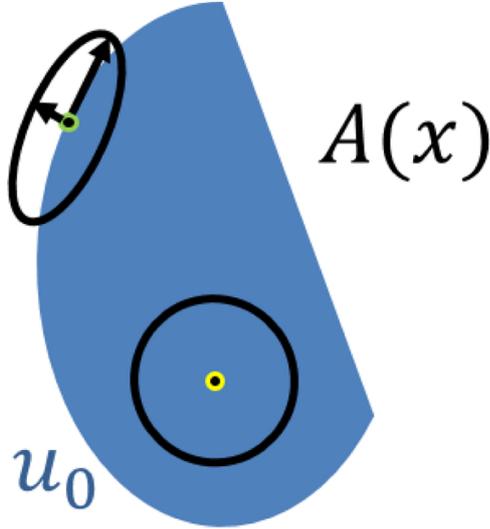


Fig. 3. An illustration of the tensor $A(\mathbf{x})$. At any point the tensor rotates and rescales the coordinate system in an image-driven manner. It assumes some approximation \mathbf{u}_0 of the data exists. The tensor is designed such that lower regularization is applied across edges (top left ellipse) whereas in flat regions regularization is applied in an isotropic manner (bottom right circle).

where $\mathbf{x} = (x, y) \in \Omega$ denotes position in 2D space, $A(\mathbf{x}) \in \mathbb{R}^{2 \times 2}$ is a tensor or a spatially adaptive matrix and $\nabla_A = A(\mathbf{x})\nabla$ is an ‘‘adaptive gradient’’. This functional is convex and can be optimized by standard convex solvers given the tensor $A(\mathbf{x})$. We now turn to the issue of how this tensor is constructed to allow a good regularization of vessel-like structures.

3.1.1. Constructing the tensor $A(\mathbf{x})$

We assume to have a rough approximation of the image to be reconstructed. This can be done, for instance, by having an initial least-square non-regularized approximation or a standard TV reconstruction, as in Eq. (10). We thus have an initial estimation \mathbf{u}_0 . The tensor $A(\mathbf{x})$ determines the principle and the secondary directions of the regularization at each point and their magnitude. The construction of $A(\mathbf{x})$ is performed using \mathbf{u}_0 according to the following principles: In regions in which \mathbf{u}_0 is relatively flat, i.e. $\nabla \mathbf{u}_0$ almost vanishes, $A(\mathbf{x})$ should resemble the identity matrix and have no preferred direction, thus leading to the conventional TV regularizer. In regions with dominant edges, $A(\mathbf{x})$ should capture the principle axes of the edge. See Fig. 3 for an illustration of $A(\mathbf{x})$.

Mathematically, the tensor $A(\mathbf{x})$ is defined by adapting the eigenvalues of a smoothed structure tensor of a smoothed image $\mathbf{u}_{0,\sigma}$ (with a Gaussian kernel of standard deviation σ), defined by,

$$\mathcal{J}_\rho(\nabla \mathbf{u}_{0,\sigma}) = \kappa_\rho * (\nabla \mathbf{u}_{0,\sigma} \otimes \nabla \mathbf{u}_{0,\sigma}), \quad (13)$$

where κ_ρ is a Gaussian kernel with a standard deviation of ρ , $*$ denotes an element-wise convolution and \otimes denotes an outer product. The Gaussian kernel's standard deviations σ and ρ are chosen in coordination with the noise level of the image and the smallest object resolution within the image. High standard deviation values might diminish the effect of small objects on the structure tensor while low values might lead to the unwanted scenario in which reconstruction errors have a significant effect on the structure tensor. The explicit expressions for \mathcal{J}_ρ in the 2D and 3D cases are respectively given by

$$\mathcal{J}_\rho(\mathbf{x}) = \begin{pmatrix} \kappa_\rho * u_{x,\sigma}^2(\mathbf{x}) & \kappa_\rho * u_{x,\sigma}(\mathbf{x})u_{y,\sigma}(\mathbf{x}) \\ \kappa_\rho * u_{x,\sigma}(\mathbf{x})u_{y,\sigma}(\mathbf{x}) & \kappa_\rho * u_{y,\sigma}^2(\mathbf{x}) \end{pmatrix} \quad (14)$$

and

$$\mathcal{J}_\rho(\mathbf{x}) = \kappa_\rho * \begin{pmatrix} u_{x,\sigma}^2(\mathbf{x}) & u_{x,\sigma}(\mathbf{x})u_{y,\sigma}(\mathbf{x}) & u_{x,\sigma}(\mathbf{x})u_{z,\sigma}(\mathbf{x}) \\ u_{x,\sigma}(\mathbf{x})u_{y,\sigma}(\mathbf{x}) & u_{y,\sigma}^2(\mathbf{x}) & u_{y,\sigma}(\mathbf{x})u_{z,\sigma}(\mathbf{x}) \\ u_{x,\sigma}(\mathbf{x})u_{z,\sigma}(\mathbf{x}) & u_{y,\sigma}(\mathbf{x})u_{z,\sigma}(\mathbf{x}) & u_{z,\sigma}^2(\mathbf{x}) \end{pmatrix}. \quad (15)$$

The structure tensor matrix has eigenvectors corresponding to the direction of the gradient and tangent at each point \mathbf{x} ; and eigenvalues corresponding to the magnitude of each direction. In order to preserve structure, we should change the relation between those eigenvalues so that for flat-like areas in the image we will smooth the image in an isotropic way, while for edge-like areas, we will perform more smoothing in the tangent direction rather the gradient one. For 2D, we begin by looking in the eigen-decomposition of the structure-tensor,

$$\mathcal{J}_\rho = \mathbf{V}\mathbf{D}\mathbf{V}^{-1} \quad (16)$$

Where \mathbf{V} is a matrix whose columns, $v_1, v_2 \in \mathbb{R}^2$, are the eigenvectors of \mathcal{J}_ρ , and \mathbf{D} is a diagonal matrix with eigenvalues in the diagonal,

$$\mathbf{D} = \begin{pmatrix} \mu_1 & 0 \\ 0 & \mu_2 \end{pmatrix}; \quad \mathbf{V} = (v_1|v_2) \quad (17)$$

Assuming $\mu_1 \geq \mu_2$. The spatially adaptive matrix $A(\mathbf{x})$, used in Eq. (12), is constructed from a modification of the eigenvalue matrix, denoted by $\tilde{\mathbf{D}}$, as follows [34]:

$$A = \mathbf{V}\tilde{\mathbf{D}}\mathbf{V}^{-1}, \quad (18)$$

where $\tilde{\mathbf{D}}$ is given by

$$\tilde{\mathbf{D}} = \begin{pmatrix} c(\mu_1/\mu_{1,\text{avg}}; k) & 0 \\ 0 & 1 \end{pmatrix}. \quad (19)$$

In Eq. (19), $\mu_{1,\text{avg}}$ is the average value of μ_1 across the image and $c(\cdot; \cdot)$ is a function of two parameters defined as follows:

$$c(s; k) = \begin{cases} 1, & s \leq 0 \\ 1 - \exp\left(-\frac{c_m}{\left(\frac{s}{k}\right)^m}\right), & s > 0 \end{cases} \quad (20)$$

where the chosen values for the parameter are the ones recommended in Ref. [34]: $c_m = 3.31488$ and $m = 4$. The parameter $k \leq 1$ determines which regions in the image will be regularized anisotropically and is chosen based on the desired level of anisotropy in the reconstructed image, as shown in Sections 4 and 5. In pixels in which $s \ll k$, i.e. $\mu_1/\mu_{1,\text{avg}} \ll k$, we will obtain $c \approx 1$, and $\tilde{\mathbf{D}}$ will be reduced to the unitary matrix. Accordingly, for regions in which the image gradient is sufficiently small from the average image gradient, as regulated by k , the $A^2\text{TV}$ functional (Eq. (12)) is reduced to the standard, isotropic TV functional (Eq. (9)). In the rest of the image, where c is sufficiently smaller than 1, regularization is performed more strongly in the direction of the eigenvector v_2 , i.e. the direction in which the image gradient is smaller, thus enhancing the anisotropy in those regions.

It is worth noting that in the 3D case, assuming $\mu_1 \geq \mu_2 \geq \mu_3$, the only modification to the analysis above is that $\tilde{\mathbf{D}}$ accepts the following form:

$$\tilde{\mathbf{D}} = \begin{pmatrix} c(\mu_1/\mu_{1,\text{avg}}; k) & 0 & 0 \\ 0 & c(\mu_2/\mu_{1,\text{avg}}; k) & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (21)$$

The structure of $\tilde{\mathbf{D}}$ in Eq. (21) will be highly anisotropic for tube-like structures and will enforce variation-reducing regularization along the structure length (eigenvector v_3), while maintaining low regularization in the cross-section plane (eigenvectors v_1 and v_2).

3.2. Reconstruction based on $A^2\text{TV}$

The reconstruction based on the $A^2\text{TV}$ minimizes the following functional:

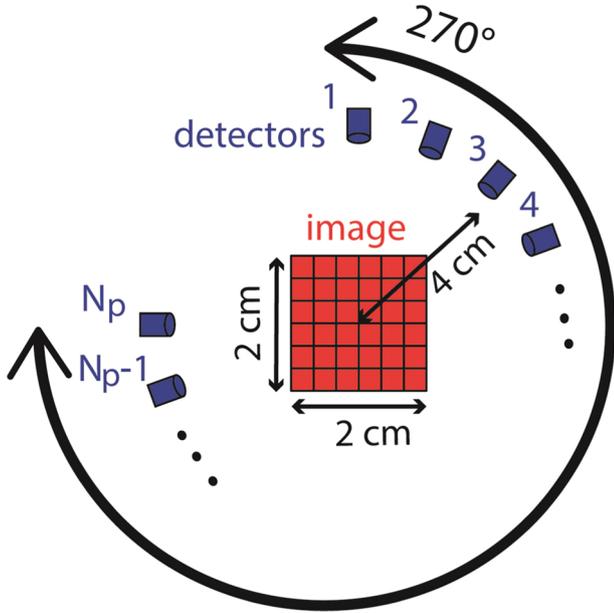


Fig. 4. An illustration of the image and detector geometry used in the numerical simulations and experimental measurements in Sections 4 and 5.

$$\mathbf{u}^* = \underset{\mathbf{u}}{\operatorname{argmin}} \mathcal{J}_{A^2TV}(\mathbf{u}) + \frac{\lambda}{2} \|\mathbf{M}\mathbf{u} - \mathbf{p}\|_2^2, \quad (22)$$

where \mathbf{M} is the model matrix and \mathbf{p} is the acoustic pressure wave. The solution of which is done by the modified Chambolle-Pock projection algorithm [36] described in Appendix A.

In the process of minimization, the tensor A is initialized as the identity matrix for all $\mathbf{x} \in \mathbb{R}^2$, which reduces the A^2TV energy to the TV one as it performs the diffusion isotropically. The tensor A is then updated according to the initial solution u_0 . This is repeated until numerical convergence is reached.

We note that while the energy of the A^2TV is convex for a fixed tensor $A(\mathbf{x})$, it is not convex when $A(\mathbf{x})$ is adaptive and depends on the imaged object. Thereby, we do not have a mathematical proof of convergence. Nonetheless, it has been shown both in our work [26] and in Refs. [34,35] that heuristically, both the image u and the tensor $A(\mathbf{x})$ converge.

4. Numerical simulations

In this section, we demonstrate the performance of A^2TV -based

inversion for the circular detection geometry illustrated in Fig. 4. The simulations were performed on a 2D vascular image of a mouse retina, obtained via confocal microscopy. The vasculature image, shown in Fig. 5a, was represented over a square grid with a size of 256×256 pixels with pixel size of 0.1×0.1 mm. The projections were simulated over a 270-degree semi-circle with a radius of 4 cm that surrounded the object, in accordance with conventional OAT systems [4]. A magnification of four square regions of the image is shown in Fig. 5b.

The reconstructions were performed using the conventional L_2 -based regularization-free approach (Eq. (5)) performed via LSQR, TV- L_1 regularization (Eq. (11)), and the proposed A^2TV regularization (Eq. (22)) method. Since the scaling of the model matrix \mathbf{M} (Eq. (4)) depends on the exact implementation of its construction [32], we normalized \mathbf{M} by $\frac{1}{160} \sqrt{\|\mathbf{M}\|_\infty \|\mathbf{M}\|_1}$ to assure that the regularization parameters are independent of the scaling of \mathbf{M} . To assess the quality of the reconstructions, we used the mean absolute distance (MAD) given by the following equation:

$$\text{MAD} = \frac{1}{N} \|\mathbf{u}_{\text{orig}} - \mathbf{u}^*\|_1, \quad (23)$$

where N is the number of pixels in the image.

Two cases were tested: In the first case, zero-mean Gaussian noise with a standard deviation of 0.6 times the maximum value of \mathbf{p} was added to the projection. The number of projections was chosen to be 256, corresponding to the geometry found in state-of-the-art optoacoustic systems [37] and sufficient for the accurate reconstruction of the tested image in the noiseless case. In the second case, the number of projections was reduced to 32, which is half the number of projections used in low-end optoacoustic systems characterized by reduced lateral resolution [37]. Accordingly, 32 projections are insufficient for producing detailed optoacoustic images using conventional reconstruction techniques. In all the examples, the number of iterations was chosen to be sufficiently high to achieve convergence.

Figs. 6 and 7 respectively show the reconstructions obtained using TV- L_1 and A^2TV , performed with 3000 iterations, for the cases of additive Gaussian noise. In both figures, 9 reconstructions are shown, corresponding to a scan in the regularization parameters. In the case of TV- L_1 , μ and α represent the strength of the L_1 and TV regularization terms Eq. (11), whereas in the case of A^2TV , λ represents the strength of the fidelity term in Eq. (22) with respect to the A^2TV term and k determines the strength of the anisotropy, where lower values of k correspond to higher anisotropy. In the A^2TV reconstructions, the standard deviations of the smoothing kernels (Eq. (14)) were $\sigma = 1.5$ pixels and $\rho = 3$ pixels. For all the reconstructions, the MAD values appear on the top-right corner of the image. In Fig. 8, we show a comparison between the regularization-free LSQR reconstruction (Fig. 8a) and the TV- L_1

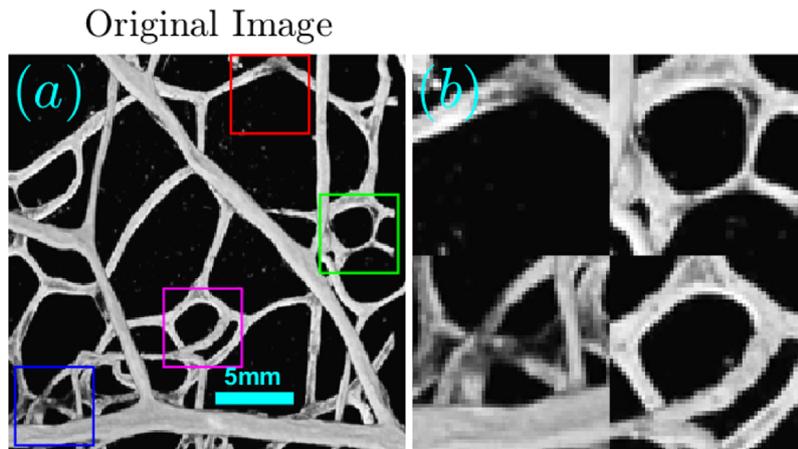


Fig. 5. (a) The originating image on which all the reconstructions were performed and (b) a binary mask that was generated from it. The bottom panel (c and d) shows magnifications of 4 marked regions in the respective images in the top panel.

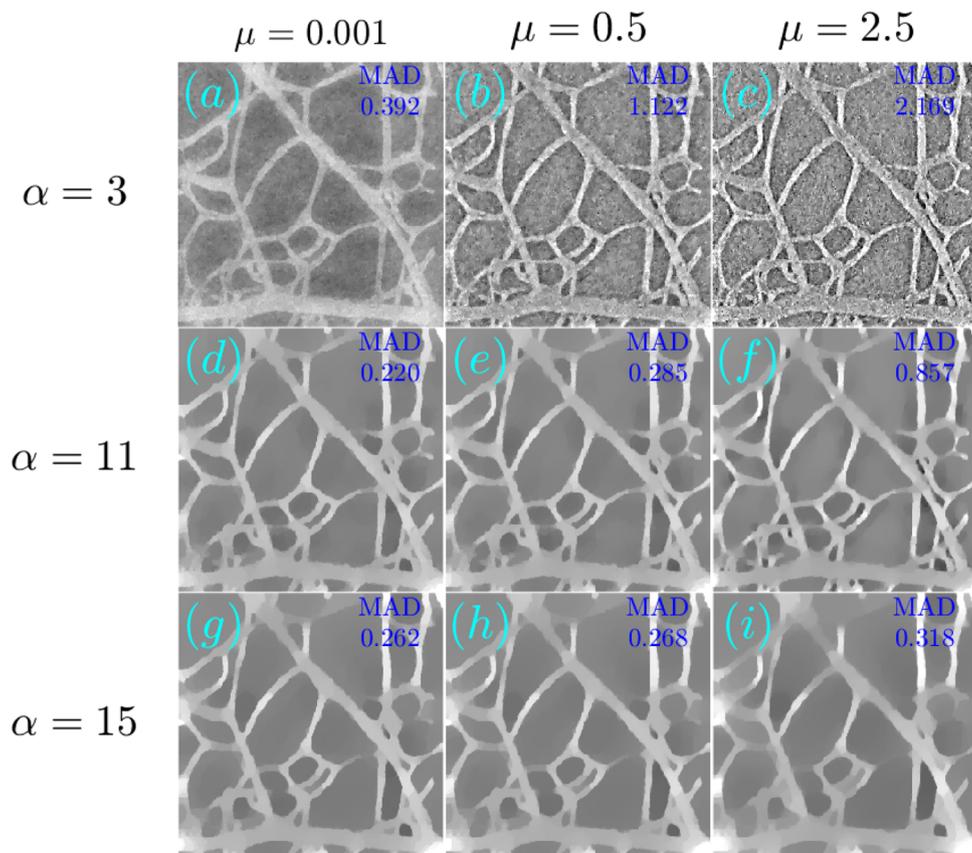


Fig. 6. (a-i) The reconstruction of the image shown in Fig. 5a for the case of additive Gaussian noise using different parameters for the TV- L_1 case. The reconstructions were performed with 3000 iterations.

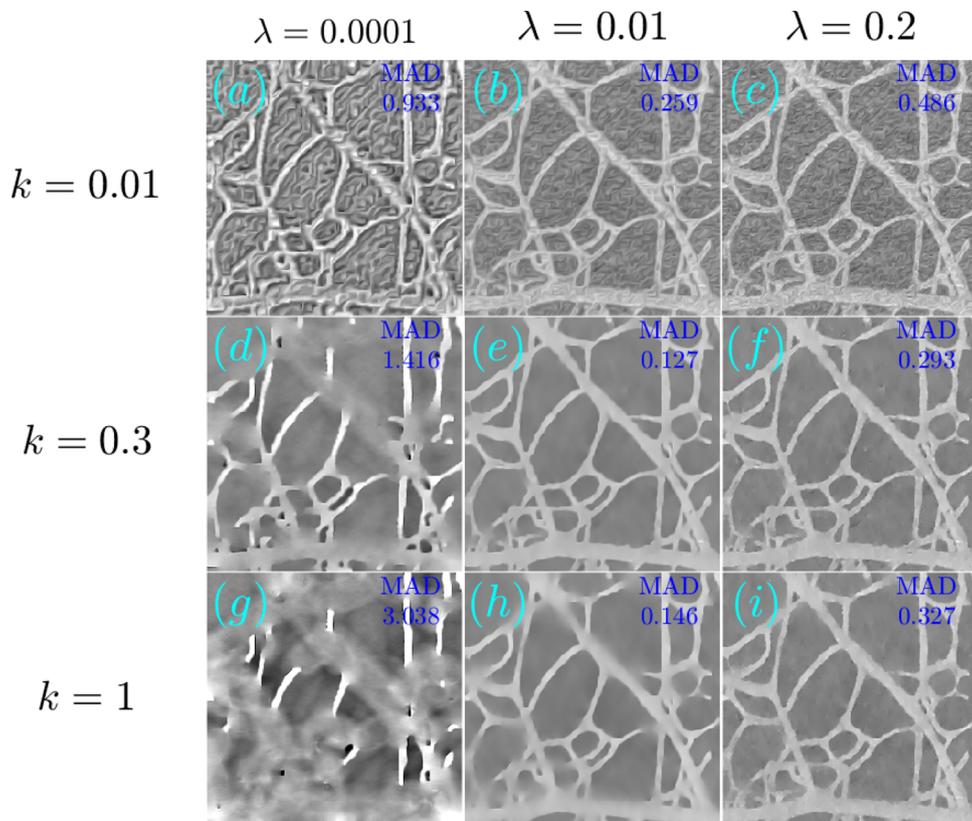


Fig. 7. (a-i) The reconstruction of the image shown in Fig. 5a for the case of additive Gaussian noise using different parameters for the A^2TV case. The reconstructions were performed with 3000 iterations.

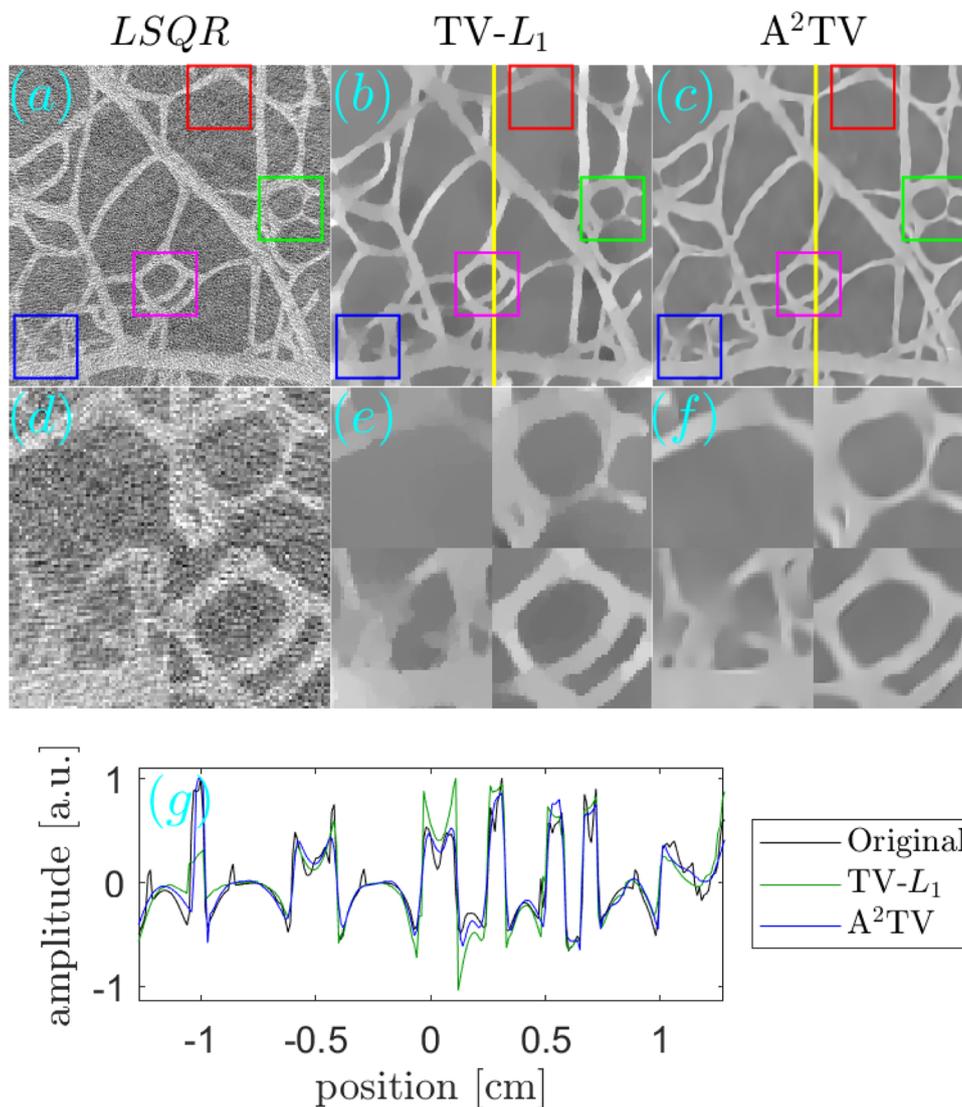


Fig. 8. (a–c) The reconstruction of the image shown in Fig. 5a for the case of additive Gaussian noise using (a) LSQR, (b) TV- L_1 , and (c) A²TV. The bottom panel (d–f) shows magnifications of 4 marked regions in the respective images in the top panel. Both the TV- L_1 and A²TV reconstructions were produced using 3000 iterations. (g) A 1D slice of the TV- L_1 and A²TV reconstructions, corresponding to the vertical lines on the top panel, in comparison to a 1D slice taken from the original image.

(Fig. 8b) and A²TV (Fig. 8c) of Figs. 6d and 7e, respectively, which correspond to the regularization parameters that achieved the lowest MAD values. The middle panel of the figure (Fig. 8d–f) shows a magnification of 4 patches taken from the reconstructions, whereas the bottom panel (Fig. 8g) presents a 1D slice taken over the yellow line in Fig. 8b and c. While both TV- L_1 and A²TV significantly improved the reconstruction quality, in A²TV more of the noise-induced texture between the blood vessels could be removed without damaging the structure of the blood vessels, thus leading to a lower MAD.

The reconstructions for the case of 32 projections are presented in a similar manner to the case of noisy data. Fig. 9, obtained with 1000 iterations, and Fig. 10, obtained with 1500 iterations, respectively show the TV- L_1 and A²TV reconstructions for a scan regularization parameters, whereas Fig. 11 shows a comparison between the LSQR and TV- L_1 and A²TV reconstructions that achieved the lowest MAD (Figs. 9d and 10e, respectively). All the A²TV reconstructions were obtained with $\sigma = 1.5$ pixels and $\rho = 1$ pixel. Fig. 11 shows that both TV- L_1 and A²TV eliminated the streak artifacts that appeared with in the LSQR reconstruction, where the lowest MAD was achieved by the TV- L_1 reconstruction. Since both regularization methods eliminated the streak artifacts, the lower MAD achieved by TV- L_1 is a result of its ability to

better preserve texture within the blood vessels, whereas in the A²TV much of the blood-vessel texture was lost. Indeed, when examining the 1D slice in Fig. 11g, it is easy to see that the TV- L_1 reconstruction captures the variations within the blood vessels better, whereas in the A²TV reconstruction, these variations are smoothed.

In both the cases studied, A²TV exhibited a higher capability than TV- L_1 to perform regularization without harming non-convex structures. Even in the case in which TV- L_1 achieved a lower MAD, the higher ability of A²TV to preserve the fine details of the blood vessel morphology can be observed when comparing Fig. 11e and f. Additionally, in all the examples, one can observe that when TV- L_1 was performed with a high level of TV regularization (bottom row in Figs. 6 and 9) significant smearing of the blood vessels was observed. In contrast, in the A²TV reconstructions, the smearing owing to over-regularization (low values of λ) could be diminished by increasing the anisotropy in the regularization, i.e. reducing the value of k . For the lowest values of k , higher levels of regularization (low values of λ) created undesirable anisotropic vessel-like texture in the reconstructed images.

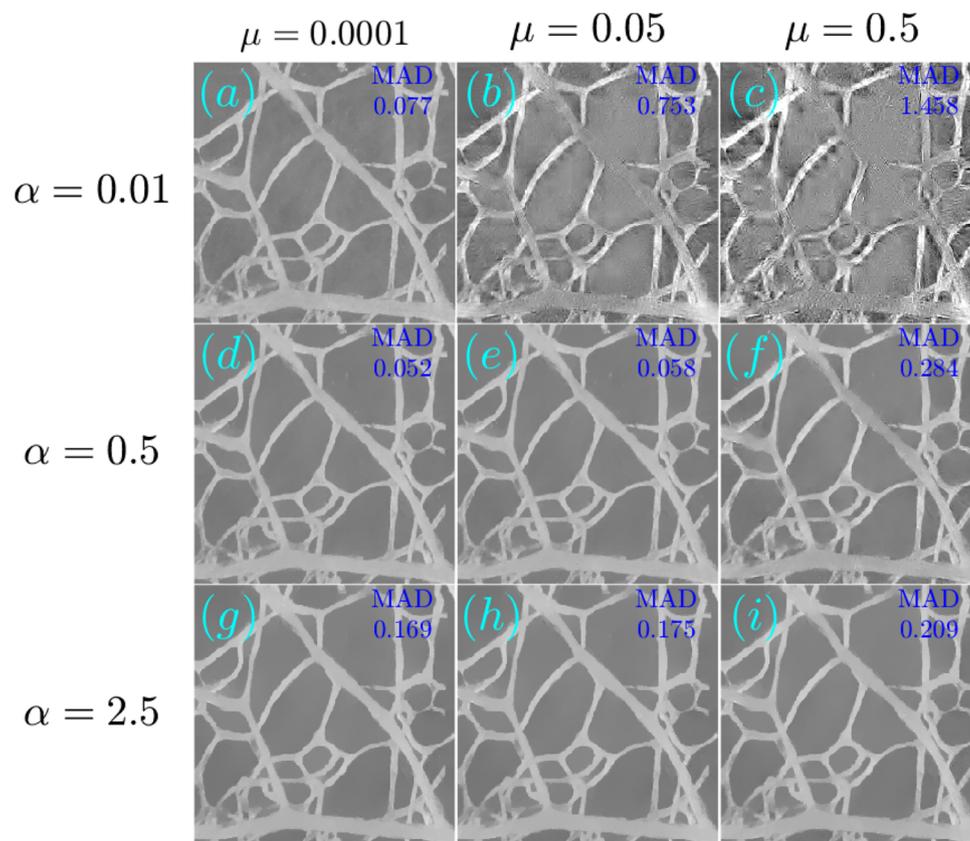


Fig. 9. (a–i) The reconstruction of the image shown in Fig. 5a for the case of under-sampled projection data using different parameters for the TV- L_1 case. The reconstructions were performed with 1000 iterations.

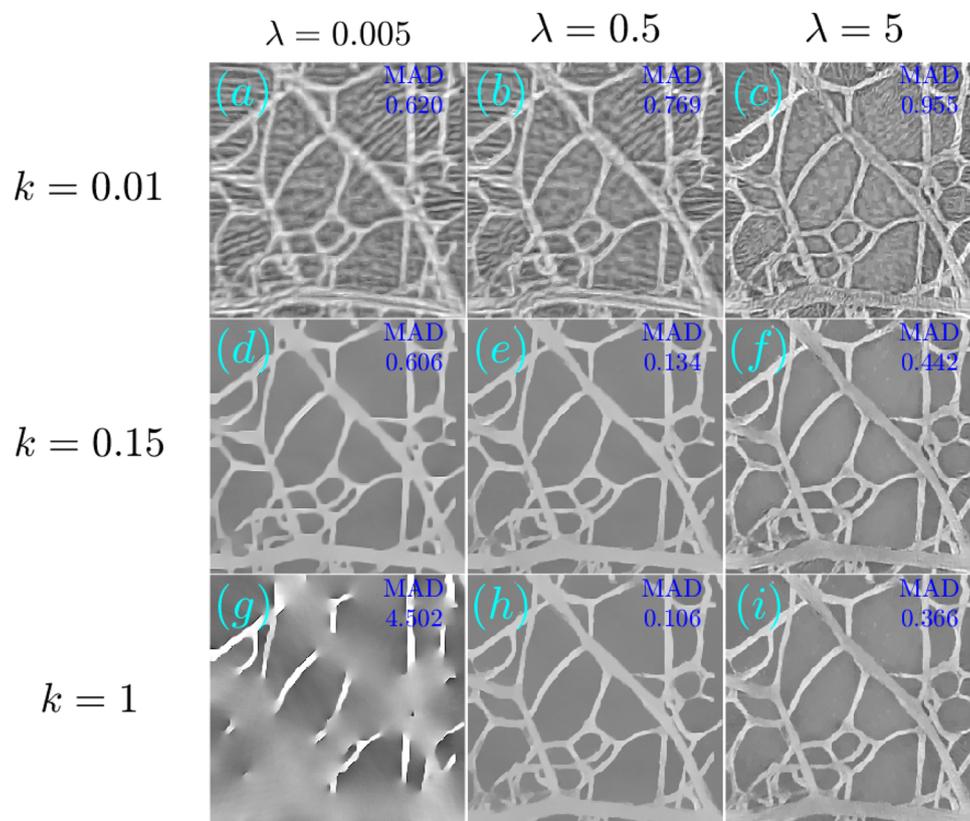


Fig. 10. (a–i) The reconstruction of the image shown in Fig. 5a for the case of under-sampled projection data using different parameters for the A²TV case. The reconstructions were performed with 1500 iterations.

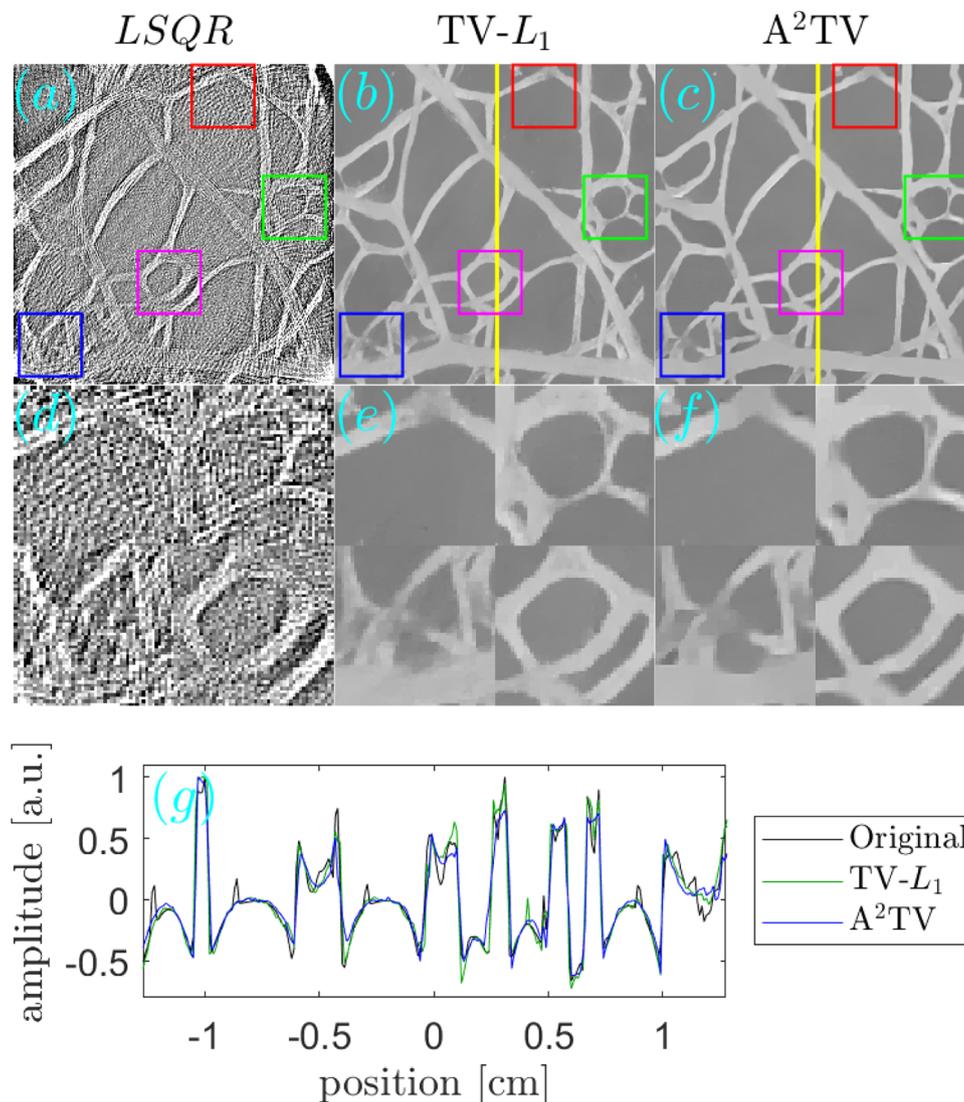


Fig. 11. (a–c) The reconstruction of the image shown in Fig. 5a for the case of under-sampled projection data using (a) LSQR, (b) TV- L_1 , and (c) A^2TV . The bottom panel (d–f) is as in Fig. 8. The TV- L_1 and A^2TV reconstructions were produced using 1000 and 1500 iterations, respectively. (g) A 1D slice of the TV- L_1 and A^2TV reconstructions, corresponding to the vertical lines on the top panel, in comparison to a 1D slice taken from the original image.

5. Experimental results

To further validate the suitability of A^2TV regularization for OAT image reconstruction, we tested its performance on experimental data. The optoacoustic setup comprised an optical parametric oscillator (OPO), which produced nanosecond optical pulses with an energy 30 mJ at a repetition rate of 100 Hz and at a wavelength of $\lambda = 680$ nm. (SpitLight DPSS 100 OPO, InnoLas Laser GmbH, Krailling, Germany). The OPO pulses were delivered to the imaged object using a fiber bundle (CeramOptec GmbH, Bonn, Germany). Ultrasound detection was performed by a 256-element annular array (Imasonic SAS, Voray sur l'Ognon, France) with a radius of 4 cm, and an angular coverage of 270 degrees, comparable to the geometry shown in Fig. 4. The ultrasound detectors were cylindrically focused to a plane, approximating a 2D imaging scenario.

The imaged object was a transparent agar phantom which contained four intersecting hairs. A photo of the phantom is shown in Fig. 12a, where Fig. 12b shows 4 magnified parts of the phantom. The phantom preparation involved mixing 1.3% (by weight) agar powder (Sigma-Aldrich, St. Louis, MO) in boiling water and pouring the solution in a cylindrical mold until solidification. To assure that all four hair strands lie approximately in the same plane, we first prepared a clear

cylindrical agar phantom, on which the hairs were placed; additional agar solution was then poured on the structure to seal the hairs.

The optoacoustic image was reconstructed from the measured data using TV- L_1 regularization with 3000 iterations and A^2TV regularization with 6000 iterations, $\sigma = 1.5$ pixels, and $\rho = 1$ pixel. Figs. 13 and 14 respectively show the images obtained via TV- L_1 and A^2TV reconstructions for a set of regularization parameters. As in the previous section, over-regularization in the TV- L_1 led to significant loss of structure, whereas for the A^2TV the ability to capture the image morphology under strong regularization (low λ) was improved when the anisotropy was increased via low values of k .

Fig. 15 compares between the unregularized LSQR reconstruction (Fig. 15a) and the TV- L_1 (Fig. 15b) and A^2TV (Fig. 15c) reconstructions, respectively taken from Figs. 13e and 14e. The second row in the figure (Fig. 15d–f) shows a magnification of 4 patches from the three reconstructions of the top row (Fig. 15a–c), whereas the bottom panel (Fig. 15g) shows a 1D slice from the vertical yellow line in Fig. 15b and c. To allow an easy comparison, the 1D slices were normalized by their maximum values. We note that the negative values in the reconstruction are a common result of the limited detection bandwidth of the ultrasound detector [6]. The figure clearly shows that the A^2TV

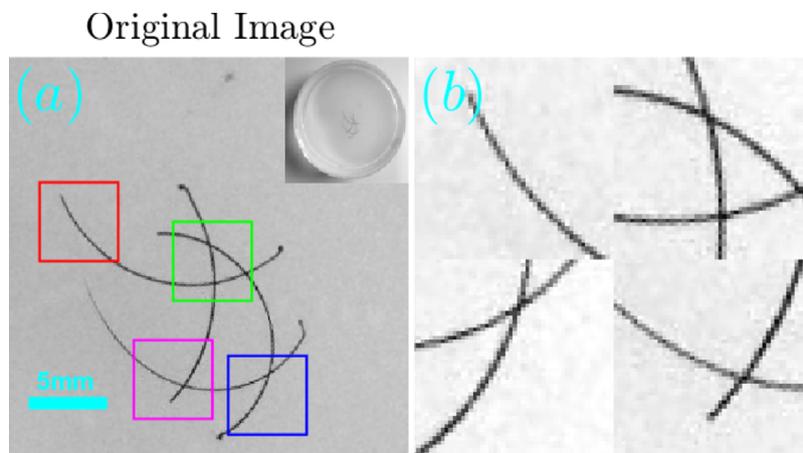


Fig. 12. (a) The originating image on which all the reconstructions were performed and (b) a binary mask that was generated from it. The bottom panel (c and d) shows magnifications of 4 marked regions in the respective images in the top panel.

reconstruction obtained the highest image quality, in particular for the weak hair structure at the image bottom. In particular, the 1D slice shows that the bottom hairs that appear around the position of 1 cm achieve a peak-to-peak signal over 4-times higher in the A^2TV reconstruction than in the $TV-L_1$ reconstruction.

6. Discussion

In this work, a novel regularization framework was developed for OAT image reconstruction. The new framework is based on an A^2TV

cost function, which represents a generalization of the conventional TV functional, that is compatible with objects that possess complex, non-convex boundaries. In contrast to TV, the A^2TV cost function is an adaptive functional whose form depends on the characteristics of the objects in the image. When using A^2TV , one first roughly defines the boundaries of the objects in the image. Then, one uses these boundaries to determine the directions in which the gradients are applied on the image. Similar to TV, A^2TV is most appropriate as a regularizer in optoacoustic image reconstruction when the images are comprised of objects with well-defined boundaries. However, A^2TV is useful also when these

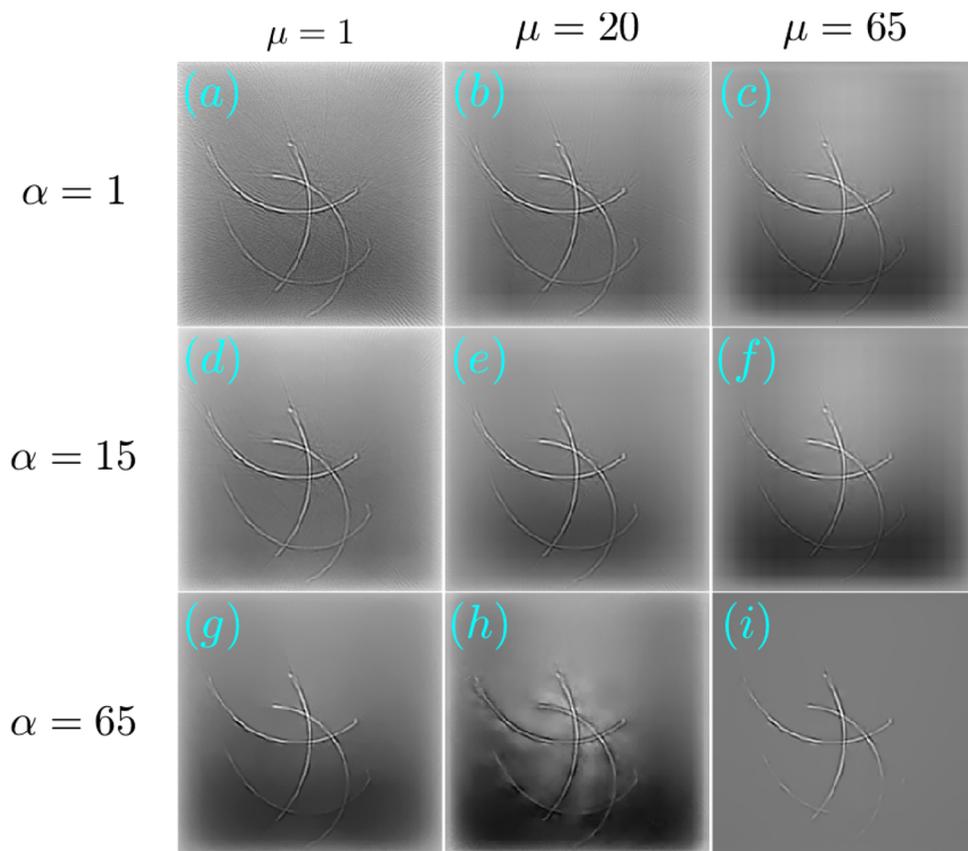


Fig. 13. (a–i) The reconstruction of the image shown in Fig. 12a for the case of experimental data using different parameters for the $TV-L_1$ case. The reconstructions were performed with 3000 iterations.

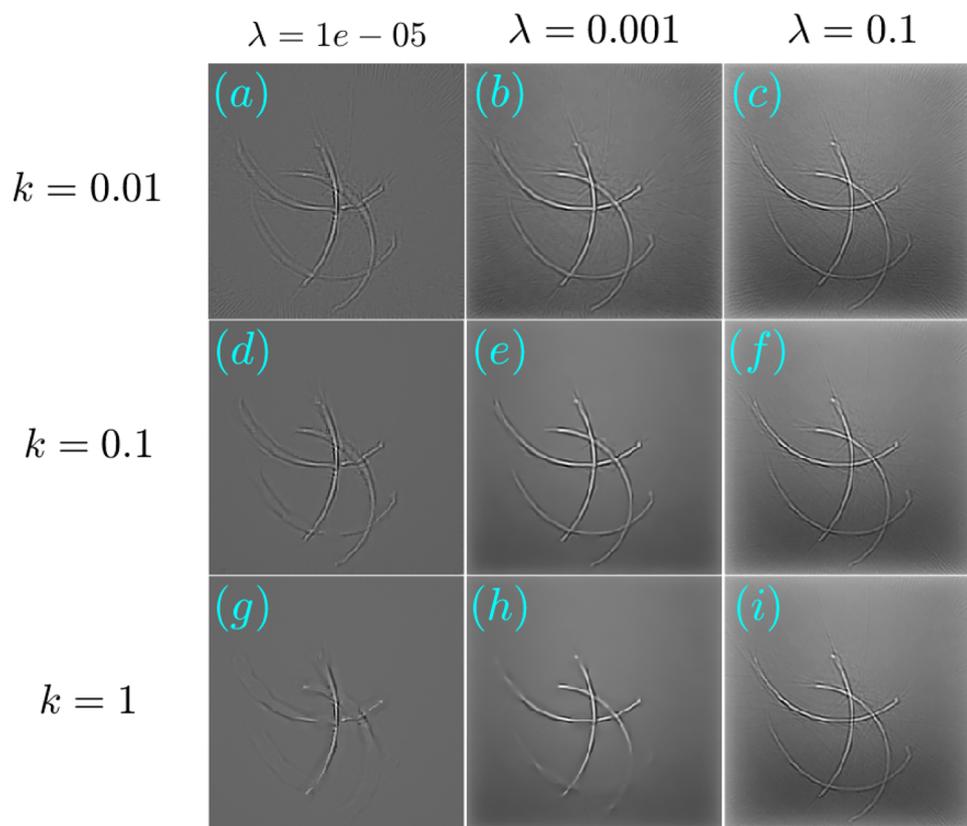


Fig. 14. (a–i) The reconstruction of the image shown in Fig. 12a for the case of experimental data using different parameters for the A^2TV case. The reconstructions were performed with 6000 iterations.

boundaries are incompatible with TV regularization due to their complexity. A common category of optoacoustic images that fits the above description is blood-vessel images. Since blood is a major source of contrast in optoacoustic imaging, OAT systems often produce images that are dominated by a complex structure of interwoven blood vessels. In particular, high-resolution images of the micro-vasculature are characterized by a complex network of arterioles, venules, and capillaries with extremely complex, nonconvex boundaries.

In our current implementation, A^2TV required setting 4 parameters: σ , ρ , λ , and k . The first two parameters, σ and ρ , determined the image smoothing used in calculating the image gradients. While smoothing reduces the noise, thus limiting the effect of reconstruction errors on the detection of the principle axes, only moderate smoothing may be used without the risk of merging the boundaries of different objects in the image. Therefore, in all our examples, the smoothing was performed with Gaussian kernels whose standard deviations, σ and ρ , were 3 pixels or less. While the choice of σ and ρ depended on the level of noise and artifacts in the regularization-free reconstruction, the parameters λ and k were determined by the amount of regularization desired in the reconstructed image, where k determined the amount of anisotropy and λ determined the strength of the regularizer. In our examples, performing over-regularization ($\lambda = 0.0001$) led to image smearing in the isotropic case ($k = 1$) and vessel-like artifacts in the case of high anisotropy ($k = 0.01$). While such artifacts are undesirable, it is worth noting that their presence did not obscure the underlying image morphology, whereas over-regularization in $TV-L_1$ led to loss of image details.

We compared the performance of the A^2TV algorithm $TV-L_1$ algorithm for both numerically simulated data and experimental data. In the numerical simulations, an image of blood vessels was reconstructed for the cases additive noise and sparse sampling of the projection data. In the experimental reconstructions, the imaged object was four intersecting hair strands whose structure emulated the morphology of blood vessels. In both the numerical and experimental examples, A^2TV

demonstrated a higher ability to preserve the blood-vessel morphology for high regularization parameters. Nonetheless, in the numerical example in which the reconstructions were performed with a low number of projections, $TV-L_1$ regularization achieved a lower MAD owing to the loss of texture in the A^2TV reconstruction. In the experimental example, A^2TV led to a considerable improvement in the contrast in the weak structures of the image in comparison to the $TV-L_1$ reconstruction.

The reconstruction performance demonstrated in this work suggests that A^2TV may be a useful tool for improving the ability of optoacoustic systems to perform vasculature imaging, which is a major application in the field. Since the texture within the blood vessels is affected by the random distribution of the red-blood cells within the blood vessels, its elimination by A^2TV may be considered as an acceptable price for better visualization of the blood-vessel morphology. In deep-tissue OAT systems, sub-millimeter vasculature imaging has been suggested as a potential diagnostic tool and has been demonstrated in the human extremities [38,39] and breast [40]. We note that while some OAT systems can also produce images of the tissue bulk, characterized by low frequencies and representative of the density of the microvasculature and fluence map, such systems require transducers capable of detecting ultrasound frequencies considerably below 1 MHz [6]. In high-resolution OAT systems, operating at frequencies above 1 MHz and capable of reaching resolutions better than $100\ \mu\text{m}$ [41,42], only signals from blood vessels may be detected. Finally, when performing optoacoustic imaging at resolutions better than $10\ \mu\text{m}$, e.g. using raster-scan optoacoustic mesoscopy (RSOM) [43], the image is dominated by the microvasculature and generally lacks any bulk component associated with blood.

We note that the formalism of A^2TV , in which the structure tensor matrix analysis is performed via eigenvalue decomposition Eq. (13) enables its adaptation to higher image dimensions. TV regularization has been recently performed in 4D optoacoustic reconstruction that included 3 spatial dimensions and time [44]. Since the variation of the pixels in time is generally different than the one space, it may be

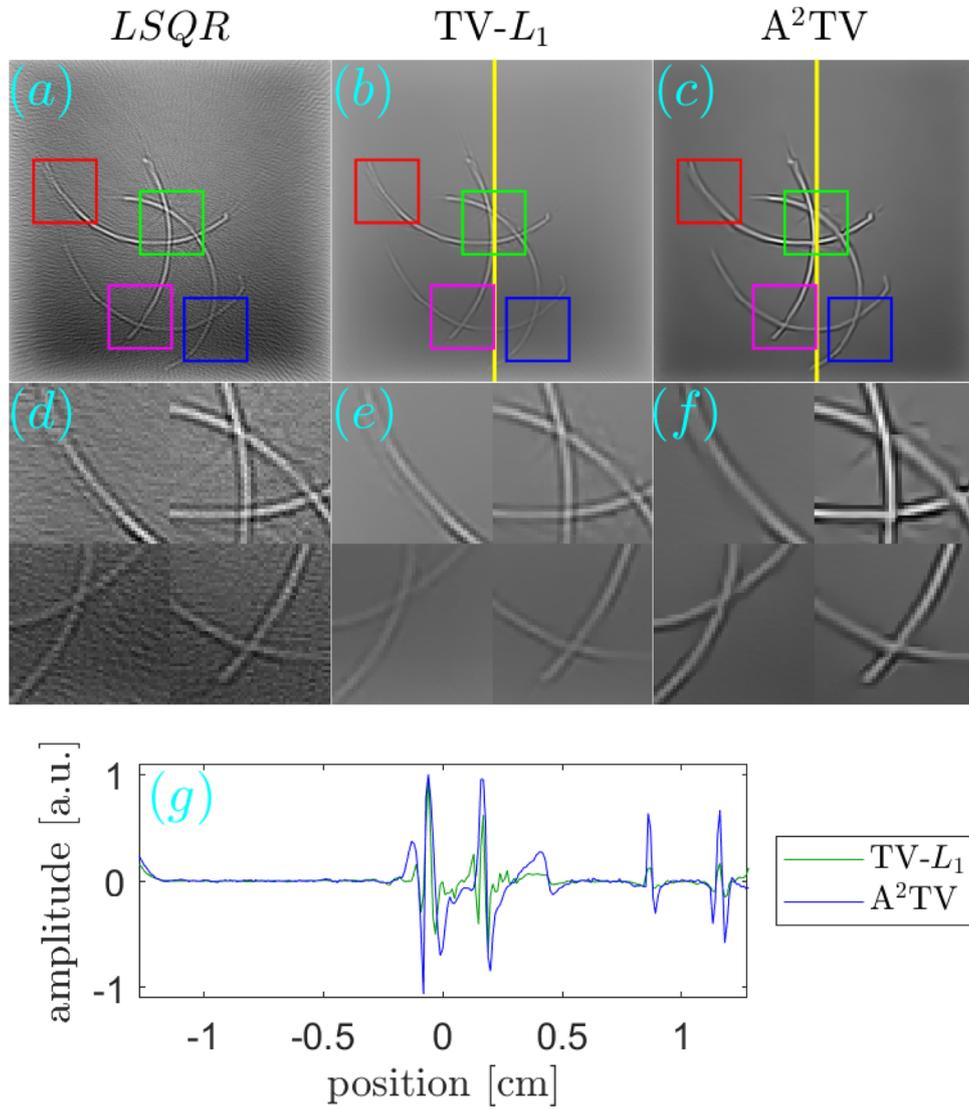


Fig. 15. (a–c) The reconstruction of the image shown in Fig. 12a for the case of experimental data using (a) LSQR, (b) TV- L_1 , and (c) A²TV. The TV- L_1 and A²TV reconstructions were produced using 3000 and 6000 iterations, respectively. (g) A 1D slice of the TV- L_1 and A²TV reconstructions, corresponding to the vertical lines on the top panel.

expected that A²TV can further improve image fidelity in such cases.

Acknowledgement

We are thankful to Anna M. Randi for providing information on the image used in Section 4. This work was supported by the Technion Ollendorff Minerva Center.

Conflicts of interest

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

Appendix A. Solving the generalized ROF model

A²TV based reconstruction amounts to solving the following generalized ROF model,

$$\mathbf{u}^* = \operatorname{argmin}_{\mathbf{u} \in X} \mathcal{E}(\mathbf{u}) = \operatorname{argmin}_{\mathbf{u} \in X} \mathcal{J}_{A^2TV}(\mathbf{u}) + \frac{\lambda}{2} \|\mathbf{M}\mathbf{u} - \mathbf{p}\|_2^2. \quad (\text{A.1})$$

Where \mathbf{M} is the model matrix, \mathbf{p} is the acoustic pressure wave and \mathbf{u}^* is the reconstructed image. The algorithm used for inversion is Chambolle-Pock, which solves the following primal form [36]:

$$\operatorname{argmin}_{\mathbf{u} \in X} F(K\mathbf{u}) + G(\mathbf{u}). \quad (\text{A.2})$$

Where, in our case,

$$F(\mathbf{y}) = \|\mathbf{y}\|_1, \quad G(\mathbf{u}) = \frac{\lambda}{2} \|\mathbf{M}\mathbf{u} - \mathbf{p}\|_2^2, \quad K = \nabla_A = \operatorname{grad}_A, \quad K^* = \nabla_A^T = \operatorname{div}_A \quad (\text{A.3})$$

Let X and Y be two finite-dimensional real vector spaces and K , a linear operator $K : X \rightarrow Y$ and its hermitian adjoint $K^* : Y \rightarrow X$. The appropriate saddle point problem is,

$$\operatorname{argmin}_{\mathbf{u} \in X} \operatorname{argmax}_{\mathbf{z} \in Y} \langle K\mathbf{u}, \mathbf{z} \rangle + G(\mathbf{u}) - F^*(\mathbf{z}). \quad (\text{A.4})$$

Where $\mathbf{z} \in Y$ is the dual variable. In our case,

$$\operatorname{argmin}_{\mathbf{u} \in X} \operatorname{argmax}_{\mathbf{z} \in Y} \langle \nabla_A \mathbf{u}, \mathbf{z} \rangle + \frac{\lambda}{2} \|\mathbf{M}\mathbf{u} - \mathbf{p}\|_2^2 - \delta_P(\mathbf{z}), \quad (\text{A.5})$$

where

$$F^*(\mathbf{z}) = \delta_P(\mathbf{z}), \quad (\text{A.6})$$

and

$$\delta_P(\mathbf{z}) = \begin{cases} 0, & \mathbf{z} \in P \\ \infty, & \mathbf{z} \notin P \end{cases}, \quad P = \{\mathbf{z} \in Y \mid \|\mathbf{z}\|_\infty \leq 1\}. \quad (\text{A.7})$$

But, because of the appearance of the model matrix, we shall adjust the saddle point problem to

$$\begin{aligned} & \operatorname{argmin}_{\mathbf{u} \in X} \operatorname{argmax}_{\mathbf{z} \in Y, \mathbf{q} \in Q} \langle \nabla_A \mathbf{u}, \mathbf{z} \rangle + \langle \mathbf{M}\mathbf{u} - \mathbf{p}, \mathbf{q} \rangle - \frac{1}{2\lambda} \|\mathbf{q}\|_2^2 - \delta_P(\mathbf{z}) \\ &= \operatorname{argmin}_{\mathbf{u} \in X} \operatorname{argmax}_{\mathbf{z} \in Y, \mathbf{q} \in Q} \langle \nabla_A \mathbf{u}, \mathbf{z} \rangle + \langle \mathbf{M}\mathbf{u}, \mathbf{q} \rangle - \langle \mathbf{p}, \mathbf{q} \rangle \\ &\quad - \frac{1}{2\lambda} \|\mathbf{q}\|_2^2 - \delta_P(\mathbf{z}) \\ &= \operatorname{argmin}_{\mathbf{u} \in X} \operatorname{argmax}_{\mathbf{z} \in Y, \mathbf{q} \in Q} \left\langle \begin{bmatrix} \mathbf{M} \\ \nabla_A \end{bmatrix} \mathbf{u}, \begin{bmatrix} \mathbf{q} \\ \mathbf{z} \end{bmatrix} \right\rangle - \langle \mathbf{p}, \mathbf{q} \rangle - \frac{1}{2\lambda} \|\mathbf{q}\|_2^2 - \delta_P(\mathbf{z}), \end{aligned} \quad (\text{A.8})$$

using

$$\frac{\lambda}{2} \|\mathbf{M}\mathbf{u} - \mathbf{p}\|_2^2 = \operatorname{argmax}_{\mathbf{q} \in Q} \langle \mathbf{M}\mathbf{u} - \mathbf{p}, \mathbf{q} \rangle - \frac{1}{2\lambda} \|\mathbf{q}\|_2^2. \quad (\text{A.9})$$

Here, the original saddle point problem is transformed into,

$$\operatorname{argmin}_{\mathbf{u} \in X} \operatorname{argmax}_{\mathbf{z} \in Y, \mathbf{q} \in Q} \langle K\mathbf{u}, \begin{bmatrix} \mathbf{q} \\ \mathbf{z} \end{bmatrix} \rangle + G(\mathbf{u}) - F^*\left(\begin{bmatrix} \mathbf{q} \\ \mathbf{z} \end{bmatrix}\right). \quad (\text{A.10})$$

Where

$$\begin{aligned} F^*\left(\begin{bmatrix} \mathbf{q} \\ \mathbf{z} \end{bmatrix}\right) &= \langle \mathbf{p}, \mathbf{q} \rangle + \frac{1}{2\lambda} \|\mathbf{q}\|_2^2 + \delta_P(\mathbf{z}), \\ G(\mathbf{u}) &= 0, \\ K &= \begin{bmatrix} K_A \\ K_{\nabla_A} \end{bmatrix} = \begin{bmatrix} \mathbf{M} \\ \nabla_A \end{bmatrix} \\ K^* &= [K_A^* \quad K_{\nabla_A}^*] = [\mathbf{M}^T \quad -\nabla_A^T]. \end{aligned} \quad (\text{A.11})$$

Now, in order to use the algorithm, we shall use the proximal operator of F^* and G ,

$$\begin{aligned} \operatorname{Prox}_{\sigma, F^*}\left(\begin{bmatrix} \tilde{\mathbf{q}} \\ \tilde{\mathbf{z}} \end{bmatrix}\right) &= \left(\frac{\tilde{\mathbf{q}} - \sigma \mathbf{p}}{1 + \frac{\sigma}{\lambda}}, \left(\frac{\tilde{\mathbf{z}}}{\max(1, |\tilde{\mathbf{z}}|)} \right)_{\text{pointwise}} \right), \\ \operatorname{Prox}_{\tau, G}(\tilde{\mathbf{u}}) &= \tilde{\mathbf{u}} \end{aligned} \quad (\text{A.12})$$

Algorithm 1. The Chambolle-Pock algorithm.

<ol style="list-style-type: none"> 1: 2: 3: 4: 5: 6: 7: 8: 9: 10: 11: 12: 13: 14: 15: 	<pre> procedure CHAMBOLLE-POCKProx$_{\sigma_n, F^*}(\cdot)$, Prox$_{\tau_n, G}(\cdot)$, K, K^* Initialize γ, τ_0 $\sigma_0 \leftarrow 1/(\tau_0 L^2)$ Initialize u^0, q^0, z^0 $i \leftarrow 0$ while $i < N$ do $[q^{n+1}, z^{n+1}] \leftarrow \operatorname{Prox}_{\sigma_n, F^*}(q^n + \sigma_n K_A(u^n), z^n + \text{xxxxxx}\sigma_n K_V(u^n))$ $[u^{n+1}, z^{n+1}] \leftarrow \operatorname{Prox}_{\tau_n, G}(u^n - \tau_n K^*([q^{n+1}, z^{n+1}]))$ $\theta_{n+1} \leftarrow 1/\sqrt{1 + 2\gamma\tau_n}$ $\tau_{n+1} \leftarrow \theta_{n+1}\tau_n$ $\sigma_{n+1} \leftarrow \sigma_n/\theta_{n+1}$ Update tensor A according to the new image u^{n+1} and its dependents K^* & K end while return u^N end procedure </pre>
--	---

The algorithm is given in Fig. 1, where in our case, we normalize the model matrix \mathbf{M} by $\frac{1}{160} \sqrt{\|\mathbf{M}\|_\infty \|\mathbf{M}\|_1}$, where $\sqrt{\|\mathbf{M}\|_\infty \|\mathbf{M}\|_1}$ is an approximation to the model matrix M Lipchitz constant. $L = L_M + L_V = 160 + 8 = 168$, $\gamma = 0.7\lambda$ and $\tau_0 = 0.5$

References

- [1] V. Ntziachristos, Going deeper than microscopy: the optical imaging frontier in biology, *Nat. Methods* 7 (2010) 603.
- [2] L.V. Wang, S. Hu, Photoacoustic tomography: in vivo imaging from organelles to organs, *Science* 335 (2012) 1458–1462.
- [3] V. Ntziachristos, D. Razansky, Molecular imaging by means of multispectral photoacoustic tomography (msot), *Chem. Rev.* 110 (2010) 2783–2794.
- [4] A. Taruttis, V. Ntziachristos, Advances in real-time multispectral photoacoustic imaging and its applications, *Nat. Photonics* 9 (2015) 219.
- [5] L.V. Wang, Multiscale photoacoustic microscopy and computed tomography, *Nat. Photonics* 3 (2009) 503.
- [6] A. Rosenthal, V. Ntziachristos, D. Razansky, Acoustic inversion in photoacoustic tomography: a review, *Curr. Med. Imaging Rev.* 9 (2013) 318–336.
- [7] A. Buehler, E. Herzog, D. Razansky, V. Ntziachristos, Video rate photoacoustic tomography of mouse kidney perfusion, *Optics Lett.* 35 (2010) 2475–2477.
- [8] X.L. Deán-Ben, S.J. Ford, D. Razansky, High-frame rate four dimensional photoacoustic tomography enables visualization of cardiovascular dynamics and mouse heart perfusion, *Sci. Rep.* 5 (2015) 10133.
- [9] Y. Xu, L.V. Wang, Time reversal and its application to tomography with diffracting sources, *Phys. Rev. Lett.* 92 (2004) 033902.
- [10] M. Xu, L.V. Wang, Universal back-projection algorithm for photoacoustic computed tomography, *Phys. Rev. E* 71 (2005) 016706.
- [11] R.J. Schulze, O. Scherzer, G. Zangerl, M. Holotta, D. Meyer, F. Handle, R. Nuster, G. Paltauf, On the use of frequency-domain reconstruction algorithms for photoacoustic imaging, *J. Biomed. Optics* 16 (2011) 086002.
- [12] A. Rosenthal, Algebraic determination of back-projection operators for photoacoustic tomography, *Biomed. Opt. Express* 9 (2018) 5173–5193.
- [13] D. Calvetti, S. Morigi, L. Reichel, F. Sgallari, Tikhonov regularization and the l-curve for large discrete ill-posed problems, *J. Comput. Appl. Math.* 123 (2000) 423–446.
- [14] Z. Yuan, C. Wu, H. Zhao, H. Jiang, Imaging of small nanoparticle-containing objects by finite-element-based photoacoustic tomography, *Optics Lett.* 30 (2005) 3054–3056.
- [15] Y. Han, S. Tzoumas, A. Nunes, V. Ntziachristos, A. Rosenthal, Sparsity-based acoustic inversion in cross-sectional multiscale photoacoustic imaging, *Med. Phys.* 42 (2015) 5444–5452.
- [16] J. Provost, F. Lesage, The application of compressed sensing for photoacoustic tomography, *IEEE Trans. Med. Imaging* 28 (2009) 585–594.
- [17] Z. Guo, C. Li, L. Song, L.V. Wang, Compressed sensing in photoacoustic tomography in vivo, *J. Biomed. Optics* 15 (2010) 021311.
- [18] D. Liang, H.F. Zhang, L. Ying, Compressed-sensing photoacoustic imaging based on random optical illumination, *Int. J. Funct. Inform. Pers. Med.* 2 (2009) 394–406.
- [19] M. Sun, N. Feng, Y. Shen, X. Shen, L. Ma, J. Li, Z. Wu, Photoacoustic imaging method based on arc-direction compressed sensing and multi-angle observation, *Optics Express* 19 (2011) 14801–14806.
- [20] J. Meng, L.V. Wang, L. Ying, D. Liang, L. Song, Compressed-sensing photoacoustic computed tomography in vivo with partially known support, *Optics Express* 20 (2012) 16510–16523.
- [21] X. Liu, D. Peng, W. Guo, X. Ma, X. Yang, J. Tian, Compressed sensing photoacoustic imaging based on fast alternating direction algorithm, *J. Biomed. Imaging* 2012 (2012) 12.
- [22] K. Wang, E.Y. Sidky, M.A. Anastasio, A.A. Oraevsky, X. Pan, Limited data image reconstruction in photoacoustic tomography by constrained total variation minimization, *Photons Plus Ultrasound: Imaging and Sensing 2011* Vol. 7899 International Society for Optics and Photonics, 2011, p. 78993U.
- [23] L. Yao, H. Jiang, Enhancing finite element-based photoacoustic tomography using total variation minimization, *Appl. Optics* 50 (2011) 5031–5041.
- [24] L. Yao, H. Jiang, Photoacoustic image reconstruction from few-detector and limited-angle data, *Biomed. Optics Express* 2 (2011) 2649–2654.
- [25] J. Wang, C. Zhang, Y. Wang, A photoacoustic imaging reconstruction method based on directional total variation with adaptive directivity, *Biomed. Eng. Online* 16 (2017) 64.
- [26] S. Biton, G. Gilboa, Adaptive Anisotropic Total Variation — A Nonlinear Spectral Analysis, (2018) arXiv:1811.11281.
- [27] F. Andreu, C. Ballester, V. Caselles, J.M. Mazón, Minimizing total variation flow, *Differ. Integral Equ.* 14 (2001) 321–360.
- [28] M. Peura, J. Iivarinen, Efficiency of simple shape descriptors, *Proceedings of the Third International Workshop on Visual Form*, vol. 443, Citeseer, 1997, p. 451.
- [29] R.A. Kruger, P. Liu, Y. Fang, C.R. Appledorn, et al., Photoacoustic ultrasound (paus)-reconstruction tomography, *Med. Phys.* 22 (1995) 1605–1609.
- [30] D.T. Blackstock, *Fundamentals of Physical Acoustics*, John Wiley & Sons, 2000.
- [31] G. Paltauf, J. Viator, S. Prahl, S. Jacques, Iterative reconstruction algorithm for photoacoustic imaging, *J. Acoust. Soc. Am.* 112 (2002) 1536–1544.
- [32] A. Rosenthal, D. Razansky, V. Ntziachristos, Fast semi-analytical model-based acoustic inversion for quantitative photoacoustic tomography, *IEEE Trans. Med. Imaging* 29 (2010) 1275–1285.
- [33] C.C. Paige, M.A. Saunders, Lsq: an algorithm for sparse linear equations and sparse least squares, *ACM Trans. Math. Softw.* 8 (1982) 43–71.
- [34] J. Weickert, *Anisotropic Diffusion in Image Processing* vol. 1, Teubner Stuttgart, 1998.
- [35] M. Grasmair, F. Lenzen, Anisotropic total variation filtering, *Appl. Math. Optim.* 62 (2010) 323–339.
- [36] A. Chambolle, T. Pock, A first-order primal-dual algorithm for convex problems with applications to imaging, *J. Math. Imaging Vis.* 40 (2011) 120–145.
- [37] A. Dima, N.C. Burton, V. Ntziachristos, Multispectral photoacoustic tomography at 64, 128, and 256 channels, *J. Biomed. Optics* 19 (2014) 036021.
- [38] P. Wray, L. Lin, P. Hu, L.V. Wang, Photoacoustic computed tomography of human extremities, *J. Biomed. Optics* 24 (2019) 026003.
- [39] Y. Matsumoto, Y. Asao, A. Yoshikawa, H. Sekiguchi, M. Takada, M. Furu, S. Saito, M. Kataoka, H. Abe, T. Yagi, et al., Label-free photoacoustic imaging of human palmar vessels: a structural morphological analysis, *Sci. Rep.* 8 (2018) 786.
- [40] M. Toi, Y. Asao, Y. Matsumoto, H. Sekiguchi, A. Yoshikawa, M. Takada, M. Kataoka, T. Endo, N. Kawaguchi-Sakita, M. Kawashima, et al., Visualization of tumor-related blood vessels in human breast by photoacoustic imaging system with a hemispherical detector array, *Sci. Rep.* 7 (2017) 41970.
- [41] X.L. Deán-Ben, H. López-Schier, D. Razansky, Photoacoustic micro-tomography at 100 volumes per second, *Sci. Rep.* 7 (2017) 6850.
- [42] J. Gateau, A. Chekkoury, V. Ntziachristos, Ultra-wideband three-dimensional photoacoustic tomography, *Optics Lett.* 38 (2013) 4671–4674.
- [43] M. Omar, M. Schwarz, D. Soliman, P. Symvoulidis, V. Ntziachristos, Pushing the optical imaging limits of cancer with multi-frequency-band raster-scan photoacoustic mesoscopy (rsom), *Neoplasia* 17 (2015) 208–214.
- [44] A. Özbek, X.L. Deán-Ben, D. Razansky, Photoacoustic imaging at kilohertz volumetric frame rates, *Optica* 5 (2018) 857–863.



Shai Biton received his B.Sc. degree from the Department of Electrical Engineering from Ort Braude, Karmiel, in 2011. Currently he pursues his MSc degree in Electrical Engineering from the Technion – Israel Institute of Technology, Haifa. His current research interests are in medical imaging, thermal imaging, image reconstruction and deep learning.



Nadav Yehonatan Arbel is an M.Sc. Candidate in the Viterbi Faculty of Electrical Engineering at the Technion – Israel Institute of Technology, where he formerly obtained a double B.Sc. in Electrical Engineering and Physics. His research interests include Matching of 3D Shapes, Medical Signal Processing and Optics.



Gilad Drozdov received his B.Sc. degree from the Department of Electrical Engineering from the Technion-Israel Institute of Technology, Haifa, in 2012, where he currently pursues his MSc degree in Electrical Engineering. His current research interests are in photoacoustic imaging, acoustic transducers modeling and image reconstruction.



Guy Gilboa received his PhD from the Electrical Engineering Department, Technion – Israel Institute of Technology in 2004. He was a postdoctoral fellow at UCLA, hosted by Prof. Stanley Osher (2004–2007). He later joined 3DV, a start-up company developing 3D sensors, which was acquired by Microsoft. He worked at Microsoft (Kinect project) and at Philips Healthcare as a senior researcher in the field of medical imaging. In 2013 he joined the Electrical Engineering Dept., Technion, for a tenure-track position. He is the main author of some highly cited papers on topics such as image sharpening and denoising, nonlocal operators theory for image processing and nonlinear spectral theory. He received several prizes including the Eshkol

prize by the Ministry of Science, Technion outstanding Ph.D. thesis, Vatat scholarship and Gutwirth prize. He serves at the editorial boards of the journals IEEE Signal Processing Letters, Journal of Mathematical Imaging and Vision and Journal of Computer Vision and Image Understanding.



Amir Rosenthal received his B.Sc. and Ph.D. degrees, both from the Department of Electrical Engineering, the Technion-Israel Institute of Technology, Haifa, in 2002 and 2006, respectively. From 2009 to 2010, he was a Marie Curie fellow at the Cardiovascular Research Center at Massachusetts General Hospital and Harvard Medical School, Boston, MA, and from 2010 to 2014 he was a group leader at the Institute of Biological and Medical Imaging (IBMI), Technische Universität München and Helmholtz Zentrum München, Germany. Since 2015 he has been an assistant professor at the Andrew and Erna Viterbi Faculty of Electrical Engineering at the Technion - Israel Institute of Technology. His research interests include optoacoustic

imaging, interferometric sensing, intravascular imaging, inverse problems, and optical and acoustical modeling.