

Journal of Biomedical Optics

BiomedicalOptics.SPIEDigitalLibrary.org

Processing methods for photoacoustic Doppler flowmetry with a clinical ultrasound scanner

Thore M. Bücking
Pim J. van den Berg
Stavroula Balabani
Wiendelt Steenbergen
Paul C. Beard
Joanna Bruncker

SPIE.

Thore M. Bücking, Pim J. van den Berg, Stavroula Balabani, Wiendelt Steenbergen, Paul C. Beard, Joanna Bruncker, "Processing methods for photoacoustic Doppler flowmetry with a clinical ultrasound scanner," *J. Biomed. Opt.* **23**(2), 026009 (2018), doi: 10.1117/1.JBO.23.2.026009.

Processing methods for photoacoustic Doppler flowmetry with a clinical ultrasound scanner

Thore M. Bücking,^a Pim J. van den Berg,^b Stavroula Balabani,^c Wiendelt Steenbergen,^b Paul C. Beard,^a and Joanna Brunker^{a,*}

^aUniversity College London, Department of Medical Physics and Biomedical Engineering, London, United Kingdom

^bUniversity of Twente, MIRA Institute for Biomedical Technology and Technical Medicine, Enschede, The Netherlands

^cUniversity College London, Department of Mechanical Engineering, London, United Kingdom

Abstract. Photoacoustic flowmetry (PAF) based on time-domain cross correlation of photoacoustic signals is a promising technique for deep tissue measurement of blood flow velocity. Signal processing has previously been developed for single element transducers. Here, the processing methods for acoustic resolution PAF using a clinical ultrasound transducer array are developed and validated using a 64-element transducer array with a -6 dB detection band of 11 to 17 MHz. Measurements were performed on a flow phantom consisting of a tube ($580\ \mu\text{m}$ inner diameter) perfused with human blood flowing at physiological speeds ranging from 3 to 25 mm/s. The processing pipeline comprised: image reconstruction, filtering, displacement detection, and masking. High-pass filtering and background subtraction were found to be key preprocessing steps to enable accurate flow velocity estimates, which were calculated using a cross-correlation based method. In addition, the regions of interest in the calculated velocity maps were defined using a masking approach based on the amplitude of the cross-correlation functions. These developments enabled blood flow measurements using a transducer array, bringing PAF one step closer to clinical applicability. © The Authors. Published by SPIE under a Creative Commons Attribution 3.0 Unported License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.JBO.23.2.026009](https://doi.org/10.1117/1.JBO.23.2.026009)]

Keywords: flowmetry; blood flow; photoacoustic Doppler effect; cross correlation; transducer array; image processing; masking.
Paper 170584PR received Sep. 4, 2017; accepted for publication Jan. 8, 2018; published online Feb. 27, 2018.

1 Introduction

Many pathologies affect the perfusion of tissues, making knowledge about the blood flow speed a crucial diagnostic aid. Measuring flow speed in the microvasculature in deep tissue (at centimeter scale depths) can be of particular benefit: for example, to monitor the perfusion of tumors, which could be useful to predict or evaluate the efficacy of drugs. Pulsed Doppler ultrasound (PD-US) is often used to image deep tissue blood flow; however, without exogenous contrast agents this modality is typically limited to relatively large vessels with diameters in the range of millimeters and larger. Smaller vessels (submillimeter in diameter) are difficult to image using ultrasound, in part because they exhibit low echogenicity.

Photoacoustic (PA) imaging is a promising candidate that has the potential to overcome the limitations of PD-US, because its contrast is based on optical absorption, offering a much better differentiation of microvessels from the surrounding tissue. There have been a number of advances in applying the PA effect to the measurement of flow.¹ One approach is based on thermal tagging of flow using laser light²⁻⁴ or high-intensity-focused ultrasound.^{5,6} Other methods exploit the Doppler effect in which motion-induced time, phase, or frequency shifts in the PA signal are used to calculate velocity. This was initially implemented using continuous wave excitation with intensity-modulated light.⁷ Analogous to Doppler ultrasound, the received signal contains a Doppler shift of the input modulation

frequency, dependent on the axial component of the velocity at which absorbers are moving with respect to the detector. This approach was later extended by using tone-burst excitation,⁸ enabling spatially resolved measurement of flow speeds. The Doppler effect can also be utilized to detect the lateral flow velocity component by measuring the broadening of the detected bandwidth.⁹ This was implemented in an optical-resolution photoacoustic microscopy (OR-PAM) setup, enabling the measurement of blood flow in the ear of a mouse.¹⁰ Other work has focused on tracking absorbers in the time-domain as opposed to measuring frequency shifts in the PA signal. This was achieved in OR-PAM using autocorrelation¹¹⁻¹³ and cross-correlation methods.¹⁴ OR-PAM-based methods are typically limited to penetrations depths of ~ 1 mm due to the diffusive nature of light transport in tissue.

Acoustic resolution photoacoustic flowmetry (AR-PAF) opens the possibility for deeper imaging depths because ultrasound is only weakly scattered in tissue. AR-PAF measurements have successfully been demonstrated using a time-domain cross-correlation strategy: this technique is based on the measurement of the displacement of discrete absorbers such as red blood cells (RBCs) moving toward or away from a transducer resulting in a time shift of the detected PA signal (Fig. 1). By measuring the change in the time of arrival of the PA signal (Δt) using two successive laser pulses, the displacement of the absorber toward (or away from) the transducer can be calculated.

Previous work has established the accuracy and scalability of AR-PAF using single element transducers.¹⁵ The time delay between laser pulses can be adjusted to ensure that absorbers

*Address all correspondence to: Joanna Brunker, E-mail: joanna.brunker.09@ucl.ac.uk




Fig. 1 Schematic of AR-PAF illustrated using a single element transducer. If an absorber (red circle, color online) moves toward the transducer in between laser pulses separated by some time T , the PA pulse will arrive Δt earlier at the transducer. By measuring Δt , and the angle between the transducer's receiving direction and the flow direction, the velocity can be estimated. The acoustic resolution mode employs diffuse, unfocused illumination with the spatial resolution defined using a focused transducer.

remain within the detector field-of-view, and this makes the technique sensitive to a wide range of flow velocities (mm/s to m/s). It was found that accurate velocity measurements become more challenging with increasing absorber density, and this difficulty was attributed to reduced absorber heterogeneity.¹⁶ Despite the challenge of low absorber heterogeneity, the first successful AR-PAF measurements of whole blood were recently reported using a transducer with a 30-MHz center frequency.¹⁷

Implementing the technique with a transducer array rather than a single element transducer is essential for advancing the technique closer to clinical applicability, because it enables imaging of flow. This has previously been demonstrated using a handheld transducer array to image a low density suspension of microspheres.¹⁸ In this work, we further develop and analyze the processing methods required for two-dimensional (2-D) AR-PAF measurements in order to facilitate automated and robust estimation of flow velocities. These methods are demonstrated using a clinical transducer array with a center frequency of 15 MHz to measure the velocity of whole human blood.

The experimental setup and acquisition of the PA images are described in Secs. 2 and 3, respectively. Section 4 describes the processing steps employed to reduce the susceptibility of the measurements to noise and outliers. These steps include filtering, displacement detection, and masking. The results are presented and discussed in Sec. 5, and Sec. 6 concludes with an outlook on the future of AR-PAF.

2 Experimental Setup

In order to mimic microscale blood flow, polyethylene tubing (580 μm inner diameter) was immersed vertically in a water bath. Experiments were performed using healthy human donor blood with 0.5 M ethylenediaminetetraacetic acid added to prevent coagulation.

The tube was illuminated with a dual-cavity Q-switched Nd:YAG Litron Nano L PIV laser, generating pulse pairs at a repetition rate of 15 Hz. The laser pulses in each pair were separated by a time $T = 1$ ms and had a pulse duration of 5 ns and a pulse energy of 150 mJ. The measurements of whole blood were performed at 1064 nm, as the comparatively




Fig. 2 Experimental setup for blood flow velocity measurements using AR-PAF. A polyethylene tube of 580 μm inner diameter serves as a vessel phantom. It is illuminated with successive laser pulses to generate PA signals. Blood flow through the tube is controlled using a syringe pump. The PA signals are detected with a 64 element ultrasound transducer array (center frequency = 15 MHz). The array is positioned such that a cross section of the tube is imaged, and the orientation at an angle θ to the tube enables motion of the cells to be observed as a displacement in the axial direction of the transducer. The phantom and the transducer were held in place using custom-made holders.

low absorption coefficient at this wavelength enabled a more homogeneous illumination of the tube than at 532 nm.¹⁷ The beam diameter at the tube was ~ 5 mm in diameter, which is large in comparison to the resolution of the reconstructed images (see Sec. 3), and thus imaging was performed in the acoustic resolution regime.

The PA waves were detected with a 128 element ultrasound probe (ESAOTE SL3116), which has a pitch of 100 μm and a bandwidth ranging from 11 to 17 MHz at the -6 dB points. The probe has 64 analog-to-digital converters (ADCs), so it recorded 64 time series per acquisition rather than 128. A single laser pulse triggered all ADCs to acquire data with a sampling frequency of 50 MHz, and these data were used to reconstruct an image. The probe was positioned at an angle of $\theta \approx 60$ deg to the tube, such that a cross section was imaged, as shown in Fig. 2. The pairs of laser pulses generated corresponding pairs of images, with the images in each pair being separated by a time T . Measurement of the displacement of absorbers between the two images allowed calculation of the flow velocity.

The flow rate was controlled using a syringe pump (Cole-Parmer WZ-74900-15), providing flow speeds in the range of 3 to 25 mm/s. For each speed setting, 200 image pairs were acquired, after which the syringe was axially rotated by 180 deg and high pressure was applied to flush the tube for a short amount of time, to prevent sedimentation inside the syringe and the tube. After flushing and setting the next flow speed, the flow was allowed to stabilize for 1 min before acquiring the next data set.

3 Image Acquisition

The acquired radio frequency (RF) data were used to reconstruct the PA images utilizing the Fourier transformation reconstruction method¹⁹ implemented using the k -wave MATLAB toolbox.²⁰ As part of the reconstruction, fourfold upsampling in the lateral direction was performed, resulting in 256 axial pressure profiles per image.

An acquired image pair consists of two discretized representations of some initial pressure map denoted as $\mathbf{p}_1(i, j)$ and $\mathbf{p}_2(i, j)$ for the first and second images, respectively. The value at point (i, j) corresponds to a pressure located at $(i\delta_x, j\delta_y)$, where $\delta_x = 25 \mu\text{m}$ and $\delta_y = 29.6 \mu\text{m}$ are the spatial extents of a pixel in the x (lateral) and y (axial) directions. A total of 200 image pairs were acquired. Therefore, in the following section, the two image stacks $\mathbf{p}_1(i, j, k)$ and $\mathbf{p}_2(i, j, k)$ will be considered, where $i \in \{1, 2, \dots, 256\}$, $j \in \{1, 2, \dots, L\}$, with $L = 477$ being the number of points in a line scan (corresponding to the number of RF data points acquired by each transducer element before image reconstruction), and $k \in \{1, 2, \dots, N\}$, with $N = 200$ being the number of image pairs. The sampling frequency of each PA signal was 50 MHz and the acquisition frequency of each image pair was 15 Hz, as determined by the laser pulse repetition frequency. The time delay between acquisition of the first and second frames of each pair was freely controllable and set to 1 ms to allow AR-PAF measurements at the physiological flow speeds used in this study.

4 Image Processing

4.1 Filtering

Two image filters were implemented in order to improve the signal-to-noise ratio (SNR): a high-pass filter and a background subtraction.

A high-pass filter removes the low frequency components of the detected signal and was applied in the axial direction. The transducer has a detection band of 11 to 17 MHz, corresponding to wavenumbers down to 73 cm^{-1} in a reconstructed image. The high-pass filter was implemented using a zero-phase infinite impulse response filter, with a stopband of 7 cm^{-1} and passband of 34 cm^{-1} (corresponding to sampling frequencies of 1 MHz and 5 MHz, respectively). This ensured that the filter only affected noise. In the future, the high-pass filter will be applied before image reconstruction.

Background subtraction is a crucial preprocessing step, analogous to clutter filtering, which is commonly implemented in Doppler US to remove the signal from the stationary and slow moving vessel walls, which can be up to 60 dB higher in amplitude than blood. Here, a simple background subtraction

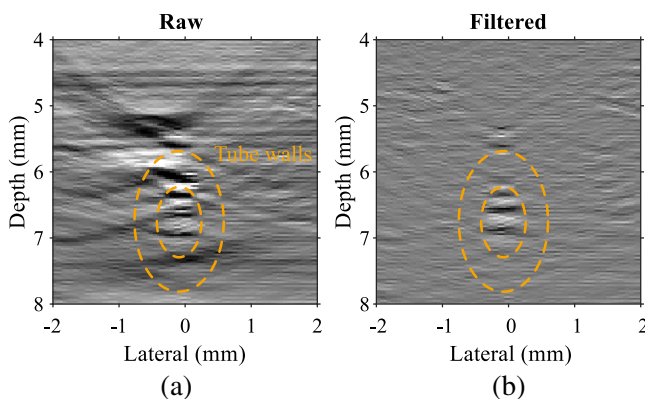


Fig. 3 Example of a reconstructed image of the cross section of a tube filled with whole blood (a) without filtering and (b) after background subtraction and high-pass filtering. The images are displayed on a linear scale (arb. units) with increased contrast for clarity. The cross section of the tube is indicated with dashed lines.

approach was implemented to reject only the stationary parts of the signal (such as those originating from the tube walls), which is sufficient in an *in vitro* setting. This is achieved by taking the average of the corresponding pixels in the image stack and subtracting that value from each pixel:

$$\text{BS}[\mathbf{p}_l(i, j, k)] := \mathbf{p}_l(i, j, k) - \frac{1}{N} \sum_{k=1}^N \mathbf{p}_l(i, j, k), \quad (1)$$

where BS is the background subtraction functional and the index $l \in \{1, 2\}$ denotes the first or second image of an image pair. In other words, background subtraction was applied separately to the data sets for $l = 1$ and $l = 2$. This was required due to different laser pulse energies of the first and second lasers in the dual-cavity system. An example of a reconstructed image before and after the filtering steps is shown in Fig. 3.

4.2 Displacement Detection

Assuming laminar flow within the tube (the Reynolds number is of the order of 10), motion can be considered to be solely along the axial direction in the images. To obtain spatially resolved velocity maps, small subsets (interrogation windows, IWs) of the vector $\mathbf{p}_l(i, j, k)$ are considered. These IWs are $1 \times W$ pixel in size (lateral \times axial) and their position with respect to $\mathbf{p}_l(i, j, k)$ is defined by a vector $\mathbf{r} = (i_0, j_0, k_0)$, with i_0 , j_0 , and k_0 being integers. So, one window is defined as

$$\mathbf{w}_{l, \mathbf{r}=(i_0, j_0, k_0)}(j - j_0) = \mathbf{p}_l(i_0, j, k_0) : j \in \{j_0 + 1, j_0 + 2, \dots, j_0 + W\}, \quad (2)$$

with W being the length of the IW in pixels (chosen to be 32) and $l \in \{1, 2\}$. IWs are constructed for all \mathbf{r} , i.e., for every possible IW location in $\mathbf{p}_l(i, j, k)$. This means that IWs are overlapping in the axial direction by $W - 1 = 31$ pixels (97%). One such IW is shown in Fig. 4.

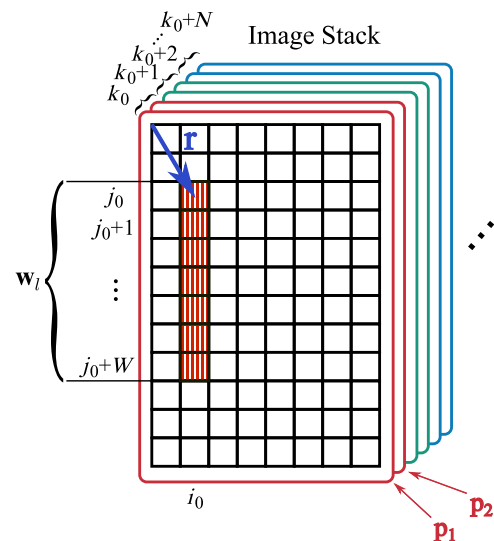


Fig. 4 Illustration of IW indexing. A window \mathbf{w}_l is a set of data points from \mathbf{p}_l with $l \in \{1, 2\}$ (red lines, color online). The location of the IW with respect to the image is defined by \mathbf{r} , which marks the window's starting point (blue arrow, color online). While actual processing was done with $W = 32$, a smaller window with a size of 7 pixels is shown here for illustrative purposes.




Fig. 5 Illustration of using IWs to find spatially resolved velocity maps using cross-correlations: in an image pair, corresponding IWs are selected from the two pressure profiles \mathbf{p}_1 and \mathbf{p}_2 (red vertical lines, color online). Cross-correlating those two signatures \mathbf{w}_1 and \mathbf{w}_2 yields $\mathbf{R}_{1,2}$ which is assigned to the location corresponding to the center of the IWs. This process is repeated for a different IW (blue horizontal lines, color online), yielding a different $\mathbf{R}_{1,2}$ corresponding to a different location. This is repeated for every pixel, resulting in a cross-correlation map with a function \mathbf{R} for every pixel location \mathbf{r} . While actual processing was done with $W = 32$, a smaller window with a size of 7 pixels is shown here for illustrative purposes.

Absorber displacement leads to changes in signal intensity distributions from one image to the next. To calculate the shift between the signals in any two corresponding IWs of an image pair, the discrete unbiased cross-correlation is employed, which is defined as

$$\mathbf{R}_{\mathbf{r}}(n) = \begin{cases} \frac{1}{W-n} \sum_{j=1}^{W-n-1} \mathbf{w}_{1,\mathbf{r}}(j+n) \mathbf{w}_{2,\mathbf{r}}^*(j) & n \geq 0 \\ \mathbf{R}_{\mathbf{r}}^*(-n) & n < 0 \end{cases}, \quad (3)$$

where \mathbf{r} defines the location of $\mathbf{w}_l(j)$ with respect to $\mathbf{p}_l(i, j, k)$, * denotes the complex conjugate, and $n \in \mathbb{N}$. In the present study, the complex conjugate can be ignored, because the acquired data have no imaginary component.

An illustration of the cross-correlation procedure described above is shown in Fig. 5. After an image pair is acquired, corresponding IWs are defined and cross-correlated, yielding $\mathbf{R}_{\mathbf{r}}$ for the selected signal pair (red vertical lines in the illustration). This process is repeated for every pixel as illustrated by a second example marked with blue horizontal lines in Fig. 5. This results in a map of cross-correlation functions $\mathbf{R}_{\mathbf{r}}$ corresponding to every possible pixel \mathbf{r} .

There is a trade-off between selecting small and large IWs. If the IW is too small, the cross-correlation is more susceptible to noise. On the other hand, small windows average information over a smaller region of interest (ROI) and hence provide a higher spatial resolution in that sense. Hence, two conditions need to be fulfilled for accurate displacement measurements: first, the IW has to be large enough to capture signal intensity variations within the IW; second, the IW has to be small enough to ensure that the velocity is spatially invariant within the IW. The first condition can only be true if the detected frequency content is high enough to allow detection of signal intensity variations within the given IW length.

The process outlined above is repeated $\forall (i, j, k)$, i.e., for all locations in each image pair and for all image pairs, resulting in a stack of 200 cross-correlation maps. To improve SNR, ensemble averaging was employed:

$$\bar{\mathbf{R}}_{\mathbf{r}=(i,j)}(n) := \frac{1}{N} \sum_k^N \mathbf{R}_{\mathbf{r}=(i,j,k)}(n). \quad (4)$$

It can be seen from the above equation that this reduces the dimensionality as the k dimension is removed. Hence, $\bar{\mathbf{R}}_{\mathbf{r}=(i,j)}(n) \forall \mathbf{r}$ represents a cross-correlation map, with one cross-correlation function corresponding to each pixel in the original image. For each point in the stack specified by i and j , the displacement Δy can be estimated using

$$\arg \max_n [R_{\mathbf{r}}(n)] = \frac{\Delta y}{\delta_y}, \quad (5)$$

where $R_{\mathbf{r}}(n): n \in \mathbb{R}$ is the cubic spline interpolation of $\mathbf{R}_{\mathbf{r}}(n): n \in \mathbb{N}$, and δ_y is the pixel length. Using Eq. (5), a map of local displacement estimates is constructed: Δy is calculated for every \mathbf{r} (i.e., every pixel) yielding a map of displacements between $t=0$ and $t=T$. The displacement map is converted into a velocity map using

$$v = \frac{\Delta y}{T \cos(\theta)}, \quad (6)$$

where θ is the angle between tube and transducer (see Fig. 2). The velocity map obtained using the method described above does not discriminate between regions of high and low SNR. This is problematic because regions of low SNR make the cross-correlation susceptible to outliers. Hence, a masking strategy based on the PA amplitude was utilized as described in the following section.

4.3 Masking

Masking is based on the assumption that regions of high PA amplitude are of interest. In this work, a masking strategy based on the amplitude of the cross-correlation function was developed. This is computationally efficient, because the

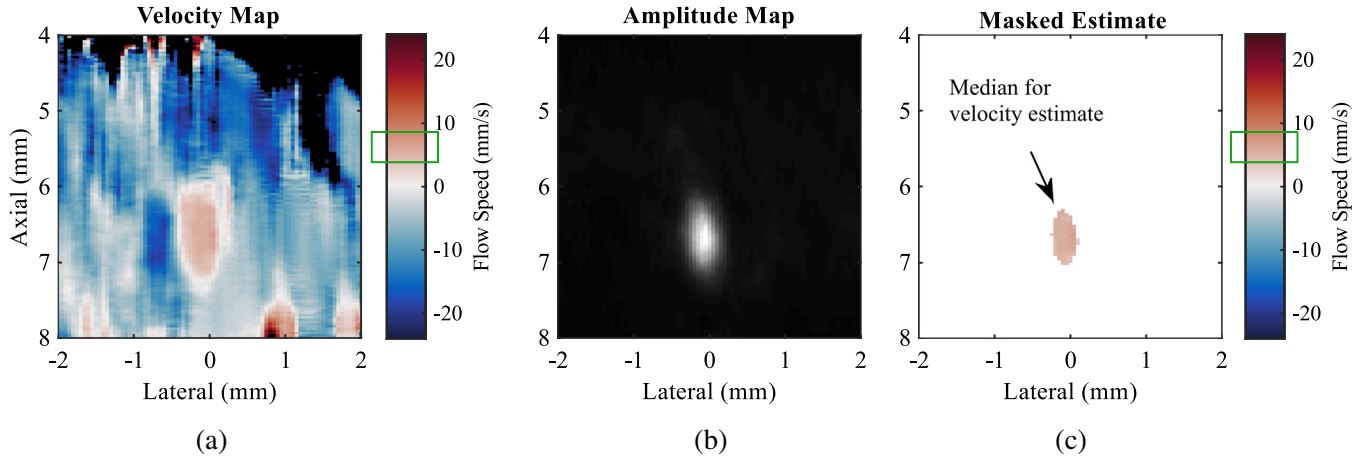


Fig. 6 Illustration of the masking applied to PA velocity maps. The velocity map (a) shows the correct flow region between a depth of 6 and 7 mm in the central region of the image (light red, highlighted with the green box in the color bar, color online). Velocities are also calculated outside the ROI, which are dominated by noise and can take random values. The amplitude map (b) corresponds well with the cross section of the tube and is thresholded to mask the velocity map (c). The median of the resulting masked estimate is taken as the final velocity estimate.

cross-correlation functions are calculated for measurement of the displacement shift, regardless of the masking scheme. Two approaches to implement this masking strategy were investigated.

The first masking approach is based on using the maximum amplitude $\bar{M}_{\mathbf{r}=(i,j)}$ from the ensemble correlation

$$\bar{M}_{\mathbf{r}=(i,j)} := \max_n [\bar{R}_{\mathbf{r}}(n)], \quad (7)$$

where $\bar{R}_{\mathbf{r}}(n): n \in \mathbb{R}$ is the cubic spline interpolation of $\tilde{R}_{\mathbf{r}}(n): n \in \mathbb{N}$.

A second masking strategy was developed employing the median rather than the mean. Specifically, the median $\tilde{M}_{\mathbf{r}=(i,j)}$ of the maxima of the cross-correlation functions was calculated:

$$\tilde{M}_{\mathbf{r}=(i,j)} := \text{med}_k [\max_n [R_{\mathbf{r}}(n)]]. \quad (8)$$

The calculations in Eqs. (7) and (8) can be performed for all \mathbf{r} , creating a 2-D mask based on the cross-correlation amplitudes. Hence, Eqs. (7) and (8) can be rewritten as

$$\bar{\mathbf{M}}^*(i,j) = \bar{M}_{\mathbf{r}=(i,j)} \quad \text{and} \quad (9)$$

$$\tilde{\mathbf{M}}^*(i,j) = \tilde{M}_{\mathbf{r}=(i,j)}, \quad (10)$$

where the \star superscript has been used to signify that the array has been calculated from cross-correlation functions. This distinction is important, because a point (i,j) in Eqs. (9) and (10) does not correspond to a point (i,j) in $\mathbf{p}(i,j)$ but rather contains information from an interrogation region located at this point [as can be seen in Eq. (3)].

This is further illustrated in Fig. 5: two overlapping IWs are shown, with their location in the image indicated by vertical red and horizontal blue lines seen on the left. The overlapping region is indicated by pixels with red and blue lines crossing each other (color online). The corresponding signals \mathbf{w}_1 and \mathbf{w}_2 are displayed in the center of the image: the signals from the red IW at the top and the signals from the blue IW at the

bottom. In this example, the largest amplitude of the signal is located in the overlapping region, so both red and blue IWs yield identical cross-correlation peaks, even though the centers of the two IWs are offset in the axial direction. Therefore, the peak cross-correlation amplitude is almost identical in these two cases despite differences in the underlying pressure profiles.

To better represent the underlying pressure profiles, $\bar{\mathbf{M}}^*(i,j)$ and $\tilde{\mathbf{M}}^*(i,j)$ were deconvolved with a rectangular function defined as

$$\Pi_W(j) = \begin{cases} 0 & \text{if } |j| > W/2, \\ 1 & \text{if } |j| \leq W/2, \end{cases} \quad (11)$$

where W is the length of the IWs. The deconvolved $\bar{\mathbf{M}}^*(i,j)$ and $\tilde{\mathbf{M}}^*(i,j)$, denoted as $\bar{\mathbf{M}}(i,j)$ and $\tilde{\mathbf{M}}(i,j)$, will subsequently be referred to as “ensemble masking” and “median masking,” respectively. After deconvolution, the amplitude maps were converted into binary masks by thresholding them at half their maximum value. A single velocity estimate was extracted from each masked velocity map by taking the median of the unmasked pixels. The median was used rather than the mean in order to reduce susceptibility to outliers. The standard deviation of these pixels was used for calculation of the error bars. The masking procedure is shown in Fig. 6. An overview of the entire proposed pipeline can be seen in Fig. 7.

5 Results and Discussion

The effects of the filtering and masking methods described in Sec. 4 were analyzed by investigating their individual impacts on the reconstructed images. When testing different filtering methods, the results were calculated after applying the median mask. For assessment of the masking methods, both the high-pass filter and the background subtraction were implemented. Velocity estimates are based on the median of the masked region and the associated error bars are calculated using the standard deviation of those values. Hence, the error bars represent the variance within a single flow map and are not indicative of fluctuations between measurements.

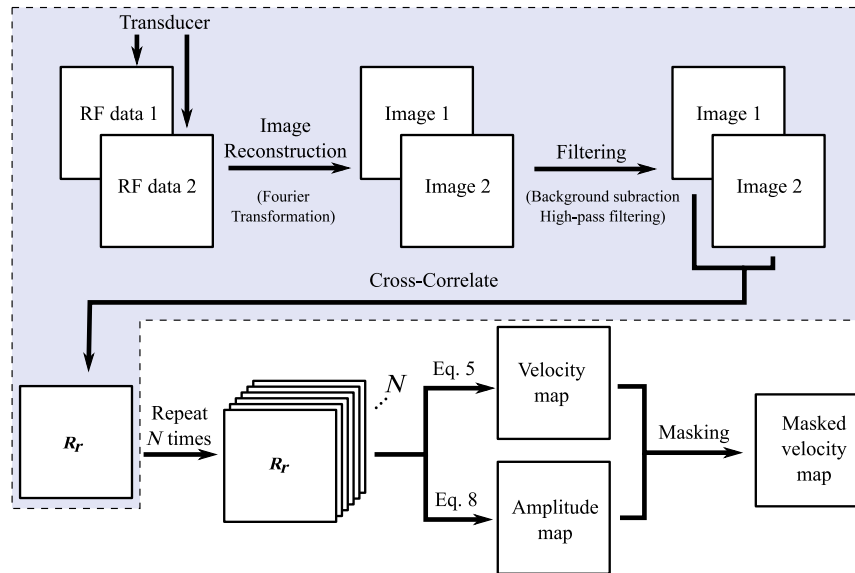


Fig. 7 Processing pipeline for extracting flow information from RF data using median masking. The processing steps in the shaded region are repeated N times to yield a stack of cross-correlation maps $\mathbf{R}_r=(i,j,k)$.

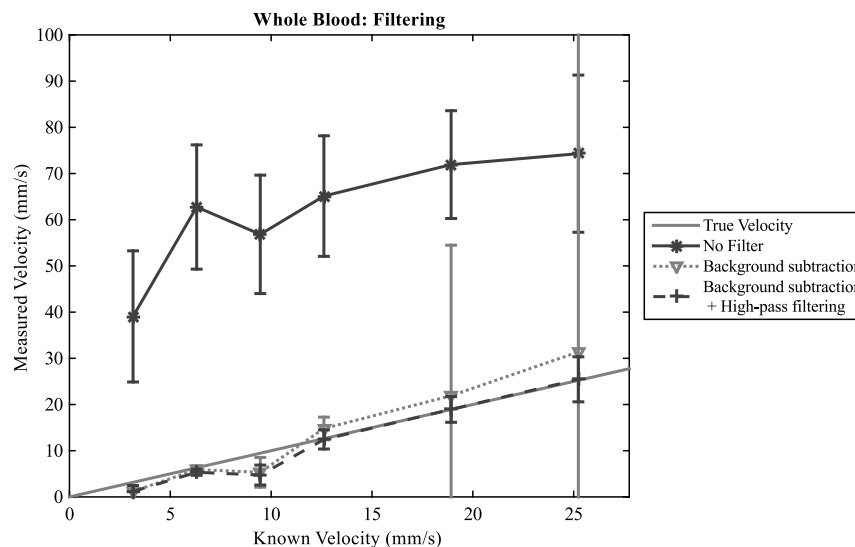


Fig. 8 AR-PAF measurements of blood flow *in vitro*, with and without image filtering. Without filtering, flow measurements were not accurate (*, solid line). Implementing background subtraction enabled good estimates of the flow speed, but outliers were present, resulting in large error bars (∇ , dotted line). Using both background subtraction and high-pass filtering enabled accurate blood flow measurements (+, dashed line).

Figure 8 shows AR-PAF measurements of blood flow *in vitro*, with and without image filtering. From this figure, it is evident that background subtraction was an essential tool, since without the filter (*, solid line), the displacement was consistently overestimated. This might appear counterintuitive, as a larger static component in the signal would suggest an underestimation of displacement. However, the systematic offset was recreated using computer simulations and was found to be due to an overcorrection in the unbiased cross-correlation method.

Figure 8 also shows that combining background subtraction with a high-pass filter provides the most effective way of correctly measuring displacements using cross correlations. Implementing background subtraction (∇ , dotted line) removes

the large DC component, making it more accurate, but some pixels in the masked region are outliers (as indicated by the large error bars at 19 and 25 mm/s). These outliers are a result of low SNR leading to local cross-correlation maxima due to noise, and therefore located anywhere on $\mathbf{R}_r(n)$ (i.e., any n). By implementing the high-pass filter, the cross-correlation function becomes more robust due to the reduction in noise, enabling the correct maxima in the ensemble correlation functions to be identified (+, dashed line).

The effect of different masking approaches was also investigated using the optimal filtering described above. Both background subtraction and high-pass filtering were employed. After calculating the flow map, the ROI was selected using three

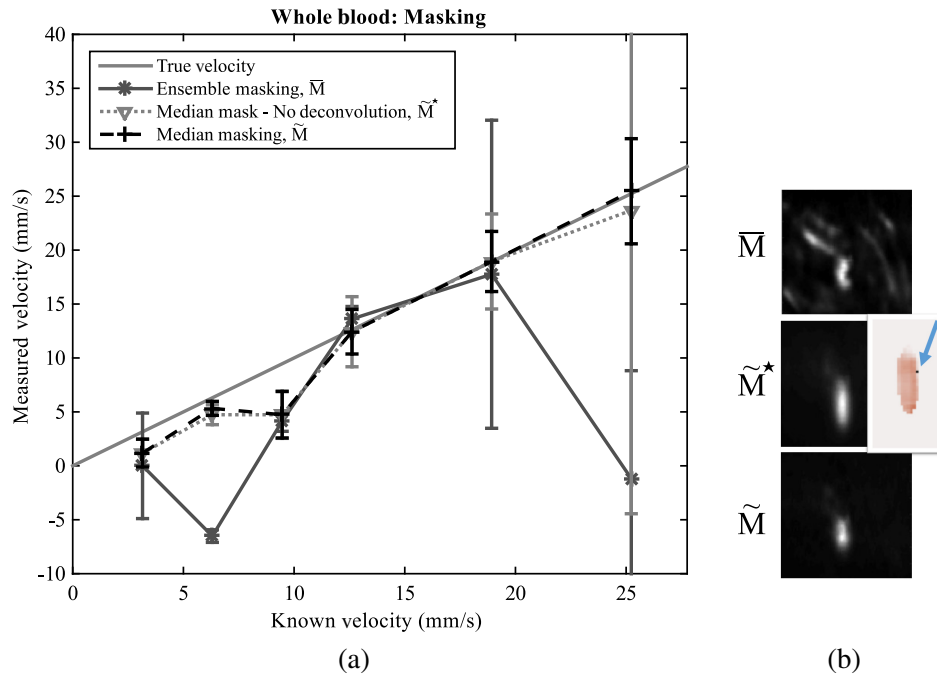


Fig. 9 (a) AR-PAF measurements of blood flow *in vitro* using different masking approaches. Ensemble masking (\bar{M}) was not successful (*, solid line). Median masking without the deconvolution step (\tilde{M}^*) succeeded in estimating the correct flow speed, but has a large error bar at the 25 mm/s data point (∇ , dotted line). Median masking with deconvolution (\bar{M}) also found the correct flow estimates, but with smaller error bars (+, dashed line). (b) Illustration of amplitude maps of \bar{M} , \tilde{M}^* , and \bar{M} . For \tilde{M}^* , the masked flow map is displayed for the 25 mm/s flow speed measurement. One outlier is in the ROI (black point, highlighted by blue arrow).

different masking strategies: ensemble masking, median masking, and median masking without deconvolution (i.e., a mask based on \tilde{M}^*).

Figure 9 shows the effect of different masking approaches on the *in vitro* blood flow measurements. It can be seen that ensemble masking exhibited large errors and underestimation (*, solid line). This is due to the higher sensitivity to outliers: ensemble masking can still find spurious ROIs if taking the mean of all the cross-correlation functions is not sufficient to eliminate outliers. This could be the case if one or several images in the image stack have a region with high transient signal not due to flow. In practice, this can happen if a strong absorber enters the field-of-view briefly during the data acquisition period. Under these conditions, the ROI is not restricted to the flow region within the imaged phantom but can be erroneously assigned to areas without flow.

Using median masking without deconvolution is less susceptible to outliers and correctly identifies the ROI (∇ , dotted line in Fig. 9). However, the ROI is stretched out along the axial direction, including pixels in the neighborhood of the ROI. This can lead to the undesirable inclusion of outliers. For example, at 25 mm/s a large error bar can be seen in the measured flow speed, because in that particular case, there happened to be one outlier near the ROI, which was included due to the stretched amplitude map. This problem is alleviated by using median masking with deconvolution (+, dashed line in Fig. 9), which produces an amplitude map that better represents the shape of the tube. Thus, the accuracy of the measurements is similar to that for median masking without deconvolution, but the error bars are smaller.

The results in Fig. 9 (+, dashed line) demonstrate the value of the processing methodology incorporating background subtraction, high-pass filtering, and median masking with deconvolution. Measuring displacements using whole blood is nontrivial because of the high concentration (about 45 vol. %) and small size (about 8 μm diameter) of RBCs, and yet, the processing methodology is robust enough to outliers to enable accurate blood flow measurements.

6 Conclusions and Outlook

This paper demonstrates the feasibility of measuring the flow velocity of whole human blood using AR-PAF. The flow was measured inside a polyethylene tube of 580 μm inner diameter using a 64 element transducer array (center frequency of 15 MHz). These are the first PA measurements of blood flow using a transducer array, and they represent an important step toward the clinical translation of AR-PAF.

A number of processing steps were developed in order to enable these challenging measurements. The key steps are image filtering, motion detection, and masking. Background subtraction and high-pass filtering were found to be critical steps for accurate motion detection. Flow speed was quantified using a windowed cross-correlation method.

Masking removes image regions that do not correspond to flow. It was found that a mask based on the amplitude of the calculated cross-correlation functions provides a more robust means of defining the ROI than one based on the intensity of the original image. Two different ways of constructing this mask were investigated: one based on taking the mean of the cross-correlation functions (ensemble masking) and one

based on taking the median of their respective maxima (median masking). High intensity pixels appeared outside the ROI on some occasions. These outliers hindered correct identification of the ROI using the ensemble masking method, whereas the median masking was more robust to such noise. Lastly, a deconvolution step was implemented to ensure improved correspondence between the shape of the ROI and the imaged cross section of the phantom.

In order to bring AR-PAF closer to *in vivo* applications, further challenges need to be overcome. First, the SNR of the signal is likely to be lower in tissue due to light scattering. Initial *in silico* studies indicate that accurate velocity measurements are still possible after reducing the SNR by a factor of two; however, it remains to be seen if these promising results are indeed indicative of robust *in vivo* performance. Second, during image acquisition, the flow is assumed to be constant and the vessels are assumed to be stationary. Acquiring 200 image frames took ~ 13 s, which is a long time in comparison to motion artifacts due to breathing, pulsatile flow, and muscle movement. A time-gating approach can be used to partially correct for these artifacts, provided that breathing and heart rate are monitored. A more straightforward solution would be to use a dual-cavity laser with a higher pulse repetition frequency (PRF): for example, a 200-Hz PRF would reduce the image acquisition time by an order of magnitude. Third, RBC motion in superficial vessels could in principle cause a moving optical shadowing effect that influences the flow measurements in deeper regions. However, this is considered unlikely as light scattering would blur out any time variations in optical fluence that could compromise the measurements.

The present study demonstrates that blood flow speed can be measured using a clinical handheld US detector. This is somewhat surprising given that the high concentration of RBCs in whole blood (about 45 vol. %) results in mean cell separations of less than 10 μm , making it impossible to resolve individual cells using frequencies of tens of MHz. In this study, the highest frequencies that were detected were in the range of 20 MHz, which in water corresponds to a wavelength of about 75 μm , and yet motion was still detected.

It is possible that the RBCs exhibit heterogeneity on larger scales than individual cells due to RBC clumping or aggregation, enabling direct tracking of RBC ensembles.¹⁷ It might also be the case that the random distribution of RBCs enables motion detection without the need for any form of cell aggregation. It is known that the random distribution of RBCs cause speckle patterns in ultrasound imaging and these speckle patterns can be used to detect motion despite much larger wavelengths than the mean spacing between individual cells. The existence of speckle in PAs is an open question,^{21,22} but it may play a role in the detection of blood flow in AR-PAF. Future work will be aimed at investigating the source of observed contrast.

Disclosures

The authors have no relevant financial interests in the manuscript and no other potential conflicts of interest to disclose.

Acknowledgments

This work was supported by the EPSRC-funded UCL Centre for Doctoral Training in Medical Imaging (EP/L016478/1),

the Department of Health's NIHR-funded Biomedical Research Centre at Moorfields Eye Hospitals, and by the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 318067. The authors would like to thank Litron for the use of its laser.

References

1. P. J. van den Berg, K. Daoudi, and W. Steenbergen, "Review of photoacoustic flow imaging: its current state and its promises," *Photoacoustics* **3**(3), 89–99 (2015).
2. A. Sheinfeld and A. Eyal, "Photoacoustic thermal diffusion flowmetry," *Biomed. Opt. Express* **3**, 800–813 (2012).
3. R. Zhang et al., "In vivo optically encoded photoacoustic flowgraphy," *Opt. Lett.* **39**(13), 3814–3817 (2014).
4. W. Liu et al., "Photoacoustic thermal flowmetry with a single light source," *J. Biomed. Opt.* **22**(9), 096001 (2017).
5. L. Wang et al., "Ultrasonically encoded photoacoustic flowgraphy in biological tissue," *Phys. Rev. Lett.* **111**(20), 204301 (2013).
6. L. L. V. Wang et al., "Ultrasound-heated photoacoustic flowmetry," *J. Biomed. Opt.* **18**(11), 117003 (2013).
7. H. Fang, K. Maslov, and L. V. Wang, "Photoacoustic Doppler effect from flowing small light-absorbing particles," *Phys. Rev. Lett.* **99**(18), 184501 (2007).
8. A. Sheinfeld, S. Gilead, and A. Eyal, "Photoacoustic Doppler measurement of flow using tone burst excitation," *Opt. Express* **18**(5), 4212–4221 (2010).
9. J. Yao et al., "In vivo photoacoustic imaging of transverse blood flow by using Doppler broadening of bandwidth," *Opt. Lett.* **35**(9), 1419–1421 (2010).
10. J. Yao et al., "Label-free oxygen-metabolic photoacoustic microscopy in vivo," *J. Biomed. Opt.* **16**(7), 076003 (2011).
11. S.-L. Chen et al., "Photoacoustic correlation spectroscopy and its application to low-speed flow measurement," *Opt. Lett.* **35**(8), 1200–1202 (2010).
12. S.-L. Chen et al., "In vivo flow speed measurement of capillaries by photoacoustic correlation spectroscopy," *Opt. Lett.* **36**, 4017–4019 (2011).
13. B. Ning et al., "Simultaneous photoacoustic microscopy of microvascular anatomy, oxygen saturation, and blood flow," *Opt. Lett.* **40**(6), 910–913 (2015).
14. Y. Zhou et al., "Calibration-free in vivo transverse blood flowmetry based on cross correlation of slow time profiles from photoacoustic microscopy," *Opt. Lett.* **38**, 3882–3885 (2013).
15. J. Brunner and P. Beard, "Pulsed photoacoustic Doppler flowmetry using time-domain cross-correlation: accuracy, resolution and scalability," *J. Acoust. Soc. Am.* **132**(3), 1780–1791 (2012).
16. J. Brunner and P. Beard, "Acoustic resolution photoacoustic Doppler velocimetry in blood-mimicking fluids," *Sci. Rep.* **6**, 20902 (2016).
17. J. Brunner and P. Beard, "Velocity measurements in whole blood using acoustic resolution photoacoustic Doppler," *Biomed. Opt. Express* **7**(7), 2789–2806 (2016).
18. P. J. van den Berg, K. Daoudi, and W. Steenbergen, "Pulsed photoacoustic flow imaging with a handheld system," *J. Biomed. Opt.* **21**(2), 026004 (2016).
19. K. P. Koestli et al., "Temporal backward projection of photoacoustic pressure transients using Fourier transform methods," *Phys. Med. Biol.* **46**, 1863–1872 (2001).
20. B. E. Treeby and B. T. Cox, "k-wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave fields," *J. Biomed. Opt.* **15**(2), 021314 (2010).
21. Z. Guo, L. Li, and L. V. Wang, "On the speckle-free nature of photoacoustic tomography," *Med. Phys.* **36**(9), 4084–4088 (2009).
22. E. Hysi, D. Dopsa, and M. C. Kolios, "Photoacoustic tissue characterization using envelope statistics and ultrasonic spectral parameters," *Proc. SPIE* **8943**, 89432E (2014).

Biographies for the authors are not available.