PROCEEDINGS OF SPIE

Robustness of a partially learned photoacoustic reconstruction algorithm

Yoeri E. Boink, Christoph Brune, Srirang Manohar

SPIE.

Robustness of a partially learned photoacoustic reconstruction algorithm

Yoeri E. Boink[1,2], Christoph Brune[2], and Srirang Manohar[1]

[1]Biomedical Photonic Imaging, Technical Medical Centre, University of Twente, the Netherlands
[2]Department of Applied Mathematics, Technical Medical Centre, University of Twente, the Netherlands

ABSTRACT

Classical reconstruction algorithms for photoacoustic tomography (PAT) are mathematically proven to converge, but can be very slow and inadequate with respect to model and data assumptions. With the help of neural networks, learned reconstruction algorithms have recently been developed. These learned algorithms have shown to surpass the reconstruction quality of non-learned ones, but their mathematical analysis is challenging, therefore convergence and stability are not guaranteed. In this work, we combine the structure of a well-known model-based algorithm with the efficiency of a data-driven neural network. We show its robustness in simulation against uncertainty and changes in PAT system settings. This is done by varying the placement and calibration of detectors, as well as varying properties of the synthetic phantom, such as geometrical structure and intensity. This yields a robust and computationally efficient algorithm that outperforms classical methods and can be applied to the PAT problem in a universal setting.

Keywords: photoacoustic tomography, iterative image reconstruction, neural networks, robustness.

1 INTRODUCTION

Photoacoustic tomography (PAT) is a useful technique for imaging vascular networks, e.g. for detection of breast cancer[1–3] or monitoring rheumatoid arthritis.[4] When measurements are made according to the Nyquist criterion, direct analytical reconstruction methods such as filtered backprojection (FBP) are available.[5] In practice, it is often not possible or too time-consuming to take sufficient samples, which makes the reconstruction problem challenging.[6] This is referred to as the limited angle or limited sampling problem. On top of this, uncertainties are ubiquitous: measured data is corrupted by noise, the modelling of the photoacoustic operator is not perfect and the type of images that are reconstructed is not always known beforehand.

Under these circumstances, direct reconstruction methods fail to give acceptable results. Regularised model-based methods are better able to cope with uncertainty,[7,8] but manually selecting the regulariser and high computational costs still lead to suboptimal methods. In this work, we make use of the learned primal-dual[9] (L-PD) algorithm, which can be seen as a learned adaption of the widely used primal-dual hybrid gradient algorithm (PDHG):[10] part of the PDHG algorithm has been replaced with a convolutional neural network (CNN) to improve quality of the reconstruction and reduce computational costs. We demonstrate ways to train the learned algorithm in such a way that it is robust against uncertainties in the image, data and system settings.

Recently there have been other works that combine CNNs with classical reconstruction algorithms,[11,12] also with the application of PAT reconstruction.[13,14] However, for the majority of these methods, stability has not be proven by using mathematical analysis. We explore the robustness of L-PD in an empirical way. This yields a robust method that gives an exceptional reconstruction quality for photoacoustic tomography in a universal setup.

---

e-mail: y.e.boink@utwente.nl

2 LEARNED PRIMAL-DUAL ALGORITHM

We compare Algorithm 1 (PDHG) with Algorithm 2, a learned primal-dual (L-PD) algorithm.[9] The inverse problem that is solved with PDHG is generally defined as $\min_u F_f(Au) + G(u)$, where A is a representation of the photoacoustic operator and F is a data-fidelity functional which makes sure that the reconstructed image u fits the data f. Regularisation could be implemented via the functional G. Some examples of the application of PDHG to PAT with different regularisers, such as TV, are given in a previous work.[8] It can be seen in Algorithm 1 that the inverse problem is solved by alternatingly performing a dual (q) update in the data domain and a primal (u) update in the image domain. The photoacoustic operator A and its adjoint connect both updates.

In the L-PD approach, the same general structure is used: both updates take previous iterates of both the primal and dual into account, as well as the data. Since the physics behind the photoacoustic operator A is known, it is explicitly used in the learned algorithm. However, the functionals F and G are not chosen explicitly, but instead the best update formula for each iteration is learned. This is achieved by a CNN, here represented by the nonlinear functions Γ_{Θ_n} and Λ_{Θ_n}, in which Θ_n describes the learned weights for iteration n. Note that weights of the network can be different for every iteration $n \in \{1, \ldots, N\}$. Instead of updating one channel of the primal and dual, we allow the network to use k channels, which may help the algorithm to efficiently make use of information from previous iterates. Besides obtaining higher quality reconstructions, a huge advantage is that we can enforce L-PD to only use a small amount of iterations, wheras in PDHG, one has to wait till convergence. The output of the L-PD reconstruction is the first channel of the last primal iterate: u_1^{N+1}.

---

for $n \leftarrow 1$ to N **do**
$\quad q^{n+1} = \text{prox}_{\sigma F_f^*}\Big(q^n + \sigma A\left[(1+\theta)u^n - \theta u^{n-1}\right]\Big),$
$\quad u^{n+1} = \text{prox}_{\tau G}\Big(u^n - \tau A^* q^{n+1}\Big).$
end for

Algorithm 1: Non-learned (PDHG).

---

for $n \leftarrow 1$ to N **do**
$\quad q_{\{1,\ldots,k\}}^{n+1} = q_{\{1,\ldots,k\}}^n + \Gamma_{\Theta_n}\Big(q_{\{1,\ldots,k\}}^n, Au_1^n, f\Big),$
$\quad u_{\{1,\ldots,k\}}^{n+1} = u_{\{1,\ldots,k\}}^n + \Lambda_{\Theta_n}\Big(u_{\{1,\ldots,k\}}^n, A^* q_1^{n+1}\Big).$
end for

Algorithm 2: Learned (L-PD).

---

3 ACOUSTIC FORWARD MODEL

The learned algorithm is a general formulation, which can be applied for many PAT-systems, like a spherical, cylindrical or slice-based system, all having many different settings. The photoacoustic operator A is generally an approximation of the real physical process: many approximations are available, including ones that also take the optical process into account.[15,16] In this work, we choose to use a ray-based approach to the acoustic forward process in a 2D-setup:

$$p(\mathbf{x}, t) = |\mathbf{x} - \mathbf{x}_p| \left(\frac{1}{t} \iint_{|\mathbf{x} - \tilde{\mathbf{x}}| = ct} u(\tilde{\mathbf{x}}) d\tilde{\mathbf{x}} \right) *_t p_{\text{cal}}(\mathbf{x}, t'). \tag{1}$$

the measured pressure $p(\mathbf{x}, t)$ is the result of a spherical mean transform[17] applied to the initial pressure u convolved with a calibration measurement $p_{\text{cal}}(\mathbf{x}, t')$. This calibration measurement is approximated by measuring the pressure generated by a strong absorbing point source. We assume the possibly inhomogeneous sound speed c to be known. After pre-processing, we arrive at the simplified equation

$$f = Au := \iint_{|\mathbf{x} - \tilde{\mathbf{x}}| = ct} u(\tilde{\mathbf{x}}) d\tilde{\mathbf{x}}. \tag{2}$$

Above expression for the acoustic forward operator A is used in this work. For a more detailed explanation of the calibration measurement and pre-processing, we refer to a previous work.[8] Other choices of the acoustic operator A can be made, such as a k-space pseudospectral method. The L-PD algorithm can also be applied to a PAT reconstruction problem in 3D with different detector setup.[14]

4 SIMULATION SETUP

A great interest for photoacoustics lies in the imaging of vascular structures, e.g. in breast imaging.[1–3] Therefore, images with a high-intensity vascular structure on a low-intensity background are considered in this work. These reflect photoacoustic images under homogeneous illumination. A total of 768 training images and 192 test images have been obtained by preprocessing patches from the openly available DRIVE-dataset.[18] The size of all images is 192×192 pixels. An example is shown in the left of Figure 3. Data (sinograms) is simulated with the forward model as described in section 3. This is done for a setting of 32 photoacoustic detectors, uniformly distributed on a circle.

4.1 Neural network architecture

Two L-PD algorithms with two sets of parameters are trained: a small network and a somewhat larger network. In the first one, we choose the number of iterations $N = 5$ and the number of primal- and dual channels $k = 2$ (cf. Algorithm 2). In the second one, we choose $N = 10$ and $k = 5$. Both networks have 32 channels in 2 hidden layers. The filter size for the convolutions is 3×3 and ReLu's are chosen as activation functions. We make use of the MSE-loss on the difference between ground truth and reconstruction, combined with the Adam optimiser.[19] Batch size increases from 2 to 16 in three steps for a total of 200 Epochs.

4.2 Robustness to image changes

One problem in learned reconstruction methods is that mathematical analysis on the stability of the neural networks is not available. We will empirically show the stability of L-PD to variations in the photoacoustic image; this knowledge can be used to create a learned reconstruction method that is robust to uncertainty in the image. For this purpose both the training and the test set, as described in section 4, have been altered in different ways. This way, 8 classes of images are obtained in which one or multiple properties have been changed; the classes are shown below, with class 0 being the standard class.

class (c_i)	sensitivity to	explanation
0	nothing	*for an example image see left image in Figure 3*
1	diameter	increase in diameter randomly chosen from the set $\{-2, -1, 0, 1, 2\}$
2	contrast	contrast randomly chosen from the set $\{0.5, 0.7, 1, 1.4, 2\}$
3	background	with $p = \frac{1}{2}$ uniform increase, with $p = \frac{1}{2}$ speckled background
4	coverage	with $p = \frac{1}{3}$ removal of smallest vessels, with $p = \frac{1}{3}$ nothing, with $p = \frac{1}{3}$ doubling of vessels
5	structure	inhomogeneous foreground scaled between $[m, 1]$, where m from $\{0.2, 0.4, 0.6, 0.8, 1\}$
6	noise level	Gaussian noise level randomly chosen from $[0, 3\sigma]$, where σ is the standard noise level
7	all of the above	above changes have all been applied with same probabilities

The L-PD algorithm is trained on each of these classes. For testing, the reconstruction quality of L-PD trained on the test class is compared with the reconstruction quality of L-PD trained on the standard class.

4.3 Robustness to system changes

Many imaging systems have multiple settings of which some are determined by its user and some are calibrated. Generally, the number of detectors used for the measurement can be chosen. In our model, as described in section 3, a calibration is performed before every measurement and will be slightly different every time. It is very time-consuming to retrain L-PD for every detector setting and every calibration. Therefore we analyse the robustness of L-PD for system changes in two ways:

1. Firstly, L-PD is trained for a setting of 32 uniformly placed detectors and tested on a setting with less detectors, more detectors and in a limited view setting. In practice, this means that the operator A that was used during training is replaced with a different operator, since every detector setting has its unique operator. Secondly, L-PD is trained on a training set where the detectors were randomly chosen from a possibility of 128 detectors: the algorithm is trained for robustness in the specific choice of detector setting.

2. Synthetic data was obtained by simulating with a specific calibration (c.f. section 3). We train L-PD on synthetic data from this underlying calibration. Next, we test it on data from a different underlying calibration. Differences in the calibration occur due to e.g. a different impulse response at different times or a different size of the absorbers used for the calibration. This means that we look at robustness to changes in preprocessed input data, which is effected by the calibration of the imaging system. We make use of three calibrations, which were obtained from three different experimental phantoms containing a strong absorbing suture in a scattering agar/intralipid mixture (see our previous work[8] for more information).

5   ROBUSTNESS RESULTS

5.1   Robustness to image changes

In the left of Figure 1, one class 3 image of the test set is shown. The classical non-learned reconstruction methods FBP and TV do not give an adequate result, due to the limited number of detectors and noisy data. It can be seen that both L-PD reconstructions are of a higher quality result. However, when trained on the standard class 0, the background is almost completely removed. This means that low-absorbing parts of biological tissue will not be detected if L-PD was not trained for such scenarios. In the right of figure 1 it can be seen that the background is correctly identified when L-PD is trained on class 3, which has more variety in background intensity.
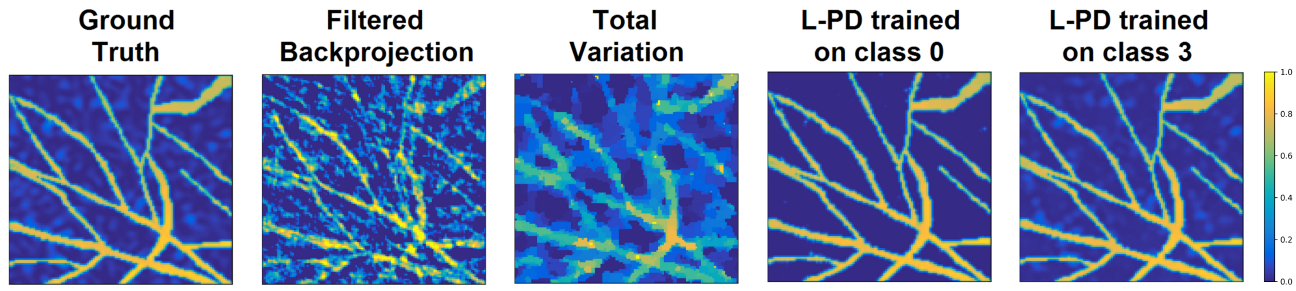

Figure 1: Several reconstructions of an image from class 3: L-PD trained on class 0 already outperforms non-learned methods; for correct background identification, training on class 3 has a positive effect.

The improved performance of L-PD by training on more variety in training images is quantified in Figure 2: it can be seen that the peak signal-to-noise ratio (PSNR) is higher when trained on the same class (class i) as the test class. However, not every variation in the image has the same effect: L-PD trained on the standard class already turns out to be robust to geometrical changes (classes 1 and 4) and different noise levels (class 6); this is not the case for robustness to intensity changes (classes 2,3, and 5). Training on the correct class gives a clear benefit for these classes. Additionally, the effect from the size of the network can be seen in Figure 2: a larger network gives a better image quality, but on average this effect is smaller than that of choosing a training set with the right amount of variety (choosing the right training class).
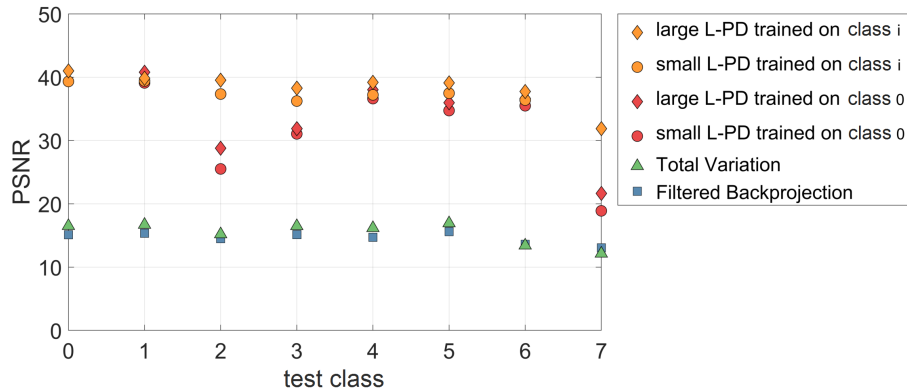

Figure 2: L-PD trained class 0 performs better than non-learned methods. Training on the same class as the test class gives robustness under image changes. The reconstruction quality is slightly improved by taking a larger neural network.

5.2   Robustness to system changes

In Figure 3 the effect of changes in detector setting can be seen on a test image of class 0: while the standard setting of 32 detectors gives almost perfect results, this is not the case for other settings. Still, in the case of less detectors or limited angles (all detectors are located at the top), the L-PD algorithm is able to remove some artefacts and noise. It is not clear why this is not the case when more detectors are used.
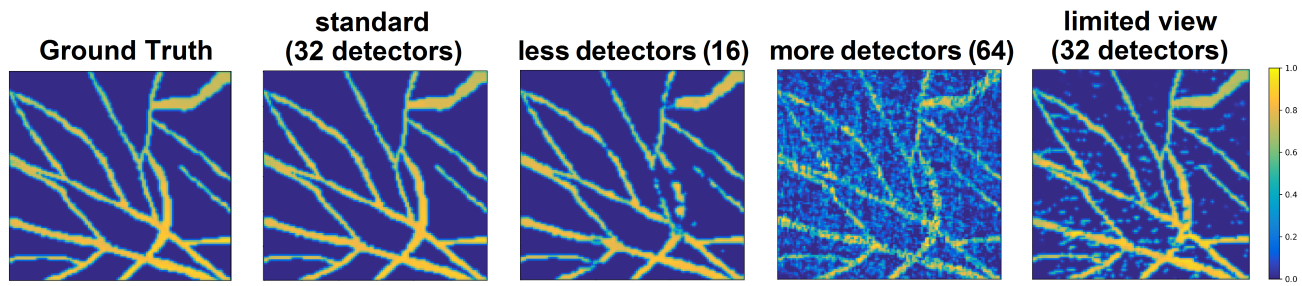
Figure 3: L-PD trained on a fixed system setting (32 detectors): applying it to a test image using different detector settings gives reconstruction artefacts.

In Figure 4 it is shown that by training for a random selection of detectors, we obtain robustness to detector settings. One drawback is that this is still a finite set of detectors, which means that retraining is still needed for full freedom in detector settings.
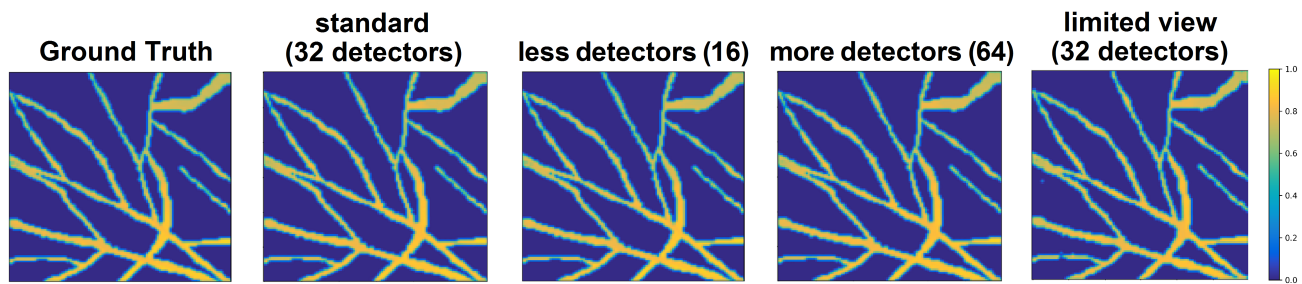

Figure 4: L-PD trained on a variable system setting (1-128 detectors) gives robustness to different detector settings.

As can be seen in Figure 5, the effect of using a different calibration for the test set compared to the training set is very minor: the standard calibration gives an almost perfect reconstruction (c.f. left image of Figure 4 for the ground truth), while the second calibration gives minor reconstruction artefacts in the form of four small high intensity spots, and the third calibration shows a slightly increased overall intensity. The latter can be easily explained by the fact that the absorbed intensity for different calibration phantoms are not exactly the same.
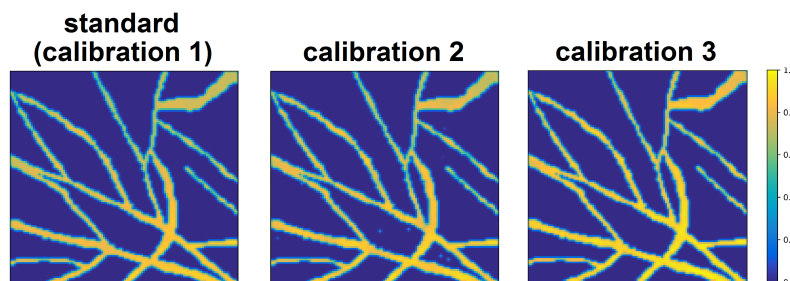

Figure 5: Using preprocessed data with different underlying calibrations has little effect on the reconstruction quality.

6   CONCLUSIONS AND FUTURE WORK

In this work we have investigated the robustness of the learned primal-dual algorithm to variations in the image and in detector settings, as well as uncertainty in the data via the measured calibration. We have shown that:

- reconstruction quality improves if the iterative algorithm has a trainable CNN structure in each iteration: the algorithm is adapted to the type of photoacoustic input data;

- L-PD is robust against variations in the geometry of image structures. To obtain robustness against intensity variations of images, more variety should be put in the training set;

- robustness to different detector settings is not readily obtained: the training set should capture this variety explicitly;

- small changes in the preprocessed data, generated by a uncertainty in the calibration, do not have a big effect on the reconstruction quality of L-PD.

We have only considered photoacoustic images with a homogeneous fluence rate, which are not realistic for many photoacoustic tomography applications. In future work, we will therefore adapt the training set to images generated by using a realistic light propagation model. Furthermore, experiments will be carried out to show the potential of the algorithm in a realistic setting. Finally, we will explore the possibility to perform photoacoustic reconstruction and segmentation simultaneously using a partially learned algorithm.[20] As a basis for this future work, the robustness analysis performed here demonstrates that it is important to have enough variety in image intensities in the training set.

ACKNOWLEDGEMENTS

REFERENCES

[1] Ermilov, S. A., Khamapirad, T., Conjusteau, A., Leonard, M. H., Lacewell, R., Mehta, K., Miller, T., and Oraevsky, A. A., "Laser optoacoustic imaging system for detection of breast cancer," *Journal of Biomedical Optics* **14**, 14 – 14 – 14 (2009).

[2] Heijblom, M., Piras, D., Brinkhuis, M., van Hespen, J. C. G., van den Engh, F. M., van der Schaaf, M., Klaase, J. M., van Leeuwen, T. G., Steenbergen, W., and Manohar, S., "Photoacoustic image patterns of breast carcinoma and comparisons with Magnetic Resonance Imaging and vascular stained histopathology," *Sci. Rep.* **5**(May), 11778 (2015).

[3] Toi, M., Asao, Y., Matsumoto, Y., Sekiguchi, H., Yoshikawa, A., Takada, M., Kataoka, M., Endo, T., Kawaguchi-Sakita, N., Kawashima, M., Fakhrejahani, E., Kanao, S., Yamaga, I., Nakayama, Y., Tokiwa, M., Torii, M., Yagi, T., Sakurai, T., Togashi, K., and Shiina, T., "Visualization of tumor-related blood vessels in human breast by photoacoustic imaging system with a hemispherical detector array," *Sci. Rep.* **7**(December 2016), 41970 (2017).

[4] Chamberland, D., Jiang, Y., and Wang, X., "Optical imaging: new tools for arthritis.," *Integr. Biol.* **2**(10), 496–509 (2010).

[5] Xu, M. and Wang, L. V., "Universal back-projection algorithm for photoacoustic computed tomography," *Phys. Rev. E - Stat. Nonlinear, Soft Matter Phys.* **71**(1), 1–7 (2005).

[6] Paltauf, G., Nuster, R., Haltmeier, M., and Burgholzer, P., "Experimental evaluation of reconstruction algorithms for limited view photoacoustic tomography with line detectors," *Inverse Probl.* **23**(6), S81 (2007).

[7] Dean-Ben, X. L., Buehler, A., Ntziachristos, V., and Razansky, D., "Accurate model-based reconstruction algorithm for three-dimensional optoacoustic tomography," *IEEE Transactions on Medical Imaging* **31**(10), 1922–1928 (2012).

[8] Boink, Y. E., Lagerwerf, M. J., Steenbergen, W., van Gils, S. A., Manohar, S., and Brune, C., "A framework for directional and higher-order reconstruction in photoacoustic tomography," *Physics in Medicine & Biology* **63**(4), 045018 (2018).

[9] Adler, J. and Öktem, O., "Learned Primal-dual Reconstruction," *IEEE Trans. Med. Imaging* **0062**(3), 1–11 (2018).

[10] Chambolle, A. and Pock, T., "A first-order primal-dual algorithm for convex problems with applications to imaging," *J. Math. Imaging Vis.* **40**(1), 120–145 (2011).

[11] Kobler, E., Klatzer, T., Hammernik, K., and Pock, T., "Variational networks: connecting variational methods and deep learning," *German Conference on Pattern Recognition* **10496**, 281–293 (2017).

[12] Meinhardt, T., Moeller, M., Hazirbas, C., and Cremers, D., "Learning Proximal Operators: Using Denoising Networks for Regularizing Inverse Imaging Problems," *Proc. IEEE Int. Conf. Comput. Vis.* **2017-Octob**, 1799–1808 (2017).

[13] Antholzer, S., Schwab, J., and Haltmeier, M., "Deep learning versus ℓ^1 -minimization for compressed sensing photoacoustic tomography," in [*2018 IEEE International Ultrasonics Symposium (IUS)*], 206–212 (Oct 2018).

[14] Hauptmann, A., Lucka, F., Betcke, M., Huynh, N., Adler, J., Cox, B., Beard, P., Ourselin, S., and Arridge, S., "Model-Based Learning for Accelerated, Limited-View 3-D Photoacoustic Tomography," *IEEE Trans. Med. Imaging* **37**(6), 1382–1393 (2018).

[15] Tarvainen, T., Cox, B. T., Kaipio, J. P., and Arridge, S. R., "Reconstructing absorption and scattering distributions in quantitative photoacoustic tomography," *Inverse Probl.* **28**(8), 084009 (2012).

[16] Haltmeier, M., Neumann, L., and Rabanser, S., "Single-stage reconstruction algorithm for quantitative photoacoustic tomography," *Inverse Probl.* **31**, 065005 (jun 2015).

[17] Kruger, R., Liu, P., Fang, Y., and Robert, A., "Photoacoustic ultrasound (PAUS)—Reconstruction tomography," *Med. Phys.* **22**(10), 1605 (1995).

[18] Staal, J., Abràmoff, M. D., Niemeijer, M., Viergever, M. A., and Van Ginneken, B., "Ridge-based vessel segmentation in color images of the retina," *IEEE Trans. Med. Imaging* **23**(4), 501–509 (2004).

[19] Kingma, D. P. and Ba, J., "Adam: A method for stochastic optimization," *CoRR* **abs/1412.6980** (2014).

[20] Boink, Y. E., Manohar, S., and Brune, C., "A partially learned algorithm for joint photoacoustic reconstruction and segmentation," *in preparation* (2019).