

Large-scale Photoacoustic Tomography: A Trade-off Between Image Quality and Computational Cost

Gabriel Bordovský^[0000-0003-2561-6322]¹, Ben Cox^[0000-0001-7296-4093]²,
Felix Lucka^[0000-0002-8763-5177]^{2,3}, and Jiří Jaroš^[0000-0002-0087-8804]¹

¹ Brno University of Technology, Faculty of Information Technology,
Centre of Excellence IT4Innovations,
Bozotechnova 2, 612 00, Brno, Czech Republic
{ibordovsky, jarosjir}@fit.vutbr.cz

² University College London, Medical Physics and Biomedical Engineering,
WC1E 6BT London, United Kingdom
b.cox@ucl.ac.uk

³ Centrum Wiskunde & Informatica,
Science Park 123,1098XG, Amsterdam, The Netherlands
Felix.Lucka@cwi.nl

Abstract. Iterative image reconstruction with accurate acoustic simulation shows promising potential for 3D photoacoustic tomography. To obtain images of high resolution and quality, the employed acoustic simulation method needs to support a wide frequency spectrum and accurate models of both the sensors and the heterogeneous acoustic properties of the tissue. However, in particular simulating a broad frequency spectrum requires considerable computational resources and time and iterative image reconstruction methods run multiple of such simulations sequentially. To reconstruct an image of a human breast at a resolution of 0.2mm, the computations may take from a couple of days to several weeks to finish even on large CPU clusters. We compare the performance and computational cost of an iterative image reconstruction method using different computational grids supporting frequencies up to 0.31, 0.61 and 1.13 MHz corresponding to grid spacings of 0.8, 0.4 and 0.2 mm, respectively. Even though the highest possible resolution is desired, potential advantages and utilization of the coarser grids are discussed, including multi-grid approaches that could significantly speed up the reconstruction process.

Keywords: Photoacoustic tomography · large-scale computing.

1 Introduction

Photoacoustic tomography (PAT) is based on the fact that chromophores in tissue, such as haemoglobin in veins and tumours, absorb light. Since tissue is usually highly optically scattering, imaging to high resolution using purely

optical means is difficult beyond a few optical scattering lengths. However, for acoustic waves, even up to tens of megahertz, the scattering is considerably lower. In PAT, the absorbed energy generates an acoustic wave as it thermalises, and this can be captured by an array of ultrasound sensors surrounding the object. The goal in PAT image reconstruction is to construct an image from this data showing the source of the detected ultrasound (US) waves, which will be proportional to the absorbed optical energy density. Therefore PAT combines the high resolution of ultrasound (hundreds of microns) with the high contrast offered by optical absorption [4]. Several PAT devices are in development, with a variety of geometries [6, 5, 2, 8]. Our work concerns bowl-shaped whole-breast scanners which surround the breast with acoustic sensors, and provide enough data for a 3D image. To construct a scanner with a fine sensor density - high enough for sub-millimetre resolution - without losing sensitivity is technologically difficult. One workaround is to use larger, more sensitive, sensors, and then move the sensor array to capture multiple measurements/views. We can combine the measurements and work with a virtual sensor array with much higher sensor density than we are currently able to construct. This approach brings experimental challenges such as the prevention or compensation of patient movement.

In our setup, a hemispherical scanner with radius of 10 centimeters is considered. This scanner employs 512 evenly distributed US detectors to gather the signal generated by 40 optical fibres outputs providing homogenous illumination by the laser light. Furthermore, the scanner rotates around the vertical axis around the breast and scans at 1000 position to provide denser coverage. The breast will be held in position by a cup which allows all the measurements to be combined and considered as a single measurement by a scanner with 512,000 US detectors. The work described in this paper is not limited to this geometry and will work with different kinds of scanners using similar coverage of the breasts such as ring arrays or cylindrical scanners.

A huge amount of measured photoacoustic data (tens to hundreds of GB) is required to produce a detailed 3D breast image with a resolution under 0.5 mm. The 3D image itself using single precision datatype occupies over 2 GB of memory on 0.2 mm resolution. This represents only one from over 10 matrices of same size used during the computation and renders use of the computational graphics cards cumbersome. There are currently no graphic cards available with sufficient on-board memory to manage such computations.

The image reconstruction used here is a variational, optimization-based, approach, chosen for its generality and flexibility. An k-space pseudo-spectral time domain acoustic wave solver is used as the forward model in the optimization problem, in which the difference between the measured acoustic data recorded by the scanner and the signal predicted from the reconstructed image is minimized [11]. This optimization requires multiple evaluation of this error functional, and therefore many ultrasound simulations. Furthermore, the photoacoustic effect generates a broadband spectrum of frequencies with the fine details of the image carried by high frequencies. Typically, the bandwidth of the generated signals is wider than the US detector bandwidth, so the simulations have

to support the full frequency range of the detectors, ie. the computational grid spacing must be less than half the shortest detectable wavelength. Since the volume of interest is quite large, tens of centimetres, and the shortest wavelengths of the order of hundreds of microns, the reconstruction process is computationally expensive.

In this article, we investigate the trade-off between the quality of the reconstructed image and the computational cost given by the number of core-hours used, using simulated data. Aside from the frequency content determining the minimum spacing of the computational grid, other factors affect the required resources. One is the technological limit of the parallel I/O system used. The computational grid has to cover the volume of $20 \times 20 \times 13$ cm including both the scanner and the chest wall above. We computed the reconstruction with three different grids with spacing between points 0.2 mm, 0.4 mm and 0.8 mm, and analyzed the artefacts in the resulting images.

2 Photoacoustic Tomography

Wherever the optical energy from the laser is absorbed, it is transformed into heat which is accompanied by a corresponding increase in pressure. This pressure, called the *initial acoustic pressure distribution*, p_0 , then propagates in the form of an acoustic wave to an array of sensors which record ultrasound signals, f . This can be written as

$$f = Ap_0 + \epsilon, \quad (1)$$

where A is the forward operator and ϵ is noise. A models the physics of acoustic wave propagation, sampling of the wavefield, and any sensor effects such as frequency and directional response. The aim of photoacoustic image reconstruction is to estimate p_0 given the data f and an approximation of the forward operator $\mathcal{A} \approx A$. Eq. (1) could be rewritten as

$$f = \mathcal{A}p_0 + \hat{\epsilon}, \quad (2)$$

where $\hat{\epsilon} = (A - \mathcal{A})p_0 + \epsilon$ now contains the modelling error as well as the noise. It has been shown [12] that p_0 can be recovered from f in principle so long as the sensors are placed such that the support of p_0 lies in a ‘visible region’. A point is ‘visible’ when rays passing through it in all directions eventually reach sensors. For example, the interior of a closed measurement surface is therefore a visible region, as is the convex hull of a bowl-shaped measurement surface. When sufficient data has been measured, in this sense, and when \mathcal{A} is a good approximation to A , then the inverse problem of finding a good estimate of p_0 from f is well-posed. However, this does not guarantee that there is a simple algorithm for estimating p_0 given f . For instance, linear back-propagation-type algorithms typically assume that the measurements are made on a surface completely surrounding $\text{supp}(p_0)$. Instead, we will formulate the reconstruction as a

least-squares optimization problem

$$p_{\text{image}} \in \underset{p}{\operatorname{argmin}} \{ \|\mathcal{A}p - f\|_2^2 \} \quad (3)$$

One way to solve Eq. (3) for p_{image} , is by iterating a gradient descent scheme

$$p^{(k+1)} = p^{(k)} - \eta \mathcal{A}^*(\mathcal{A}p^{(k)} - f) \quad (4)$$

where \mathcal{A}^* is the adjoint operator to \mathcal{A} and η is a step-size. The step $\eta \in (0, 2/\theta)$ with θ being the largest singular eigenvalue of the normal operator $\mathcal{A}^*\mathcal{A}$ ensures the convergence towards a minimizer (see [3]).

In practice, \mathcal{A} may not be a good approximation of A . For example, the model \mathcal{A} may assume that the sound speed is constant everywhere when it isn't, or that the sensors are omni-directional or broadband when they aren't, or that the measurement surface has a particular geometry when it doesn't quite. In addition, the data measured may be insufficient to determine p_{image} due to undersampling or the limited aperture problem. Then, Eq. (3) does not have a unique solution and/or may be ill-conditioned. In such cases, *regularizing* the problem by supplying additional a-priori information on the solution can help. In *variational regularization*, one amends Eq. (3) by a functional $\mathcal{J}(p)$ that penalizes undesired features of p such as noise-like high frequency components:

$$p_{\text{image}} = \underset{p}{\operatorname{argmin}} \{ \|\mathcal{A}p - f\|_2^2 + \beta \mathcal{J}(p) \}, \quad (5)$$

where β , the regularisation parameter, controls the balance between the emphasis put on the data and on the prior. If $\mathcal{J}(p)$ is differentiable, the regularised problem can be solved by adding its gradient to Eq. (4). In the very common case that it is convex but not differentiable, a more general framework called *proximal gradient descent* has to be used [3]

$$p^{(k+1)} = \operatorname{prox}_{\eta\beta\mathcal{J}} \left\{ p^{(k)} - \eta \mathcal{A}^*(\mathcal{A}p^{(k)} - f) \right\} \quad (6)$$

where the proximal operator $\operatorname{prox}_{\gamma\mathcal{J}}$ solves the sub-problem

$$\operatorname{prox}_{\gamma\mathcal{J}}(y) = \underset{x}{\operatorname{argmin}} \left\{ \gamma \mathcal{J}(x) + \frac{1}{2} \|x - y\|_2^2 \right\}. \quad (7)$$

Here, we will only use the prior information that p_0 cannot be negative which follows from physical arguments. Then, the proximal operator becomes a projection Π_+ onto the non-negative reals, $\mathbb{R}_{\geq 0}^N$, and the iteration becomes a *projected gradient descent*:

$$p^{(k+1)} = \Pi_+ \left\{ p^{(k)} - \eta \mathcal{A}^*(\mathcal{A}p^{(k)} - f) \right\}. \quad (8)$$

An accelerated form of the gradient descent was proposed by Nesterov [9, 3] in which a combination of the last two image approximations are used as the latest estimate:

$$p^{(k+1)} = \Pi_+(y - \eta \mathcal{A}^*(\mathcal{A}y - f)), \quad y = p^{(k)} + \frac{k-1}{k+2}(p^{(k)} - p^{(k-1)}). \quad (9)$$

Eq. (9) can be rewritten as

$$p^{(k+1)} = \Pi_+(y - \eta \mathcal{A}^* \mathcal{A} y + \eta \mathcal{A}^* f), \quad (10)$$

where the last term can be pre-computed and the normal operator $\mathcal{A}^* \mathcal{A}$ can sometimes be implemented more efficiently as one operator rather than as separate calls to \mathcal{A} and \mathcal{A}^* .

Here, the k-Wave toolbox [11] was used to apply both the forward operator \mathcal{A} and its adjoint \mathcal{A}^* . The number of computed timesteps, T , was set to ensure the waves had time to propagate from one corner of the simulation domain to the opposite corner.

The signals corresponding to the grid points belonging to the scanner surface are mapped to the individual sensors with a specific position, size and geometry. These pressure time series were then filtered by the sensor frequency response. For the adjoint evaluation, the signals were re-ordered in time (see Appendix A in [1]), then mapped from the individual sensors to the scanner surfaces and the simulated as an additive pressure source. The pressure field present in the whole domain after T timesteps is then returned.

3 Reconstruction Setup

The reconstruction requirements are determined by several factors, mainly by the scanner geometry and the frequency response of the sensors. In our case, the scanner is a hemispherical US detector array with a radius of 10 cm. Since it is desired to provide imaging close to the chest wall we add additional 3 cm to get dimensions $20 \times 20 \times 13$ cm. This volume is transformed into a computational grid with spacing between grid points fulfilling the Nyquist–Shannon sampling theorem of at least two points per shortest detectable wavelength. This corresponds to the highest frequency which is expected in the tissue and which the sensors are able to measure. In our case, sensors with 0.5 MHz central frequency and 100% bandwidth are chosen, sensitive up to 1.13 MHz. To provide at least two grid points per wavelength in the computational grid representing tissue with fat (which has the slowest sound speed, 1470 m s^{-1}) we would need the distance between grid points to be 0.647 mm. Since two points per wavelength are at the edge of the sampling theorem not all effects of the acoustic wave propagation are accurately represented. Based on an acoustic reconstruction done in [10] we chose 0.2 mm which is slightly over six points per wavelength. To use the pseudo-spectral simulation methods on a non-periodic domain we use the perfectly matched layer constructed by adding a few points on each side of the grid. Using finer resolution than 0.2 mm should not provide any benefits for PAT using such sensors.

We compare reconstructions using this native/exact computational grid with 0.2 mm spacing against grids with spacings of 0.4 mm and 0.8 mm. These grids are not able to cover the full measured frequency range, so the question is how much this impacts the reconstructed pressure field and how much faster are we able to obtain an image. To maintain the same CFL number, each grid size

requires a different number of timesteps to simulate the wave propagation: the finer the spacing between points, the more timesteps required. Table 1 shows the dimensions of each grid in terms of the number of grid points, $N_{\{x,y,z\}}$, the number of timesteps N_t , and maximum frequency supported at three grid points per wavelength, F_{\max} .

Table 1: Computational grid dimensions in number of points and the maximal supported frequency considered with six points per wavelength.

	N_x, N_y	N_z	N_t	F_{\max} [MHz]
0.2 mm	1024	672	5220	1.13
0.4 mm	528	350	2616	0.61
0.8 mm	280	192	1320	0.31

One of the challenging parts of the computation is managing the signals produced by the simulation. The signals are composed of the pressure at the position of the scanner surface in each timestep. The captured signals can be represented as a 2D array with one dimension being timesteps N_t and the second being the number of grid points used to describe the scanner’s hemispherical sensor array. The memory required to store the captured pressure on the sensor array is displayed in the Table 2. The signals are then mapped from grid points

Table 2: Memory requirements for signals produced during forward simulation before mapping to sensors.

	N_t	# of points	Size [GB]
0.2 mm	5220	1 882 836	36.6
0.4 mm	2616	786 395	7.7
0.8 mm	1320	218 487	1.1

to sensors. This mapping is required as the resolution of the reconstruction and simulations is sub-millimetre but the size of the sensor is larger. In our case, we model each sensor by a disc with a radius of 3 mm. We use 512 sensors for 1000 measurements which are then combined into one virtual sensor array with 512,000 overlapping sensors. The size of the mapped signal is displayed in Table 3.

The signal mapping from the grid point space to the sensor space is done as a matrix multiplication with a sparse mapping matrix. The size of the matrix itself is negligible at a couple of megabytes. Due to insufficient support for sensors larger than the grid resolution, the signals produced by the simulation have to be post-processed using sufficient computational resources. The mapping of signals reduces the required memory footprint 3.6 times for 0.2 mm and a native support in the simulation binaries could lead to a reduction of output files and therefore

reduction of I/O operations. On the other hand, for a coarse computational grid, the mapped signal is more than two times larger and such support could have a negative effect.

Table 3: Memory requirements for signals mapped on sensors.

	N_t	# of sensors	Size [GB]
0.2 mm	5220	512 000	10
0.4 mm	2616	512 000	5.0
0.8 mm	1320	512 000	2.5

One of the shortcomings of the acoustic solver used is in its implementation of I/O. Due to the usage of MPI-I/O the size of any dataset is limited by the number of MPI processes used. As the MPI limits the I/O operations to 2 GB per process we can compute how many MPI processes are required for the simulation. The other limitation is the use of the Fast Fourier Transform operating best on the number of processes with low prime factors. The closest higher power of 2 was therefore used for 0.2 mm simulation. The number of processes used for the coarser grids is scaled with its spacing dx .

The simulation setup uses one MPI process per CPU. The supercomputer cluster Salomon was used, equipped with two CPUs per computational node with a total of 24 cores. The number of core-hours billed for each computation hour is shown in the last column of the Table 4.

Table 4: Computational requirements given by I/O subsystem.

dx	size	req. MPI	used MPI	core-hours/hour
0.2 mm	36.6 GB	19	32	768
0.4 mm	7.6 GB	4	16	384
0.8 mm	1.1 GB	1	8	192

4 Measurements

Our reconstruction framework takes measured time series from the scanner as an input. The series are resampled to match the number of samples of the ultrasound computational grid (N_t in the Table 1). The adjoint of the resampled signal $\mathcal{A}^* f$ is pre-computed. The reconstruction follows Eq. (10) and after each forward simulation, \mathcal{A} , the signal produced is compared against the resampled measured signal. The normalized difference of these two signals we call normalized relative

residuum and it is a metric of convergence of the iterative reconstruction. The residuum is smaller if the signal generated by $p^{(k)}$ is closer to the measured one.

We evaluated 25 iterations of Eq. (10) and compared the residuum and the images produced. The workload for each iteration is divided into five parts representing the tasks for the supercomputer. These tasks are classified as either memory intensive, requiring a huge amount of RAM on single node, or computationally intensive using many computational nodes. The jobs doing simulation are computationally intensive and use a parallel solver with the computational requirements described by Table 4. The remaining jobs require more than the 128 GB of memory available on the common computational node of Salomon supercomputer for native resolution 0.2 mm, and have to employ the fat node with 3 TB of memory. The table below contains the average time for each workload.

Table 5: Time required to process each part of the simulation. Total time for all 25 iterations and the cost of the reconstruction in core-hours.

Workload	0.2 mm		0.4 mm		0.8 mm	
Forward preparation	3 min 45 s		1 min 50 s		30 s	
Forward simulation	4 h	36 min	1 h	23 min	8 min	15 s
Adjoint preparation	31 min 5 s		16 min 20 s		3 min	
Adjoint simulation	4 h	15 min	1 h	24 min	8 min	5 s
Post-processing	4 min 15 s		4 min 45 s		3 min 45 s	
Per iteration	9 h	30 min 5 s	3 h	9 min 55 s	21 min	20 s
Per 25 iterations	9.89 d		3.33 d		8.9 h	
Price [core-hours]	182 292,48		30 849,12		1708,8	

Initially only 25 iterations were computed with each grid. The image reconstruction for the native resolution required almost ten days of computational time on dedicated computational resources. In case of shared resources of the supercomputer, the reconstruction time varied greatly depending on the activity of other users. The fat node proved to be the bottleneck of our reconstruction as other users could allocate it fully for up to two days which caused the time per iteration to extend up to a week. This issue could be prevented in the future by separating memory-intensive computations manually so they would fit into the regular computation node.

The convergence of the normalized relative residuum for each grid is displayed in the Figure 1. The result on the computational grid with spacing 0.2 mm in 25 iterations was similar to the result on 0.4 mm so we ran another 5 iterations to emphasize its potential to further provide a better result.

From the figure the 0.4 mm grid converges to residuum around 20 % when 0.8 mm around only 70 %. An interesting observation came from the simulation with native resolution, which achieved the same residuum in 25 iterations as the 0.4 mm grid. The trend of the residuum suggested the potential to progress further towards better match with the original signal. Therefore another 5 iter-




Fig. 1: The normalized residuum of the image reconstruction in iteration of the accelerated gradient descent.

ations were computed to demonstrate the potential. The question of how far is the native grid able to converge the residuum remains, but due to limited time and excessive computational requirements this was not pursued.




Fig. 2: The normalized residuum of the reconstructed image in time corresponding to figure 1 and scaled by the iteration time of each computational grid. The native grid is cut off for clarity.

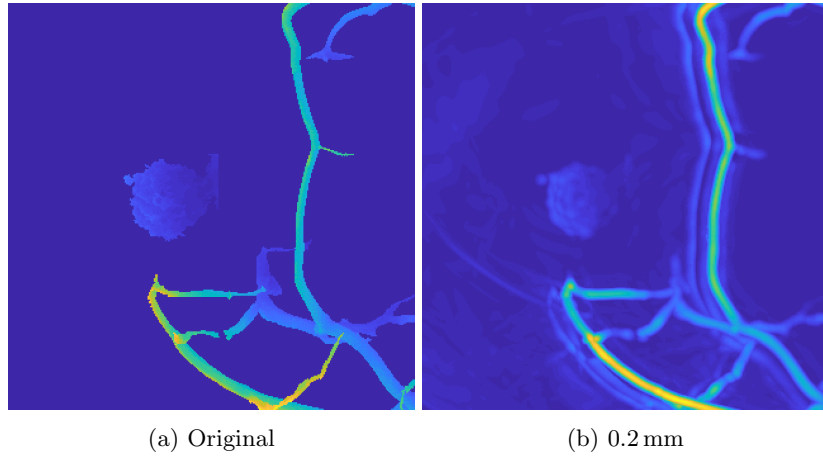


Fig. 3: Frontal maximum intensity projections on the pressure distribution used for generating the acoustic signal and the reconstructed picture using 0.2 mm computational grid.

The convergence of the residuum in time is displayed in Figure 2. (A comparison by the total computational cost would look similar to Figure 2.) The 0.4 mm grid takes 3 times less time to reach 20 % residuum and uses half of the resources. This makes it 6 times cheaper. The grid with the lowest resolution is then 9 times faster against the 0.4 mm grid, but does not progress below 69 % of residuum. It again uses only half of the resources and is therefore 18 times cheaper.

The metric of the residuum in the generated signal does not say much about the image quality itself. Other metrics also do not provide adequate comparison between the true and the reconstructed images, mostly due to the noise and artefacts in the reconstructions. We choose an area of 64×64 mm around the centre of each reconstructed image to demonstrate the results. Figure 3 displays maximum intensity projections of the image of initial pressure distribution when the image was reconstructed on the native grid. The initial pressure field was generated by simulating the optical absorption of a breast phantom derived from MRI scans [7] with an artificially placed tumour.

This image was used to generate the data used in the reconstruction. The loss of higher frequencies causes the image to be more blurred. Even the native resolution is not able to represent a sharp interface between blood veins and the surrounding. This is not directly related to the grid spacing. The frequency response of the used sensors does not cover the high frequencies necessary to generate such sharp interfaces. The other difference noticeable in the image at right is the "halo" around the reconstructed veins. For example, the vein going from the top edge down has aside from the main vein structure another line following it, which is not present in the original pressure field. These halos could be caused by the imperfect geometry of the used scanner, insufficient representation

of sensors in the computational grid or some other yet unknown aspects. This needs further investigation. An image post-processing technique, such as image filters, can be further used to reduce the noise in the reconstructed pressure distribution. These post-processing techniques are expected to have insignificant computational cost against the reconstruction itself and are not discussed here.

Figure 4 compares the images produced on all three computational grids. The top line represents frontal and lateral views. As expected, the image reconstructed on the grid with the more coarse resolution is blurred as a result of neglecting the high-frequency component of the signal. For the 0.4 mm grid, veins are less visible, their contrast against the background is lower (mainly in lateral view) and it is getting harder to distinguish between true position of the vein and the phantom vein. The lateral view also shows more noise in the two-thirds of the picture towards the left edge. This could be caused by the positioning of the scanner which covers the breast from the left side. The tumour is still fairly visible in 0.4 mm reconstruction.

The right-most image raises the question of whether it would still be possible to notice the tumour. As these 2D images are the maximum intensity projections




Fig. 4: Maximum intensity projections of the pressure distribution in frontal (a-d) and lateral (e-h) view. The original pressure distribution used for generating of the measured signal and reconstructed pictures using 25 iterations. The images display an area of 6.4 cm around the centre of the reconstruction containing the tumor phantom and partially vein system. See Appendix-A for full images. The peak-signal-to-noise ratio (PSNR) in the reconstructed pictures is 46.4 dB for grid using 0.2 mm. The original pressure was re-sampled to match the other two reconstructions and compared. The PSNR is 29.6 dB and 11.9 dB for 0.4 mm and 0.8 mm reconstruction grid, respectively.

of the 3D images, the noise from the whole domain reduces the visibility of the tumour contours. On the other hand, the tumour is located in the middle of the domain and the noise is distributed over a larger area. Suitable image filtering and/or threshold could improve its visibility. Each of the used reconstructed images is generated using 25 iterations of the reconstruction framework. For comparison between 25 and 30 iterations on the native resolution and bigger images see appendix A.

5 Conclusion

The required computational cost for PAT image reconstruction of whole breast using an iterative approach is high. We presented a reconstruction using a computational grid able to represent all the frequencies present in the data, and evaluated 25 iterations of an accelerated gradient descent scheme. The image produced sufficiently represents the original pressure distribution used for signal generation. It contains some noise and artefacts but the veins and tumour are visible in the frontal and lateral maximal projection of the 3D image. These impurities could potentially be reduced by the use of an appropriate image filtering. We have also shown results of reconstruction with the use of the computational grids with larger grid spacings. These provided worse images, with the expected blurring and increased levels of background noise. The veins are still visible, but their size and position are not exact. The visibility of the tumour in the coarsest computational grid is low.

These results were related to two cost metrics. The first is the time required to reconstruct the images, which is of interest to the physicians. The native computational grid with point spacing 0.2 mm required 9.89 days of computation time to produce the final image. The coarser computational grid provided the result in only 3.3 days. This image is without doubt a worse representation of the true solution, but important features as tumour and veins are still visible. The crudest computational grid required less than 10 hours to produce the final image. The tumour is no longer easily visible but the placement of the veins is still noticeable. In term of the computational cost, the exact reconstruction requires more than 180 thousand core-hours. The reconstruction using the 0.4 mm grid uses only half of the CPUs for a third of the time and requires 30 thousand core-hours. The worst image was generated consuming 1 708 core-hours which is more than 100 times less than the best image. The reconstruction using 0.4 mm achieved similar relative residuum towards the measured signal as the native reconstruction, but the image does not reflect this and is worse. This could be caused by the initial down-sampling of the input signal to the number of samples matching the computational domain resulting in slightly different references for residuum estimation. This needs further investigation.

Both insufficiently spaced computational grids were used to their maximal potential. The reconstruction with them does not converge fast enough after the 25 iterations to be worth any further computation. The convergence limit of the native resolution was not found in 30 iterations. Due to the difference in the

computational costs, one of our next goals is to investigate the usability of the progressive grid refinement. We would start with coarse grid and use its output for the finer grid progressively moving towards the one with spacing supporting all frequencies. This could reduce the cost of the reconstruction while achieving good results.

Other planned works include reduction of the I/O operation and support of GPUs to accelerate the reconstruction. GPUs have great computational power but are usually equipped with limited memory. This rules them out for usage with full computational grid, covering the whole scanner and supporting all captured frequencies. The reduced images could still be computed reasonably fast and used as seed for the full reconstruction and/or to establish if the costly reconstruction is indeed needed.

6 Acknowledgement

This work was supported by The Ministry of Education, Youth and Sports from the National Programme of Sustainability (NPU II) project “IT4Innovations excellence in science - LQ1602” and by the IT4Innovations infrastructure which is supported from the Large Infrastructures for Research, Experimental Development and Innovations project “IT4Innovations National Supercomputing Center - LM2015070”. This project has received funding from the European Union’s Horizon 2020 research and innovation programme H2020 ICT 2016-2017 under grant agreement No 732411 and is an initiative of the Photonics Public Private Partnership.

References

1. Arridge, S.R., Betcke, M.M., Cox, B.T., Lucka, F., Treeby, B.E.: On the adjoint operator in photoacoustic tomography. *Inverse Problems* **32**(11), 115012 (2016)
2. Asao, Y., Hashizume, Y., Suita, T., Nagae, K.i., Fukutani, K., Sudo, Y., Matsushita, T., Kobayashi, S., Tokiwa, M., Yamaga, I., et al.: Photoacoustic mammography capable of simultaneously acquiring photoacoustic and ultrasound images. *Journal of biomedical optics* **21**(11), 116009 (2016)
3. Burger, M., Sawatzky, A., Steidl, G.: First order algorithms in variational image processing. In: *Splitting Methods in Communication, Imaging, Science, and Engineering*, pp. 345–407. Springer (2016)
4. Cox, B.T., Kara, S., Arridge, S.R., Beard, P.C.: k-space propagation models for acoustically heterogeneous media: Application to biomedical photoacoustics. *The Journal of the Acoustical Society of America* **121**(6), 3453–3464 (2007)
5. Ermilov, S.A., Khamapirad, T., Conjusteau, A., Leonard, M.H., Lacewell, R., Mehta, K., Miller, T., Oraevsky, A.A.: Laser optoacoustic imaging system for detection of breast cancer. *Journal of biomedical optics* **14**(2), 024007 (2009)
6. Lin, L., Hu, P., Shi, J., Appleton, C.M., Maslov, K., Li, L., Zhang, R., Wang, L.V.: Single-breath-hold photoacoustic computed tomography of the breast. *Nature communications* **9**(1), 2352 (2018)

7. Lou, Y., Zhou, W., Matthews, T.P., Appleton, C.M., Anastasio, M.A.: Generation of anatomically realistic numerical phantoms for photoacoustic and ultrasonic breast imaging. *Journal of biomedical optics* **22**(4), 041015 (2017)
8. Manohar, S., Vaartjes, S.E., van Hespén, J.C., Klaase, J.M., van den Engh, F.M., Steenbergen, W., Van Leeuwen, T.G.: Initial results of in vivo non-invasive cancer imaging in the human breast using near-infrared photoacoustics. *Optics express* **15**(19), 12277–12285 (2007)
9. Nesterov, Y.: A method for unconstrained convex minimization problem with the rate of convergence $O(1/k^2)$. In: *Doklady an ussr*. vol. 269, pp. 543–547 (1983)
10. Robertson, J.L., Cox, B.T., Jaros, J., Treeby, B.E.: Accurate simulation of transcranial ultrasound propagation for ultrasonic neuromodulation and stimulation. *The Journal of the Acoustical Society of America* **141**(3), 1726–1738 (2017)
11. Treeby, B.E., Cox, B.T.: k-wave: Matlab toolbox for the simulation and reconstruction of photoacoustic wave fields. *Journal of biomedical optics* **15**(2), 021314 (2010)
12. Xu, Y., Wang, L.V., Ambartsoumian, G., Kuchment, P.: Reconstructions in limited-view thermoacoustic tomography. *Medical physics* **31**(4), 724–733 (2004)

7 Appendix A: Full Images

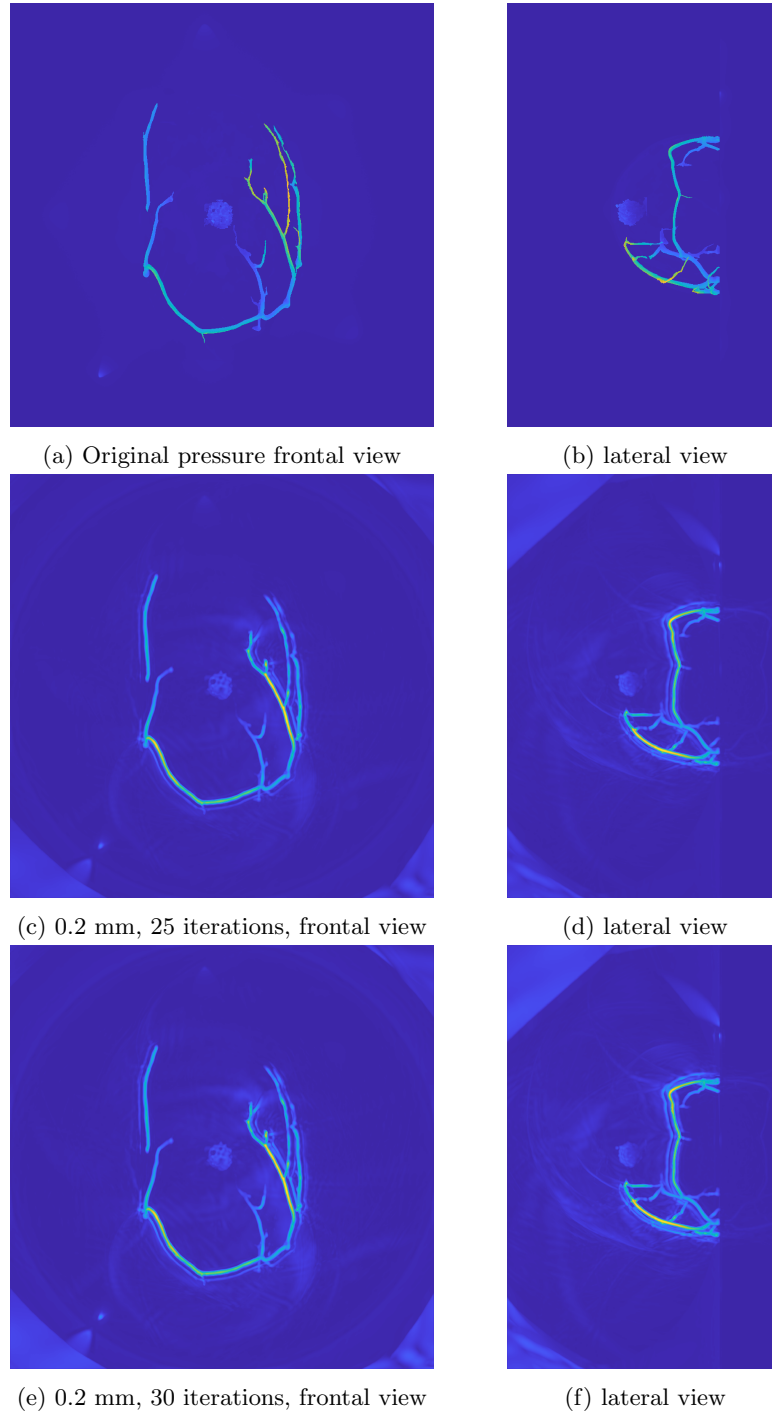


Fig. 5: Maximum intensity projection of the initial pressure distribution used for generating of the measured signals. The maximal acoustic pressure caused by the illumination is 9.3 kPa. After 25 iteration, the maximal pressure reaches up to 4.5 kPa. The image after 30 iteration look similar to the one after 25 iteration, but its maximal value 4.9 kPa is closer to the true solution.

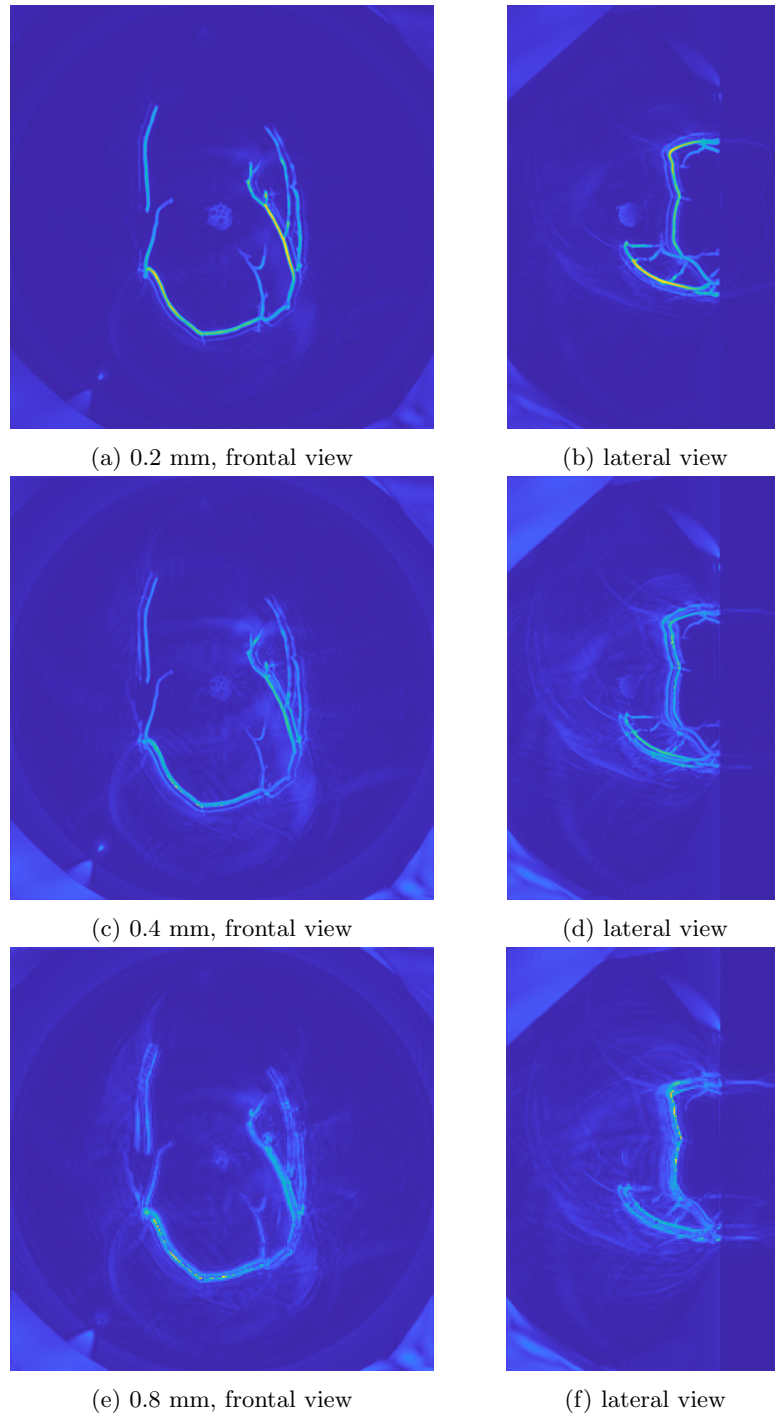


Fig. 6: Maximum intensity projection of reconstruction result after 25 iterations on different computational grids.