

# Photoacoustic Velocimetry



Caitlin Smith

Department of Physics  
The University of Auckland

Supervisors: Jami Shepherd & Kasper van Wijk

A dissertation submitted in partial fulfilment of the requirements for the degree of BSc(Hons)  
in Physics, The University of Auckland, 2021.





# Abstract

Photoacoustic imaging is an up-and-coming medical diagnostic modality based on the recordings of high frequency pressure waves that are generated by the absorption of light inside the human body. These waves can be recorded by clinical ultrasonic transducers and interpreted to locate the position of optical absorbers. By repeating the experiment, movements of these absorbers can be tracked with the help of two flow algorithms, namely Colour Doppler and Vector Flow. This study is known as Photoacoustic velocimetry (PAV).

In this dissertation, photoacoustic simulations of moving absorbers were studied to determine the optimum flow algorithm which correctly estimates the flow speed, and, in the case of PAV, the direction. It was found that using an f-number reconstruction gave flow speeds that were closer to the true values than a method using discrete sub-apertures. This optimised method was applied to experimental data generated from a flow phantom, utilising a variety of blood-mimicking fluids and ultrasonic transducers. Singular value decomposition was used to remove image clutter from stationary absorbers in the reconstructed images before the flow algorithms are implemented.

Over a range of four pump speeds, the calculated flow speeds are proportional to the true flow rates. Flow maps are generated showing the distribution of particle movement across an image however both of the flow algorithms consistently under-estimate the flow speeds. PAV calculations were able to obtain flow speeds in both the lateral and axial directions, while colour Doppler could only determine the axial velocity component, with less accuracy than vector flow could.

Whilst PAV is not used clinically, it has several benefits over standard ultrasound methods for flow determination. While ultrasound's contrast is determined by impedance mismatch, photoacoustics utilises optical contrast, which has higher contrast for observation of chromophores such as those found in red blood cells. PAV has the potential to be better at measuring low flow speeds which would otherwise get lost in traditional ultrasound. Blood perfusion measurements are important for the diagnosis and treatment of a variety of diseases, motivating the development of this novel technique.



# Contents

<b>Abstract</b>	<b>1</b>
<b>1 Introduction</b>	<b>5</b>
1.1 Photoacoustic Velocimetry . . . . .	5
1.1.1 Overview . . . . .	7
<b>2 Theory</b>	<b>9</b>
2.1 Image reconstruction . . . . .	9
2.2 Colour Doppler flow . . . . .	10
2.3 Photoacoustic Velocimetry . . . . .	12
2.4 Flow speed considerations . . . . .	13
2.4.1 Maximum flow speed observable . . . . .	13
2.4.2 Minimum flow speed observable . . . . .	14
<b>3 Methodology</b>	<b>15</b>
3.1 Speed of sound optimisation . . . . .	15
3.2 Singular value decomposition . . . . .	16
3.3 Masking . . . . .	18
3.4 Simulated PA data . . . . .	18
3.5 Experimental set-up . . . . .	20
3.5.1 Light source . . . . .	21
3.5.2 Photoacoustic detection . . . . .	21
3.5.3 Flow phantom . . . . .	22
3.5.4 Blood mimicking fluid . . . . .	22
3.5.4.1 Graphite spheres . . . . .	23
3.6 Photos . . . . .	23
3.7 Data processing pipeline . . . . .	23
3.7.1 Sub-aperture reconstruction . . . . .	24
3.7.2 F-number reconstruction . . . . .	25
3.7.3 Colour Doppler computation . . . . .	27
3.7.4 PAV computation . . . . .	27
<b>4 Results</b>	<b>29</b>
4.1 Simulation results using colour Doppler . . . . .	29
4.2 Simulation results using PAV . . . . .	30
4.2.1 Low frequency simulation . . . . .	30
4.2.2 High frequency simulation . . . . .	31
4.3 Experimental results . . . . .	34
4.3.1 L11-5V probe . . . . .	34
4.3.1.1 Colour Doppler flow maps . . . . .	34
4.3.1.2 PAV . . . . .	36

4.3.2	L22-14xvLF probe . . . . .	37
4.3.2.1	Colour Doppler . . . . .	37
4.3.2.2	PAV . . . . .	38
4.3.2.3	Calculated flow velocity over a range of pump speeds with PAV	40
<b>5</b>	<b>Discussion</b>	<b>43</b>
5.1	Simulation results . . . . .	43
5.2	Experimental results . . . . .	44
5.3	Improvements & future work . . . . .	45
<b>6</b>	<b>Conclusions</b>	<b>47</b>
	<b>References</b>	<b>49</b>
<b>A</b>	<b>Summary of all experiments</b>	<b>53</b>
A.1	Other blood-mimicking fluids . . . . .	53
A.1.1	Blood-mimicking particle suspension . . . . .	53
A.1.2	Chocolate milk . . . . .	53
A.1.3	Indian ink . . . . .	53
A.2	Discussion of other experimental set-ups . . . . .	54

# Chapter 1

## Introduction

Photoacoustic imaging is a direct application of the Photoacoustic (PA) effect first described by Alexander Bell [7], where PA waves are generated in a sample through excitation by temporally modulated light. In a biological context, when light is absorbed by chromophores (part of a molecule responsible for its colour), it creates a small temperature rise, less than 0.1K. This produces an initial increase in pressure which relaxes via the emission of broadband low-amplitude acoustic waves (less than 10 kPa) [6]. These ultrasonic waves can then be detected at the surface of the sample and used to reconstruct an image of variations in optical absorption.

PA imaging can be seen as a variation on traditional ultrasound (US). In US imaging, an ultrasonic pulse is sent into a sample, where it is reflected and refracted off of tissue boundaries as a result of the acoustic impedance mismatch. This acoustic impedance mismatch depends on the difference between the speed of sound and density on either side of the boundary. These reflected ultrasonic waves can then be detected at the surface by the same transducer which initiated them. Alternatively, in PA imaging, a pulse of laser light enters the tissue and is specifically absorbed by chromophores receptive to the incoming light's wavelength. These tissues then emit broadband ultrasonic waves, which can be detected by a transducer on the surface. PA imaging is only made possible due to recent developments in laser technology enabling the generation of short-pulsed monochromatic light. The duration of laser excitation is effectively instantaneous compared to the time taken for PA waves to travel to the surface for detection. As a result, PA imaging is a modality that is based on an optical contrast while US imaging is based on a mechanical contrast. Some examples of optically active tissue components include haemoglobin, water and lipids [6]. This study will focus on blood as the optical absorber as it has a very high optical absorption coefficient relative to surrounding molecules in the visible to near-IR spectrum [6]. As a result, this method generates high-contrast images using endogenous contrast. There is a wide range of diseases that PA imaging could help diagnose and monitor [3].

A benefit of PA imaging is that pre-existing clinical US probes can be used for the detection of PA waves [12]. As US has been developed over many decades, these probes are fast and highly sensitive, making PA imaging a viable technique for use in a clinical setting soon [6].

### 1.1 Photoacoustic Velocimetry

Each optical absorber acts as a source of ultrasonic waves. When this source moves relative to the observer, a Doppler shift will be observed. The Doppler-shifted frequency depends on the location of the observer relative to the source's position and its speed. For example, when a moving particle that is emitting PA waves is travelling towards the receiver, the broadband emission is detected at higher frequencies, indicating a Doppler shift has occurred. Here the

term photoacoustic velocimetry, PAV, is used to describe the technique used to estimate the flow rate and direction using PA imaging. Meanwhile, PA flowmetry describes the determination of the flow speed when the angle of flow is known *a priori*.

In flowmetry and PA imaging as a whole there are two regimes used to localise the PA signal, they are Optical Resolution (OR) and Acoustic Resolution (AR) methods. OR methods rely on having the spot size of the laser tightly focused to an area whose size is comparable to the red blood cells we are measuring. As this regime deals with individual absorbers flowing through a tiny illuminated region, it is possible to track individual particles, resulting in the suspension appearing spatially heterogeneous. Due to the requirement of a small spot size, it can only probe up to a depth of approximately 1 mm before optical scattering creates a beam too diffuse for flowmetry. OR methods have been used successfully to determine transverse flow in phantoms and *in vivo* for studies on mouse ears and chicken embryos [35, 37].

The AR regime is still in its scientific infancy, but shows much promise for clinical applications. AR methods do not rely on a narrowly focused spot size to localise the PA signal. Instead, the origin of the PA signals is determined using a focused transducer or an array of receivers. The laser light is only required to supply diffuse illumination, enabling greater penetration than in the OR regime. The issue with this regime is in ensuring there is detectable heterogeneity of the flowing particles. If the separation between the particles is much smaller than the minimum detectable PA wavelength, then the fluid will be observed by the transducer as a stationary absorbing medium [24]. Despite this challenge, flow has been successfully measured in phantom studies in the AR regime [9, 10, 25, 31].

There are several new methods to extract the Doppler shift, and subsequently measure the speed of the moving PA source [32]. The first involves looking at the frequency shift resulting from continuous illumination in the OR regime. This method can not reconstruct images, but determines flow speed using a range-gating technique [16, 15]. Structured illumination, another OR technique, creates fringes to spatially modulate illumination while simultaneously modulating it with a pulsed laser. This method was shown to determine flow speed from 5-20mm/s in bovine blood [36]. The bandwidth broadening of the frequency spectrum can also be used to extract flow information in the OR regime [25, 35, 37]. As the first two methods listed above are trying to harness information by varying the light (through focusing or spatial modulation), these techniques lack the necessary penetration to study tissue at centimetre depths. Density tracking or displacement determination are techniques often used in the AR regime, which do not attempt to observe a Doppler shift. Instead, flow is determined using a pulsed laser to excite particles with a small time difference between them. The time-of-flight between pulses is extracted by cross-correlation of sequential pairs of waveforms, which is used to determine the flow speed, [8, 9, 11, 10, 12, 31]. Other flowmetry methods do exist and are discussed by Van Den Berg et al. [32].

The downside to these methods is that the angle of flow relative to the probe must be known in order for the flow velocity to be correctly determined. This is a significant challenge to overcome as, in a clinical setting, prior knowledge about the flow angle can not be assumed. US has developed advanced methods of vector-flow mapping as a way to overcome this [38, 39]. Another challenge is that, unlike US where the frequency content of our signal is known and determined by the probe used, the centre frequency of the PA wave is not inherently known and may change depending on the size, depth and speed of absorbers. This is an important quantity that must be accurately known to obtain accurate results from our flowmetry calculations.

To date, there has been one PAV study where the components of the velocity vector ( $v_z, v_x$ )

have been determined without *a priori* knowledge of the angle of flow [41]. This paper achieves what we hope to do, but using a very different method, where the speed is determined by studying the decorrelation of PA signal over time.

Even though US methods are already used clinically to measure and map flow, the nature of PAV has the potential to be more powerful in some scenarios. In PAs, a high-contrast acoustic signal is emitted by moving blood cells, whereas in US the acoustic waves are weakly reflected off of them. Therefore, PAV inherently offers a better SNR than US. In low flow-rate scenarios ( $< 50$  mm/s), US waves reflected by red-blood cells can be overpowered by much larger back-scattered signal from tissue moving due to respiratory or cardiac motion. In PAs, this problem is less prominent as PA waves are emitted primarily from blood flow and not tissue walls, resulting in less signal clutter. As a result, PAV is a promising technique to map the flow speeds in micro-vasculature, which has eluded US to date.

In a clinical setting, there is an array of applications that PAV may be suitable for. The ability to measure blood flow is important in the diagnosis and treatment of many diseases such as burns, port-wine stains, and cancers [33]. In the case of burns, the amount of perfusion (how much blood reaches a certain region) can indicate how quickly a burn will heal and whether surgery is required [18]. Furthermore, PAV could be used to study the perfusion of tumours as a way to evaluate the efficacy of drugs [6] as well as detecting angiogenesis (creation of new blood vessels) which is often a result of tumour growth or inflammatory disease such as rheumatoid arthritis [17].

### 1.1.1 Overview

In this dissertation, I will discuss the development of a new technique, termed photoacoustic velocimetry (PAV), which computes both the magnitude and direction of flow automatically. This is achieved by adapting vector flow methods from US to create a new technique for PA flow imaging. Using both simulated and experimental data from bench-top experiments to validate the procedure, both the magnitude and direction of flow can be determined and mapped for each pixel.





# Chapter 2

## Theory

The general recipe to obtain PAV maps begins with reconstructing images from raw PA signals. The value of a pixel over many subsequent frames is used to determine its phase shift. Finally, by changing the effective receiver angle in post-processing of the data, the velocity of each pixel in a series of reconstructed images can be obtained.

### 2.1 Image reconstruction

The “industry standard” delay and sum method was used to convert a series of pressure wave recordings from elements in a transducer array to an image. A mesh grid is created where the spatial arrangement of the sensors in the probe is on the x-axis, while the depths to be probed are on the y-axis. Then, for every sensor position along the x-axis, the Pythagorean theorem is applied to determine the distance from the sensor element to each pixel. The time delay is found by dividing this distance by the speed of sound determined for that medium as seen in Equation 2.1. The signal data is then interpolated at this delay time for each pixel relative to that specific sensor. Each sensor’s data is combined in a matrix which is resized to match the size of the image mesh. The envelope of the image is taken and then plotted to give our reconstructed image.

The time taken for a PA wave to travel from a pixel to the sensor is given by the following equation, where  $c$  is the speed of sound in the medium and the  $x/z$  coordinates are the locations of the pixel and the receiver.

$$t = \frac{\sqrt{(x_{\text{sensor}} - x_{\text{pixel}})^2 + (z_{\text{sensor}} - z_{\text{pixel}})^2}}{c} \quad (2.1)$$

Our reconstruction is implemented in C++ to speed up reconstruction to around 1 s per frame or less depending on the probe’s frequency, imaging grid size and resolution

Ideally, for PAV we want to create maps showing the spatial distribution of PA absorbers and their flow velocity. The sound detected at an element on our US probe is a superposition of pressure waves originating from wave sources originating from all points at a fixed radial distance to the element. Just from looking at the raw PA A-scan data it is impossible to determine where a given frequency originated from. This is shown in Figure 2.1.




Figure 2.1: The diagram on the left shows a simple pressure pulse arriving at a single receiving element. Given the speed of sound is known, we can calculate that it must have originated 1.5 cm away from the receiving element. This diagram on the right shows the possible origins of the signal, highlighting that in the raw PA data, it is impossible to know where certain frequencies (i.e those indicating flow) originated from.

In order to determine the origins of the PA waves, images need to be reconstructed to reflect the location of PA emitters. Discussion around reconstruction methods can be found in section 3.7.

## 2.2 Colour Doppler flow

Colour Doppler is a commonly used US flowmetry method which I have adapted from the literature to be used in a PA setting [14, 40, 5].

For a particle travelling parallel with the probe, the detected frequency shift in the ultrasonic PA wave is given by:

$$f_d = f_0 \frac{v}{c}, \quad (2.2)$$

where  $f_0$  is the centre PA frequency generated by a particle if it was stationary as observed by the detector,  $f_d$  is the shift in the centre frequency caused by the movement of the particles towards or away from the detector,  $c$  is the speed of sound in the medium separating the probe and the detector while  $v_z$  is the speed the particle is moving at perpendicular to the detector. Experimentally, the centre frequency is taken to be the frequency of the spectrum with the greatest amplitude.

When the flow is at an angle,  $\alpha$ , to the probe, where  $\alpha$  is the angle between the normal to the surface and the direction of flow, the Doppler shift is given by:

$$f_d = f_0 \frac{v}{c} \cos \alpha. \quad (2.3)$$

In colour Doppler, the flow angle  $\alpha$  is known so the component of the particle's velocity perpendicular to the probe can be determined. When  $\alpha = 0$ , the axial component of flow is given by:

$$\begin{aligned} v_z &= \frac{c f_d}{f_0} \\ &= \frac{c \omega}{2\pi f_0}. \end{aligned} \quad (2.4)$$

When measuring flow rates that are constant with time, the angular frequency of the Doppler shift,  $\omega$ , is taken to be the average angular frequency shift.

The reconstructed images show the spatial distribution of PA emitters (optical absorbers), and the strength of their emissions (pixel amplitude). However, each reconstructed image only shows a snapshot in time, and no temporal information is present. As a result, a large sequence of frames must be taken to create a stack of images, to extract the PA wave emission strength as a function of time for each pixel. A visual representation of a stack of scans is shown in Figure 2.2.

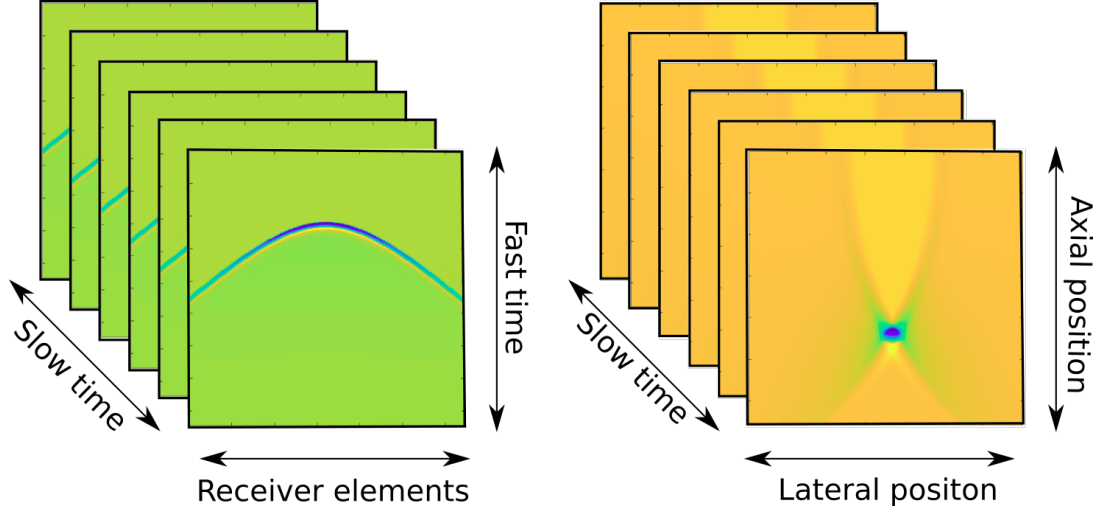


Figure 2.2: These diagrams show the dimensions of the different data stacks. The diagram on the left shows the stack containing the raw PA signals obtained as RF signals. This stack is comprised of many scans across slow-time (time between individual scans), with each scan having a time series (fast-time) for each receiving element. The diagram on the right is a stack of images generated from reconstructing each of the PA frames. Each scan has dimensions controlled by the size of the imaging grid.

Now, we can look at a pixel's amplitude as a function of time through the various frames and obtain information about the Doppler frequency spectrum originating at this point. The only unknown value we need to calculate the flow speed from Equation 2.4 is the angular frequency,  $\omega$ . When  $\omega$  is known at every pixel in our image a flow map can be generated. The angular frequency can be expressed in the following way:

$$\omega = \frac{d\phi}{dt} \approx \frac{\phi_i - \phi_{i-1}}{T}. \quad (2.5)$$

Where  $\phi_i$  is the phase of frame  $i$ , and  $T$  is the time between frames, also known as inverse of the pulse repetition rate (PRF). To find the phase of the pixel as a function of time, the real signal must be converted into an analytic signal of the form  $i(t) + jq(t)$ . The components of this signal are also known as the in-phase magnitude  $I$  and the quadrature magnitude  $Q$ . Discussion of how this occurs is presented in the methodology section (Section 3). Once we have an in-phase and quadrature signal for each pixel as a function of time, the instantaneous phase of a given pixel can be calculated as follows:

$$\begin{aligned} \tan \phi_i &= \frac{Q_i}{I_i} \\ \phi_i &= \arctan \frac{Q_i}{I_i}. \end{aligned} \quad (2.6)$$

Using the following trigonometric identity;  $\arctan x - \arctan y = \arctan\left(\frac{x-y}{1+xy}\right)$ , the phase shift between two subsequent scans is as follows:

$$\begin{aligned}
\phi_i - \phi_{i-1} &= \arctan\left(\frac{Q_i}{I_i}\right) - \arctan\left(\frac{Q_{i-1}}{I_{i-1}}\right) \\
&= \arctan\frac{\frac{Q_i}{I_i} - \frac{Q_{i-1}}{I_{i-1}}}{1 + \frac{Q_i Q_{i-1}}{I_i I_{i-1}}} \\
&= \arctan\frac{I_i I_{i-1} \left(\frac{Q_i}{I_i} - \frac{Q_{i-1}}{I_{i-1}}\right)}{I_i I_{i-1} + Q_i Q_{i-1}} \\
&= \arctan\frac{Q_i I_{i-1} - Q_{i-1} I_i}{I_i I_{i-1} + Q_i Q_{i-1}}.
\end{aligned} \tag{2.7}$$

Now, as we are taking multiple pairs of sequential waveforms over slow-time, we can obtain a mean angular frequency for a given pixel. This is calculated by summing over the number of pulse pairs like so:

$$\begin{aligned}
\omega &= PRF \times \arctan\left[\frac{Q_i I_{i-1} - Q_{i-1} I_i}{I_i I_{i-1} + Q_i Q_{i-1}}\right] \\
\bar{\omega} &= PRF \times \arctan\left[\frac{\sum_{i=1}^n Q_i I_{i-1} - Q_{i-1} I_i}{\sum_{i=1}^n I_i I_{i-1} + Q_i Q_{i-1}}\right].
\end{aligned} \tag{2.8}$$

Equation 2.8 is also known as the lag-one autocorrelator commonly used in US Doppler [14, 5, 22]. For every pixel,  $\bar{\omega}$  can be estimated with Equation 2.8, so the axial velocity can be calculated like so:

$$v = \frac{c\bar{\omega}}{2\pi f_0}. \tag{2.9}$$

One major drawback of colour Doppler is that it can only determine the flow component perpendicular to the probe. This means that flow information in the plane of the image (along the B-scan) can only be obtained if the angle is manually chosen by the operator and assumed constant within the soft-tissue. Therefore, a method that automatically determines the flow direction as well as the speed at each pixel is highly desirable.

## 2.3 Photoacoustic Velocimetry

In the colour Doppler flow description, we explained that when the flow is at an angle,  $\alpha$ , to the probe, the Doppler shift is given by:

$$f_d = f_0 \frac{v}{c} \cos \alpha. \tag{2.10}$$

As shown in Figure 2.3, when the receiving angle,  $\phi_n$ , is no longer perpendicular to the surface, the effective angle between the receiving beam and the direction of flow yields the following:

$$f_d = f_0 \frac{v}{c} \cos(\phi_n - \alpha). \tag{2.11}$$

Using the difference formula for cosines ( $\cos(A - B) = \cos(A)\cos(B) + \sin(A)\sin(B)$ ), the equation above can be rewritten as follows:

$$\frac{f_d c}{f_0} = v \cos \alpha \cos \phi_n + v \sin \alpha \sin \phi_n. \tag{2.12}$$

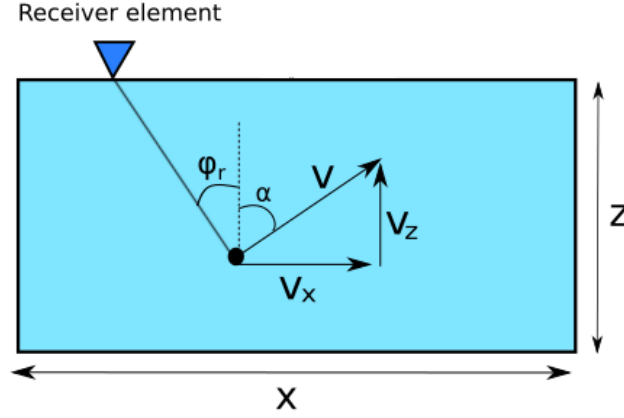


Figure 2.3: Diagram depicting the origin angles used in a PAV calculation.

Now we can substitute some of these quantities for the flow velocity components in the  $x$  and  $z$  directions,  $v_x$ ,  $v_z$ :

$$\begin{aligned} v_x &= v \cos \alpha \\ v_z &= v \sin \alpha. \end{aligned} \quad (2.13)$$

Now the equation 2.12 can be rewritten as:

$$v_x \cos \phi_n + v_z \sin \phi_n = \frac{f_n c}{f_0} = \frac{c \bar{\omega}}{2\pi f_0}. \quad (2.14)$$

Where  $f_n$  is the Doppler frequency shift for the  $n$ th receive angle. In matrix form:

$$\begin{bmatrix} \cos \phi_1 & \sin \phi_1 \\ \vdots & \vdots \\ \cos \phi_n & \sin \phi_n \end{bmatrix} \begin{bmatrix} v_x \\ v_z \end{bmatrix} = \frac{c}{2\pi f_0} \begin{bmatrix} \bar{\omega}_1 \\ \vdots \\ \bar{\omega}_n \end{bmatrix}. \quad (2.15)$$

For each pixel and receiving angle,  $\phi_n$ , an over-determined linear system can be solved to determine  $v_x$  and  $v_z$ :

$$\begin{bmatrix} v_x \\ v_z \end{bmatrix} = \left( \begin{bmatrix} \cos \phi_1 & \dots & \cos \phi_n \\ \sin \phi_1 & \dots & \sin \phi_n \end{bmatrix} \begin{bmatrix} \cos \phi_1 & \sin \phi_1 \\ \vdots & \vdots \\ \cos \phi_n & \sin \phi_n \end{bmatrix} \right)^{-1} \begin{bmatrix} \cos \phi_1 & \dots & \cos \phi_n \\ \sin \phi_1 & \dots & \sin \phi_n \end{bmatrix} \begin{bmatrix} c\bar{\omega}_1/2\pi f_0 \\ \vdots \\ c\bar{\omega}_n/2\pi f_0 \end{bmatrix}. \quad (2.16)$$

Using the same methodology as the colour Doppler method,  $\bar{\omega}$  can be estimated for each pixel with Equation 2.8. Subsequently, the components  $v_x$  and  $v_z$  can be solved using a least-squares optimisation of Equation 2.16. The only difference is that  $c\bar{\omega}/2\pi f_0$  is calculated for every pixel and every receiver angle, so the equation above may be used to determine both the axial and lateral flow components.

## 2.4 Flow speed considerations

### 2.4.1 Maximum flow speed observable

To find the range of velocities to measure in the experiment, we had to work out the maximum velocity we could measure. This was limited by the temporal separation of images formed,

which was limited by the laser repetition rate and the Nyquist limit,

$$f_{max,Doppler} = \frac{PRF}{2}.$$

In Equation 2.3 the Doppler shift for PA modulation is :

$$f_d = f_0 \frac{v_{flow}}{c} \cos \theta.$$

For example, to find the maximum flow velocity for the L11-5V probe, the maximum Doppler shift corresponds to  $f_{max,Doppler}$  above. By equating these two relations and rearranging for  $v_{flow}$  with  $\theta = 0$ , the maximum lateral flow measurable is given by;

$$v_{flow,max} = \frac{f_{d,max}}{f_0 \cos \theta} c = \frac{PRF/2}{f_0 \cos \theta} c = \frac{20Hz \times 1500m/s}{2 \times 7.5 \times 10^6 Hz} = 2mm/s.$$

If we try to detect a flow speed faster than the  $v_{max}$  permissible by the Nyquist limit aliasing will occur and the correct flow speed can not be determined.

If the angle  $\theta$  is increased, the maximum flow speed measurable also increases. This is because this limit,  $v_{flow,max}$  above applies to the lateral component of flow only. Therefore by measuring flow at an angle, we can measure a greater range of speeds without aliasing.

#### 2.4.2 Minimum flow speed observable

The minimum observation time depends on the time needed to observe one period of the Doppler-shifted frequency,

$$T = \frac{1}{f_d} < \frac{N}{PRF}.$$

Where N is the number of frames in a scan, PRF is the pulse repetition frequency and  $f_d$  is the Doppler frequency. Using the formula for the Doppler shift (Equation 2.3) we can insert it into the inequality:

$$\frac{1}{f_d} = \frac{c}{f_0 v_{min} \cos \theta} < \frac{N}{PRF}.$$

Rearranging for the flow speed, we find the minimum flow speed is given by:

$$v_{min} = \frac{cPRF}{f_0 \cos \theta N}. \quad (2.17)$$

Therefore, by including more frames (samples in slow-time) we are able to measure smaller flow speeds.

# Chapter 3

## Methodology

### 3.1 Speed of sound optimisation

To reconstruct the clearest images, the speed of sound used in the image reconstruction must be chosen correctly. To do this I used multiple focusing criteria described in the literature [30]. A 3D stack of images is reconstructed like a book, where each of the pages is an image reconstructed using a different speed of sound. Each method in the focusing function is plotted to show how the different methods determine the focus for each speed of sound. An example is shown in Figure 3.1 for the reconstruction of a section of tubing from PA data. These methods should converge on the true speed of sound to give the best image resolution.

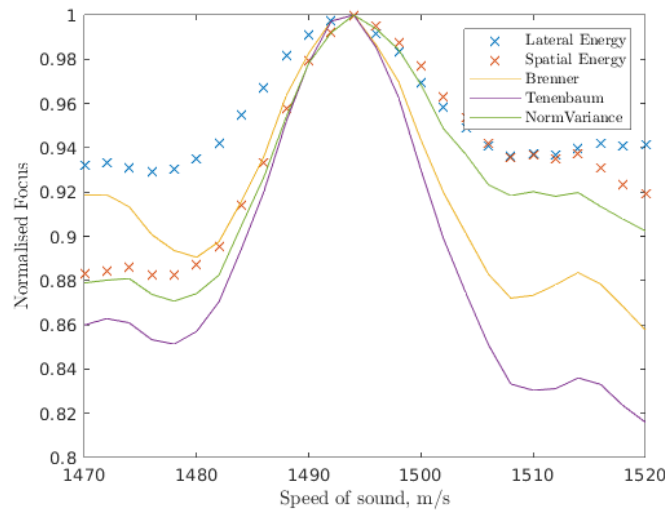


Figure 3.1: Example of a speed of sound optimisation output for a reconstructed tube along the long axis. The normalised focus as determined by the various focusing functions are plotted as a function of the speed of sound. In this example, the focusing functions converge at 1494 m/s.

We found that using power images instead of the regular images converged on the correct speed of sound most repeatably. This is because the power image is just the reconstructed image squared so noise will be suppressed relative to the signal. When determining the focus we only want to observe how the region of interest comes into focus, not the noise around it.

### 3.2 Singular value decomposition

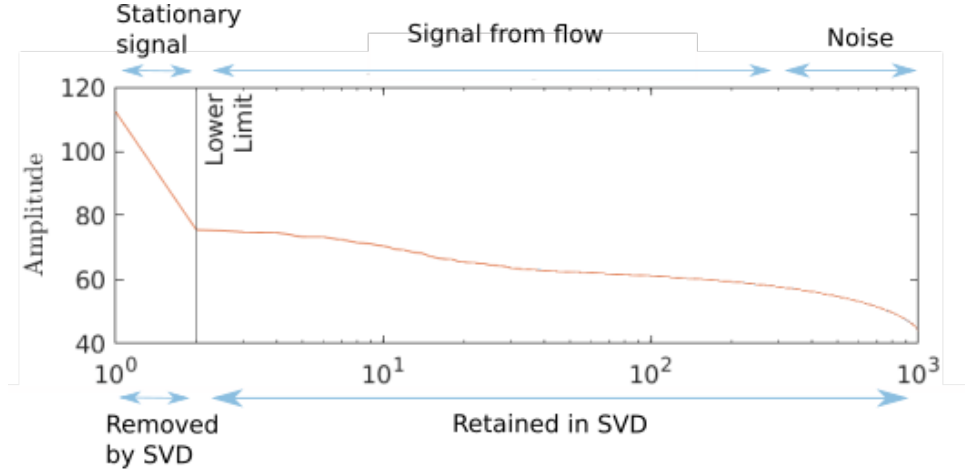


Figure 3.2: This plot shows the amplitude of singular values order in decreasing size from an experiment described later in this thesis. The singular values (on the x-axis) are shown in order of their amplitude. The vertical line on the left shows the boundary between singular values corresponding to stationary signal (high spatio-temporal coherence) and those that exhibit low spatio-temporal variability (corresponding to flow and noise). No cut off has been shown for noise as it is low-powered so has minimal effect on the resulting images.

Singular value decomposition (SVD) is used to reject clutter in our images[13, 41, 34, 4]. It utilises the different properties of our samples components including spatio-temporal coherence to distinguish moving absorbers from the stationary background medium. If a pixel's spatio-temporal coherence varies greatly over  $N$  frames it is attributed to blood flow, while if it has a high spatio-temporal coherence it is attributed to stationary media.

A 3D matrix comprised of a stack of 2D images over time is reformed into a 2D Casorati matrix, where each column of an image is added end to end to create a column vector. Each of these columns from the  $N$  frames is placed side by side to make a 2D array. This Casorati matrix (known as  $S$ , with dimensions equal to  $n_x \times n_z, n_t$ ), has a singular value decomposition given by:

$$S = U\Delta V^*,$$

where  $\Delta$  is a non-square diagonal matrix with the same dimensions as  $S$ , while the columns of  $U$  and  $V$  correspond to the spatial and temporal singular vectors of  $S$ . SVD is a way of decomposing a matrix,  $S$ , into a weighted, ordered sum of separable matrices  $A_i$ , such that  $S$  can be written as:

$$S = \sum_i \lambda_i A_i = \sum_i \lambda_i U_i V_i,$$

Where  $\lambda_i$  are ordered singular values. Here each column  $V_i$  corresponds to a temporal signal with length  $n_t$  and each column of  $U_i$  corresponds to a spatial signal of length  $n_x \times n_z$ .  $U_i$  also describes a spatial image,  $I_i$ . Overall our original 3D matrix of image frames at each time interval can now be written as:

$$s(x, z, t) = \sum_{i=1}^{\text{rank}(S)} \lambda_i I_i(x, z) V_i(t),$$

There are three groups that the spatial images,  $I_i$ , can be separated into based on spatio-temporal coherence. Images corresponding to low  $\lambda$  values have high spatio-temporal coherence



and correspond to stationary signal. Images corresponding to flow have low spatio-temporal coherence, while noise exhibits no spatio-temporal coherence, so it will correspond to the highest  $\lambda$  values. By removing images with high  $\lambda$  values, the remaining images containing only flow and noise. Choosing which  $\lambda$  values to included is done using a plot like the one in Figure 3.2.

Figure 3.2 is an example of these singular values plotted in order of size, with the y-axis representing the power of that singular value for all singular values ordered from largest to smallest along the x-axis. The first vertical line marks the sudden transition in power between the stationary signal and temporally varying signal. In this example, this occurs between the largest and second-largest singular value. We could go one step further and add a threshold at the lower power singular values, which corresponds to noise. Removing these values doesn't have a huge effect as these singular values have low powers so they do not change the resulting image significantly.

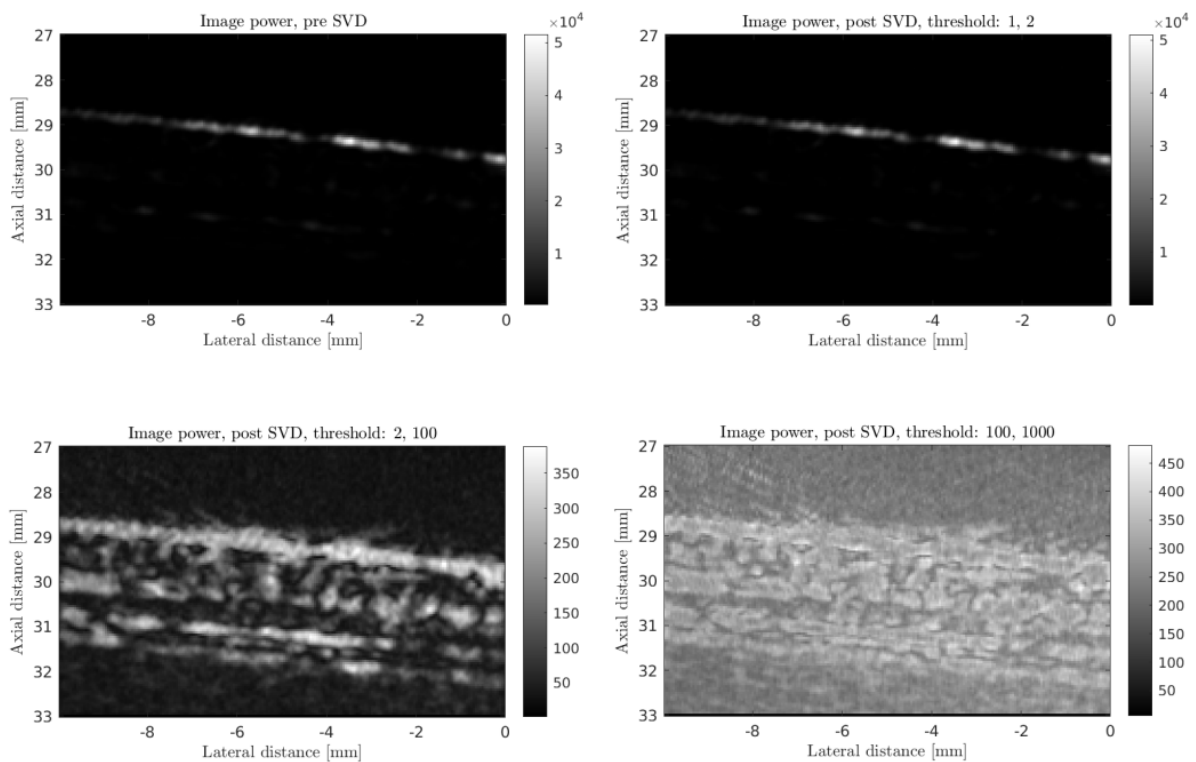


Figure 3.3: Power images with SVD applied. Top left; image power for a section of tubing before SVD. The image is swamped by the signal from the top of the tube, the rest of the tube can't be resolved. Top right; The first singular value only; this contains most of the image's power and corresponds to the stationary signal. Bottom left: result with SVD, the predominant stationary signal has been removed as well as some of the highly variable noise. This image displays the depth of the tube and has good contrast. Bottom right; power image with too much low amplitude noise included and flow signal removed, resulting in poor contrast.

To demonstrate how the inclusion/exclusion of certain singular values changes the resulting power image, four images with different threshold singular values are shown in Figure 3.3. The power image is calculated to be the magnitude squared of every frame added together. The top left is the image with all singular values included (threshold = 1:1000 for an image with 1000 frames). The top right image corresponds to the largest singular values representing the stationary signal. This singular value is removed to show where flow is occurring. Note how the amplitude of the top row's images are similar, showing that most of the original image's

power comes from the stationary component. The bottom left image excludes singular values corresponding to stationary signal and noise. Notice how the inside of the tube is now clearer than in the previous two images. Finally, the bottom right image is just the lowest powered singular values, largely dominated by noise. This image has a similar power to the one before despite being composed of 10 times more singular values. This shows that this is primarily low power noise.

### 3.3 Masking

As our signal should only originate from optically absorbent particles, we only want to isolate signals from flowing particles that are not multiply reflected. To do this we can either manually draw a mask or make a mask based on pixel power (power mask). A power mask is generated from a stack of reconstructed images whose pixel values have been squared to create a map showing image power. High amplitude regions are where blood mimicking fluids are present, so by determining a threshold power, we can identify optically absorbent regions and use this to only study the flow in the relevant locations. Power masks are susceptible to including high powered artefacts that are generated by multiply-reflected PA waves. The manual mask is generated by the operator who selects the area where flow is expected to be in the image, often based on a B-mode image. This prevents the inclusion of artefacts as the operator can compare the reconstructed image to the phantom and only include regions that should be present.

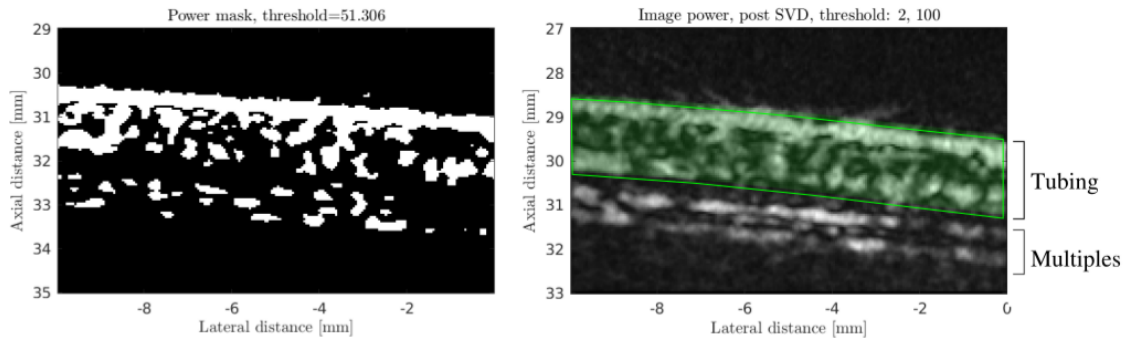


Figure 3.4: Mask created using the power mask and manual mask method shown for the power image of a section of straight tubing. In this example, the power mask is ineffective as it doesn't include low power regions of the tube, but it does include some artefacts created by reflections in the tubing. The manual mask on the right is much better as the operator can ensure the mask is the correct size as the diameter of the tubing is known.

Figure 3.4 shows the two methods for making a mask based on the mean power of a reconstructed image of a tube. On the left, a power mask has been created which does not include relevant low-power signals and includes multiples caused by PA reflections in the tube. The image on the right shows a manual mask, which can be constructed using the known dimensions of the tubing to separate real signals from multiples.

### 3.4 Simulated PA data

Using the k-Wave package for time-domain simulation of acoustic wave fields, we can simulate PA flow experiments in Matlab [28, 29, 27]. The major advantage of simulating PA data is the quality of the data that can be obtained. There is no noise and multiple reflections are eliminated by simulating a homogeneous medium. Also, by specifying the flow speed in the simulation, where it comes to determining the flow speed of reconstructed images, the accuracy

of our methods can be quantified by comparison to a known ground truth.

To simulate flow I started with particles randomly distributed in a tube shape as seen in Figure 3.5. These particles are defined as acoustic emitters with a given initial pressure value and diameter. After each time step, the emitters are shifted a set distance in the desired direction of flow and the simulation is repeated. These distances are a quantised number of grid points so the particles move the exact same distance in each time interval. When the particles move off the edge of the imaging grid, they are regenerated back at the other end of the tube, so there is always the same number of emitters in the grid. In the simulation, I not only define the location, size and number of emitters, but also the location and characteristic of the probe. In the two simulations trialled, I mimicked two readily-available transducers used in our experiments by defining the sensor length, number of elements and spectral characteristics. The specifications of these simulations are outlined in Table 3.1 below.

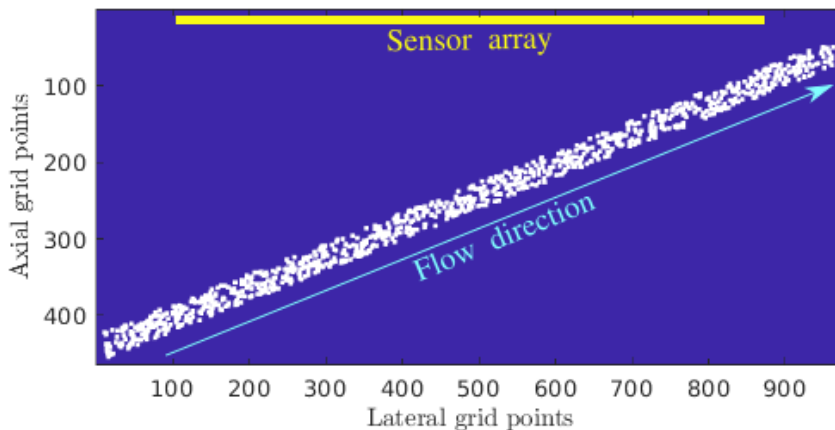


Figure 3.5: Simulation grid with a sensor array and particles arranged within a tube. This was the set-up for the low resolution simulation.

Initially, I started with a high resolution simulation to simulate discs roughly the size of red blood cells (on the order of  $10\mu\text{m}$ ) as if detected by our L11-5V probe with a centre frequency of 7.5 MHz. The characteristic of this probe and the others used in the experiments can be found in Table 3.2. However, to simulate objects that small over a 2 cm by 4 cm large imaging grid, the number of grid points was large, making each loop of the simulation take a long time (around 3 hours per frame).

Another thing to consider is the stability, which is described by the one-dimensional Courant-Friedrichs-Lewy (CFL) number. This is defined as the ratio of the distance a wave can travel in one time step to the grid spacing [26]:

$$CFL = \frac{c_0 \Delta t}{\Delta x} < 0.3.$$

The CFL number can be considered as a non-dimensionalised time step, which is useful for defining the maximum permissible time step for sampling in fast-time [26]. For the simulations to be stable, k-Wave recommends using a CFL of less than 0.3, so for the high resolution simulation, the sampling frequency must be at least 500 MHz. Given the maximum depth in the medium is roughly 2cm, I needed to collect around 8500 samples per receiver element per frame. This, coupled with the high spatial resolution and a high sampling frequency, makes it is easy

to see why this simulation is so computationally intensive, taking around 3 hours to simulate tens of microseconds of wave propagation.

The other constraint from k-Wave is that it must have at least two grid points per wavelength [26].

$$\frac{\lambda}{\Delta x} = \frac{c_0}{f_0 \Delta x} < 2$$

In the high resolution simulation, I used roughly 20 grid points per wavelength (significantly more than the minimum recommended number of two), indicating that the product of the centre frequency  $f_0$  and the resolution ( $\Delta x$ ) could be increased further. This was implemented by creating a lower resolution simulation that would run more quickly while meeting CFL and grid point per wavelength constraints. To do this I started by mimicking the lower resolution probe, the P4-1, with a grid size that had 1/5th of the spatial resolution compared to the high frequency simulation. This limits the smallest possible size a particle can be to 100  $\mu\text{m}$ , much larger than a red blood cell; however it resulted in a significant decrease in computation time. With this probe's centre frequency I could sample speeds of up to 6 mm/s, meaning it was simple to find combinations of quantised steps that would give speeds of around 6 mm/s.

Finally, a perfectly matched layer (PML) is added around the imaging grid. As the simulation uses periodic boundary conditions, PA waves can be seen leaking out of one side of the simulation and appearing on the other. By adding a PML around it, these PA waves are attenuated so there is nothing to re-emerge on the other side.

	Simulation 1	Simulation 2
Probe mimicked	L11-5V	P4-1
Probe centre frequency	7.6 MHz	2.5 MHz
Spatial Resolution	10 $\mu\text{m}$	50 $\mu\text{m}$
Sampling Frequency	650 MHz	150 MHz
Number of temporal samples per frame	8500	2500
Grid size (x $\times$ z)	4.05 cm $\times$ 2.00 cm	4.88 cm $\times$ 2.32 cm
Grid points (x $\times$ z)	4048 $\times$ 2000	976 $\times$ 464
Frames per second (PRF)	20	20
Number of particles	1000	1000
Particle diameter, mm	20 $\mu\text{m}$	100 $\mu\text{m}$
Axial Speed, mm/s	0.4 mm/s	2 mm/s
Lateral Speed, mms	1 mm/s	5 mm/s
Magnitude of speed, mm/s	1.044 mm/s	5.077 mm/s
Computation time per frame	$\approx$ 3 hours	$\approx$ 3 minutes
Number of frames simulated	50	200

Table 3.1: Summary of the differences between the two different simulations carried out to validate PAV and colour Doppler methods.

### 3.5 Experimental set-up

The experimental set-up consists of two main components; the laser and the ultrasound probe and is shown in Figure 3.6.




Figure 3.6: PAV bench-top set-up. The red box is the pump which pushes a syringe full of blood-mimicking fluid through the tubing. The US probe is the white box which is suspended into the water bath from above. The fibre laser output is the grey tube which is fixed to the US probe by a 3D printed probe holder in our actual experiments. The fibre is at an angle to the US probe which is determined by a 3D printed probe holder not shown in this schematic.

### 3.5.1 Light source

The laser is an optical parametric oscillator (OPO, Oportek RADIANT 532 LD Tunable Optical Parametric Oscillator Laser, Carlsbad, CA) which allows the wavelength to be tuned from 680-2600 nm. It emits pulsed laser light at a repetition frequency of 20 Hz with a pulse duration of approximately 5 ns. This repetition rate limits the flow speeds which can theoretically be observed with our set-up. The derivation of the maximum flow speed observable for a given laser/transducer is in Section 2.4.

The laser light is coupled to a fibre which enables it to be directed at the region of interest. The fibre used produces a circular beam with a spot size of around 5 mm at the output. The limiting safety factor of PA imaging is the laser pulse energy [6]. As a result, a pulse energy output of around 20 mJ/cm<sup>2</sup> was chosen, which is the limit for nanosecond pulsed light under 1000 nm entering soft tissue [2][1]

The choice of the laser wavelength depends on what is being imaged. Ideally we want to maximise absorption by the blood-mimicking fluid while minimising absorption by water which mimics the surrounding soft tissue. In the wavelength range available to the OPO, the lowest absorption by water occurs at the smallest wavelength available at 680 nm [19]. This is the wavelength used across all the experiments mentioned in this thesis.

### 3.5.2 Photoacoustic detection

The other key component of our set up is the US probe. As probes are only detecting the emitted PA waves and not generating US, there are no safety concerns regarding this apparatus [6]. Table 3.2 below summarises the transducer characteristics:

Probe	Centre Frequency	Number of elements	Sampling Frequency	Length	Maximum detectable flow at $\theta = 0$
Verasonics L11-5V	7.5 MHz	128	31250000	38.4 mm	2 mm/s
P4-1	2.5 MHz	96	10000000	28.32 mm	6 mm/s
Verasonics L22-14vXLF	18 MHz	128	62500000	12.8 mm	0.83 mm/s

Table 3.2: Transducer characteristics

The starting point of these experiments was the L11-5V probe, as it is our mid-frequency range probe, which offers a better resolution than the P4-1 probe, but also enables decent flow speeds (above 1 mm/s) to be determined. The resolution of the reconstructed image is determined by the sampling frequency of the transducer; the more frequently the pressure is sampled the more precise the reconstructed images will be. The lower centre frequency of the P4-1 probe enabled higher flow speeds to be observed, which prevented the blood mimicking fluids from getting stuck or dropping out of suspension in the tubing. The downside of this is the decrease in resolution due to its lower sampling frequency (1/3 the resolution of that from the L11-5V probe). On the other hand, the L22-14vXLF probe provided the highest resolution with a reasonable focal depth of 20mm. This probe has twice the resolution as the L11-5V, however, the downside of this is that the maximum flow speed observable is very slow. This probe also featured in recent PA flowmetry studies [41] [34]. The Verasonics Vantage 128 High Frequency system was used to control the imaging system and data acquisition. Matlab is used to program the US probe and laser such that each laser pulse had its own recorded PA response. As a result, all further image processing and analysis is undertaken in Matlab.

### 3.5.3 Flow phantom

The flow phantom used was created using polymer tubing (OD 1/8" (3.18 mm), ID 1/16" (1.5875 mm)) attached to a syringe pump (NE-4000 Programmable 2 Channel Syringe Pump, New Era Pump Systems Inc.). The tubing is placed in a Pyrex dish of water with the probe suspended from above to allow the transmission of the PA waves. The laser fibre is attached to the US probe from above using a 3D printed probe holder. One end of the syringe pump is propped up by a small plate to ensure air bubbles don't flow towards the tip of the syringe. The tubing was slightly angled using metal plates to prop up one end. A benefit of having it slightly angled is that colour Doppler, which measures towards/away from the probe, will have some movement in the axial direction to measure.

To run the experiment, at least 1000 frames at each of the flow speeds are recorded. For the L11-5V probe, the theoretical minimum observable flow speed for this many frames is  $4\mu\text{m/s}$  from Equation 2.17. To take an image, a short laser pulse is fired and the US probe's sensors record the arrival of PA waves at the detector. The number of samples and the sampling rate depends on the probe (see Table 3.2). The laser repetition rate is 20Hz, so for each of the probes used, 1000 frames took 50 seconds to record.

### 3.5.4 Blood mimicking fluid

Blood-mimicking fluids are used instead of real blood as it enables repeatable experiments for the development of our methods. A good blood mimicking fluid has to following characteristics:

- Strong optical absorption where water has a weak optical absorption (in the visible region of the spectrum).
- Composed of tiny absorbers in solution, around the same size as red-blood cells (about 7  $\mu\text{m}$  wide and 2 $\mu\text{m}$  thick ).



- Has well-known optical and mechanical properties.
- If the mimic is a suspension it will have some mechanism to prevent particles from aggregating, getting stuck in the tubing or falling out of suspension.

While this thesis only describes the set-up which gave good flow results, many other experiments were undertaken to get to this point. A table summarising these can be found in the Appendix.

### 3.5.4.1 Graphite spheres

The successful trials used a blood-mimicking fluid composed of carbon particles suspended with a volume fraction of 15%, where the particle diameters distributed from 2 to 12 micrometres (carbon glassy spherical powder, 484 164, Sigma-Aldrich) as utilised by Fang *et al* [15]. These particles are a good size match to red-blood cells. Sodium polytungstate powder is added to increase the density of the solution to prevent the carbon particles from settling out. At the desired density, roughly  $\sim 15$  g had been added to a 30 mL solution. This increased the viscosity of the blood-mimic dramatically, making it difficult for the pump to push it through the narrow tubing. Tween-20 (1%, Sigma-Aldrich) is added to the solution to prevent aggregation of the graphite particles. As we used our entire supply of graphite powder to make the first batch of the blood-mimic, it was saved after each experiment to use in the next. Each time it was used, 1 mL of water was stirred in to break up any congealed lumps. As a result, the concentration of the graphite powder in the blood mimicking fluid is in the range of 10 to 20 %.

## 3.6 Photos




Figure 3.7: Flow experiment set-up with labels. In this set up the 3D printed probe holder is used to ensure a fixed relation between the excitation laser fibre and the US probe. This is the arrangement used for the two successful flow experiments.

## 3.7 Data processing pipeline

Each data set is loaded into Matlab. The 3D matrix has dimensions given by the number of samples, the number of receiving elements in the transducer and the number of frames. Next, the first samples from each receiving element were removed as the tip of the US probe was also illuminated so emits its own PA signal. These initial samples are referred to as the lens signal, and the length of this signal depends on which transducer is used. Removing the lens signal effectively selects when  $t=0$  is and ensures reconstructions are accurate. At this point it is possible to apply a bandpass filter to the data. I used Matlab's `fir1` function, which is used

to smooth the impulse response within a window [21]. It is used in a bandpass capacity, where it preserves frequencies closest to the centre frequency and minimise signals further out. The bandpass filter is used, as for the flow estimates to be accurate, the centre frequency must be well known, however the PA signal is fairly broadband. The cut-off frequencies are defined as  $0.5 \times \frac{f_0}{f_s} < f < 1.5 \times \frac{f_0}{f_s}$ , so that it scales with the centre frequency and sampling frequency of the probe.

The data is reconstructed in two ways: one to implement colour Doppler and one for PAV. Both methods perform the reconstructions using two different methods for defining the angles between the pixels and the relevant receiver. The reason a number of different angles are needed is for the over-determined linear equation in Equation 2.15. The more angles used, the more phase estimates obtained for a pixel and the more accurately the velocity components can be estimated.

### 3.7.1 Sub-aperture reconstruction

This method splits the linear transducer into smaller segments and reconstructs the imaging grid using only those elements in a segment at a given time. This has the effect of changing the angle between the centre of the receiving elements and the imaging grid, as shown in Figure 3.8 below. The angle between a given pixel and the receiver is set to be the angle between a pixel and the centre of the sub-aperture, so the smaller the sub-aperture, the more accurate the angle value is. This code returns the I/Q reconstructed image from each of the sub-apertures, which can be fed into the PAV algorithm. Alternatively, by combining the I/Q images over all of the sub-apertures, I/Q images can be obtained as reconstructed with the conventional delay and sum method. These total I/Q images are used as an input to the colour Doppler algorithm.

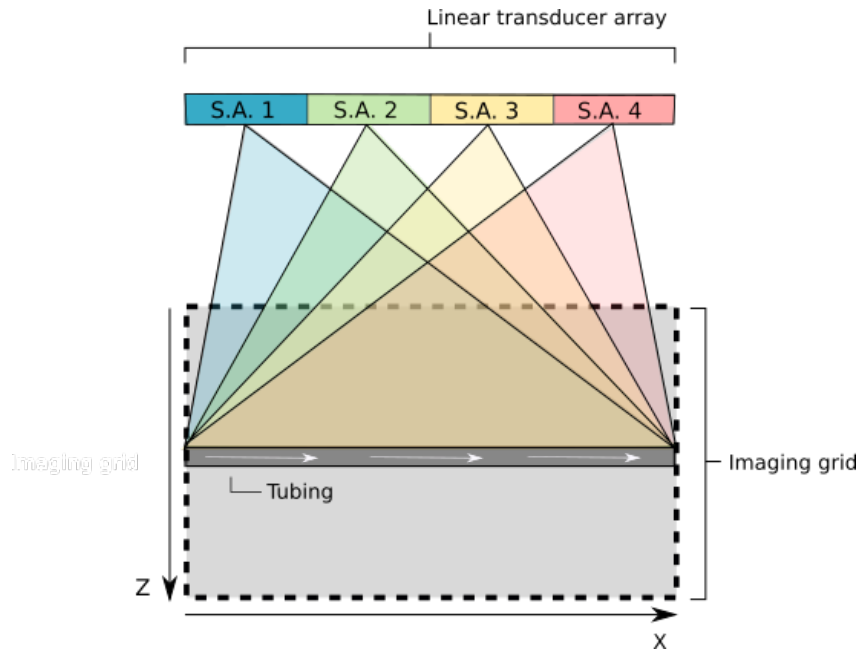


Figure 3.8: This image shows how dividing the linear transducer array into smaller parts called sub-apertures enables the PA signal from different angles to be detected. Each cone shows the field-of-view of each sub-aperture, and it demonstrates the difference in effective angle between the different components.



### 3.7.2 F-number reconstruction

The other reconstruction method for PAV is referred to as the f-number approach. In optics, the f-number is used to describe the ratio between a system's focal length and the size of the aperture. To vary the f-number of our acoustic receiver (probe), the number of elements in the linear array used are varied. In a normal US B-scan, the f-number determines the spatial resolution in the lateral direction. When used to estimate the Doppler shift, it enables the spectral broadening of the Doppler spectrum to be controlled. In the Doppler PA formula, Equation 2.3, it is assumed the US hits the target from a single angle, however, when using a linear transducer array, it really is insonified from a range of angles. If the f-number is increased, the beam is wider at its focus but it is more uni-directional/ collimated, meaning the range of transmitting and receiving angles is reduced. One downside is that as the f-number increases, the number of elements receiving decreases so there is a SNR penalty. In addition, as the f-number is increased there is a decrease in lateral resolution.

To get a number of different phase shifts for each pixel to put into the least-squares equation (Eq.2.15), a range of angles are defined which determines the elements that are used for a pixel's reconstruction. This ensures that the angle between the pixel and the relevant receiver elements is more accurately known. As the f-number is fixed, if a deeper pixel is observed, then the ideal focal length will be longer. To keep the f-number fixed, the aperture needs to be increased, which is done by using more receiver elements in the probe for the reconstruction. A diagram depicting the angles and elements used for reconstruction of a pixel using the f-number reconstruction method is shown in Figure 3.9. The result is an I/Q image stack for every angle used.




Figure 3.9: Diagram depicting the angles used for the f-number reconstruction. The yellow box is the pixel of interest, and the three boxes in the linear transducer are the elements used to reconstruct the pixel from different angles. As the f-number is fixed, there are fewer elements in the central block as it is closer to the pixel than the side blocks.

For PAV, both the sub-aperture and the f-number method will provide I/Q image stacks that can be implemented in the PAV algorithm.

Several parameters must be defined for accurate reconstruction:

- **Transducer parameters:** These include the centre frequency of the probe, the sampling frequency and the pulse repetition rate. The separation between elements in the transducer is defined. This is known as the pitch.
- The speed of sound in the medium is determined using the method described in section 3.1 to find the speed of sound which optimises the clarity of the reconstructed image.
- Properties of the lens (tip of the US probe) are also corrected for. These include the speed of sound here and the thickness of the lens so the acoustic signal before entering the lens for detection can be accommodated.
- **Imaging parameters,** in particular the half-open angle. In Figure 2.1 the maroon arc shows the possible pulse origin given the arrival time of a pulse. The half-open angle defines the angle from an element which should be considered in the reconstruction. In that Figure 2.1, the angle shown is  $180^\circ$ , however for reconstruction I would normally only consider  $60^\circ$ . This is mainly to do with the directionality of the transducer elements as they are more sensitive to US pulses arriving normal to the probe and exhibit side lobe artefacts at larger angles. The greater the angular deviation, the less sensitive the detector is.
- **Definition of imaging grid.** Initially, the spatial position of the transducer elements needs to be defined, then the region of interest and the resolution of the reconstructed image can be defined. The higher the resolution, the better the image quality, but the longer the reconstruction takes. There comes a point where an increase in resolution makes a negligible difference in image quality, so it is about finding a sweet spot which has sufficient resolution and doesn't take too long to compute. As for choosing the depth and width of an image, I would reconstruct the first frame fully and then decide where regions of interest are in the image, and only reconstruct those regions going forward. This corresponds to where the illumination is the best. For the resolution typically used 6 points per wavelength, however when computing power becomes a limiting factor, I would go down to 4.
- **Reconstruction method selection:** This determines whether the sub-aperture or f-number method is used.
- To mimic rotation of the receiver for PAV using the sub-aperture method, the linear receiver is split into small sections, called sub-apertures. The image is reconstructed using only elements in this small section. Each of these sub-apertures looks at the imaging grid from a different angle, so it simulates rotating the receiver. A diagram showing this is shown in Figure 3.8. For the sub-aperture reconstruction, the number of sub-apertures needs to be defined by the operator. I've been using four as it is enough to over-estimate the linear system sufficiently to arrive at a reasonable map without doing too many sub-apertures (angles), which would increase the computation time and decrease the SNR of each sub-aperture's data.
- For f-number reconstruction, a range of angles needs to be defined, normally spread above and below  $\phi = 0$ . Three different angles were chosen:  $-10^\circ, 0^\circ, 10^\circ$ . The minimum f-number for all the different angles needs to be defined. This determines the maximum number of elements used in a given reconstruction depending on the depth of the pixel and the receiving angle used.

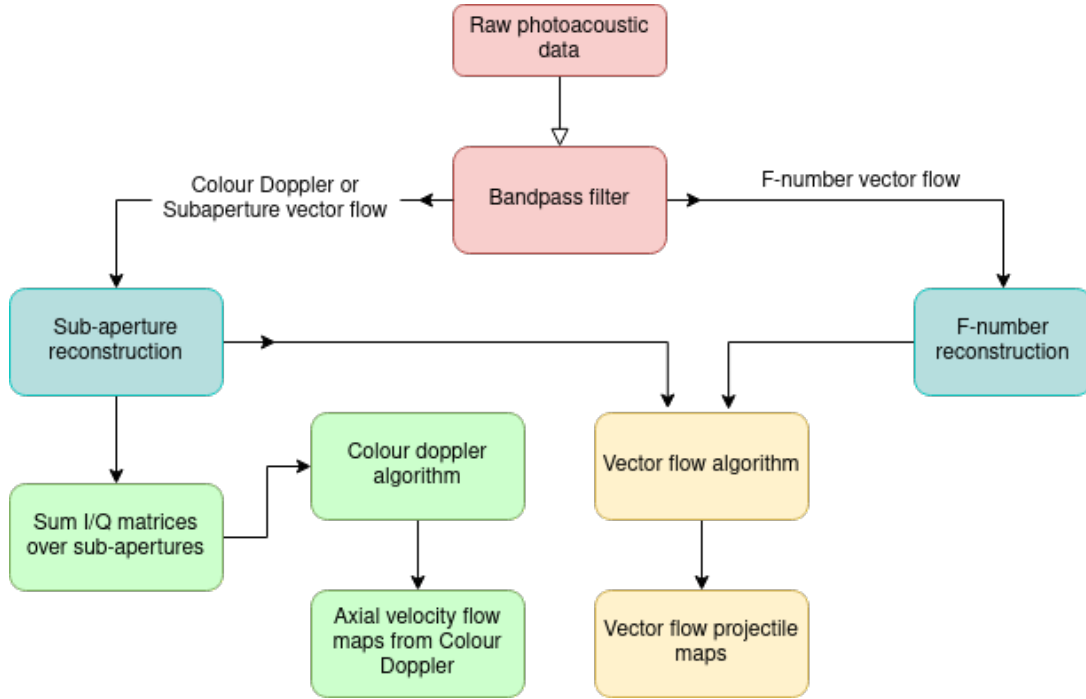


Figure 3.10: Flow chart showing an outline of the data processing pipeline based on the map type/reconstruction method chosen.

The data from each frame is converted into an analytic signal, before it is acted upon by the reconstruction function. This function takes in this I/Q signal data and the parameters above and returns a vast array of information. For each frame, the reconstructed image envelope is returned, as well as the I and Q images for each frame for each sub-aperture/f-number angle. These I/Q images are the product of image reconstruction taking place using the in-phase/quadrature signal instead of the real-valued signal. Another matrix is returned, which shows the angles of the rays at the location of each pixel for every sub-aperture/f-number reconstruction. This is a very memory-heavy computation, so Memory Mapping is used to store data as it is generated to limit stress on the RAM.

### 3.7.3 Colour Doppler computation

For colour Doppler, the I/Q image stack from the sub-aperture reconstruction is taken and summed over each of the different receive sub-apertures, to get stacks of I/Q images utilising all of the receiving elements at one time. Then, SVD is performed on the I/Q data to get filtered I/Q stacks. It was sufficient to remove only one singular value to isolate the flow signal.

Once we have a stack of SVD-filtered I/Q images, the average angular frequency for each pixel must be found. This method looks at the in-phase and quadrature pixel value over slow-time and uses Equation 2.8 to find the average phase shift of the pixel over all of the frames. This is then multiplied by the terms in Equation 2.9 to determine the axial velocity for a given pixel. This is repeated for every pixel to create a map showing axial velocity flow speeds at each pixel in the imaging grid.

### 3.7.4 PAV computation

Determining the PAV field requires the I/Q reconstructed images observed at the different sub-apertures/f-number angles, the angle between each pixel and the centre of the elements used for its reconstruction, the pulse repetition rate, speed of sound and centre frequency. These I/Q

images can be filtered using SVD or not. The phase shift for each pixel is found using Equation 2.8. This uses the I and Q pixel values in slow-time to determine the average phase shift per time step like in the colour Doppler example. This is then multiplied by the PRF to obtain the average angular frequency of the pixel.

Once the average angular frequency of each pixel of a given sub-aperture/f-number is acquired, it is multiplied by  $c/2\pi f_0$  to get the right-hand side of the linear equation, Equation 2.15. The rest of the function follows the methodology in Section 2.3, where the over-determined linear equation is established and solved using data from the different sub-apertures/f-numbers angles used. The result is two  $n_x * n_z$  matrices showing the velocity in the axial and lateral direction for each pixel. To display a velocity vector, the magnitude of the velocity ( $\sqrt{v_x^2 + v_z^2}$ ) is plotted as a heat map, and arrows with length and direction corresponding to the velocity vectors are plotted periodically on top of the heat map.

# Chapter 4

## Results

Throughout the results section, I have chosen to express each pixel’s difference to the ground truth / preset flow speeds as a percentage difference. This is found by subtracting the theoretical value from each pixel’s measured value and expressing it as a percentage of the theoretical value. This way, the error maps show where the flow was under and over-estimated and by how much. By using the relative error instead of the absolute error it is easier to compare errors from different flow speeds.

### 4.1 Simulation results using colour Doppler

For the two simulations, the axial velocity speeds were determined using colour Doppler. To determine if the application of a bandpass filter would be helpful, the colour Doppler axial velocity was determined with and without it. The speeds and average pixel difference from the theoretical speed are summarised in Table 4.1. The high frequency simulation has a smaller

Colour Doppler	Low frequency simulation			High frequency simulation		
	Ground Truth	No BPF	BPF	Ground Truth	No BPF	BPF
Axial velocity speed, mm/s	-2	-0.383	-1.259	-0.4	-0.356	-0.363
% Difference		-80.9	-37.0		-11.1	-9.1

Table 4.1: This table summarises the axial flow speeds found in the simulated data using Colour Doppler. The columns “No BP” and “BPF” corresponded to whether or not the data was passed through the bandpass filter, described earlier, before reconstruction. Across all the combinations, axial velocity is always underestimated by this method. The best accuracy is found in the high frequency simulation where the bandpass filter was applied.

average percentage difference from the theoretical values, such that this simulation’s results were closer to the true values. In both simulations, implementing a bandpass filter increased the measurement accuracy.

The resulting flow maps for both simulations with the bandpass filter are shown in Figure 4.1 below. A mask has been used to only show flow that is in the tube and that is also near the centre of the image, as pixels on the edge of the image tend to be more unreliable. This also aligns well with the pixels which were reconstructed in every sub-aperture (overlapping cones in Figure 3.8).

From Figure 4.1, the high frequency simulation is much more accurate as the flow within the simulated tube is uniform and in one direction, while in the low frequency simulation, the measured flow direction in the tube changes around a lateral distance of 10 mm, which physically makes no sense. As a result, there is a large difference for most of the pixels between their axial

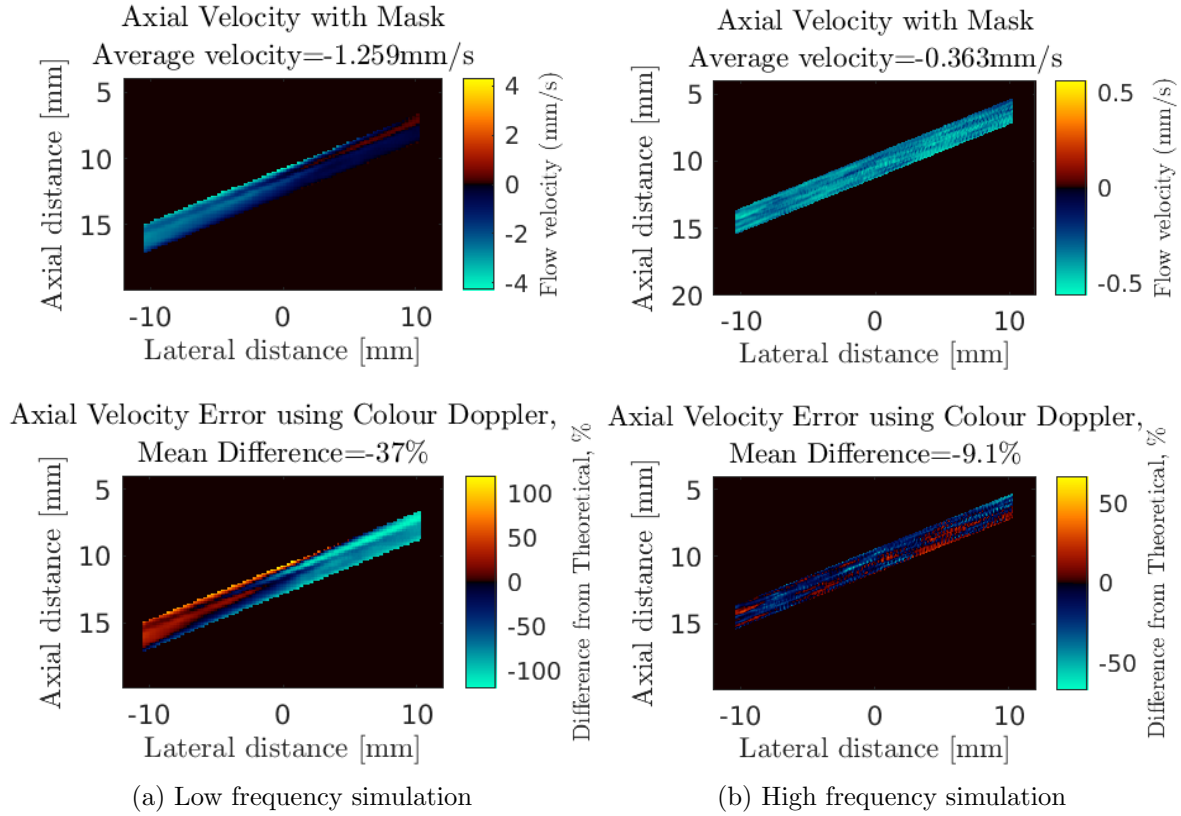


Figure 4.1: Colour Doppler maps and error maps obtained from the simulated data. The maps in the top row show the axial velocity map for each pixel included in the mask, while the maps in the bottom row show percentage difference between each pixel's axial velocity and the theoretical axial speed. Masking eliminates regions outside the tube, on the outskirts of the imaging grid and where flow is not detected.

speeds and the ground truth. On the other hand, the high frequency simulation's differences to the theoretical values are evenly distributed around the tubing, suggesting the variability is primarily caused by noise and nothing systematic.

## 4.2 Simulation results using PAV

### 4.2.1 Low frequency simulation

The PAV results using the reconstruction method with discrete sub-apertures is compared to the reconstruction using a fixed f-number at different angles. The results with and without the bandpass filter applied are also compared. These results are summarised in Table 4.2.

The method with the lowest average difference from the theoretical values is the f-number reconstruction with a bandpass filter, with each component of flow (axial, lateral and total) having an average difference of less than 10%. The resulting flow and error maps for each component of flow processed using the f-number reconstruction on bandpass filtered data are shown in Figure 4.3.

The plots in Figure 4.3 indicate the simulation and subsequent processing is very successful, as the components of flow are uniform within the mask, and the vectors in the bottom left diagram are all parallel to the tube. The differences are evenly distributed above and below the theoretical data indicating this results has minimal systematic errors present. There is

		No Bandpass filter			Bandpass filter	
		Ground Truth	Sub-ap.	F-Num.	Sub-ap.	F-Num.
Axial Velocity	Speed, mm/s	-2	-1.677	-1.461	-1.851	-1.827
	%Difference		-16.1	-26.9	-7.5	-8.7
Lateral Velocity	Speed, mm/s	5	4.002	5.528	4.366	5.188
	%Difference		-19.9	10.6	-12.7	3.7
Velocity Magnitude	Speed, mm/s	5.385	4.349	5.720	4.729	5.503
	%Difference		-19.2	6.2	-11.8	2.2

Table 4.2: This table summarises the average flow speeds obtained using data from the low frequency simulation. The column headings “Sub-ap.” and “F-Num.” refer to the image reconstruction methods used (sub-aperture and f-number respectively). The most accurate results occurred when the f-number reconstruction method was used on bandpass filtered data. The average percentage difference between the theoretical values and those calculated here was less than 10% for the different components.

some banding present in the difference plots which may be due to particles not being evenly distributed in this region of the tubing.

#### 4.2.2 High frequency simulation

The same methodology as was applied to the low frequency simulation above was repeated on the high frequency simulation data set. The results are summarised in Table 4.3 below.

		No Bandpass filter			Bandpass filter	
		Ground Truth	Sub-ap.	F-Num.	Sub-ap.	F-Num.
Axial Velocity	Speed, mm/s	-0.4	-0.534	-0.376	-0.564	-0.377
	%Difference		33.4	-6.1	40.9	-5.7
Lateral Velocity	Speed, mm/s	1	0.245	0.816	0.386	0.819
	%Difference		-65.5	-18.4	-61.4	-18.1
Velocity Magnitude	Speed, mm/s	1.077	0.671	0.901	0.720	0.905
	%Difference		-37.7	-16.4	-33.2	-16.0

Table 4.3: This table summarises the flow speeds obtained using data from the high frequency simulation. The column headings “Sub-ap.” and “F-Num.” refer to the image reconstruction methods used (sub-aperture and f-number respectively). The most accurate results occurred when the f-number reconstruction was applied to bandpass filtered data. The average percentage difference between the theoretical values and those calculated here was less than 20% for the different components.

Once again, the data set processed with the bandpass filter and the f-number reconstruction had the lowest average difference in flow values compared to the ground truth programmed into the simulation. In this case, the bandpass filter doesn’t significantly increase in accuracy unlike the low frequency simulation’s trials. As before, the resulting flow and error maps for each component of flow processed using the f-number reconstruction on bandpass filtered data are shown in Figure 4.2.

The flow maps for the different components of the velocity vector are less uniform in this simulation’s data compared to the low frequency one. The components of flow are, on average, being under-estimated by the PAV algorithm.




Figure 4.2: These plots show the resulting maps from the low-frequency simulation using PAV with the f-number reconstruction method. The left column shows the flow maps for the axial, lateral and magnitude of the velocity components. The bottom left map shows arrows indicating the direction and magnitude of flow at points in the simulated tubing. The right column shows the average percentage difference between the theoretical and calculated flow components.






Figure 4.3: These plots show the resulting maps from the high-frequency simulation using PAV with the f-number reconstruction method. The left column shows the flow maps for the axial, lateral and magnitude of the velocity components. The bottom left map shows arrows indicating the direction and magnitude of flow at points in the simulated tubing. The right column shows the average percentage difference between the theoretical and calculated flow components.

### 4.3 Experimental results

Various combinations of blood-mimicking fluids and probes were tested and it was found that flow was observed in two experiments, both of which used the graphite powder suspension as the blood mimic and the circular fibre output. The only difference is that one used the L11-5V probe and the other the L22-14xvLF probe. The high frequency probe's data set trialled four different flow speeds, while the L11-5V only looked at one.

For the experimental data, I chose not to apply the bandpass filter to simplify the data processing and isolate the results using only the flow algorithms. Furthermore, when using SVD I only removed the first singular value corresponding to the high energy stationary signal. Finally, due to RAM limitations on my computer, only 1000 frames were processed for each speed of the L22-14xvLF data set and 1500 frames for the L11-5V probe's data.

#### 4.3.1 L11-5V probe

Table 4.4 summarises the flow values calculated using colour Doppler and PAV on a data set obtained using the graphite powder blood mimic recorded by the L11-5V probe. Only one flow speed was trialled (1 mm/s).

	Theoretical speed, mm/s	Average pixel velocity, mm/s	% Error
Axial velocity, colour Doppler	0.326	0.157	-49.4
Axial velocity, PAV	0.326	0.350	7.3
Lateral velocity, PAV	-0.946	-1.275	34.8
Velocity magnitude, PAV	1.000	1.362	36.2

Table 4.4: Table summarising the calculated flow speeds at the average pixel error using both colour Doppler and PAV methods. For determining the axial velocity, PAV was more accurate than colour Doppler.

##### 4.3.1.1 Colour Doppler flow maps

To begin with, the axial velocity was calculated using colour Doppler for 1 mm/s flow in the phantom. This flow map with and without a mask applied is shown in Figure 4.4.

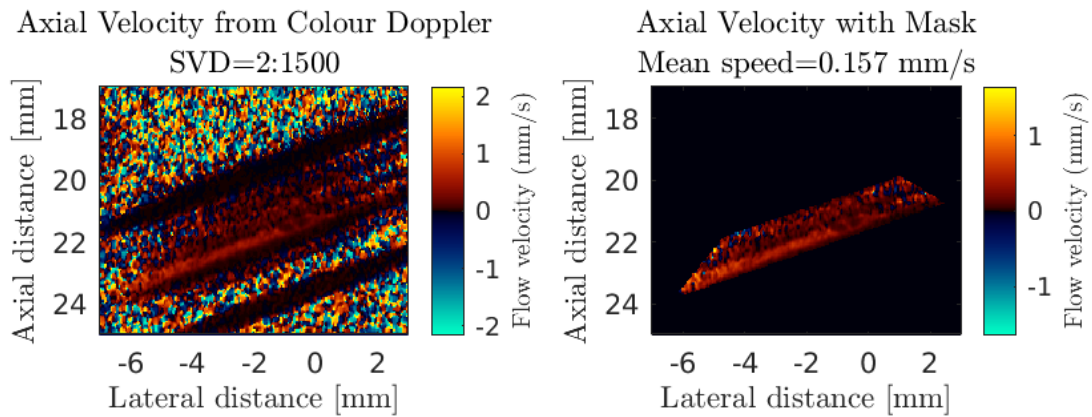


Figure 4.4: Axial velocity map for the 1 mm/s flow rate generated using colour Doppler. Only the first singular value was removed using SVD. The left shows the entire field of view, while the right map is masked to only include flow.

In Figure 4.4, I manually drew the mask based on where the strongest flow results were located. The masked region was selected to eliminate signals from multiples and isolate flow. This mask will be used for the rest of this data set's to ensure the same pixels are studied each time. In the un-masked map in Figure 4.4 there are two strong lines of no flow surrounding the flowing region. Without knowledge of the phantom, I would say these two lines would correspond the stationary signal from the tubing walls. However as the tubing is only  $\sim 1.5$  mm wide, the lowest line corresponds to a reflected signal from the tubing walls arriving later. The flow in the masked region is predominantly in the same direction, but it does vary within the tube, as there appears to be faster flow along the bottom left of the tube.

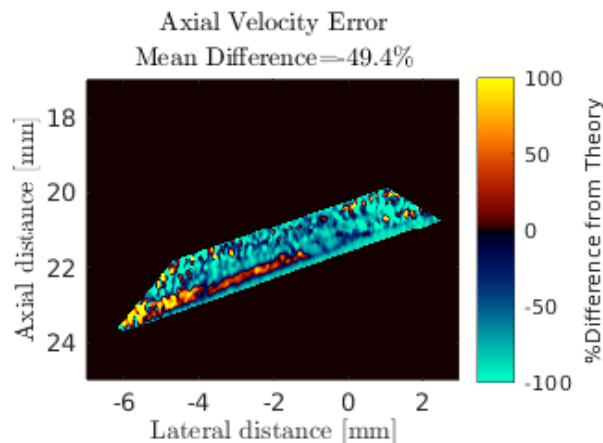


Figure 4.5: Percentage difference between each pixel's calculated axial flow speed and the theoretical value for a programmed flow speed of 1 mm/s. The majority of the pixels have a difference greater -100% to the theoretical values, which is very poor.

Figure 4.5 shows the difference of each pixel from the preset pump speed. The majority of the pixels are underestimated, shown by the large blue regions. Only the fast region in the bottom left of the tube is above the theoretical mean flow speed, suggesting this is where the centre of the tube is, so it is where the fastest flow should be observed.

## 4.3.1.2 PAV

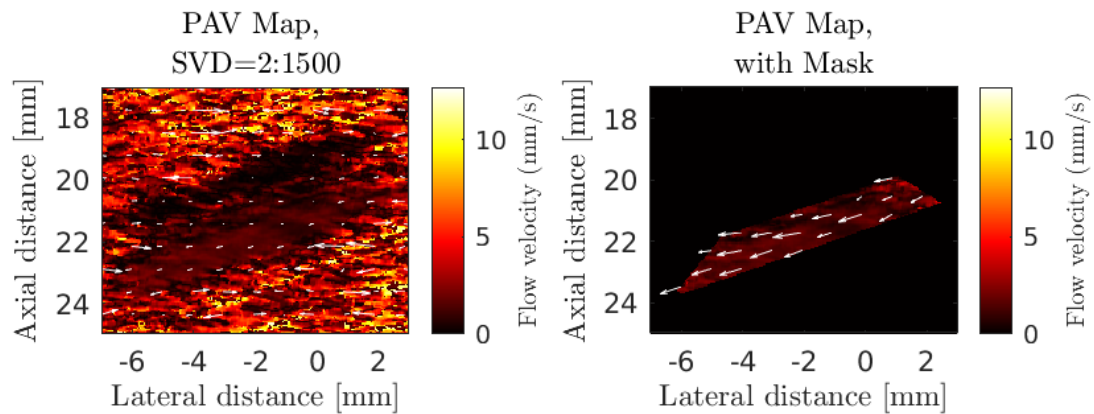


Figure 4.6: PAV maps generated from the 1 mm/s data set using f-number reconstruction. Only the first singular value was removed using SVD. The left diagram shows the PAV of the entire grid, while the right includes a mask to isolate the region where flow is present.

PAV using the f-number method was applied to the same data set as above to obtain the components of flow velocity vector at each pixel. The resulting maps are shown in Figures 4.6, 4.7 and 4.8 for a preset speed of 1 mm/s.

In Figure 4.6 above, the flow region is easier to isolate as the flow region is more uniform than the background. When the mask is applied, the vectors within the tube are mainly found to be parallel to the tubing.

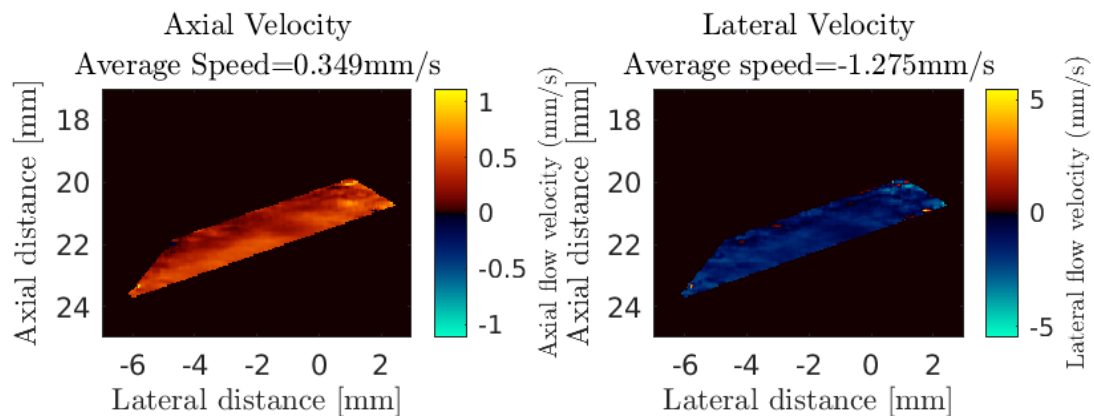


Figure 4.7: PAV maps generated from the 1 mm/s data set using F-number reconstruction. The axial component (left) and lateral component (right) of the flow are shown within the masked region.

In Figure 4.7 the axial and lateral flow components are fairly consistent within the masked region. In both maps, the lower half appears to be flowing faster than the area above, suggesting the centre of the tubing is located here.




Figure 4.8: These three plots show the percentage difference for each pixel between the experimental flow speeds and the theoretical values based on the preset flow speed of 1mm/s. From left to right, these plots show the difference in the axial velocity, the lateral velocity and the magnitude of the velocity vector.

Finally, Figure 4.8 shows the difference in the masked region compared to the theoretical values for the components of the flow velocity vector. In the axial velocity plot, the bottom of the tube is a lot faster than the top, while in the lateral and magnitude plots, the differences from the theoretical values are more evenly distributed throughout the masked region.

### 4.3.2 L22-14xvLF probe

#### 4.3.2.1 Colour Doppler

Initially, the axial velocity was calculated using colour Doppler for the images with the fastest preset flow speed. These maps are shown below in Figures 4.9 and 4.10

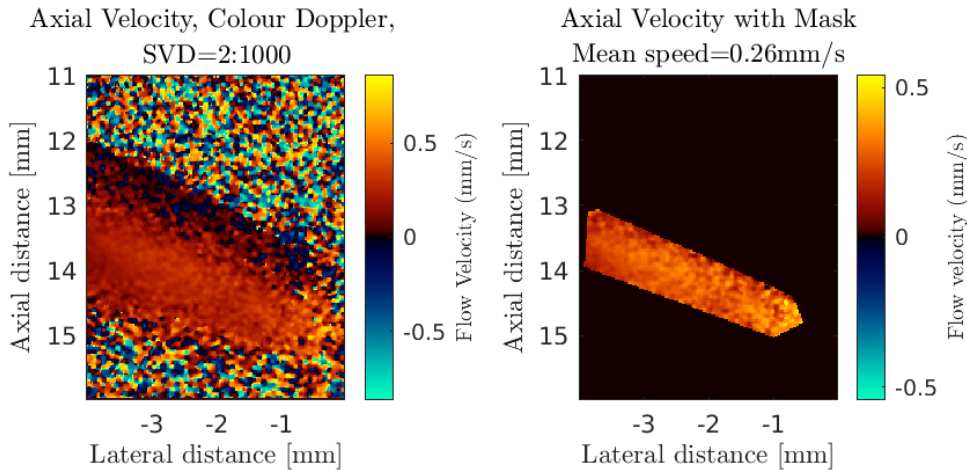


Figure 4.9: Axial velocity map for the 0.887 mm/s flow rate generated using colour Doppler. Only the first singular value was removed using SVD. The left plot shows the entire field of view, while the right map is masked to only include flow.

In Figure 4.9 above, I manually drew the mask based on where the strongest flow results were located. This mask will inevitably exclude regions where the tube is and where the image power was too low to accurately determine the flow. This mask will be used for the rest of the experimental data set to ensure the same pixels are studied in each trial.

The flow in the masked region is all in the same direction, but it does vary within the tube. One interesting observation is that the flow is fastest in the middle of the tube, corresponding to parabolic flow. This parabolic flow was not accounted for when comparing the results to the preset flow speed (i.e the flow speed is assumed to be uniform everywhere). Parabolic flow is expected in these experimental results but absent in the simulated ones. As the particle movement was constant in the simulation, it did not account for interactions between the flowing particles and the walls of the tubing.

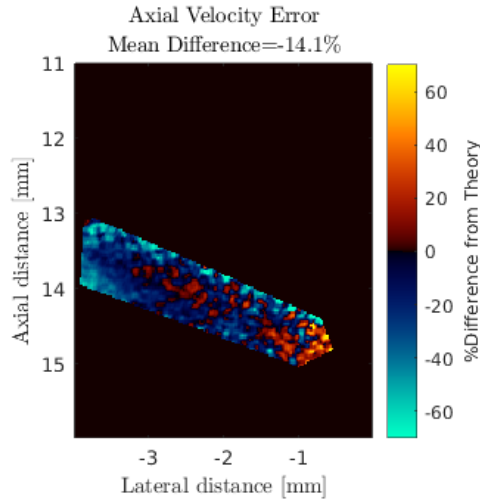


Figure 4.10: Percentage difference between each pixel's calculated axial flow speed and the programmed flow speed of 0.887mm/s.

Figure 4.10 shows the average difference from the theoretical values within the tube. The flow is significantly faster in the bottom right of the tube than at the top. Unlike in the simulated data, these plots are less uniform, which is most likely due to non-uniform illumination.

#### 4.3.2.2 PAV

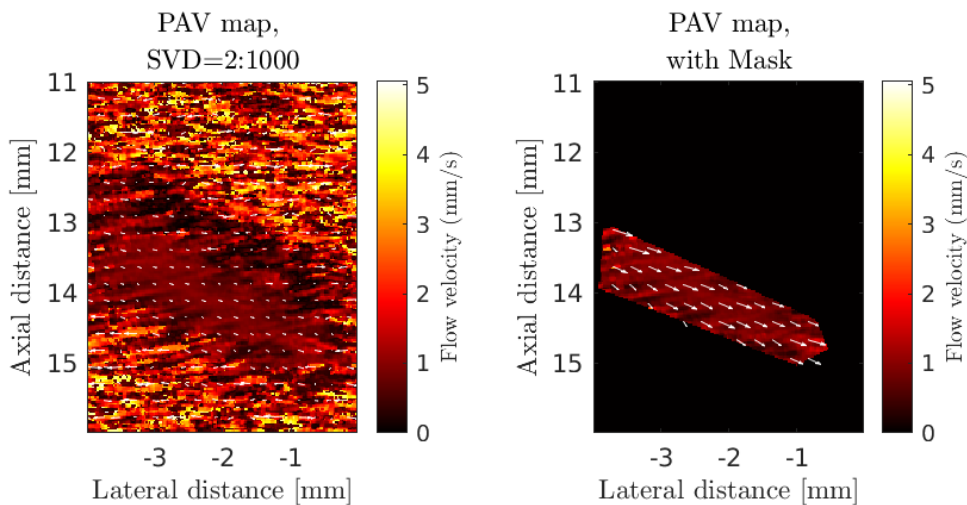


Figure 4.11: PAV maps generated from the 0.887 mm/s data set using f-number reconstruction. Only the first singular was removed using SVD. The left diagram is the PAV of the entire grid, while the right includes a mask to isolate the region where flow is present.

PAV using the f-number method was applied to the data set as above to obtain the vector components of flow. The resulting maps are shown in Figures 4.11, 4.12 and 4.13 for the fastest preset flow speed trialled (0.887 mm/s).

In Figure 4.11 above, the flow region is easily isolated as the flow regions are more uniform than the background. When the mask is applied, the vectors within the tube are parallel to the tubing.

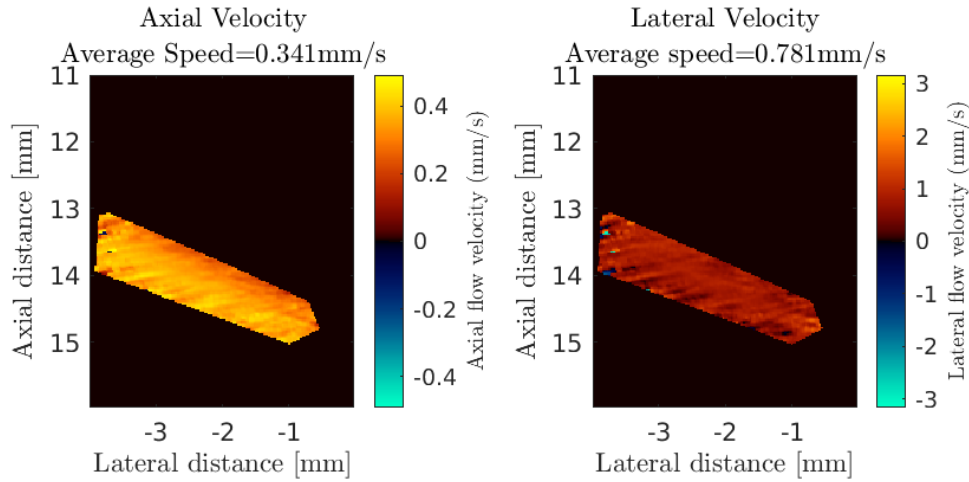


Figure 4.12: PAV maps generated from the 0.887 mm/s data set using f-number reconstruction. The axial component (left) and lateral component (right) of the flow are shown within the masked region.

In Figure 4.12, the axial and lateral flow components are shown within the masked region. The flow is fairly uniform throughout, though the flow is slightly faster near the middle in both subplots.




Figure 4.13: These three plots show the percentage difference for each pixel between the experimental flow speeds and the preset flow speed of 0.887 mm/s. From left to right, these plots show the difference in the axial velocity, the lateral velocity and the magnitude of the velocity vector.

Finally, Figure 4.13 shows the difference in the masked region compared to the theoretical values for the components of the flow velocity vector. In the axial plot, the top of the tubing appears to have slower flow than the bottom. The difference is randomly distributed in the lateral and total flow plots.



### 4.3.2.3 Calculated flow velocity over a range of pump speeds with PAV

Given the speed that the pump was set to and the angle of the tubing is known, the mean values for the different components of flow can be compared to preset flow speeds using the colour Doppler and f-number algorithms. Table 4.5 summarises the mean pixel velocity for the different components and the average pixel difference from the preset flow speed.

		Preset flow rate, mm/s			
		0.222	0.443	0.665	0.887
Axial velocity	Speed, mm/s	0.022	0.092	0.173	0.341
	%Difference	-62.9	-36.9	-22.9	12.0
Lateral velocity	Speed, mm/s	0.027	0.185	0.362	0.781
	%Difference	-83.3	-54.3	-41.4	-6.2
Magnitude of velocity	Speed, mm/s	0.105	0.260	0.445	0.881
	%Difference	-50.6	-40.5	-32.6	-0.7

Table 4.5: Table summarising the average value for the different velocity components and the average percentage difference between each pixel in the mask and the preset flow speed.

As both PAV and colour Doppler methods allow the axial flow speed to be measured, the quality of the two methods can be compared over a range of flow speeds. This is shown in Figure 4.14.

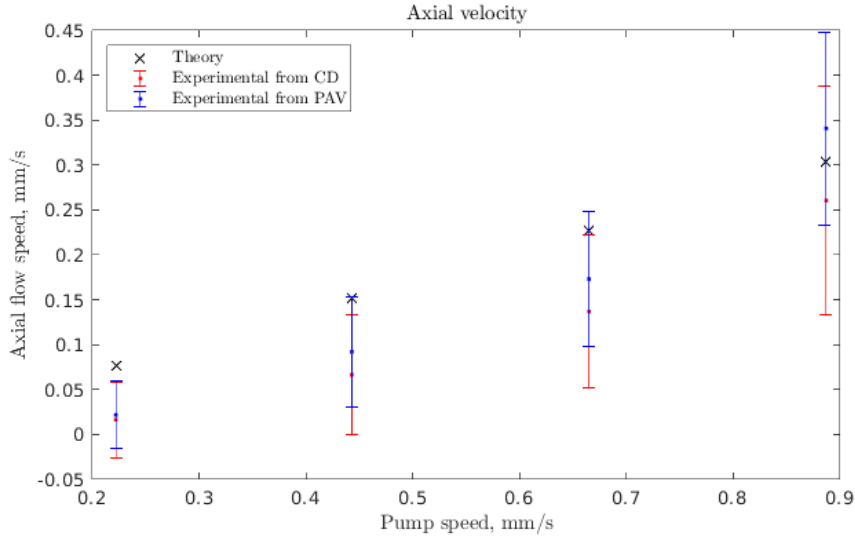


Figure 4.14: Plot comparing the calculated axial flow speeds for four different pump flow rates. The blue points correspond to the axial velocity calculated using f-number PAV, while the red points are the axial velocity speeds found using colour Doppler. The black points show the preset pump speeds. The error bars represent two standard deviations from the mean.

From Figure 4.14 the PAV results are closer to the theoretical values for the axial velocity calculation. For the faster flow rates, the results match the theoretical values much better. Aside from the 0.887 mm/s reading, all of the mean values are below the theoretical values. Finally, it is worth clarifying that negative flow velocity does make sense in this context as it just means the flow could be in the opposite direction. There is also a clear proportional relationship between the preset and measured flow speed.



The same methodology can be applied to the other components from PAV to observe their similarity to the theoretical flow speeds over the four different pump speeds. This result is shown in Figure 4.15.

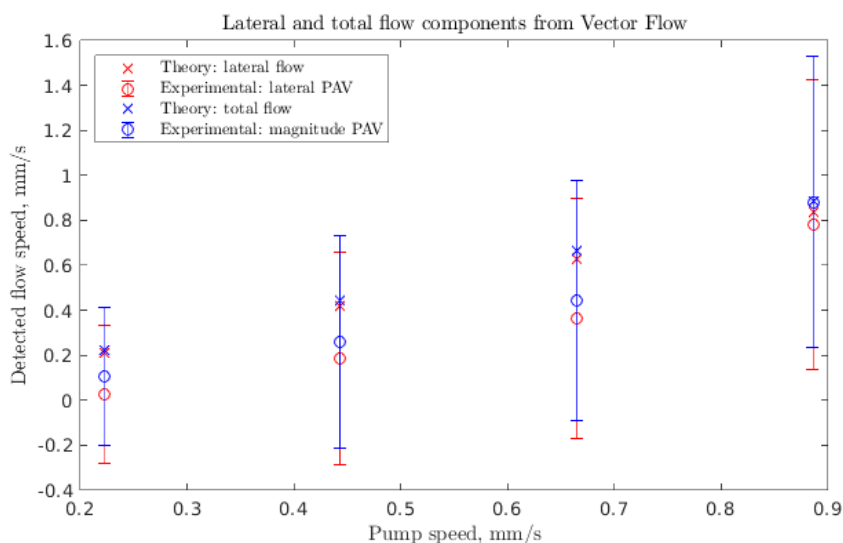


Figure 4.15: Plot comparing the calculated lateral velocity and magnitude of the velocity vector for four different pump flow rates. The red point corresponds to the lateral velocity, while the blue points are the magnitude of the velocity vector components, both calculated using f-number PAV. The crosses corresponded to the preset pump speeds for these components. The error bars represent two standard deviations from the mean.

Unlike the axial velocity plots, the theoretical values are encompassed by the error bars in Figure 4.15. With these results, it is verified that the flow speed is linearly related to the pump speed.



# Chapter 5

## Discussion

### 5.1 Simulation results

With the simulations, it was clear that the most accurate flow speeds were derived when the reconstruction was undertaken using the f-number method compared to the sub-aperture one. This is due to the fact that in the f-number code, the angle from each pixel to the receiver is well defined, whereas in the sub-aperture code, it is taken to be the angle between each pixel and the middle of the sub-aperture used, when in reality this angle spans from the leftmost element in the sub-aperture to the rightmost. Another reason why the flow velocity components derived from the f-number reconstruction are more accurate is because in every f-number orientation, all the pixels in the imaging grid are reconstructed, while in the sub-aperture version, only pixels encapsulated in the cone shape from the sub-aperture's elements are reconstructed. As a result, when it comes to solving the over-determined linear equation, each pixel is included in three different f-number estimates compared to anywhere between one and four times for the sub-aperture variant. The more phase shift estimates for a given pixel, the more accurately the velocity components can be calculated.

The bandpass filter also increases the accuracy as it narrows the frequency spectrum of the detected PA waves. Every flow equation includes the centre frequency  $f_0$  (e.g. Equation 2.8) as a factor, however, the detected PA waves are broadband and do not have a delta-like spectrum at  $f_0$ . By narrowing the spectral content around the centre frequency it ensures the velocity estimate is more accurate.

In the simulations, the mask not only removed regions outside the tubing, it also removed parts of the image  $\pm 1$  cm away from the centre of the image laterally. When a pixel's reconstruction is based on more sensor elements on one side of the pixel than the other, the resulting image is prone to artefacts. By only including pixels near the centre of the image, only pixel's whose reconstructed values are more robust are observed.

For PAV using the same reconstruction methodology, the low frequency simulation was more accurate than the high frequency one. This may be because the low frequency simulation used 200 frames, while the high frequency only had 50 as it was more computationally intensive to run. The more frames used, the longer the I/Q signal for a pixel's value is, so the more accurately the phase shift for each pixel can be determined. Another reason may be that the particle size in the high frequency simulation was a lot smaller than the low frequency simulation, so the total area of the tube occupied by emitters is smaller in the high frequency simulation, meaning the overall density of particles was much lower in the high frequency simulation so there were regions where flow was only present periodically, biasing the result towards a no-flow result.

## 5.2 Experimental results

Discussion around the various bench-top experiments using different blood-mimicking fluids can be found in the Appendix.

Despite running many different experiments, only two showed the desired flow we needed. They both used the graphite suspension, composed of particles between a 2 and 12  $\mu\text{m}$ , which is very similar to the dimensions of a red blood cell, indicating that a great blood-mimicking fluid for our experiments has been found.

The mask removed regions of the tubing that we know contained flow, but it could not be quantified due to poor optical illumination in this area. If the illumination is poor in that region, then there will be minimal absorption of light resulting in a low PA signal from moving particles here. These PA signals will be swamped by stronger stationary signal such that this flow can not be determined. Regions outside tubing shouldn't emit PA waves as it is water, which has a very low optical absorbance compared to the blood mimic used at 680 nm. As a result, the pixel values around the tubing are based on noise, so will vary randomly from frame to frame. When the phase shift is calculated from this noise, it is completely random. Hence, when looking at colour Doppler or velocity vector components outside of the tubing the calculated flow speed is random (like sprinkles).

An exciting observation from the colour Doppler results is the emergence of what appears to be parabolic flow, where the flow speed near the edge of the tubing is slower than in the middle. Unlike how I have calculated the ideal flow speeds (as the average speed everywhere), the flow near the edges of the tube will be slower due to friction with the walls compared to the centre of it. In future trials with more targeted illumination, it could be interesting to see if parabolic flow can be seen along the entire cross-section of the tubing.

The centre frequency used for the flow calculations has a proportional effect on the flow speeds measured. Initially, I was using the probe's prescribed centre frequency, however, when I looked at the spectrum, it was often skewed to one side of the defined centre frequency. This is because, unlike US, the frequency content detected at the transducer depends on the optical emitters, as well as the probe detection bandwidth, not the transmitted wave in US. I then tried using the centroid (centre of mass of the frequency spectrum) before landing on the frequency with this highest amplitude (or the maximum frequency). This was used, as when it was compared to the other two centre frequency options, it gave the closest flow speeds to the theoretical values when applied to both the high and low frequency simulation's data.

In the PAV results with the L11-5V probe, the measured components of the velocity vector were faster than the preset pump flow speeds. This indicates that the centre frequency used was too low, systematically increasing the calculated flow speeds. For this and all trials, the centre frequency used was the frequency with the maximum amplitude when looking at the acoustic spectrum to the detected PA waves. If the spectrum was skewed significantly, then it may have made sense to use the centroid frequency instead. Additionally, a bandpass filter could be applied to narrow the acoustic spectrum around the centre frequency.

For the L22-14xvLF data over different pump speeds, it was always the fastest speed which was calculated with the most accuracy. One reason is that the faster speeds make it easier for SVD to distinguish between stationary and moving signal. Another reason is to do with the pump used for these experiments. At faster speeds, the pump motion was more continuous so I am more confident that the preset flow speeds were actually administered. Generally speaking,

for this probes data set, the calculated flow speeds were lower than they should have been, which could be because some pixels included in the masked area are primarily noise due to lower optical illumination. These pixels will have a more random calculated flow speed in either direction. This will have the effect of decreasing the average flow speed in this area. The closet mean pixel value to the preset flow speed was found for the 0.887mm/s trial using the f-number reconstruction. The average pixel difference between each pixel's calculated flow speed and the pump's preset speed was -0.7% for the magnitude of the velocity vector.

It is difficult to compare the accuracy of the velocity components obtained from the two probes because different pump speeds were used. In the L22-14xvLF trials, the maximum permissible flow speed was used and the calculated velocity components were very accurate, with an average error of -0.7% for the magnitude of the velocity vector. Meanwhile, the L11-5V probe only looked at a flow speed that was half of its maximum determinable flow speed and had a much poorer accuracy, with an average error of 36.2% for the magnitude of the velocity vector. If the accuracy of the two probes at half of their respective maximum measurable flow speeds are compared then the average difference between the two is much more similar (-40.5% for the L22-14xvLF compared to 36.2% for the L11-5V). I would expect the high frequency probe to have better accuracy due to the higher resolution of the reconstructed images, however further investigation is required to see if this truly is the case.

In processing the experimental data I applied the same SVD thresholds to remove only the highest amplitude singular value. This was primarily to keep the processing consistent between the different flow speeds as picking the optimal SVD threshold can be somewhat of an art. An improvement would be to optimise the SVD thresholds to obtain the best flow values possible. This would include removing singular values corresponding to noise.

While my metric for showing the difference between theoretical speeds and those calculated for pixels is good at showing regions where flow is faster/slower than the theoretical values, taking the average is not a very good measure of variance. From many of these percentage difference plots, such as those in Figure 4.13, expressing an average percentage difference is not a good way to express the error as it doesn't give information about the range of pixel variation. Hypothetically, if half of the pixel's flow rates are 100% smaller than the theoretical values and the other half had flow rates 100% faster than the theoretical values, then this metric would say the mean difference is 0%, when in reality, there is a huge variance in pixel speed which is not expressed. A better way to express variance would be to take the mean difference squared for each pixel and average these.

### 5.3 Improvements & future work

In these experiments, the number of frames used was often limited by my computer. Above 1500 frames, functions such as SVD caused Matlab to crash as there wasn't enough RAM. As a result, I only ever used 1000 frames for the L22-14xvLF data and 1500 for the L11-5v. However, if more were used, the SVD thresholds could be tuned to better to isolate flow from stationary signal, as the number of singular values is equal to the number of frames used. Furthermore, by increasing the number of frames, the length of each I/Q signal for a given pixel would increase, resulting in a better estimate for the phase shift using the lag-one auto-correlation method. Eventually, it would be ideal to use Python instead of Matlab for better performance and reliability, however, this would be a significant undertaking as Matlab is the industry standard for ultrasound processing, so it would be time-consuming to convert our codes over.

In the experimental results, I chose not to apply a bandpass filter to show how the PAV

and colour Doppler algorithms perform on unprocessed data. However, in the L11-5v results, the measured flow speeds were higher than the preset pump speeds, which could be due to an inaccurate estimate of the centre frequency. In the future, I would like to try implementing a bandpass filter in the experimental data, which will narrow the spectrum around a centre frequency making the velocity vector measurements more accurate.

When using the different PAV reconstruction codes I used the default parameters, however in the f-number reconstruction, the number of angles and the minimum f-number can be changed. As for the sub-aperture method, the number of sub-apertures can be changed. By changing these parameters, the accuracy for the velocity estimate would be altered. In future studies, it would be beneficial to run a parametric study to optimise the PAV method.

One issue that we had was providing sufficient illumination to the tubing for PA emission from moving particles to occur. However, with a spot size of 5 mm, illuminating the entire tubing region below the probe is impossible. One way to improve the distribution of the illumination would be to use a rectangular fibre output, like the one recently acquired by our research group. This has two rectangular outputs which can be placed on either side of the tubing. This would illuminate a greater area, leading to a greater area in the reconstructed images where the flow can be determined. One challenge in the future will be to improve illumination (area and intensity) while still staying below the recommended laser power levels [1, 2].

While the graphite powder blood-mimic was used to show flow, it would be better to make it more blood-like. This blood mimic was more optically absorbent than real blood so the light can not penetrate as far into the tubing. This is because significant absorption occurs on the outside of the tubing, meaning less light is available to be absorbed near the centre of the tubing, so this region isn't sampled as well by PAV. The mimic was also very thick, making it difficult to pump through such thin tubing. To improve this, I would like to make up different blood mimics by varying the concentration of the graphite powder. By making the mimic less optically absorbent, the light will be able to penetrate deeper into the phantom to illuminate flow that is not on the surface. I would also like to add less sodium polytungstate to make the mimic less viscous and see if the graphite particles do fall out of suspension.

One piece of equipment in our set-up that was unreliable was the pump. We suspect that at low flow speeds (like those used in the L22-14xvLF trials), the pump's motion wasn't constant. In a later experiment with a non-successful blood mimic, we borrowed a more reliable pump (Cetoni Low Pressure Syringe Pump neMESYS 290N) from another lab and were more confident that the preset flow speeds were actually what was administered. In future trials, it would be ideal to use this more reliable pump again with a blood-mimic we know will work.

As I had some difficulty determining where the tubing was, it would be beneficial to overlay PA data with US B-mode images to verify the location of the flow. This will also make it easier to distinguish and remove multiple reflections from the mask used in the PAV maps. Because the probes used can record both PA and US images, the probe and the phantom can remain in the same position for both scan types, so their images line up.

Eventually, we would like to work towards using PAV to measure blood flow in bone. To work towards this goal, it would be great to optimise the set-up for flow speeds similar to those seen in bone, which is similar to the speeds trialled in this thesis. Furthermore, to make the phantom more realistic it could be beneficial to add acoustic and optical scatterers into the water bath to mimic soft-tissue better.

## Chapter 6

# Conclusions

In this investigation I have adapted vector flow methodology from US to create PAV, a modality that utilises a least-squares approach to automatically determine the direction and magnitude of flowing optical absorbers. Data was acquired from simulations and bench-top experiments to validate our approach. Using different reconstruction methods, I determined that synthetically changing the receiving angle using an f-number approach provided the most accurate PAV measurements. When applied to experimental results using a graphite powder suspension travelling at 0.887 mm/s with a high frequency US probe, the average pixel difference from the theoretical value was less than 1% for the magnitude of flow. The magnitude of these errors suggests that clinical applications are worth pursuing.

## Acknowledgements

Firstly I would like to thank Jami Shepherd for all your wisdom and support during this project. I have learnt so much from you over the past year and I'm incredibly excited and grateful to pursue this project into a PhD with you. Also, a shout-out to Cora for the extra motivation to get this thesis written in a timely manner; you are in excellent hands.

I would like to thank Kasper for initially getting me interested in physics research through a summer project two years ago. Without that insight I would never have known what it is to do physics outside of courses and I'm grateful to still be a part of your research group today.

Thanks must go out to my parents for all the love and support over the years. Dad, thank you for always showing an interest in my work and always offering to proofread things for me, no matter how long they may be. And to my mum, thank you for looking after me and supplying me with many knitted items to keep me cosy during later nights studying. I'd also like to acknowledge my sister Vic.

I would also like to thank all the friendly faces I've been so fortunate to interact within the physics department, particularly the folks of 722. Even though the Christmas tree was deemed inappropriate, may festive times be had by all in there. And to Liam, the friendliest face of them all, thank you for all your support and encouragement. I'm incredibly grateful to have you in my life.





# References

- [1] American National Standard for Safe Use of Lasers ANSI Z136. 1–2014. 2014.
- [2] AS/NZS IEC 60825.1:2014 Safety of laser products Part I: Equipment classification and requirements. 2014.
- [3] Amalina Binte Ebrahim Attia, Ghayathri Balasundaram, Mohesh Moothanchery, US Din-  
ish, Renzhe Bi, Vasilis Ntziachristos, and Malini Olivo. A review of clinical photoacoustic  
imaging: Current and future trends. *Photoacoustics*, 16:100144, 2019.
- [4] Jérôme Baranger, Bastien Arnal, Fabienne Perren, Olivier Baud, Mickael Tanter, and  
Charlie Demené. Adaptive spatiotemporal svd clutter filtering for ultrafast doppler imag-  
ing using similarity of spatial singular vectors. *IEEE transactions on medical imaging*,  
37(7):1574–1586, 2018.
- [5] William D Barber, Jeffrey W Eberhard, and Steven G Karr. A new time domain technique  
for velocity measurements using doppler ultrasound. *IEEE Transactions on Biomedical  
Engineering*, (3):213–229, 1985.
- [6] Paul Beard. Biomedical photoacoustic imaging. *Interface Focus*, 1(4):602–631, August  
2011.
- [7] Alexander Graham Bell. On the production and reproduction of sound by light. *Journal  
of the Society of Telegraph Engineers*, 9(34):404–426, 1880.
- [8] J Brunker and P Beard. Pulsed photoacoustic doppler flowmetry using a cross correlation  
method. In *Photons Plus Ultrasound: Imaging and Sensing 2010*, volume 7564, page  
756426. International Society for Optics and Photonics, 2010.
- [9] Joanna Brunker and Paul Beard. Pulsed photoacoustic Doppler flowmetry using time-  
domain cross-correlation: Accuracy, resolution and scalability. *The Journal of the Acous-  
tical Society of America*, 132(3):1780–1791, September 2012.
- [10] Joanna Brunker and Paul Beard. Acoustic resolution photoacoustic Doppler velocimetry  
in blood-mimicking fluids. *Scientific Reports*, 6(1):20902, February 2016.
- [11] Joanna Brunker and Paul Beard. Velocity measurements in whole blood using acoustic  
resolution photoacoustic Doppler. *Biomedical Optics Express*, 7(7):2789, July 2016.
- [12] Thore M. Bücking, Pim J. van den Berg, and Stavroula Balabani. Processing methods for  
photoacoustic Doppler flowmetry with a clinical ultrasound scanner. *Journal of Biomedical  
Optics*, 23(02):1, February 2018.
- [13] Charlie Demené, Thomas Deffieux, Mathieu Pernot, Bruno-Félix Osmanski, Valérie Biran,  
Jean-Luc Gennisson, Lim-Anna Sieu, Antoine Bergel, Stephanie Franqui, Jean-Michel Cor-  
reas, et al. Spatiotemporal clutter filtering of ultrafast ultrasound data highly increases  
doppler and fultrasound sensitivity. *IEEE transactions on medical imaging*, 34(11):2271–  
2285, 2015.

- [14] DH Evans. Colour flow and motion imaging. *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, 224(2):241–253, 2010.
- [15] Hui Fang, Konstantin Maslov, and Lihong V. Wang. Photoacoustic Doppler Effect from Flowing Small Light-Absorbing Particles. *Physical Review Letters*, 99(18):184501, October 2007.
- [16] Hui Fang, Konstantin Maslov, and Lihong V. Wang. Photoacoustic Doppler flow measurement in optically scattering media. *Applied Physics Letters*, 91(26):264103, December 2007.
- [17] Judah Folkman. Angiogenesis in cancer, vascular, rheumatoid and other disease. *Nature medicine*, 1(1):27–30, 1995.
- [18] Amín D Jaskille, Jeffrey W Shupp, Marion H Jordan, and James C Jeng. Critical review of burn depth assessment techniques: Part i. historical review. *Journal of burn care & research*, 30(6):937–947, 2009.
- [19] Jami L. Johnson, Mervyn Merrilees, Jeffrey Shragge, and Kasper van Wijk. All-optical extravascular laser-ultrasound and photoacoustic imaging of calcified atherosclerotic plaque in excised carotid artery. *Photoacoustics*, 9:62–72, March 2018.
- [20] Jing Lu, Nurit Argov-Argaman, Jeni Anggrek, Sjef Boeren, Toon van Hooijdonk, Jacques Vervoort, and Kasper Arthur Hettinga. The protein and lipid composition of the membrane of milk fat globules depends on their size. *Journal of dairy science*, 99(6):4726–4738, 2016.
- [21] MathWorks. FIR Filter Design. <https://au.mathworks.com/help/signal/ug/fir-filter-design.html>.
- [22] Anthony S Podkowa, Michael L Oelze, and Jeffrey A Ketterling. High-frame-rate doppler ultrasound using a repeated transmit sequence. *Applied Sciences*, 8(2):227, 2018.
- [23] Laurent Sandrin, Sébastien Manneville, and Mathias Fink. Ultrafast two-dimensional ultrasonic speckle velocimetry: A tool in flow imaging. *Applied physics letters*, 78(8):1155–1157, 2001.
- [24] Adi Sheinfeld and Avishay Eyal. Photoacoustic thermal diffusion flowmetry. *Biomedical optics express*, 3(4):800–813, 2012.
- [25] Adi Sheinfeld, Sharon Gilead, and Avishay Eyal. Photoacoustic Doppler measurement of flow using tone burst excitation. *Optics Express*, 18(5):4212, March 2010.
- [26] Bradley Treeby, Ben Cox, and Jiri Jaros. k-wave a matlab toolbox for the time domain simulation of acoustic wave fields user manual. *Manual Version 1. 0. 1*, 2012.
- [27] Bradley E Treeby and Ben T Cox. Modeling power law absorption and dispersion for acoustic propagation using the fractional laplacian. *The Journal of the Acoustical Society of America*, 127(5):2741–2748, 2010.
- [28] Bradley E Treeby and Benjamin T Cox. k-wave: Matlab toolbox for the simulation and reconstruction of photoacoustic wave fields. *Journal of biomedical optics*, 15(2):021314, 2010.
- [29] Bradley E Treeby, Jiri Jaros, Alistair P Rendell, and BT Cox. Modeling nonlinear ultrasound propagation in heterogeneous media with power law absorption using ak-space pseudospectral method. *The Journal of the Acoustical Society of America*, 131(6):4324–4336, 2012.

- [30] Bradley E. Treeby, Trond K. Varslot, Edward Z. Zhang, Jan G. Laufer, and Paul C. Beard. Automatic sound speed selection in photoacoustic image reconstruction using an autofocus approach. *Journal of Biomedical Optics*, 16(9):090501, 2011.
- [31] Pim J. van den Berg, Khalid Daoudi, and Wiendelt Steenbergen. Pulsed photoacoustic flow imaging with a handheld system. *Journal of Biomedical Optics*, 21(02):1, February 2016.
- [32] PJ Van Den Berg, Khalid Daoudi, and Wiendelt Steenbergen. Review of photoacoustic flow imaging: its current state and its promises. *Photoacoustics*, 3(3):89–99, 2015.
- [33] Peter Vaupel, Friedrich Kallinowski, and Paul Okunieff. Blood flow, oxygen and nutrient supply, and metabolic microenvironment of human tumors: a review. *Cancer research*, 49(23):6449–6465, 1989.
- [34] Sergey Vilov, Guillaume Godefroy, Bastien Arnal, and Emmanuel Bossy. Photoacoustic fluctuation imaging: theory and application to blood flow imaging. *Optica*, 7(11):1495–1505, 2020.
- [35] Junjie Yao. Transverse flow imaging based on photoacoustic Doppler bandwidth broadening. *Journal of Biomedical Optics*, 15(2):021304, March 2010.
- [36] Junjie Yao, Rebecca C. Gilson, Konstantin I. Maslov, Lidai Wang, and Lihong V. Wang. Calibration-free structured-illumination photoacoustic flowgraphy of transverse flow in scattering media. *Journal of Biomedical Optics*, 19(4):046007, April 2014.
- [37] Junjie Yao, Konstantin I. Maslov, Yunfei Shi, Larry A. Taber, and Lihong V. Wang. In vivo photoacoustic imaging of transverse blood flow by using Doppler broadening of bandwidth. *Optics Letters*, 35(9):1419, May 2010.
- [38] Billy YS Yiu and CH Alfred. Least-squares multi-angle doppler estimators for plane-wave vector flow imaging. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 63(11):1733–1744, 2016.
- [39] Billy Y.S. Yiu, Simon S.M. Lai, and Alfred C.H. Yu. Vector Projectile Imaging: Time-Resolved Dynamic Visualization of Complex Flow Patterns. *Ultrasound in Medicine & Biology*, 40(9):2295–2309, September 2014.
- [40] George York and Yongmin Kim. Ultrasound processing and computing: Review and future directions. *Annual review of biomedical engineering*, 1(1):559–588, 1999.
- [41] Reza Pakdaman Zangabad, Sophinese Iskander-Rizk, Pim van der Meulen, Bram Meijlink, Klazina Kooiman, Tianshi Wang, Antonius FW van der Steen, and Gijs van Soest. Photoacoustic flow velocity imaging based on complex field decorrelation. *Photoacoustics*, page 100256, 2021.



# Appendix A

## Summary of all experiments

### A.1 Other blood-mimicking fluids

#### A.1.1 Blood-mimicking particle suspension

The first fluid we tried was a suspension of red polystyrene microspheres with a  $100\mu\text{m}$  diameter. These are much larger than red blood cells which are approximately  $7\mu\text{m}$  across and  $2\mu\text{m}$  wide. The concentration of these spheres was 5%, which compared to red blood cells in blood is significantly lower (around 45-50%). Even with these “blood” mimicking spheres, they are significantly larger and more dilute than what is actually found in human blood, making them a poor substitute. Despite this shortcoming, the larger size should make it easier to detect flow, while the red colour provides the correct absorption similar to what would be expected in real blood. One study has used these particles further diluted to 3%, which is what will be used as a starting point for these experiments [6].

#### A.1.2 Chocolate milk

Milk is commonly used as a blood-mimicking fluid in US phantoms [23]. Milk consists of fat globules which have a heterogeneous size distribution with diameters ranging from 0.1 to  $15\mu\text{m}$ [20]. For PA studies, chocolate milk seemed like it could be a good blood mimic as a suspension of distinct particles that make it good for US phantoms, but it also has a darker colour resulting in an enhanced optical absorption due to the chocolate flavouring. We also decided to try lime milk as 680 nm corresponds to red light, so green should give better absorption compared to brown chocolate milk. After performing the experiment, we realised it might be the coloured syrup surrounding the milk lipids that absorbs the light and not the globules themselves. If this is the case, then chocolate milk is not a good blood mimic for PA studies as the optical absorber needs to also be the US sound emitter, i.e. the milk globules needed to absorb the light and not the surrounding solution.

#### A.1.3 Indian ink

At the end of my Summer research project in 2020, I tried to observe flow in undiluted Indian ink using SVD, however, flow could not be detected as this fluid was too homogeneous, meaning the individual absorbers were too small and concentrated. More recently, diluted Indian ink ( $\approx 2\%$ ) has been used in flow phantoms with good results by Zangabad *et Al* [41], so it was worth seeing if dilution would help. The downside of Indian ink is that it is not made to a scientific standard; that is, we have no idea about the distribution of absorber sizes within the ink, so we could be mimicking blood cells or something of a very different size. On the other hand, a significant benefit of the diluted Indian ink is how thin and homogeneous the fluid is

as thick or viscous blood mimics tend to get stuck and resist flow in our thin phantom, or need higher pressures from the pump to flow.

## A.2 Discussion of other experimental set-ups

Despite running many different experiments, only two showed the desired flow. They both used the graphite suspension, composed of particles between a 2 and 12  $\mu\text{m}$ , which is very similar to the dimensions of a red blood cell. The polystyrene particles were too big and got stuck in the tubing, while the chocolate milk results weren't reanalysed as we were uncertain if it was the fat globules absorbing the light or the flavoured fluid they are suspended in. When flow maps were generated from the Indian ink data set, they showed no flow in the tubes with the same random "sprinkle-like" pattern on the outside. This means that the signal from inside the tube didn't change between frames due to the movement of the ink. This indicates the fluid was too homogeneous, such that the individual absorbers could not be distinguished.

Finally, there were two experiments that used the L11-5V probe with the graphite powder, but flow could only be quantified in the one where the normal circular fibre output was used and not the rectangular one. As the rectangular fibre delivered less optical energy than the circular one to the phantom, we can conclude that for flow to be detected, there must be sufficient light hitting the moving absorbers. This idea also links to why the mask inevitably crops out part of the tubing, even though there must be moving particles in there. If the illumination is poor in that region, then there will be minimal absorption of light and emission of PA waves from moving particles here. These will be swamped by stronger stationary signal such that this flow can not be determined.

	Blood mimic	Ultrasound probe used	Flow Speeds, mm/s	Orientation	Set-up notes	Feedback
1	100 $\mu\text{m}$ polystyrene microsphere suspension, 3% (v/v) and 2.5%w soap	L11-5V	0.5,1 & 2	Parallel to probe		Microspheres got stuck on tubing walls, need smaller particles or faster flow. Can't test colour Doppler on flow parallel to the probe.
2	Chocolate and lime flavoured milk	L11-5V	0.5,1 & 2	Parallel to probe		Uncertain as to whether fat globules are absorbing light or if the surrounding suspension is. Can't test colour Doppler on flow parallel to the probe
3	15 % Graphite particles (2-12 $\mu\text{m}$ ) in water with sodium polytungstate and Tween (1%)	L11-5V	2	Parallel to probe & slightly angled ( $\approx 20^\circ$ )		Dust in the water bath could be observed moving but this obscured flow in the phantom.
		P4-1	6	Parallel to probe & slightly angled ( $\approx 20^\circ$ )		Poor image resolution with this probe
4		L11-5V*	1 & 2	Parallel to probe & slightly angled ( $\approx 20^\circ$ )	Rectangular fibre used	Poor illumination resulting in low PA signal
5		L22-14vXLF	0.222, 0.443, 0.665 & 0.886	Slightly angled ( $\approx 20^\circ$ )		Uncertain as to whether the pump is flowing continuously or not at low flow speeds. Blood mimic is too thick, unsure if it was actually flowing through the phantom
6	Indian ink (2%) in water	L22-14vXLF	0.5, 1, 1.5 & 2	Slightly angled ( $\approx 20^\circ$ )	New Cetoni pump used, glass capillary tubing	Glass capillary tube has a greater impedance contrast than polymer tubing, many reflections in the received signal.

Table A.1: Summary of experimental trials. As we have made several modifications to the experiments as they have been undertaken, here is a table summarising the different variants.

\* This experiment was repeated by my supervisor using the normal fibre output at a flow speed of 1 mm/s.