

Toward an automated multimodal breast ultrasound imaging system

by

Corey Kelly

B.Sc., Memorial University of Newfoundland, 2009

M.Sc., The University of Guelph, 2011

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate and Postdoctoral Studies

(Biomedical Engineering)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

December 2021

© Corey Kelly 2021

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the dissertation entitled:

Toward an automated multimodal breast ultrasound imaging system

submitted by **Corey Kelly** in partial fulfillment of the requirements for the degree of **Doctor of Philosophy in Biomedical Engineering**.

Examining Committee:

Septimiu E. Salcudean, Professor, Electrical and Computer Engineering, Biomedical Engineering, UBC
Supervisor

Purang Abolmaesumi, Professor, Electrical and Computer Engineering, UBC
Supervisory Committee Member

John Madden, Professor, Electrical and Computer Engineering, UBC
University Examiner

Mu Chiao, Professor, Mechanical Engineering, UBC
University Examiner

Additional Supervisory Committee Members:

Karen Cheung, Professor, Electrical and Computer Engineering, Biomedical Engineering, UBC
Supervisory Committee Member

Lukas Chrostowski, Professor, Electrical and Computer Engineering, UBC
Supervisory Committee Member

Abstract

Breast cancer affects millions of women worldwide each year, and is responsible for hundreds of thousands of deaths. Mammography, the standard screening method for these cancers, underperforms for women with dense breast tissue, who account for half of all women under the age of 50. Given that these women are also at higher risk of developing breast cancer, there is a significant need for better screening methods to serve this population.

In this thesis, we develop a breast imaging platform which combines several modalities that have been clinically shown to aid in the detection and staging of breast malignancies. This operator-independent, completely automated scan, simultaneously acquires B-mode ultrasound, absolute elasticity, Doppler flow, and photoacoustic tomography of the entire breast in 20 minutes. We describe the hardware and software components which comprise each of these imaging subsystems, and conduct a preliminary study testing the combined system by imaging a phantom which we designed to incorporate inclusions which are uniquely visible in either elasticity or photoacoustic imaging.

The photoacoustic tomography system constitutes the most significant contribution, and as such is the primary focus of the thesis. We have designed, built, and tested a custom, fiber-based tissue illuminator to accommodate the unique scanning geometry of the automated breast ultrasound scanner upon which we have based our system. We have also developed a novel data reconstruction scheme which can account for the spatial non-uniformity of this illumination. We tested this system *in vitro*, as well as a purpose-built wire phantom.

Finally, we developed a data processing pipeline which uses generative adversarial networks to improve the signal-to-noise ratio of our raw photoacoustic data, and implemented a state-of-the-art regularized reconstruction scheme to remove imaging and reconstruction artifacts. We tested this method using several phantoms, including an anatomically realistic blood vessel phantom generated from real breast imaging data.

Lay Summary

Mammography is the gold standard for breast cancer screening, but is known to be less effective for women with dense breast tissue, who account for a significant portion of the population, and are at a higher risk of developing breast cancer. Ultrasound imaging can be used to supplement mammography, but has its own shortcomings such as operator dependence, and a high false positive rate, resulting in unnecessary biopsies.

We have developed the first system which can simultaneously map four different types of structural and anatomical information for the entire breast. This scan takes 20 minutes per breast and is painless, fully automated (therefore highly reproducible), and involves no ionizing radiation or contrast agents. The chosen anatomical markers have all been clinically shown to improve breast cancer detection and staging accuracy.

This thesis describes the design and construction of this system, and tests its performance using a variety of imaging phantoms.

Preface

The imaging results in Chapter 2 were presented as a poster, and published in IEEE proceedings as

Corey Kelly, Julio Lobo, Mohammad Honarvar, Yanan Shao, Septimiu Salcudean (2018). An Automated Breast Ultrasound Scanner with Integrated Shear Wave Elastography, Doppler Flow Imaging and Photoacoustic Tomography. In 2018 IEEE International Ultrasonics Symposium (IUS) (Vol. 2018-October, pp. 1–4). IEEE.
<https://doi.org/10.1109/ULTSYM.2018.8580074>

I wrote the code to handle the multiplexed transducer, designed and built the phantoms, and conducted all data acquisition and analysis. I also wrote the manuscript. Dr. Honarvar designed and build the tissue exciter for the SonixEmbrace system. Julio Lobo integrated the shear wave absolute vibroelastography (S-WAVE) software into the SonixEmbrace system, and assisted with debugging. Dr. Shao implemented the scan conversion code for the SonixEmbrace transducer. Dr. Salcudean helped with revisions.

A version of Chapter 3 was initially presented orally, and published in the proceedings of SPIE as

Corey Kelly, Septimiu Salcudean (2018). Automated full-breast photoacoustic tomography with non-uniform illumination. Photons Plus Ultrasound: Imaging and Sensing 2018, 1049405(February), 5. <https://doi.org/10.1117/12.2292274>

then expanded, revised, and published in its final form as

Corey Kelly, Amir Refaee, Septimiu Salcudean (2020). Integrating photoacoustic tomography into a multimodal automated breast ultrasound scanner. Journal of Biomedical Optics, 25(11), 1–18. <https://doi.org/10.1117/1.JBO.25.11.116010>

I designed and built the illumination system reported in the paper, and pioneered the idea of including a correction for non-uniform illumination

in the reconstruction. I wrote the reconstruction and analysis code, and performed all data acquisition and analysis. Amir Refaee ported sections of my Python code to C++/CUDA to accelerate image reconstruction. Dr. Salcudean advised throughout and helped edit the manuscript.

Chapter 4 has been published as

Amir Refaee*, Corey Kelly*, Hamid Moradi, Septimiu Salcudean (2021). Denoising of pre-beamformed photoacoustic data using generative adversarial networks. *Biomedical Optics Express*, 12(10), 6184. (*equal authorship contribution)
<https://doi.org/10.1364/BOE.431997>

My prior experimental work motivated this paper, and I developed the gold standard technique this study was compared against. I designed the experiments, acquired the data, and performed the data analysis. Amir Refaee and I worked together on formulating the deep learning method used in the paper and contributed equally to the writing. Amir implemented the network and the iterative reconstruction solver. Drs. Moradi and Salcudean provided guidance on the overall structure of the manuscript, and helped with editing.

Table of Contents

Abstract	iii
Lay Summary	iv
Preface	v
Table of Contents	vii
List of Tables	x
List of Figures	xi
List of Abbreviations	xiii
Acknowledgements	xvi
Dedication	xvii
1 Introduction and Motivation	1
1.1 Motivation	1
1.2 Ultrasound Imaging of the Breast	1
1.3 Automated Breast Ultrasound Scanners	2
1.3.1 Semi-automated	2
1.3.2 Paddle Style	4
1.3.3 Ultrasound computed tomography (USCT) systems	4
1.3.4 Dome Style	6
1.3.5 Outstanding Issues	7
1.4 Photoacoustic Tomography	9
1.4.1 Overview	9
1.4.2 Illumination Methods	11
1.4.3 Theory and Reconstruction	12
1.4.4 Denoising Photoacoustic Data	14
1.5 Elastography	14

Table of Contents

1.6	Contributions and Thesis Overview	16
2	System Overview and Preliminary Results	19
2.1	SWAVE System	19
2.2	Photoacoustic System	21
2.2.1	SonixEmbrace Transducer	21
2.2.2	Light Delivery	23
2.2.3	Hardware Synchronization	28
2.2.4	Control Software	28
2.3	Summary	31
2.4	Preliminary Study	34
2.4.1	Methods	34
2.4.2	Results	38
2.4.3	Discussion	42
2.4.4	Conclusions	43
3	Integrating Photoacoustic Tomography into a Multimodal Automated Breast Ultrasound Scanner	44
3.1	Overview	44
3.2	Introduction	44
3.3	Methods	47
3.3.1	The SonixEmbrace ABUS	47
3.3.2	Laser and Optics	48
3.3.3	Data Acquisition and Pre-processing	51
3.3.4	Reconstruction	52
3.3.5	Illuminator Fluence Measurement	57
3.3.6	Reconstruction Implementation	58
3.3.7	Phantoms	59
3.4	Results	61
3.4.1	SVD Denoising	61
3.4.2	Transducer Directivity	63
3.4.3	Combined Imaging	63
3.4.4	Illuminator Spatial Profile	66
3.4.5	Reconstruction Comparison	66
3.5	Discussion	72
3.6	Conclusions	74
4	Denoising of Pre-beamformed Photoacoustic Data Using Gen- erative Adversarial Networks	76
4.1	Overview	76

Table of Contents

4.2	Methods	77
4.2.1	Photoacoustic Model	78
4.2.2	Reference RF Dataset	78
4.2.3	GAN Denoising	79
4.2.4	Reconstructions	87
4.2.5	Metrics	89
4.2.6	Data Acquisition	90
4.3	Results	91
4.3.1	Model Hyperparameter Tuning	95
4.3.2	RF Data	95
4.3.3	Reconstructed Images	95
4.4	Discussion	99
4.5	Conclusion	103
5	Conclusions and Future Work	104
5.1	Summary and Discussion	104
5.2	Implications	105
5.3	Limitations and Future Work	106
	Bibliography	109
 Appendices		
A	Supporting Information	131
A.1	Denoising methods comparison	131
A.2	GAN reference dataset - number of averaged frames	133
A.3	Number of FISTA iterations	134
B	ABUS illuminator geometric optics calculations	135
B.1	Light Propagation From Optics to Dome	135
B.2	Reflective Losses	136
C	Photoacoustic reconstruction kernel implementation	138
D	Photoacoustic trigger box code	141

List of Tables

1.1	ABUS systems comparison	8
3.1	Registration errors between imaging modalities.	66
3.2	Reconstruction comparison on chicken breast phantom	68
4.1	Knot phantom RF denoising comparison	96
4.2	Vessel phantom RF denoising comparison	96
4.3	Knot phantom reconstruction comparison	99

List of Figures

1.1	ABUS systems comparison part 1	3
1.2	ABUS systems comparison part 2	5
1.3	SonixEmbrace B-mode imaging	7
1.4	Photoacoustic imaging schematic	10
1.5	Clinical breast PAT systems compared	13
1.6	Elastography schematic	15
2.1	SWAVE shaker photos and schematics	20
2.2	SonixEmbrace transducer	22
2.3	Photoacoustic resolution limits of our system	24
2.4	SonixEmbrace imaging dome photos	25
2.5	SonixEmbrace imaging dome modifications	27
2.6	Synchronization circuit photos	29
2.7	Triggering algorithm schematic	30
2.8	Multimodal system block diagram	32
2.9	Custom system software screenshot	33
2.10	CIRS breast elasticity phantom	35
2.11	ABUS Doppler flow imaging experiment	36
2.12	B-mode and elasticity of CIRS breast phantom	39
2.13	ABUS Doppler flow imaging results	40
2.14	Custom multimodal phantom photo and imaging results	41
3.1	SonixEmbrace system photo	47
3.2	ABUS illuminator schematics	49
3.3	PAT system block diagram	50
3.4	Transducer directivity experiment	55
3.5	Illuminator coordinate system	57
3.6	Reconstruction method schematic	58
3.7	Calibration wire phantom photos	60
3.8	SVD denoising example	62
3.9	Transducer directivity data	64
3.10	B-mode and PAT of calibration wire phantom	65

List of Figures

3.11	Illuminator spatial profile estimate and measurement	67
3.12	Reconstruction technique comparison on chicken breast phantom	69
3.13	Tomographic reconstruction of calibration wire phantom	70
3.14	Reconstruction schemes compared in terms of image error vs. depth	71
4.1	Reference data generation process	79
4.2	Example RF data frame pair for GAN training	80
4.3	SVD denoising example	81
4.4	Example RF data patch pair for GAN training	83
4.5	GAN network architecture	84
4.6	GAN training process schematic	85
4.7	GAN testing process schematic	86
4.8	Custom phantoms used for GAN training	92
4.9	Vessel phantom used for GAN testing	93
4.10	GAN validation metrics and epoch choice	94
4.11	Knot phantom RF denoising comparison	97
4.12	Vessel phantom RF denoising comparison	98
4.13	Knot phantom filtered backprojection reconstruction comparison	100
4.14	Knot phantom regularized reconstruction comparison	101
A.1	Data denoising method comparison	132
A.2	RF quality vs. number of averaged frames	133
A.3	Iterative reconstruction minimization objective	134

List of Abbreviations

2D	two-dimensional
3D	three-dimensional
ABUS	automated breast ultrasound scanner
ABVS	automated breast volume scanner
ARFI	acoustic radiation force impulse
AUC	area under the curve
AWBU	automated whole-breast ultrasound
CAD	computer-aided design
cGANs	conditional generative adversarial networks
CT	X-ray computed tomography
DAS	delay-and-sum
DCIS	ductal carcinoma in situ
EIR	electrical impulse response
FBP	filtered backprojection
FDA	Food and Drug Administration
FISTA	fast iterative shrinkage-thresholding algorithm
FSIM	feature similarity index measurement
FWHM	full width at half maximum
GANs	generative adversarial networks
GPU	graphics processing unit
GUI	graphical user interface

List of Abbreviations

Hb	deoxyhemoglobin
HbO₂	oxyhemoglobin
HHUS	handheld ultrasound
IDC	invasive ductal carcinoma
MRI	magnetic resonance imaging
MSE	mean square error
OPO	optical parametric oscillator
PA	photoacoustic
PAI	photoacoustic imaging
PAT	photoacoustic tomography
PSNR	peak signal-to-noise ratio
QPAT	quantitative photoacoustic tomography
RF	radio-frequency
RMSE	root mean square error
RTE	radiative transfer equation
S-WAVE	shear wave absolute vibroelastography
SDK	software development kit
SIR	spatial impulse response
SNR	signal-to-noise ratio
SO₂	blood oxygen saturation
SSIM	structural similarity index measurement
SVD	singular value decomposition
SWE	shear wave elastography
TTL	transistor–transistor logic
TV	total variation

List of Abbreviations

US	ultrasound
USB	universal serial bus
USCT	ultrasound computed tomography
XRM	X-ray mammography

Acknowledgements

This thesis would not have been possible without the support of my advisor, Tim Salcudean. Thank you for your guidance and endless patience. I also wish to extend my deepest thanks to my thesis committee for lending me their expertise and wisdom, and helping to point me in the right direction on many occasions.

I extend my gratitude to NSERC and the Walter C Sumner Foundation for providing financial support for this work.

To my colleagues who supported me and this work through discussion, moral support, or contributions to the research — our collaborations were the best times I had in the lab, and I hope we'll get to work together again soon!

Everyone who was there to celebrate the wins, and take the edge off the losses — my friends and family, my kitty, and my bird — I truly couldn't have done this without you, and I love you all.

For Vern. I know you would have read the whole thing.

Chapter 1

Introduction and Motivation

1.1 Motivation

In 2020, female breast cancer surpassed lung cancer as the most diagnosed cancer globally, accounting for 2.2 million new cases, and a total of 689 996 deaths [1]. While breast cancer morbidity is dropping, incidence is increasing as the global population ages, and higher quality health care and screening standards become more accessible [2]. For this reason, diagnostic imaging of the breast has remained an active area of research since the introduction of the mammogram 40 years ago [3].

X-ray mammography (XRM) remains the gold standard for the detection of breast cancer. While the technique has an overall sensitivity of up to 90%, this is highly variable across populations, and can be as low as 30% for women with radiographically dense breast tissue, who account for one half of women under the age of 50, and one third of women over the age of 50 [4,5]. Women with dense breast tissue are also at higher risk of developing breast cancer. Combined, these factors result in up to 17.8 times higher incidence of interval cancers (those detected less than 12 months after a negative mammogram) when compared to women with predominantly fatty breast tissue [6].

1.2 Ultrasound Imaging of the Breast

Ultrasound (US) imaging serves as a supplement to XRM for peripheral sites, as a “second-look” modality for identifying cysts and distinguishing benign and malignant solid masses [7,8], and for screening women with dense tissue and other risk factors [9,10]. While US avoids the ionizing radiation and painful compression of XRM, it is known to suffer from a high false positive rate, and a significant operator variability [7,11,12], preventing its widespread use as a primary screening modality.

When classifying tissue as benign or malignant, clinicians are accustomed to having XRM images available to visualize amorphous calcifica-

tions [13]. While these are not necessarily visible in standard B-mode ultrasound, they can be detected using state-of-the-art transducers and reconstruction algorithms [10, 14], or by using more sophisticated modalities of ultrasound such as colour Doppler [15]. Another key distinguishing factor of malignancies is that they are typically stiffer than benign lesions [16]. Ultrasound-based tissue stiffness measurements using techniques such as shear wave elastography (SWE) (most commonly using acoustic radiation force impulse (ARFI)) have shown significant promise in replicating and supplementing palpation for this purpose [17, 18, 19, 20]. Finally, there is increasing consensus that multimodal ultrasound-based imaging is both feasible and clinically relevant, with many combinations of imaging modalities being explored: US/photoacoustic [21, 22, 23, 24], US/Doppler [25], photoacoustic/elastography [26], US/elastography [27, 28, 29], US/elastography/Doppler [30, 31], and US/photoacoustic/elastography [32].

1.3 Automated Breast Ultrasound Scanners

Automated ultrasound scanning of the breast has become a promising alternative to handheld ultrasound (HHUS), especially with respect to the issues of operator dependence and poor repeatability [33, 34, 35, 36]. To explore this potential, several automated breast ultrasound scanners (ABUSs) systems have become commercially available. Figure 1.1, Figure 1.2, and Table 1.1 summarize these systems, with a focus on features relevant to this thesis. In the following sections we will briefly discuss each of these systems, highlighting significant clinical results. We have broadly classified the systems based on the scanning geometry and patient position.

1.3.1 Semi-automated

SonoCiné

The SonoCiné automated whole-breast ultrasound (AWBU) (SonoCiné, Reno, NV, USA) is unique in that it can be used with any HHUS system and transducer. The system consists of a transducer mount guided by a robotic arm to obtain parallel image slices in several overlapping scans, while a technologist holds the transducer, ensuring proper placement at the start of each scan. These images can then be played back in a loop (or ciné) for review. Figure 1.1 shows the SonoCiné system, as well as the scanning procedure.

The SonoCiné system was the focus of an early clinical trial investigating the diagnostic value of ABUS scanning as an adjunct to XRM for women

1.3. Automated Breast Ultrasound Scanners

Hybrid System



Paddle Style Systems



Figure 1.1: An overview of commercially available automated breast ultrasound scanners. Further details are provided in Table 1.1 and in the text. Continued in Figure 1.2

with dense breast tissue, or other risk factors. The study included 4419 women, and found that ABUS provided an increase in diagnostic yield from 3.6 per 1000 to 7.2 per 1000 compared to XRM alone [37]. While this improvement is similar to that achieved with HHUS, the authors highlight that the other benefits of ABUS (reproducibility, flexibility in review time) make it the preferred technique.

1.3.2 Paddle Style

GE Invenia ABUS

Several clinical studies were performed using the SomoVu Automated Breast US System by U-Systems (a subsidiary of GE Healthcare) [38, 39] which scans the patient in the supine position using a paddle transducer with a bandwidth of 14 MHz to 16 MHz. This system has since been superceded by the GE Invenia ABUS, which includes a patented Reverse Curve transducer and automated compression built into the scan head. Figure 1.1 shows the system footprint, as well as the scan head and scanning position.

U-Systems conducted a pivotal retrospective reader clinical study in 2012 using this system, exploring the use of ABUS in conjunction with XRM for the detection of cancers in women with dense breasts. Examining 164 total cases, consisting of 133 non-cancers and 31 biopsy-confirmed cancers, the study found the addition of ABUS significant improved overall sensitivity from 38.8% to 63.1%, with no statistically significant decrease in specificity (78.0% to 76.0%) [40]. This result paved the way for Food and Drug Administration (FDA) approval of the Invenia system in 2013.

Siemens Acuson S2000

The Acuson S2000 automated breast volume scanner (ABVS) from Siemens Healthcare features a 5 MHz to 14 MHz paddle style transducer on an arm which is lightly compressed against the breast with the patient in the supine position. A single scan acquires up to 448 axial slices, and the workstation uses multiplanar reconstruction in post-processing to generate coronal and sagittal views. By conducting additional HHUS scans, this system can also perform ARFI elastography to measure tissue elasticity [17]. Figure 1.1 shows the system footprint, as well as the scanning position.

The Acuson system is the most widely used of the systems discussed here [11, 13, 41, 42, 43]. A meta-analysis conducted by Zhang *et al.* of 1376 patients over nine studies found that compared to HHUS, ABUS improves diagnostic sensitivity from 90% to 93% and specificity from 82% to 86% [34]. Of the 22 studies included in all of their analyses, 20 were conducted using the Acuson system.

1.3.3 Ultrasound computed tomography (USCT) systems

Table 1.1 also contains two commercially available USCT systems, the SoftVue (Delphinus Medical, Novi, MI, USA) and the QTscan (QT Imaging,

1.3. Automated Breast Ultrasound Scanners

Ultrasound Computed Tomography Systems



Dome Style Systems

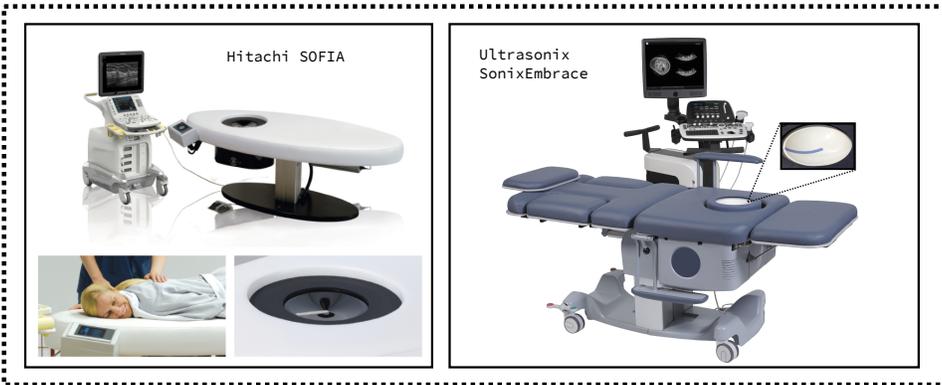


Figure 1.2: Continued from Figure 1.1. An overview of commercially available automated breast ultrasound scanners. Further details are provided in Table 1.1 and in the text.

Novato, CA, USA). These scanners are also shown in Figure 1.2. In contrast to standard US imaging which measures only reflected acoustic waves, USCT also measures transmitted waves. Through inversion algorithms similar to X-ray computed tomography (CT), these measurements can be used to reconstruct the speed of sound and acoustic attenuation throughout the breast. Volumetric speed of sound measurement has been shown to be a good surrogate for breast density [44, 45], and the combination of speed of sound and attenuation can be used to infer relative tissue stiffness [46].

While distinct from ABUS, USCT systems constitute a significant portion of the literature regarding *US-based* automated scanning of the breast, and are similar to dome-style ABUS systems in their prone scanning geometry and concave transducers.

1.3.4 Dome Style

Hitachi SOFIA

The SOFIA system was developed by iVu Imaging (Southlake, TX, USA), and is now distributed by Hitachi Healthcare Americas (Twinsburg, OH, USA). The scanner consists of a platform with an inset scanning interface, containing the transducer. The patient lies prone with one breast against the transducer, which rotates as it acquires US images. The scan takes 30 seconds, and produces volumetric ultrasound data of the entire breast.

Figure 1.2 shows the SOFIA system, illustrating the prone scanning position and the transducer area.

A pilot study of 63 patients using this system achieved similar lesion detection capabilities to HHUS, but resulted in 16.67% more unnecessary second-look recalls [47]. The authors of the study also report issues breast sizes of a D cup or larger, as they exceeded the size of the scan area, resulting in incomplete image data.

SonixEmbrace ABUS

This thesis will focus on the the SonixEmbrace automated breast ultrasound scanner (Ultrasonix Medical Corporation, Richmond, BC, Canada). The SonixEmbrace consists of a large, concave ultrasound transducer embedded in a spherical dome. The patient lies prone on a padded platform with one breast lying against a gel-based custom coupling pad within the imaging dome. A motor rotates the dome through 360°, acquiring B-mode images at 0.5° intervals, which are then interpolated to produce volumetric ultrasound [48]. A photo of the this system is included in Figure 1.2. The SonixEmbrace is similar to the SOFIA system in that the scan is performed with the patient in the prone position using a rotating transducer. The SOFIA system, however, uses a linear transducer embedded in a cone-shaped scanning surface, and also is designed to image the breast while compressed by the body weight of the patient. The SonixEmbrace system has several advantages outlined by Zahiri Azar *et al* [49], including high lateral resolution due to the concave transducer array design, a fixed scanning volume that facilitates registration to other imaging modalities such as magnetic resonance imaging (MRI), and a rapid volumetric acquisition time of 2 min. The breast does not move nor is it deformed during the acquisition, and therefore its volumetric image can be re-sliced for analysis and diagnosis. Thus operator-dependence in image acquisition is effectively removed, and diagnosis by a qualified radiologist can be facilitated. An example of

1.3. Automated Breast Ultrasound Scanners

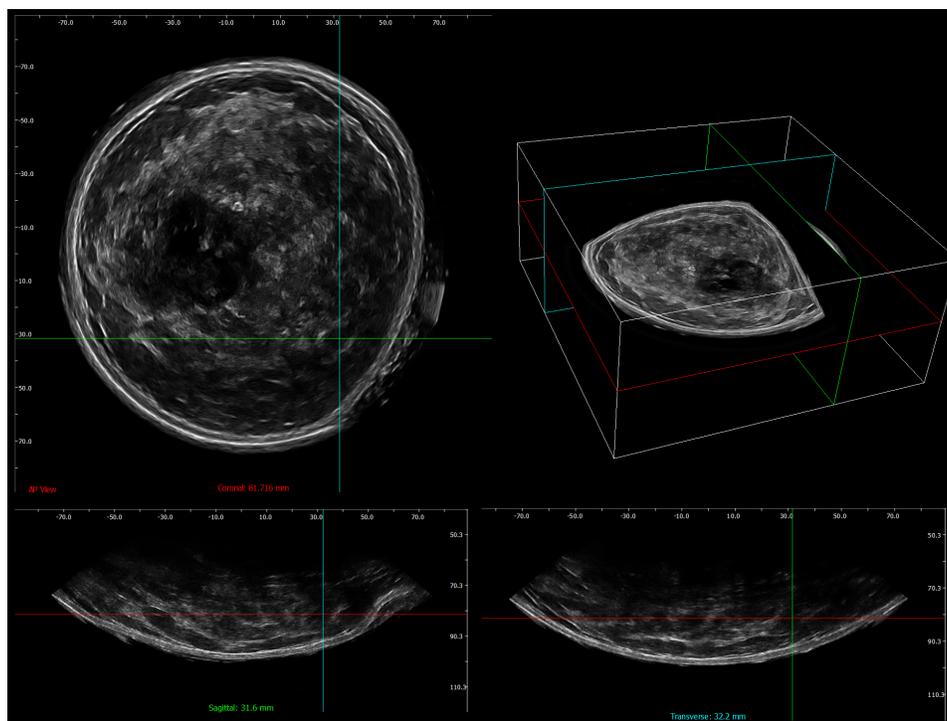


Figure 1.3: Volumetric B-mode US data of a healthy volunteer acquired with the SonixEmbrace ABUS, demonstrating the viewer software and slicing capability.

volumetric B-mode data acquired with the SonixEmbrace system is shown in Figure 1.3.

Early results with this system showed promise in the detection of microcalcifications in the breast [14], which serve as an early warning sign of malignancy, especially in cases of ductal carcinoma in situ (DCIS) [50, 51].

1.3.5 Outstanding Issues

While ABUS systems have been generally seen as valuable diagnostic tools [11, 13, 38], there have been several shortcomings preventing their widespread usage.

One concern is that in contrast to HHUS imaging, there currently exists no standardized approach to interpreting ABUS image data [41]. Solving this problem is a matter of further clinical study with the the express pur-

Device Type	System	Manufacturer	Patient position	Transducer Type	Clinical Results
Semi-automated	SonoCiné	SonoCiné	Supine	Any HHUS probe	Multi-center trial, 6425 examinations [37]
Paddle Style ABUS	Invenia ABUS	GE Healthcare	Supine	15.4 cm concave, linear scanning	Multi-center trial, 15 318 women [33]
	Acuson S2000	Siemens Healthcare	Supine	15 cm linear array, linear scanning	[34, 52]
Dome Style ABUS	SOFIA	Hitachi Medical	Prone	Inset disc with “trapezoidal linear probe”, rotational scanning	[47, 53]
	SonixEmbrace	Ultrasonix Medical Corp.	Prone	Inset dome with concave transducer, rotational scanning	None
Ultrasound Tomography	QTscan	QT Imaging	Prone	Two transmission arrays, three reflection arrays	[54, 55]
	SoftVue	Delphinus Medical Technologies	Prone	22 cm ring array with axial scanning	[45, 56, 57]

Table 1.1: A comparison of ABUS systems illustrated in Figures 1.1 and 1.2. Further details and discussion provided in the text.

pose of establishing these guidelines, either by suitably modifying HHUS evaluation protocols, or building a standardization specifically for ABUS.

For several reasons, ABUS image reviewers in a clinical setting will request “second look” HHUS imaging be performed, negating the time (and cost) benefits of automated scanning. Since ABUS technology is still in its infancy, lack of training is a problem, and studies report a high rate of image artifacts, which can often resemble lesions [38, 41]. ABUS images are reviewed by a specialist after the scan, so the ability to modify compression, transducer position, and scanner settings is lost [11, 13]. Wojcinski *et al.* also reported the inability to use ultrasound modalities such as Doppler imaging or sonoelastography as a limitation [11] (this study was performed using the Siemens Acuson S2000 system).

One clinical study using a paddle style ABUS system reported difficulty imaging breasts larger than cup size D [11], while a study performed using a dome style transducer mentioned no such limitation [58]. ABUS systems are also unable to easily examine the axillary region, and the posterior region of larger breasts [10, 41].

With the exception of the SonoCiné, which can use any standalone US system, none of the currently available ABUS systems include integrated multimodal imaging in their automated scan.

In this thesis, we develop a multimodal breast imaging system based upon the SonixEmbrace ABUS. We believe that the rigid volume scanning scheme and tomographic geometry of the SonixEmbrace provide an excellent platform for this type of application, and the high quality of the ultrasound images produced by the system serves as an excellent starting point. Through the use of the SonixDAQ and the Sonix software development kit (SDK), pre-beamformed radio-frequency (RF) data is accessible for this system, which is necessary to perform photoacoustic imaging. Finally, our group has extensive expertise integrating elastography with Ultrasonix hardware, making multimodal imaging with this system more easily attainable.

1.4 Photoacoustic Tomography

1.4.1 Overview

One of the most promising emerging techniques for the diagnosis and staging of breast cancer is photoacoustic imaging (PAI). PAI is a biomedical imaging technique that leverages the advantages of both ultrasound and optical imaging to non-invasively image vasculature and other optically absorbing

1.4. Photoacoustic Tomography

targets [59,60]. Significant progress both in hardware and data reconstruction techniques over the past decade have moved this modality from the lab to the clinic [61], where it has shown promise in the detection and staging of cancer [62,63], the detection of inflammatory conditions [64,65], and for intraoperative guidance [66,67].

The past several years have seen a multitude of clinical studies demonstrating the utility of this technique for the detection [23,26,68,69,70], staging [71,72], and pre-operative margin assessment [73] of breast malignancies.

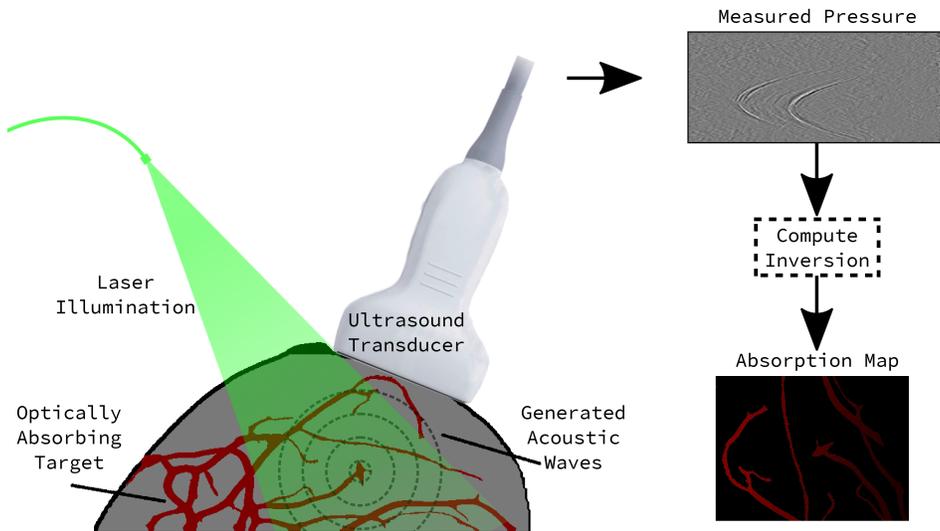


Figure 1.4: A schematic illustration of photoacoustic imaging. A laser pulse incident upon the sample is absorbed, causing rapid thermal expansion. This generates acoustic waves which can be detected with an US transducer. Finally, by processing these measured acoustic pressure data, the spatial distribution of optical absorption in the sample can be reconstructed.

When a pulse of laser light is incident upon tissue, it causes rapid, local thermal expansion. The resultant rise in pressure creates acoustic waves which propagate through the tissue and can be detected by standard ultrasound transducers. This technique is known as PAI, and is illustrated schematically in Figure 1.4. Since the energy deposition depends on local optical absorption, the measured ultrasound signal will contain physiological information about the targeted area. In particular, since the hemoglobin in blood is a strong optical absorber relative to tissue, PAI can visualize

vasculature in tissue (angiography). The exponential growth of cancerous tissue requires much more vasculature than healthy tissue, making dense local vasculature an excellent indicator of malignancy [74, 75, 76]. These factors combine to make PAI an excellent technique for early breast cancer diagnosis and directed treatment [58, 77, 78, 79]. An extension of PAI known as photoacoustic spectroscopy involves irradiating the tissue with different wavelengths of light to measure wavelength-dependent absorption. This is particularly useful in the imaging of vasculature, as the oxygenation states of hemoglobin, oxyhemoglobin (HbO_2) and deoxyhemoglobin (Hb) have different peak absorption wavelengths, allowing their relative concentrations, the blood oxygen saturation (SO_2), to be measured [80, 81, 82]. This technique has been shown to be highly effective in detecting and staging breast cancer [23, 75, 83] and clinical results using an automated scanner for this purpose have been promising [84].

When a three-dimensional (3D) arrangement of transducers is used to generate volumetric photoacoustic data, this is known as photoacoustic tomography (PAT). PAT is a well-suited modality for an ABUS, as the immobilization of the breast and careful control of the detector will allow for much higher imaging resolution [58]. One difficulty in the integration of PAT into an ABUS is that the entire transducer must be coupled to the tissue with a water bath, requiring that the scan be performed with the patient in the prone position.

1.4.2 Illumination Methods

The state of the art in breast PAT was recently summarized by Manohar *et al.* [12], and we refer the reader to this excellent review for a detailed description of common scanning and illumination geometries currently in use. Figure 1.5, reproduced from this work, concisely highlights the information relevant to this thesis. In particular, such systems are carefully designed from the ground up to balance the density of detection elements against the ease of delivering light to the sample. Illuminating the entire surface of the tissue with a collimated or divergent beam is often preferable, and this geometry has seen widespread use for the imaging of small animals [85, 86] and human extremities [87]. This type of illumination has the advantage of providing a relatively uniform fluence at the tissue surface, such that the fluence *within* the imaging volume can either be assumed constant, or accounted for with a depth-based attenuation model [23]. In the case of the SonixEmbrace, where the ultrasound hardware is a given and we wish to retrofit the illumination, we are more limited in our options. In particular,

for reasons elaborated upon in Chapter 2, a collimated beam illuminating the entire surface of the breast is not easily accomplished.

1.4.3 Theory and Reconstruction

PAT reconstruction amounts to solving for the energy absorption in the sample, $A_e(\mathbf{r})$ from the resultant photoacoustic (PA) signal, $p(\mathbf{r}, t)$. These quantities are related via the initial pressure distribution, $p_0(\mathbf{r}, t)$ by the differential equation

$$\left(\frac{\partial^2}{\partial t^2} - v_s^2 \nabla^2\right) p(\mathbf{r}, t) = \frac{\partial}{\partial t} p_0(\mathbf{r}, t); \quad (1.1)$$

$$p_0(\mathbf{r}, t) = \frac{v_s^2 \beta}{C_v} \Psi(\mathbf{r}, t) A_e(\mathbf{r}) [D(\mathbf{r}) \cdots] \quad (1.2)$$

where β , C_v , and v_s are the coefficient of thermal expansion, heat capacity, and acoustic speed in the irradiated medium, respectively, and $\Psi(\mathbf{r}, t)$ is the energy distribution of the laser pulse [81]. $D(\mathbf{r})$, the directional response of the transducer, and other spatially varying terms can be optionally included, depending on the sophistication of the reconstruction technique being employed [88]. Equation (1.1) assumes acoustic homogeneity, and ignores thermal diffusion and viscosity [60]. Under the further assumption of a delta function-like excitation pulse $\Psi(\mathbf{r}, t) = \Psi(\mathbf{r})\delta(t)$, a Green's function solution to Equation (1.1) describes the measured pressure data $p_d(\mathbf{r}, t)$ in terms of the initial pressure distribution $p_0(\mathbf{r})$ as

$$p_d(\mathbf{r}_0, t) = \frac{\partial}{\partial t} \left[\frac{t}{4\pi} \int \int_{|\mathbf{r}_0 - \mathbf{r}| = v_s t} p_0(\mathbf{r}) d\Omega \right] \quad (1.3)$$

where $d\Omega$ is the solid angle element of the source point \mathbf{r} relative to the measurement position \mathbf{r}_0 .

There are many ways to proceed from this point. The form of equation Equation (1.3) suggests standard tomography techniques that may employ a Radon transform, which have indeed been implemented by several groups [89, 90]. Radon transform approximations have proven successful in cases of enclosed spherical detector geometry with the sample being much smaller than the detection volume, but introduce artifacts when these conditions are not met.

Another approach, borrowed from traditional ultrasound image reconstruction, is to backproject each measured signal onto spherical surfaces

1.4. Photoacoustic Tomography

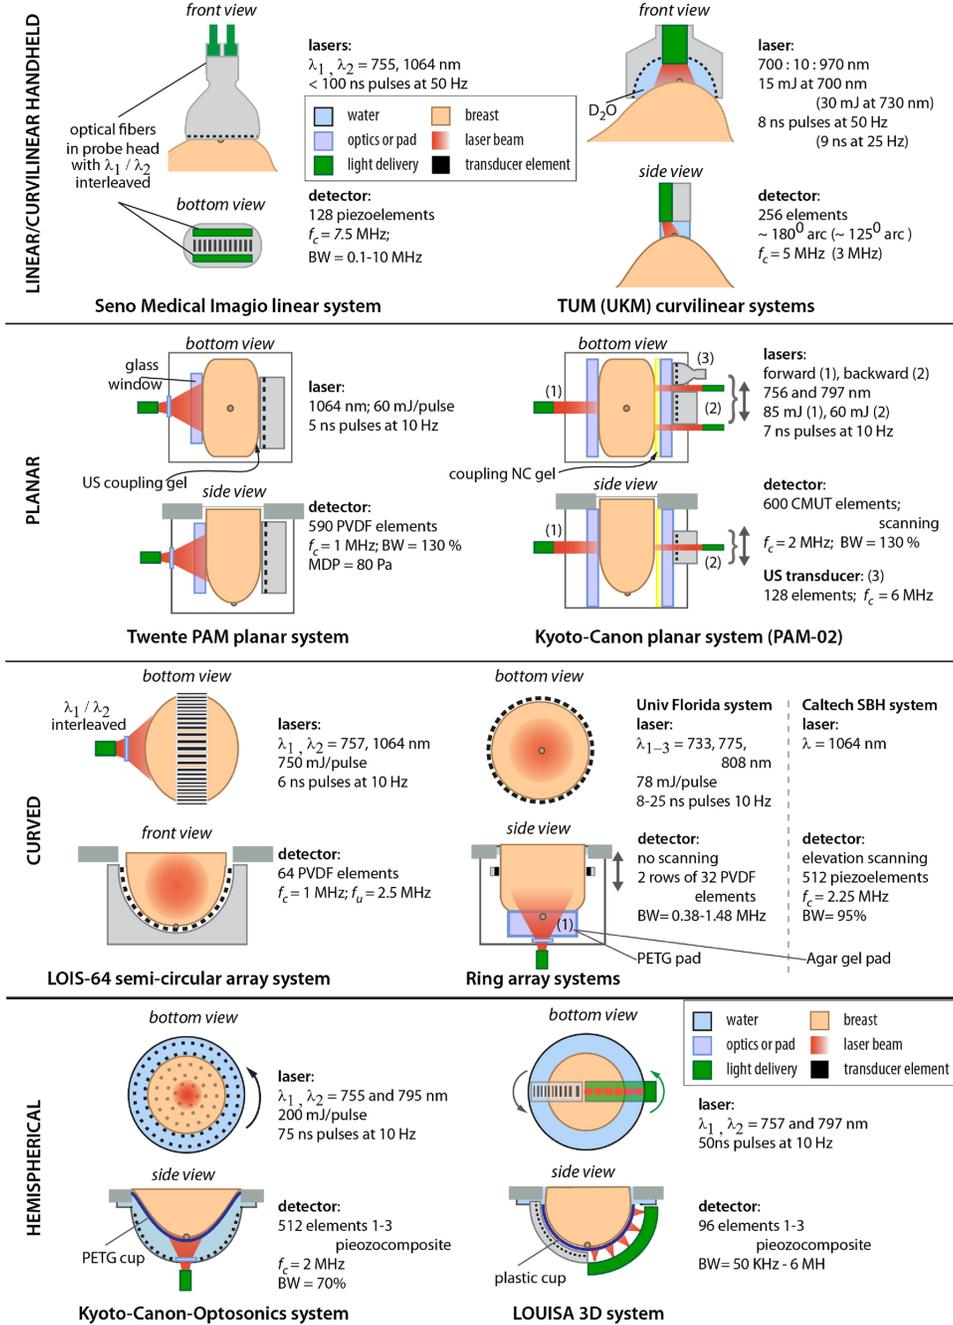


Figure 1.5: Illumination and scanning geometries in use for PAT and PAI of the breast. Reproduced from Manohar *et al.* [12].

with radii given by v_st . This is referred to as Poisson reconstruction or simply backprojection, and is commonly implemented using a delay-and-sum (DAS) [91,92] or synthetic aperture [93,94] algorithm. Extending these typically linear algorithms to spherical or cylindrical geometries has been accomplished using Fourier transform-based techniques [77,95] similar to those used in ultrasound reflectivity imaging [96]. This approach ultimately amounts to performing a backprojection with various filtering and weighting schemes.

Once $p_0(\mathbf{r})$ is obtained, it is commonly assumed that the illumination is spatially uniform, $\Psi(\mathbf{r}) = \Psi$, and the energy absorption $A_e(\mathbf{r})$ can be directly obtained via Equation (1.2). If the illuminating light is non-uniform, more sophisticated techniques such as iterative schemes [97], model-based inversion [98], or deconvolution-based inversion [88] must be employed.

1.4.4 Denoising Photoacoustic Data

Since photoacoustic signal intensity is proportional to the dose of laser light reaching the imaging target [60], and the laser light is heavily attenuated in tissue, photoacoustic tomography often involves the recovery of very weak signals. A fundamental tradeoff exists between maximizing this intensity while minimizing the laser dose to which the patient is exposed [99]. Although different reconstruction schemes ranging from filtered back-projection [100] to iterative approaches [101,102,103] have been developed to enhance the quality of the reconstructed images, it is still common to average the raw radio-frequency (RF) ultrasound data generated from multiple consecutive laser pulses, reducing stochastic noise and improving the signal-to-noise ratio (SNR) [104,105]. This approach not only increases the scan time proportionally to the number of frames acquired at each imaging position, lowering the frame rate of the system, but it also increases the total laser dose. More specific noise-reduction techniques such as bandpass filtering to the working range of the ultrasound transducer [100], or the removal of laser-induced noise using singular value decomposition (SVD) denoising [106] have also been employed to improve the quality of photoacoustic RF data.

1.5 Elastography

In the context of medical imaging, elastography is the measurement of the mechanical, or more specifically, viscoelastic, properties of tissue. This can

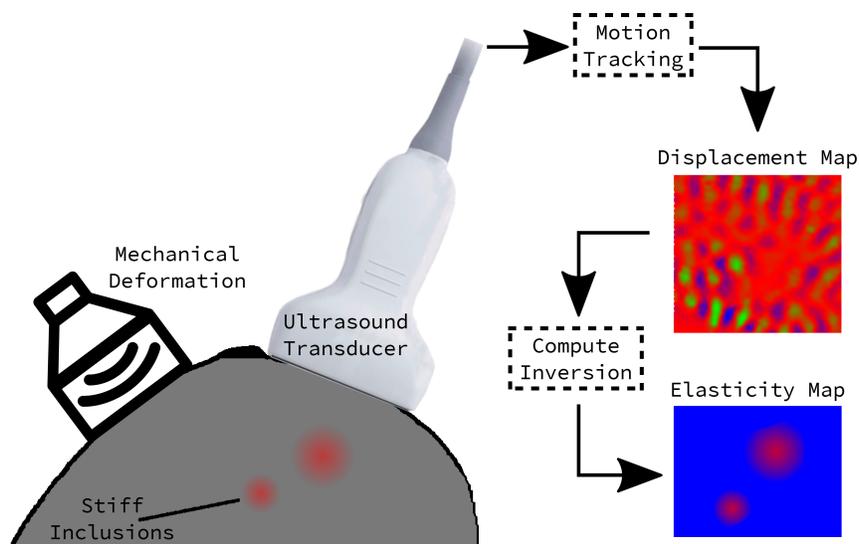


Figure 1.6: A schematic diagram of a general elastography technique.

offer a digital analogue to palpation, allowing a clinician to distinguish stiff tumours from the softer tissue surrounding them.

There are many elastography methods, but they broadly consist of three steps. First, some sort of mechanical deformation is applied to the tissue. Then, the resultant tissue displacement is tracked using some imaging or point measurement technique. Finally, the displacement information is used to calculate the elasticity of the tissue, subject to some mathematical model. Figure 1.6 illustrates this general process schematically.

This thesis makes use of a technique known as shear wave absolute vibroelastography (S-WAVE), a technique which was first introduced as “sono-elastography” by Parker *et al.* [107], in which a low-frequency harmonic excitation is applied externally to the tissue, and the resultant time-dependent displacement is tracked using US imaging. A mathematical model for the elastic response of the tissue is used to derive an inverse problem which can be solved to reconstruct an elastogram from the displacement data [108, 109, 110].

For a more detailed theoretical and technical overview of the S-WAVE technique, the reader is referred to the theses of Ali Baghani [111], Mohammad Honarvar [112], and Jeff Abeysekera [113].

Our group has previously used this technique for imaging the liver [114, 115], prostate [116, 117], *ex vivo* placenta [118], and *ex vivo* kidneys [119].

Eskandari *et al.* conducted a preliminary study using S-WAVE to differentiate malignant from benign breast lesions in 20 patients [120]. Using pathology of core-needle biopsy as ground truth, they reported that while invasive ductal carcinoma (IDC) could be distinguished from benign masses, more sophisticated data analysis would be necessary to avoid mis-classifying stiff benign lesions such as fibroadenoma. Recently, the work of Shao [121] *et al.* combined B-mode US, S-WAVE, and time-series RF US data and developed a machine learning pipeline to distinguish malignant and benign breast lesions, achieving an area under the curve (AUC) of $(95.00 \pm 1.45)\%$, with 95% sensitivity and 93% specificity. This work provides a very strong case for multi-modal imaging of the breast as a diagnostic tool.

1.6 Contributions and Thesis Overview

The objective of this thesis is to integrate new ultrasound-based imaging modalities into an existing SonixEmbrace ABUS system, with the ultimate goal of developing a multimodal platform for the imaging of breast malignancies. In particular, we seek to simultaneously acquire

- volumetric ultrasound,
- shear-wave elastography,
- Doppler flow imaging,
- and photoacoustic tomography

of the entire breast. Further, to preserve the clinical viability and automated nature of the scanner, we hope to avoid or minimize

- change to the existing system, such that its ultrasound imaging is unaffected,
- increase in scan time,
- and the introduction of any operator-dependence.

Toward this goal, this thesis consists of the following contributions:

1. A complete photoacoustic tomography system, consisting of
 - hardware to synchronize ultrasound acquisition with the laser illumination,

- software to control said hardware, as well as the ABUS transducer motor and multiplexing electronics,
 - a novel, fiber-based illuminator to deliver laser light to the tissue surface,
 - and a unique wire phantom design to test this imaging system in a robust yet repeatable way.
2. A photoacoustic reconstruction and data processing pipeline including
 - the implementation of a regularized iterative reconstruction scheme which improves image quality and reduces common imaging and reconstruction artifacts,
 - an image reconstruction scheme which accounts for the inherently non-uniform optical fluence introduced by our custom illuminator,
 - a novel data pre-processing method which can enhance extremely weak photoacoustic signals through the use of generative adversarial networks (GANs),
 3. A prototype system for shear-wave elastography and Doppler flow imaging with the SonixEmbrace ABUS, consisting of
 - an exciter to induce steady-state vibrations in the breast during imaging,
 - implementation of the necessary ultrasound sequencing for the multiplexed SonixEmbrace transducer,
 - implementation of scan-conversion for the unique transducer and motor geometry of the SonixEmbrace system.
 4. Preliminary testing and characterization of this combined system by imaging a custom, in-house multimodal phantom, as well as a commercially available multimodal breast phantom.

Chapter 2 provides a detailed overview of the novel hardware and software contributions which form the basis of the following chapters. We also include preliminary imaging results combining S-WAVE, Doppler flow, and PAT acquired with the SonixEmbrace system.

Chapter 3 describes in detail the photoacoustic imaging system. This includes a hardware overview, the design of the tissue illuminator, and a reconstruction method which accounts for spatially non-uniform illumination. This chapter also introduces a custom wire phantom design which is

imaged as a means to quantify the performance of the system in terms of signal-to-noise ratio (SNR) and achievable resolution.

Chapter 4 presents a method we have devised to remove both stochastic and structured noise in pre-beamformed photoacoustic data using GANs. This chapter also outlines our implementation of a regularized reconstruction method to remove imaging artifacts. These two techniques allow us to greatly improve our imaging results with no hardware changes.

Chapter 5 summarizes and discusses the results of the thesis, and provides suggested paths forward for this research.

Chapter 2

System Overview and Preliminary Results

In this chapter, we will describe the hardware and software additions and alterations made to the SonixEmbrace ABUS to facilitate PAT and S-WAVE imaging, including the system constraints which informed these design decisions. We will first discuss the S-WAVE sub-system, then the PAT sub-system. Finally, we present preliminary imaging results acquired with the combined system, and discuss how these results direct and motivate the remainder of the thesis.

2.1 SWAVE System

The main hardware addition required for S-WAVE imaging is a shaker which can excite periodic, low-frequency vibrations in the breast. Figure 2.1 shows the tissue shaker designed and built for the SonixEmbrace by Dr. Mohammad Honarvar. The shaker consists of a ring-shaped frame containing three equally-spaced voice coil actuators. The frame clamps around the upper portion of the cylindrical enclosure containing the imaging dome, allowing the weight of the patient to provide good coupling of the chest wall to the actuators. The current prototype uses a plastic plate to interface the actuators with the tissue, as shown in Figure 2.1e.

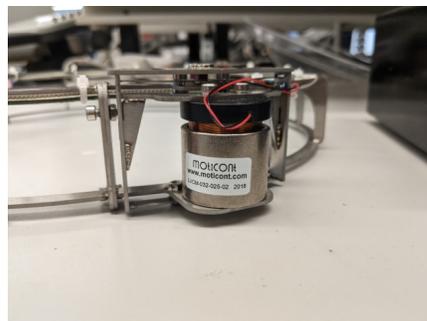
Since the dome enclosure can be moved up and down to accommodate different breast sizes, the shaker was designed to have a low profile, such that it is hardly noticeable once retracted below the padded surface of the ABUS scanning platform. This padded surface is removed in Figure 2.1.

In addition to the shaker itself, the S-WAVE system requires a controller to generate the excitation signal, and synchronize the data acquisition with the phase of said signal. The system also uses *eScan*, software developed by our group, for control of image acquisition and display.

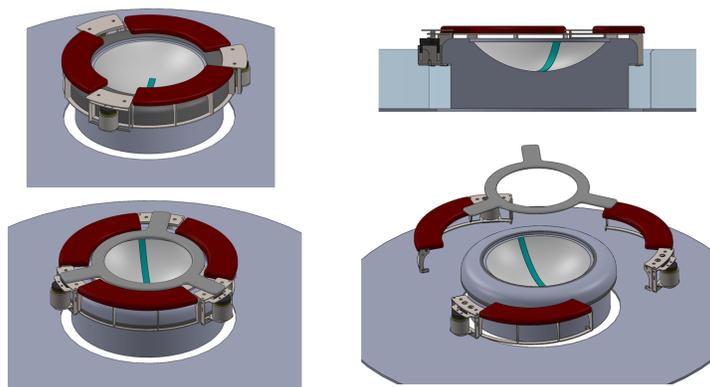
2.1. SWAVE System



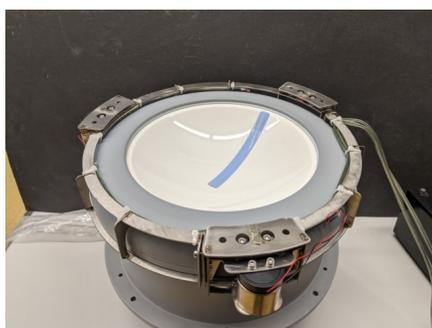
(a) The detached shaker, with the three excitation motors and electrical connection visible.



(b) Close-up of one of the three voice coil actuators.



(c) Design schematics for the tissue exciter.



(d) Attachment of the shaker to the imaging dome.



(e) Imaging configuration, including a plastic probe cover on the shaker, and the prototype plastic coupling plate.

Figure 2.1: Photos and schematics of the S-WAVE tissue exciter for the SonixEmbrace system.

2.2 Photoacoustic System

Manohar *et al.*, in their recent review titled “*Current and Future Trends in Photoacoustic Breast Imaging*” [12], described a generic photoacoustic breast imager as consisting of the following sub-systems:

1. patient-instrument interface,
2. ultrasound detector array,
3. data acquisition system
4. light source,
5. light delivery system,
6. computer running control and image reconstruction software.

The SonixEmbrace provides the first two items. A SonixDAQ unit (DAQ - Ultrasonix, Richmond, BC) constitutes our data acquisition system, and a Continuum Surelite II laser (Continuum Inc., Santa Clara, CA, USA) provides our light source. This thesis is concerned primarily with the final two components.

The following sections will discuss these system components, focusing on those which constitute contributions to the present thesis.

2.2.1 SonixEmbrace Transducer

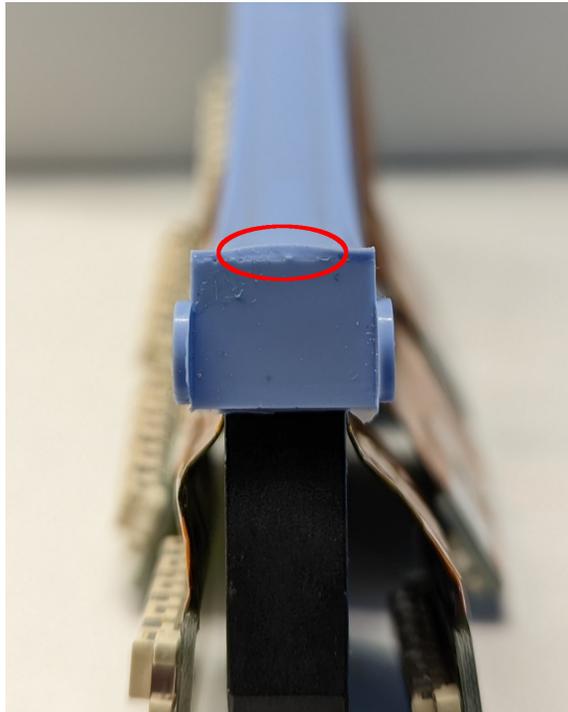
Given that the SonixEmbrace transducer was designed for US imaging, it is important to assess its suitability as a photoacoustic detector.

Figure 2.2 shows the SonixEmbrace transducer, separated from the imaging dome. The 3D scanning geometry and concave shape of the transducer make it well-suited for PAT, since a large portion of the dome volume has significant solid angle coverage. This means that a given imaging voxel will be “seen” by multiple transducer elements, from multiple directions, resulting in a higher SNR. Since the transducer extends past the rotational axis, the areolar region (which is known to be difficult to image [122, 123, 124]) is particularly well covered, with voxels being visible from thousands of distinct views in a typical scan. An ideal transducer would have a very wide acceptance angle, further increasing the number of voxels to which each element contributes [125]. As seen in Figure 2.2b, the SonixEmbrace transducer is limited in this respect by the inclusion of an acoustic lens. In Chapter 3

2.2. Photoacoustic System



(a) Side-on view of the SonixEmbrace transducer. The transducer curvature is -12 cm, and the total length is 11.5 cm.



(b) End-on view of the SonixEmbrace transducer, with the acoustic lens being visible as the slightly curved upper surface, indicated by a red ellipse.

Figure 2.2: Close-up photos of the SonixEmbrace transducer, removed from the imaging dome for clarity.

we measure the angular sensitivity of the transducer, and discuss how to account for this non-uniform response during image reconstruction.

The primary disadvantage of this transducer with regards to PAT is the working frequency range. Since photoacoustic signals are broadband, with the frequency content depending on the size of the absorbing structure, optimal detection requires a transducer with a carefully chosen centre frequency and working range [125, 126]. To resolve features with diameters ranging from a_{min} to a_{max} in a medium with acoustic velocity ν , the frequency range of the transducer, (f_{min}, f_{max}) is given by [127]

$$f_{min} = \frac{0.32\nu}{a_{max}} \qquad f_{max} = \frac{3\nu}{a_{min}}. \quad (2.1)$$

The SonixEmbrace transducer has a centre frequency of 10 MHz with a 90% bandwidth, giving a working range of 5.5 MHz to 14.5 MHz. Equation (2.1) suggests that this transducer would be able to resolve features ranging from 33 μm to 818 μm . Figure 2.3 compares this range to the sizes of various types of vasculature, indicating that our system can theoretically resolve most arteries and veins. Breast tumours begin to undergo exponential growth during the vascular phase, when their diameters are in the 1 mm to 2 mm range [74]. While a working range extending into the kHz would be ideal for resolving these features faithfully, the higher frequency content will still be able to resolve feature boundaries.

Figure 2.3 also shows the so called ‘‘stress confinement’’ limit for our laser system. At or below this limit, stress waves propagate significantly in the target during the laser pulse duration. The physical model governing PAT reconstruction does not hold in this regime, and significant blurring of the resultant images can occur [60]. The physics of PAT are described in detail in Chapter 3. With a laser pulse width of 5 ns, our system is above this limit for features greater than 10 μm in size.

2.2.2 Light Delivery

A custom solution was required to deliver the laser illumination to the imaging volume. While proof-of-principle phantom imaging could be conducted using a collimated beam from above the dome, this obviously would not work in a realistic imaging situation, where the dome is covered by the prone patient. This problem can be split into two parts. First, getting the light from the laser to beneath the dome, and second, getting the light past the dome to the imaging volume.

2.2. Photoacoustic System

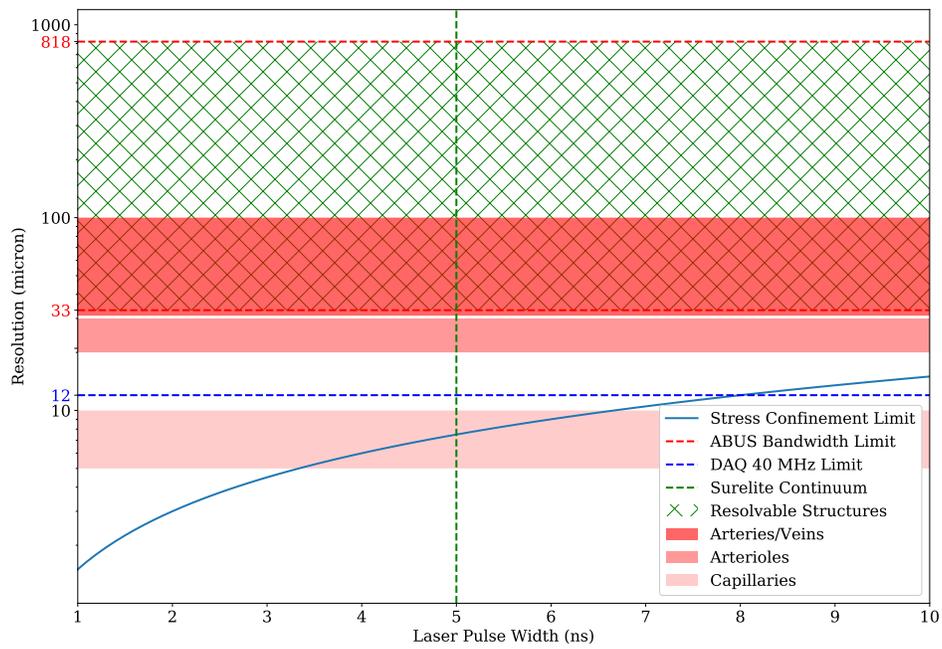


Figure 2.3: The diameter of resolvable structures as a function of the illumination laser pulse width in PAT imaging, demonstrating the limits of our system.

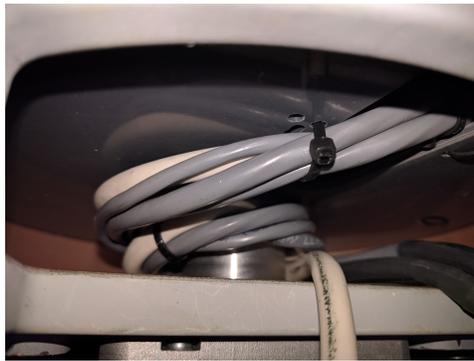
2.2. Photoacoustic System



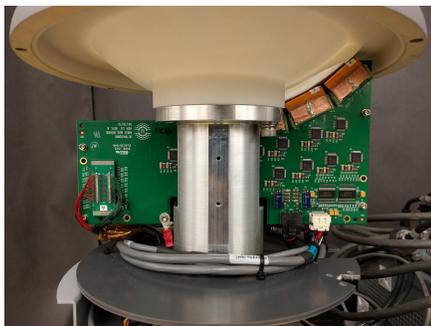
(a) Top down view on the transducer dome.



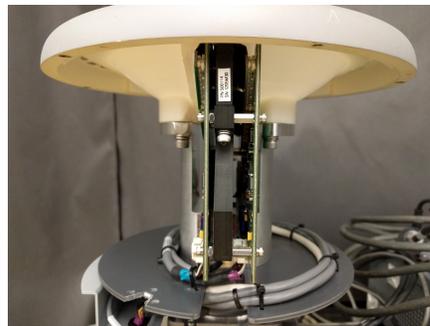
(b) Side-on view of the transducer dome.



(c) Under the ABUS dome cylinder, showing the path of electrical connections to the transducer.



(d) Outer cylinder removed, side-on view. One multiplexing board visible.



(e) Outer cylinder removed, end-on view. Multiplexing boards and axle profile visible.

Figure 2.4: Photos showing the geometry of the SonixEmbrace imaging dome and transducer electronics.

2.2. Photoacoustic System

The SonixEmbrace imaging dome is contained in a cylindrical enclosure which provides rigidity for the rotational motor, and supports the weight of the patient. This enclosure is shown in Figures 2.4a and 2.4b. The only opening in this dome is a panel through which the transducer electrical connections pass. As such, fiber-coupled light delivery is the only option which doesn't require extensive modification to the enclosure. Figure 2.4c shows this opening, with the motor axle and cable strain relief visible. While optical fiber bundles could maintain this radius of curvature without damage, these tend to be expensive, easily damaged, and difficult or impossible to polish or re-connectorize. For the purposes of this thesis, we will use less flexible, but more robust, silica-core multimode optical fibers. The larger radius of curvature of these fibers requires removing the dome enclosure entirely, but they are still able to coil 360° around the axle to allow full rotation of the transducer dome. Figures 2.4d and 2.4e show the transducer dome with the enclosure removed, revealing the transducer electronics.

Getting the light from the fiber output to the tissue surface is more difficult. Our detailed solution to this issue is the focus of Chapter 3, but we will discuss briefly here the modifications we have made to the Sonix-Embrace system to accommodate our illuminator. Figure 2.5 shows these modifications, as compared to Figure 2.4.

Since light is being delivered from beneath, and the imaging dome is opaque, an optical window was added. Rather than modify the existing dome, we used the computer-aided design (CAD) files for the transducer dome to 3D print a new dome with a rectangular hole alongside the transducer. A piece of polycarbonate plastic was affixed in the hole with epoxy, maintaining the water-tightness of the dome. Figures 2.5a and 2.5b show this window from above and below, respectively.

The optics for the illuminator (again, see Chapter 3 for details) are contained in a 3D printed enclosure which mounts to the motor axle, fixing alignment relative to the transducer dome. The position of this enclosure can be seen in Figures 2.5c and 2.5d. By including all optical elements in the CAD process, the positions of the illuminator enclosure and optical window were both optimized to maximize the amount of light reaching the tissue surface.

This design could in principle be extended to include a second illuminator placed symmetrically on the other side of the transducer (with a second optical window). This would however necessitate splitting the fiber output, requiring more optical elements, and as such was not explored in the present study.

2.2. Photoacoustic System



(a) Top-down view of the new 3D-printed imaging dome, showing the position of the optical window.



(b) Under the imaging dome, to illustrate the space constraints upon the position of the optical window and illuminator beam. The illuminator is removed in this photo for clarity.



(c) Side-on view demonstrating the position of the illuminator, indicated by a red circle, under the imaging dome.



(d) End-on view demonstrating the position of the illuminator, indicated by a red circle, under the imaging dome.

Figure 2.5: Photos of the SonixEmbrace imaging dome, illustrating the optical window and the position of the illuminator optics.

2.2.3 Hardware Synchronization

The various hardware elements of the PAT system must be simultaneously controlled for imaging to be possible. While some of this is accomplished in software, the synchronization of the illumination laser and ultrasound acquisition requires more precision. We designed and built a custom “trigger box” to handle all timing-critical aspects of the system in a robust and reproducible way. This circuit must provide precision across multiple time scales — \sim ns transistor–transistor logic (TTL) pulse widths, \sim μ s Q-switch delays, \sim ms laser repetition rate — and stability over the course of experiments which can last hours.

Photos of the trigger box are shown in Figure 2.6. Based on an Arduino Uno, the system is electronically simple, consisting mainly of digital input/output pins, switches, and a universal serial bus (USB) connection for serial communication with the PC. The box is designed to operate the Continuum Surelite II laser in a steady-state mode, increasing power stability, and minimizing any jitter in timing. The algorithm used is schematically illustrated in Figure 2.7. Essentially, when enabled, the circuit provides flashlamp triggers to the laser at its optimal 9.8 Hz pulse repetition frequency, and when data is requested, Q-switch triggers are sent without timing interruption to induce lasing and generate illumination pulses. Synchronized with these pulses are triggers to begin data acquisition by the DAQ, followed by triggers to the US sequencer to update the transducer multiplexing. The Arduino Uno has a clock speed of 16 MHz, theoretically allowing timing precision into the nanosecond range.

Figure 2.7 also demonstrates how this triggering integrates asynchronously with the motor control to generate a volumetric scan. The full Arduino code implementing this algorithm is included in Appendix D.

2.2.4 Control Software

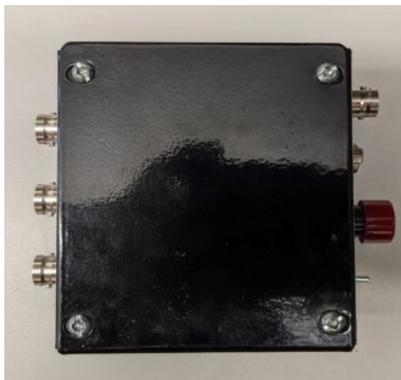
Since our photoacoustic system contains multiple devices which must be synchronized and communicated with, a purpose-built software solution is required. This central piece of software serves to

- allow control and monitoring of the trigger box, SonixEmbrace motor, and laser power meter,
- change experimental parameters such as volumetric scan settings, and laser power,

2.2. Photoacoustic System



(a) The trigger box with lid removed, revealing the Arduino board and wiring.



(b) Top-down view of the trigger box.



(c) Side-on view: USB and power connections.



(d) Side-on view: single pulse button, laser enable switch, and 9-pin connector for laser triggering.



(e) Side-on view: TTL trigger connections.

Figure 2.6: Photos of the Arduino-based trigger circuit used to control the Q-switched laser, DAQ acquisition, and transducer multiplexing.

2.2. Photoacoustic System

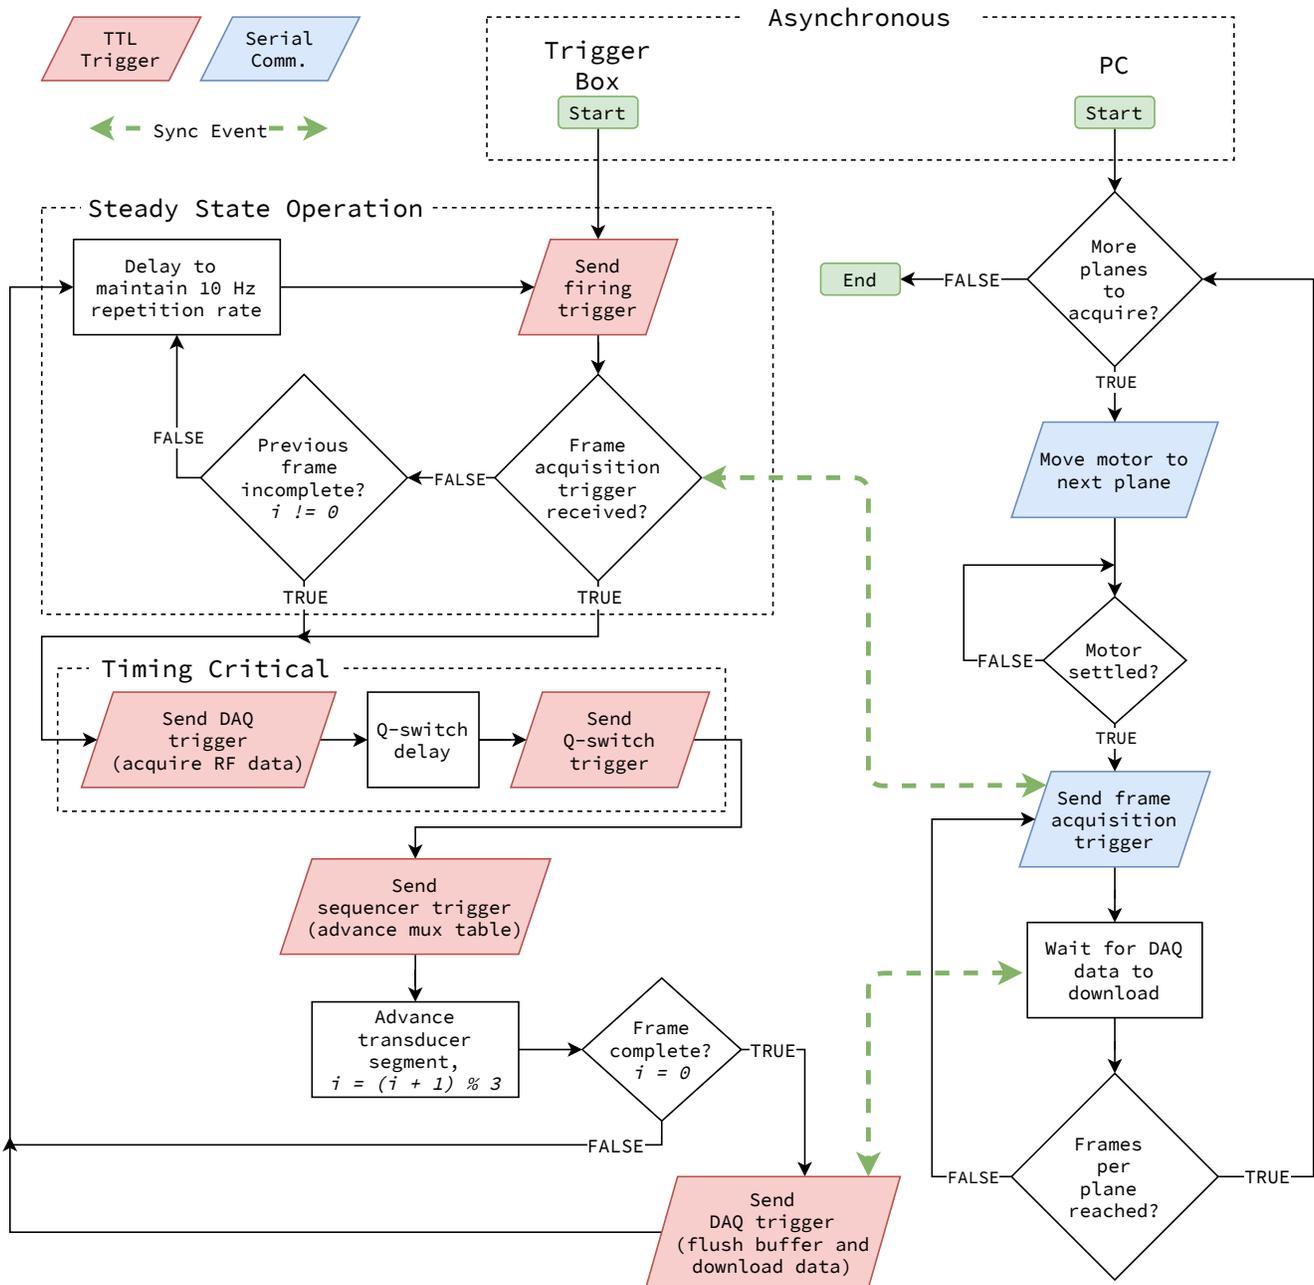


Figure 2.7: An outline of the triggering algorithm for stable operation of the Q-switched laser source during PAT acquisition.

2.3. Summary

- collate acquired data in a format suitable for storage and further processing,
- compute image reconstructions for subsequent display.

The *ABUS Photoacoustic Interface* was developed in Python 3 using *PyQt5* for the graphical user interface (GUI), *pyserial* for all device communication over USB, and *pyCUDA* for graphics processing unit (GPU) accelerated image reconstruction. See Figure 2.9 for a screenshot of the software during data acquisition. All experimental and reconstruction parameters can be changed in the GUI, and controls are available for common tasks such as homing the ABUS motor, firing single laser pulses for optics alignment and testing, and manually sending serial communication commands for debugging purposes. Following a volumetric scan, all data is combined into a custom binary file format containing all necessary information to generate image reconstructions, which can be performed either immediately, or at a later time.

We note that the SonixDAQ is initialized and programmed via the Sonix SDK, but once running, can be controlled entirely using 5 V TTL triggering signals.

2.3 Summary

Figure 2.8 shows the components of the combined imaging system, divided into two sub-systems. The photoacoustic system consists of the illuminating laser and associated optics, as well as the synchronization hardware. The S-WAVE system is simpler, consisting of the tissue shaker, and the controller that drives it and synchronizes the US acquisition. Doppler flow imaging requires no additional hardware or data acquisition, and only requires additional processing of the raw RF data used to generate the S-WAVE images.

The two imaging modules operate completely independently, allowing acquisition of any combination of modalities.

The remainder of this chapter consists of a preliminary study exploring the imaging capabilities of this system.

2.3. Summary

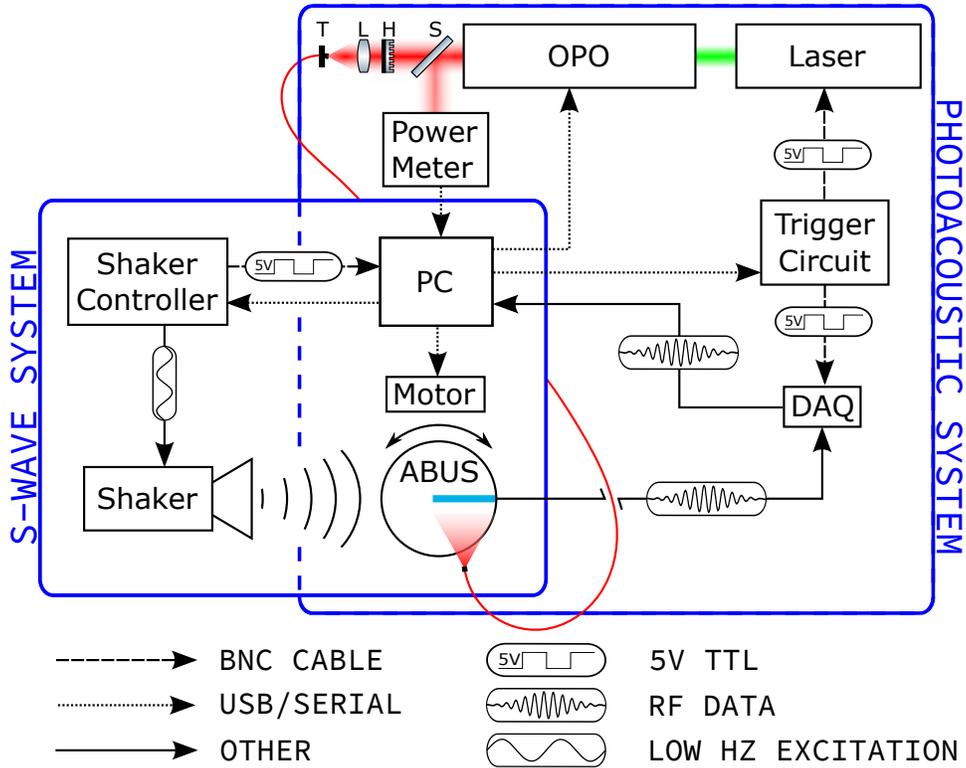


Figure 2.8: Block diagram of the multimodal imaging system showing the essential components and their interactions. Optical elements S, H, L, and T are a beam splitter, homogenizing lens array, fiber coupling lens, and translation stage, respectively.

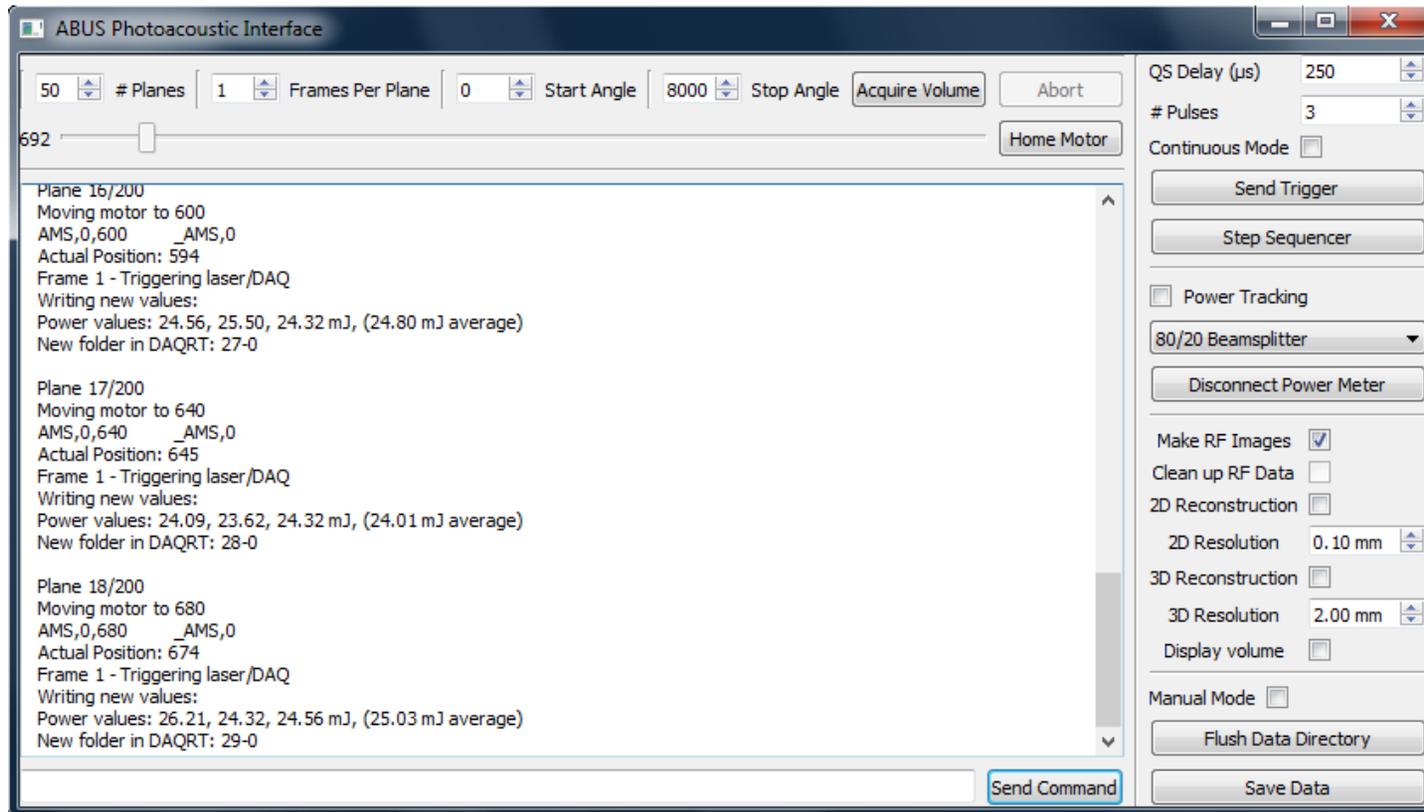


Figure 2.9: The ABUS photoacoustic acquisition software. This interface allows control over volume scan parameters, laser power, and reconstruction parameters (for online reconstruction). It also facilitates communication with the trigger box.

2.4 Preliminary Study¹

2.4.1 Methods

Ultrasound System

We performed all imaging using the SonixEmbrace ABUS. The ABUS is designed to capture B-mode ultrasound images of the breast using a rotating concave transducer while the patient lies in the prone position. The ABUS transducer is concave and has a radius of 12 cm, with 384 piezoelectric elements and a pitch of 0.3 mm. The elements have a center frequency of 10 MHz (bandwidth 90%) [49]. The arc of the transducer extends past the centre of the ABUS dome by 2.5 cm such that the sampling density is twice as high near the rotational axis of the dome. The motor which rotates the ABUS transducer is controlled over a Universal Serial Bus (USB) connection, and can achieve an angular accuracy of 0.5° over 360°.

Elasticity Measurement

S-WAVE measures absolute tissue stiffness by externally exciting the tissue with a periodic vibration and precisely tracking the resultant displacement using ultrasound [113,120]. We developed a custom tissue exciter compatible with the geometry of the ABUS. It consists of a circular steel frame which clamps onto the top of the transducer dome. The frame houses three motors connected via magnets to a rigid plastic plate, which serves to couple the motor motion to the chest of the patient, resulting in a uniform excitation in the patient’s breast. These motors are driven by an elastography module which synchronizes their phase with the ultrasound acquisition. Photos of the shaker are included in Figure 2.1.

To test the S-WAVE system, we used a CIRS model 059 breast biopsy training phantom (CIRS Inc., Norfolk, VA, USA), which is shown in Figure 2.10. This phantom contains dense masses ranging from 3 mm to 10 mm in diameter. The manufacturer cites a background elastic modulus of 20 ± 5 kPa, and masses which are “at least twice as stiff as the background material”. While a calibrated phantom would be a more rigorous test, the scanning geometry of the ABUS precludes the use of most commercially available options.

¹This section is previously published [128], and has been modified and expanded upon for this thesis.



Figure 2.10: CIRS model 059 breast elasticity phantom.

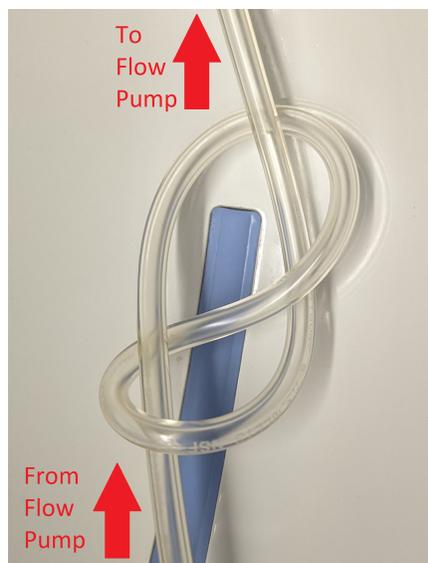
Doppler Flow

The same time domain correlation data used to calculate the tissue motion in the S-WAVE technique can be used to generate a flow magnitude, or *power Doppler*, signal [129]. Figure 2.11 shows a knotted tube flow phantom used to test this capability. This phantom consists of 6.4 mm inside diameter Tygon™ tubing, tied in a figure-of-eight knot to create flow in various directions within the dome. The tube was connected to a peristaltic pump and water with 1% cellulose by weight was pumped through at a rate of 200 mL min^{-1} . This corresponds to an average flow velocity of 10 cm s^{-1} , which is comparable to flow rates measured in the axillary artery [130]. The entire tube was submerged in water for acoustic coupling to the transducer.

All S-WAVE data presented here were acquired using 250 Hz, 300 Hz, and 350 Hz excitation simultaneously.

For both S-WAVE and Doppler, plane (2D) images were reconstructed and displayed alongside B-mode in real time during data acquisition. These data can then be scan converted offline to generate volumetric images.

2.4. Preliminary Study



(a) Top-down view of the knotted tubing flow phantom, indicating direction of flow, as well as the position of the tubing within the scanning dome.



(b) Side-on view of the knotted tubing flow phantom.

Figure 2.11: A proof-of-principle Doppler flow imaging setup for the Sonix-Embrace ABUS. During imaging, the scanning dome is filled with water which has been removed here for clarity.

Photoacoustic Illumination and Reconstruction

The photoacoustic sub-system and associated data reconstruction used in this study is described in full in Chapter 3. Briefly, our illumination system consists of a Continuum Surelite II Q-switched Nd:YAG laser (Continuum Inc., Santa Clara, CA, USA) coupled to an optical parametric oscillator (OPO) from the same manufacturer to control the output wavelength. All data presented here were acquired with an illumination wavelength of 700 nm, with a pulse energy of approximately 20 mJ. This corresponds to a peak fluence at the sample surface of 3.0 mJ cm^{-2} . For photoacoustic imaging, RF data from the transducer elements is acquired using a Sonix-DAQ module (Ultrasonix), which collects pre-beamformed transducer data at 40 MHz. This acquisition happens in parallel with standard ultrasound acquisition, facilitating multi-modal imaging. The DAQ acquisition was synchronized with the laser illumination pulses by a purpose-designed triggering circuit, controlled by a serial connection from the ultrasound PC. Reconstructions were performed using a delay-and-sum algorithm implemented in Python. Both 2D and 3D reconstructions are performed with a 0.1 mm resolution, and weighted to account for the spatially non-uniform illumination (see Chapter 3 for details).

Multimodal Phantom Construction

This phantom was designed to test elasticity and photoacoustic imaging simultaneously. It consists of a homogeneous background of 5% gelatin with 25% whole milk as an optical scatterer, and 1% by weight cellulose to provide ultrasound speckle. This background should have an elastic modulus of approximately 31 kPa [131], which is similar to that of healthy breast tissue (around 20 kPa for fat, and anywhere from 30 kPa to 60 kPa for glandular tissue) [132]. Optically, we expect an absorptive attenuation coefficient, μ_a of $3.75 \times 10^{-4} \text{ mm}^{-1}$, and a scattering attenuation coefficient, μ_s of 0.75 mm^{-1} [133]. This provides similar scattering to breast tissue, which has $\mu_s = 1.0 \text{ mm}^{-1}$, and while $\mu_a = 2.0 \times 10^{-3} \text{ mm}^{-1}$ [134] for breast tissue differs significantly, this contributes minimally compared to scattering.

Cylindrical inclusions with diameters of 13 mm and 20 mm, targeting both imaging modalities were suspended in the background. The photoacoustic inclusions contained black acrylic paint throughout to act as an optical absorber. The elasticity inclusions had 10% gelatin by weight, for an elastic modulus of 81 kPa [131], and red food colouring to make them more easily identified by eye. These inclusions do not mimic any particular breast

malignancy (the elastic moduli of which can vary from 20 kPa to 600 kPa, depending on type and measurement technique [132]), they just need to be measurably stiffer than the background. The phantom is moulded to fit the spherical cap shape of the ABUS dome. A cross-sectional view of this phantom can be seen in panel a) of Figure 2.14.

2.4.2 Results

Elasticity

Figure 2.12 shows B-mode, 2D phasor, and absolute elasticity images acquired of the CIRS breast phantom. When averaged over the three excitation frequencies, we measured the elastic modulus of the background to be 17 ± 1 kPa and masses to be 32 ± 3 kPa, with uncertainty defined as the standard error of the mean.

Doppler Flow

Figure 2.13 shows volumetric Doppler imaging of the knotted flow phantom described in Section 2.4.1, wherein the shape of the knot can be clearly discerned. Since this data is generated from interpolating 2D images acquired at different angular positions of the transducer, an “axial” direction cannot be uniquely defined in the volumetric data. While it would be possible to locally compute a flow direction at each position the volume, this information is much less useful than the flow magnitude in the context of quantitative flow imaging. As is, our data approximately represents flow toward or away from the dome surface, shown in red and blue respectively. This is particularly apparent in the side-on view shown in Figure 2.13b. The measured flow magnitude was largely in the 15 cm s^{-1} to 20 cm s^{-1} .

Multimodal Phantom

Figure 2.14 shows a representative 2D plane of the multimodal phantom, with B-mode, elasticity, and photoacoustic images. Also shown is a photograph of the corresponding cross-section of the phantom, taken after imaging to illustrate the inclusion geometry. The line parallel to the transducer in the B-mode image is a gel pad used to acoustically couple the phantom to the transducer while allowing the transducer dome to rotate freely. The diagonal line at the bottom of the B-mode image corresponds to the other surface of the phantom.

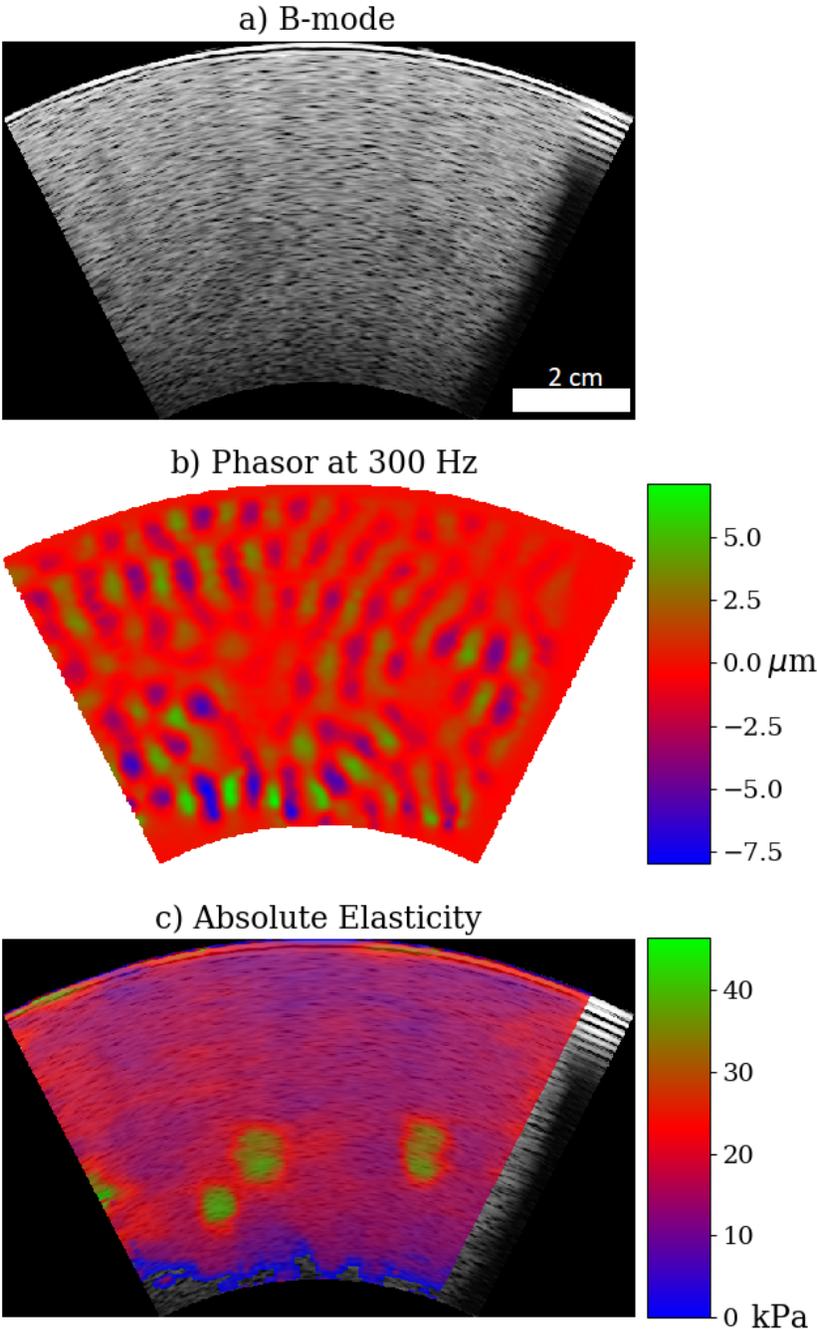
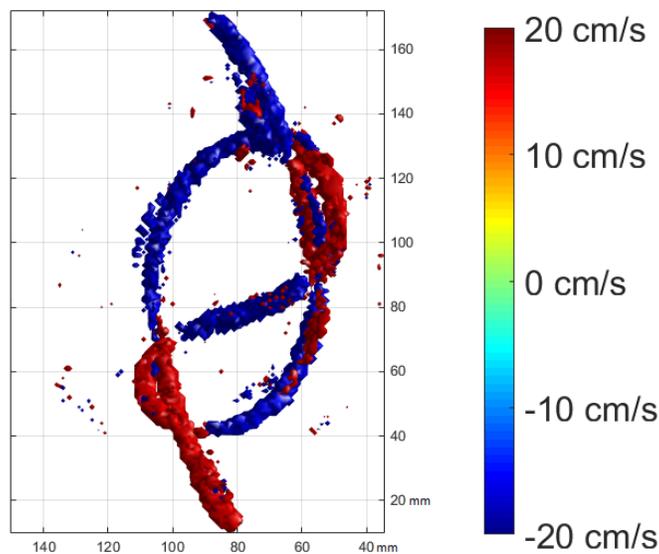
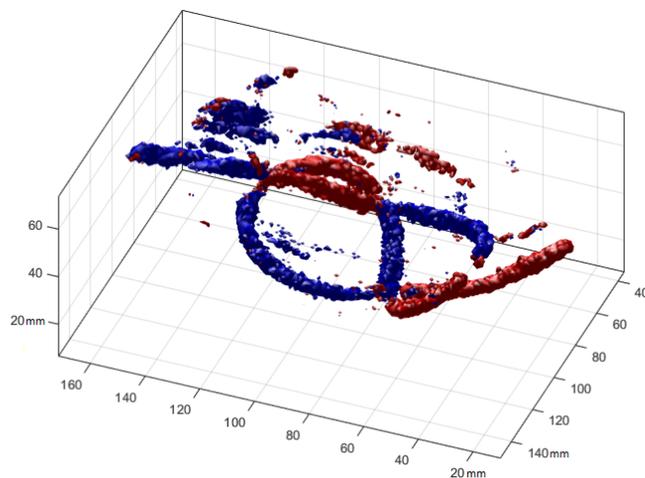


Figure 2.12: B-mode and absolute elasticity images of the CIRS model 059 breast phantom shown in Figure 2.10. Total imaging depth is 6 cm.

2.4. Preliminary Study



(a) Top-down projection of Doppler imaging of the knotted flow phantom. This view direction corresponds to Figure 2.11a. Colorbar indicates magnitude of flow, and applies to both panels.



(b) Side-on projection of Doppler imaging of the knotted flow phantom. This view direction corresponds to Figure 2.11b.

Figure 2.13: Volumetric Doppler flow imaging of the knotted flow phantom shown in Figure 2.11. Colour indicates the magnitude of the axial component of the flow (see discussion in text).

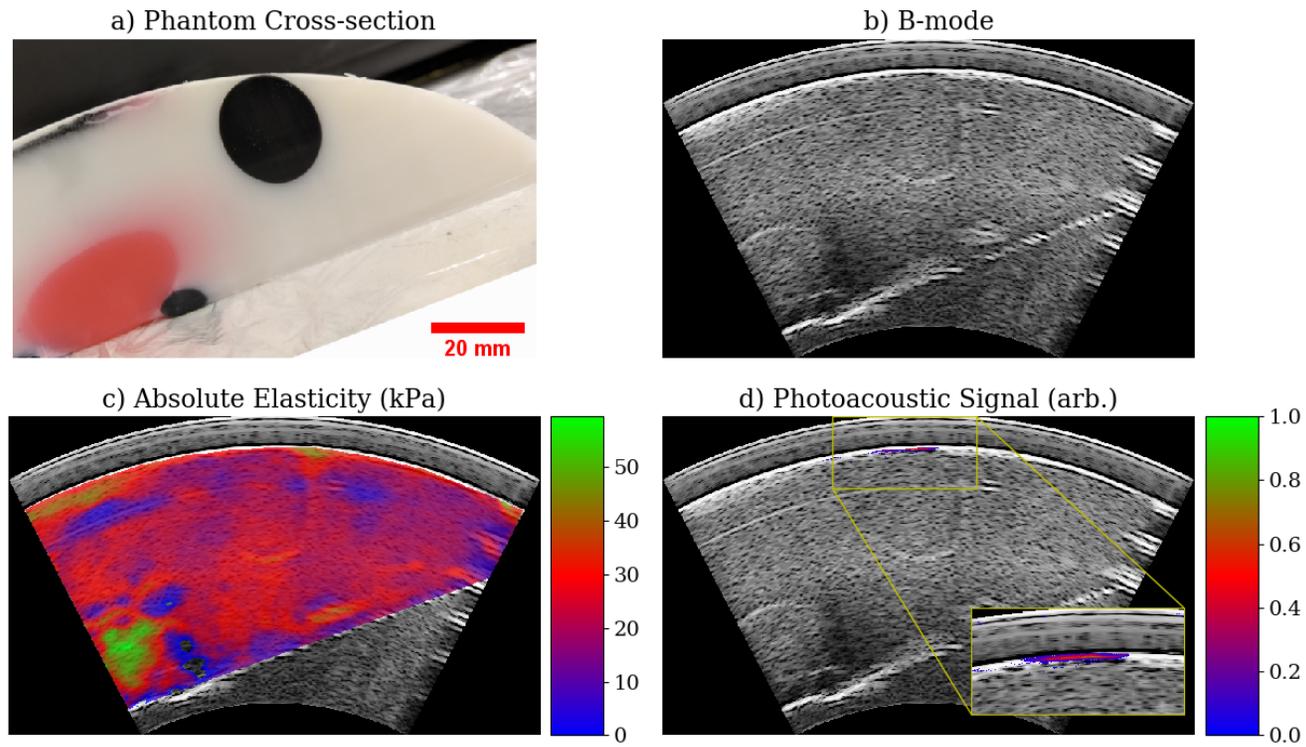


Figure 2.14: A representative 2D plane of the multimodal phantom. Panel a) shows a photograph of a cross-section of the phantom, showing the black (photoacoustic) and pink (stiffness) inclusions. Panel b) is a B-mode image. Panel c) is an S-WAVE elastogram. Panel d) is photoacoustic data, and is thresholded below 0.1. Total imaging depth is 6 cm.

2.4.3 Discussion

The absolute elasticity measurements of the CIRS breast phantom serve as a good initial test of this system. While the manufacturer listed values are imprecise, our measurements fell within the expected ranges. We measured the elastic modulus of the background to be 17 ± 1 kPa, which agrees with the expected value of 20 ± 5 kPa. The inclusions should be twice as stiff, which suggests a range of 30 kPa to 50 kPa, given the provided uncertainty in the background value. Our measured elastic modulus for the inclusions was 32 ± 3 kPa, which does fall within this range. That these values agree shows that we have successfully integrated the S-WAVE hardware with the SonixEmbrace ABUS. Qualitatively, this phantom also provides an excellent example showing the potential of multimodal imaging. The boundaries of the stiff inclusions are only barely visible in the B-mode image. Elasticity allows visualization of isoechoic regions, even deep within tissue.

The knotted tube phantom was a simple way to test our Doppler flow measurements. This distinct geometry also provides a visual confirmation that our 3D scan conversion is working properly. Since this method only approximates the axial component of the flow, without moving to a more sophisticated reconstruction method, it serves more to indicate the presence of flow, rather than the magnitude or direction. That said, we measured the magnitude of the flow to be in the 15 cm s^{-1} to 20 cm s^{-1} , which is comparable to the known flow rate of 10 cm s^{-1} .

The multimodal phantom, as shown in Figure 2.14, demonstrates the utility of simultaneously acquired elasticity and photoacoustic tomography data. The elasticity image correctly suggests that the pink inclusion is stiffer than the background, as expected. More precisely, we measure the elastic modulus of the inclusion to be approximately 50 kPa, and that of the background to be 20 to 25 kPa. Note that the black inclusion, which has the same gelatin concentration as the white background, has visible boundaries in B-mode but does not provide contrast in the elasticity image. In the photoacoustic image, on the other hand, a high contrast signal is seen at the boundary of the black inclusion nearest to the transducer. There is very little photoacoustic signal elsewhere in the plane, but this high contrast signal from a strong optical absorber is at least qualitatively what we would expect.

2.4.4 Conclusions

Since there is little precedent for combining these modalities of imaging, we have many options to explore further, and several ways to improve our system in the hopes of moving toward *in vivo* imaging.

We did note that in the multimodal phantom, the photoacoustic signal from inclusions deeper than 2 cm was very weak. This is not surprising, and is a common problem in photoacoustic tomography. In our particular case, one source of fluence loss was reflection at the surface of the phantom, and at the interface of the acoustic coupling pad. While the coupling pads appear transparent, we have not yet precisely characterized their optical properties. Choosing optically index-matched materials can help to mitigate this source of loss.

Further, in the interest of minimizing acquisition time, all data for this study were acquired without averaging over multiple laser illumination pulses. Our fluence of 3.0 mJ cm^{-2} is still well below the ANSI safety limit of 20 mJ cm^{-2} , so we can average over multiple pulses to increase the signal-to-noise ratio in future studies.

Based on these results, our photoacoustic system stands to gain the most from further development. In Chapter 3, we describe and characterize this system in greater detail, and work toward improving its capabilities.

Chapter 3

Integrating Photoacoustic Tomography into a Multimodal Automated Breast Ultrasound Scanner²

3.1 Overview

PAT is a promising emergent modality for the screening and staging of breast cancer. To minimize barriers to clinical translation, it is common to develop PAT systems upon existing ultrasound hardware, which can entail significant design challenges in terms of light delivery. This often results in inherently non-uniform fluence within the tissue, and should be accounted for during image reconstruction. In this chapter, we aim to integrate photoacoustic tomography into an ABUS with minimal change to the existing system.

To achieve this, we designed and implemented an illuminator which directs spatially non-uniform light to the tissue near the acquisition plane of the imaging array. We developed a GPU-accelerated reconstruction method which accounts for this illumination geometry by modelling the structure of the light in the sample. We quantified the performance of this system using a custom, modular photoacoustic phantom, and graphite rods embedded in chicken breast tissue.

3.2 Introduction

This rigid volume scanning scheme and tomographic geometry of the Sonix-Embrace provides an excellent platform for multimodal imaging. Integrating photoacoustic tomography into this scanning geometry without sacrificing

²This chapter is previously published [135], and has been modified and expanded upon for this thesis.

3.2. Introduction

quality of the other modalities presents challenges in hardware and optics design, and in data reconstruction.

The state of the art in breast PAT was recently summarized by Manohar *et al* [12], and we refer the reader to this review for a detailed description of common imaging and illumination geometries currently in use. See also Figure 1.5, which summarizes the systems discussed in this review. Such systems are carefully designed from the ground up to balance the density of detection elements against the ease of delivering illumination to the sample. Illuminating the entire surface of the tissue with a collimated or divergent beam is often preferable, and this geometry has seen widespread use for the imaging of small animals [85, 86] and human extremities [87]. This type of illumination has the advantage of providing a relatively uniform fluence at the tissue surface, such that the fluence *within* the imaging volume can either be assumed constant, or accounted for with a depth-based attenuation model [23]. In the case of the SonixEmbrace, where the ultrasound hardware is a given and we wish to retrofit the illumination, we are more limited in our options. In particular, due to the design constraints discussed in Section 1.3, a collimated beam illuminating the entire surface of the breast is not possible.

Taking inspiration from the ever-growing body of work integrating illumination hardware into freehand ultrasound systems [23, 136, 137], we explored diffusive, beam-shaping optics which can be placed near the transducer to provide fluence in the acoustic acquisition plane. In this study, we demonstrate the use of a fan-shaped diffuse beam for this purpose. This illumination approach is novel amongst tomographic breast PAT systems, and in the present study we provide a complete characterization of both our hardware and software, as well as a new GPU-accelerated reconstruction scheme. At the same conference, Oraevsky *et al* presented the LOUISA-3D system [24], which is superficially similar to our own in its use of a curved transducer array, and a moving illuminator, but differs in several significant ways. Most relevant is that their system was designed to include a transparent imaging dome allowing the illuminator to be placed nearer to the tissue, and to illuminate a larger area with a single laser pulse. This illumination system also moves independently of the transducer array, whereas these are coupled in our system. The authors describe the illuminator itself as an arc-shaped “paddle” with multiple fiber optic segments, which is distinct from our simpler single-fiber solution.

We also note that our design is distinct from existing systems which employ diode light bars or rectangular shaping diffusive elements [137, 138] in that the light diverges significantly in the space between the optics and the

tissue, resulting in inherently non-uniform fluence at the tissue boundary. While these optics offer flexibility in terms of the light delivery, this spatial non-uniformity of the fluence must be accounted for during data reconstruction [98, 139, 140].

Spatial fluence modelling is commonly used in photoacoustic simulations to provide an accurate ground truth for experimental validation [138, 139], and is usually accomplished via a Monte Carlo simulation of the radiative transfer equation (RTE). It is also often necessary in quantitative photoacoustic tomography (QPAT) to fully model the fluence such that the optical parameters of the tissue can be recovered quantitatively. See, for example the recent work of Hänninen *et al* [141]. In motivating their work, the authors highlight that a significant shortcoming of the standard Monte Carlo method is the necessary computation time, which can be prohibitive for large tomographic problems such as ours. This is further compounded in the case of non-stationary illumination since the fluence must be computed separately for each tomographic acquisition angle. In the present study, we employ a first-order approximation to the RTE to lower this computational burden.

We note the unique recent work of Park *et al* which estimates and accounts for non-uniform illumination and attenuation as a part of the image reconstruction [142]. Our work differs from this approach in that we do not need to compute the location of the breast surface, and that we do not assume constant fluence at the sample surface.

In this study we present a novel, fan-shaped illuminator for photoacoustic tomography of the breast, and we characterize and quantify the performance of this illumination system in terms of the fluence it can deliver to the tissue surface. We also explore the use of an approximate, model-based method for deconvolving spatially-varying illumination from measured photoacoustic signals, and quantify the benefits of using such a method. To our knowledge, this study constitutes the first attempt to integrate photoacoustic tomography into a pre-existing multimodal ABUS system. By attempting to add PAT to this system with minimal changes to the existing design, we have had to make compromises relative to the current state of the art. In particular, our transducer is a conventional array that is designed primarily for B-mode imaging.

3.3 Methods

3.3.1 The SonixEmbrace ABUS

The SonixEmbrace transducer array consists of 384 elements on a 11.5 cm long circular arc with a curvature of -12 cm. It has a center frequency of 10 MHz and 90% bandwidth. The transducer dome has a radius of curvature of 12 cm and a diameter of 18 cm. A photograph of the dome is presented in Figure 3.1. This large transducer requires multiplexing (mux) hardware



Figure 3.1: The SonixEmbrace ABUS system, consisting of the ultrasound PC and scanning platform. Inset shows a top-down view of the transducer dome, with the blue transducer array visible.

which occupies a significant amount of space directly behind the transducer array under the dome (see Figure 3.2), limiting conventional illumination

options for breast PAT such as a collimated, expanded beam [24,143]. In further contrast to these systems, the SonixEmbrace uses a thin (~ 5 mm) gel pad for acoustic coupling, as opposed to a flexible membrane or liquid-filled chamber. This means that there is little provision for any divergent optics to expand to cover a significant area of the tissue surface. We explored the possibility of using diode-based illumination, as the smaller footprint of such systems could potentially place the illumination closer to the surface. These systems would provide lower pulse power, so we chose to use a laser-based system in the interest of maximizing fluence, and therefore the attainable penetration depth. We can also use an OPO to perform multi-spectral PAT without changes to the output optics, whereas a diode system would require one illuminator per wavelength.

The mechanical constraints imposed by the configuration of the ABUS electronics and motor led us to consider a narrow, fan-shaped beam that could fit between the multiplexer boards and the motor axle. We 3D printed an identical transducer dome in poly(lactic acid) (PLA; Afinia, Chanhassen MN, USA) which has a polycarbonate optical window epoxied parallel to the transducer, as shown in Figure 3.2.

3.3.2 Laser and Optics

Our illumination system consists of a Continuum Surelite II LASER (Continuum Inc, Santa Clara CA, USA) coupled to an OPO from the same manufacturer to control the output wavelength in the range of 675 nm to 2500 nm. We use a 80/20 beam splitter (68-376, Edmund Optics, Barrington NJ, USA — note that this is the highest-transmission splitter the manufacturer offers for the visible to near-infrared wavelength range) to direct 20% of the OPO output to a USB power meter (EnergyMax J-50MB-YAG - Coherent Inc., Santa Clara, CA, USA), allowing us to measure per-pulse energy while imaging. The remainder of the beam has a diameter of 9.5 mm, and is subsequently homogenized using a cross cylindrical lens array (Nr.18-00142, SUSS MicroOptics, Hauterive, Switzerland) and coupled by means of a plano-convex spherical lens (LA1608, Thorlabs, Newton, United States) into a 1 mm diameter silica core optical fiber. Our beam homogenization and coupling system is described in further detail in [144]. This system provides 5 ns pulses at 700 nm with a per-pulse energy at the output end of the fiber exceeding 50 mJ.

At the output end of the fiber is an optical assembly contained in a 3D-printed housing to maintain alignment and direct the diffuse output toward the optical window in the ABUS dome. The fiber is connected to the optical

3.3. Methods

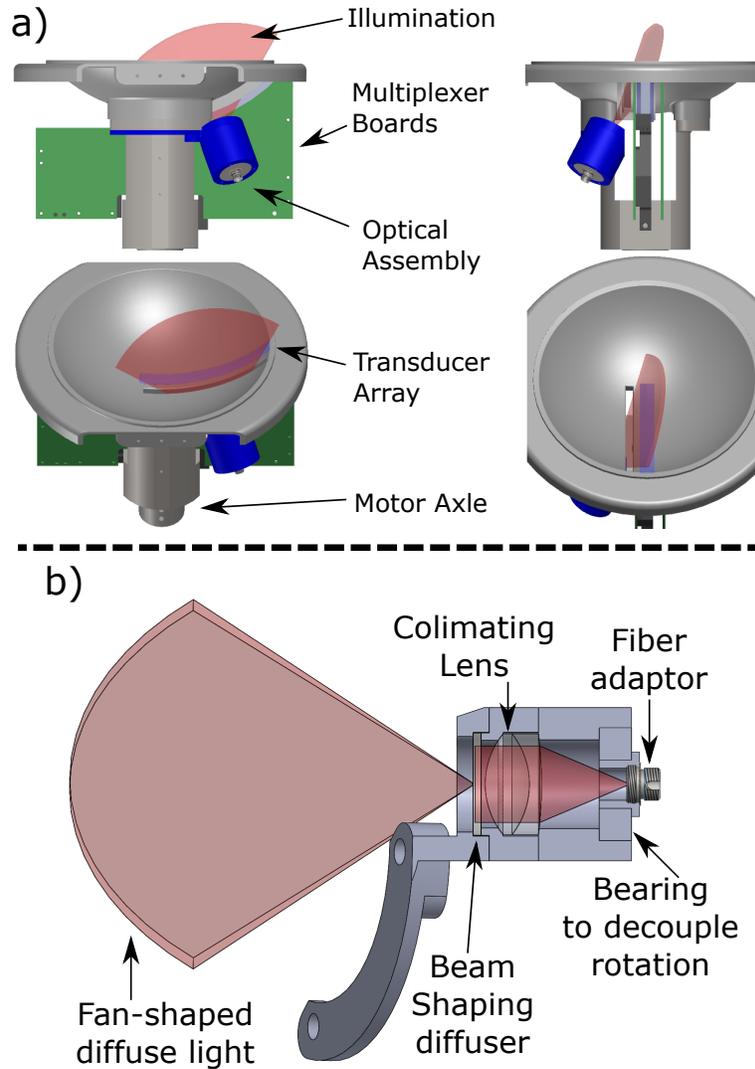


Figure 3.2: Schematic diagrams showing a) the position of the ABUS illuminator and the electromechanical components beneath the ABUS imaging dome and, b) the illuminator, with cutaway view revealing the optical elements. Note that the fan shape of the illumination within the dome in panel a) is only shown here to emphasize the beam geometry. In reality, this structure would be disrupted at the dome interface by optical scattering within the tissue. This is described in detail in Section 3.3.4.

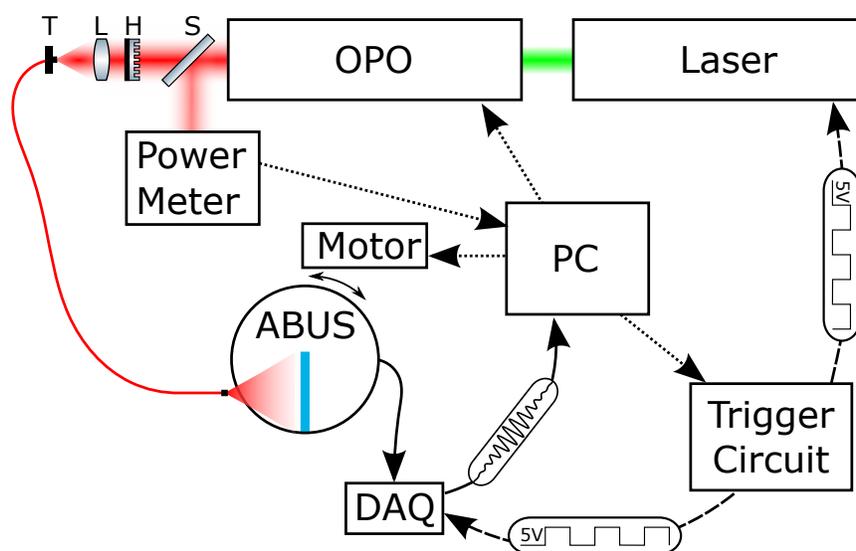


Figure 3.3: The experimental setup used in this study. Optical elements T, L, H, and S are a translation stage, lens, beam homogenizer, and beam splitter, respectively. Dotted lines represent serial communications, and dashed lines are synchronization triggers.

assembly by means of an FC fiber bulkhead connector (Thorlabs, Newton NJ, USA) and a bearing to decouple the fiber from the dome rotation. The fiber output first passes through an achromatic collimating doublet (65-438, Edmund Optics, Barrington NJ, USA) and is then incident upon an engineered “line diffuser” (EDL-100x0.4, RPC Photonics, Rochester NY, USA) which shapes the collimated beam into a fan. This fan has a 90° divergence along one axis, and a 0.24° divergence along the other. This shape permits the beam to pass through the optical window, while illuminating its entire 11.5 cm length. A cross-sectional view of this assembly is shown in Figure 3.2.

3.3.3 Data Acquisition and Pre-processing

Radio frequency (RF) data from the ABUS transducer was acquired using a SonixDAQ module (BK Medical). The SonixDAQ has a sampling rate of at 40 MHz, a 12-bit resolution, and a -10 dB sensitivity in the range 2 MHz to 20 MHz. Since the SonixDAQ can only acquire 128 channels at a time, each plane acquisition requires three acquisitions to cover all of the transducer elements. Control of the ABUS motor, as well as all triggering and synchronization of the optics and DAQ acquisition was accomplished using a custom Arduino-based circuit controlled over USB from a Windows PC.

This circuit allows flashlamp triggers to continually be sent to the laser at its pulse repetition rate of 10 Hz, while illumination pulses (via Q-switch triggers) can be requested via hardware or software triggers. We have found that this “steady-state” operation decreases inter-pulse energy variability. For the pulse energies we employ in this study, this variability is on the order of 10 to 15%. To further minimize the effect of this variability, we use the power meter data to normalize each 128-element acquisition by the relative pulse energy. A diagram summarizing the system configuration is shown in Figure 3.3.

Several steps are needed to pre-process the RF data prior to photoacoustic reconstruction. First, each A-line is low-pass filtered to remove high-frequency components outside the working range of the transducer. This was performed in software using a fifth-order Butterworth filter with a cutoff at 14.5 MHz. We retain the low-frequency components at this stage as they are essential for the following processing step: the application of singular value decomposition (SVD) denoising to remove cross-channel noise bands [106]. To perform this denoising, we compute the singular value decomposition (SVD) of a 128×2080 matrix where each row corresponds to

the RF data for one transducer element. In this representation, this band-shaped noise dominates the first few singular value components. We can isolate these components, and revert to the original representation, producing a matrix consisting of only the noise. This can then be subtracted from the original data, resulting in the denoised matrix. Since these noise bands are stochastic, this decomposition and denoising must be performed for each 128-element acquisition. By using the GPU implementation of SVD available in the *scikit-cuda* Python package [145], this denoising takes only 100 ms for each 128×2080 array. Any low-frequency components outside the working range of the transducer will be removed by a ramp filter described in Equation (3.3).

We explored several other denoising approaches including template matching, directional filtering [146] (both to the radio-frequency and image data), and temporal averaging (i.e. acquiring multiple frames of RF data per transducer position). Of these, temporal averaging was most successful at reducing the prominence of the noise bands, although it comes at a cost of increasing acquisition time. Our 10 Hz laser system and SonixDAQ create significant overhead for each additional acquisition, especially since the ABUS array requires three 128-element acquisitions per transducer position. For a volumetric scan of 200 angular positions, acquiring a single frame per position resulted in a 20 minute scan time, whereas acquiring 10 frames per position increased the scan time to well over 1 hour.

The other denoising approaches were either very sensitive to input parameters, or were unacceptably detrimental to the image quality. We have found that the SVD-based approach works on a wide variety of data with the same input parameters, removes the noise bands with minimal effect on the photoacoustic signals, and is computationally efficient to apply to the data. Appendix A.1 includes an example comparing the effectiveness of directional filtering and SVD-based denoising on our data.

3.3.4 Reconstruction

When a laser pulse with spatial energy distribution $\Psi(\mathbf{r}, t)$ at time t and positions \mathbf{r} is incident upon a sample with coefficient of thermal expansion β , heat capacity C_v , acoustic speed v_s , and optical absorption coefficient $\mu_a(\mathbf{r})$, an acoustic wave is generated with initial pressure distribution

$$p_0(\mathbf{r}, t) = \frac{v_s^2 \beta}{C_v} \Psi(\mathbf{r}, t) \mu_a(\mathbf{r}). \quad (3.1)$$

3.3. Methods

Assuming the medium properties β , C_v and v_s are spatially homogeneous and the laser pulse is temporally short, $\Psi(\mathbf{r}, t) = \Psi(\mathbf{r})\delta(t)$, we have that

$$p_0(\mathbf{r}) \propto \Psi(\mathbf{r})\mu_a(\mathbf{r}) \quad (3.2)$$

at time $t = 0$, when the laser is incident upon the sample. The goal of PAT is to reconstruct $\mu_a(\mathbf{r})$ for the set of positions \mathbf{r} , based on a set of time-varying pressure values, $p_D(\mathbf{r}_D, t)$ measured at discrete detector positions \mathbf{r}_D . This involves first inverting the measurements $p_D(\mathbf{r}_D, t)$ to find the initial pressure distribution $p_0(\mathbf{r})$. For a spherical scanning geometry, Xu *et al* [100] expressed this inversion in a back-projection form as

$$p_0(\mathbf{r}) = \frac{2}{\Omega_0} \int_{\Omega_0} d\Omega_0 \left[p_D(\mathbf{r}_D, t) - t \frac{\partial p_D(\mathbf{r}_D, t)}{\partial t} \right]_{t=\frac{|\mathbf{r}-\mathbf{r}_D|}{v_s}} \quad (3.3)$$

where Ω_0 is the solid angle of the detection surface containing the points \mathbf{r}_D , and $d\Omega_0$ is the solid angle of the surface element at a location \mathbf{r}_D relative to a sample point \mathbf{r} . While this form is exact only in the case in which this surface completely encloses the sample ($\Omega_0 = 4\pi$), the $\frac{d\Omega_0}{\Omega_0}$ term serves as a weight which will somewhat mitigate the effects of the well-known *partial view problem* [100].

Finding p_0 in Equation (3.3) is typically accomplished using back-projection algorithms such as delay-and-sum, and is a well-studied problem [102]. More sophisticated reconstruction algorithms take into account the fact that the *measured* pressure data are not exactly equal to the actual pressure incident on the detector due to factors such as the electromechanical and spatial impulse responses of the detector [102, 147]. We can generalize the measured pressure data as

$$p_D(\mathbf{r}_D, \mathbf{r}, t) = E(t) * S(\mathbf{r}_D, \mathbf{r}, t) * p(\mathbf{r}_D, t) \quad (3.4)$$

where $p(\mathbf{r}_D, t)$ is the actual time-varying pressure incident on the detector at \mathbf{r}_D , $S(\mathbf{r}_D, \mathbf{r}, t)$ is the spatial impulse response (SIR) of the detector at \mathbf{r}_D relative to measurement point \mathbf{r} , $E(t)$ is the electrical impulse response (EIR), and $*$ denotes linear convolution with respect to time. Here we consider only the spatial sensitivity of the transducer, i.e. we assume a flat frequency response, giving

$$\begin{aligned} E(t) &= \delta(t), \\ p_D(\mathbf{r}_D, \mathbf{r}, t) &= S(\mathbf{r}_D, \mathbf{r}, t) * p(\mathbf{r}_D, t). \end{aligned} \quad (3.5)$$

3.3. Methods

Since the SonixEmbrace transducer was designed for ultrasound imaging, it includes an acoustic lens for beam focusing, and as such it would be a poor assumption to fully neglect the SIR, as is common in many PAT reconstruction schemes. It is known, however, that accounting for the SIR is computationally difficult [102, 147]. In an attempt to strike a balance, we separate the SIR into a spatial portion, describing the directional sensitivity, and a temporal portion, describing the averaging effect due to the finite size of the transducer element

$$S(\mathbf{r}_D, \mathbf{r}, t) = S_0(\mathbf{r}_D, \mathbf{r}) S_1(t), \quad (3.6)$$

and proceed under the assumption that $S_1(t) = \delta(t)$. This implies that the elements are point-like, but still have a non-uniform directional response

$$p_D(\mathbf{r}_D, \mathbf{r}, t) = S_0(\mathbf{r}_D, \mathbf{r}) p(\mathbf{r}_D, t). \quad (3.7)$$

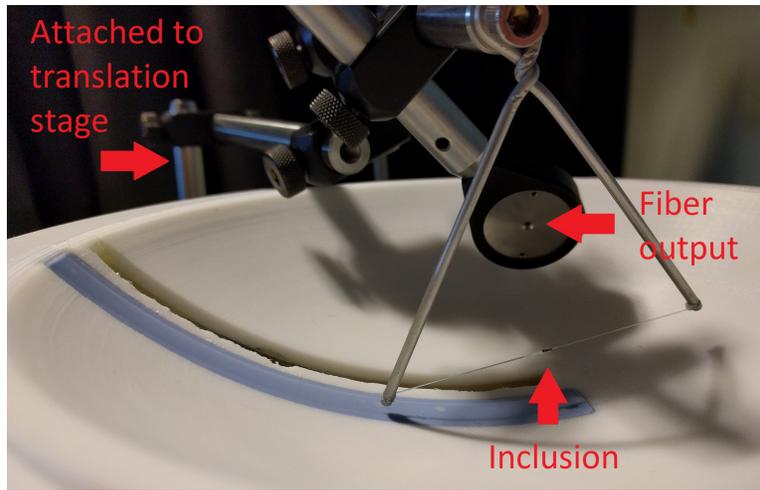
Under this assumption, $S_0(\mathbf{r}_D, \mathbf{r})$ has no time dependence, and as such does not depend on $|\mathbf{r} - \mathbf{r}_D|$ when substituted into Equation (3.3). This assumption also implies that the transducer has a flat frequency response. In the coordinate system of one transducer we can define θ_L and θ_E as the angles between the element normal (the *axial* direction) and the sample point in the lateral (in-plane) and elevational (out-of-plane) directions, respectively, such that

$$p_D(\mathbf{r}_D, \mathbf{r}, t) = S_0(\theta_L, \theta_E) p(\mathbf{r}_D, t). \quad (3.8)$$

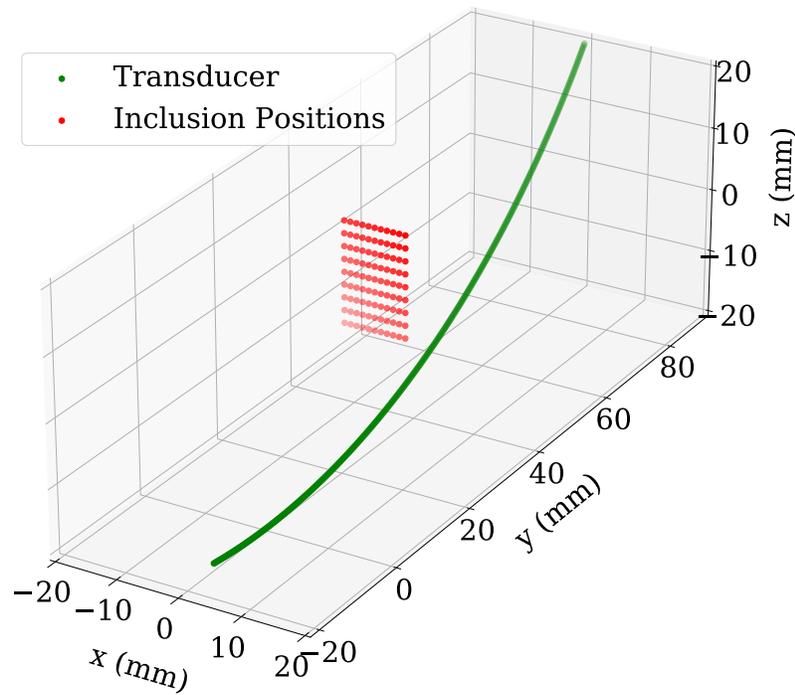
Since we are already computing the total angle between the element normal and the sample point to calculate the solid angle element $d\Omega_0$ in Equation (3.3), $S_0(\theta_L, \theta_E)$ does not appreciably change the computational effort required. Since θ_L and θ_E must be computed for every pair of \mathbf{r}_D and \mathbf{r} ($10^5 \cdot 10^7 \sim 10^{12}$), this approach lends itself well to parallelization using GPUs.

The manufacturer was unable to provide detailed characteristics of the acoustic lens for the purpose of modelling the directional response, $S_0(\theta_L, \theta_E)$, so we measured it experimentally. This was accomplished by attaching a 2 mm long, 0.7 mm diameter graphite rod on a piece of fishing line to a 3 axis translation stage such that it was suspended in the water-filled ABUS dome. We coupled our laser system into the optical fiber as described in Section 3.3.2, and mounted the output end on the translation stage, directed toward the graphite inclusion. This fixes the fiber position relative to the inclusion, ensuring relatively constant illumination as the inclusion is translated. This configuration is shown in Figure 3.4a. Photoacoustic

3.3. Methods



(a) Experimental setup for measuring transducer directivity using photoacoustic imaging.



(b) Inclusion locations for directivity measurement, shown relative to the position of the transducer. Only every second row and every second column of inclusion points are shown here for clarity.

Figure 3.4: The experimental setup used to measure the directional sensitivity of the SonixEmbrace transducer.

3.3. Methods

data were acquired at a total of 200 points, over 1.0 cm in the elevational direction, and 1.6 cm in the axial direction. The lateral position was fixed near the midpoint of the transducer, since there are sufficient elements to constitute a range of lateral positions. These points are shown relative to the transducer position in Figure 3.4b.

We assume that the elevational and lateral portions of the directivity are separable

$$S_0(\theta_L, \theta_E) = S_L(\theta_L) S_E(\theta_E) \quad (3.9)$$

and further, that each portion can be approximated by a Gaussian function of the corresponding angle

$$S_0(\theta_L, \theta_E) = \exp\left(-\frac{1}{2} \left[\frac{\theta_L}{\sigma_L}\right]^2\right) \exp\left(-\frac{1}{2} \left[\frac{\theta_E}{\sigma_E}\right]^2\right) \quad (3.10)$$

where σ_L and σ_E are standard deviations of the lateral and elevational components, respectively.

Non-uniform Illumination

If the illumination source is spatially uniform within the sample, i.e. $\Psi(\mathbf{r}) = \Psi$, Equation (3.2) implies that $\mu_a(\mathbf{r})$ is directly proportional to the initial pressure $p_0(\mathbf{r})$. If the illumination is non-uniform, we must deconvolve these two spatial distributions. Since strong light attenuation in tissue limits imaging depth, it is common in PAT to apply an exponential weighting to the reconstructed $p_0(\mathbf{r})$, to enhance signals further from the surface. Our illumination is not collimated (uniform) at the tissue surface, but can still be accounted for in a similar manner.

With sufficient geometric constraints, we can model $\Psi(\mathbf{r})$ as a solution of the radiative transfer equation for homogeneous medium with temporally short pulses [81]. For an illumination source given by \mathbf{r}_i , the optical fluence at a point \mathbf{r} will be a function of the distance $r = |\mathbf{r} - \mathbf{r}_i|$ from the source:

$$\Psi(\mathbf{r} - \mathbf{r}_i) = \frac{1}{4\pi D r} \exp(-\mu_{\text{eff}} r) \quad (3.11)$$

where the effective attenuation coefficient μ_{eff} , and the photon diffusion coefficient D , are defined as

$$\mu_{\text{eff}} = \sqrt{\frac{\mu_{a0}}{D}} \quad D = \frac{1}{3(\mu_{a0} + \mu_{s'})} \quad \mu_{s'} = \mu_s(1 - g) \quad (3.12)$$

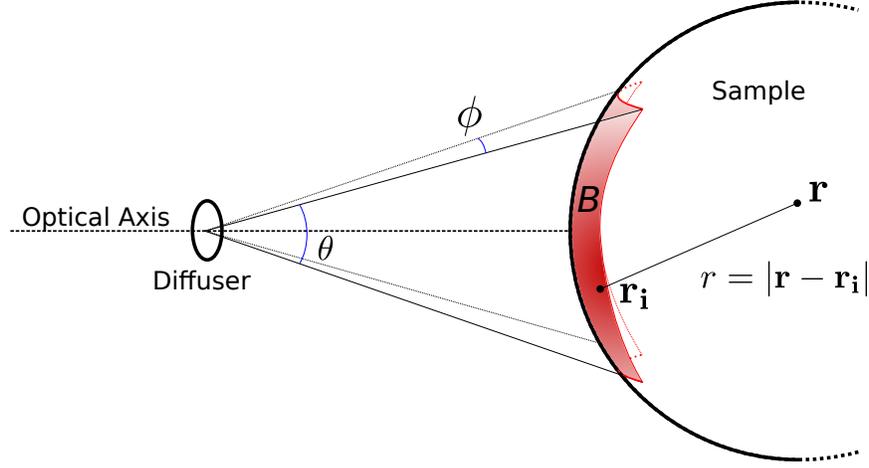


Figure 3.5: The coordinate system for the illumination system, indicating the divergence angles of the beam, and the intersection with the sample surface.

with μ_{a0} and μ_s being the bulk absorptive and scattering attenuation coefficients, and g being the tissue anisotropy. Equation (3.11) holds in the diffusion limit, when $\mu_{a0} \ll \mu_s$. For human breast tissue at 700 nm, $\mu_s = 10.0 \text{ cm}^{-1}$ and $\mu_{a0} = 0.2 \text{ cm}^{-1}$ [134]. Assuming constant optical properties in the tissue constitutes a first order approximation to the true fluence, which would require either prior knowledge of, or joint reconstruction of the scattering and absorption coefficients.

The geometry of the illuminator and the resultant fluence at the sample surface is shown in Figure 3.5. For a given position in the sample, \mathbf{r} , the fluence can be computed by integrating Equation (3.11) over values of \mathbf{r}_i on B , the section of the dome surface bounded by the angles θ and ϕ corresponding to the two diverging axes of the fan-shaped beam:

$$\Psi(\mathbf{r}) = \iint_B \frac{1}{4\pi D r} \exp(-\mu_{\text{eff}} r) d\theta d\phi. \quad (3.13)$$

Additional calculations for the geometric optics describing the diffusive optics, and reflective losses at the dome interfaces are included in Appendix B.

3.3.5 Illuminator Fluence Measurement

To estimate the overall losses in our fiber coupling and diffusing system, and to check our calculated illumination map at the dome boundary, we

measured the fluence at the sample side of the optical window. This was accomplished by scanning the same power meter described in Section 3.3.2 along the arc of the window in 0.25 mm increments, and measuring the pulse energy at each position.

3.3.6 Reconstruction Implementation

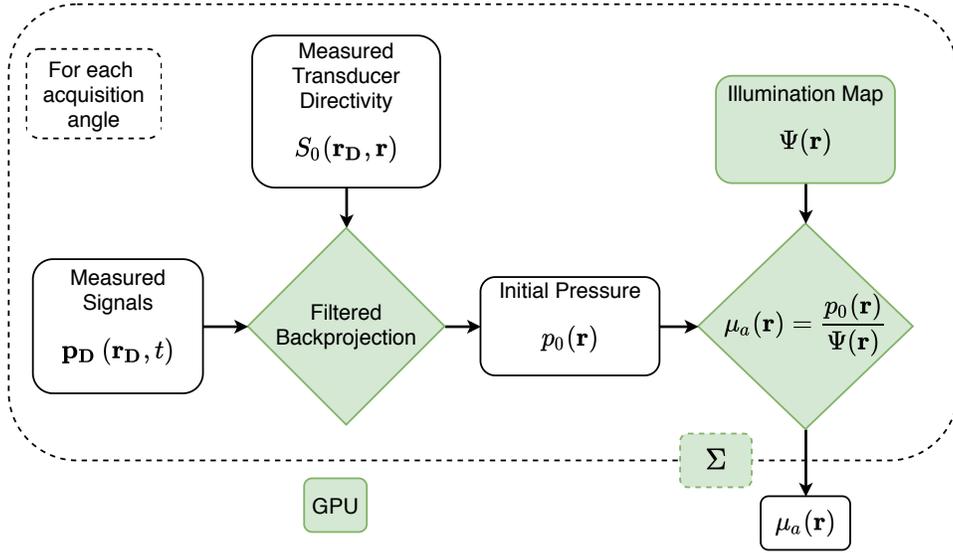


Figure 3.6: A schematic outline of our reconstruction method. Green boxes indicate operations computed on GPUs.

Figure 3.6 shows a summary of our reconstruction scheme. We first filter and back-project the measured signals to recover the initial pressure, taking into account the spatial sensitivity of the transducer array. We then compute the illumination map and use it to normalize the initial pressure, recovering the optical absorption map for the current illuminator position. It is important to reiterate that since the illuminator is fixed relative to the transducer and not the sample, $\Psi(\mathbf{r})$ must be recalculated for each acquisition angle. By applying this method for each angle, and combining the results, the optical absorption coefficient in the entire 3D tissue volume can be reconstructed. This accumulation over acquisition angles means that the final estimate of μ_a can not be linearly separated into Ψ and p_0 as would be the case with systems employing stationary illumination.

For all results in the present study, we collected RF data at 200 planes with a constant angular spacing of 1.8° . We structured our GPU reconstruction into a number of blocks equal to the number of image voxels, $i = 1, \dots, N$, with each 384 element acquisition plane divided into 384 threads. This was performed sequentially for each of the planes, with the results being accumulated in GPU memory.

We compute the illumination map given by Equation (3.13) at each of the N image voxels by dividing the surface B into small patches and summing their contributions to the fluence within the sample. These contributions are independent and can thus be computed in parallel on a GPU. For the present study, we divided B into 200 patches, each of which was computed in a separate GPU thread, resulting in a $200\times$ reduction in computation time. The number of patches was chosen empirically as a compromise between the number of patch coordinates which need to be pre-computed and stored, and eliminating discretization artifacts in the illumination map, especially near the dome surface.

Reconstruction was implemented in Python, with the following computationally intensive portions implemented in CUDA: back-projection, filtering, SVD denoising, and calculation of the illumination map $\Psi(\mathbf{r})$. Reconstruction was performed on a 64-bit Windows 10 PC with 16 GB of memory, an Intel Core i7-7700 processor (Intel Corporation, Santa Clara CA, USA), and a GeForce GTX 1060 3GB GPU (Nvidia Corporation, Santa Clara, CA, USA). The CUDA implementation of the back-projection is included in Appendix C.

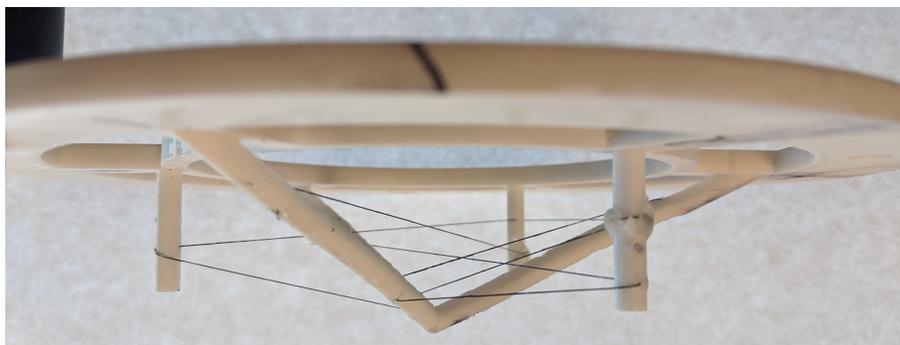
All reconstructions in this paper have a resolution of 0.5 mm and assume a constant acoustic velocity of 1500 m s^{-1} . To calculate the illumination profile using Equations (3.12) and (3.13), we used the scattering and absorptive attenuation coefficients of milk at 700 nm [133] - $\mu_s = 30 \text{ cm}^{-1}$ and $\mu_{a0} = 0.015 \text{ cm}^{-1}$ for the calibration phantom, and human breast tissue at 700 nm [134] - $\mu_s = 10.0 \text{ cm}^{-1}$ and $\mu_{a0} = 0.2 \text{ cm}^{-1}$ for the chicken breast phantom.

3.3.7 Phantoms

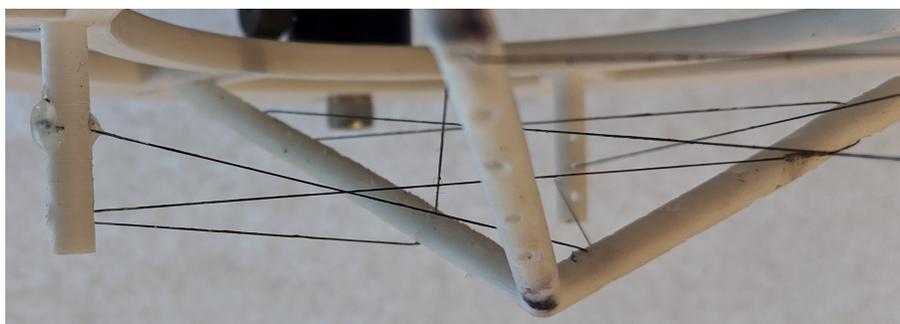
To provide a robust and reproducible way to test our entire imaging system and reconstruction scheme, we designed and built a modular photoacoustic phantom.

As an absorber, we used 0.25 mm diameter monofilament fishing line, coated with black spray paint. This provides a strong photoacoustic signal and the small size serves as a test of the system resolution. To suspend the

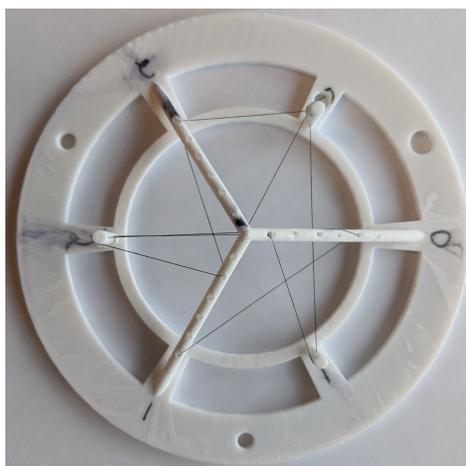
3.3. Methods



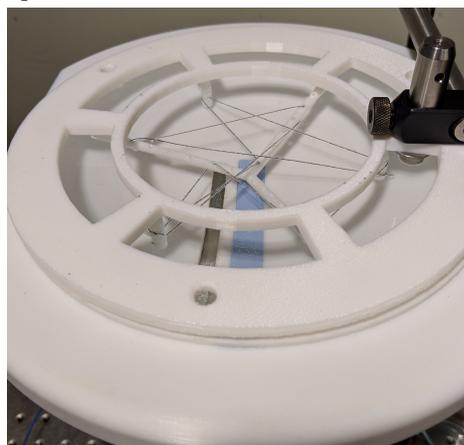
(a) Side-on view.



(b) Close-up view.



(c) Top-down view.



(d) In place on the transducer dome during an experiment.

Figure 3.7: Images of the calibration phantom used in this study. The white template is 3D printed plastic, and black painted fishing line serves as a photoacoustic source. The circular portion of the template has an outer diameter of 19 cm.

fishing line in the imaging volume, we designed and 3D printed a scaffold offering multiple mounting points between which the line can be extended. The scaffold consists of a base which precisely aligns the phantom with the imaging volume, and several cylindrical columns with regularly-spaced holes through which the fishing line can be threaded. This scaffold is shown in Figure 3.7. As shown in Figure 3.7, the fishing line can be threaded through multiple points, providing a complex inclusion geometry which is nonetheless simple to describe and reproduce. It is important to note that only the black surface of the wires will function as a photoacoustic absorber, rather than the entire cross-section which would be the case for a solid coloured inclusion, or a blood vessel. Since our primary quantification metric in this study is SNR rather than resolution, this is sufficient for our purposes.

To image the calibration phantom, the ABUS dome was filled with a 50/50 mixture of water and milk, to provide optical scattering. To account for this mixture, we scaled the scattering and absorptive attenuation coefficients in our reconstruction by a factor of two, giving $\mu_s = 15 \text{ cm}^{-1}$ and $\mu_{a0} = 0.0075 \text{ cm}^{-1}$.

An additional advantage to using a 3D printed template is we can easily use the design file to generate the inclusion geometry for use in simulations and validation of our reconstructed data.

To test our system in a more highly-scattering medium, we imaged a piece of 0.5 mm diameter pencil graphite embedded 2 cm deep in a piece of chicken breast tissue.

3.4 Results

3.4.1 SVD Denoising

Figure 3.8 shows an example of the effect of the denoising method described in Section 3.3.3 on representative RF data. For our data, discarding the first 10 singular values produced the best balance between removing noise and maintaining signal. This will vary depending on the imaging target, and is determined empirically by examining an A-line containing both noise bands, and a known inclusion signal. An example of this process is included later, in Figure 4.3.

Note that in the resultant “denoised” data, while the absolute photoacoustic signal intensity is increased relative to the noise, the structure of the photoacoustic is significantly changed. Namely, the negative portions have been removed, and the oscillating signal has merged into one broader peak.

3.4. Results

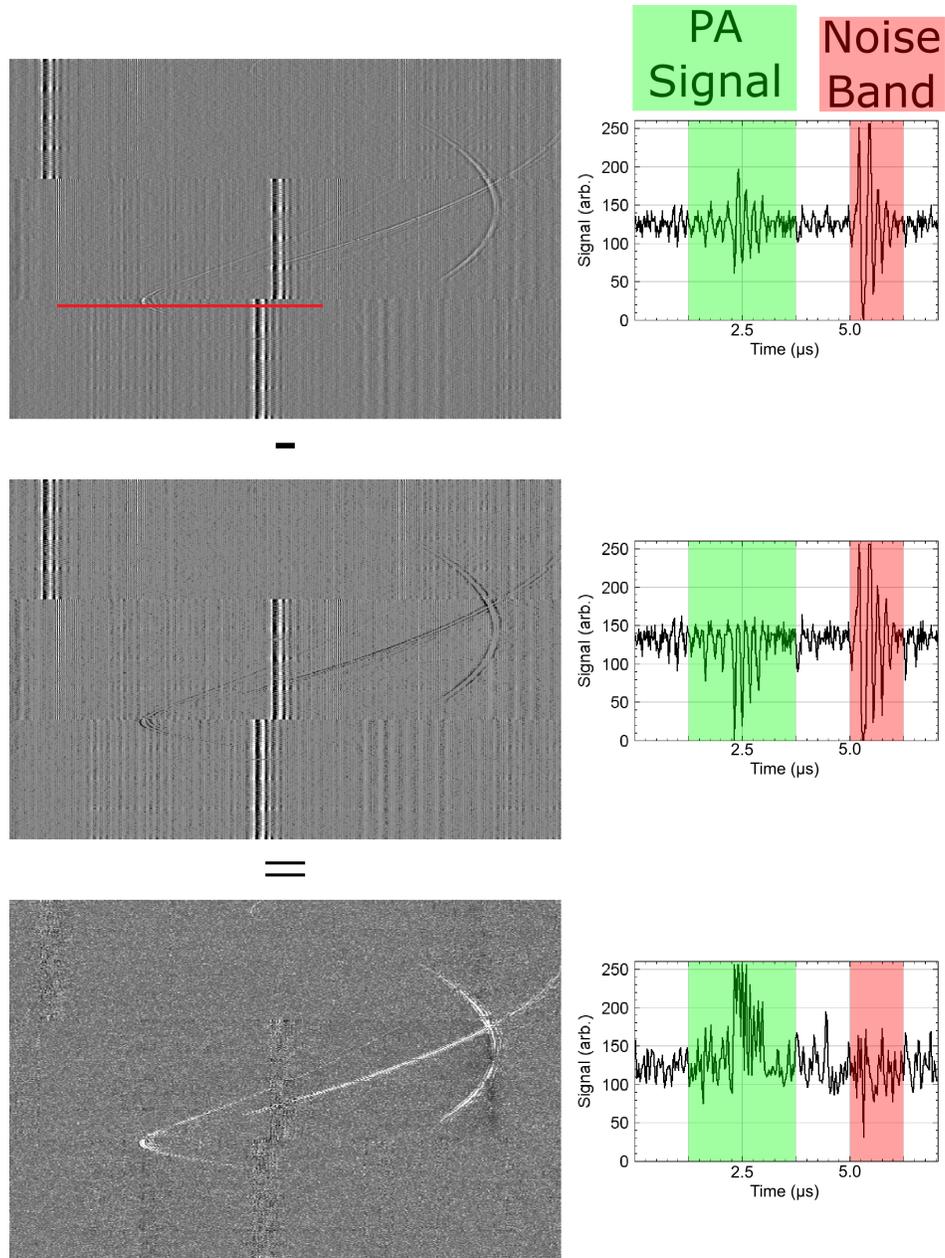


Figure 3.8: The effect of SVD denoising on the RF data. Right panels show the signal along the red line for each of the three images, with a representative photoacoustic signal, and a representative noise band indicated in green and red, respectively. Noise signal corresponds to the first 10 singular values. The top row shows the raw data without processing. The middle row shows the isolated noise. The third row is the subtraction of the two. 62

3.4.2 Transducer Directivity

We observed a Gaussian dependence of the signal intensity on the angle between the transducer normal and the inclusion position, both in the lateral and elevational direction. The measured standard deviations of these histograms, σ_L and σ_E , corresponding to the lateral and elevational angles θ_L and θ_E were 18.5° and 6.2° . These data and the associated fits are included in Figure 3.9. While this simple Gaussian model fits the overall shape of each histogram, it does deviate in two notable ways. The lateral distribution has a shoulder near 10° , and might be better described using a model with two regimes. The elevational data is of low quality for very small angles, potentially overlooking a more complex dependence. The chosen inclusion points allowed both of these measurements to be performed simultaneously, but the elevational data would benefit from an additional, more precise experiment focused in the range of -5° to 5° .

3.4.3 Combined Imaging

Figure 3.10 shows representative in-plane images of the calibration phantom, demonstrating coregistered B-mode and photoacoustic imaging. For clarity of display when overlaid on other modalities, the photoacoustic data were thresholded such only pixels with amplitudes in the top 10% of the image histogram are visible, removing most of the background noise. The full dynamic range of the non-thresholded photoacoustic image was 10 dB. Panel c) of this figure shows model inclusion data derived from the design file for the 3D printed wire phantom, which further confirms that the designed structure of the phantom matches what we see when imaging. Note that the diameter of the inclusions in the model data in Panel c) are exaggerated $10\times$ to be more easily visible. There are two circular features in the model data, corresponding to wires oriented normal to the imaging plane. These features, as well as some sections of the X-shaped feature, are not visible in the photoacoustic data. This could be due to shadowing of the illumination light by the phantom support structures.

We used two measures to quantify the agreement between any two of these registered modalities. The first was the root mean square error (RMSE), defined as

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (A_i - B_i)^2} \quad (3.14)$$

where A_i and B_i are the value of the i 'th pixel in the two images, and N is

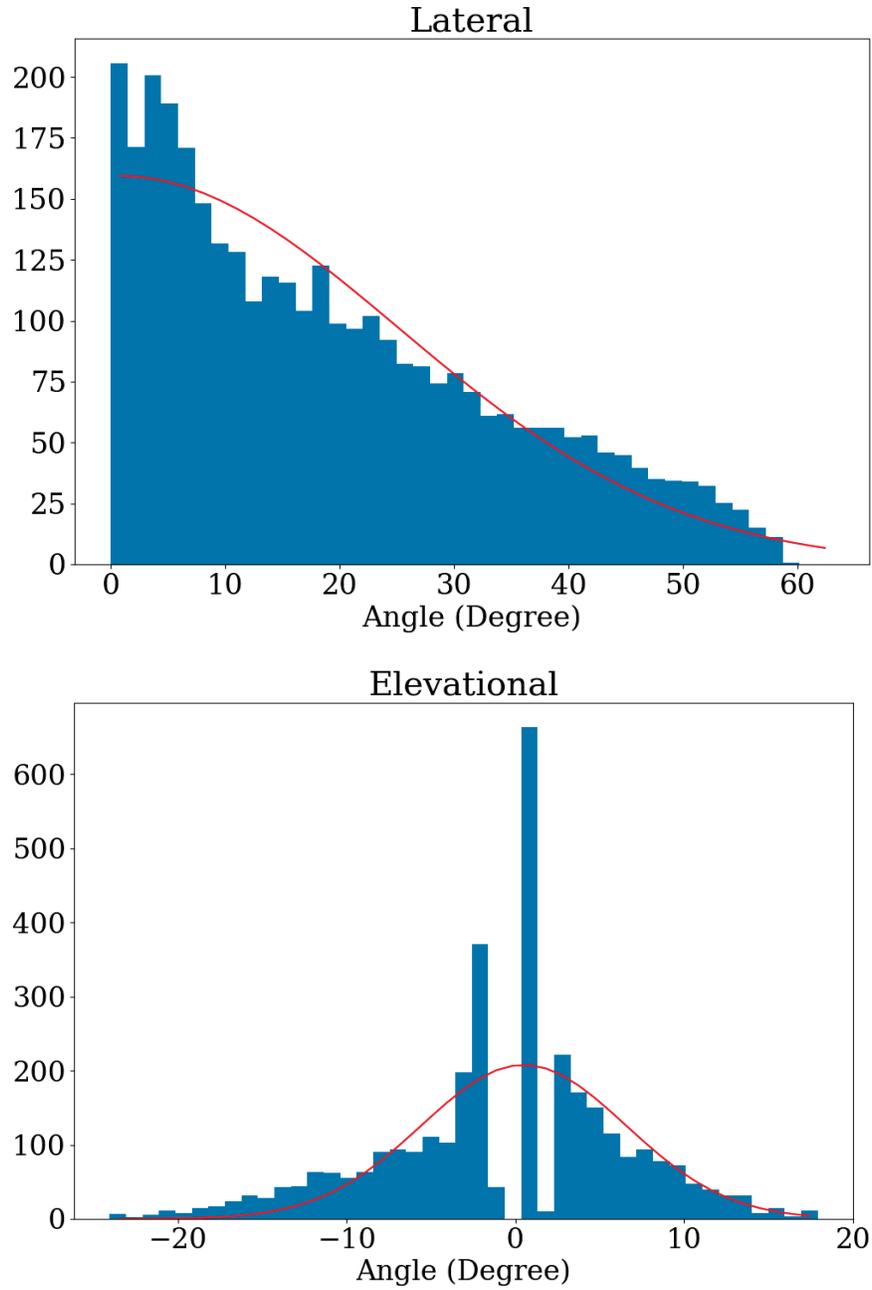


Figure 3.9: The sensitivity of the SonixEmbrace transducer elements as a function of the elevational and lateral angle between the element normal and the imaging point. Red lines indicate Gaussian fit.

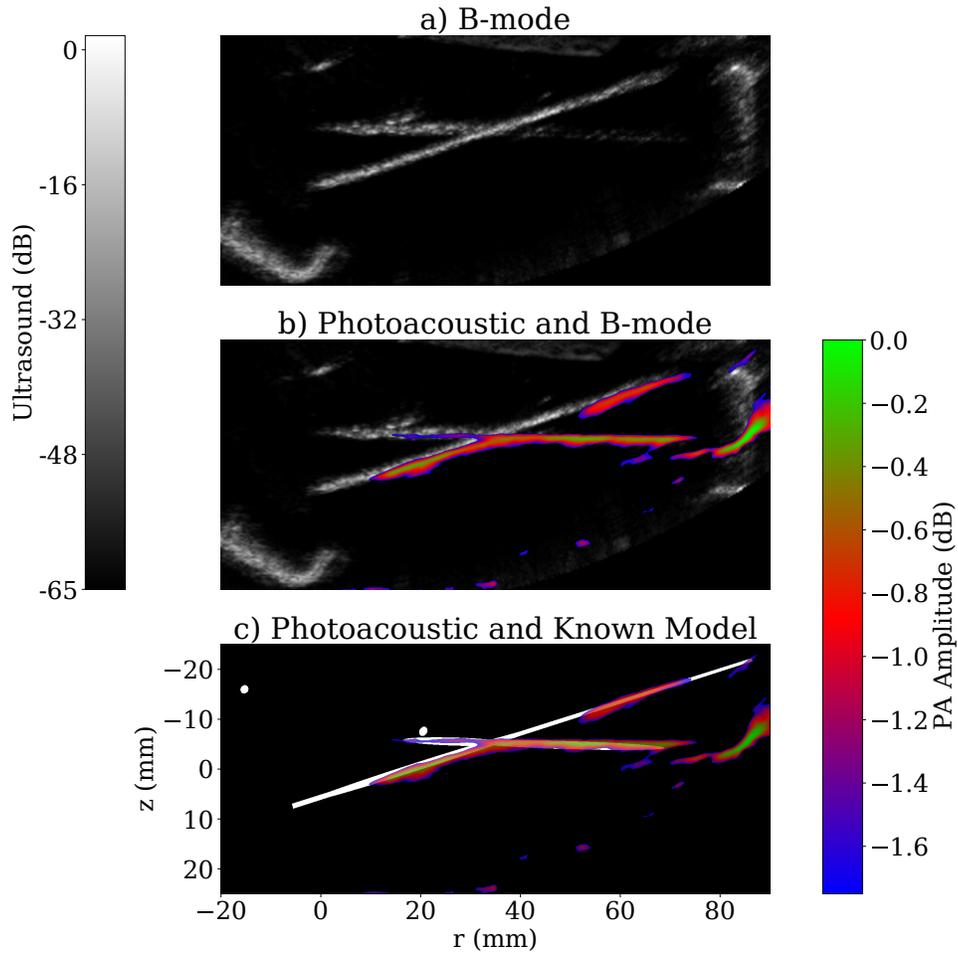


Figure 3.10: B-mode and photoacoustic data for a representative plane of the calibration phantom, indicating registration with modelled inclusion positions derived from the wire phantom design file. Photoacoustic data only shows the top 10% of the image histogram, such that it does not fully obscure the underlying images.

3.4. Results

the number of pixels in the image. The second metric used was the structural similarity index measurement (SSIM) defined by Wang *et al* [148]. We used the Python implementation of this function from the *scikit-image* package [149]. Support structures from the phantom template (see Figure 3.7) visible in the bottom left corner and far right side of the B-mode image were cropped out before performing these measurements. Otherwise, these measurements are performed on the entire region of the images visible in Figure 3.10. These comparisons are summarized in Table 3.1.

	B-mode vs. Model	PA vs. Model	B-mode vs. PA
RMSE	2.66	2.33	3.31
SSIM	0.48	0.78	0.40

Table 3.1: Registration errors between imaging modalities.

3.4.4 Illuminator Spatial Profile

Figure 3.11 Panel a) shows the calculated illumination structure within the sample, as per Equation (3.13), while Panel b) shows the measured fluence at the sample side of the optical window alongside the modeled fluence for the same location. The broader shape of the measured data is due to the 1 cm aperture of the power meter used. Integrating the fluence over the entire window indicates a total energy per pulse of around 35 mJ, down from 100 mJ at the OPO output. If this same laser pulse were collimated to the size of the dome (surface area 250 cm²), the average fluence at the tissue surface would be 0.14 mJ cm⁻².

3.4.5 Reconstruction Comparison

To test our reconstruction scheme, we reconstructed our measured data both with and without accounting for transducer directivity and non-uniform illumination. We then measured SNR of a known inclusion for each method. We measure SNR here by extracting the photoacoustic amplitude along a line perpendicular to the long axis of the inclusion. We find the maximum of this signal, and define the “peak” as the region where the amplitude is greater than 50% of this maximum. We then define the SNR as the mean amplitude of the signal in this region, divided by the standard deviation of

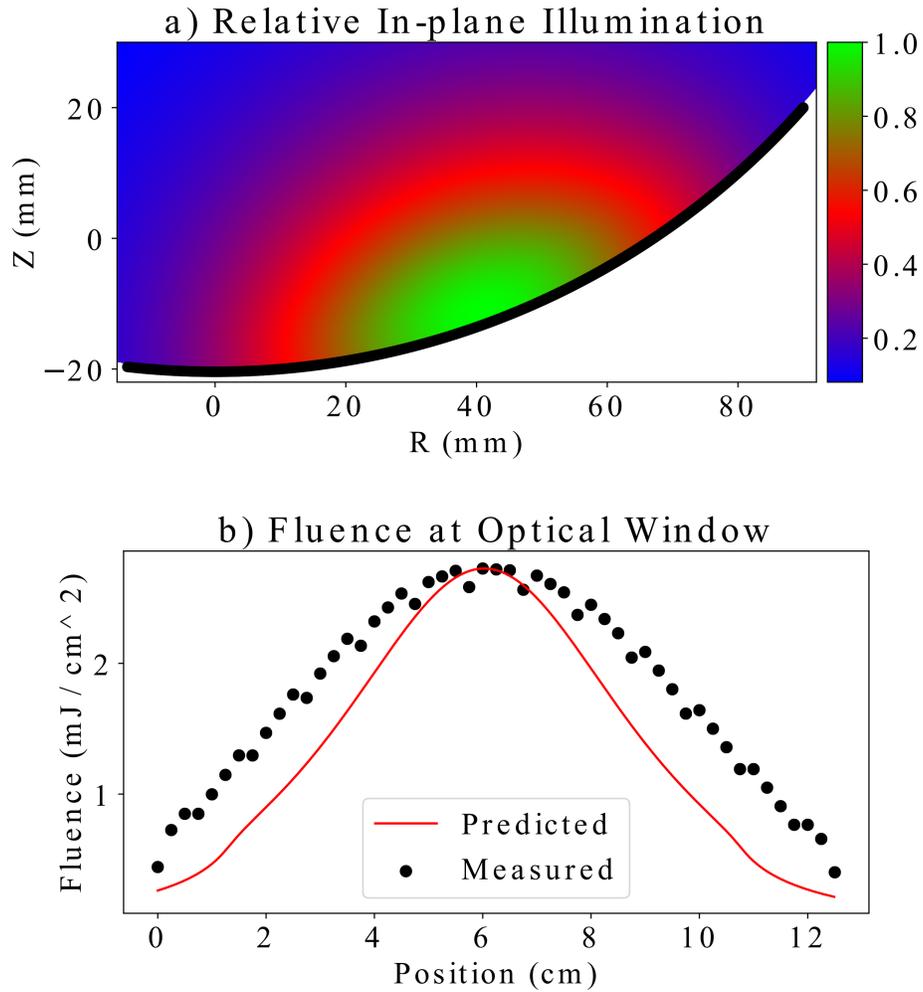


Figure 3.11: Optical fluence for one acquisition plane. Panel a) shows the two dimensional structure of the illumination, relative to the transducer face, which is shown in black. Panel b) shows the measured fluence at the sample side of the optical window in black, compared to its predicted value, as a function of the position along the arc of the window.

3.4. Results

the amplitude over the entire image [150]. We found that this definition, using a peak region rather than a single maximum value, and using the entire image for the standard deviation rather than a representative background selection, most accurately represented the qualitative changes visible in the images. This method was also least sensitive to the choice of perpendicular line position, therefore not inadvertently favouring one reconstruction method over another. The width of the peak region is the full width at half maximum (FWHM) and provides a measure of the resolving power of the system. Figure 3.12 shows an example of this measurement from the chicken breast phantom. A comparison of the measured values is given in Table 3.2. The four panels in Figure 3.12 have been scaled to the same dynamic range for display to make qualitative differences in the images more clear.

	No Directivity	Directivity
No Illumination Compensation	SNR: 6.3 dB FWHM: 1.8 mm	SNR: 7.1 dB FWHM: 1.3 mm
Illumination Compensation	SNR: 7.3 dB FWHM: 1.5 mm	SNR: 8.0 dB FWHM: 1.2 mm

Table 3.2: Comparison of reconstruction methods applied to chicken breast phantom (see Figure 3.12)

To further quantify the performance of our reconstruction scheme, we compared the image error in this reconstruction as a function of imaging depth in our calibration phantom. Using the known model for the inclusions, M_i , we define image error as

$$\text{ERR} = \frac{|A_i - M_i|}{|M_i|} \quad (3.15)$$

where A_i is the reconstructed absorption value of the i 'th voxel, and $|\dots|$ is the two-norm. The model data is derived from the design files for the phantom template and the ABUS transducer dome. This provides a list of coordinates for all potential wire attachments points, in the frame of reference of the transducer. By drawing lines between these points, a 3D binary mask can be generated to indicate which voxels would fall within the wire diameter. This binary mask represents the model data, M_i . We performed this measurement both with and without accounting for the directivity of the transducer, and with and without accounting for the non-uniformity of the

3.4. Results

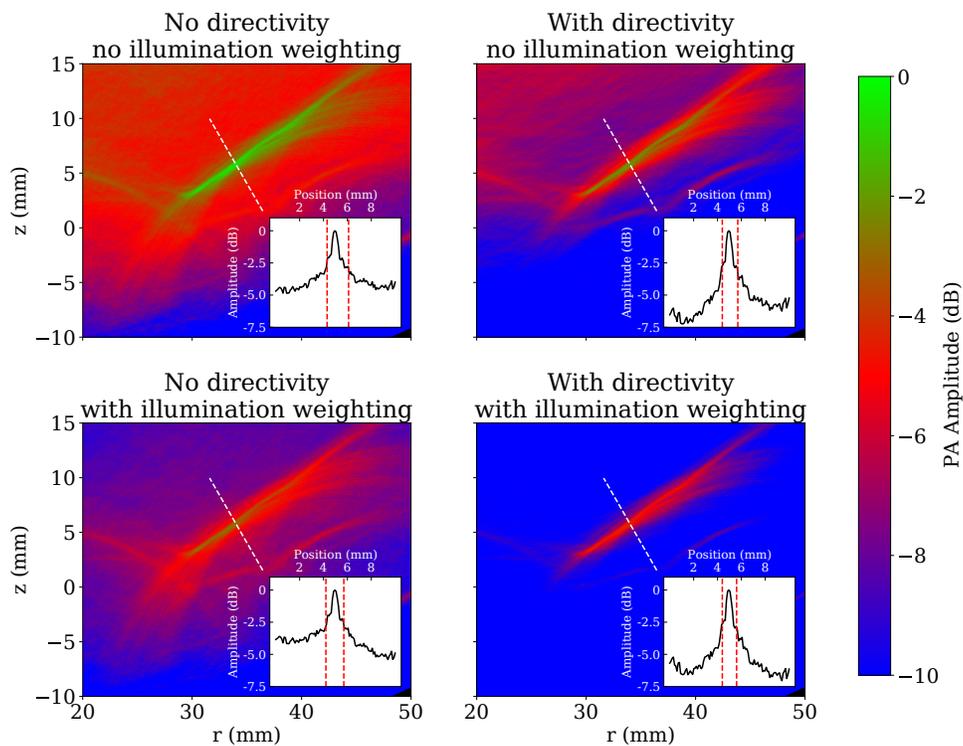


Figure 3.12: Comparison of reconstruction methods applied to chicken breast phantom. Insets show the PA amplitude along the dashed white line, normalized to the same dynamic range across all four images, with vertical dashed red lines indicating the FWHM of the peak. Corresponding measurements are summarized in Table 3.2.

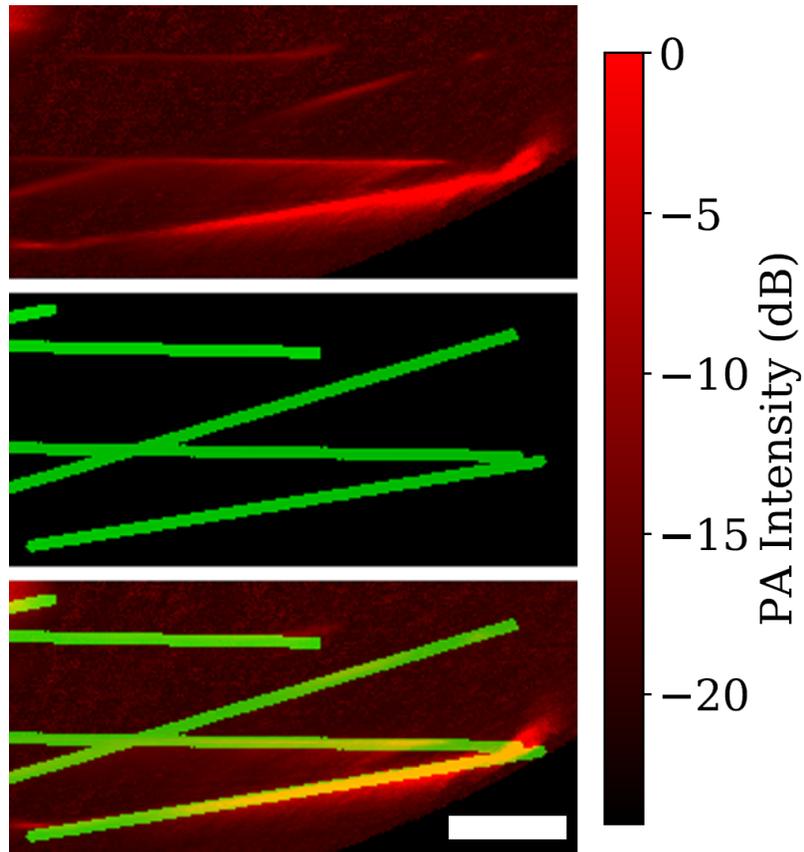


Figure 3.13: Maximum intensity projections of a tomographic reconstruction of the calibration phantom. The top panel shows the data reconstructed with directivity and illumination compensation, the middle panel shows the model data, and the bottom panel shows an overlay of the two other panels, to demonstrate the registration of the tomographic data to the model. White scale bar indicates 1 cm.

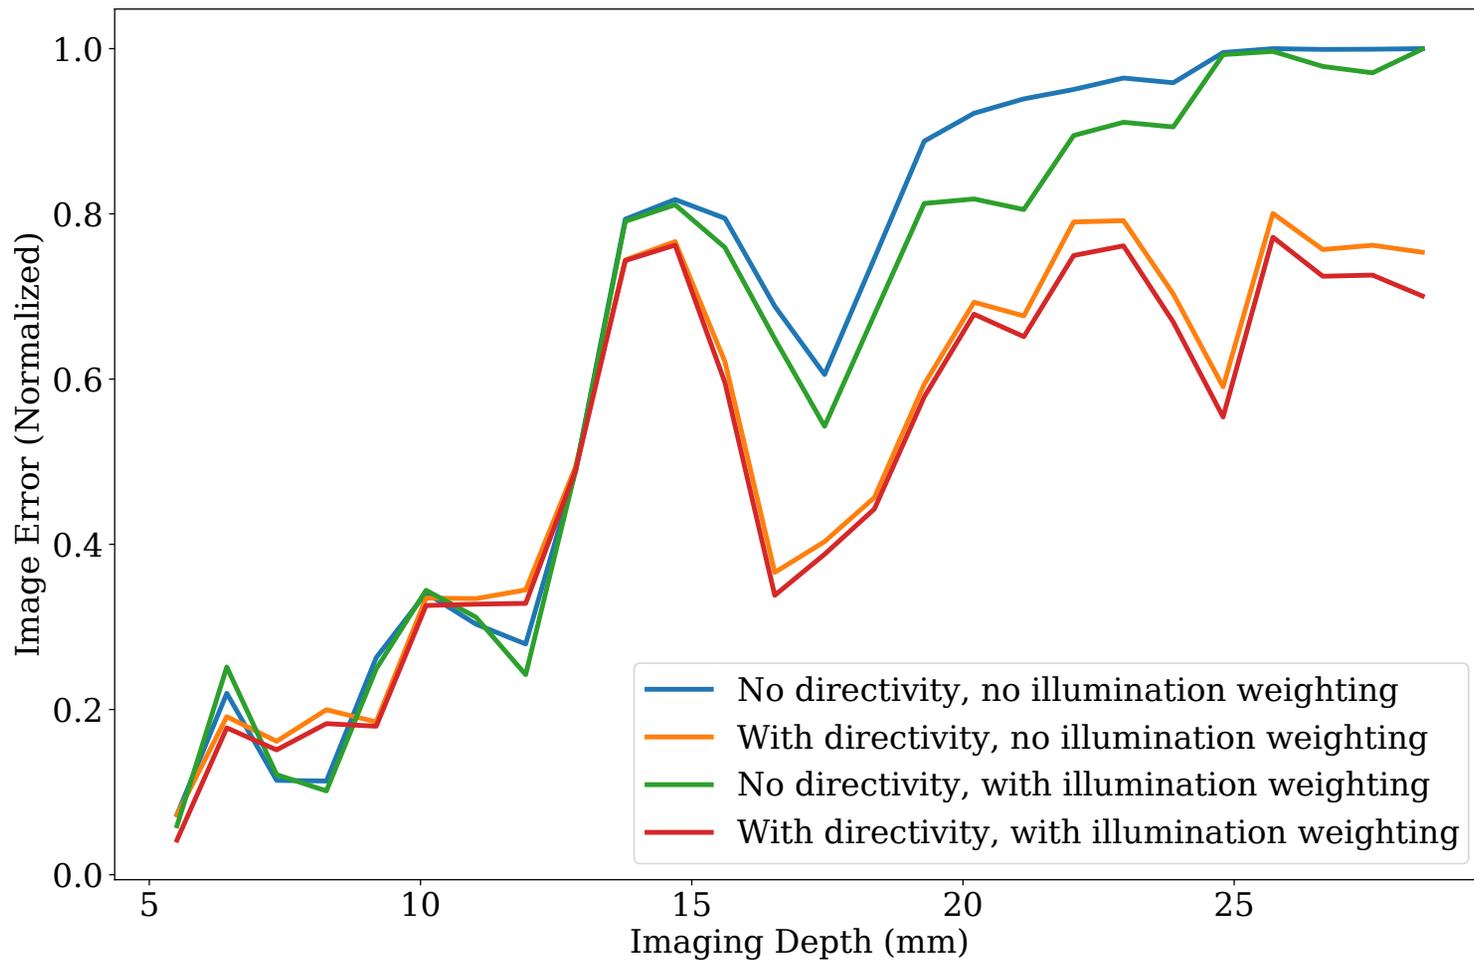


Figure 3.14: Image error relative to known model data, M_i as a function of imaging depth, computed by segmenting the image into 1 mm thick, constant depth shells. Compared for reconstruction schemes with and without transducer directivity and/or non-uniform illumination compensation.

illumination. To measure the depth dependence, we divide the imaging volume into 1 mm thick concentric spherical shells which each correspond to a particular distance from the transducer surface. We then compute the image error in each of these shells using Equation (3.15). These data are summarized in Figure 3.14. Figure 3.13 shows maximum intensity projections of a tomographic reconstruction of our calibration phantom, alongside the model data predicted from our design files. We note a broader dynamic range displayed in this data since the calibration phantom was imaged in a medium with lower light attenuation than the chicken breast tissue. The dynamic range here is also broader than the photoacoustic portion of the multimodal imaging shown in Figure 3.10 because that 2D data was reconstructed from a single transducer position, whereas in this tomographic reconstruction, multiple transducer positions (and therefore many more transducer elements) will contribute to a given voxel, improving the SNR.

3.5 Discussion

Our combined B-mode and photoacoustic imaging results, as shown in Figure 3.10 show qualitatively a strong agreement between both modalities and the model data, indicating no significant errors between the designed phantom template and the final image data. This reinforces a primary advantage of the ABUS as a multi-modal imaging platform, as no registration procedure was required between these data. Figure 3.10 does not, however, show a perfect alignment of the photoacoustic image with either of the other datasets. Given that the photoacoustic signal is consistently shifted onto the side of the inclusion nearest to the transducer, this is potentially due to our use of coated rather than solid absorbers. There is also the potential for positional errors in the wires due to imperfections in the 3D printing process, or small offsets at the attachment points.

As indicated by the registration metrics in Table 3.1, we see that the photoacoustic data is in closer agreement with the model, evidenced by a lower RMSE, and a higher structural similarity value. We expect this is likely due to the photoacoustic data being similarly sparse to the model data, in comparison with the B-mode data, which significantly broadens the wire diameter, and contains more background speckle.

Comparing our imaging results in chicken breast tissue, we see that our reconstruction method offers a significant improvement both to the qualitative appearance of the images in Figure 3.12, and a quantitative improvement to the SNR values and resolution of an inclusion at 2.0 cm imaging

depth, as summarized in Table 3.2. We see that including either transducer directivity or illumination compensation increases SNR and improves resolution, whereas the best results are achieved when both are included.

Our best-case measured SNR of 8.0 dB is comparable to the results of Wray *et al* [151], who achieved SNR values of 6.1 dB and 9.2 dB at depths of 1.5 cm and 1.8 cm respectively in human muscular tissue. This study employed a transducer optimized for photoacoustic imaging. As such they are able to attain a much higher resolution of 255 μm , compared to our resolution of 1.2 mm. While muscular tissue is more highly attenuating than fatty tissue such as the breast, we chose this study because the authors had a similar constraint of a single laser pulse per acquisition.

Our calibration phantom image error shown in Figure 3.14 shows at imaging depths greater than 1.5 cm, both directivity and illumination compensation offer an improvement over the standard reconstruction. The effect of directivity is more significant, but the combination of the two consistently produces the smallest error. Adding both compensation terms improves the image error by an average of 18% over all imaging depths, and by an average of 31% over imaging depths greater than 1.5 cm. It is to be expected that our method would offer the greatest improvement at larger imaging depths, since that is where the actual fluence will differ most from the assumption of uniform illumination.

We note that since the illumination map Ψ must be computed for each acquisition angle, it adds significant computation time to the reconstruction. For a tomographic reconstruction of the entire ABUS volume from 200 transducer planes, adding non-uniform illumination compensation increases the computation time from ~ 5 minutes to ~ 1 hour.

Our calibration phantom design has the advantage that its physical description is fully determined by the attachment points of the wires in the 3D printed structure, which are known from the CAD design. While a vascular structure, mimicking actual breast images as presented by Schoustra *et al* [143], for example, would have been more relevant to the ABUS, uncertainty in the inclusion locations would have made the computation of metrics such as the image error much more difficult. Our next step is to acquire real images from healthy volunteers as performed in that study.

While our PAT system is complete and functional, our images still fall well short of the state of the art, with a relatively low SNR being only modestly increased by our improved reconstruction method described here. Since we have few straightforward options for upgrading our imaging or illumination hardware within the scope of this thesis, our best path forward is through improved data pre-processing and image reconstruction.

Despite the effectiveness of the SVD denoising method, it does inherently remove some portion of the desired signals in the process. As such, we put significant effort into investigating the source and nature of the noise bands, in the hope of eliminating them at the source. We ruled out most sources of electronic interference, including the laser, which is a common source of similar noise in photoacoustic data [106], by turning these sources off, and using ferrite beads on the transducer cabling. We also ruled out any contribution from the SonixDAQ by replacing it with a Verasonics Vantage system (Verasonics, Inc, Kirkland WA, USA) using their Ultrasonix transducer adapter (UTA 156-U).

Through a process of elimination, removing or replacing each piece of hardware in turn, we found that these bands are somehow characteristic of Ultrasonix transducers, and not only the SonixEmbrace. We also found these noise bands present when using an Ultrasonix L-14 linear probe. We expect this is largely overlooked since it only becomes apparent when measuring extremely weak signals, as we are attempting to do.

Further compounding this issue, our initial attempts at implementing a more sophisticated regularized reconstruction scheme were ineffective due to these noise bands being of similar intensity to the photoacoustic signals.

The other potential avenue for improving our data is to average over multiple acquisitions, decreasing stochastic noise. However, since our per-breast scan time is 20 min when only acquiring a single frame per transducer position, acquiring multiple frames would multiply this scan time, and reduce the potential for eventual clinical translation of this system.

To this end, Chapter 4 focuses on developing a more effective denoising method, and testing it using our implementation of a state-of-the-art regularized reconstruction scheme.

3.6 Conclusions

We have developed an illumination scheme for photoacoustic tomography which can be implemented in cases where there is insufficient space to illuminate the entire surface of the breast, as is the case with the SonixEmbrace ABUS. This illuminator can attain a fluence of up to 2.5 mJ cm^{-2} at the tissue surface. We have demonstrated the use of this PAT system to image the 0.25 mm features of our calibration phantom at depths of up to 3 cm in a background with a scattering coefficient $\mu_s = 15 \text{ cm}^{-1}$, which is similar to the value of 10 cm^{-1} in human breast tissue.

We have also devised a reconstruction scheme which can account for this

3.6. Conclusions

non-uniform illumination, attaining a 25% increase in SNR at 2 cm in chicken breast tissue when compared to a reconstruction scheme which assumes spatially uniform illumination. We have quantified the improvement this reconstruction scheme affords, such that we can apply it on a case-by-case basis depending on whether image fidelity or reconstruction time needs to be prioritized. Implementing this algorithm on a GPU provided a $200\times$ speed improvement over the same algorithm when executed on a CPU, for an identical imaging configuration (number of acquisition planes, image grid size, etc.).

Chapter 4

Denoising of Pre-beamformed Photoacoustic Data Using Generative Adversarial Networks³

4.1 Overview

Since the intensity of generated photoacoustic signals is proportional to the dose of laser light reaching the imaging target [60], a fundamental tradeoff exists between maximizing this intensity while minimizing the laser dose to which the patient is exposed [99]. Although different reconstruction schemes including modified delay-and-sum techniques [153], filtered back-projection [100] and iterative approaches [101, 102, 103] have been developed to enhance the quality of the reconstructed images, it is still common to average the raw RF ultrasound data generated from multiple consecutive laser pulses, reducing stochastic noise and improving the SNR [104, 105]. This approach not only increases the scan time proportionally to the number of frames acquired at each imaging position, lowering the frame rate of the system, but it also increases the total laser dose. More specific noise-reduction techniques such as bandpass filtering to the working range of the ultrasound transducer [100], or the removal of laser-induced noise using SVD denoising [106] have also been employed to improve the quality of photoacoustic RF data. In this paper we propose a deep learning based method that is capable of reducing Gaussian background noise similar to averaging multiple acquisition frames, while simultaneously replicating more sophisticated denoising, in particular the SVD denoising method of Hill *et al.* [106], which

³This chapter is previously published [152], and has been modified and expanded upon for this thesis.

we have previously found to be particularly effective for our data [135].

GANs [154] and conditional generative adversarial networks (cGANs) [155] emerged in 2014, and were quickly adopted as the state-of-the-art deep learning models for a variety of tasks in different fields [156]. cGANs capable of mapping different image domains, such as the Pix2Pix model of Isola *et al.* [157] for instance, can be used for tasks such as segmentation, grayscale to RGB transformation, and super-resolution reconstruction and rendering. In recent years, there have been numerous studies investigating applications of deep learning in PAI [158, 159, 160, 161]. More specifically, GANs have been applied to PAI both for artifact removal and as a substitution for iterative solutions [162, 163] where the input and output of the GANs resemble the initial guess and final iterative solution provided by iterative reconstruction algorithms respectively. However, these studies have mainly focused on post-beamforming processing of photoacoustic data for enhancing reconstruction. In our case, we have noisy input images, and idealized (or “gold standard”) output images, generated by averaging many acquisitions frames and applying SVD denoising. To reproduce this mapping without the correspondingly long acquisition time, a generative supervised learning approach such as a GAN is ideally suited.

In this study, we have developed a GAN-based method using the Pix2Pix model for denoising pre-beamformed RF data that can mimic both the behaviour of multi-frame temporal averaging noise reduction using only one frame of data while also being able to remove sensor-specific artifacts. Since the Pix2Pix model was designed for image-to-image translation tasks involving highly structured graphical outputs, we hypothesized that it might work as an alternative or supplement to our SVD denoising, which excels at removing structured noise.

We achieved comparable results both in terms of raw RF data and also in the resultant reconstructed images using the denoised data. To the best of our knowledge, this study is the first to use GANs as a pre-processing step in photoacoustic imaging to enhance the quality of the raw RF data and consequently the resultant reconstructed images.

4.2 Methods

In the following sections, scalars are represented by lower-case letters, data matrices by bold lower-case letters, and operators by upper-case italic letters. We will begin by describing the photoacoustic model and defining our reference dataset. We then propose our method of GAN denoising, as well

as our reconstruction methods and metrics which will be used for quantitative assessment. Finally we describe our data acquisition hardware and experimental parameters.

4.2.1 Photoacoustic Model

The photoacoustic effect describes the generation of propagating acoustic waves by a temporally short pulse of laser light. These waves can be detected by standard ultrasound transducers, and the received signals can be processed to reconstruct the initial pressure distribution p_0 which, under the assumption of an acoustically homogeneous sample and spatially uniform illumination, will be proportional to the optical absorption in the sample [60]. When the measured time-varying pressure $p_D(\mathbf{r}_D, t)$ is acquired at multiple positions \mathbf{r}_D in 3D space and $p_0(\mathbf{r})$ is reconstructed at sample points \mathbf{r} , this is referred to as PAI reconstruction [100]. In imaging systems, the discretized versions of $p_0(\mathbf{r})$ and $p_D(\mathbf{r}_D, t)$ are denoted by \mathbf{p}_0 and \mathbf{p}_D , respectively.

4.2.2 Reference RF Dataset

At each imaging location across a volume, a sensor is used to acquire a frame of RF data with the general shape of (m, n) where m is the number of detection elements of the sensor and n is the number of time samples taken, which correspond to how far away is the source from the detection element. This noisy frame of RF data is denoted by $\mathbf{p}_{D, \text{noisy}}$. We preprocess these frames by first performing temporal averaging at each location where n_f frames ($n_f = 20$ in our example) are averaged together denoted by $\mathbf{p}_{D, n_f\text{-avg}}$. We then perform SVD denoising to remove cross-channel noise bands from the data [106, 135]. These steps will result in a *reference frame* for the experiments performed in this paper, which we will denote by $\mathbf{p}_{D, \text{ref}}$. Additionally, SVD denoising applied to a single noisy frame of RF data (as opposed to a temporal average of frames) will be denoted by $\mathbf{p}_{D, \text{SVD}}$.

For example, a frame of our RF data contains $n = 1792$ time samples across $m = 384$ elements of the transducer, where for each 128-element, 1792-sample segment of a frame of RF data, we construct a 128×1792 matrix and compute its SVD (our transducer array of 384 elements is multiplexed into 128 channels; three acquisitions are performed at each transducer location). By isolating the first k singular value components and reverting to the original representation, we are left with only the noise, which dominates those k singular values. We subtract this noise from the original matrix, resulting in the denoised data (the quality of which we will assess quantita-

tively in Section 4.3.2). This process is illustrated in Figure 4.1. An example frame pair of $\mathbf{p}_{\mathbf{D},\text{noisy}}$ and $\mathbf{p}_{\mathbf{D},\text{ref}}$ can also be seen in Figure 4.2.

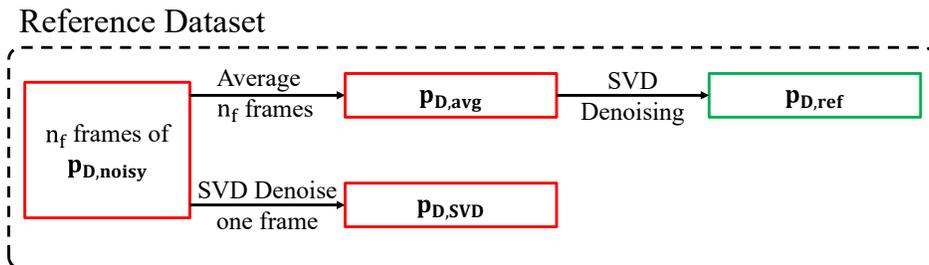


Figure 4.1: Preparation of a *reference frame*. After imaging one location n_f times and performing temporal averaging, the resultant frame is denoised using SVD denoising. We also denoise each frame with the SVD approach.

In SVD denoising, as k , the number of discarded singular values increases, lower-frequency structures are suppressed at the expense of an increased level of background noise, which can be observed in Figure 4.3. To strike the best balance between removing the noise bands and preserving the photoacoustic signals, the data in the present study was denoised using $k = 15$, obtained empirically by examining the inset A-line plots in Figure 4.3, and choosing the smallest k such that the noise band signals are no stronger than the background fluctuations.

4.2.3 GAN Denoising

We use the Pix2Pix Model [157] which trains a cGAN to learn a mapping from an input image and random noise vector \mathbf{z} , to an output image. Due to the size of $\mathbf{p}_{\mathbf{D},\text{noisy}}$ and $\mathbf{p}_{\mathbf{D},\text{ref}}$ in our example, (384, 1792), we cannot use these frames at once as an input/output pair and therefore we need to divide them into smaller patches denoted by ${}^i\mathbf{p}_{\mathbf{D},\text{noisy}}$ and ${}^i\mathbf{p}_{\mathbf{D},\text{ref}}$ with i indicating the location of the patch within each frame (see Figure 4.4 for an example). Letting $G(\cdot, \cdot)$ and $D(\cdot, \cdot)$ represent the outputs of the generator and discriminator respectively, the training objective of the model presented in [157] will be

$$L_{cGAN}(G, D) = \mathbb{E}_{i_{\mathbf{p}_{\mathbf{D},\text{noisy}}, i_{\mathbf{p}_{\mathbf{D},\text{ref}}}}} [\log D({}^i\mathbf{p}_{\mathbf{D},\text{noisy}}, {}^i\mathbf{p}_{\mathbf{D},\text{ref}})] + \mathbb{E}_{i_{\mathbf{p}_{\mathbf{D},\text{noisy}}, \mathbf{z}}} [\log (1 - D({}^i\mathbf{p}_{\mathbf{D},\text{noisy}}, G({}^i\mathbf{p}_{\mathbf{D},\text{noisy}}, \mathbf{z})))] \quad (4.1)$$

$$L_{L1}(G, D) = \mathbb{E}_{i_{\mathbf{p}_{\mathbf{D},\text{noisy}}, i_{\mathbf{p}_{\mathbf{D},\text{ref}}}, \mathbf{z}}} [\|{}^i\mathbf{p}_{\mathbf{D},\text{ref}} - G({}^i\mathbf{p}_{\mathbf{D},\text{noisy}}, \mathbf{z})\|_1] \quad (4.2)$$

4.2. Methods

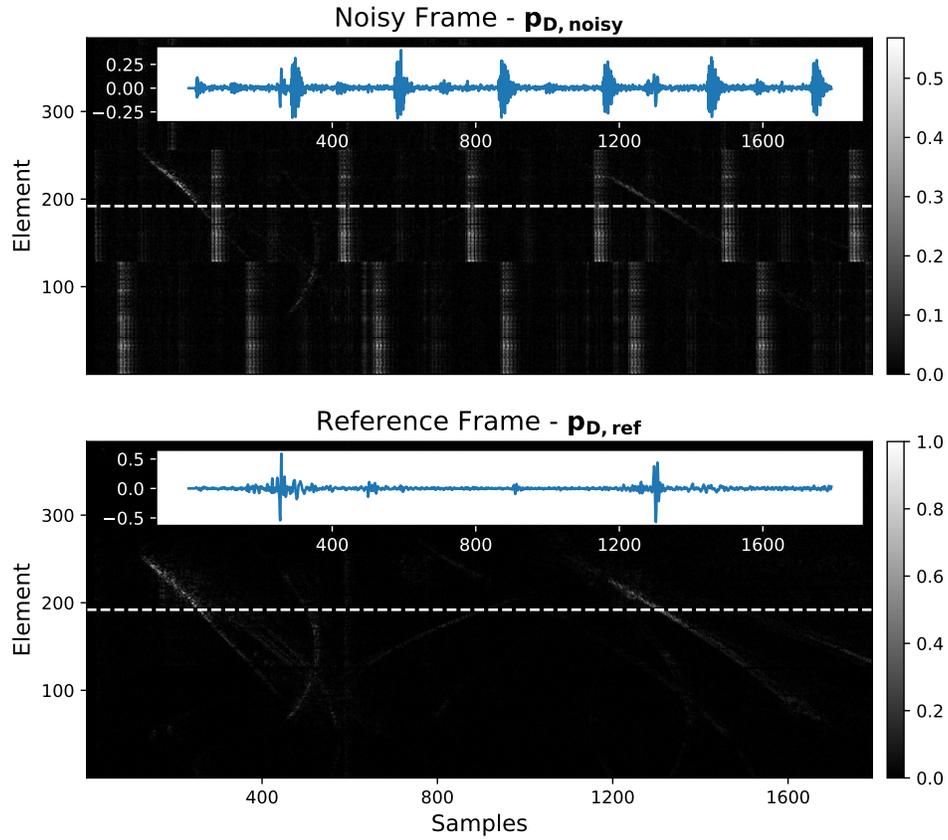


Figure 4.2: RF data training frame pair for the GAN. These 384×1792 frames are divided into 128×128 patches to be used with the GAN. The *reference frame* is the result of averaging 20 frames followed by SVD denoising. The line plot in each frame is an A-line profile of the frame at the location of the dashed white line.

4.2. Methods

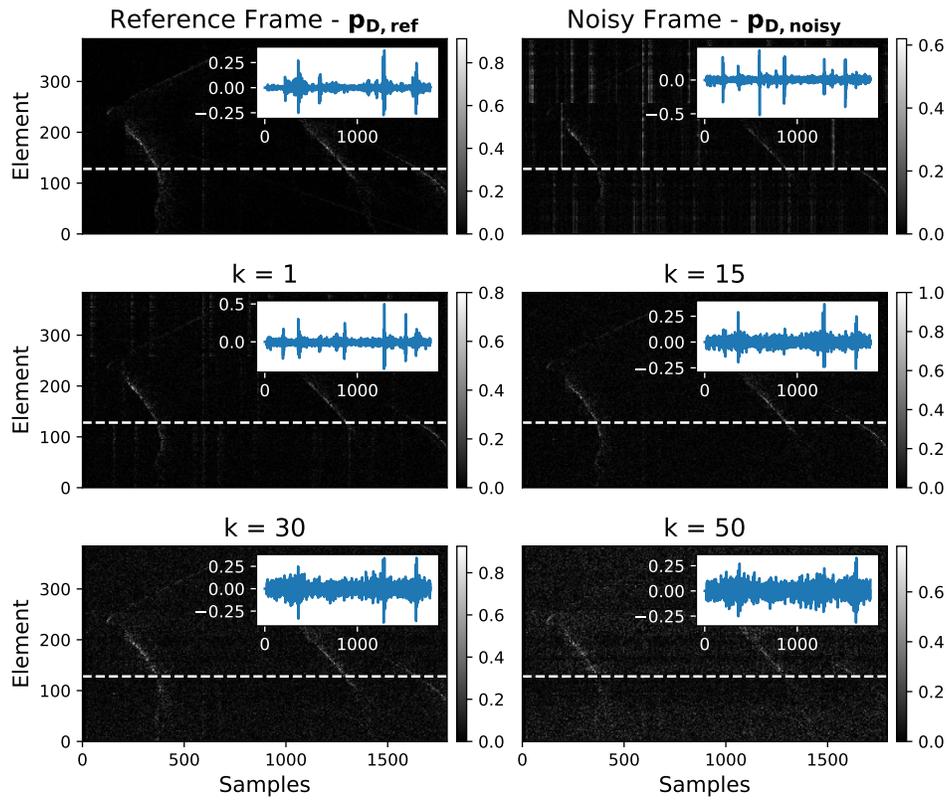


Figure 4.3: Example results of SVD denoising, illustrating the suppression of signals as the number of discarded singular values, k , increases.

$$G^* = \arg \min_G \max_D \{L_{cGAN}(G, D) + \lambda_1 L_{L1}(G, D)\} \quad (4.3)$$

where L_{cGAN} denotes the cGAN loss, L_{L1} the L1 distance norm, and \mathbb{E} the expectation values [157]. Combining Equations (4.1) and (4.2) results in G^* which is the cGAN loss wherein the generator tries to minimize the objective and the discriminator tries to maximize it, in addition to a sparsifying distance which encourages a more focused image controlled by the parameter λ_1 .

In training the Pix2Pix model, the loss function does not provide a descriptive source for determining the quality of training [157], therefore in addition to having a training dataset and a testing dataset, we employ a validation dataset to tune the parameters of the model after each training trial. While this requires performing an additional experiment, it ensures that the model training is completely blind to the validation data, unlike cross-validation methods such as k-fold or Monte Carlo. We found that we achieved optimal performance using our validation dataset with a patch size of (128, 128), without overlap, a batch size of $n_b = 5$, a learning rate of $\alpha = 0.0002$, and a regularizer $\lambda_1 = 100$. We used the Adam optimizer with the model weights being initialized randomly from a Gaussian distribution with mean 0 and standard deviation 0.02. With the exception of the patch size, which we changed to match the number of RF channels in our data (and therefore the size of the noise bands), all other parameters were the default values recommended in [157]. As seen from Equation (4.3), training occurs with the generator learning to output images that are closer to the reference, while the the discriminator learns to distinguish between real images and fake images from the generator. This training process is outlined in Figure 4.6. Our choice of hyperparameters and input/output images sizes will result in the GAN architecture illustrated in Figure 4.5 where the generator is a U-Net architecture [164] and the discriminator is a PatchGAN classifier [157]. Our model was trained using a single Tesla V100 GPU on an NVIDIA DGX-1 system.

Performing inference on our unseen testing data using the GAN is similar and requires the frames to be first divided into sub-patches, denoised, and subsequently put back together as outlined in Figure 4.7. A frame of RF data denoised using the GAN will be denoted by \mathbf{PD}_{GAN} . During inference, only the generator is used to output denoised patches using noisy patch inputs, while the discriminator is left unused.

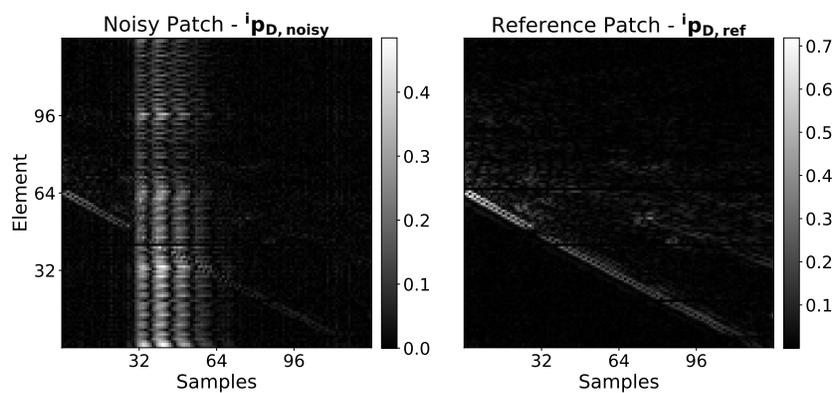
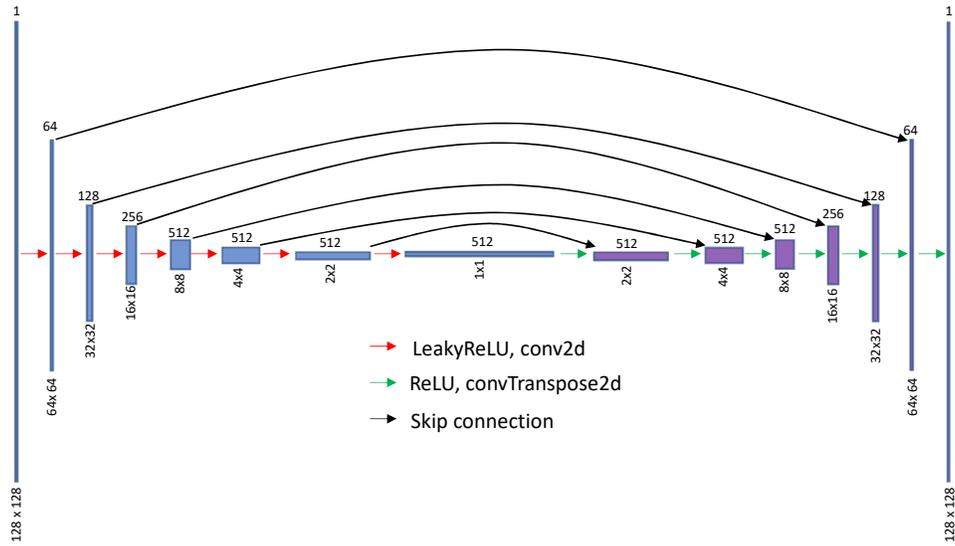
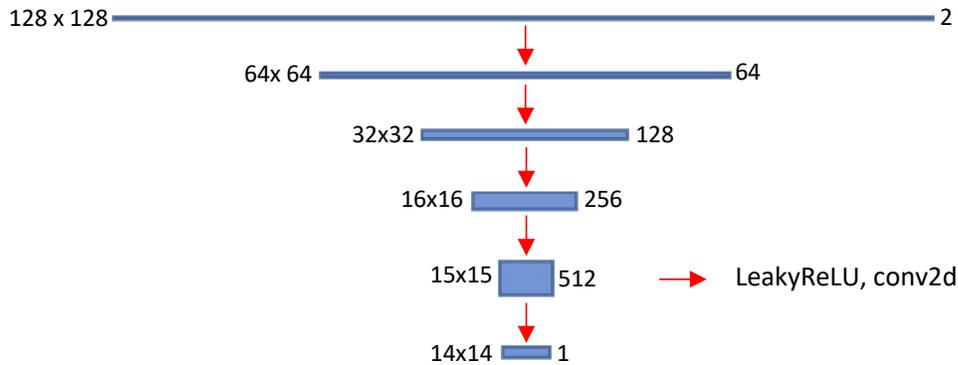


Figure 4.4: RF data training patch pair for the GAN. In our example, these 128×128 patches are extracted from 384×1792 RF frames. The patches have been absolute-valued for display only, to increase visibility of the signals. The *reference* patches are the result of averaging 20 frames followed by SVD denoising. Note the increased intensity of the main (diagonal) signal, as well as some weaker signals to its right which were not visible at all in the noisy data.

4.2. Methods



(a) Generator: UNet



(b) Discriminator: Convolutional PatchGAN Classifier

Figure 4.5: GAN architecture portraying the dimensions of the hidden layers fusing 128×128 patches. If a different patch size were used, the size of the layers would be adjusted accordingly. Each output of the generator is subsequently used as an input to the discriminator which has to determine if it is a real patch or a fake patch.

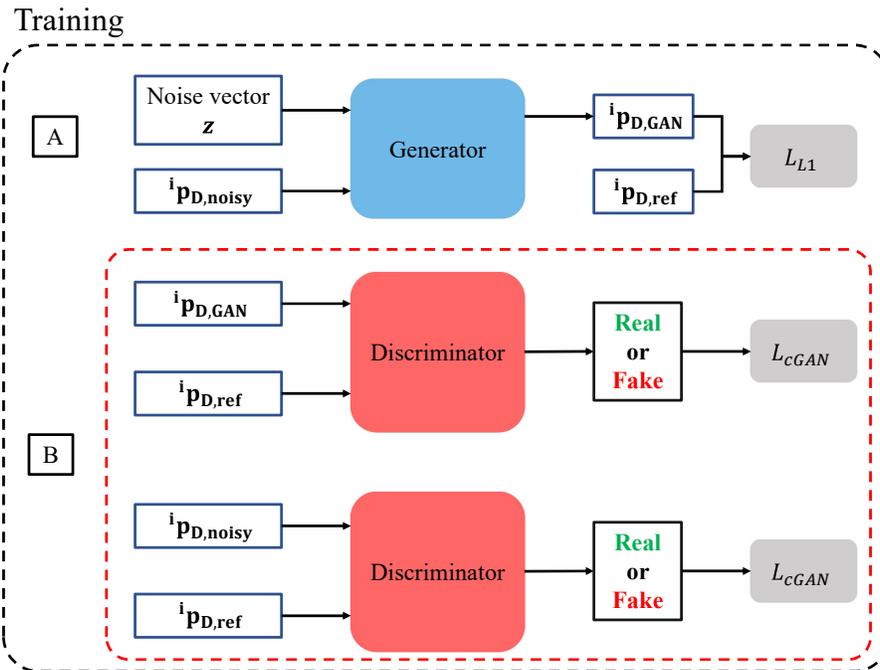


Figure 4.6: The training process of the GAN involves **A** the generator learning to output more realistic data and **B** the discriminator learning to distinguish between fake and real images. The combined objective function $L_{cGAN} + \lambda_1 L_{L1}$ is back-propagated through both networks, updating their respective weights.

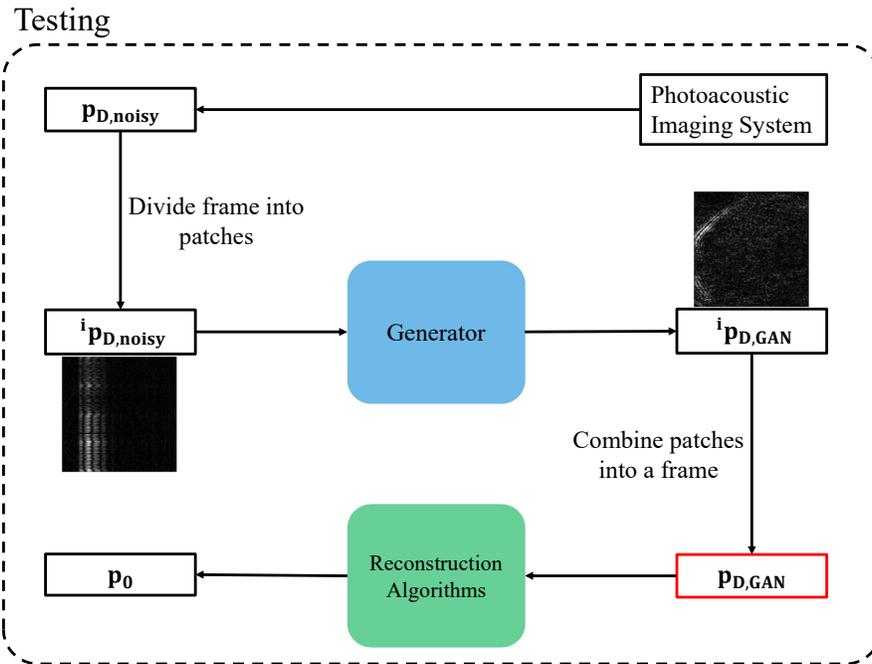


Figure 4.7: Once the noisy RF data $\mathbf{p}_{D,\text{noisy}}$ is acquired using the imaging system, each frame is divided into ${}^i\mathbf{p}_{D,\text{noisy}}$ sub-patches (128×128 in our examples). These patches are then denoised using the trained generator model making ${}^i\mathbf{p}_{D,\text{GAN}}$ which are combined into a denoised $\mathbf{p}_{D,\text{GAN}}$ frame subsequently used by the reconstruction algorithms to approximate the initial pressure distribution \mathbf{p}_0 .

4.2.4 Reconstructions

Improving the pre-processing and denoising of RF data ultimately serves to improve the quality of reconstructed photoacoustic images. We therefore need to assess the reconstructions that use our GAN-denoised RF dataset compared to reconstructions using temporal averaging and SVD denoising.

Filtered Back-Projection

In this study, we use the model of Xu et al. [100] for spherical scanning geometry, which relates $p_0(\mathbf{r})$ to $p_D(\mathbf{r}_D, t)$ in a backprojection form as

$$p_0(\mathbf{r}) = \frac{2}{\Omega_0} \int_{\Omega_0} d\Omega_0 \left[p_D(\mathbf{r}_D, t) - t \frac{\partial p_D(\mathbf{r}_D, t)}{\partial t} \right]_{t=\frac{|\mathbf{r}-\mathbf{r}_D|}{v_s}} \quad (4.4)$$

where v_s is the acoustic speed in the sample, Ω_0 is the solid angle of the surface containing the detection points \mathbf{r}_D , and $d\Omega_0$ is the solid angle of the surface element at a location \mathbf{r}_D relative to a sample point \mathbf{r} . While this form is exact only in the case in which this surface completely encloses the sample ($\Omega_0 = 4\pi$), the $\frac{d\Omega_0}{\Omega_0}$ term serves as a weight to mitigate the effects of the well-known *partial view problem* [100]. We have previously described our reconstruction scheme in detail [135] but we note that in the present study we assume spatially uniform illumination. We refer to this single-shot reconstruction as filtered backprojection (FBP) in the later sections.

Fast Iterative Shrinkage Thresholding Algorithm

To implement the iterative reconstruction algorithm, we also need a forward model describing the generated p_D as a function of p_0 . If we define a matrix A^* which describes the action of Equation (4.4) on a discrete sampling of p_D in vector form \mathbf{p}_D (the *RF data*) as

$$A^* \mathbf{p}_D = \mathbf{p}_0 \quad (4.5)$$

where \mathbf{p}_0 is a vector of initial pressure values in the sample (the *image*), then we can also define a matched [101, 165] forward (adjoint) operator A such that

$$A \mathbf{p}_0 = \mathbf{p}_D \quad (4.6)$$

where $A^* = A^T$ since A contains real-valued elements [101]. We will refer to A and A^T as the projection and backprojection operators, respectively.

4.2. Methods

Our CUDA implementation of the action of A^T upon \mathbf{p}_D is included in Appendix C.

Iterative frameworks for solving the acoustic inverse problem in PAI have been well studied in the past [102]. In this study, we have chosen the fast iterative shrinkage-thresholding algorithm (FISTA) [166] with a total variation (TV) regularizer [167] which is used widely in the field of PAI [101,103]. As we are dealing with the limited-view problem [100], we can never exactly compute \mathbf{p}_0 , therefore formulating PAI as an inverse problem and following the steps outlined in [101,103,166,167] with our best possible estimation to \mathbf{p}_0 being denoted by $\hat{\mathbf{p}}_0$, the minimization objective will be

$$\min_{\hat{\mathbf{p}}_0} \{F(\hat{\mathbf{p}}_0) = \|A\hat{\mathbf{p}}_0 - \mathbf{p}_D\|^2 + \lambda_2 \|\hat{\mathbf{p}}_0\|_{TV}\} \quad (4.7)$$

Algorithm 1: Monotone FISTA

```

j = 0;
t_0 = 1;
n = number of trials;
 $\hat{\mathbf{p}}_{0,0}$  = FBP Reconstruction;
while j  $\in$  [1, n] do
     $\mathbf{q}_j = \text{TV}_{2\lambda_2/L_j} \{ \mathbf{g}_j - \frac{2}{L_j} A^T (A\mathbf{g}_j - \mathbf{p}_D) \}$  (Alg 1.1)
     $t_{j+1} = \frac{1 + \sqrt{1 + 4t_j^2}}{2}$  (Alg 1.2)
     $\hat{\mathbf{p}}_{0,j} = \{ \mathbf{q}_j \text{ if } F(\mathbf{q}_j) \leq F(\hat{\mathbf{p}}_{0,j-1}) \text{ else } \hat{\mathbf{p}}_{0,j-1} \}$  (Alg 1.3)
     $\mathbf{g}_{j+1} = \hat{\mathbf{p}}_{0,j} + \frac{t_j}{t_{j+1}} (\mathbf{q}_j - \hat{\mathbf{p}}_{0,j-1}) + \frac{t_j - 1}{t_{j+1}} (\hat{\mathbf{p}}_{0,j} - \hat{\mathbf{p}}_{0,j-1})$  (Alg 1.4)

```

Algorithm 1 outlines the monotone version of FISTA (MFISTA) outlined in [167], which prevents divergence of the cost function F due to the lack of an exact solution for the TV denoising problem [167]. We iteratively update $\hat{\mathbf{p}}_{0,j}$, the j -th guess, from our initial guess of $\hat{\mathbf{p}}_{0,0}$ being the FBP reconstruction to converge to $\hat{\mathbf{p}}_0$. $\text{TV}_{2\lambda_2/L_j}$ is a solution to the TV denoising problem, implemented using the method of Chambolle [168], and L_j is the Lipschitz constant, used as the denoising weight found using backtracking line search [167]. Note that \mathbf{p}_D in Algorithm 1 may be any of the RF datasets

explained, i.e. Noisy, Reference, GAN Denoised, or SVD Denoised. In summary, as the algorithm progresses, $\hat{\mathbf{p}}_{0,\mathbf{j}}$ converges to $\hat{\mathbf{p}}_0$, a TV-regularized solution to Equation (4.5).

4.2.5 Metrics

In this section we present the metrics we use to both assess the result of the GAN denoising directly using the RF data, as well as indirectly by inspecting the resultant reconstructions. In the formulas presented in this section, \mathbf{y} refers to our ground-truth and expected output and $\hat{\mathbf{y}}$ refers to our estimated output. Due to the dual purpose of our assessment, \mathbf{y} and $\hat{\mathbf{y}}$ can be either the RF data or the reconstructed images in the formulas below. We begin by using the mean square error (MSE), defined as

$$MSE = \frac{1}{n_p} \sum_{i=1}^{n_p} (y_i - \hat{y}_i)^2 \quad (4.8)$$

where n_p is total number of pixels in an image and y_i and \hat{y}_i the value of the pixel i in the expected and estimated images respectively.

The second metric we use to quantify the performance is the SSIM [148], which is given by

$$SSIM(\mathbf{y}, \hat{\mathbf{y}}) = \frac{(2\mu_{\mathbf{y}}\mu_{\hat{\mathbf{y}}} + c_1)(2\sigma_{\mathbf{y}\hat{\mathbf{y}}} + c_2)}{(\mu_{\mathbf{y}}^2 + \mu_{\hat{\mathbf{y}}}^2 + c_1)(\sigma_{\mathbf{y}}^2 + \sigma_{\hat{\mathbf{y}}}^2 + c_2)} \quad (4.9)$$

where μ is the mean intensity and σ the standard deviation of the signal, and $\sigma_{\mathbf{y}\hat{\mathbf{y}}}$ is the correlation coefficient of \mathbf{y} and $\hat{\mathbf{y}}$

$$\sigma_{\mathbf{y}\hat{\mathbf{y}}} = \frac{1}{n_p - 1} \sum_{i=1}^{n_p} (y_i - \mu_{\mathbf{y}})(\hat{y}_i - \mu_{\hat{\mathbf{y}}}) \quad (4.10)$$

The constant terms c_1 and c_2 are used to avoid ill-defined values in Equation (4.9); we use the same values presented in [148]. The final metric used for assessing the quality of the denoised RF data directly and indirectly is the feature similarity index measurement (FSIM) [169,170] which is designed to replicate subjective (i.e. human) evaluations, by assessing the similarity of low-level features in the images. FSIM is defined as

$$FSIM = \frac{\sum_{\mathbf{x} \in \Omega} S_L(\mathbf{x}) \cdot PC_m(\mathbf{x})}{\sum_{\mathbf{x} \in \Omega} PC_m(\mathbf{x})} \quad (4.11)$$

where the vector \mathbf{x} represents the spatial domain of the the images \mathbf{y} and $\hat{\mathbf{y}}$, and Ω is the entire spatial domain. $PC_m(\mathbf{x})$ is the maximum phase congruency [171]—a dimensionless quantity providing an absolute measure of significance of feature points, between the two images. $S_L(\mathbf{x})$ is a similarity measure between the two images calculated using

$$S_L(\mathbf{x}) = S_{PC}(\mathbf{x}) \cdot S_G(\mathbf{x}) \quad (4.12)$$

where

$$S_{PC}(\mathbf{x}) = \frac{2PC_1(\mathbf{x}) \cdot PC_2(\mathbf{x}) + s_1}{PC_1(\mathbf{x})^2 + PC_2(\mathbf{x})^2 + s_1} \quad (4.13)$$

$$S_G(\mathbf{x}) = \frac{2G_1(\mathbf{x}) \cdot G_2(\mathbf{x}) + s_2}{G_1^2(\mathbf{x}) + G_2^2(\mathbf{x}) + s_2} \quad (4.14)$$

with $G_i(\mathbf{x})$ being the gradient magnitude and $PC_i(\mathbf{x})$ the phase congruency of image i in domain \mathbf{x} . s_1 and s_2 are constants to avoid ill-defined values of $S_{PC}(\mathbf{x})$ and $S_G(\mathbf{x})$, for which we chose the same values used in [169,170].

4.2.6 Data Acquisition

All RF data were acquired using the SonixEmbrace ABUS (Ultrasonix Medical Corporation, Richmond, BC, Canada). This scanner consists of a 384-element transducer with -12 cm radius of curvature, 10 MHz centre frequency, and 90% bandwidth. The transducer is embedded in a spherical dome attached to a motor, which rotates through 360° to collect volumetric data. A SonixDAQ module (DAQ - BK Medical, Peabody, MA) was used to acquire pre-beamformed RF data at 40 MHz. Our illumination source consists of a Continuum Surelite II laser (Continuum, Santa Clara, CA) pumping an OPO from the same manufacturer. The OPO output is homogenized [144] and coupled into a 1 mm silica-core optical fiber, which is coupled to a custom illuminator designed to deliver a fan-shaped beam of diffuse illumination to the sample surface through a window parallel to the ABUS transducer [135]. Synchronization of the RF data acquisition with the laser illumination was accomplished using a custom Arduino-based circuit controlled over USB. With the data download from the SonixDAQ to the PC being the rate-limiting step, we are able to acquire one frame per transducer position at 50 equally spaced positions in about 20 minutes.

Our training data were acquired from a custom, modular wire phantom [135], consisting of black spray-painted monofilament fishing line suspended between attachment points on a 3D printed template as seen in Figure 4.8a. To validate our GAN, we imaged a commercial photoacoustic

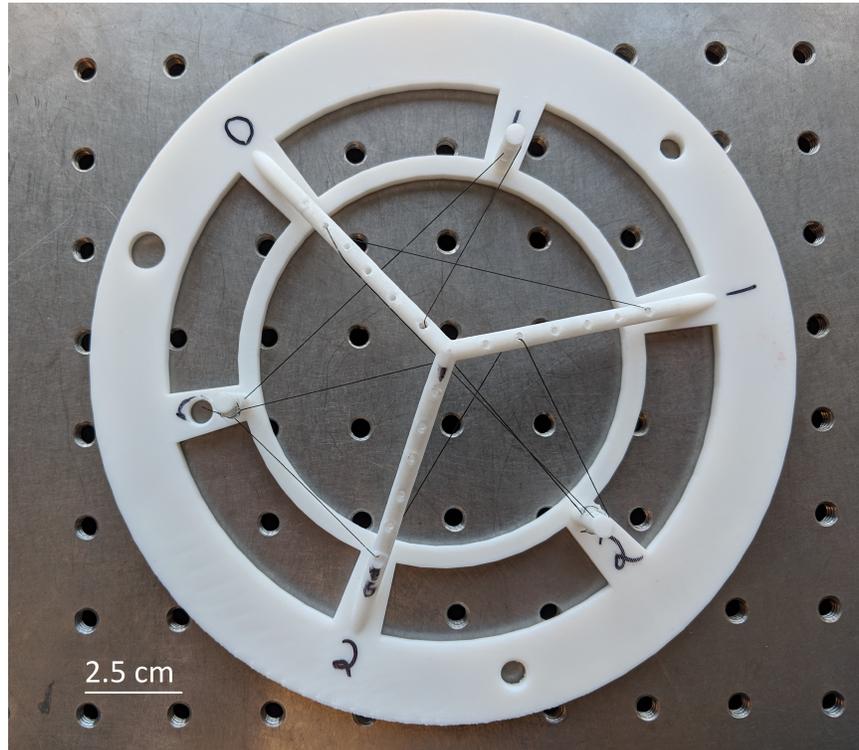
phantom containing optically absorbing spheres purchased from Computerized Imaging Reference Systems Incorporated (CIRS - Norfolk, VA). For testing and quantification, we used the same black fishing line described above, tied into a small figure-of-eight knot, as shown in Figure 4.8b, which provides several features at various angles within a small imaging volume. It is important to note that only the black surface of the wires will function as a photoacoustic absorber, rather than the entire cross-section which would be the case for a solid coloured inclusion, or a blood vessel.

To test how the model performs on out-of-distribution data, we also developed an anatomically realistic 3D printed vessel phantom using labelled magnetic resonance angiography patient data made publicly available by Lou *et al.* [172]. Beginning with dataset “*Neg_35_Left*”, we extracted voxels labelled as blood vessels. We added a rectangular base for the vessels to attach to, and discarded any vessel segments which were not attached to either the base or another vessel, as these would be unsupported in the final print. Finally, we converted the voxel data to a surface mesh which was exported to STL format for 3D printing. The phantom was printed on a Form 2 3D printer (FormLabs, Somerville, MA) with a layer height of 0.1 mm, using “Tough 2000” resin from the same manufacturer. Finally, the print was cleaned up and spray painted black. While some of the smallest vessels were not faithfully reproduced, and others still were lost during the cleanup process, the phantom still provides realistic geometry, and vessels ranging in diameter from 0.2 mm (the voxel size in the original dataset) up to 5 mm. The numerical data, as well as a photo of the final phantom are shown in Figure 4.9.

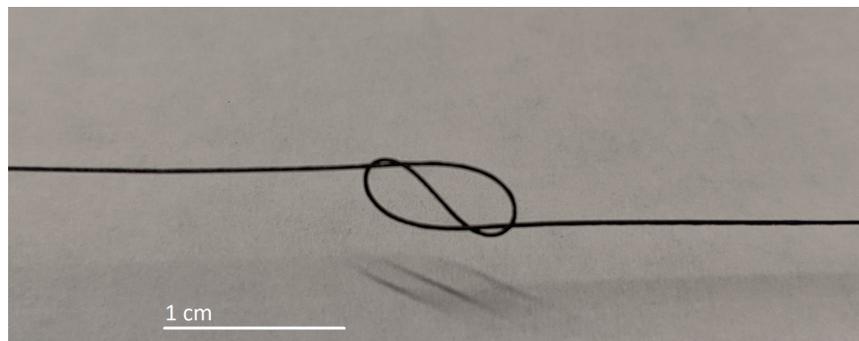
The training phantom was imaged at 100 equally-spaced transducer positions in the ABUS dome, the validation phantom at 10, and the two different testing phantoms at 50 positions each. Each location was imaged 20 times in all three cases. Since we train our GAN using (128, 128) patches, this corresponds to 4300, 430, and 2150 patch pairs for the training, validation and hyperparameter optimization, and testing datasets, respectively.

4.3 Results

In this section, we first provide an example of a training trial for the model. We then provide our results by first assessing the pre-beamformed RF data followed by the resulting reconstructions.



(a) Wire phantom used for training the GAN.



(b) Knot phantom used for testing the GAN.

Figure 4.8: Custom phantoms used for training and testing the GAN.

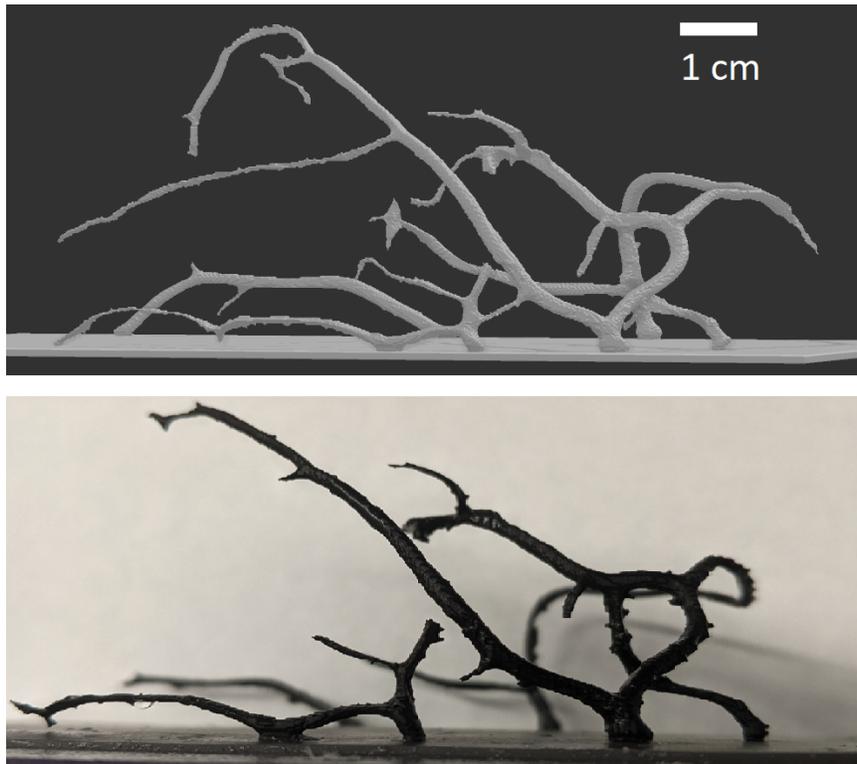


Figure 4.9: Vessel phantom used for testing the GAN. The top panel shows the numerical data, and the bottom panel is a photo of the 3D print.

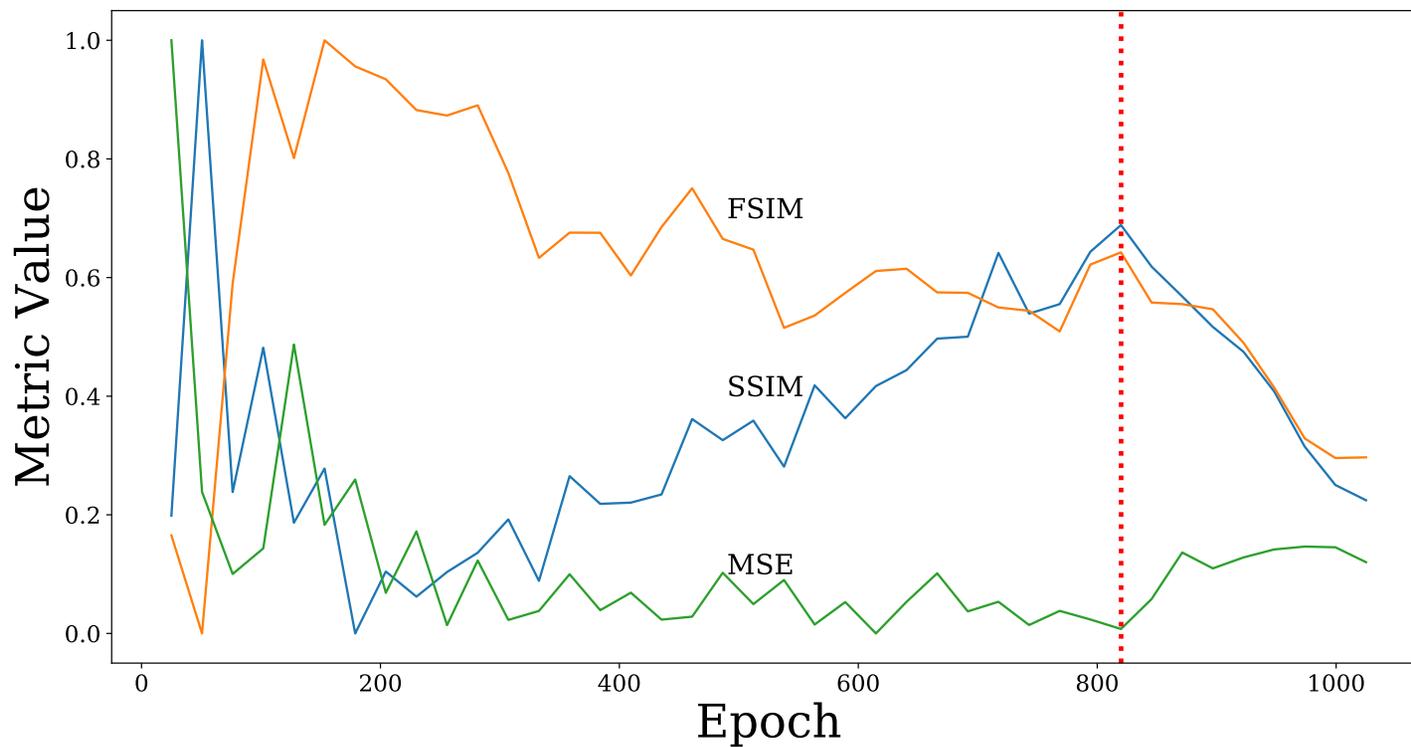


Figure 4.10: Metrics for the validation dataset. The dashed red line indicates the epoch chosen to be applied to our test dataset. Although this epoch is not at the maximum or minimum value of any of the metric, it achieves an appropriate balance between the three. Note that it is well known that GAN optimum performance is not necessarily monotone with the number of epochs [173].

4.3.1 Model Hyperparameter Tuning

As mentioned previously, we use a validation dataset to tune the hyperparameters of the GAN. Here we used data acquired of the CIRS phantom described in Section 4.2.6. Figure 4.10 is an example of this process, demonstrating the choice of optimal number of training epochs. Ideally we would want the minimum MSE value and the maximum FSIM and SSIM to occur at the same epoch, however, Figure 4.10 suggests that considering all three metrics, epoch 820 (dashed red line) provides a good balance of a low MSE value alongside high SSIM and FSIM values which are desirable for our task. This is due to the fact that while low MSE corresponds to the removal of the background noise, SSIM and FSIM correspond to the fidelity of the output signals' shape compared to the true signals. These 820 training epochs took an average of 28 seconds each, for a total training time of about 6.5 hours.

4.3.2 RF Data

In applying our metrics, we compared our reference RF dataset to the other possible approaches of only frame-averaging, only SVD denoising, and finally GAN denoising. The results are the mean values from 50 imaging locations along the ABUS dome. We have also added a 10 frame averaged case for comparison purposes to the 20 frames averaged case. These results are summarized in Table 4.1 for the knot phantom, and in Table 4.2 for the vessel phantom. Figure 4.11 and Figure 4.12 illustrates sample denoised frames of the knot and vessel phantom, respectively, using the different approaches for one of the aforementioned 50 imaging locations. We note that increasing the number of averaged frames beyond 20, we found diminishing returns with respect to our quality metrics which did not justify the proportional increase in scan time. A plot of this relationship is included in Appendix A.2.

4.3.3 Reconstructed Images

Table 4.3 summarizes the reconstruction results of the knot phantom using the different RF data available with reconstructions resultant from the reference RF dataset as ground-truth. The 2D reconstructions are in the axial/lateral plane of the ABUS transducer at a resolution of 0.1 mm. We achieved convergent FISTA reconstructions with the regularizer parameter $\lambda_2 = 0.01$ after 20 iterations of Algorithm 1. By examining the minimization objective as a function of iteration number, we can see that there is little improvement beyond 20 iterations. An example of this relationship is included in Appendix A.3. As before, these results are the mean of 50

4.3. Results

Denoising Category	MSE (%)	SSIM	FSIM
Noisy Data	0.28	0.50	0.54
10 Frames Averaged	0.07	0.72	0.71
20 Frames Averaged	0.06	0.78	0.74
SVD Denoised	0.17	0.37	0.75
GAN Denoised	0.05	0.76	0.85

Table 4.1: Comparison of denoising methods for the knot phantom RF data. Each row represents the mean of 50 imaging locations around the ABUS dome. Best results are bolded.

Denoising Category	MSE (%)	SSIM	FSIM
Noisy Data	0.31	0.43	0.50
10 Frames Averaged	0.16	0.47	0.69
20 Frames Averaged	0.17	0.46	0.71
SVD Denoised	0.88	0.15	0.71
GAN Denoised	0.17	0.44	0.84

Table 4.2: Comparison of denoising methods for the vessel phantom RF data. Each row represents the mean of 50 imaging locations around the ABUS dome. Best results are bolded.

4.3. Results

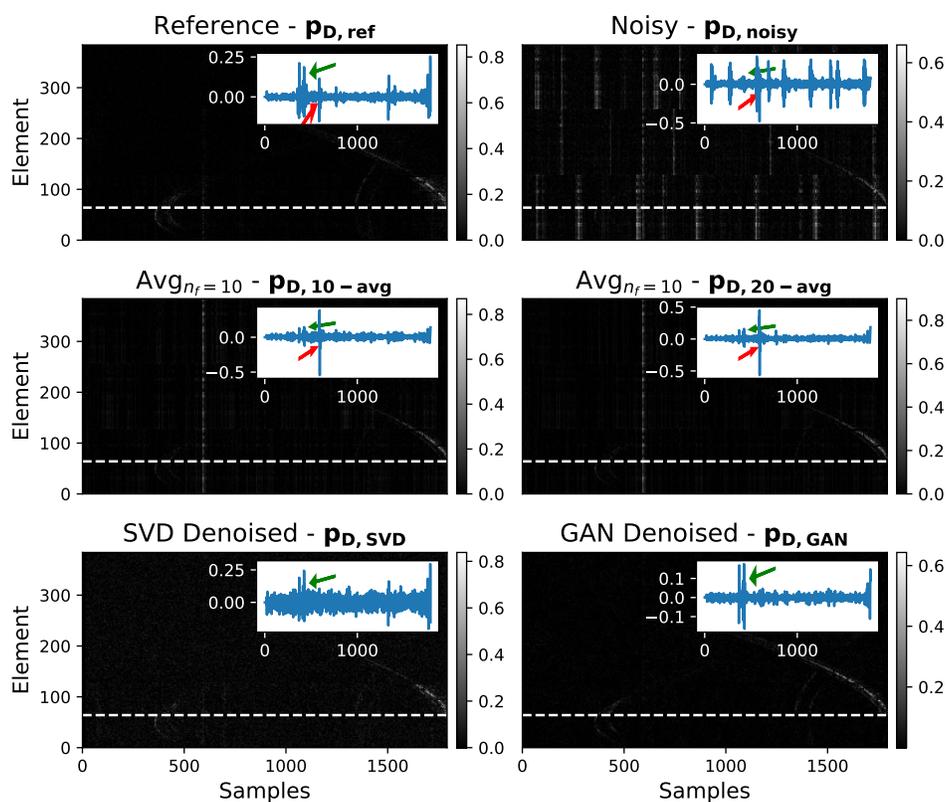


Figure 4.11: Testing data showing the denoising techniques considered applied to the knot phantom. The blue plot in each frame is an A-line profile of the frame at the location of the dashed white line. The absolute values of the RF data are displayed. The A-lines portray the real range of the data. The green arrow points to the signal at sample 400 for comparison. The red arrow points to the straight line artifact in denoising cases where it has not been removed. Note that the *SVD denoised* and *GAN denoised* are single-frame results without temporal averaging.

4.3. Results

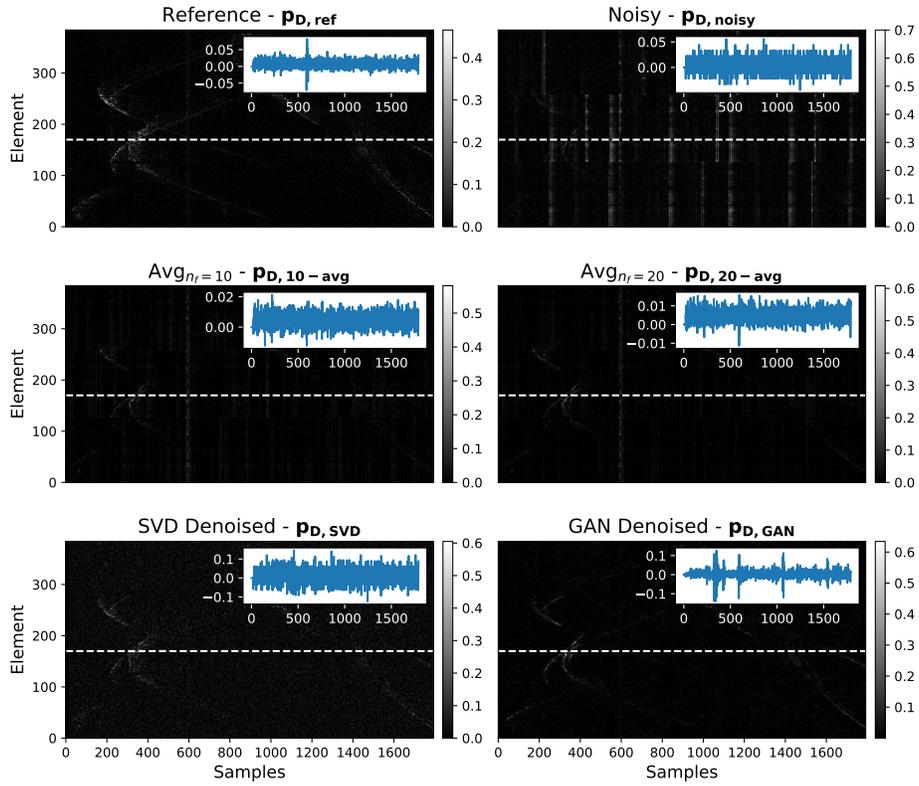


Figure 4.12: Testing data showing the denoising techniques considered applied to the vessel phantom. The blue plot in each frame is an A-line profile of the frame at the location of the dashed white line. The absolute values of the RF data are displayed. The A-lines portray the real range of the data. Note that the *SVD denoised* and *GAN denoised* are single-frame results without temporal averaging.

4.4. Discussion

imaging locations with Figures 4.13 and 4.14 providing examples for FBP and iterative reconstructions at one of these imaging locations.

Denoising Category	MSE(%)		SSIM		FSIM	
	FBP	FISTA	FBP	FISTA	FBP	FISTA
Noisy Data	2.00	0.51	0.40	0.48	0.69	0.69
10 Frames Averaged	0.85	0.11	0.43	0.64	0.75	0.74
20 Frames Averaged	0.51	0.08	0.49	0.73	0.79	0.76
SVD Denoised	0.15	0.05	0.58	0.71	0.84	0.78
GAN Denoised	0.18	0.05	0.66	0.80	0.83	0.80

Table 4.3: Comparison of denoising methods and their effects on the reconstructed images of the knot phantom. Each row represents the mean of 50 imaging locations around the ABUS dome. Best results are bolded.

4.4 Discussion

For the knot phantom, as shown in Table 4.1 the GAN outperforms all of the other denoising cases except in its SSIM value which is slightly below that of 20 frames averaged. This is mainly due to the GAN being trained on the SVD and 20 frames averaged combination which is the theoretical upper limit of its performance. While 20 frames was sufficient for our system (see Appendix A.2), this will vary depending on the attainable SNR of the system, and the specifics of the “structural” noise (if any) being targeted. As such, the necessary training dataset will vary for each system. Fortunately, since the training time is only several hours, it would not be prohibitive to train on multiple datasets to find what works best.

In Table 4.3, we can see that the reconstructions of the knot phantom data denoised using the GAN outperform all of the other cases, both in FBP and FISTA. This is significant since all of these performance benefits come at a fraction of a time of the frame averaged cases due to the GAN needing only a single frame of data per imaging location. These results are consistent with characteristics of the metrics used. While MSE mainly focuses on the absolute error between images, SSIM and FSIM focus on the structural differences and similarities [174].

Additionally, if we take a closer look at Figure 4.11, we can see that the GAN and SVD are the only two cases that have managed to remove the

4.4. Discussion

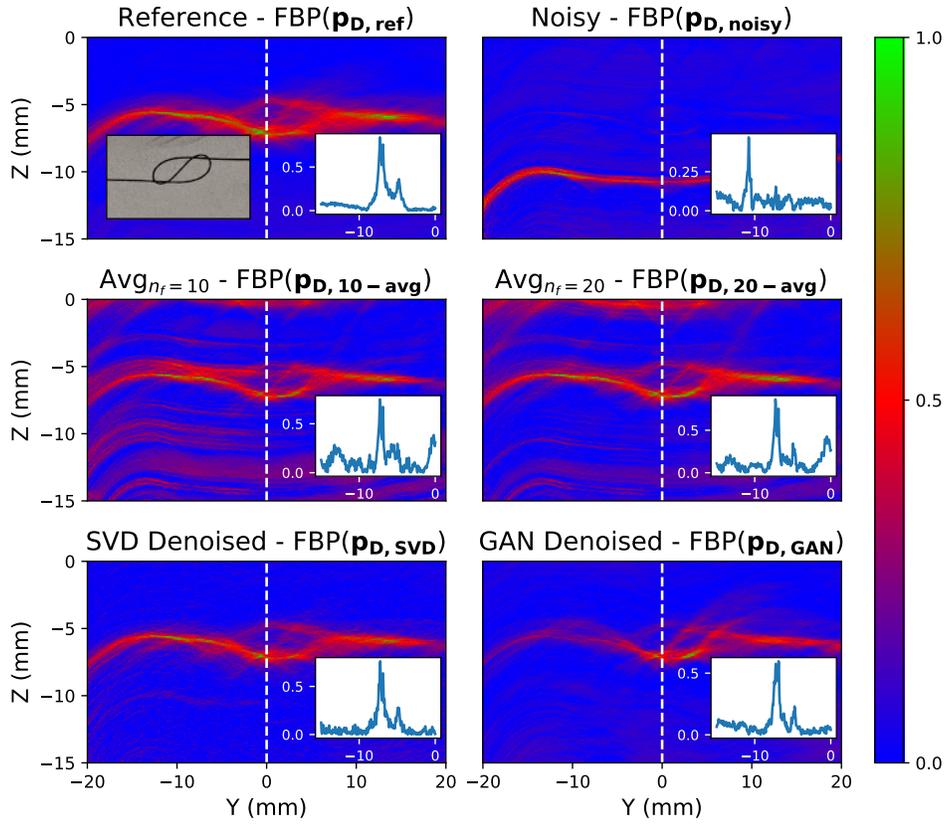


Figure 4.13: 2D FBP reconstruction of the knot phantom data after the RF data has been processed using different denoising approaches. The line plot in each image is a profile plot along the dashed white line. *Reference* refers to reconstructions using the Reference RF dataset.

4.4. Discussion

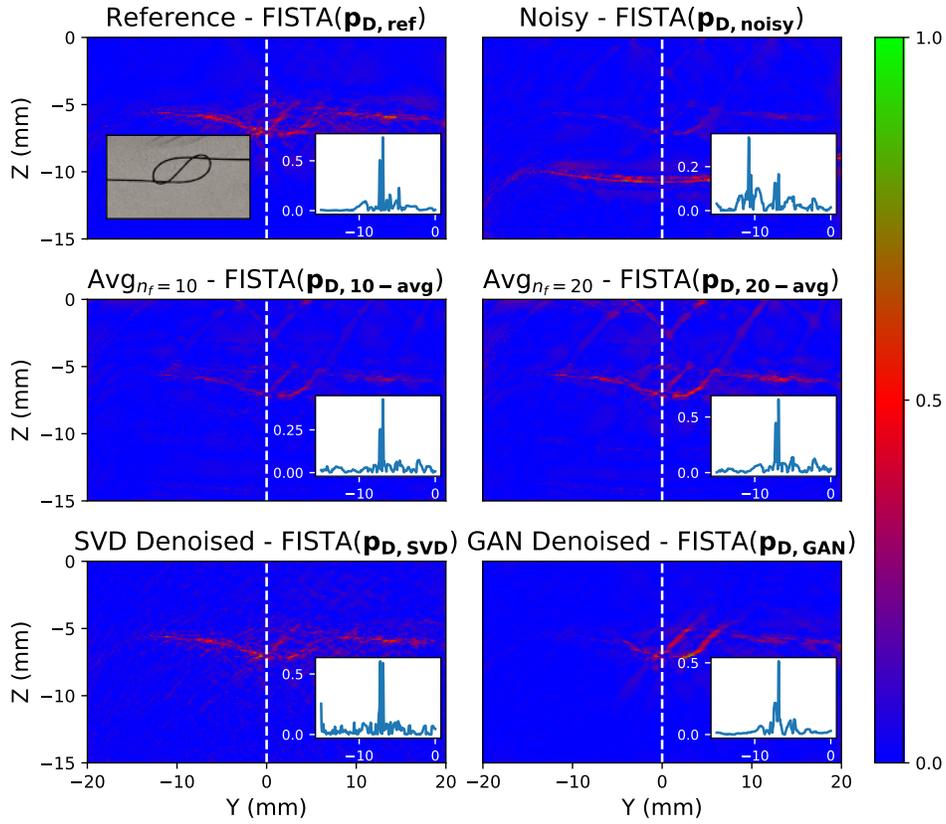


Figure 4.14: 2D FISTA reconstruction of the knot phantom data after the RF data has been processed using different denoising approaches. The line plot in each image is a profile plot along the dashed white line. *Reference* refers to reconstructions using the reference RF dataset.

straight line noise band visible at sample 550 (indicated in the A-lines in Figure 4.11 with a red arrow). However, the GAN displays less background noise which is consistent with the results of Table 4.1, one such example being the signal at sample 400 (indicated in the A-lines in Figure 4.11 with a green arrow) which is strongest in the GAN output. Once again this is due to the fact that the GAN learns the optimal behaviour from both frame averaging and SVD denoising since it has been trained on their combination. Similar behaviour is also seen in the reconstruction results in Figure 4.13 and Figure 4.14. The effects of the leftover noise bands in the RF data can clearly be seen in the FBP($\mathbf{pD},10\text{-avg}$), FBP($\mathbf{pD},20\text{-avg}$), FISTA($\mathbf{pD},10\text{-avg}$), and FISTA($\mathbf{pD},20\text{-avg}$) cases as streaks in the lower left of the images, which are running parallel to the transducer face. Further, we can see examples of significant background noise in the FBP(\mathbf{pD},\mathbf{SVD}) and FISTA(\mathbf{pD},\mathbf{SVD}) cases.

Despite being significantly different from the training data, we see that the GAN performed similarly well on the vessel phantom data. Qualitatively, we see again that the GAN denoising removed the worst noise bands from the RF data, as is clear in the A-line plots in Figure 4.12. Table 4.2 shows that the GAN performed best with respect to FSIM, and very near to the best result with respect to MSE and SSIM.

The timing benefits of using only a single frame of data for imaging will result in a better PAI frame rate as the GAN model can be loaded prior to the imaging session and denoising a frame takes on average 0.3 seconds on an Nvidia GeForce GTX 1060 GPU using the PyTorch library. This computation time will be reduced once inference-only libraries like TensorRT are used in the clinical deployment stage [175]. This is similar to our GPU implementation of SVD denoising, and since our 384-element transducer requires three 128-element acquisitions per frame (with one laser pulse each), we are still limited by the 10 Hz repetition rate of our laser. Additionally, averaging several frames per imaging location will most likely cause distortions in the data as clinical imaging subjects, i.e. human organs, will move throughout the data acquisition, therefore imaging only a single frame per location provides additional benefits.

Future work will benefit from a set of training data that is independent of both the SVD and frame averaging. One possible approach would be to use simulation data and add noise for training purposes; however, adding sensor specific noise to simulation data correctly will bring forth a new set of challenges. Ultimately, this will enable the GAN to overcome the shortcomings that might arise from using either the SVD or averaging. Further, while the effectiveness of this denoising method is fundamentally limited by the

SNR of the training data, it would be useful to explore the absolute limits of this technique in terms of the minimum SNR which still allows recovery of photoacoustic signals. We are exploring the application of this technique to photoacoustic data acquired using a laser diode illuminator, where significant frame averaging is usually required to attain an SNR similar to pulsed laser systems.

While this technique is promising based on the metrics we have chosen, it will be important moving forward to further study the specific effects of this method on the RF data and corresponding reconstructions to ensure, for example, that the attainable resolution is not affected, or that certain reconstruction methods are not rendered less effective. It would, for example, be useful to characterize the effect of this technique on the frequency content of the photoacoustic signals to ensure that certain feature sizes are not preferentially suppressed or enhanced.

Finally, for this method to be useful, it must be tested on *in vivo* data, which raises the question of how one would gather sufficient training data for such a study. While it would be prohibitively time consuming to acquire such a dataset from a single patient, we are hopeful that smaller scans of a modest cohort of patients or volunteers (10 to 15) would provide a sufficient dataset to train our model. Whether or not such a model would be robust enough for the high variability of *in vivo* data will be the ultimate test of this technique. Our tests on the vessel phantom suggest that some out-of-distribution data can be effectively denoised, providing some promise in this regard.

4.5 Conclusion

In this paper we have shown that using a GAN to denoise pre-beamformed RF data is a viable option both for removing general background noise for which frame-averaging is standard [104]. The approach can be used to remove artifacts where algorithms such the SVD denoising are traditionally employed. Using SVD and frame-averaging as our training standard, we have improved upon the performance of single frame SVD output, producing results similar to SVD and frame averaging combined. Additionally we have shown that improving the quality of the RF data results in improvements in the corresponding reconstructions.

Chapter 5

Conclusions and Future Work

This chapter first provides a summary of the results of this thesis, and discusses their relevance in the context of the current state of the art. We then highlight shortcomings and limitations of these results, and provide suggestions to improve and further develop our multimodal system in future work.

5.1 Summary and Discussion

Breast cancer accounts for the deaths of almost 700 000 women globally each year. Improved early detection, especially in high risk groups, has the potential to save many of these lives.

In this thesis, we have built a system for automated multimodal imaging of the breast, with the hope of improving the detection and staging of breast cancer. Our ultimate goal is to conduct a single, fast, operator-independent scan which can provide clinicians with a number of volumetrically-resolved anatomical markers, informing an accurate and efficient diagnosis. We have focused on imaging modalities which are supported by prior clinical results in the field, with both our choice of modalities, and many aspects of our system design being completely novel. We have used the SonixEmbrace ABUS as a starting point, and incorporated the new modalities without compromising the high quality US images this system produces.

In Chapter 2, we described the hardware and software of each imaging subsystem, including the S-WAVE system which provides tissue elasticity measurements, and the PAT system, which provides tissue vascularization and oxygenation state. We rationalized our design decisions, and described the inter-connectivity of the various system components. We performed S-WAVE imaging of a commercially available elasticity phantom, accurately measuring the absolute stiffness values consistent with the product specifications. Finally, we reported the first combined imaging with this system,

acquiring B-mode US, S-WAVE, and PAT of a custom-designed multimodal phantom containing features which are visible uniquely in each of the modalities.

In Chapter 3 we described our complete PAT system in greater detail. In particular, we reported our novel illumination system and reconstruction software, and quantified their capabilities. Using this system, we successfully resolved PA inclusions embedded in chicken breast tissue, which is commonly used to mimic the scattering and absorption properties of human breast tissue. We also designed, built, and imaged a custom wire phantom designed to test imaging depth and reconstruction accuracy. Our results confirmed not only that this illumination method is feasible, but that it can provide comparable imaging depth to similar systems in the literature.

Building upon the results of Chapter 3, Chapter 4 further improved our photoacoustic data processing method. We developed and implemented a machine learning framework for denoising photoacoustic data, and tested its performance both on the raw RF data, and on images reconstructed with a state-of-the-art iterative reconstruction scheme. This data analysis pipeline was shown to greatly improve the quality of our images, which were acquired with the minimum possible scan time, given the current hardware restrictions. With the system complete, and this analysis pipeline in place, we are now able to generate PAT on par with leading groups in the field.

In conclusion, this thesis sought to develop and quantify the performance of an automated scanning platform for the detection of breast cancer. This type of imaging has the potential to detect breast cancers earlier and with greater accuracy, especially in high risk populations, such as women with dense breast tissue. Screening with such a system could lower cost compared to the current gold standard of XRM, reduce unnecessary biopsies, and most importantly, prevent deaths. The majority of the thesis focused on the promising, but technically complex modality of PAT, but also provides preliminary results simultaneously acquiring B-mode US, S-WAVE, Doppler flow, and PAT of a custom multimodal phantom.

5.2 Implications

While the photoacoustic portion of our system still falls well short of clinical application, we have solved many of the engineering difficulties associated with developing this system, and have hopefully provided an excellent starting point for fruitful studies in the future.

One of the primary concerns at the outset of this work was that the

geometric constraints of the SonixEmbrace system would make illuminating the breast for PAT technically infeasible, or even impossible. Our approach of using diffusive optics and modeling the illumination is general enough to be used in other photoacoustic imaging applications, especially those attempting to minimize cost [138], image disease sites which are difficult to illuminate by conventional means such as the prostate [176, 177], and guidance of minimally invasive surgical procedures where space constraints often lead to inherently non-uniform illumination [67].

The problem of denoising PA data is one that affects every system. Acquiring more data reduces random noise, but increases scan time and the total laser dose to which the patient is exposed. This balance is one of the fundamental choices made in PA system design, and as such, a technique that can improve SNR without acquiring additional frames is a very valuable contribution to the field.

S-WAVE imaging with the SonixEmbrace ABUS would be much easier to translate to *in vivo* imaging, given the simpler hardware requirements and more extensive clinical results supporting such a system. A modest volunteer study could serve to further explore the capabilities of the system, and inform further development of the system. The recent work of Shao *et al* [121] using multimodal breast data to effectively distinguish benign and malignant breast lesions, offers a very promising path forward for this work. Combining automated multimodal scanning with a sophisticated classification pipeline such as this could provide an extremely powerful screening tool.

5.3 Limitations and Future Work

There are several outstanding issues and limitations of our system which could be barriers to clinical translation.

1. Relatively slow PAT scan time of 20 min per breast. This is primarily due to the slow data transfer rate of the SonixDAQ module. To mitigate this, we could switch to a Verasonics Vantage system, or upgrade our laser system to improve on our current pulse repetition rate of 10 Hz. We are currently exploring whether our denoising method could make a high repetition rate (kHz) diode-based system more feasible, despite each pulse having significantly lower power than a laser-based system. We also currently stop the motor for each acquisition, but a faster data transfer rate could enable continuous scanning, further decreasing the overall scan time.

2. The SonixEmbrace transducer is not ideal for PAT. For the imaging of breast cancers undergoing angiogenesis, a center frequency of 2.25 MHz with sensitivity extending into the kHz range is ideal [127]. Further, the inclusion of an elevational acoustic lens on the SonixEmbrace transducer limits the acceptance angle of the detection elements, leading to imaging artifacts and decreased sensitivity [140]. We intend to explore supplementing the SonixEmbrace transducer with a more broadband, purpose-built photoacoustic array. Capacitive micromachined ultrasound transducers (CMUTs) have proven very promising in terms of frequency response and angular sensitivity [178, 179] but suffer from expensive and complex fabrication, and are known to be delicate and therefore prone to degradation with use. Recent developments in fabrication techniques may alleviate or eliminate some of these issues, facilitating rapid prototyping of thin, flexible, transparent CMUT arrays [180, 181].
3. With Ultrasonix Medical Corporation being acquired and subsequently shut down, research-level support for the SonixEmbrace, SonixDAQ, and other hardware and software central to our system will likely soon be non-existent. While the transducer and US system can be replaced with similar offerings from other vendors, access to the pre-beamformed RF data required for PAT imaging is not commonly made available outside of research systems such as the Verasonics Vantage. To use the Vantage system for human studies requires first obtaining regulatory clearance for the combined device, introducing further barriers.
4. More in-depth analysis will be required to determine the effectiveness of our GANs-based denoising on real breast data. While we were able to show good performance on out-of-distribution data, the high variability of *in vivo* data will require the most robust training data possible, necessitating at least a small volunteer study to acquire.
5. Given that we presented only preliminary multimodal data, many questions remain about exactly how these data could be combined, displayed, and analysed. Providing information in a format which can be easily and consistently interpreted by clinicians is an ever-present issue in medical imaging. There are existing workflows and protocols for two-dimensional (2D) US, elasticity, and Doppler imaging, but displaying and interpreting volumetric data is more difficult. Much of the current clinical research in PA imaging is focused on this problem,

5.3. *Limitations and Future Work*

and work is underway to standardize PA data toward this goal [182].

Bibliography

- [1] Hyuna Sung, Jacques Ferlay, Rebecca L. Siegel, Mathieu Laversanne, Isabelle Soerjomataram, Ahmedin Jemal, and Freddie Bray. Global Cancer Statistics 2020: GLOBOCAN Estimates of Incidence and Mortality Worldwide for 36 Cancers in 185 Countries. *CA: A Cancer Journal for Clinicians*, 71(3):209–249, 2021.
- [2] Amy Maxmen. The hard facts. *Nature*, 485(7400):S50–S51, 2012.
- [3] F. Kilburn-Toppin and S. J. Barter. New Horizons in Breast Imaging. *Clinical Oncology*, 25(2):93–100, 2013.
- [4] Regina J Hooley, Kathryn L Greenberg, Rebecca M Stackhouse, Jaime L Geisel, Reni S Butler, and Liane E Philpotts. Screening US in patients with mammographically dense breasts: initial experience with Connecticut Public Act 09-41. *Radiology*, 265(1):59–69, oct 2012.
- [5] Paul C Stomper, D J D’Souza, P A DiNitto, and Mark a Arredondo. Analysis of parenchymal density on mammograms in 1353 women 25-79 years old. *American Journal of Roentgenology*, 167(5):1261–1265, nov 1996.
- [6] Norman F. Boyd, Helen Guo, Lisa J. Martin, Limei Sun, Jennifer Stone, Eve Fishell, Roberta A. Jong, Greg Hislop, Anna Chiarelli, Salomon Minkin, and Martin J. Yaffe. Mammographic Density and the Risk and Detection of Breast Cancer. *New England Journal of Medicine*, 356(3):227–236, jan 2007.
- [7] Wendie a Berg. Tailored supplemental screening for breast cancer: what now and what next? *AJR. American journal of roentgenology*, 192(2):390–9, feb 2009.
- [8] Thomas M Kolb, Jacob Lichy, and Jeffrey H Newhouse. Comparison of the performance of screening mammography, physical examination, and breast US and evaluation of factors that influence them: an analysis of 27,825 patient evaluations. *Radiology*, 225(1):165–75, oct 2002.

- [9] Wendie A. Berg, Zheng Zhang, Daniel Lehrer, Roberta A. Jong, Etta D. Pisano, Richard G. Barr, Marcela Böhm-Vélez, Mary C. Mahoney, W. Phil Evans, Linda H. Larsen, Marilyn J. Morton, Ellen B. Mendelson, Dione M. Farria, Jean B. Cormack, Helga S. Marques, Amanda Adams, Nolin M. Yeh, and Glenna Gabrielli. Detection of breast cancer with addition of annual screening ultrasound or a single screening MRI to mammography in women with elevated breast cancer risk. *JAMA - Journal of the American Medical Association*, 307(13):1394–1404, 2012.
- [10] Regina J Hooley, Leslie M Scutt, and Liane E Philpotts. Breast ultrasonography: state of the art. *Radiology*, 268(3):642–59, sep 2013.
- [11] Sebastian Wojcinski, Andre Farrokh, Ursula Hille, Jakub Wiskirchen, Samuel Gyapong, Amr a Soliman, Friedrich Degenhardt, and Peter Hillemanns. The Automated Breast Volume Scanner (ABVS): initial experiences in lesion detection compared with conventional handheld B-mode ultrasound: a pilot study of 50 cases. *International journal of women's health*, 3:337–46, jan 2010.
- [12] Srirang Manohar and Maura Dantuma. Current and Future Trends in Photoacoustic Breast Imaging. *Photoacoustics*, 16(April):100134, 2019.
- [13] Hyunji Kim, Joo Hee Cha, Ha-Yeun Oh, Hak Hee Kim, Hee Jung Shin, and Eun Young Chae. Comparison of conventional and automated breast volume ultrasound in the description and characterization of solid breast masses based on BI-RADS features. *Breast cancer (Tokyo, Japan)*, oct 2012.
- [14] B. Zhuang, T. Chen, C. Leung, K. Chan, J. Dixon, K. Dickie, and L. Pelissier. Microcalcification enhancement in ultrasound images from a concave Automatic Breast Ultrasound Scanner. *2012 IEEE International Ultrasonics Symposium*, pages 1662–1665, oct 2012.
- [15] G J Rizzatto. Towards a more sophisticated use of breast ultrasound. *European radiology*, 11(12):2425–35, jan 2001.
- [16] B S Garra, E I Cespedes, J Ophir, S R Spratt, R A Zuurbier, C M Magnant, and M F Pennanen. Elastography of breast lesions: initial clinical results. *Radiology*, 202(1):79–86, 1997. PMID: 8988195.

- [17] Wei Meng, Guangchen Zhang, Changjun Wu, Guozhu Wu, Yan Song, and Zhaoling Lu. Preliminary results of acoustic radiation force impulse (ARFI) ultrasound imaging of breast lesions. *Ultrasound in medicine and biology*, 37(9):1436–43, sep 2011.
- [18] Gelareh Sadigh, Ruth C. Carlos, Colleen H. Neal, and Ben A. Dwamena. Accuracy of quantitative ultrasound elastography for differentiation of malignant and benign breast abnormalities: A meta-analysis. *Breast Cancer Research and Treatment*, 134(3):923–931, 2012.
- [19] Hui Yang, Yongyuan Xu, Yanan Zhao, Jing Yin, Zhiyi Chen, and Pintong Huang. The role of tissue elasticity in the differential diagnosis of benign and malignant breast lesions using shear wave elastography. *BMC Cancer*, 20(1):1–10, 2020.
- [20] Marvin M. Doyley and Kevin J. Parker. Elastography: General principles and clinical applications. *Ultrasound Clinics*, 9(1):1–11, 2014. Oncologic Ultrasound.
- [21] Yasufumi Asao, Yohei Hashizume, Takahiro Suita, Ken ichi Nagae, Kazuhiko Fukutani, Yoshiaki Sudo, Toshikazu Matsushita, Shuichi Kobayashi, Mariko Tokiwa, Iku Yamaga, Elham Fakhrehajani, Masae Torii, Masahiro Kawashima, Masahiro Takada, Shotaro Kanao, Masako Kataoka, Tsuyoshi Shiina, and Masakazu Toi. Photoacoustic mammography capable of simultaneously acquiring photoacoustic and ultrasound images. *Journal of Biomedical Optics*, 21:116009, 11 2016.
- [22] Erin I. Neuschler, Reni Butler, Catherine A. Young, Lora D. Barke, Margaret L. Bertrand, Marcela Böhm-Vélez, Stamatia Destounis, Pamela Donlan, Stephen R. Grobmyer, Janine Katzen, Kenneth A. Kist, Philip T. Lavin, Erini V. Makariou, Tchaiko M. Parris, Kathy J. Schilling, F. Lee Tucker, and Basak E. Dogan. A pivotal study of optoacoustic imaging to diagnose benign and malignant breast masses: A new evaluation tool for radiologists. *Radiology*, 287:398–412, 5 2018.
- [23] A.A. A. Oraevsky, B. Clingman, J. Zalev, A.T. T. Stavros, W.T. T. Yang, and J.R. R. Parikh. Clinical optoacoustic imaging combined with ultrasound for coregistered functional and anatomical mapping of breast tumors. *Photoacoustics*, 12(July):30–45, dec 2018.
- [24] Alexander A. Oraevsky, Richard Su, Ha Nguyen, James Moore, Yang Lou, Sayantan Bhadra, Mark Anastasio, Luca Forte, and Wei Yang.

- Full-view 3D imaging system for functional and anatomical screening of the breast. In Alexander A. Oraevsky and Lihong V. Wang, editors, *Photons Plus Ultrasound: Imaging and Sensing 2018*, page 248. SPIE, apr 2018.
- [25] Afaf F. Moustafa, Theodore W. Cary, Laith R. Sultan, Susan M. Schultz, Emily F. Conant, Santosh S. Venkatesh, and Chandra M. Sehgal. Color doppler ultrasound improves machine learning diagnosis of breast cancer. *Diagnostics*, 10, 9 2020.
- [26] Li Lin, Peng Hu, Junhui Shi, Catherine M. Appleton, Konstantin Maslov, Lei Li, Ruiying Zhang, and Lihong V. Wang. Single-breath-hold photoacoustic computed tomography of the breast. *Nature Communications*, 9(1), 2018.
- [27] Tomoyuki Fujioka, Mio Mori, Kazunori Kubota, Yuka Kikuchi, Leona Katsuta, Mai Kasahara, Goshi Oda, Toshiyuki Ishiba, Tsuyoshi Nakagawa, and Ukihide Tateishi. Simultaneous comparison between strain and shear wave elastography of breast masses for the differentiation of benign and malignant lesions by qualitative and quantitative assessments. *Breast Cancer*, 26:792–798, 11 2019.
- [28] Xiang Zhang, Ming Liang, Zehong Yang, Chushan Zheng, Jiayi Wu, Bing Ou, Haojiang Li, Xiaoyan Wu, Baoming Luo, and Jun Shen. Deep learning-based radiomics of b-mode ultrasonography and shear-wave elastography: Improved performance in breast mass classification. *Frontiers in Oncology*, 10, 8 2020.
- [29] Daniela Elia, Daniele Fresilli, Patrizia Pacini, Sara Cardaccio, Giorgia Polti, Olga Guiban, Iliaria Celletti, Eriselda Kutrolli, Carlo De Felice, Rossella Occhiato, Corrado De Vito, Maria Ida Amabile, Alessandro De Luca, Vito D’Andrea, Massimo Vergine, Federica Pediconi, Ferdinando D’Ambrosio, and Vito Cantisani. Can strain us-elastography with strain ratio (sre) improve the diagnostic accuracy in the assessment of breast lesions? preliminary results. *Journal of Ultrasound*, 24:157–163, 6 2021.
- [30] Comparison of ultrasound elastography and color doppler ultrasonography for distinguishing small triple-negative breast cancer from fibroadenoma. *Journal of Ultrasound in Medicine*, 37:2135–2146, 9 2018.

- [31] Chadaporn Keatmanee, Utairat Chaumrattanakul, Kazunori Kotani, and Stanislav S. Makhanov. Initialization of active contours for segmentation of breast cancer via fusion of ultrasound, doppler, and elasticity images. *Ultrasonics*, 94:438–453, 4 2019.
- [32] Heechul Yoon, Yiyang I. Zhu, Steven K. Yarmoska, and Stanislav Y. Emelianov. Design and demonstration of a configurable imaging platform for combined laser, ultrasound, and elasticity imaging. *IEEE Transactions on Medical Imaging*, 38:1622–1632, 7 2019.
- [33] Rachel F. Brem, László Tabár, Stephen W. Duffy, Marc F. Inciardi, Jessica A. Guingrich, Beverly E. Hashimoto, Marla R. Lander, Robert L. Lapidus, Mary Kay Peterson, Jocelyn A. Rapelyea, Susan Roux, Kathy J. Schilling, Biren A. Shah, Jessica Torrente, Ralph T. Wynn, and Dave P. Miller. Assessing improvement in detection of breast cancer with three-dimensional automated breast US in women with dense breast tissue: The somoinsight study. *Radiology*, 274(3):663–673, 2015.
- [34] Xiaohui Zhang, Juan Chen, Yidong Zhou, Feng Mao, Yan Lin, Songjie Shen, Qiang Sun, and Zhaolian Ouyang. Diagnostic value of an automated breast volume scanner compared with a hand-held ultrasound: A meta-analysis. *Gland Surgery*, 8(6):698–711, 2019.
- [35] Mengmeng Jia, Xi Lin, Xiang Zhou, Huijiao Yan, Yaqing Chen, Peifang Liu, Lingyun Bao, Anhua Li, Partha Basu, Youlin Qiao, and Rengaswamy Sankaranarayanan. Diagnostic performance of automated breast ultrasound and handheld ultrasound in women with dense breasts. *Breast Cancer Research and Treatment*, 181(3):589–597, 2020.
- [36] Yujing Xin, Xinyuan Zhang, Yi Yang, Yi Chen, Yanan Wang, Xiang Zhou, and Youlin Qiao. A multicenter, hospital-based and non-inferiority study for diagnostic efficacy of automated whole breast ultrasound for breast cancer in China. *Scientific Reports*, 11(1):1–9, 2021.
- [37] Kevin M Kelly, Judy Dean, W Scott Comulada, and Sung Jae Lee. Breast cancer detection using automated whole breast ultrasound and mammography in radiographically dense breasts. *European Radiology*, 20(3):734–742, 2010.

- [38] E Wenkel, M Heckmann, M Heinrich, S a Schwab, M Uder, R Schulz-Wendtland, W a Bautz, and R Janka. Automated breast ultrasound: lesion detection and BI-RADS classification—a pilot study. *RöFo : Fortschritte auf dem Gebiete der Röntgenstrahlen und der Nuklearmedizin*, 180(9):804–8, sep 2008.
- [39] Jung Min Chang, Woo Kyung Moon, Nariya Cho, Jeong Seon Park, and Seung Ja Kim. Radiologists’ performance in the detection of benign and malignant masses with 3D automated breast ultrasound (ABUS). *European Journal of Radiology*, 78(1):99–103, 2011.
- [40] Maryellen Giger, Dave Miller, Jeremy Brown, Marc Inciardi, Charles Metz, Yulei Jiang, Rachel Brem, Robert Nishikawa, Alexandra Edwards, and John Papaioannou. Clinical reader study examining the performance of mammography and automatic breast ultrasound in breast cancer screening. In *Radiological Society of North America 2012 Scientific Assembly and Annual Meeting*, 11 2012.
- [41] Elizabeth Kagan Arleo, Marwa Saleh, Dana Ionescu, Michele Drotman, Robert J Min, and Keith Hentel. Recall rate of screening ultrasound with automated breast volumetric scanning (ABVS) in women with dense breasts: a first quarter experience. *Clinical imaging*, pages 8–13, apr 2014.
- [42] Young Wook Kim, Seon Kwang Kim, Hyun Jo Youn, Eun Jung Choi, and Sung Hoo Jung. The clinical utility of automated breast volume scanner: a pilot study of 139 cases. *Journal of breast cancer*, 16(3):329–34, sep 2013.
- [43] Hong-Yan Wang, Yu-Xin Jiang, Qing-Li Zhu, Jing Zhang, Qing Dai, He Liu, Xing-Jian Lai, and Qiang Sun. Differentiation of benign and malignant breast lesions: a comparison between automatically generated breast volume scans and handheld ultrasound examinations. *European journal of radiology*, 81(11):3190–200, nov 2012.
- [44] Elizabeth A.M. O’Flynn, Jeremie Fromageau, Araminta E. Ledger, Alessandro Messa, Ashley D’Aquino, Minouk J. Schoemaker, Maria Schmidt, Neb Duric, Anthony J. Swerdlow, and Jeffrey C. Bamber. Ultrasound tomography evaluation of breast density: A comparison with noncontrast magnetic resonance imaging. *Investigative Radiology*, 52(6):343–348, 2017.

- [45] Neb Duric, Mark Sak, Shaoqi Fan, Ruth M. Pfeiffer, Peter J. Littrup, Michael S. Simon, David H. Gorski, Haythem Ali, Kristen S. Purrington, Rachel F. Brem, Mark E. Sherman, and Gretchen L. Gierach. Using Whole Breast Ultrasound Tomography to Improve Breast Cancer Risk Assessment: A Novel Risk Factor Based on the Quantitative Tissue Property of Sound Speed. *Journal of Clinical Medicine*, 9(2):367, 2020.
- [46] Neb Duric, Peter Littrup, Cuiping Li, Olivier Roy, Steve Schmidt, John Seamans, Andrea Wallen, and Lisa Bey-Knight. Whole breast tissue characterization with ultrasound tomography. *Medical Imaging 2015: Ultrasonic Imaging and Tomography*, 9419(March 2015):94190G, 2015.
- [47] André Farrokh, Harika Erdönmez, Fritz Schäfer, and Nicolai Maass. SOFIA: A Novel Automated Breast Ultrasound System Used on Patients in the Prone Position: A Pilot Study on Lesion Detection in Comparison to Handheld Grayscale Ultrasound. *Geburtshilfe und Frauenheilkunde*, 78(05):499–505, 2018.
- [48] Thomas Kuiran Chen, Corina Leung, Reza Zahiri Azar, Kwun-Keat Chan, Bo Zhuang, Kris Dickie, John Dixon, Linda Pendziwol, and Laurent Pelissier. Importance of transducer position tracking for automated breast ultrasound: Initial assessments. In *2012 IEEE International Ultrasonics Symposium*, pages 2623–2626. IEEE, oct 2012.
- [49] Reza Zahiri Azar, Corina Leung, Thomas Kuiran Chen, Kris Dickie, John Dixon, Kwun-Keat Chan, and Laurent Pelissier. An automated breast ultrasound system for elastography. In *2012 IEEE International Ultrasonics Symposium*, volume d, pages 1–4. IEEE, oct 2012.
- [50] Takeshi Nagashima, Hideyuki Hashimoto, Keiko Oshida, Shigeharu Nakano, Naoto Tanabe, Takashi Nikaido, Keiji Koda, and Masaru Miyazaki. Ultrasound demonstration of mammographically detected microcalcifications in patients with ductal carcinoma in situ of the breast. *Breast Cancer*, 12(3):216–220, 2005.
- [51] Malai Muttarak, P. Kongmebhol, and N. Sukhamwang. Breast calcifications: Which are malignant? *Singapore Medical Journal*, 50(9):907–914, 2009.
- [52] M. Lagendijk, E. L. Vos, K. P. Ramlakhan, C. Verhoef, A. H.J. Koning, W. van Lankeren, and L. B. Koppert. Breast and Tumour Volume

Bibliography

- Measurements in Breast Cancer Patients Using 3-D Automated Breast Volume Scanner Images. *World Journal of Surgery*, 42(7):2087–2093, 2018.
- [53] G. Gatta, G. Pace, D. Donatello, G. Falco, A. Reginelli, S. Cappabianca, R. Grassi, and G. Di Grezia. Role of a prone automated three-dimensional ultrasonography in breast cancer screening: preliminary results. In *European Congress of Radiology*, pages C–0430, 2019.
- [54] Michael Andre, James Wiskin, David Borup, Steven Johnson, Haydee Ojeda-Fournier, and Linda Olson. Quantitative volumetric breast imaging with 3D inverse scatter computed tomography. *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, pages 1110–1113, 2012.
- [55] Elaine Iuanow, Kathleen Smith, Nancy A. Obuchowski, Jennifer Bullen, and John C. Klock. Accuracy of Cyst Versus Solid Diagnosis in the Breast Using Quantitative Transmission (QT) Ultrasound. *Academic Radiology*, 24(9):1148–1153, 2017.
- [56] Peter J Littrup, Nebojsa Duric, Mark Sak, Cuiping Li, Olivier Roy, Rachel F Brem, and Mary Yamashita. The Fat-glandular Interface and Breast Tumor Locations: Appearances on Ultrasound Tomography Are Supported by Quantitative Peritumoral Analyses. *Journal of Breast Imaging*, 3(4):455–464, 2021.
- [57] Neb Duric, Peter Littrup, Mark Sak, Cuiping Li, Di Chen, Olivier Roy, Lisa Bey-Knight, and Rachel Brem. A Novel Marker, Based on Ultrasound Tomography, for Monitoring Early Response to Neoadjuvant Chemotherapy. *Journal of Breast Imaging*, 2(6):569–576, 2020.
- [58] Robert a Kruger, Cherie M Kuzmiak, Richard B Lam, Daniel R Reinecke, Stephen P Del Rio, and Doreen Steed. Dedicated 3D photoacoustic breast imaging. *Medical physics*, 40(11):113301, nov 2013.
- [59] Robert A. Kruger, Daniel R. Reinecke, and Gabe A. Kruger. Thermoacoustic computed tomography-technical considerations. *Medical Physics*, 26(9):1832–1837, sep 1999.
- [60] Minghua Xu and Lihong V. Wang. Photoacoustic imaging in biomedicine. *Review of Scientific Instruments*, 77(4):041101, 2006.

- [61] Li Lin, Peng Hu, Xin Tong, Shuai Na, Rui Cao, Xiaoyun Yuan, David C Garrett, Junhui Shi, Konstantin Maslov, and Lihong V Wang. High-speed three-dimensional photoacoustic computed tomography for preclinical research and clinical translation. *Nature Communications*, 12(1):882, dec 2021.
- [62] Nikhila Nyayapathi and Jun Xia. Photoacoustic imaging of breast cancer: a mini review of system design and image features. *Journal of Biomedical Optics*, 24(12):1, 2019.
- [63] Xueding Wang, William W. Roberts, Paul L. Carson, David P. Wood, and J. Brian Fowlkes. Photoacoustic tomography: a potential new tool for prostate cancer. *Biomedical Optics Express*, 1(4):1117, 2010.
- [64] Justin Rajesh Rajian, Gandikota Girish, and Xueding Wang. Photoacoustic tomography to identify inflammatory arthritis. *Journal of Biomedical Optics*, 17(9):0960131, 2012.
- [65] Zhihua Xie, Yanqing Yang, Yaqiong He, Chengyou Shu, Dong Chen, Jinke Zhang, Jingqin Chen, Chengbo Liu, Zonghai Sheng, Huadong Liu, Jie Liu, Xiaojing Gong, Liang Song, and Shaohong Dong. In vivo assessment of inflammation in carotid atherosclerosis by noninvasive photoacoustic imaging. *Theranostics*, 10(10):4694–4704, 2020.
- [66] Lei Xi, Stephen R. Grobmyer, Lei Wu, Ruimin Chen, Guangyin Zhou, Luke G. Gutwein, Jingjing Sun, Wenjun Liao, Qifa Zhou, Huikai Xie, and Huabei Jiang. Evaluation of breast tumor margins in vivo with intraoperative photoacoustic imaging. *Optics Express*, 20(8):8726, 2012.
- [67] Muyinatu A. Lediju Bell. Photoacoustic imaging for surgical guidance: Principles, applications, and outlook. *Journal of Applied Physics*, 128(6), 2020.
- [68] Toshiyuki Kitai, Masae Torii, Tomoharu Sugie, Shotaro Kanao, Yoshiki Mikami, Tsuyoshi Shiina, and Masakazu Toi. Photoacoustic mammography: Initial clinical results. *Breast Cancer*, 21(2):146–153, 2014.
- [69] Gaël Gael Diot, Stephan Metz, Aurelia Noske, Evangelos Liapis, Barbara Schroeder, Saak V. Ovsepian, Reinhard Meier, Ernst Rummeny, and Vasilis Ntziachristos. Multi-Spectral Photoacoustic Tomography (MSOT) of human breast cancer. *Clinical Cancer Research*, 23(22):6912–6922, 2017.

- [70] Anne Becker, Max Masthoff, Jing Claussen, Steven James Ford, Wolfgang Roll, Matthias Burg, Peter J. Barth, Walter Heindel, Michael Schäfers, Michel Eisenblätter, and Moritz Wildgruber. Multispectral optoacoustic tomography of the human breast: characterisation of healthy tissue and malignant lesions using a hybrid ultrasound-optoacoustic approach. *European Radiology*, 28(2):602–609, feb 2018.
- [71] Gisela L.G. Menezes, Ruud M. Pijnappel, Carla Meeuwis, Robertus Bisschops, Jeroen Veltman, Philip T. Lavin, Marc J. Van De Vijver, and Ritse M. Mann. Downgrading of breast masses suspicious for cancer by using optoacoustic breast imaging. *Radiology*, 288(2):355–365, 2018.
- [72] H. Y. Choi, Y. M. Sohn, and M. Seo. Comparison of 3D and 2D shear-wave elastography for differentiating benign and malignant breast masses: focus on the diagnostic performance. *Clinical Radiology*, 72(10):878–886, 2017.
- [73] André Farrokh, Nicolai Maass, Linn Treu, Thorsten Heilmann, and Fritz K.W. Schäfer. Accuracy of tumor size measurement: comparison of B-mode ultrasound, strain elastography, and 2D and 3D shear wave elastography with histopathological lesion size. *Acta Radiologica*, 60(4):451–458, 2019.
- [74] Gabriele Bergers and Laura E Benjamin. Tumorigenesis and the angiogenic switch. *Nature reviews. Cancer*, 3(6):401–10, jun 2003.
- [75] Bruce J. Tromberg, Brian W. Pogue, Keith D. Paulsen, Arjun G. Yodh, David a. Boas, and Albert E. Cerussi. Assessing the future of diffuse optical imaging technologies for breast cancer management. *Medical Physics*, 35(6):2443, 2008.
- [76] C G Hoelen, F F de Mul, R Pongers, and a Dekker. Three-dimensional photoacoustic imaging of blood vessels in tissue. *Optics letters*, 23(8):648–50, apr 1998.
- [77] Robert a. Kruger, Richard B. Lam, Daniel R. Reinecke, Stephen P. Del Rio, and Ryan P. Doyle. Photoacoustic angiography of the breast. *Medical Physics*, 37(11):6096, 2010.
- [78] Srirang Manohar, Susanne E Vaartjes, Johan C G van Hespén, Joost M Klaase, Frank M van den Engh, Wiendelt Steenbergen, and Ton G van Leeuwen. Initial results of in vivo non-invasive cancer

- imaging in the human breast using near-infrared photoacoustics. *Optics express*, 15(19):12277–85, sep 2007.
- [79] Srirang Manohar, Alexei Kharine, Johan C. G. van Hespén, Wiendelt Steenbergen, and Ton G. van Leeuwen. Photoacoustic mammography laboratory prototype: imaging of breast tissue phantoms. *Journal of Biomedical Optics*, 9(6):1172, 2004.
- [80] Jan Laufer, Dave Delpy, Clare Elwell, and Paul Beard. Quantitative spatially resolved measurement of tissue chromophore concentrations using photoacoustic spectroscopy: application to the measurement of blood oxygenation and haemoglobin concentration. *Physics in medicine and biology*, 52(1):141–68, jan 2007.
- [81] Lihong V. Wang, editor. *Photoacoustic imaging and spectroscopy*. CRC Press, Boca Raton, 1st edition, 2009.
- [82] Lihong V Wang and Song Hu. Photoacoustic tomography: in vivo imaging from organelles to organs. *Science (New York, N.Y.)*, 335(6075):1458–62, mar 2012.
- [83] Alexander A. Oraevsky, Elena V. Savateeva, Sergey V. Solomatin, Alexander A. Karabutov, Valeri G. Andreev, Zoran Gatalica, Tuenchit Khamapirad, and P. M. Henrichs. <title>Optoacoustic imaging of blood for visualization and diagnostics of breast cancer</title>. In Alexander A. Oraevsky, editor, *Proceedings of SPIE*, pages 81–94. International Society for Optics and Photonics, jun 2002.
- [84] Xavier Intes. Time-domain optical mammography SoftScan: initial results. *Academic radiology*, 12(8):934–47, aug 2005.
- [85] Nikolaus Knauer, Xose Luis Dean-Ben, and Daniel Razansky. Spatial compounding of volumetric data enables freehand optoacoustic angiography of large-scale vascular networks. *IEEE Transactions on Medical Imaging*, 39(4):1160–1169, 2020.
- [86] Jerome Gateau, Miguel Angel Araque Caballero, Alexander Dima, and Vasilis Ntziachristos. Three-dimensional optoacoustic tomography using a conventional ultrasound linear detector array: whole-body tomographic system for small animals. *Medical physics*, 40(1):013302, jan 2013.

- [87] Pim J. van den Berg, Khalid Daoudi, Hein J. Bernelot Moens, and Wiendelt Steenbergen. Feasibility of photoacoustic/ultrasound imaging of synovitis in finger joints using a point-of-care system. *Photoacoustics*, 8:8–14, 2017.
- [88] Hamid Moradi, Shuo Tang, and Septimiu E. Salcudean. Deconvolution based photoacoustic reconstruction for directional transducer with sparsity regularization. In Alexander A. Oraevsky and Lihong V. Wang, editors, *Proc. SPIE*, page 97082D, mar 2016.
- [89] Pingyu Liu. The P -transform and photoacoustic image reconstruction. *Physics in Medicine and Biology*, 43(3):667–674, mar 1998.
- [90] Valeri G. Andreev, Dmitriy a. Popov, Dmitriy V. Sushko, Alexander a. Karabutov, and Alexander a. Oraevsky. Image reconstruction in 3D optoacoustic tomography system with hemispherical transducer array. In Alexander A. Oraevsky, editor, *Proceedings of SPIE*, volume 4618, pages 137–145, jun 2002.
- [91] C G Hoelen and F F de Mul. Image reconstruction for photoacoustic scanning of tissue structures. *Applied optics*, 39(31):5872–83, nov 2000.
- [92] K P Köstli, Martin Frenz, H P Weber, G Paltauf, and H Schmidt-Kloiber. Optoacoustic tomography: time-gated measurement of pressure distributions and image reconstruction. *Applied optics*, 40(22):3800–9, aug 2001.
- [93] Dazi Feng, Yuan Xu, Geng Ku, and Lihong V. Wang. Microwave-induced thermoacoustic tomography: Reconstruction by synthetic aperture. *Medical Physics*, 28(12):2427, 2001.
- [94] C K Liao, M L Li, and P C Li. Optoacoustic imaging with synthetic aperture focusing and coherence weighting. *Optics letters*, 29(21):2506–8, nov 2004.
- [95] K.P. Kostli, Daniel Frauchiger, J.J. Niederhauser, Günther Paltauf, H.P. Weber, and Martin Frenz. Optoacoustic imaging using a three-dimensional reconstruction algorithm. *IEEE Journal of Selected Topics in Quantum Electronics*, 7(6):918–923, 2001.
- [96] S J Norton and M Linzer. Ultrasonic reflectivity imaging in three dimensions: exact inverse scattering solutions for plane, cylindrical, and spherical apertures. *IEEE transactions on bio-medical engineering*, 28(2):202–20, feb 1981.

- [97] Peng Shao, Tyler Harrison, and Roger J Zemp. Iterative algorithm for multiple illumination photoacoustic tomography (MIPAT) using ultrasound channel data. *Biomedical optics express*, 3(12):3240–9, dec 2012.
- [98] Thomas Jetzfellner, Amir Rosenthal, Andreas Buehler, Alexander Dima, Karl Hans Englmeier, Vasilis Ntziachristos, and Daniel Razansky. Tomographic optoacoustic inversion in dynamic illumination scenarios. In Alexander A. Oraevsky and Lihong V. Wang, editors, *Photons Plus Ultrasound: Imaging and Sensing 2011*, page 78991Y, feb 2011.
- [99] Allan Rosencwaig and Allen Gersho. Theory of the photoacoustic effect with solids. *Journal of Applied Physics*, 47(1):64–69, 1976.
- [100] Minghua Xu and Lihong Wang. Universal back-projection algorithm for photoacoustic computed tomography. *Physical Review E*, 71(1):016706, jan 2005.
- [101] Chao Huang, Kun Wang, Liming Nie, Lihong V. Wang, and Mark A. Anastasio. Full-wave iterative image reconstruction in photoacoustic tomography with acoustically inhomogeneous media. *IEEE Transactions on Medical Imaging*, 32(6):1097–1110, 2013.
- [102] Joemini Poudel, Yang Lou, and Mark A. Anastasio. A survey of computational frameworks for solving the acoustic inverse problem in three-dimensional photoacoustic computed tomography. *Physics in Medicine and Biology*, 64(14):14TR01, 2019.
- [103] Joemini Poudel, Shuai Na, Lihong V. Wang, and Mark A. Anastasio. Iterative image reconstruction in transcranial photoacoustic tomography based on the elastic wave equation. *Physics in Medicine and Biology*, 65(5), 2020.
- [104] Jingsong Li, Benli Yu, Weixiong Zhao, and Weidong Chen. A review of signal enhancement and noise reduction techniques for tunable diode laser absorption spectroscopy. *Applied Spectroscopy Reviews*, 49(8):666–691, 2014.
- [105] Rayyan Manwar, Matin Hosseinzadeh, Ali Hariri, Karl Kratkiewicz, Shahryar Noei, and Mohammad R.N. Avanaki. Photoacoustic signal enhancement: Towards utilization of low energy laser diodes in real-time photoacoustic imaging. *Sensors (Switzerland)*, 18(10), 2018.

- [106] Emma R Hill, Wenfeng Xia, Matthew J Clarkson, and Adrien E Desjardins. Identification and removal of laser-induced noise in photoacoustic imaging using singular value decomposition. *Biomedical optics express*, 8(1):68–77, 2017.
- [107] K. J. Parker, S. R. Huang, R. A. Musulin, and R. M. Lerner. Tissue response to mechanical vibrations for "sonoelasticity imaging". *Ultrasound in Medicine and Biology*, 16(3):241–246, 1990.
- [108] Ali Baghani, Hani Eskandari, Weiqi Wang, Daniel Da Costa, Mohamed Nabil Lathiff, Ramin Sahebjavaher, Septimiu Salcudean, and Robert Rohling. Real-time quantitative elasticity imaging of deep tissue using free-hand conventional ultrasound. In Nicholas Ayache, Hervé Delingette, Polina Golland, and Kensaku Mori, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2012*, pages 617–624, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- [109] Jeffrey Abeysekera, Robert Rohling, and Septimiu Salcudean. Vibro-elastography: Absolute elasticity from motorized 3D ultrasound measurements of harmonic motion vectors. *2015 IEEE International Ultrasonics Symposium, IUS 2015*, pages 3–6, 2015.
- [110] M. Honarvar, J. Lobo, O. Mohareri, S. E. Salcudean, and R. Rohling. Direct vibro-elastography FEM inversion in Cartesian and cylindrical coordinate systems without the local homogeneity assumption. *Physics in Medicine and Biology*, 60(9):3847–3868, 2015.
- [111] Ali Baghani. *A Wave Equation Approach to Ultrasound Elastography*. PhD thesis, University of British Columbia, 2010.
- [112] Mohammad Honarvar. *Dynamic Elastography With Finite Element-Based Inversion*. PhD thesis, University of British Columbia, 2015.
- [113] Jeffrey Michael Abeysekera. *Three Dimensional ultrasound imaging*. PhD thesis, University of British Columbia, 2016.
- [114] Caitlin Schneider, Mohammad Honarvar, Julio Lobo, Edward V Tam, Robert Rohling, and Septimiu Salcudean. S-wave for liver elastography: Repeatability and in vivo study. In *Hepatology*, volume 68, pages 612A–613A. WILEY 111 RIVER ST, HOBOKEN 07030-5774, NJ USA, 2018.

- [115] Qi Zeng, Mohammad Honarvar, Julio Lobo, Caitlin Schneider, Robert Rohling, Anup Agarwal, Gerard Harrison, Changhong Hu, Scott Di-anis, James Jago, and Septimiu E. Salcudean. 3D Liver Shear Wave Absolute Vibro-Elastography with an xMATRIX Array - A Healthy Volunteer Study. In *2018 IEEE International Ultrasonics Symposium (IUS)*, pages 1–9. IEEE, oct 2018.
- [116] Mehdi Moradi, S. Sara Mahdavi, Guy Nir, Omid Mohareri, Anthony Koupparis, Louis Olivier Gagnon, Ladan Fazli, Rowan G. Casey, Joseph Ischia, Edward C. Jones, S. Larry Goldenberg, and Septimiu E. Salcudean. Multiparametric 3D in vivo ultrasound vibroelastography imaging of prostate cancer: Preliminary results. *Medical Physics*, 41(7), 2014.
- [117] Julio Lobo, Ali Baghani, Hani Eskandari, Sara Mahdavi, Robert Rohling, Larry Goldenberg, William James Morris, and Septimiu Salcudean. Prostate vibro-elastography: Multi-frequency 1D over 3D steady-state shear wave imaging for quantitative elastic modulus measurement. *2015 IEEE International Ultrasonics Symposium, IUS 2015*, pages 3–6, 2015.
- [118] Jeffrey M. Abeysekera, Manyou Ma, Mehran Pesteie, Jefferson Terry, Denise Pugash, Jennifer A. Hutcheon, Chantal Mayer, Lutz Lampe, Septimiu Salcudean, and Robert Rohling. SWAVE Imaging of Placental Elasticity and Viscosity: Proof of Concept. *Ultrasound in Medicine and Biology*, 43(6):1112–1124, 2017.
- [119] Caitlin Schneider, Mohammad Honarvar, Julio Lobo, Robert Rohling, Tim Salcudean, Samir Bidnur, and Christopher Nguan. Blood pressure dependent elasticity measurements of porcine kidney ex-vivo. *IEEE International Ultrasonics Symposium, IUS*, 2016-Novem:16–19, 2016.
- [120] Hani Eskandari, Septimiu E Salcudean, Robert Rohling, Ali Baghani, Samuel Frew, Paula B Gordon, and Linda Warren. Identifying malignant and benign breast lesions using vibroelastography. In *2013 IEEE International Ultrasonics Symposium (IUS)*, pages 25–28. IEEE, jul 2013.
- [121] Yanan Shao, Hoda Hashemi, Paula Gordon, Linda Warren, Z.Jane Wang, Robert Rohling, and Tim Salcudean. Breast Cancer Detection using Multimodal Time Series Features from Ultrasound Shear Wave

- Absolute Vibro-Elastography. *IEEE Journal of Biomedical and Health Informatics*, 2194(c):1–1, 2021.
- [122] M. Heijblom, D. Piras, W. Xia, J.C.G. van Hespén, J.M. Klaase, F.M. van den Engh, T.G. van Leeuwen, W. Steenbergen, and S. Manohar. Visualizing breast cancer using the Twente photoacoustic mammoscope: What do we learn from twelve new patient measurements? *Optics Express*, 20(11):11582, 2012.
- [123] G Zografos, D Koulocheri, P Liakou, M Sofras, S Hadjiagapis, M Orme, and V Marmarelis. Novel technology of multimodal ultrasound tomography detects breast lesions. *European radiology*, 23(3):673–83, mar 2013.
- [124] Ioana Boca (Bene), Anca Ileana Ciurea, Cristiana Augusta Ciortea, and Sorin Marian Dudea. Pros and Cons for Automated Breast Ultrasound (ABUS): A Narrative Review. *Journal of Personalized Medicine*, 11(8):703, 2021.
- [125] Wenfeng Xia, Daniele Piras, Mithun K. a. Singh, Johan C. G. van Hespén, Spiridon van Veldhoven, Christian Prins, Ton G. van Leeuwen, Wiendelt Steenbergen, and Srirang Manohar. Design considerations for ultrasound detectors in photoacoustic breast imaging. In Alexander A. Oraevsky and Lihong V. Wang, editors, *Photons Plus Ultrasound: Imaging and Sensing 2013*, volume 8581, pages 171–177, mar 2013.
- [126] Valeriy G Andreev, Alexander a Karabutov, and Alexander a Oraevsky. Detection of ultrawide-band ultrasound pulses in optoacoustic tomography. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 50(10):1383–90, oct 2003.
- [127] Wenfeng Xia, Daniele Piras, Johan C G van Hespén, Spiridon van Veldhoven, Christian Prins, Ton G van Leeuwen, Wiendelt Steenbergen, and Srirang Manohar. An optimized ultrasound detector for photoacoustic breast tomography. *Medical physics*, 40(3):032901, mar 2013.
- [128] Corey J. Kelly and Septimiu E. Salcudean. Automated full-breast photoacoustic tomography with non-uniform illumination. *Photons Plus Ultrasound: Imaging and Sensing 2018*, 1049405(February):5, 2018.

- [129] Reza Zahiri-Azar and Septimiu E Salcudean. Motion estimation in ultrasound images using time domain cross correlation with prior estimates. *IEEE Trans. on BioMedical Engineering*, 53(10):1990–2000, 2006.
- [130] Carolina Nascimben Matheus and Elaine Caldeira de Oliveira Guirro. Change in blood flow velocity demonstrated by Doppler ultrasound in upper limb after axillary dissection surgery for the treatment of breast cancer. *Breast Cancer Research and Treatment*, 127(3):697–704, jun 2011.
- [131] H. Sato, M. Katano, T. Takigawa, and T. Masuda. Estimation for the elasticity of vascular endothelial cells on the basis of atomic force microscopy and Young’s modulus of gelatin gels. *Polymer Bulletin*, 47(3-4):375–381, dec 2001.
- [132] Elastic Moduli of Breast and Prostate Tissues under Compression. *Ultrasonic Imaging*, 20(4):260–274, oct 1998.
- [133] Sabrina Stocker, Florian Foschum, Philipp Krauter, Florian Bergmann, Ansgar Hohmann, Claudia Scalfi Happ, and Alwin Kienle. Broadband Optical Properties of Milk. *Applied Spectroscopy*, 71(5):951–962, 2017.
- [134] Steven L Jacques. Optical properties of biological tissues: a review. *Physics in Medicine and Biology*, 58(11):R37–R61, 2013.
- [135] Corey Kelly, Amir Refaee, and Septimiu E Salcudean. Integrating photoacoustic tomography into a multimodal automated breast ultrasound scanner. *Journal of Biomedical Optics*, 25(11):1–18, nov 2020.
- [136] Khalid Daoudi, P J van den Berg, O Rabot, A Kohl, S Tisserand, P Brands, and W Steenbergen. Handheld probe integrating laser diode and ultrasound transducer array for ultrasound/photoacoustic dual modality imaging. *Optics express*, 22(21):26365–26374, 2014.
- [137] Andreas Buehler, Marcin Kacprowicz, Adrian Taruttis, and Vasilis Ntziachristos. Real-time handheld multispectral optoacoustic imaging. *Optics Letters*, 38(9):1404, 2013.
- [138] Kalloor Joseph Francis, Yoeri E. Boink, Maura Dantuma, Mithun Kuniyil Ajith Singh, Srirang Manohar, and Wiendelt Steenbergen. Tomographic imaging with an ultrasound and LED-based photoacoustic system. *Biomedical Optics Express*, 11(4):2152, 2020.

- [139] Yang Lou, Kun Wang, Alexander A. Oraevsky, and Mark A. Anastasio. Impact of nonstationary optical illumination on image reconstruction in optoacoustic tomography. *Journal of the Optical Society of America A*, 33(12):2333, 2016.
- [140] Wonseok Choi, Donghyeon Oh, and Chulhong Kim. Practical photoacoustic tomography: Realistic limitations and technical solutions. *Journal of Applied Physics*, 127(23):230903, jun 2020.
- [141] Niko Hänninen, Aki Pulkkinen, Alekski Leino, and Tanja Tarvainen. Application of diffusion approximation in quantitative photoacoustic tomography in the presence of low-scattering regions. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 250, 2020.
- [142] Seonyeong Park, Alexander A. Oraevsky, Richard Su, and Mark A. Anastasio. Compensation for non-uniform illumination and optical fluence attenuation in three-dimensional optoacoustic tomography of the breast. In Alexander A. Oraevsky and Lihong V. Wang, editors, *Photons Plus Ultrasound: Imaging and Sensing 2019*, page 180. SPIE, mar 2019.
- [143] Sjoukje M Schoustra, Daniele Piras, Roeland Huijink, Tim J P M Op 't Root, Laurens Alink, Wouter Muller Kobold, Wiendelt Steenbergen, and Srirang Manohar. Twente Photoacoustic Mammoscope 2: system overview and three-dimensional vascular network images in healthy breasts. *Journal of biomedical optics*, 24(12):1–12, 2019.
- [144] Min Ai, Weihang Shu, Tim Salcudean, Robert Rohling, Purang Abolmaesumi, and Shuo Tang. Design of high energy laser pulse delivery in a multimode fiber for photoacoustic tomography. *Optics Express*, 25(15):17713, 2017.
- [145] Lev E Givon, Thomas Unterthiner, N Benjamin Erichson, David Wei Chiang, Eric Larson, Luke Pfister, Sander Dieleman, Gregory R Lee, Stefan van der Walt, Bryant Menn, Teodor Mihai Moldovan, Frédéric Bastien, Xing Shi, Jan Schlüter, Brian Thomas, Chris Capdevila, Alex Rubinsteyn, Michael M Forbes, Jacob Frelinger, Tim Klein, Bruce Merry, Nate Merrill, Lars Pastewka, Li Yong Liu, S Clarkson, Michael Rader, Steve Taylor, Arnaud Bergeron, Nikul H Ukani, Feng Wang, Wing-Kit Lee, and Yiyin Zhou. scikit-cuda 0.5.3: a Python interface to GPU-powered libraries, may 2019.

- [146] A. Manduca, D. S. Lake, S. A. Kruse, and R. L. Ehman. Spatio-temporal directional filtering for improved inversion of MR elastography images. *Medical Image Analysis*, 7(4):465–473, 2003.
- [147] H. Moradi, S. Tang, and S. E. Salcudean. Deconvolution based photoacoustic reconstruction with sparsity regularization. *Optics Express*, 25(3):2771, 2017.
- [148] Zhou Wang, Alan Conrad Bovik, Hamid Rahim Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [149] Stéfan Van Der Walt, Johannes L. Schönberger, Juan Nunez-Iglesias, François Boulogne, Joshua D. Warner, Neil Yager, Emmanuelle Gouillart, and Tony Yu. Scikit-image: Image processing in python. *PeerJ*, 2014(1):1–18, 2014.
- [150] John C Russ and F Brent Neal. *The Image Processing Handbook*. CRC Press, Inc., USA, 7th edition, 2015.
- [151] Parker Wray, Li Lin, and Peng Hu. Photoacoustic computed tomography of human extremities. *Journal of Biomedical Optics*, 24(02):1, 2019.
- [152] Amir Refaee*, Corey J. Kelly*, Hamid Moradi, and Septimiu E. Salcudean. Denoising of pre-beamformed photoacoustic data using generative adversarial networks. *Biomedical Optics Express*, 12:6184, oct 2021. *equal authorship contribution.
- [153] Maryam Basij, Moein Mozaffarzadeh, Roya Paridar, Mahdi Orooji, and Mohammad Mehrmohammadi. Delay-multiply-and-standard-deviation weighting factor improves image quality in linear-array photoacoustic tomography. In Alexander A. Oraevsky and Lihong V. Wang, editors, *Photons Plus Ultrasound: Imaging and Sensing 2019*, pages 601–607. SPIE, feb 2019.
- [154] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Networks, jun 2014.
- [155] Mehdi Mirza and Simon Osindero. Conditional Generative Adversarial Nets. *ArXiv e-prints*, pages 1–7, nov 2014.

- [156] Antonia Creswell, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A. Bharath. Generative Adversarial Networks: An Overview. *IEEE Signal Processing Magazine*, 35(1):53–65, jan 2018.
- [157] Phillip Isola, Jun Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, 2017-Janua:5967–5976, 2017.
- [158] Janek Gröhl, Melanie Schellenberg, Kris Dreher, and Lena Maier-Hein. Deep learning for biomedical photoacoustic imaging: A review. *Photoacoustics*, 22(January):100241, jun 2021.
- [159] Emran Mohammad Abu Anas, Haichong K. Zhang, Chloé Audigier, and Emad M. Bector. Robust Photoacoustic Beamforming Using Dense Convolutional Neural Networks. In *Simulation, Image Processing, and Ultrasound Systems for Assisted Diagnosis and Navigation*, pages 3–11. Springer International Publishing, Cham, 2018.
- [160] Emran Mohammad Abu Anas, Haichong K. Zhang, Jin Kang, and Emad Bector. Enabling fast and high quality LED photoacoustic imaging: a recurrent neural networks based approach. *Biomedical Optics Express*, 9(8):3852, aug 2018.
- [161] Ali Hariri, Kamran Alipour, Yash Mantri, Jurgen P. Schulze, and Jesse V. Jokerst. Deep learning improves contrast in low-fluence photoacoustic imaging. *Biomedical Optics Express*, 11(6):3360, 2020.
- [162] Hengrong Lan, Kang Zhou, Changchun Yang, Jun Cheng, Jiang Liu, Shenghua Gao, and Fei Gao. Ki-GAN: Knowledge Infusion Generative Adversarial Network for Photoacoustic Image Reconstruction In Vivo. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, pages 273–281. Springer, Cham, 2019.
- [163] Tri Vu, Mucong Li, Hannah Humayun, Yuan Zhou, and Junjie Yao. A generative adversarial network for artifact removal in photoacoustic computed tomography with a linear-array transducer. *Experimental Biology and Medicine*, 245(7):597–605, apr 2020.
- [164] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medi-*

cal Image Computing and Computer-Assisted Intervention (MICCAI), volume 9351, pages 234–241, Switzerland, 2015. Springer.

- [165] Kun Wang, Chao Huang, Yu Jiun Kao, Cheng Ying Chou, Alexander A. Oraevsky, and Mark A. Anastasio. Accelerating image reconstruction in three-dimensional optoacoustic tomography on graphics processing units. *Medical Physics*, 40(2), 2013.
- [166] Amir Beck and Marc Teboulle. A Fast Iterative Shrinkage-Thresholding Algorithm. *Society for Industrial and Applied Mathematics Journal on Imaging Sciences*, 2(1):183–202, 2009.
- [167] Amir Beck and Marc Teboulle. Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems. *IEEE Transactions on Image Processing*, 18(11):2419–2434, 2009.
- [168] Antonin Chambolle. An Algorithm for Total Variation Minimization and Applications. *Journal of Mathematical Imaging and Vision*, 20(1-2):89–97, 2004.
- [169] Lin Zhang, Lei Zhang, Xuanqin Mou, and David Zhang. FSIM: A Feature Similarity Index for Image Quality Assessment. *IEEE Transactions on Image Processing*, 20(8):2378–2386, aug 2011.
- [170] M. U. Müller, N. Ekhtiari, R. M. Almeida, and C. Rieke. Super-Resolution of Multispectral Satellite Images Using Convolutional Neural Networks. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 5(1):33–40, 2020.
- [171] Peter Kovesi. Image Features from Phase Congruency. *Videre: Journal of Computer Vision Research*, 1(3):1–26, 1999.
- [172] Yang Lou, Weimin Zhou, Thomas P. Matthews, Catherine M. Appleton, and Mark A. Anastasio. Generation of anatomically realistic numerical phantoms for photoacoustic and ultrasonic breast imaging. *Journal of Biomedical Optics*, 22(4):041015, 2017.
- [173] Ian Goodfellow. NIPS 2016 Tutorial: Generative Adversarial Networks, dec 2016.
- [174] Lin Zhang, Lei Zhang, Xuanqin Mou, and David Zhang. A comprehensive evaluation of full reference image quality assessment algorithms. In *2012 19th IEEE International Conference on Image Processing*, pages 1477–1480. IEEE, sep 2012.

- [175] Han Vanholder. Efficient inference with tensorrt, 2016.
- [176] Muyinatu A. Lediju Bell, Xiaoyu Guo, Danny Y. Song, and Emad M. Boctor. Transurethral light delivery for prostate photoacoustic imaging. *Journal of Biomedical Optics*, 20(3):036002, 2015.
- [177] Min Ai, Jong-in Youn, Septimiu E. Salcudean, Robert Rohling, Purang Abolmaesumi, and Shuo Tang. Photoacoustic tomography for imaging the prostate: a transurethral illumination probe design and application. *Biomedical Optics Express*, 10(5):2588, 2019.
- [178] Maëva Vallet, François Varray, Jérôme Boutet, Jean Marc Dinten, Giosuè Caliano, Alessandro Stuart Savoia, and Didier Vray. Quantitative comparison of PZT and CMUT probes for photoacoustic imaging: Experimental validation. *Photoacoustics*, 8:48–58, 2017.
- [179] Johannes Rebling, Omri Warshavski, Cyril Meynier, and Daniel Razansky. Broadband optoacoustic characterization of cMUT and PZT transducer directivity in receive mode. In Neb Duric and Brecht Heyde, editors, *Medical Imaging 2017: Ultrasonic Imaging and Tomography*, page 101391K, mar 2017.
- [180] Carlos D. Gerardo, Edmond Cretu, and Robert Rohling. Fabrication and testing of polymer-based capacitive micromachined ultrasound transducers for medical imaging. *Microsystems and Nanoengineering*, 4(1), 2018.
- [181] Zhenhao Li, Afshin Kashani Ilkhechi, and Roger Zemp. Transparent capacitive micromachined ultrasonic transducers (CMUTs) for photoacoustic applications. *Optics Express*, 27(9):13204, 2019.
- [182] Sarah E Bohndiek, Ben Cox, James Joseph, Janek Grohl, Lina Hacker, Stefan Morscher, and William C. Vogt. IPASC: a Community-Driven Consensus-Based Initiative Towards Standardisation in Photoacoustic Imaging. In *2020 IEEE International Ultrasonics Symposium (IUS)*, pages 1–4. IEEE, sep 2020.

Appendix A

Supporting Information

A.1 Denoising methods comparison⁴

Figure A.1 shows a comparison of denoising methods applied to our raw RF data, along with peak signal-to-noise ratio (PSNR) measurements for an A-line which passes through both a known photoacoustic signal and a noise band. We define PSNR here as the peak amplitude of the known signal (the green region in Figure A.1) divided by the standard deviation of the entire A-line.

First, we use the spatio-temporal filtering method of Manduca *et al* [146], with the filter oriented to remove vertical features. While this method did significantly reduce the prominence of the noise bands relative to the photoacoustic signals, as evidenced by the PSNR improvement, the bands are still clearly visible in the data. We expect this is due to these bands being non-uniform across different elements, resulting in significant frequency content passing the filter.

With the SVD-based method of Hill *et al* [106], as described in the main text, we see that the noise bands are completely suppressed, and the PSNR is higher than that achieved with the directional filtering method.

⁴This section originally included as supplemental material for the published version of Chapter 3

A.1. Denoising methods comparison

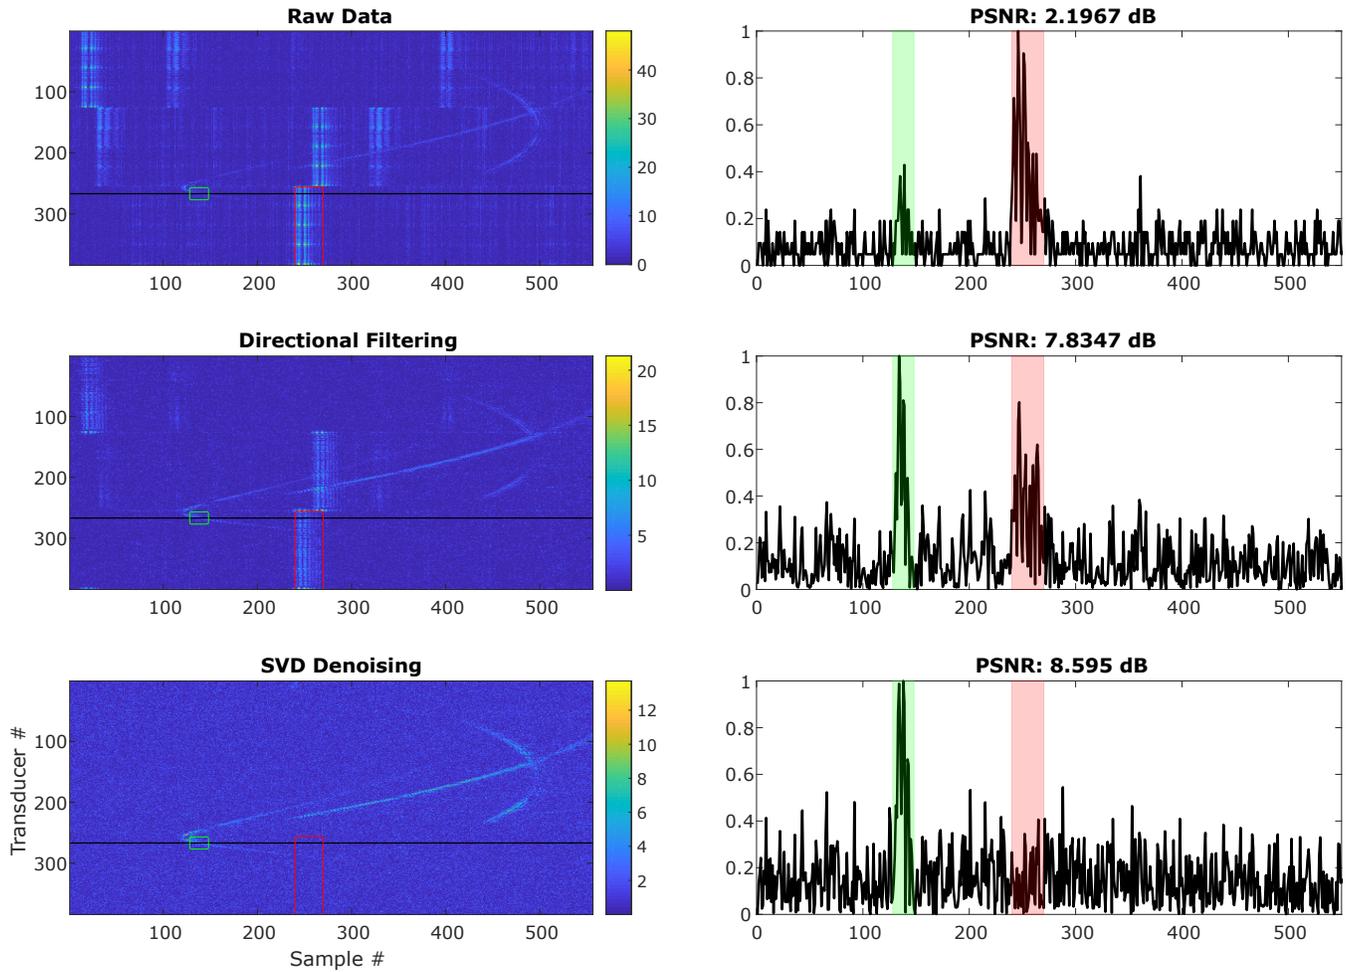


Figure A.1: Comparison of data denoising methods. Images on the right show RF data for each transducer element. Plots on the right are the RF signal along the black line in the image plot, with actual photoacoustic signal highlighted in green and a representative noise band highlighted in red.

A.2 GAN reference dataset - number of averaged frames⁵

In choosing the number of RF frames to average for our reference dataset, we sought a balance between data quality and scan time. Figure A.2 illustrates this payoff for a representative plane of our training data, by measuring SSIM and FSIM as a function of the number of frames averaged, when compared to the average of 100 frames total. For the present study, we found that 20 frames offered the best balance.

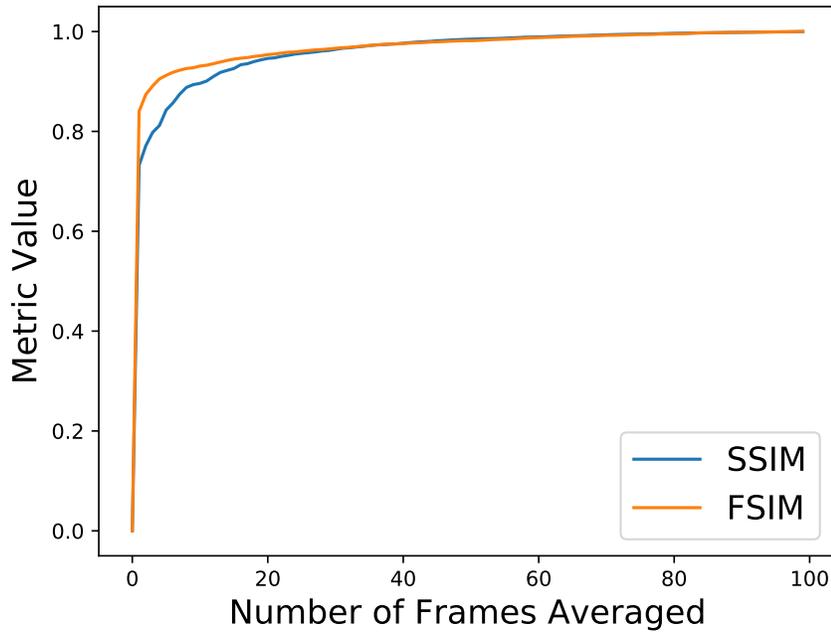


Figure A.2: Quality metrics as a function of the number of averaged RF frames per scan position.

⁵This section originally included as supplemental material for the published version of Chapter 4

A.3 Number of FISTA iterations⁶

To determine how many iterations of our reconstruction algorithm are required, we can plot the minimization objective, defined in Equation (4.7), as a function of the iteration number. This relationship is shown in Figure A.3 for a representative imaging plane from our testing dataset. We chose to stop all reconstructions in the present study after 20 iterations, as the objective function changes very little beyond this point.

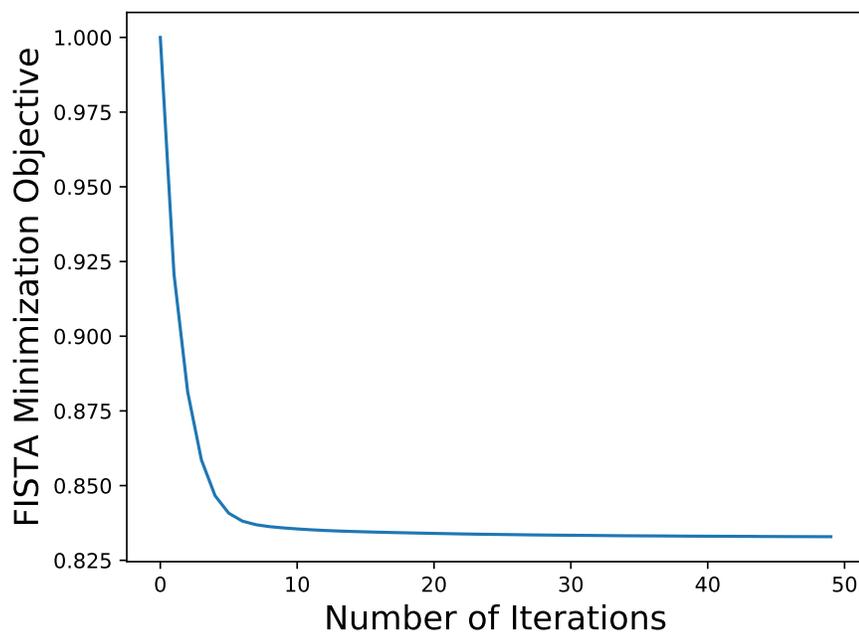


Figure A.3: FISTA minimization objective as a function of iteration number, which was used to choose the stopping point for our image reconstruction.

⁶This section originally included as supplemental material for the published version of Chapter 4

Appendix B

ABUS illuminator geometric optics calculations

B.1 Light Propagation From Optics to Dome

The surface (set of points), \mathbf{r}_s where the fan beam intersects the outer surface of the ABUS dome can be calculated by generating a set of rays, \mathbf{r}_1 , originating at \mathbf{U}_0 , bounded by the angular envelopes of the illuminator. The long axis of the fan beam corresponds to rotations of the optical axis, $\hat{\mathbf{U}}_z$, by angle θ around $\hat{\mathbf{U}}_x$, with $\theta \in \{-45^\circ, 45^\circ\}$. The short axis of the fan beam corresponds to rotations of the optical axis by angle ϕ around $\hat{\mathbf{U}}_y$, with $\phi \in \{-0.5^\circ, 0.5^\circ\}$. This gives

$$\hat{\mathbf{r}}_1(\theta, \phi) = \mathbf{R}(\theta, \hat{\mathbf{U}}_x) \cdot \mathbf{R}(\phi, \hat{\mathbf{U}}_y) \cdot \hat{\mathbf{U}}_z \quad (\text{B.1})$$

where $\mathbf{R}(\beta, \hat{\mathbf{v}})$ is the transformation matrix corresponding to a rotation of angle β about an axis defined by unit vector $\hat{\mathbf{v}}$. The ray $\mathbf{r}_1(\theta, \phi)$ with origin \mathbf{U}_0 will intersect the the sphere with centre \mathbf{c} and radius R at a point \mathbf{r}_d given by

$$\mathbf{r}_s(\theta, \phi) = \mathbf{U}_0 + d(\theta, \phi) \cdot \hat{\mathbf{r}}_1(\theta, \phi) \quad (\text{B.2})$$

where we want the intersection point closest to \mathbf{U}_0 , given by

$$\begin{aligned} d(\theta, \phi) &= \min \left(-\hat{\mathbf{r}}_1(\theta, \phi) \cdot (\mathbf{U}_0 - \mathbf{c}) \pm \sqrt{R^2 + ((\mathbf{U}_0 - \mathbf{c}) \cdot \hat{\mathbf{r}}_1(\theta, \phi))^2 - |\mathbf{U}_0 - \mathbf{c}|^2} \right) \\ &= -\hat{\mathbf{r}}_1(\theta, \phi) \cdot (\mathbf{U}_0 - \mathbf{c}) - \sqrt{R^2 + ((\mathbf{U}_0 - \mathbf{c}) \cdot \hat{\mathbf{r}}_1(\theta, \phi))^2 - |\mathbf{U}_0 - \mathbf{c}|^2}. \end{aligned} \quad (\text{B.3})$$

Since the illuminator produces a divergent beam it will illuminate an area A_s of the sample surface with energy total energy Ψ_{tot} . The fluence (energy per area) incident upon a patch dA_s centered on a point $\mathbf{r}_s(\theta, \phi)$ is

$$I(\mathbf{r}_s(\theta, \phi)) = \frac{\Psi_0}{d(\theta, \phi)^2} \quad (\text{B.4})$$

where Ψ_0 is a normalization constant such that

$$\iint_{\theta, \phi} I(\mathbf{r}_s(\theta, \phi)) dA_s = \Psi_{\text{tot}}. \quad (\text{B.5})$$

B.2 Reflective Losses

Light intensity loss will occur at both interfaces of the optical window (air/window and window/sample) depending on the angle of incidence and the indices of refraction. We first apply Snell's law, i.e. for a ray incident upon a boundary at an angle θ_i relative to the surface normal, the corresponding transmission ray will have angle θ_t given by

$$n_1 \sin \theta_i = n_2 \sin \theta_t \quad (\text{B.6})$$

where n_i and n_t are the indices of refraction of the incident and refractive media, respectively. The optical window has a thickness of $t = 1.5$ mm and both surfaces have radius of curvature $R = 120$ mm. In this limit ($t \ll R$), the curvature of the window can be neglected over the propagation length of the ray, such that θ_t for the air-window interface will equal θ_i for the window-sample interface.

Having calculated the refraction angles, we can use the Fresnel equations to compute the fraction of reflected and transmitted power at the interfaces. The fraction of power reflected at a boundary is

$$R_{s,p} = \left| \frac{n_1 \cos \theta_{i,t} - n_2 \cos \theta_{t,i}}{n_1 \cos \theta_{i,t} + n_2 \cos \theta_{t,i}} \right|^2 \quad (\text{B.7})$$

where s, p are the s- and p-polarized components of the incident ray. Our incident light is diffuse and unpolarized, so we take the average

$$R_{\text{eff}} = \frac{1}{2}(R_s + R_p). \quad (\text{B.8})$$

The fraction of transmitted power is thus given by conservation of energy as

$$T_{\text{eff}} = 1 - R_{\text{eff}}. \quad (\text{B.9})$$

It is also worth noting that the window-sample interface satisfies $n_2 < n_1$, and there will be a *critical angle* of incidence θ_c , beyond which no light is transmitted, defined via Equation B.6 as

$$\theta_c = \arcsin \frac{n_2}{n_1}. \quad (\text{B.10})$$

B.2. Reflective Losses

At room temperature and $\lambda = 700$ nm, our indices of refraction are

$$n_{\text{air}} = 1.00 \quad n_{\text{polycarbonate}} = 1.58 \quad n_{\text{water}} = 1.33, \quad (\text{B.11})$$

resulting in a critical angle of 57° . Because of the shallow angle at which the optical “fan” meets the dome (the optical axis is not normal to the dome surface), the incident angle at the air-window interface is relatively large (18° to 67°). Fortunately, since the window has a higher index of refraction, $\theta_t < \theta_i$, it provides a “focusing” effect, such that the refracted angle (and the incident angle upon the second interface) is in the range 11° to 36° , and does not exceed the critical angle.

We find that the fraction of power from the incident beam that reaches the sample is 83% to 94%. While this may be small in terms of total power loss, it does significantly change the spatial distribution of power at the surface of the sample, which will subsequently change the fluence at a given point *in* the sample.

Appendix C

Photoacoustic reconstruction kernel implementation

```
1  __global__ void AT(float* Img,           // Image grid
2                      float* RF,         // RF data
3                      float* Omega,      // Solid angle
4                      float* Y_pos,      // Transducer coordinates
5                      float* Y_orn,      // Transducer unit normals
6                      int iter,         // Frame index
7                      float x0,         // Image x coordinates
8                      float y0,         // Image y coordinates
9                      float z0,         // Image z coordinates
10                     float freq,       // Sampling frequency
11                     float c,          // Speed of sound
12                     int n_samples,    // Samples per A-line
13                     float dpMin,      // Directivity dot product cutoff
14                     int angleCheck,   // Use transducer directivity?
15                     float sigmaE,     // Elevational directivity
16                     float sigmaL,     // Lateral directivity
17                     float res         // Reconstruction resolution
18                 )
19 {
20     // Image voxel and transducer element indices
21     int index = blockIdx.z + gridDim.z * (blockIdx.y + gridDim.y * blockIdx.x);
22     int el_idx = iter*blockDim.x+threadIdx.x;
23
24     // -----
25     // Voxel interpolation calculations
26     float voxel_x_pos_min = fmaf(blockIdx.x,res, x0); //
27     float voxel_y_pos_min = fmaf(blockIdx.y,res, y0); // Voxel coordinate lower bound
28     float voxel_z_pos_min = fmaf(blockIdx.z,res, z0); //
29
30     float voxel_x_pos_max = fmaf(blockIdx.x+1,res, x0); //
31     float voxel_y_pos_max = fmaf(blockIdx.y+1,res, y0); // Voxel coordinate upper bound
32     float voxel_z_pos_max = fmaf(blockIdx.z+1,res, z0); //
33
34     // Determining closest/furthest corners of voxel
35     float xs[] = {voxel_x_pos_min - Y_pos[el_idx*3 + 0],
36                  Y_pos[el_idx*3 + 0]-voxel_x_pos_max};
37     float ys[] = {voxel_y_pos_min - Y_pos[el_idx*3 + 1],
```

Appendix C. Photoacoustic reconstruction kernel implementation

```
38         Y_pos[el_idx*3 + 1]-voxel_y_pos_max};
39 float zs[] = {voxel_z_pos_min - Y_pos[el_idx*3 + 2],
40              Y_pos[el_idx*3 + 2]-voxel_z_pos_max};
41
42 float rx_min = fmaxf(fmaxf(xs[0], 0), xs[1]);
43 float rx_max = fminf(fminf(xs[0], 0), xs[1]);
44
45 float ry_min = fmaxf(fmaxf(ys[0], 0), ys[1]);
46 float ry_max = fminf(fminf(ys[0], 0), ys[1]);
47
48 float rz_min = fmaxf(fmaxf(zs[0], 0), zs[1]);
49 float rz_max = fminf(fminf(zs[0], 0), zs[1]);
50
51 // Min/max/avg radii for voxel corners
52 float r_min = norm3df (rx_min, ry_min, rz_min);
53 float r_max = r_min + res;
54 float r_avg = (r_min+r_max)/2;
55
56 // How many samples in this range?
57 int num_samples_between = abs(floor((freq/c)*r_min)-ceil((freq/c)*r_max));
58 // -----
59
60 // Axial unit vector components
61 float x_A = Y_orn[el_idx*9 + 0];
62 float y_A = Y_orn[el_idx*9 + 1];
63 float z_A = Y_orn[el_idx*9 + 2];
64
65 // Elevational unit vector components
66 float x_E = Y_orn[el_idx*9 + 3];
67 float y_E = Y_orn[el_idx*9 + 4];
68 float z_E = Y_orn[el_idx*9 + 5];
69
70 // Lateral unit vector components
71 float x_L = Y_orn[el_idx*9 + 6];
72 float y_L = Y_orn[el_idx*9 + 7];
73 float z_L = Y_orn[el_idx*9 + 8];
74
75 // Vector from transducer element to image voxel
76 float dx = (voxel_x_pos_min+0.5*res - Y_pos[el_idx*3 + 0]);
77 float dy = (voxel_y_pos_min+0.5*res - Y_pos[el_idx*3 + 1]);
78 float dz = (voxel_z_pos_min+0.5*res - Y_pos[el_idx*3 + 2]);
79
80 // Voxel angle dot products
81 float dotPE = (x_E*dx + y_E*dy + z_E*dz)/r_avg;
82 float dotPA = (x_A*dx + y_A*dy + z_A*dz)/r_avg;
83 float dotPL = (x_L*dx + y_L*dy + z_L*dz)/r_avg;
84
85 // Is voxel within acceptance cone?
86 if (dotPA > dpMin){
```

Appendix C. Photoacoustic reconstruction kernel implementation

```
87     // Clip values
88     if(abs(dotPA)>1.0)
89         dotPA = dotPA/abs(dotPA);
90     if(abs(dotPE)>1.0)
91         dotPE = dotPE/abs(dotPE);
92     if(abs(dotPL)>1.0)
93         dotPL = dotPL/abs(dotPL);
94
95     // Angles
96     float angE = asin(dotPE);
97     float angL = asin(dotPL);
98
99     // Solid angle calculations
100    float omegai = dotPA/(r_avg*r_avg);
101
102    // Total weighting factor
103    float weight=1;
104
105    // Apply transducer directivity
106    if(angleCheck){
107        float weightE = exp(-0.5*(pow(angE / sigmaE, 2)));
108        float weightL = exp(-0.5*(pow(angL / sigmaL, 2)));
109
110        weight = weightE * weightL;
111    }
112
113    // Relevant RF data start index
114    int ind = (int)ceil((freq/c)*r_min);
115
116    // Apply solid angle weighting
117    float val = weight*omegai/Omega[index];
118
119    // Sum RF samples that fall within this vowel
120    float RF_sum = 0;
121    for (int i=0; i<num_samples_between; i++){
122        int curr_idx = (el_idx)*n_samples + ind+i;
123
124        if(ind + i < n_samples)
125            RF_sum += RF[curr_idx];
126
127    }
128    // Multiply by weight and add to image
129    atomicAdd(&Img[index], val*RF_sum);
130 }
131 }
```

Appendix D

Photoacoustic trigger box code

```
1 // Enable Debug Mode
2 const int DEBUG_ON = 0;
3
4 // Assign output pins
5 const int fireOut = 3; // Lasing trigger output
6 const int daqOut = 4; // DAQ acquire trigger output
7 const int extTrGnd = 5; // Pulse request trigger ground pin
8 const int buttonGnd = 6; // Pulse request button ground pin
9 const int switchPow = 7; // On/off switch power
10 const int seqOut = 8; // US sequencer trigger output
11 const int qsOut = 9; // Q-switch trigger output
12 const int buttonIn = 10; // Pulse request button input
13 const int extTrig = 11; // Pulse request trigger input
14 const int switchIn = 12; // On/off switch input
15 const int ledOut = 13; // Status LED output
16
17 // Pulse frequency -> default 9.8 Hz
18 float freq = 9.8;
19 int loopPeriod = round(1000 / freq); // milliseconds
20
21 // Pulse width -> default 12
22 int pWidth = 12; // microsecond
23
24 int nFrames = 3;
25 int qsDelay = 250; // microseconds
26
27 int frameCount = 0; // Frame count accumulator
28 int toggle = 0;
29
30 int LED = HIGH; // Is the LED indicator on?
31
32 long blinkInterval = 1000; // milliseconds
33 long t0 = 0;
34
35 long loopDelay = 500; // Delays the triggering loop
36 long tLoopOut = 0;
37
```

Appendix D. Photoacoustic trigger box code

```
38 // Serial stream buffer
39 int newVals1 = 0;
40 int newVals2 = 0;
41 int checkVal = 0;
42 int serReq = 0;
43
44 // "continuous" mode
45 int contMode = 0;
46
47 // should the sequencer be advanced?
48 int stepSeq = 0;
49
50 unsigned long tNow = 0;
51
52 // Only trigger Q-switch/DAQ on every Nth loop
53 int decimation = 1;
54 int loopDecimation = 0;
55
56 void setup() {
57     // initialize pins
58     pinMode(extTrGnd, OUTPUT);
59     digitalWrite(extTrGnd, LOW);
60
61     pinMode(switchPow, OUTPUT);
62     digitalWrite(switchPow, HIGH);
63
64     pinMode(buttonGnd, OUTPUT);
65     digitalWrite(buttonGnd, LOW);
66
67     pinMode(fireOut, OUTPUT);
68     digitalWrite(fireOut, HIGH);
69
70     pinMode(qsOut, OUTPUT);
71     digitalWrite(qsOut, HIGH);
72
73     pinMode(daqOut, OUTPUT);
74     digitalWrite(daqOut, LOW);
75
76     pinMode(seqOut, OUTPUT);
77     digitalWrite(seqOut, LOW);
78
79     pinMode(buttonIn, INPUT_PULLUP);
80     pinMode(extTrig, INPUT);
81     pinMode(switchIn, INPUT);
82
83     pinMode(ledOut, OUTPUT);
84     digitalWrite(ledOut, HIGH);
85
86     // Initialize serial connection
```

Appendix D. Photoacoustic trigger box code

```
87     Serial.begin(9600);
88     Serial.setTimeout(500);
89
90 }
91
92 void sendVals() {
93     Serial.write((serReq << 7) |
94                 (stepSeq << 6) |
95                 (toggle << 5) |
96                 (contMode << 4) |
97                 nFrames);
98     Serial.write(qsDelay - 200 + 128); // offset to fit range into one byte
99 }
100
101 void serialEvent() {
102     // Serial read/write loop. Timeout is minimized,
103     // but this will still delay the main loop slightly
104     while (Serial.available() > 1) {
105         newVals1 = Serial.read(); // first byte
106         newVals2 = Serial.read(); // second byte
107
108         // Decode
109         serReq = ((newVals1 >> 7) & 0x01);
110         if (!serReq) {
111             stepSeq = ((newVals1 >> 6) & 0x01);
112             toggle = ((newVals1 >> 5) & 0x01);
113             contMode = ((newVals1 >> 4) & 0x01);
114
115             checkVal = (newVals1 & 0x0F);
116             if (checkVal > 0) {
117                 nFrames = checkVal;
118             }
119
120             // Offset to fit range into one byte
121             qsDelay = newVals2 + 200 - 128;
122
123             if (DEBUG_ON) { // for debugging
124                 Serial.println(stepSeq);
125                 Serial.println(toggle);
126                 Serial.println(contMode);
127                 Serial.println(nFrames);
128                 Serial.println(qsDelay);
129             }
130         }
131         // Send values back to confirm receipt
132         sendVals();
133
134         // reset buffer
135         checkVal = 0;
```

Appendix D. Photoacoustic trigger box code

```
136     newVals1 = 0;
137     newVals2 = 0;
138 }
139 }
140
141 void loop() { // Main (triggering) loop
142     tNow = millis();
143     loopDecimation = (loopDecimation + 1) % decimation;
144
145     if (digitalRead(switchIn) == HIGH) { // If switched on
146
147         if (tNow - tLoopOut > loopDelay) {
148             if (LED == LOW) {
149                 digitalWrite(ledOut, HIGH);
150                 LED = HIGH;
151             }
152
153             // fire lasing pulse
154             digitalWrite(fireOut, LOW);
155             digitalWrite(fireOut, LOW);
156             digitalWrite(fireOut, LOW); // Burn a few clock cycles
157             digitalWrite(fireOut, HIGH);
158
159             // For accurate Q-switch timing,
160             // account for pulse width
161             delayMicroseconds(qsDelay - pWidth);
162
163             // If external trigger was received,
164             // or button was pressed
165             if ((digitalRead(extTrig) == HIGH) ||
166                 (digitalRead(buttonIn) == LOW)) {
167                 toggle = 1;
168             }
169
170             // Has data been requested?
171             if (toggle && (loopDecimation == 0)) {
172                 if ((frameCount < nFrames) || contMode) {
173                     frameCount += 1;
174                     // fire Q-switch, and DAQ triggers
175                     digitalWrite(daqOut, HIGH);
176                     digitalWrite(qsOut, LOW);
177                     digitalWrite(qsOut, LOW);
178                     digitalWrite(qsOut, LOW); // Burn clock cycles
179                     digitalWrite(qsOut, HIGH);
180                     digitalWrite(daqOut, LOW);
181                     stepSeq = 1;
182                 } else if (frameCount == nFrames) {
183                     // one more to flush the buffer
184                     digitalWrite(daqOut, HIGH);
```

Appendix D. Photoacoustic trigger box code

```
185     digitalWrite(daqOut, LOW);
186     frameCount = 0; //reset
187     toggle = 0;
188   }
189 }
190 delay(loopPeriod); // Maintain repetition rate
191 }
192 } else { // If the switch is off, blink LED
193   toggle = 0;
194   frameCount = 0;
195   stepSeq = 0;
196   if (tNow - t0 > blinkInterval) {
197     t0 = tNow;
198     if (LED == LOW) {
199       LED = HIGH;
200     } else {
201       LED = LOW;
202     }
203     digitalWrite(ledOut, LED);
204   }
205   tLoopOut = millis();
206 }
207
208 if (stepSeq) { // step the sequencer
209   digitalWrite(seqOut, HIGH);
210   digitalWrite(seqOut, LOW);
211   stepSeq = 0;
212 }
213 }
```