

NEW IMAGE PROCESSING METHODS FOR
ULTRASOUND MUSCULOSKELETAL APPLICATIONS

A Dissertation

by

XU YANG

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Chair of Committee, Raffaella Righetti
Committee Members, Riccardo Bettati
Steven M. Wright
Tie Liu

Head of Department, Miroslav M. Begovic

May 2018

Major Subject: Electrical Engineering

Copyright 2018 Xu Yang

ABSTRACT

In the past few years, ultrasound (US) imaging modalities have received increasing interest as diagnostic tools for orthopedic applications. The goal for many of these novel ultrasonic methods is to be able to create three-dimensional (3D) bone visualization non-invasively, safely and with high accuracy and spatial resolution. Availability of accurate bone segmentation and 3D reconstruction methods would help correctly interpreting complex bone morphology as well as facilitate quantitative analysis. However, in vivo ultrasound images of bones may have poor quality due to uncontrollable motion, high ultrasonic attenuation and the presence of imaging artifacts, which can affect the quality of the bone segmentation and reconstruction results.

In this study, we investigate the use of novel ultrasonic processing methods that can significantly improve bone visualization, segmentation and 3D reconstruction in ultrasound volumetric data acquired in applications in vivo. Specifically, in this study, we investigate the use of new elastography-based, Doppler-based and statistical shape model-based methods that can be applied to ultrasound bone imaging applications with the overall major goal of obtaining fast yet accurate 3D bone reconstructions. This study is composed to three projects, which all have the potential to significantly contribute to this major goal.

The first project deals with the fast and accurate implementation of correlation-based elastography and poroelastography techniques for real-time assessment of the mechanical properties of musculoskeletal tissues. The rationale behind this project is that,

in the future, elastography-based features can be used to reduce false positives in ultrasonic bone segmentation methods based on the differences between the mechanical properties of soft tissues and the mechanical properties of hard tissues. In this study, a hybrid computation model is designed, implemented and tested to achieve real time performance without compromise in elastographic image quality .

In the second project, a Power Doppler-based signal enhancement method is designed and tested with the intent of increasing the contrast between soft tissue and bone while suppressing the contrast between soft tissue and connective tissue, which is often a cause of false positives in ultrasonic bone segmentation problems. Both in-vitro and in-vivo experiments are performed to statistically analyze the performance of this method.

In the third project, a statistical shape model based bone surface segmentation method is proposed and investigated. This method uses statistical models to determine if a curve detected in a segmented ultrasound image belongs to a bone surface or not. Both in-vitro and in-vivo experiments are performed to statistically analyze the performance of this method.

I conclude this Dissertation with a discussion on possible future work in the field of ultrasound bone imaging and assessment.

DEDICATION

To my mother and father for their encouragement and to my wife for her loving support.

ACKNOWLEDGEMENTS

I would like to thank my committee chair, Dr. Righetti, for her continuous support and guidance throughout my research. I would also like to thank my committee members, Dr. Liu, Dr. Bettati, Dr. Wright, for their advice and feedbacks throughout the course of this research.

I thank our collaborators Dr. Krouskop, Dr. Ferrari and Dr. Tasciotti for their help and contribution to this work. Many thanks and acknowledgements to my former and current colleagues and friends in the Ultrasound Imaging Laboratory at Texas A&M University: Anuj, Biren, Mohan, Sanjay, Shafeeq, Songyuan, Srinath and Sthiti. With them I had fruitful discussions and a productive collaboration, which greatly inspired me and positively affected my research. I would also like to thank the staff in the Dept. of Electrical and Computer Engineering at Texas A&M University, who helped me with the administrative requirements for my PhD.

The expertise, efforts and support of all these people have made my doctoral experience at Texas A&M University perfect.

CONTRIBUTORS AND FUNDING SOURCES

Contributors

This work was supported by a dissertation committee consisting of Professor Raffaella Righetti, Professor Steven M. Wright and Professor Tie Liu of the Department of Electrical and Computer Engineering and Professor Riccardo Bettati of the Department of Computer Science. All work conducted for the dissertation was completed by the student independently.

Funding Sources

This work was partially supported by funding from the Department of Defense under Grant Number W81XWH-14-1-0600 and W81XWH-15-1-0718.

Its contents are solely the responsibility of the authors and do not necessarily represent the official views of the Department of Defense.

TABLE OF CONTENTS

	Page
ABSTRACT	ii
DEDICATION	iv
ACKNOWLEDGEMENTS	v
CONTRIBUTORS AND FUNDING SOURCES.....	vi
TABLE OF CONTENTS	vii
LIST OF FIGURES.....	ix
LIST OF TABLES	xii
1 INTRODUCTION AND LITERATURE REVIEW	1
2 A HYBRID CPU-GPGPU APPROACH FOR REAL-TIME ELASTOGRAPHY	7
Introduction	7
Method.....	10
Result.....	25
Discussion	37
Conclusion.....	42
3 BONE SURFACE ENHANCEMENT IN ULTRASOUND IMAGING USING A NEW DOPPLER-BASED ACQUISITION/PROCESSING METHOD	43
Introduction	43
Method.....	47
Results	55
Discussion	63
Conclusion.....	67
4 BONE SURFACE SEGMENTATION FOR FREE HAND ULTRASOUND IMAGE BASED ON STATISTICAL SHAPE MODEL.....	68
Introduction	68
Method.....	70
Results	79
Discussion	83
Conclusion.....	86

5 CONCLUSIONS	87
Summary	87
Future work	88
REFERENCES	91

LIST OF FIGURES

	Page
Figure 1 Architecture of G92 GPGPU (adapted from NVIDIA.com).	14
Figure 2 Time sequences of hybrid thread code and single thread code. (a) Hybrid approach showing how CPU and GPGPU cooperate to minimize computational time. (b) Time sequence of single thread code.	18
Figure 3 Time profile of hybrid version of the elastography algorithm detailing the time spent executing each function in the algorithm	25
Figure 4 (a) Time cost of each elastographic frame computed using the hybrid implementation as a function of the cross-correlation window length for different sampling frequencies. (b) Corresponding time cost of each elastographic frame computed using the CPU implementation as a function of the cross-correlation window length for different sampling frequencies. (c) Estimated GFLOPS for the hybrid implementation as a function of the cross-correlation window length for different sampling frequencies. The RF data size is 1039*128 for 20 MHz sample frequency, 2078*128 for the 40 MHz sample frequency and 4156*128 for 80 MHz sample frequency.	29
Figure 5 (a) Time cost of each elastographic frame computed using the hybrid implementation as a function of the cross-correlation window length for different number of A-lines. (b) Corresponding time cost of each elastographic frame computed using the CPU implementation as a function of the cross-correlation window length for different number of A-lines. (c) Estimated GFLOPS for the hybrid implementation as a function of the cross-correlation window length for different number of A-lines. The RF data size in the axial direction is 2078.	30
Figure 6 (a-c) Examples of simulated ideal strain images (from FEM) obtained by simulating an elastic medium containing a cylindrical inclusion at different applied strain levels. (b-d) Corresponding simulated axial strain elastograms obtained using the hybrid version of the algorithm. (e-g) Corresponding simulated axial strain elastograms obtained using the CPU version of the algorithm. For the elastograms, the cross-correlation window length was fixed at 2 mm. No averaging was used.	31
Figure 7 (a-c) Examples of simulated axial strain elastograms obtained using the hybrid version of the algorithm as applied to the simulation model containing a cylindrical inclusion at a 2% applied strain. (d-f) Corresponding simulated axial strain elastograms obtained using the CPU version of the algorithm. Data were processed using cross-correlation	

window lengths ranging from 1 mm to 3 mm. No averaging was used. W stands for window size in mm.	32
Figure 8 Results of SNRe study: (a) CPU results; (b) GPGPU results; and (c) NRMS error results.	33
Figure 9 Results of CNRe study: (a) CPU results; (b) GPGPU results; and (c) NRMS error results.	34
Figure 10 Example of experimental axial strain elastograms obtained using the Hybrid.AxialOnly version of the algorithm (b and e). For comparison, the corresponding B-mode images (a and d) and CPU axial strain elastograms (c and f) are also shown. Data refer to a gelatin phantom containing a cylindrical inclusion approximately two times stiffer than the background. Each elastogram is an average of 10 elastograms. Images (a)-(c) refer to data acquired with the transducer perpendicular to the axis of the cylinder. Images (d)-(f) refer to data acquired with the transducer parallel to the axis of the cylinder.	36
Figure 11 Example of experimental axial strain elastogram obtained using the hybrid approach (b). For comparison, the corresponding B-mode image (a) is also shown. Data refer to a meat sample containing a HIFU-induced lesion (circled area in the elastograms) scanned in free-mode acquisition. No averaging was applied.	37
Figure 12 Example of experimental axial strain elastogram obtained using the hybrid approach (b). For comparison, the corresponding B-mode image (a) is also shown. Data refer to a liver sample containing a HIFU-induced lesion (indicated by the arrow) scanned in free-mode acquisition. No averaging was applied.	37
Figure 13 Simplified schematic of a bi-layer tissue compressed from the top using an ultrasound probe (with a compressor attached to it) at a constant speed v . (a) Model before compression: the transducer is moving with speed v towards the bone surface. (b) Model after compression: both the thicknesses of the soft tissue layer and bone layer reduce, but the reduction in the thickness of the bone layer is negligible compared to the one of the soft tissue layer.	48
Figure 14. Signal processing procedure to obtain enhanced bone images using the proposed Doppler-based technique.	50
Figure 15 A picture of the setup used for the experiments.	52
Figure 16 Simulated ultrasound signal from one A-line, used to illustrate how the SNR is computed in our application.	55

Figure 17 Power Doppler images created using different surface-transducer relative speeds and the same scale: (a) 1cm/s (b) 2cm/s (c) 3cm/s.	56
Figure 18 Selected in-vitro experimental results. Top panel refers to a chicken sample, while bottom panel refers to a sheep leg. In both cases both the B-mode image and the corresponding Power Doppler image are shown.....	57
Figure 19 Selected results obtained from an intact adult sheep tibia in vivo.	58
Figure 20 Selected results obtained from an intact adult sheep tibia in vivo.	59
Figure 21 Images from adult sheep II tibia with fracture in vivo.	60
Figure 22 Comparison of normalized B-mode and Doppler signals obtained from columns of the image shown in Figures 19a (dotted lines) and 19b (solid lines).	61
Figure 23 SNR Results: left column shows average contrast at 80.5 (+56.5,-28.5) for regular B-mode data; right column shows average contrast at 1172.9 (+1397.0,-191.5) for Power Doppler data.	62
Figure 24 Steps for image preprocessing.	73
Figure 25 Illustration of complete segment steps.....	74
Figure 26 Segmentation results of ex-vivo experiments.....	79
Figure 27 Different bone surface segmented using our algorithm.....	81

LIST OF TABLES

	Page
Table 1 Profile of main operations occurring in the hybrid algorithm to compute the elastogram. In the table, GST stands for global store and GLD stands for global load. The GST/GLD efficiency is the proportion of coalesced memory operations with respect to the total number of operations.....	26
Table 2 Algorithm detection rate for the two studies	82

1 INTRODUCTION AND LITERATURE REVIEW

Clinical diagnosis and assessment of bone abnormalities is achieved through the use of imaging modalities. Recently, three-dimensional (3D) imaging methods have been proposed to improve diagnosis of skeletal pathologies and the assessment of treatments and to provide anatomical insight before or during orthopedic procedures. Computerized Tomography (CT) is currently the gold standard for 3D imaging of skeletal tissues. CT methods provide superior sensitivity, specificity and spatial resolution in the detection of bone abnormalities. However, CT methods require the use of ionizing radiations as well as large and expensive equipment.

In the recent past, ultrasound imaging modalities have received considerable interest in bone imaging applications. Ultrasound imaging modalities have been proposed and tested for the detection of bone defects, as intra-operative aiding tools in orthopedic procedures and to non-invasively monitor bone healing in regenerative medicine applications (Ricciardi, Perissinotto et al. 1993, Allen and Wilson 1999, Weiss, Jacobson et al. 2005) Ultrasound imaging techniques have some unique properties that could provide several advantages with respect to traditional bone imaging means. These include: portability, safety (no radiation), non-invasiveness, cost-effectiveness, real-time feedback and good spatial resolution and penetration depth. Recent developments in the field of magnetic position tracking or optical position tracking also allow accurate recording of the position of the ultrasound probe to facilitate 3D imaging (Barratt, Davies et al. 2001, Franz, Marz et al. 2013, Giannetti, Petrella et al. 2017). Additionally,

some 3D ultrasound probes are now commercially available. These recent developments combined with the availability of new image and processing methods to extract bone surface information from ultrasound volumetric data are now making ultrasound imaging a more usable and reliable modality for 3D bone imaging. Long bone assessment using ultrasound remains a prominent research area as the most commonly fractured bones are the distal radius, humerus, forearm shaft and tibial shaft (Cheng and Shen 1993). While most of the interest is in long bones, other ultrasound imaging applications such as on the knee, the rotator cuff and the spine are currently being investigated (de Jesus, Parker et al. 2009, Teefey, Petersen et al. 2009, Yu, Tan et al. 2014, Slane, Slane et al. 2017) . One potential and important application of ultrasound imaging methods is for the monitoring of bone healing processes. In such application, a patient may necessitates of multiple examinations during bone growth, which would be difficult to achieve using CT or Magnetic Resonance Imaging (MRI) methods. This is particularly critical in the case that the patient is a child or if the bone healing process involves the use of regenerative medicine devices (Protopappas, Vavva et al. 2008, Fayaz, Giannoudi et al. 2011). Another important application of ultrasound imaging methods is as aiding tool in orthopedic surgeries. Since ultrasound systems are portable and provide real-time feedback, they are particularly suitable in intra-operatively scenarios. Additionally, due to the high spatial resolution, ultrasound imaging methods have shown to have comparable performance than CT and MRI methods in finding small tendon tears, tiny calcifications and foreign body on a bone surface . Thus, they have

been proposed as an enhanced early bone pathology detection tool for surgeons (de Jesus, Parker et al. 2009, Chen, Lin et al. 2014, Wang, Wang et al. 2016).

The majority of the work retrievable in the literature on ultrasound bone imaging refers to the use of quantitative low-frequency ultrasound (QUS) methods to assess bone density or detect bone abnormalities (Hijazy, Al-Smoudi et al. 2006, Laugier and Haët 2011). By measuring low frequency (0.1-1 MHz) ultrasound absorption and travel time through the bone structure using a specifically designed probe, QUS methods have been used to quantitatively estimate bone density and composition (Gluer and Barkmann 2003). QUS methods have been found very useful for diagnosis of osteoporosis and bone fragility (Heaney, Avioli et al. 1989, Gnudi, Malavolta et al. 1996) and for the assessment of bone metabolism (Guglielmi G 2010). More recently, high frequency ultrasound methods have been extensively investigated for bone imaging (Hagiwara, Saijo et al. 2011, Giannetti, Petrella et al. 2017). However, while 3D ultrasound imaging methods are widely deployed for surface structural and functional imaging in several medical areas, most notably, perhaps, being obstetrics and cardiology, their utility for skeletal imaging applications remains largely unexplored. Only a few studies have been reported to date that investigate the use of 3D ultrasound techniques for bone applications (Barratt, Penney et al. 2006, Hacıhaliloglu, Abugharbieh et al. 2008).

3D imaging of bones using ultrasound require the design and implementation of accurate and robust bone surface segmentation and 3D reconstruction methods. Bones typically manifest as a hyper-echoic regions in ultrasound images. However, after envelope detection and dynamic range compression, the echogenicity contrast between

bone and soft tissue may no longer be significant. Additionally, ultrasonic attenuation and imaging artifacts can make accurate detection of bone surfaces in ultrasound images very challenging. It is now evident that acoustic intensity alone is insufficient to differentiate between bone surface and other soft tissue structures (Alfiansyah, Streichenberger et al. 2006, Foroughi, Boctor* et al. 2007, Hacihaliloglu, Abugharbieh et al. 2008) and that more imaging features need to be used in order to identify a bone surface in an ultrasound image. One of the features that has been proposed is the shadow region below the bone surface (Foroughi, Boctor* et al. 2007, Shajudeen and Righetti 2017). In some studies, a priori information about the bone shape has also been used. For example, Thomas et al. have discussed the possibility to segment a fetus femur using purely morphological operations, which is ideal for feature extraction if prior knowledge on the shape of the bone anatomy is established (Thomas, Peters et al. 1991). Shape information has also been exploited to utilize active contours for the crux of the segmentation process (He and Zheng 2001, Alfiansyah, Streichenberger et al. 2006). I. Hacihaliloglu et al. proposed ultrasound bone surface localization and registration methods using local phase tensor and statistical shape models (Hacihaliloglu, Abugharbieh et al. 2009, Hacihaliloglu, Abugharbieh et al. 2011, Hacihaliloglu, Rasoulia et al. 2014) .

Ultrasound elastography is a non-invasive technique that has been used to estimate the mechanical behavior of soft tissues (Ophir 1991). With the recent introduction on the market of ultrasound scanners with elastographic capabilities, elastography is becoming a widely accepted cost-effective tool for the diagnosis of a number of pathologies (Garra, Cespedes et al. 1997, Schaar, Korte et al. 2003, Souchon,

Rouviere et al. 2003). Several types of elastography techniques have been proposed, and the main ones used in the clinical practice are: quasi-static elastography, shear wave elastography (Sebag, Vaillant-Lombard et al. 2010), and acoustic radiation force impulse (ARFI) elastography(Cho, Lee et al. 2010). Our laboratory has been actively involved in the development of quasi-static elastography techniques, including novel poroelastography methods. The goal of quasi-static elastography techniques is to generate maps of the local strains that a tissue experiences due to the application of a mechanical stimulus. The resulting strain images are called “elastograms”. More specifically, an axial strain elastogram is an image of the strain tensor component along the beam axis. Other types of elastograms can also be generated by combining different strain components (Konofagou and Ophir 1998, Righetti, Ophir et al. 2005, Konofagou, Harrigan et al. 22001). Specifically, axial shear strain elastography (ASSE) (Konofagou and Ophir 1998, Konofagou et al 2000, Thitaikumar et al 2007a) is a recent development in this field that has been used to assess the degree of connectedness between different tissues. Recent work from our laboratory has demonstrated that ASSE may be used as a monitoring tool for bone imaging applications.

While elastography is primarily used for the assessment of soft tissues, recent studies have suggested the possibility of incorporating elastographic information to improve bone detection and assessment and to better differentiate between soft tissue and bone tissue based on their distinctive mechanical properties.

In this study, I investigate and propose new methods to improve bone surface enhancement and segmentation in ultrasound images, with particularly emphasis on techniques that could be readily implementable in clinical in vivo applications necessitating free-hand acquisition settings and real-time feedback.

In Chapter 2, I propose, design and test new methods to perform elastographic imaging of musculoskeletal structures in real-time but without the loss of image quality. These

methods might be used in the future to select new features to be incorporated in a bone surface detection algorithm for clinical applications. In Chapter 3, I propose, design and test a new Doppler-based bone surface enhancement method that might be used in the future to facilitate bone segmentation problems. This method gives a way to enhance signal from bone surface while suppress signal from soft tissue and connective tissue. It can be integrated into existing US Doppler module easily. In Chapter 4, I propose, design and test a statistical shape method to identify bone surfaces in ultrasound images. The advantages of this method include low computation cost for training and classification with high segment accuracy.

2 A HYBRID CPU-GPGPU APPROACH FOR REAL-TIME ELASTOGRAPHY¹

INTRODUCTION

Ultrasound (US) elastography is a non-invasive imaging modality that maps local strains to color-coded tissue information (Ophir 1991). With the recent introduction on the market of US scanners with elastographic capabilities, elastography is becoming a widely accepted cost-effective tool for the diagnosis of a number of pathologies (Garra, Cespedes et al. 1997, Schaar, Korte et al. 2003, Souchon, Rouviere et al. 2003). In clinical applications, the availability of high quality elastograms combined with real-time tissue visualization is usually a fundamental requirement for correct and early diagnosis. Yet, achievement of both these requirements in vivo simultaneously is a difficult task, which depends on a number of parameters related to the US system, the processing methods and noise (Varghese, Ophir et al. 2003).

In the past years, several real-time elastography methods have been proposed. These include: low computationally intensive estimation methods such as sum of square difference (SSD)(Hall, Zhu et al. 2003), spectral strain estimation (Konofagou, Varghese et al. September 1999), phase root seeking (Lindop, Treece et al. 2006), zero-crossing (Srinivasan and Ophir 2003) and sample tracking (Zahiri-Azar and Salcudean 2008) and cross-correlation with prior estimates (Zahiri-Azar and Salcudean 2006). These methods

¹ Reprinted with permission from "A hybrid CPU-GPGPU approach for real-time elastography" by Xu Yang, Sthiti Deka and Raffaella Righetti, IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control, 2011 Dec 58(12):2631-45

achieve real-time frame rates, but usually their improvement in computational speed is accompanied with losses in the quality of the resulting elastograms. Conventional cross-correlation elastography algorithms usually provide high quality images (Hoyt, Forsberg et al. 2006), but they are computationally intensive and not suitable for real-time applications.

The problem of achieving high speed computational powers while preserving image quality in clinical environments extends well beyond elastography applications. In fact, this capability not only may allow fast and improved diagnosis but it may be essential to monitor treatments in real-time for individualized therapies. Recent developments in the biomedical imaging field have suggested that this issue might be solved using a hardware/software approach (Luebke 2008). In this connection, General Purpose Graphics Processing Unit (GPGPU) has become an attractive platform for bioscience and bioengineering problems of this kind (Owens, Luebke et al. 2007). Recently, GPGPU has been used in several medical imaging application (Xu and Mueller 2005, Samant, Xia et al. 2008, Jia, Lou et al. 2010) and in ultrasonic imaging applications that involve computationally intensive algorithms, such as real-time 3D reconstruction (Dai 2010), real-time 2D temperature imaging (Liu 2010), and real-time displacement estimation (Rosenzweig 2011). Our group has reported the first work retrievable in the literature that suggests that the use of GPGPU may be advantageous in ultrasound elastography applications (Deka, Yang et al. 2009).

The tremendous computational power of GPGPU is supported by the SIMT (Single Instruction Multiple Thread) architecture, which was originally intended for 3D

rendering applications. In addition, fast context switching and high memory bandwidth also contribute to GPGPU's high performance. To fully exploit the high computational horsepower of GPGPU, an implementation has to follow a number of restrictions to maximize device utilization, memory throughput and instruction throughput. In general, GPGPU makes computation faster in most applications and, equally importantly, frees CPU for other high priority computations. Therefore areas ranging from complex simulations to medical imaging are actively investigating the use of GPGPU to accelerate complex algorithms that could easily saturate the CPU (Owens, Luebke et al. 2007).

While in the past graphics hardware had to be programmed through multiple software layers and complex low level languages, the availability of high level programming languages make the GPGPU readily accessible to developers for general purpose computation. Current popular programming models such as NVIDIA's Compute Unified Device Architecture (CUDA)(NVIDIA 2010) and OpenCL(ATI 2010) are designed to abstract the hardware and offer language APIs that are simple extensions to the familiar C99 standard. For example, CUDA enables application software to transparently scale its parallelism to leverage the increasing number of processor cores. A compiled CUDA program can be executed on any available NVIDIA processor cores because the runtime system will directly schedule the parallelism leaving programmers the only task of parallelizing the algorithm (NVIDIA 2010).

In this paper, we investigate the use of GPGPU hardware accelerators for US elastography applications. We provide a comprehensive analysis of the performance,

potentials and fundamental limitations of GPGPU methods as applied to classical elastography estimations problems. The driving motivation at the basis of this study is to tackle the speed-quality orthogonal problem in elastography by employing a new approach, which is cost effective, efficient and allows real-time performance with no loss in image quality.

METHOD

Elastography image formation process

The goal of US elastography techniques is to generate maps of the local strains that a tissue experiences due to the application of a mechanical stimulus (Ophir 1991). The resulting strain images are called “elastograms”. More specifically, an axial strain elastogram is an image of the strain tensor component along the beam axis, while a lateral strain elastogram is an image of the strain tensor component orthogonal to the beam axis within the scanning plane. Other types of elastograms can also be generated by combining different strain components (Konofagou and Ophir 1998, Righetti, Ophir et al. 2005, Konofagou, Harrigan et al. 2001). These local strains (or combination thereof) are related to the underlying tissue mechanical properties (Srinivasan, Krouskop et al. 2004).

Below, we provide a detailed explanation of the elastographic image formation process, since it is fundamental to understand the computational issues related to elastographic estimations. While elastography is fundamentally a three dimensional problem, in this study we focus primarily on axial displacement and strain estimation.

However, the concepts covered in this paper can be easily extended to lateral and elevational displacement and strain estimations as well (Konofagou and Ophir 1998). To generate an elastogram, a set of digitized RF echo-lines (pre-compression frame) is acquired from the region of interest inside the tissue. The tissue is then compressed by a small amount (typically 0.1-1% strain) usually along the ultrasonic radiation axis and a second set of digitized RF echo-lines (post-compression frame) is acquired from the same region of interest inside the tissue. Congruent echo-lines are globally stretched to correct for echo-decorrelation noise (Varghese and Ophir 1996) and divided into temporal windows. Time delay estimation (TDE) techniques are applied to the windowed signals, to obtain a local displacement estimate (Cespedes, Huand et al. 1995). The temporal windows are then shifted along the radiation axis and the computation is repeated for all depths. This process allows generation of a displacement map. The strain image is computed as the gradient of the local displacements (after median filtering to remove salt-and-pepper noise).

Each pixel in the elastogram represents the local value of the strain experienced by the tissue as a result of the compression. Local displacements may be obtained using TDE techniques, and several of these methods can be found in the literature. While each of the proposed methods has his own advantages and limitations, cross-correlation based methods are known to provide the highest quality estimates albeit at the expense of computational time (Hoyt, Forsberg et al. 2006). A recent study conducted in our laboratory demonstrates that cross-correlation methods provide the highest image quality

when compared to other algorithms based on minimization of cost functions (Sambasubramanian August 2010).

In this study, we implemented a generalized cross-correlation algorithm in C, using the following steps: 1) temporally stretch the post-compression echo signals; 2) divide data into temporal windows; 3) cross-correlate corresponding windowed pre- and post-compression echo-signals; 4) estimate displacement from the position of the peak of cross-correlation function; 5) repeat steps 3 and 4 for the entire depth of the echo-line and for every echo-line; 6) median filter displacements; 7) estimate strains. In our implementation, the cross-correlation is performed in the frequency domain, where it is computed as the inverse Fourier transform of the product of the FFT of the pre-compression window and the conjugation of the FFT of the post-compression window. While cross-correlation can be performed either in the time-domain or in the frequency domain, FFT implementation takes logarithmic time complexity, while it would take quadratic time complexity ($O(n^2)$) to perform cross-correlation in the time domain. The position of the peak of the cross-correlation function yields the time-delay estimate. From the displacement maps, the strain elastograms are computed using a staggered strain estimation technique, which is a multi step strain estimation method that provides elastograms of improved quality with respect to those obtained using a simple gradient method (Srinivasan, Ophir et al. 2002).

The GPGPU architecture

To investigate the feasibility of using GPGPU for elastography applications, we implemented the elastographic estimator described in the previous section on the Nvidia

GeForce 8800 GT. The Nvidia GeForce 8800 GT is a graphics card equipped with a G92 graphics processing unit (GPU) based on Tesla unified graphics and computing architecture. Tesla architecture can support NVIDIA's CUDA, which allows developers to easily implement data-parallel algorithms based on SIMT model. In the following paragraphs, we discuss the features of CUDA and G92 that are most relevant for the problem of accelerating ultrasonic elastography.

Figure 1 shows a schematic of the architecture of the G92 GPGPU. The G92 GPGPU has 14 streaming multiprocessors (SMs), each containing 8 streaming processors (SPs) running at 1.5GHz. The SMs use SIMT unit to organize threads into groups of 32 parallel threads (warps). In addition, each SM has two special functional units shared among 8 SPs to perform reciprocal, square root, sine and cosine with low latency. Each SP has one arithmetic unit that performs single precision floating point operations and 64-bit integer operations. For each clock cycle, a SP can perform one multiply-add operation and a SFU can perform one complex floating point operation. The combination of the 8 arithmetic units and the 2 SFUs allow 18 floating point operations per clock cycle for each SM, and the 14 SMs carry out 378 GFLOPS peak theoretical performance for the G92 ($14 \text{ SM} * 18 \text{ FLOP/SM} * 1.5 \text{ GHz}$).

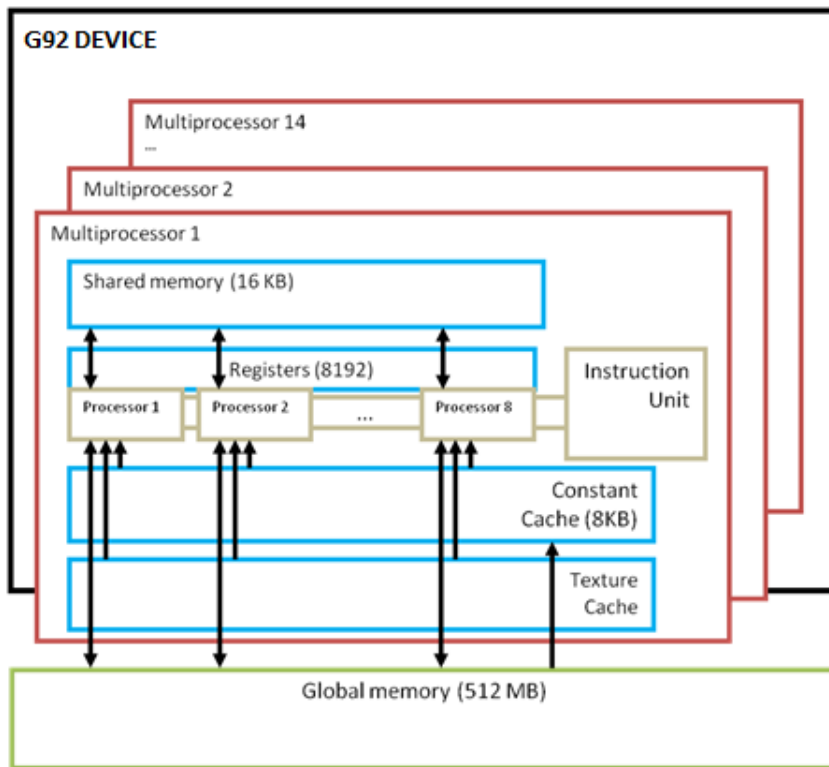


Figure 1 Architecture of G92 GPGPU (adapted from NVIDIA.com).

There are two kinds of memory on graphics card: the on-chip memory and the off-chip memory. The on-chip memory is integrated in the GPU chip and has low latency compared to the off-chip memory, which is loaded on the graphics card as main data storage. The GeForce 8800 GT has 57.6 GB/s of bandwidth on its 512 MB global memory. Although the available bandwidth is high, off-chip global memory suffers from high-latency. Besides, the global memory has no cache, so the theoretical bandwidth can only be achieved under strict access patterns. Therefore as all GPGPUs based on Tesla architecture, G92 has four types of on-chip memory designed to reduce demands for global memory access. The fastest on-chip memory is the register in SMs. For G92, each

SM has 8192 shared registers. Additionally, each SM has 8kB constant memory cache that speeds up read from 64 kB off-chip constant memory, and 16kB on-chip shared memory that is optimized for parallel data accessing. Finally, Tesla architecture provides the off-chip texture memory and the on-chip texture caches optimized for 2D spatial locality.

Elastography on GPGPU

To study the usefulness of the GeForce 8800 GT's architectural features in elastography applications, we implemented a hybrid CPU-GPGPU version of the conventional elastographic estimator detailed in section A. The hybrid implementation was achieved through iterative developments, during which four major versions of the algorithm were generated. It should be noted that, apart from parallelizing the algorithm to improve efficiency, the major issues that we need to address when implementing elastography on GPGPU are related to memory and control flow problems. Memory-related issues are present in all steps of the algorithm (see section A), from transferring and storing RF data to saving and accessing FFT results and cross-correlation values for the estimation of displacements. Control flow problems are present mainly during operations such as summation and sorting cross-correlation values to find the peak of the cross-correlation.

The various iterations of the algorithm implemented in this study differ from each other with respect to the optimization level used for the memory bandwidth and GPGPU occupancy. The final hybrid version allows CPU and GPGPU to work simultaneously to overlap waiting time when processing continuous windows. Below, we provide a brief description of the technical characteristics of each iteration implemented in this study.

GPGPU.Base

The GPGPU.Base is a direct implant from CPU implementation with no optimization methods employed to conserve memory bandwidth or tolerate long latency loads. This simple base version executes data on the GPGPU in a parallel fashion, which

results in frequent data transfer between host memory and graphics memory. In this version, only the cross-correlation computations are performed on the GPGPU. For the FFT computation, CUFFT library is used.

GPGPU.H2DMem

In the second version, GPGPU.H2DMem, the transferring of RF data from host memory to device global memory is optimized in several ways. First, the RF data is stored, A-line by A-line (instead of row by row as in the base version). This modification improves cache misses when dividing A-line into windows. It also makes coalesced memory access possible for executing computations on the GPGPU. Second, coalesced memory transaction is used to make the use of global memory bandwidth more efficiently. Data should be aligned and accessed in sequence by threads, so that memory transactions of threads in a half-warp can be coalesced into one transaction. Third, RF data is saved in page-locked memory before transferring. Page-locked memory has higher bandwidth than regular pageable host memory.

GPGPU.On-ChipMem

In the version, GPGPU.On-ChipMem, data for displacement and cross-correlation estimation are loaded into on-chip shared memory to reduce device memory transaction. While in the previous versions, to estimate the displacements threads need to access cross-correlation functions saved in the device global memory repeatedly, in this version cross-correlation functions are loaded into the on-chip shared memory before actual computation. This allows avoiding the high latency of the device global memory transaction. Shared memory is divided into banks, and any memory request of n addresses that fall into n distinct memory banks can be combined into one transaction. If two addresses fall into one bank, a bank conflict happens and the request is serialized. Although additional synchronization operations are needed for the loading process, this step allows saving computational time because the on-chip shared memory is very fast (32 bits per two clock cycle) when there is no bank conflict, compared to 400 to 800 cycles for the off-chip memory. Furthermore, in this version, all computations except the

median filtering and the strain estimation are executed on the GPGPU to reduce frequent data exchange between host memory and graphics memory.

GPGPU.ControlFlow

The fourth version, GPGPU.ControlFlow, additionally improves performance by using parallel reduction (Harris 2007). According to NVIDIA, any flow control instruction (if, switch, while, for) can significantly impact the effective instruction throughput by causing threads of the same warp to diverge (NVIDIA 2010). As a consequence of such divergence, the different execution paths need to be serialized, increasing the total number of instructions executed. Since for the displacement estimation, only the maximum value of the cross-correlation is of importance, there is no need to perform a complete sorting. Therefore a parallel reduction method (Harris 2007) is used to reduce divergence and find the maximum value of the cross-correlation. As a common data parallel primitive, parallel reduction cuts off runtime effectively by using tree-based approach within each block. This can be easily implemented in a loop.

One important technique to implement parallel reduction in CUDA is the unrolling of the last six loops. As reduction proceeds, the number of active threads decreases by half in each step. Due to warp synchronization, all threads within a warp have the same instruction. Therefore when the number of effective threads in a step is equal to or below the warp size, there is no need to have synchronization and judgment of active threads in the step. So unrolling of the last six loops makes the last six step divergence free.

Hybrid multithread

This final version uses Consumer/Producer Pattern to create a hybrid computation model, which employs both CPU and GPGPU. In this version, the median filtering and the strain computation are performed on the CPU. Execution of these operations on the GPGPU is deliberately avoided because they involve uncoalesced memory accesses, bank conflict and control flow that cap the performance. In addition, the CPU computational time required by the execution of these two steps is so small that it can be overlapped with the computational time taken by the GPGPU to execute other operations. Therefore, the Consumer/Producer Pattern approach is used to solve the problem. In our approach, we let CPU compute the median filter

and strain of frame n while the GPGPU is calculating the displacement for frame $n+1$. Thus, for continuous frames, the time used to estimate the strain will be completely overlapped with the time used to compute the displacement. Figure 2 shows a schematic of the proposed approach vs. a single thread approach. As it can be observed from this schematic, the hybrid thread code allows an effective reduction of the computational time by overlapping the processing times of the different stages of the algorithm.

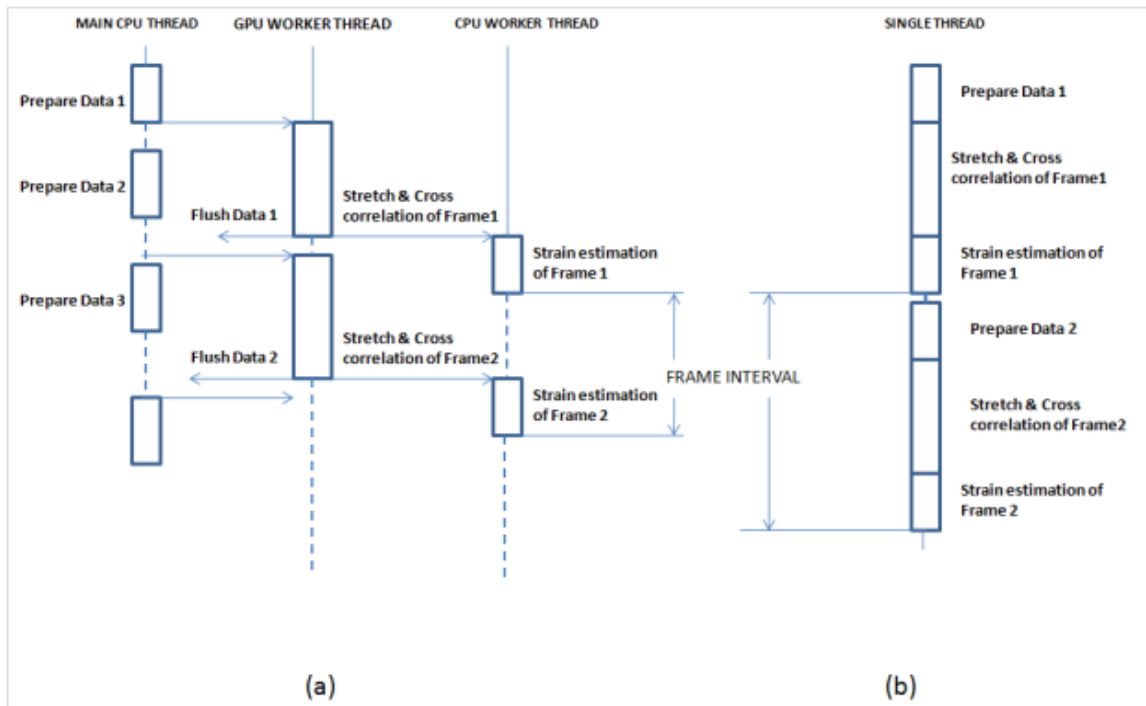


Figure 2 Time sequences of hybrid thread code and single thread code. (a) Hybrid approach showing how CPU and GPGPU cooperate to minimize computational time. (b) Time sequence of single thread code.

Performance evaluation

The CPU-GPGPU hybrid implementation was compiled using speed maximize option (-O2, CUDA version 3.2) and tested on GeForce 8800 GT installed in a system with a 2.2 GHz Core Duo E4500 CPU. Both CPU and GPGPU versions were tested on this platform.

All tests were performed on simulated data. We simulated a 40 by 40 mm² media using a simulation software developed in our lab and previously described (Desai, Krouskop et al. 2010). In all simulated models, the media were compressed from the top by applying a constant axial strain (ranging from 0.01% to 10%). We assumed slip boundary conditions at the bottom and a plane-strain state model. In all models, a scatter density of at least 40 scatters/pulse-width was used, satisfying the requirement for obtaining fully developed Rayleigh backscatters. The speed of sound was fixed at 1540 m/s. The sonographic signal-to-noise ratio (SNR_s) was set at 20 dB.

We simulated an ultrasound system that resembled our experimental ultrasound system. The simulated ultrasound transducer had 128 elements, frequency bandwidth between 5-14 MHz, a 6.6 MHz center frequency and 50% fractional bandwidth at -6 dB. The transducer's beamwidth was assumed to be approximately 1 mm at 6.6 MHz.

To evaluate the effect of data size (both axial and lateral) on the computational speed, two sets of data were simulated. The first set was used to test the impact of axial data size on the computational speed. The second set was used to test the impact of lateral data size on the computational speed. For statistical analysis, 50 independent realizations were simulated for each data configuration.

The number of axial pixels in an elastogram depends on the acquisition depth, window length and overlap between successive windows and can be computed as $\frac{d}{l_w(1-\alpha)}$, where d is acquisition depth, l_w is window length and α is the overlap. Here, d and l_w can be in units of mm or pixel (as long as they have the same unit). The number of lateral pixels in an elastogram is equal to the number of A-lines. For each pixel in the elastogram, the problem size is determined by the number of samples in each window, which is given by $2 \frac{f_s l_w}{c}$, where f_s is the sampling frequency and c is the speed of sound (in this formula l_w has the same unit of measurement of c).

To test the effect of axial data size on the computational speed, three values of sampling frequency were considered: 20 MHz, 40 MHz and 80 MHz while the number of A-lines was set at 128. This resulted in simulated RF frames of size 1039 x 128 pixels², 2078 x 128 pixels² and 4156 x 128 pixels² respectively. To test the effect of

lateral data size on the computational speed, three values of A-lines were considered: 128, 256 and 512 while the sampling frequency was set at 40 MHz. This resulted in simulated RF frames of size 2078×128 pixels², 2078×256 pixels² and 2078×512 pixels² respectively.

Finally, we also investigated the effect of the cross-correlation window size on the computational speed. Window sizes ranging from 0.4 mm to 4.0 mm were used for this study (for a 2% applied strain - note that the applied strain does not affect computational time). For all simulation cases, the window overlap was set at 80% (note that the window overlap has only a linear effect on computational speed).

Evaluation of computational speed

For the evaluation of the speed performance we used simulation data corresponding to a uniform medium. The computational speed of the elastography algorithm was assessed by computing the execution time and GFLOPS for RF data sets described in previous section. The execution time was calculated as the average processing time it takes to complete each trial.

To compute the number of floating point operations in each step of the elastographic algorithm, the following procedure was used. Let each RF frame have dimension $L \times N$ pixels², where L is the number of samples in the axial direction and N is the number of samples in the lateral direction. Let K be the number of pixels in each cross-correlation window. Then, the elastogram size is $M \times N$ pixels² (where M can be computed as specified above). The number of floating point operations in each step of the elastographic algorithm is estimated as follows:

- a. The *stretching* operation is performed on the RF data. Stretching is performed using linear interpolation: given two points (x_0, y_0) and (x_1, y_1) , for each $x \in (x_0, x_1)$ the corresponding y is given by $y = y_0 + (x - x_0) \frac{y_1 - y_0}{x_1 - x_0}$. Thus, linear interpolation costs 5 operations for each sample, and stretching costs $5LN$ floating point operations.
- b. The *cross-correlation* cost $MN(15K \log_2(K) + 6K + 4K)$ floating point operations. The number of FFT floating point operations for each pixel in the elastogram is given by $15K \log_2(K)$ because 2 FFT and 1 inverse FFT operations are needed for each pixel and

each operation has roughly $5K \log_2(K)$ operations (according to the radix-2 Cooley-Tukey algorithm)(Cooley and Tukey 1965). Additionally, 6K operations are needed to multiply complex FFT results, and additional 4K operations are needed to calculate the sum of squares in each pair of pre-compression and post-compression windows.

c. *Displacement estimation* uses $MN(K + 11)$ floating point operations. To find the peak of the cross-correlation value in a window, K floating point operations are needed using a parallel reduction method, and 11 more operations are used for sub-sample estimation (cosine interpolation around the peak of the cross-correlation function). Similarly as for point a), the y value here can be calculated as $y = y_0(1 - \mu) + y_1\mu$, where $\mu = (1 - \cos(\frac{x-x_0}{x_1-x_0}\pi))/2$.

d. *Staggered strain estimation* costs $MN(J \log_2(J) + 4)$, where J is the median filter size. Staggered strain calculation is somewhat similar to linear interpolation and roughly costs 4 floating point operations for each pixel in the elastogram. The strain value in strain map is calculated as $s_{i,j} = |\frac{d_{i,j} - d_{i,j+W}}{d_{shift} \times W}|$, where $s_{i,j}$ is strain value at position (i,j) and $d_{i,j}$ is displacement at (i,j) in displacement map, d_{shift} is the window shift and W is staggered size (Srinivasan, Ophir et al. 2002).

It is important to note that the above calculations are not an actual flop counts, because all operations (addition, subtraction, multiplication, division and trigonometric functions) are counted as one operation, as it is often the case in GFLOPS calculations.

In addition to the execution time, for the hybrid approach we also analyzed the percentage of time spent on the different functions of the algorithm during the execution of the elastograms. Using NVIDIA Compute Visual Profiler, we evaluated the profile of warp serialization, the profile of branch and divergent branch and the profile of coalesced and uncoalesced global memory loads.

Image quality assessment

In order to evaluate any eventual loss in resulting image quality, we statistically compared the performance of the CPU/GPGPU hybrid approach implementation and the

CPU implementation of the elastography algorithm. This performance comparison between GPGPU and CPU implementations of the algorithm is necessary because the GPGPU employed for this study is limited by single precision floating point arithmetic, as opposed to the CPU's double precision capacities.

Image quality tests were performed on two simulation models. For the SNR_e tests, we used a uniform simulated medium (elastic modulus = 2 kPa and Poisson's ratio = 0.495). For the CNR_e tests, we used a simulated medium containing a cylindrical inclusion (background: elastic modulus = 2 kPa and Poisson's ratio = 0.49; inclusion: elastic modulus = 10 kPa and Poisson's ratio = 0.495). The mechanical properties of the background and the inclusion materials were chosen based on previous elastography studies (Kallel and al. 2001, Righetti, Ophir et al. 2002).

The SNR_e was computed as $SNR_e = \frac{s}{\sigma_1}$, where s is the mean value of the estimated strain and σ_1 is the standard deviation of the strain (Varghese, Ophir et al. 2003). The CNR_e was computed as $CNR_e = \frac{2(s_t - s_b)^2}{\sigma_t^2 + \sigma_b^2}$, where s_t and s_b are the mean values of the estimated strain in the target and the background respectively, and σ_t and σ_b are the respective standard deviations of the estimated strains (Varghese and Ophir 1998). The SNR_e was measured on simulated uniform cubic phantoms. The CNR_e was measured on simulated cubic phantoms containing a 10 mm diameter cylindrical inclusion. These image quality factors were assessed using a set of 50 independent trials.

To further evaluate any statistical difference on a pixel-by-pixel basis between the CPU implementation and the hybrid implementation, the normalized root mean square error (NRMS) was also calculated. The NRMS error was computed as $e =$

$\frac{1}{(s_{max} - s_{min})} \sqrt{\frac{\sum_{i=1}^M \sum_{j=1}^N (\hat{s}_{ij} - \check{s}_{ij})^2}{MN}}$, where M and N are the axial size and lateral size of the elastograms (as specified earlier), \hat{s}_{ij} is the strain estimation at position (i,j) obtained using the CPU implementation and \check{s}_{ij} is strain estimation at the same position obtained using the hybrid implementation.

Experiments

Experiments were performed to demonstrate the technical feasibility for generating elastograms using the CPU/GPGPU hybrid approach as applied to noisy, real elastographic data. For comparison, the same sets of data were also processed using the CPU implementation of the elastography algorithm.

Experiments on gelatin phantoms

For the first set of experiments, we used gelatin phantoms containing a cylindrical inclusion. Five gelatin phantoms were used for these experiments. The gelatin phantoms were constructed following the procedure indicated in Kallel et al.(Kallel and al. 2001). Briefly, the gelatin blocks were fabricated by mixing 5% by weight of gelatin and 3% by weight of agar to boiling water and following the procedure described in Kallel et al. (2001). After solidification, a cylindrical cavity of diameter equal to 1cm was created at the center of the phantom. The cylindrical inclusion was the filled with a gelatin mix fabricated by mixing 12% by weight of gelatin and 3% by weight of agar. This allowed generation of a phantom with a cylindrical inclusion approximately two times stiffer than the background. It should be noted that similar gelatin phantoms serve well as ultrasonic tissue-mimic phantoms and are commonly used in elastography applications for the purpose of evaluating image quality of elastographic estimators.

Each gelatin sample was then placed on a rubber plate and compressed from the top using an apparatus that allows application of a constant strain to the sample(Desai, Krouskop et al. 2010). The custom-build compression-loading frame held the ultrasonic transducer in contact with the sample being studied. A compressor plate was attached to the transducer to approximate a uniform stress distribution. Before elastographic data collection, each sample was initially pre-compressed in order to assure good contact (the amount of pre-compression depended on the curvature of the muscles). Subsequently a 2% axial strain was applied to the sample.

The samples were scanned using a 38 mm real-time linear array scanner Sonix RP (Ultrasonix, Richmond, BC, Canada) that has 128 elements, a bandwidth between 5-14 MHz, a center frequency of 6.6 MHz, 50% fractional bandwidth at -6 dB, sampling frequency of 40 MHz, and 1 mm beamwidth at the focus. After data acquisition, the RF data are then transferred in an external computer for processing.

Experiments on ex vivo tissues

For the second set of experiments, we used eight samples of *ex vivo* tissues, which included samples of fresh bovine muscle (5) and porcine liver (3) purchased at a local market. The samples were casted into gelatin blocks approximately 50 mm thick and 50 mm wide. Prior elastographic data collection, the samples were exposed to High Intensity Focused Ultrasound (HIFU), which provided a controlled means to locally modify the stiffness properties of the materials and generate stiffer inclusions inside the tissue samples. It should be noted that the use of HIFU to create localized and coagulated areas within tissues has been well documented in the literature (Hill and terHaar 1995, Madersbacher, Kratzik et al. 1995). Righetti et al. (Righetti, Kallel et al. 1999) demonstrated the feasibility of using HIFU to create localized stiff lesions in canine livers. By varying the amount of ultrasound exposure and acoustic power, they were able to create a large range of stiff lesions having different size and elastic modulus. In addition, the field of HIFU treatments represents a possible medical area where the application of real-time elastographic methods could prove clinically relevant.

The in-house HIFU system used for this study comprises a HIFU transducer that has a diameter of 6 cm, a nominal center frequency of 1 MHz and, -6dB beamwidth of 1 mm and a focal length equal to 76 mm. The HIFU acoustic power can reach up to 120 W, which combined with the focal spot size measurements results in a focal intensity of 1582.1 W/cm². For the specific study, one inclusion was induced in each meat sample using a 75 W acoustic power and 30 s exposure time.

Following the HIFU experiments, the sample was cast in a gelatin block, which was used with the sole scope of supporting the phantom during elastographic data collection.

Each sample was then placed on a rubber plate and compressed from the top. A compressor plate was attached to the transducer to approximate a uniform stress distribution. Before elastographic data collection, each sample was initially pre-compressed in order to assure good contact (the amount of pre-compression depended on the curvature of the muscles). Data were acquired while compressing the sample in free-mode configuration. Approximately, a 2% total axial strain was applied to the sample.

RESULT

Speed performance evaluation

Figure 3 shows the percentage of time spent by the hybrid algorithm in the various stages that compose the elastography algorithm.




Figure 3. Time profile of hybrid version of the elastography algorithm detailing the time spent executing each function in the algorithm

Table 1 shows a detailed performance analysis of warp serialization, branch and global memory efficiencies for the hybrid version of the algorithm. In Table 1, we can observe that the global load efficiency is Not Applicable (N/A) for all the functions since all functions take advantage of texture cache during loading data from the global memory. Some functions have branch problems, but divergent branches are largely avoided. In the axial displacement estimation and cross correlation normalization, reduction is used in sorting and summation. This generates a high number of branches. However, the proportion of divergent branches with respect to the total branches in these two functions is low. The product computation (i.e., the product between the FFT of the pre-compression window and the FFT of the post-compression window) has very low

global store efficiency due to the dimension of the output arrays, which are determined by the size of the FFT. Since the FFT length is padded to be 2^n , the length of the FFT result is $2^{n-1}+1$, which is an odd number and therefore causes large uncoalesced memory access. The axial displacement estimation also has low global store efficiency but the size of stored data in this case is small. Being the most time consuming portion of the algorithm, the FFT function has several issues except global loading.

Table 1 Profile of main operations occurring in the hybrid algorithm to compute the elastogram. In the table, GST stands for global stor. The GST efficiency is the proportion of coalesced memory operations with respect to the total number of operations.

	GPU%	Branches	Divergent branches	Warp serializations	GST efficiency	
Data transfer from device to host	0.37%	N/A	N/A	N/A	N/A	
Stretching	0.55%	1616	9	0	1	
Data transfer from host to device	6.56%	N/A	N/A	N/A	N/A	
Estimation of cross correlation peak & axial displacement	18.44%	224460	860	0	0.11	
Zero padding and normalization	20.38%	275040	1719	0	0.99	
Cross-correlation computation	FFT & IFFT	41.14%	61024	1718	44438	0.36
	Product	12.56%	0	0	0	0.012

Figure 4 shows the effect of sampling frequency on computational speed, and figure 5 shows the effect of A-lines on computational speed. As the FFT size is padded to 2^n and the size of the FFT depends on the window length, there is a jump in performance when the padded length changes due to the window length increase in both figures. In elastography, smaller window sizes lead to improved resolutions but lower

SNRe and CNRe (Varghese, Ophir et al. 2003). Computationally, smaller window size means smaller number of axial samples in each window but larger number of windows in the elastograms (since the depth is fixed). The sampling frequency is proportional to number of axial samples in each window. The graph in Figs. 4a and 5a shows that the hybrid method gains more speed for smaller RF data size (lower sample frequencies) and larger window sizes. GFLOPS increases with RF data size and window size (Fig. 4c and 5c). For the hybrid implementation, computational time increases sub-linearly (Figs. 4a and 5a) as the data size increases (laterally and axially) while CPU computation time increases (Figs. 4b and 5b).

From Fig. 4a we observe that the sampling frequency has practically no effect on the time cost for window lengths below 1mm. This is because for window length less than 1mm (corresponding to a FFT length of 64 samples or lower), the number of windows is more relevant than the window length itself and the algorithm spends more time in calling the functions than in the computation itself. Since the depth is fixed, the number of windows changes with the window length but does not change with the sampling frequency. In Fig. 5a, instead, for each lateral data size, the window length is the same but the number of windows is different. The CPU results (Figs. 4b and 5b) show a similar trend as the GPU results (Figs. 4a and 5a) even though in this case the computational time is significantly higher than the corresponding computational time required by the hybrid approach (using t-test at a 95% confidence level). As previously mentioned graphs mentioned it should be noted that, if single precision CPU are used instead of double-precision CPU, computation in CPU is about 1.8 times faster than

double precision computation for FFT which occupies around 90% of the computation time in CPU.

We finally observe that increase in the lateral data size causes a higher number of GFLOPS than the increase in axial data size (Fig. 5c vs. Fig. 4c). This is because our kernel function computes the displacement and correlation peak values for a row of windows each time the function is called. So the data size expansion in the lateral direction causes CUDA to start more threads each time, hence more latency can be hidden (when a batch of threads is waiting for data, GPGPU can switch to work on another batch that is ready for computation). This would indicate that the hybrid approach has the potential to increase GFLOPS and maximize the occupancy of GPU by computing more windows in different rows each time the function is called.

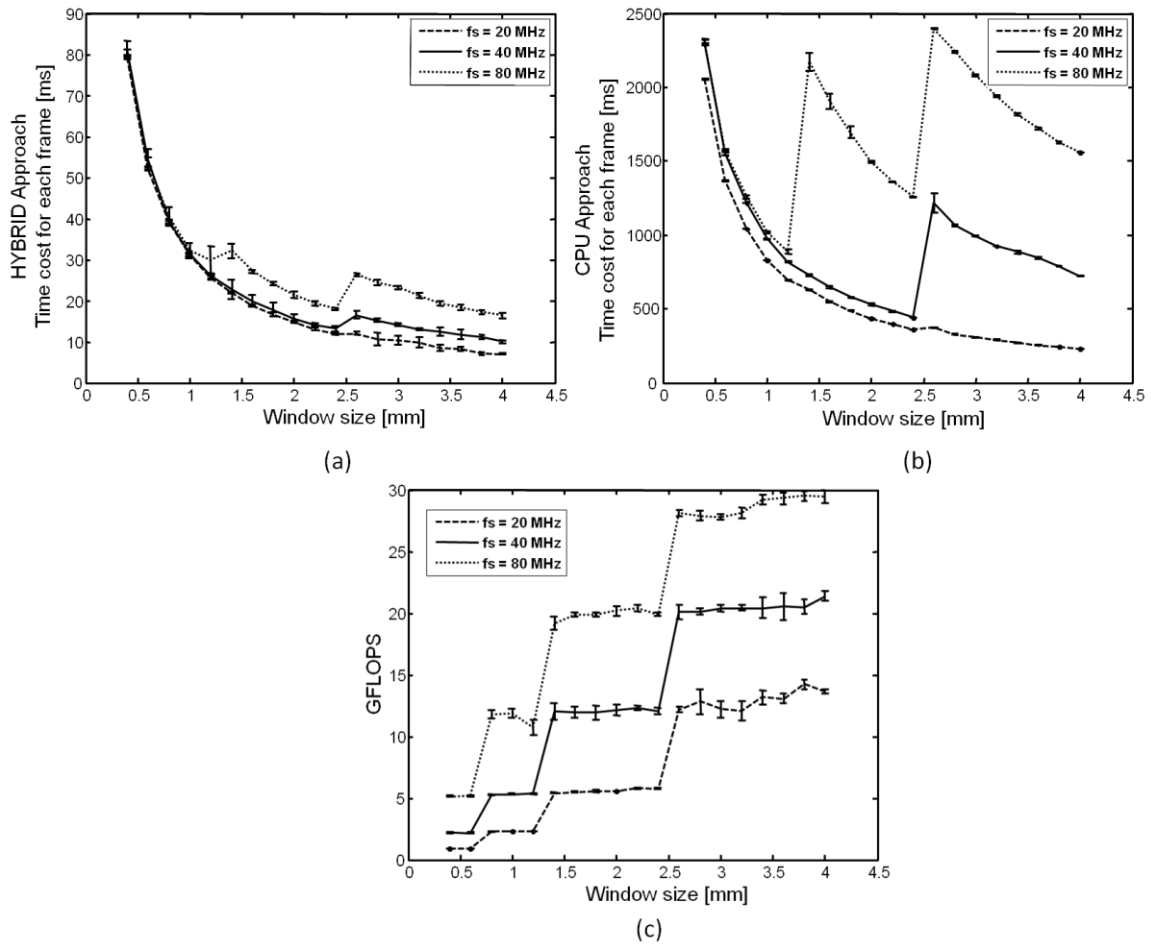


Figure 4 (a) Time cost of each elastographic frame computed using the hybrid implementation as a function of the cross-correlation window length for different sampling frequencies. (b) Corresponding time cost of each elastographic frame computed using the CPU implementation as a function of the cross-correlation window length for different sampling frequencies. (c) Estimated GFLOPS for the hybrid implementation as a function of the cross-correlation window length for different sampling frequencies. The RF data size is 1039×128 for 20 MHz sample frequency, 2078×128 for the 40 MHz sample frequency and 4156×128 for 80 MHz sample frequency.

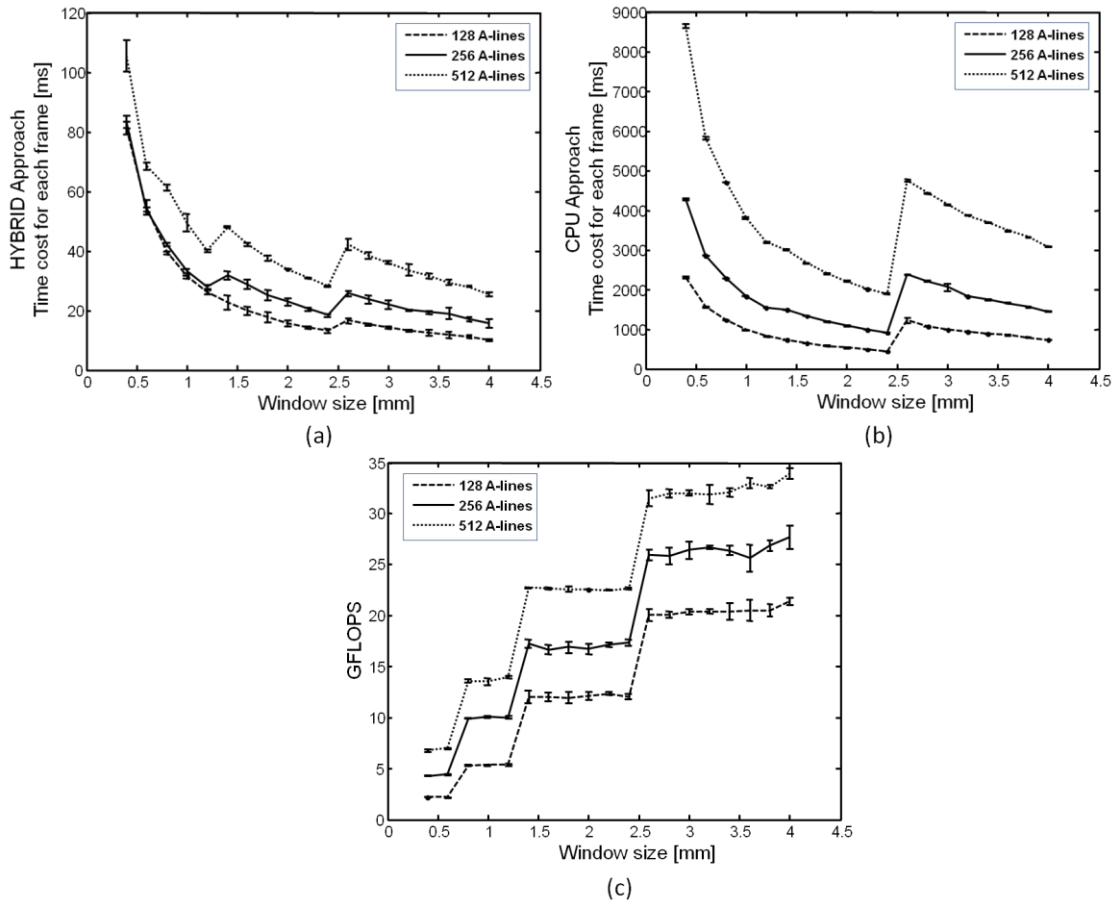


Figure 5 (a) Time cost of each elastographic frame computed using the hybrid implementation as a function of the cross-correlation window length for different number of A-lines. (b) Corresponding time cost of each elastographic frame computed using the CPU implementation as a function of the cross-correlation window length for different number of A-lines. (c) Estimated GFLOPS for the hybrid implementation as a function of the cross-correlation window length for different number of A-lines. The RF data size in the axial direction is 2078.

Image quality assessment

Figures 6-9 show the results of the image quality analysis performed using simulations. Figures 6-8 shows simulated axial strain elastograms for the simulation model containing a cylindrical inclusion. These elastograms were obtained for different values of applied strains (figure 6) and cross-correlation window length (figure 7). The

corresponding ideal images (from the FEM simulation model) and the CPU elastograms are also shown.




Figure 6 (a-c) Examples of simulated ideal strain images (from FEM) obtained by simulating an elastic medium containing a cylindrical inclusion at different applied strain levels. (b-d) Corresponding simulated axial strain elastograms obtained using the hybrid version of the algorithm. (e-g) Corresponding simulated axial strain elastograms obtained using the CPU version of the algorithm. For the elastograms, the cross-correlation window length was fixed at 2 mm. No averaging was used.




Figure 7 (a-c) Examples of simulated axial strain elastograms obtained using the hybrid version of the algorithm as applied to the simulation model containing a cylindrical inclusion at a 2% applied strain. (d-f) Corresponding simulated axial strain elastograms obtained using the CPU version of the algorithm. Data were processed using cross-correlation window lengths ranging from 1 mm to 3 mm. No averaging was used. W stands for window size in mm.

Figure 8 shows the results of the SNRe study, while figure 9 shows the results of the CNRe study as a function of the applied strain and window length. These data refer to simulated axial strain elastograms obtained using the hybrid and CPU versions of the algorithm. In both cases, the SNRe and CNRe curves corresponding to the CPU and hybrid versions of the algorithm are practically overlapping, which explains why the elastograms in figure 6 and 7 have no perceptible visual differences. The NRMS error

shows that the difference in precision between the two implementations is larger for higher strains and smaller window sizes. In both cases, the NRMS is found below 1% when the applied strain is less 5%, which is typically the case in elastography applications.

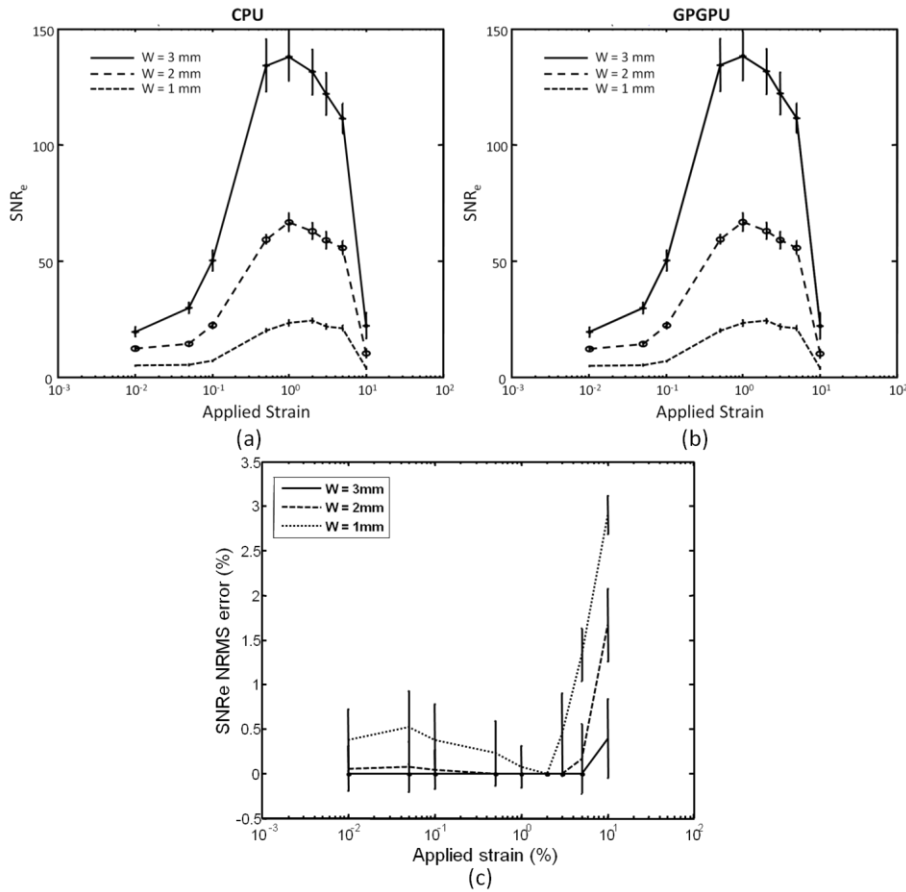


Figure 8 Results of SNRe study: (a) CPU results; (b) GPGPU results; and (c) NRMS error results.

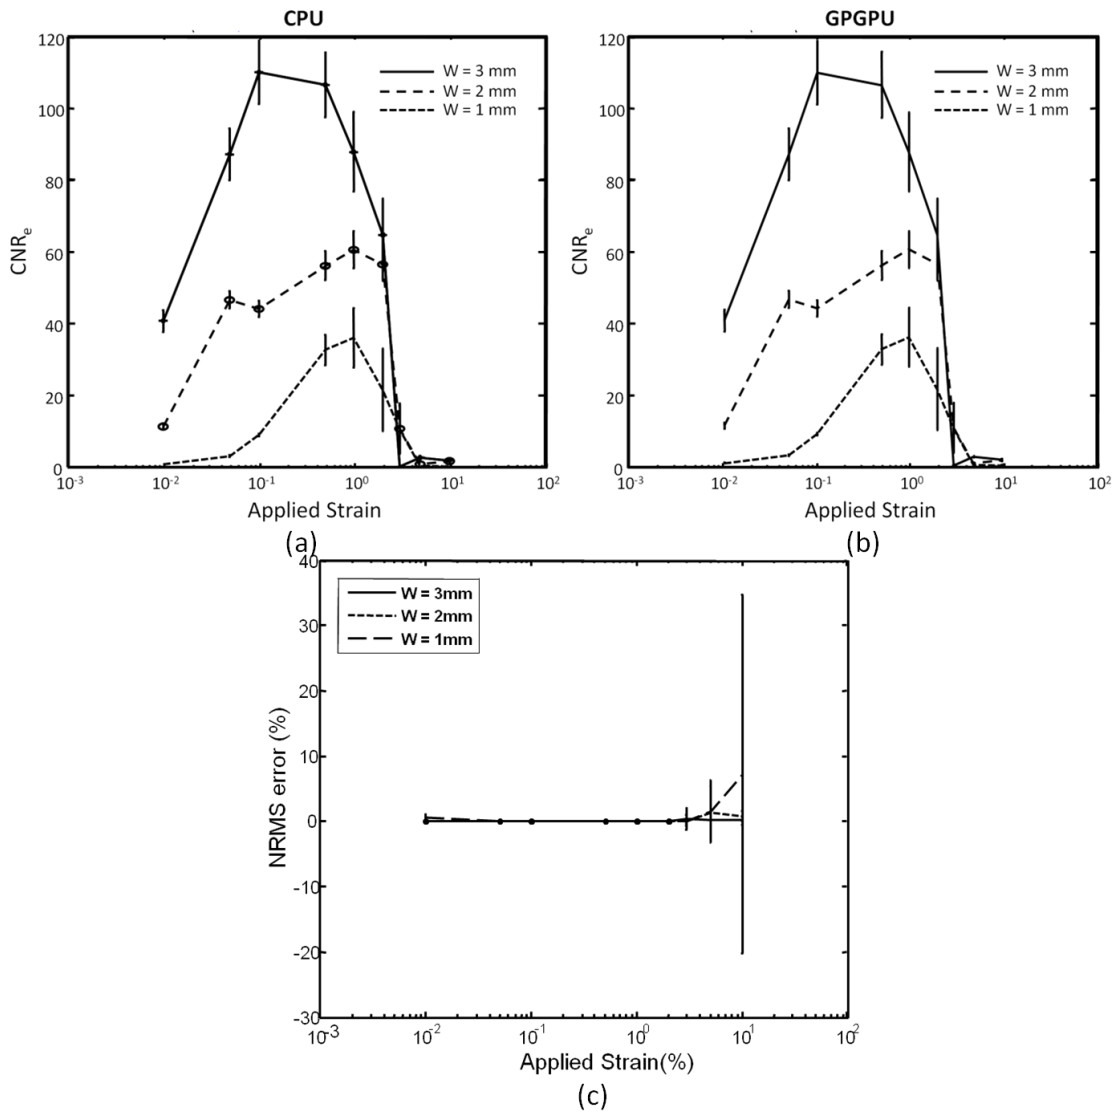


Figure 9 Results of CNRe study: (a) CPU results; (b) GPGPU results; and (c) NRMS error results.

The results of the image quality analysis carried out in this study show that there is no statistically significant difference in the analyzed image quality factors between CPU elastograms and GPGPU elastograms.

Experimental results

Figure 10 shows examples of experimental axial strain elastograms obtained from a gelatin phantom using the CPU and hybrid version of the elastography algorithm. As it can be observed by the results shown in figure 11, there is no appreciable difference between the experimental elastograms obtained using the CPU- and GPGPU-based versions of the algorithm, as expected from the statistical image quality analysis reported in the previous section.

As additional examples of experimental application of GPGPU elastography, Figs. 11-12 show examples of experimental axial strain elastograms obtained from a meat sample (Fig. 11) and two liver samples (Figs. 12 and 13), each of them containing a HIFU-induced lesion.

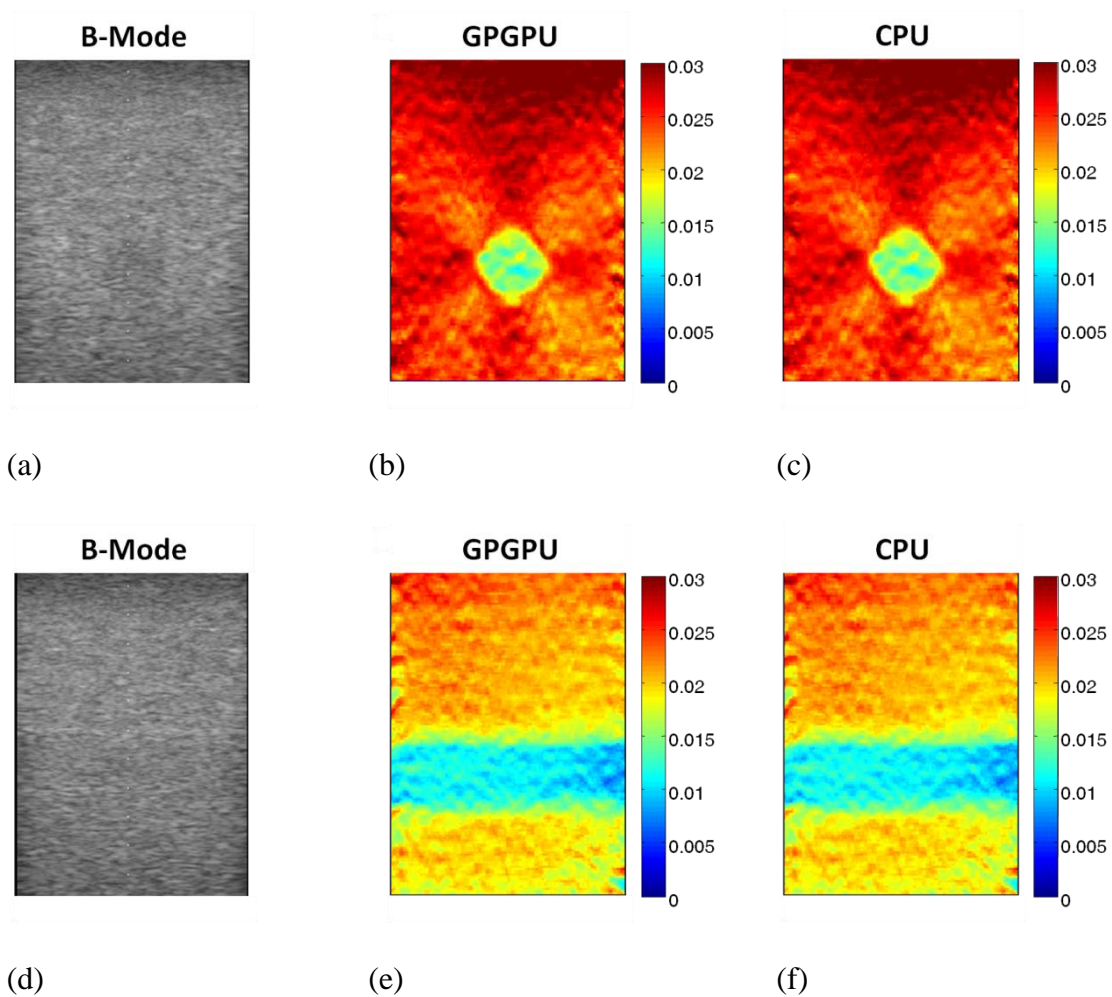


Figure 10 Example of experimental axial strain elastograms obtained using the Hybrid.AxialOnly version of the algorithm (b and e). For comparison, the corresponding B-mode images (a and d) and CPU axial strain elastograms (c and f) are also shown. Data refer to a gelatin phantom containing a cylindrical inclusion approximately two times stiffer than the background. Each elastogram is an average of 10 elastograms. Images (a)-(c) refer to data acquired with the transducer perpendicular to the axis of the cylinder. Images (d)-(f) refer to data acquired with the transducer parallel to the axis of the cylinder.

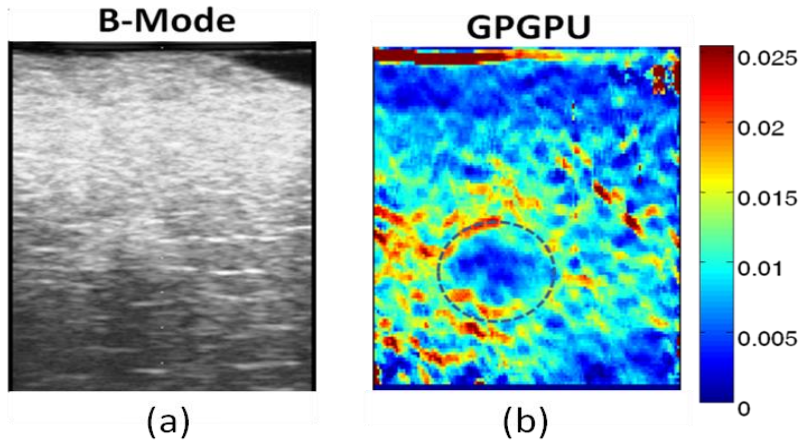


Figure 11 Example of experimental axial strain elastogram obtained using the hybrid approach (b). For comparison, the corresponding B-mode image (a) is also shown. Data refer to a meat sample containing a HIFU-induced lesion (circled area in the elastograms) scanned in free-mode acquisition. No averaging was applied.




Figure 12 Example of experimental axial strain elastogram obtained using the hybrid approach (b). For comparison, the corresponding B-mode image (a) is also shown. Data refer to a liver sample containing a HIFU-induced lesion (indicated by the arrow) scanned in free-mode acquisition. No averaging was applied.

DISCUSSION

In this paper, we have presented a new technique that allows generation of elastograms in real-time and with no loss in image quality. This technique is based on

the use of graphic cards to speed up computation. GPGPU cards are becoming popular in biomedical applications. They are in general cheap and widely available. Thus, they may provide an attractive alternative to software-based solutions for real-time elastography applications. In addition to real-time elastography, the findings reported in this manuscript might be of use for the design and implementation of other elastography techniques that deal with applications involving a large amount of data and computations (Lindop, Treece et al. 2006) and, possibly, to other biomedical engineering areas.

We have demonstrated how a hybrid computation method that uses both CPU and GPGPU can be accomplished through iterative optimization. Our study demonstrates that by using GPGPU as a co-computational device to the CPU, cross-correlation based elastography algorithms can be accelerated to deliver real-time performance. Our statistical analysis indicates that the loss in image quality when using the hybrid approach with respect to the CPU-based approach is very limited. The image quality comparison between GPGPU and CPU was necessary because of the single floating point precision of the GPGPU card employed in this study (vs. the double floating point precision of the CPU).

With respect to previous real-time strain estimators, the newly proposed hybrid algorithm performs comparably in terms of speed. For a typical 4 cm × 4 cm area sampled at 40MHz using a transducer with 128 elements, this implementation can achieve 56.4 fps when 2mm window length and 80% overlap are used (12800 pixels in the elastograms). However, it may have several important advantages with respect to previous real-time strain estimators. First, the statistical analysis performed in this study

using simulations demonstrates that the proposed elastography algorithm provides high quality elastographic estimates, with SNRe and CNRe values in the range typical for ‘slow’ elastographic estimators but significantly higher with respect to real-time strain estimator algorithms previously reported in the literature (Zahiri-Azar and Salcudean 2006). This may be expected because the majority of real-time elastography algorithms previously proposed are based on the use of sub-optimal displacement estimators or prior information to accelerate computation (Hall, Zhu et al. 2003, Zahiri-Azar and Salcudean 2006). It is indeed known that cross-correlation TDE algorithms are optimal (with respect to the mean squared error) and outperform spectral estimator, sum-of squared difference and sum-of absolute difference elastographic algorithms (Viola and Walker 2003). A recent study conducted in our laboratory confirms the statistical superiority of cross-correlation elastographic estimators both for axial displacement and lateral displacement estimations (Sambasubramanian August 2010). Second, the newly proposed GPGPU-based elastography has the advantage to minimize the utilization of CPU in the ultrasound system, which can be used for other operations.

Experiments were performed as a proof of principle of the applicability of the algorithm to real data. We performed two sets of experiments. The first set of experiments involved the use of conventional gelatin phantoms under controlled compression conditions. The second set of experiments involved the use of ex vivo tissues containing HIFU-induced lesions under free-mode acquisition conditions. The reason for these experiments can be justified as follow: first, the necessity of testing the algorithm using more complicated samples in more realistic noise and clinical

conditions; second, the large interest in studying the suitability of real-time elastographic estimators for monitoring HIFU-based treatments(Kallel, Stafford et al. 1999, Righetti, Kallel et al. 1999). The preliminary experimental results reported in this study appear to confirm the simulation findings.

By harnessing the computational horsepower of both CPU and GPGPU, we achieved real-time performance using conventional elastography for most of the acquisition and imaging configurations and without loss in image quality. The hybrid approach also shows more efficiency when RF data size increased. Beyond this, parallelization of cross-correlation based elastography algorithms is limited because of the inherent data-dependency of the algorithm.

Several challenges are encountered when implementing conventional cross-correlation based elastography algorithms on GPGPU to achieve high acceleration ratios. Divergence and uncoalesced memory access will cause efficiency penalty in SIMT architecture. Cross-correlation based elastography algorithms contain FFT, sorting and median filtering, which are destined to have these problems. In fact, the highest GFLOPS that our implementation can reach is limited by CUFFT library. As shown in our performance analysis (figure 3), cross-correlation related computation (FFT, FFT products and sum of squares) uses 53.70 % of the GPGPU time. More specifically, computation of FFT and IFFT occupy 41.14% of the GPGPU time, thus heavily affecting the overall computational speed. From our performance analysis (table 1), we can also observe that CUFFT is largely affected by warp serialization problems, high branching (although divergent branching has been largely avoided) and uncoalesced

global memory stores. Finally, the speed is also limited by the platform we are targeting. GeForce 8800 GT does not have a very strong computation power, but its characteristics in terms of power consumption fit our ultrasonic imaging system limitations. Our Sonix RP system has a total power supply of 400 W, a constraint that leads to a compromise between performance and system's safety when choosing the GPGPU to be used within the system.

For this algorithm, we see a potential improve on speed by increasing GPU occupancy. Improvement in speed could also be obtained using different TDE estimations algorithms, i.e., elastography estimators that do not use cross-correlation. However, this could lead to decreased image quality, which in clinical application is not desirable (Sambasubramanian August 2010). Further improvement in speed can be obtained only with higher performance hardware and more power supply from host. A G200 based Quadro FX 5600 (4 cores), executing multi-threaded version of our software, will deliver 4 times more performance than a single G92 based GeForce 8800 GT. However, GPGPUs with more computation resources are more expensive and have power consumption that can easily exceed 200 W, which might not be suitable for many US systems (unless a high output power supply unit is used).

The above observations lead to considering the heterogeneous approach of computing where CPU, GPGPU and other compute devices are allotted computational load that agree best to their underlying architecture, for achieving performance optimization at little extra cost. On this regard, the OpenCL standard and the availability of its libraries could make load-balancing through heterogeneous computing feasible.

CONCLUSION

Using GPU as a co-computational device to the CPU, we have demonstrated that cross correlation based elastography can be accelerated to deliver real time performance at 3fps against an initial benchmark of 0.1fps, with no loss in visual or statistical image quality. There are limitations to parallelizing cross-correlation based elastography further because of the inherent data-dependency in the algorithm. Further improvement in speed can be obtained only with higher computation power hardware. A 4-GPU G200 based Quadro FX 5600, executing multi-threaded version of our software, will deliver 4 times more performance than a single GPU G80 based GeForce 8800M GTS. However, with cost to performance ratio being the goal, we soon hit the price wall. This leads to considering the heterogeneous approach of computing where CPU, GPU and other compute devices are allocated computational load that agree best to their underlying architecture, for achieving performance optimization at little extra cost. The OpenCL standard and the availability of its libraries can make load-balancing through heterogeneous computing a reality.

Integration of our implementation with an ultrasonic diagnostic system equipped with GPU hardware and testing of the software with in-vivo tissue is part of our future work.

3 BONE SURFACE ENHANCEMENT IN ULTRASOUND IMAGING USING A NEW DOPPLER-BASED ACQUISITION/PROCESSING METHOD ²

INTRODUCTION

Ultrasound (US) imaging has been shown to be a potentially useful tool for orthopedic surgeries (Barratt, Penney et al. 2006) and to visualize the healing process of bone fractures (Li, Le et al. 2013). While US imaging has many advantages compared to more traditional bone imaging methods such as safety (Fayaz, Giannoudi et al. 2011) and lack of radiations, US bone assessment is limited by noise, presence of speckle, attenuation, hand-held motion and imaging-related artifacts. Unlike computed tomography (CT), the tissue and bone signal intensities in US images cannot be calibrated. In many cases, depending on the incident angle between the US beam propagation and a bone surface, US images of bones are heavily affected by reverberation artifacts. Consequently, complex signal processing and pattern recognition methods need to be used in bone segmentation algorithms to obtain meaningful quantitative analysis or for 3D ultrasonic bone reconstructions (Carson, Oughton et al. 1977, Sehgal, Lewallen et al. 1988, Lasaygues, Ouedraogo et al. 2005, Hacihaliloglu, Abugharbieh et al. 2008, Laugier and Haät 2011). These methods are computationally intense, slow down processing speed and often offer marginal improvement to the bone segmentation process.

² Reprinted with permission from "Bone surface enhancement in ultrasound images using a new Doppler-based acquisition/processing method " by Xu Yang, Songyuan Tang, Ennio Tasciotti and Raffaella Righetti. Phys Med Biol. 2018 Jan 17;63(2).

The localization of a bone surface in a US image is normally characterized by different features such as a high shadow profile underneath the location of the soft tissue/bone interface and high intensity across the bone surface (Blankstein 2011). In the past, acoustic signal-to-noise ratio at the bone surface has been measured to be of the order of 10 dB (Culjat, Choi et al. 2008), depending on a number of factors including the properties of the US system used for data acquisition and the location of the bone. However, it is well known that ultrasonic bone segmentation methods based solely on intensity information perform poorly due to the limited contrast between the bone and soft tissue and the presence of other imaging artifacts (Hacihaliloglu, Abugharbieh et al. 2008). In some applications, a priori information about the bone is used to improve bone localization and segmentation (Thomas, Peters et al. 1991). For example, Thomas et al. have discussed the possibility to segment a fetus femur using purely morphological operations, which requires prior knowledge of the shape of the bone anatomy (Thomas, Peters et al. 1991). Shape information has been exploited by using active contours for the crux of the segmentation process (He and Zheng 2001, Alfiansyah, Streichenberger et al. 2006). Fuzzy logic has also been used for bone segmentation with prior knowledge about the osseous interface (Vincent, Tonetti et al. 2004). Kowal et al. proposed a method that employs adaptive threshold of the low intensity regions coupled with morphological filtering to reduce noise emanating from echoes (Kowal, Amstutz et al. 2007). The aforementioned methods are image processing techniques that operate on the B-mode images. Methods that use 'raw' ultrasonic data, also referred to as radio frequency (RF) data, for bone segmentation have also been developed. Xu Wen used a

power reflection detector to locate the bone surface and found that the location of the bone coincided with the maximum reflection ratio (Wen and Salcudean 2007). Another proposed method is based on the comparison between pre- and post-compression RF data and application of a rigid bone position technique (Doctor, Vondenbusch et al. 2011). Local phase features have been recently used for bone segmentation and fracture detection purposes (Hacihaliloglu, Abugharbieh et al. 2008, Hacihaliloglu, Abugharbieh et al. 2009, Hacihaliloglu, Abugharbieh et al. 2011) and demonstrated many advantages over intensity-based methods. However, calculation of phase features is typically computationally intense. Reconstruction of a 3D volume may need segmentation of the bone surface in hundreds of planes (images), and the accumulated processing time can be significant. Ideally, US elastography (T and J 1997) should be a reliable tool to separate a bone surface from the surrounding soft tissue since the Young's Modulus of bones can be several orders of magnitude greater than the Young's Modulus of soft tissues. While strain images can be used to obtain some mechanical information about the soft tissue in proximity of the bone surface (Tang, Chaudhry et al. 2017), in its present form and without any user-input information, strain elastography alone cannot be used to automatically detect the exact location of a bone surface. This is due to a number of reasons, which include the poor reliability of the ultrasonic signal below the bone surface due to high attenuation and other artifacts and the presence of soft tissue structures generating similar strain contrast and patterns to the ones generated by the bones in the elastographic images (Parmar, Yang et al. 2015). While, theoretically, the strain contrast corresponding to the soft tissue-bone interface should be much higher

than the strain contrast between different soft tissue layers, in practice, the strain concentration pattern at the soft tissue-bone interface may not stand out significantly with respect to other strain concentration patterns especially in heterogeneous tissue environments due to contrast-transfer-efficiency (CTE) limitations (Kallel and al. 2001). The use of parametric elastographic techniques and/or inversion methods to generate maps of underlying mechanical parameters is limited by the lack of suitable, realistic and sufficiently complex theoretical models, which would need to consider a priori information about the bone and boundary conditions as well as non-linearity, anisotropy and various degrees of connectedness between soft tissue and bones and would require high computation time.

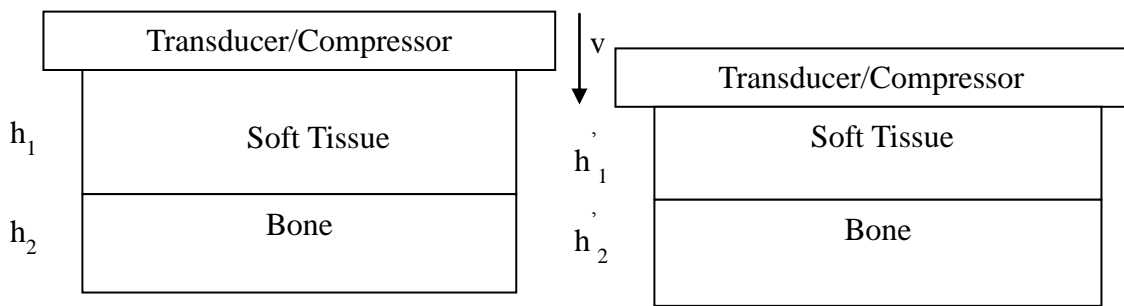
In this study, a novel and fully automated bone surface enhancement method for US is presented, tested and discussed. This method greatly increases the contrast between bone surface and soft tissue in a US image. Therefore, it should significantly facilitate the segmentation of bone surfaces, which may be performed using simpler and faster methods. The proposed method is based on Tissue Doppler Imaging (TDI) and knowledge that the mechanical and acoustic properties of bones are significantly different than those of soft tissues. TDI is typically used for blood flow measurements, but it has also been used to measure soft tissue motion (Agarwal, Gosain et al. 2012) such as myocardial motion (Smiseth, Stoylen et al. 2004, Gorcsan and Tanaka 2011) and for muscle assessment (Al Naimi, Fittschen et al. 2014). The proposed method is real-time and does not require specialized US systems. The real-time property makes it

particularly suitable for 3D applications. Results obtained from ex-vivo and in-vivo experiments are reported and discussed.

METHOD

When a tissue is compressed, the particles inside the tissue are displaced according to their mechanical properties. During compression, these particles move with a certain speed relative to the compressor, which, in an elastography experiment, is typically attached to the ultrasound transducer itself. To illustrate how these basic concepts can be used to enhance bone surfaces in US bone imaging applications, we use the schematic shown in Figure 13(a). Here, the overall tissue system is modeled as composed by two layers: the top layer is the soft tissue and the bottom layer is the bone. Before compression, the soft tissue has a thickness of h_1 , while the bone layer has a thickness of h_2 . After compression, the thickness of the soft tissue becomes h'_1 and the thickness of the bone layer becomes h'_2 . If the total thickness of the tissue is L before compression ($L = h_1 + h_2$) and L' after compression ($L' = h'_1 + h'_2$), the total strain experienced by this bi-layer system is $\epsilon = \frac{L-L'}{L}$. Let us assume that the compression occurs in a given time interval with constant speed, and that the compression is very small, i.e., in the range 1%-5%, as it is typically the case for elastography applications. Since the bone is several orders of magnitude stiffer than the tissue, most of the applied compression is taken up by the soft tissue layer while the thickness of the bone layer does not change appreciably (i.e., $h_2 \sim h'_2$). This results in very little strain inside the bone. Given the fact that the relative displacement between the transducer and the tissue

increases monotonically with depth and that ultrasonic penetration below the bone surface is very small, the peak of this displacement occurs at the bone surface. In terms of Doppler parameters, the Doppler frequency shift generated by the soft tissue/bone interface is the highest - any tissue above the bone surface would cause a lower Doppler frequency shift during the compression.



(b)

Figure 13 Simplified schematic of a bi-layer tissue compressed from the top using an ultrasound probe (with a compressor attached to it) at a constant speed v . (a) Model before compression: the transducer is moving with speed v towards the bone surface. (b) Model after compression: both the thicknesses of the soft tissue layer and bone layer reduce, but the reduction in the thickness of the bone layer is negligible compared to the one of the soft tissue layer.

In comparison with other soft tissue interfaces, the soft tissue-bone interface would typically have a larger reflection, so that the Doppler signal generated at the soft tissue-bone interface would have the maximal power as well. The ultrasound incident angle (i.e., the angle between the ultrasound beam propagation and the bone surface) usually depends on the compression direction (since the transducer itself is typically used to compress the tissue). When the tissue is subjected to uni-axial compression, the

incident angle is close to 90 degrees (Figure 13), and, since the motion of the particles in proximity of the bone surface is predominantly along the axis of compression, the Doppler angle between the direction of the particle motion and the axis of compression is close to zero degrees. In our technique, we use Power Doppler, so the angle information is not needed for the estimation of the Doppler frequency shift. Note that, unlike elastography that depends on signal penetration, the Doppler signals used in our model mainly relate to the reflection occurring at the soft tissue-bone interface. Therefore, these Doppler signals can, in theory, be generated more reliably than elasticity information using RF signals in the vicinity of the bone surface.

Figure 14 shows a schematic of the acquisition/processing scheme used in our model to extract the Doppler shift occurring at the soft tissue/bone interface. The transducer is used both to compress the tissue (1%-5% strain (T and J 1997)) and to acquire the RF data. The RF data are acquired while the transducer is moved towards the bone surface at a constant speed v_t . From the acquired RF data, displacement estimation is first performed using cross-correlation between two Color RF frames. This is used to find the maximum displacement in each A-line. As previously mentioned, the maximum displacement theoretically occurs at the soft tissue/bone interface. The central frequency of the Doppler signal generated by the relative motion between the transducer and the bone interface during the compression is given by $f_1 = 2 \frac{v_t}{v_c} f_0$, where v_c is the speed of sound in the tissue and f_0 is the central frequency of the transmitting pulse. Second, a phase-based mean frequency estimator is applied to the RF data (Jensen 1996). A band-pass filter centered at f_1 is then applied to the I/Q decomposed RF signals to eliminate

the Doppler signals generated from the soft tissue's structures and any reverberation artifact. The filtered signals (output of the band-pass filter) are then used to produce a Power Doppler image using autocorrelation-based methods (Kasai, Namekawa et al. 1985). Using the autocorrelation estimator (Namekawa, Kasai et al. 1983, Kasai, Namekawa et al. 1985), we can get the signal power, mean frequency of Doppler spectrum and bandwidth of Doppler spectrum. The signal power information is used to form the Power Doppler images used in this paper.

Based on the proposed model, the bone surface in the output image should be greatly enhanced in the Power Doppler image with respect to the original B-mode image, and this should facilitate bone surface segmentation.

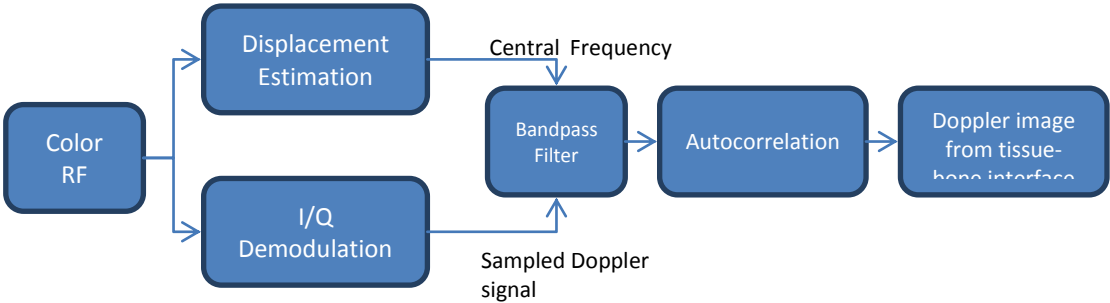


Figure 14. Signal processing procedure to obtain enhanced bone images using the proposed Doppler-based technique.

Experiments

Ex-vivo and in-vivo experiments were used to test the performance of the proposed method. Experiments were performed using a Sonix RP diagnostic ultrasound imaging system (Ultrasonix Medical Corp., Richmond, BC, Canada). Data were acquired with a 38-mm linear array transducer with 128 elements, 5 to 14 MHz bandwidth, 50% fractional bandwidth at -6dB, and 1mm beamwidth at the focus. The center frequency was set at 5-10 MHz, depending on the experiment, as detailed below. For Doppler data acquisition, pulse repetition frequency was 1428 Hz and sampling frequency was 20MHz. For the displacement estimation, we used a window size of 40 samples with an overlapping size of 4 samples. Each RF data package had 12 observations in time. These settings allowed us to obtain Doppler images with a theoretical axial resolution of approximately 1.5mm. In all samples tested in this paper, the upper bone surface (i.e., the surface closer to the transducer) was located at a depth of few cm (in all cases, < 2.5 cm). No signal could be reliably detected below the upper bone surface. A single focal zone was set at the location of the upper bone surface in the B-mode image prior data acquisition. The time gain compensation (TGC) was not used in our experiments.

Phantom experiments

Controlled phantom experiments were performed to test the sensitivity of the proposed technique. In these experiments, PVC pipes (1/2 inch × 1 inch) were embedded into gelatin blocks, fabricated as reported by Kallel et al. (Kallel and al. 2001). Each sample was then positioned in a controlled compression frame. A picture of the

compression frame used for the experiments is shown in Fig. 15. Due to the compression frame, the transducer motion was confined along the axis of compression. Three different transducer speeds were achieved: 1cm/s, 2cm/s and 3cm/s. In all cases, RF data were acquired during the compression. In this experiment, the transducer was moved for 0.1 second axially using each speed setting. Color RF data were acquired during the compression to calculate the Power Doppler image as explained in the previous section.




Figure 15 A picture of the setup used for the experiments.

Ex-vivo experiments

Ex vivo experiments were conducted on 3 chicken leg samples and 2 sheep leg samples, in all cases with intact soft tissue. The chicken samples were embedded into gelatin cases to assure stability during compression. The sheep leg samples, instead, were not embedded into gelatin cases (due to their dimensions). RF data were acquired both using the same compression apparatus used for the phantom experiments and in freehand mode. The freehand experiments were performed in preparation for the in vivo experiments, where the compression apparatus could not be used. The samples were

scanned at different locations both with the transducer parallel to the bone axis and with the transducer perpendicular to the bone axis.

In-vivo experiments

In-vivo experiments were performed on sheep legs to investigate the feasibility of the proposed method in more realistic noise conditions. Specifically, in living tissues, blood flow and wall movements could, in principle, affect the performance of the proposed method. The lower legs of two adult sheep were scanned for these experiments. The sheep samples were obtained from the Texas A&M Institute for Preclinical Studies (TIPS). This study was approved by the Texas A&M University Institutional Animal Care and Use Committee (IACUC) (ARO #60598-MS-DRP, Award W911NF-11-1-0266). In one case, the tibia was intact; in the other case, the tibia was fractured. Signals were acquired from anesthetized animals, in free-hand compression mode. Different locations of the sheep legs were imaged both with the transducer parallel to the bone axis and perpendicular to the bone axis. More details regarding the in-vivo experiment protocol can be found in Parmar et al (Parmar, Yang et al. 2015).

Statistical Analysis

For 40 in-vitro and in-vivo data acquisitions, we compared the local peak intensity contrast for each A-line (in each image) between the envelope of RF data (i.e., B-mode) and Power Doppler data. In general, local maxima are generated at the interface between different tissue layers due to reflection. While the bone-tissue interface should generate a peak in both the B-mode and Power Doppler data, additional tissue-tissue interfaces (such as soft tissue/connective tissue, etc.) can generate peaks of

intensity comparable to the ones generated by the bone-tissue interface in the B-mode data. The peaks generated by tissue-tissue interfaces can create false positives in bone segmentation, which is one of the reasons why a simple intensity thresholding method, in general, does not perform well in ultrasonic bone segmentation problems. Generally speaking, a “good” enhancement method (prior to bone segmentation) should reduce the probability of false positives, which in most cases are generated by tissue-tissue interfaces. To evaluate if the use of the proposed Power Doppler method would indeed help reducing the presence of false positives caused by tissue-tissue interfaces, we defined the following metric:

$$\text{SNR} = \frac{1}{N} \sum_{i=1}^N \frac{P^i}{P_{\text{avg}}} \quad (1)$$

where P is the local peak pixel intensity value in each column of the image’s original RF data, P_{avg} is the average peak value except the maximum peak and N is the number of A-lines (columns) in an image. This metric represents the ratio between the maximum local peak intensity and the average of other local peak intensities, and, for our specific application, it is indicative of the signal-to-noise ratio. This metric was computed for both the B-mode and the Power Doppler signals. In reference to the B-mode signal, this metric shows the contrast between the highest echogenic tissue interface and the average among other interfaces. In reference to the Power Doppler signal, it shows the ratio of the signals from the interface between materials with the highest acoustic and mechanical property contrast to those from other interfaces.

To illustrate how this parameter is computed, we can use the simulated ultrasound signal in Fig. 16. In this specific example, the signal shows 3 local peaks, with P2 being the highest. Therefore, for this specific signal, the $SNR = P2/[(P1+P3)/2]$.





Figure 16 Simulated ultrasound signal from one A-line, used to illustrate how the SNR is computed in our application.

RESULTS

Phantom experiments

Figure 17 shows selected Power Doppler results from the controlled phantom experiments. These results were obtained with the transducer moving at speed: (a) 1cm/s (b) 2cm/s (c) 3cm/s. The only purpose of these controlled experiments is to verify the sensitivity of the proposed method to speed detection. It can be observed that Power Doppler can be used to detect the gelatin-PVC interface in all cases. In addition, the contrast-to-noise ratio between gelatin/PVC in the Doppler images increases when the compression speed increases. These results are consistent both when the compressor apparatus was used and in freehand mode.



(a) (b) (c)
Figure 17 Power Doppler images created using different surface-transducer relative speeds and the same scale: (a) 1cm/s (b) 2cm/s (c) 3cm/s.

Controlled in vitro experiments

Figure 18 shows examples of bone enhancement obtained from a chicken leg embedded into gelatin (figures 18a and 18b) and an ex-vivo sheep leg (figures 18c and 18d) both in freehand compression conditions. In both cases, the samples were scanned with the transducer perpendicular to the axis of the bone. Figure 18a shows the B-mode image obtained from the chicken leg. The bone surface appears in the B-mode image as an upside-down “U” (Blankstein 2011), with intensity somewhat higher than in the soft tissue. In this image, both the tibia and the fibula are visible (white arrows). Note, however, that other soft tissue structures (yellow arrow, on the left side) may appear to have similar intensity as the bone surface. This is because the soft tissue layer itself may also be inhomogeneous (there is fat, connective tissue, muscle, etc.). At the various tissue interfaces, ultrasound reflection is generated due to the mismatch of acoustic impedance and can show up as a locally high intensity peak in the B-mode image. In some cases, these local peaks can be as high as the peak from the soft tissue/bone interface. The presence of such soft tissue structures typically causes inaccuracy in

intensity-based bone segmentation methods. Figure 18b shows the corresponding Power Doppler image obtained using the proposed technique. In this image, the bone surfaces of both the tibia and fibula are enhanced (white arrows) and have significantly higher intensity than the connective tissue (yellow arrow).

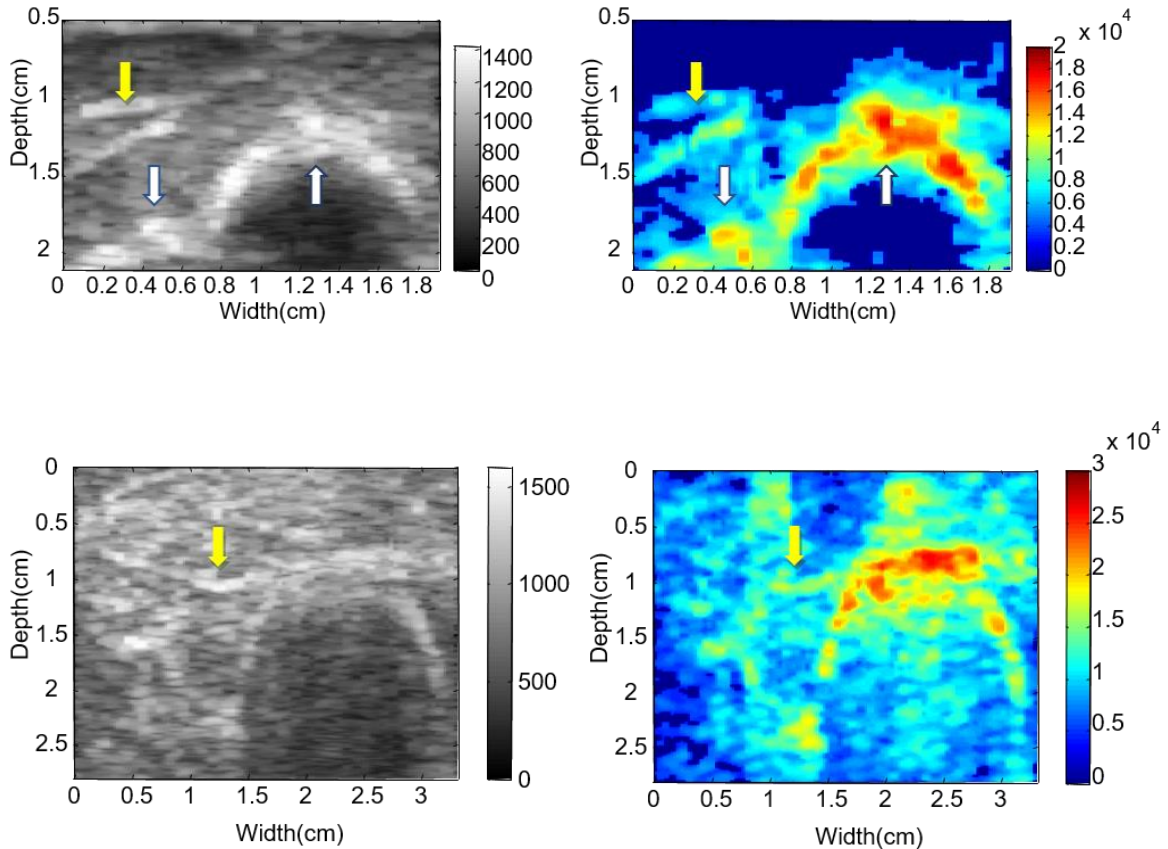


Figure 18 Selected in-vitro experimental results. Top panel refers to a chicken sample, while bottom panel refers to a sheep leg. In both cases both the B-mode image and the corresponding Power Doppler image are shown.

Figure 18c and 18d refer to data acquired from the sheep leg. In general, these images are significantly noisier than the chicken ones. This is presumably due to the geometry of the limb and the fact that the soft tissue has intensity comparable to that at

the soft tissue/bone interface. Consequently, bone surface detection is challenging in both the B-mode and the Power Doppler images. However, the contrast between soft tissue/bone interface and surrounding soft tissue is shown to be improved after the Doppler processing. Note also that the high intensity soft tissue structures are greatly reduced in the Power Doppler images (yellow arrows in 6c and 6d).

In vivo experiments

Figure 19 shows a set of results obtained from an intact adult sheep tibia in vivo. Figure 17a shows the B-mode image and Figure 19b shows the corresponding Power Doppler image. For these images, data were acquired with the transducer parallel to the bone axis. While the bone surface is clearly visible in the B-mode image (white arrow), portions of the soft tissue above the bone surface appears to have comparable intensity to that at the bone-soft tissue interface. In contrast, the signals from the soft tissue are shown to be significantly lower than the signal generated by the bone surface in the Power Doppler image (Figure 19b).

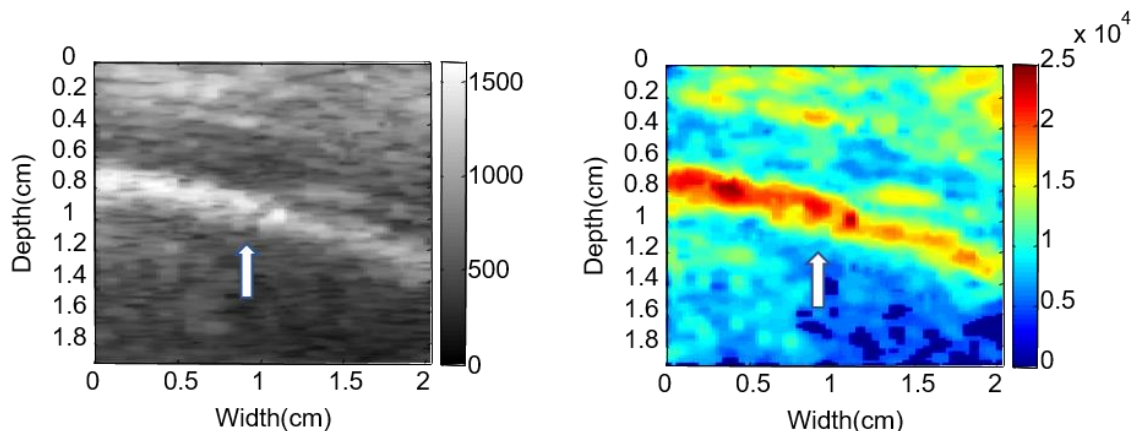


Figure 19 Selected results obtained from an intact adult sheep tibia in vivo.

Figure 20 shows another set of data obtained from the in vivo sheep study. Similarly to the previous case, the transducer was oriented parallel to the bone axis. The bone surface enhancement is clear in the Power Doppler image throughout the upper bone surface (yellow arrow in 20b). A measure of the SNR (dB) computed as explained previously is also reported. The SNR values associated with the Doppler images are significantly higher than those obtained from the B-mode images. Note that the multi-reflection artifact below the bone surface (white arrows in 20a and 20b) is reduced on the Power Doppler image.

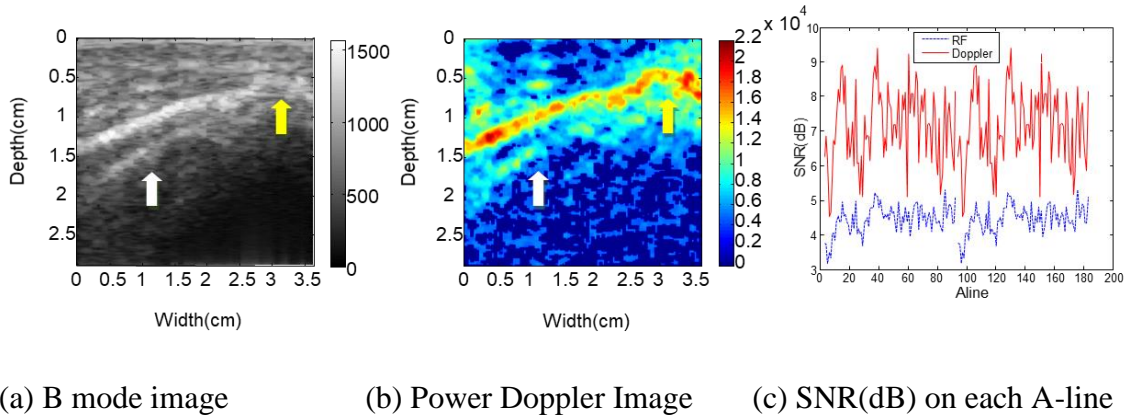



Figure 20 Selected results obtained from an intact adult sheep tibia in vivo.

The set of images shown in Figure 21 refers to a case where a tibia fracture was visible. Figure 21a shows the B-mode image. The fracture can be clearly localized in this image, but, as for the intact case, large portions of the soft tissue present with intensity comparable to that of the bone surface. Figure 21b shows the corresponding Power Doppler image. The soft tissue signals are almost entirely suppressed in this image. Figure 21c shows the Color Doppler Image, which is used here to evaluate the blood and

soft tissue wall movements' influence on the frequency shift. The Color Doppler image appears to indicate the presence of a blood vessel, but note that the bone surface is not visible in the Color Doppler image. This is because blood speed is higher than freehand compression speed, and therefore signals from the bone surface are filtered out by the filter used to generate the Color Doppler images. By comparing figures 21b and 19c, the Power Doppler signal intensity generated by the soft tissue/bone interface is 1-2 magnitudes higher than the Power Doppler signal intensity generated by the blood velocity. This is because the intensity associated to the bone surface can be orders of magnitude higher than the intensity generated by blood cell scattering. Note also that Doppler signals associated to vessel walls don't show up in figure 21b. This is because vessel wall motion usually creates Doppler signals at 50-100Hz (Kruskal, Newman et al. May 2004), and the band-pass filter applied in the proposed bone method had a highest cut-off frequency of 45 Hz. This demonstrates that Doppler signals from bone and blood manifest at different frequencies and can be separated in different Doppler images.




(a) B-mode Image


(b) Power Doppler Image

(c) Color Doppler Image

Figure 21 Images from adult sheep II tibia with fracture in vivo.



(a)



(b)

Figure 22 Comparison of normalized B-mode and Doppler signals obtained from columns of the image shown in Figures 19a (dotted lines) and 19b (solid lines).

Figure 22 shows two examples of normalized signal axial profiles obtained from the B-mode image's original RF envelope (dotted lines) shown in figure 21a and the Power Doppler image's original RF envelope (solid lines) shown in figure 21b. In the B-

mode data, the soft tissue has an echo intensity comparable to that of the bone surface. In contrast, the Power Doppler signals corresponding to the bone surface are significantly higher than those associated to the soft tissue in both cases. The SNR associated to the Power Doppler signal was found to be statistically significantly higher when compared to the one obtained from the B-mode one.

Statistical Analysis

For 40 in vitro and in vivo data acquisitions, we compared the SNR for each column (in each image) both for the B-mode and Power Doppler data. The results of our statistical analysis are shown in Figure 23. Our data show that the presence of a bone surface (which is typically highly reflective and harder than the surrounding tissue) would generate an increase in the SNR of the Power Doppler signal (1172.9) approximately 14.5 times higher with respect to the B-mode case (80.5).

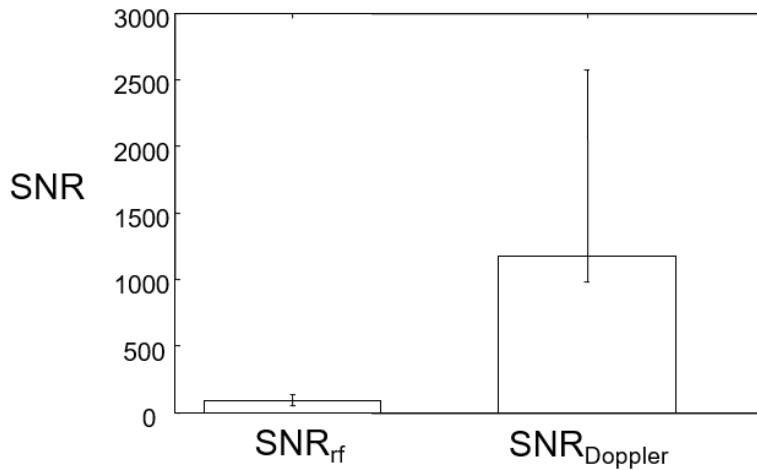


Figure 23 SNR Results: left column shows average contrast at 80.5 (+56.5,-28.5) for regular B-mode data; right column shows average contrast at 1172.9 (+1397.0,-191.5) for Power Doppler data.

DISCUSSION

In this paper, we have proposed and investigated the use of a novel bone enhancement method, which is based on the use of Power Doppler imaging and elasticity imaging concepts. Inspired by elastography, this method takes advantage of the bone's mechanical and acoustical characteristics, which are significantly different than those of soft tissues. The method was tested in a variety of experimental conditions. In all in vitro and in vivo cases analyzed for this study, the application of the proposed method allowed to enhance the bone surface in ultrasonic images and reduce the signal generated by the surrounding tissue.

To better understand the idea at the basis of the proposed method, we need to consider the following. In general, the presence of high reflective tissue is greatly reduced in Power Doppler images of bones because its average Doppler frequency is lower than that corresponding to bone surfaces. Typically, if we indicate with Δf the Doppler frequency shift, we would expect $\Delta f_{\text{soft tissue}} < \Delta f_{\text{connective tissue}} < \Delta f_{\text{bone surface}} < \Delta f_{\text{vessel wall}} < \Delta f_{\text{blood}}$. The difference in Doppler frequencies between connective tissue and bone surface depends on the thickness of the tissue, the depth of the tissue and difference in the mechanical properties between soft tissue and connective tissue. Thus, by properly selecting the frequency limits of the band-pass filter, signals not associated with the bone surface can be reduced or eliminated. Additionally, the proposed method is expected to reduce the effect of multiple reflections from the soft tissue/transducer interface or the soft tissue/bone interface. In US imaging, multiple reflections typically arise in non-homogeneous materials when the thickness of a layer is

larger than the ultrasonic wavelength, and there is a large acoustic impedance mismatch between the tissue layer and the background (Zheng, Le et al. 2007). Since the cortical bone thickness is larger than the ultrasonic wavelength and has a significant different acoustic impedance with respect to the surrounding soft tissue, imaging artifacts caused by multiple reflections are very common in US bone imaging and have become a major source of noise for bone reconstruction methods. The performance of bone segmentation methods based on RF or B-mode intensity can be greatly affected by such artifacts because, in many cases, the intensity of the artifacts can be comparable to the intensity of the actual bone surface. In our method, these artifacts can be filtered out because they would have a different (higher) Doppler shift than the one associated to the bone surface. The success of the proposed method depends on several factors. One is the synchronization between the RF signal acquisition and the compression. If the relative compression speed is not constant during acquisition, this could introduce noise in the Doppler estimation. This is because we have implemented a fixed frequency filter to extract the signal. When the compression speed matches the filter frequency, the contrast is greatly enhanced. This contrast enhancement is reduced when the speed and filter frequency do not match. In the future, the result can be improved by adaptively choosing cut-off frequencies for each line in the image if freehand compression doesn't have a constant velocity during each acquisition.

Based on the results reported in this paper, the Power Doppler method proposed for bone enhancement does not seem to be affected by the presence of blood flow, despite the fact that the velocity of blood in major blood vessels is much higher than the

transducer's compression speed. This is because the signal intensities from bone surface and blood cells are typically very different, which is the dominant factor over the velocity. In fact, Rayleigh scattering signals generated by red blood cells are very weak compared to strong reflective signals generated by bone surfaces. As for small vessels, the difference in relative speed itself can also be low.

The intensity of the Doppler signals generated by blood vessel wall motion is stronger than the intensity of the Doppler signals generated by blood motion. However, it is still not comparable to the signal intensity at the soft tissue-bone interface. In addition, common wall filters used to remove blood vessel signals operate between 50Hz and 100Hz (Rubens, Bhatt et al. Jan. 2006). In our study, we found that the band-pass filter has typically a cut-off frequency of 15-20Hz. Thus, this filter would eliminate most of the signals generated by the vessel wall motion and blood motion. In our experiments, vessel wall motion did not affect the Power Doppler bone images after the application of the low-pass filter unless the transducer was parallel to the blood vessel.

Power Doppler images are typically more sensitive to the flash artifact than Color Doppler images (Bude and Rubin 1996). In our experiments, the flash artifact was not observed or not visibly comparable to the Power Doppler signals generated by tissue interfaces. This may be due to the fact that our method focuses on signals from tissue interfaces where reflection is the dominating effect. While transducer motion can generate signals from tissue scattering, the intensity of scattered US signals is typically not comparable to the intensity of reflected US signals. In our experiments, we concentrated only on motion directions perpendicular to the bone surface. In the future,

the presence and effect of flash artifacts when imaging complex bone geometries may need to be evaluated. Aliasing is another common Doppler artifact, but it does not affect Power Doppler. So, in our experiments, we did not need to adjust the pulse repetition frequency to the different transducer compression speeds.

In our proposed method, we have used quasi-static elastography in multi-compression mode to compute the displacements and determine the frequency cut-off of the band-pass filter. In principle, other elastographic techniques could be used for this purpose. Some advantages of quasi-static elastography include that it is real-time and that it does not require the use of complex theoretical models. It should be noted, however, that the proposed method differs significantly from established elastography techniques, which are typically optimized for soft tissue imaging. No standard elastography method has proven to be suitable for automatic bone detection yet. Our proposed method is optimized to detect the bone surface and reduce the soft tissue information. This is in complete contrast to current standard elastography methods. As a result, the generated bone image is stable and the location of the bone is relatively easy to identify.

In some of the Power Doppler images collected in this study, the spatial resolution appears to be reduced with respect to the corresponding B-mode images presumably due to the filtering introduced during the Doppler estimation. Nevertheless, in the future, bone surface segmentation may be improved by combining the information from both the B-mode and Doppler images, where the Doppler image could be used to

automatically detect the location of the bone, while the actual bone segmentation could be performed on the B-mode image at that location to preserve resolution.

Perhaps, the major advantages of the method reported in this study are its relative simplicity and the limited overall computational cost since all processing steps can be performed in real-time on commercial US imaging systems. When used in conjunction with a bone segmentation technique, the proposed method has the potentials to significantly reduce the overall complexity of the bone segmentation process and improve its performance.

CONCLUSION

In this paper, a US bone enhancement method is proposed. This method is designed to work robustly for in-vivo freehand scans, and it is based on the fundamental concept that bones have distinct mechanical and acoustic properties with respect to soft tissue. The method was tested both in ex vivo and in vivo experiments. Based on the results presented in this paper, the availability of the proposed Power Doppler-based technique could improve ultrasound bone surface detection and enable fast volume data processing for 3D reconstructions.

4 BONE SURFACE SEGMENTATION FOR FREE HAND ULTRASOUND IMAGE BASED ON STATISTICAL SHAPE MODEL

INTRODUCTION

Bone fracture assessment using ultrasound imaging techniques has received increasing attention recently for a number of reasons, including the potentials of these techniques to be used in intra-operative scenarios (Beek, Abolmaesumi et al. 2008) and to non-invasively and safely monitor bony healing (Ricciardi, Perissinotto et al. 1993) (Li, Le et al. 2013) (Chen, Lin et al. 2014). Since ultrasound imaging modalities are non-ionizing, their application to monitor bone growth in children has also been actively investigated in the past few years (Chen, Kim et al. 2007, Poonai, Myslik et al. 2017). However, bone segmentation and reconstruction from ultrasound images can be affected by low image quality due to the presence of speckle, imaging and motion-related artifacts generated during free-hand scans. Therefore, robust bone image enhancement and segmentation methods must be used to achieve acceptable image quality.

The location of a bone in an ultrasound image is normally characterized by different features. Two of these features are a high shadow profile underneath the location of the bone and high intensity across the bone surface, which have been successfully used in a spine surface detection method recently proposed by our group. Other priori information about the bone can also be used collaboratively. For example, Thomas et al. has discussed the possibility to segment a fetus femur using purely morphological operations assuming a priori knowledge about the shape of the bone

anatomy (Thomas, Peters et al. 1991). Bone shape information has also been exploited to utilize active contours for the crux of the segmentation process (He and Zheng 2001, Alfiansyah, Streichenberger et al. 2006). Hacihaliloglu et al. have proposed bone surface localization and registration methods using a local phase tensor and statistical shape models (Hacihaliloglu, Abugharbieh et al. 2009, Hacihaliloglu, Abugharbieh et al. 2011, Hacihaliloglu, Rasouljan et al. 2014). However, given the irregular shape and complex structure of bone fractures, assumptions about bone surface shape may be difficult to make. Connective tissue or other soft tissue structures may manifest in the upper and lower limb sonograms with similar intensity as those at the bone surface, which can create ambiguity in pattern classification. Aside from feature selection, another challenge in bone segmentation in ultrasound lies in the free hand nature of the in vivo experiment, which usually results in images with relatively lower quality. This fact manifests itself in the following several aspects. First, the distinctive shadow pattern beneath the bone surface becomes not clear. Second, uncontrolled motion of the transducer can be a source of refraction artifacts and can affect the acquisition of aligned parallel frames necessary for 3D reconstructions. Third, 3D shape model-based methods used in other medical imaging segmentation applications (such as magnetic resonance imaging) become not applicable in ultrasound imaging because ultrasound usually doesn't give slice position information. Fourth, because of undesired transducer motion, the bone surface shape change necessitates of more information to accurately describe shape variations. Iterative shape model-based methods may be a solution to some of the aforementioned problems but suffer from high computation cost (Cootes, Taylor et al.

1995, Zheng, Barbu et al. 2008), which is a reason why methods such as marginal space learning have been developed (Zheng, Barbu et al. 2008, Kelm, Wels et al. 2013). In ultrasound imaging, however, processing speed is an important factor especially when dealing with intra- operative or other clinical applications. In this study, we propose a machine learning-based method that coordinates both intensity information and statistical shape model to segment bone surfaces in ultrasound images. This method allows to achieve fast processing and high accuracy.

METHOD

In the proposed segmentation technique, we first use phase symmetry-based methods to extract bone surface candidates (Hacihaliloglu, Abugharbieh et al. 2009, Shajudeen and Righetti 2017). By 'candidate', we mean possible areas where bone surface exists. The candidates are manually labeled to create a database of bone surface shapes. Then a statistical shape model is extracted from the database. This model can be used to score each segmentation candidate's possibility of being a bone surface. This score along with shadow intensity and size information is combined as a feature vector to describe each segmentation candidate. These features can be used to train a classifier, which can supervise the segmentation process in other ultrasound images (applicable to stacks of images for 3D bone reconstruction).

Pre-processing

In this step, we use a phase symmetry-based method to extract high intensity profiles in the ultrasound image. These extracted regions, usually with the axial thickness of one single pixel size, are bone segmentation candidates.

To enhance the bone surface contrast and decrease the intensity of non-bone surface tissues, we create a shadow value image (Shajudeen and Righetti 2017) which is tailored to the long bone segmentation application. Its intensity value for each pixel represents the average image intensity among a vertical line segment of a fixed length below that pixel in the original image. In the shadow value image, the region underneath the hyperechoic bone surface is typically hypoechoic since sound waves are mostly reflected by the bone surface and also highly attenuated inside the bone. While across the soft tissue its intensity tends to be in general higher due to significant echogenicity. The pixel intensity in the region close to the bone surface, on the other hand, is very sensitive to the selection of the line segment length parameter used for the local averaging. The pixel value of each component in the mask is calculated as:

$$SH(x, y) = \frac{1}{N} \sum_{i=k}^{k+N} I(x + i, y) \quad (1)$$

where $SH(x, y)$ is termed as the shadow value of a pixel at row x and scan line y . N is the depth of the shadow and k is set to be the thickness of bone surface in the image to get optimal performance. $I(x, y)$ represents the pixel intensity at location (x, y) in the original image. The new image obtained by subtracting the mask image from the original image is referred to as the "shadow-enhanced" (SE) image.

Subsequently, a common blob detection method based on the Laplacian of Gaussian is used to extract the bright part of the SE image. This is done by convolving an image with the Gaussian kernel given below to smooth the image and desensitize it to noise:

$$g(x, y, t) = \frac{1}{2\pi t^2} e^{-\frac{x^2+y^2}{2t^2}} \quad (2)$$

where t is a scaling constant. Then the Laplacian of the smoothed image is calculated to highlight discontinuities in the intensities in the image. In the resulting Laplacian of Gaussian image, strong negative values correspond to high intensity areas in the original image, which have a high probability of being parts of the bone surface. To regularize the image, the positive values are set to zero while the negative values' signs are reversed (Foroughi, Boctor* et al. 2007).

The intensity-based Laplacian of Gaussian image highlights the bright areas in the image. However, it may not maintain the original shape of the bright blob and is sensitive to non-uniform intensities along the bone surface, which are quite common in free hand scans. According to Kovési (Kovési 1997), the phase symmetry can be used to recognize various regions of interest in natural objects, and this information is used in our algorithm to complement the Laplacian of Gaussian image. The phase symmetry is computed as:

$$\text{Sym}(x) = \frac{\sum_n [|e_n(x)| - |o_n(x)|] - T}{\sum_n A_n(x) + \epsilon} \quad (3)$$

$e_n(x)$ and $o_n(x)$ are even symmetry component and odd symmetry component at scale n . In this paper, the two symmetry components are calculated by convolving the SE image with even- and odd-symmetric wavelets at scale n .

$$[e_n(x), o_n(x)] = [I(x) * M_n^e, I(x) * M_n^o] \quad (4)$$

$A_n(x)$ is the amplitude of the transform defined in equation 4. The phase symmetry technique is widely used in feature extraction as it provides mirror symmetry, rotation

symmetry and curve symmetry at the same time (Shajudeen and Righetti 2017). It extracts image features using phase information and, therefore, is more robust to intensity-based methods. Specific to speckle in ultrasound images, however, there is always a certain degree of local symmetry in the tissue and shadow areas. As a consequence, there will also be some remnants in the phase symmetry image from regions that are not bone surface. Since it extracts bone surface steadily and preserves its shape, phase symmetry is used in complement with the Laplacian of Gaussian image to bring the pre-processing step to a conclusion. A summary of the overall preprocessing steps is provided in:

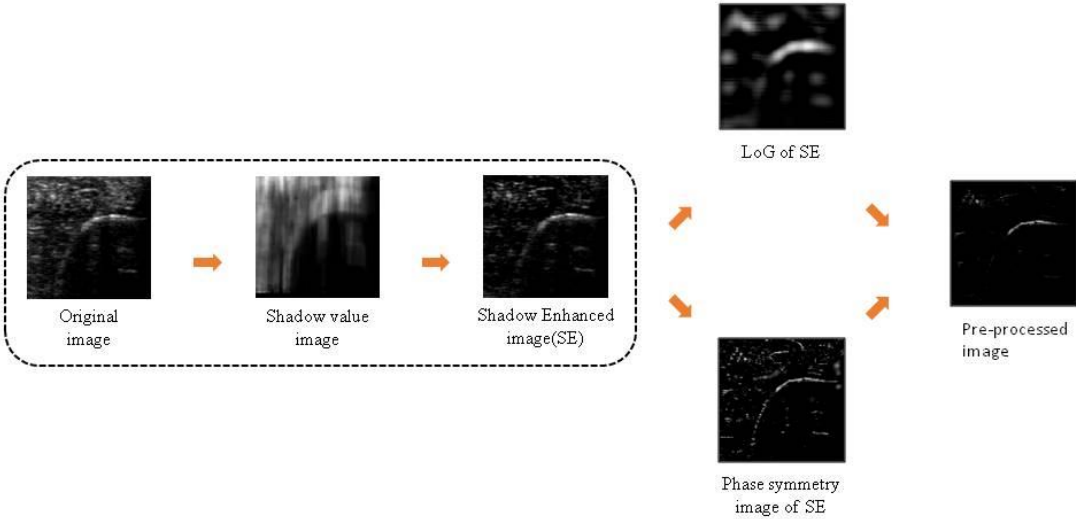


Figure 24 Steps for image preprocessing.

In Figure. 24, it is clear that the average intensity of surrounding tissues is successfully diminished in the SE image. In the phase symmetry image of SE, the bone surface is extracted. However, there is still unwanted symmetry information especially in

the soft tissue area. In LoG of SE, on the other hand, the bone surface is quite blurred and its contour is distorted from the original shape due to non-uniformity illumination along the bone surface. We note that the combination of the two methods removes most normal intensity tissue but keeps most high-intensity bone surface candidates. Hence, the pre-processed image is ready for further segmentation operations.

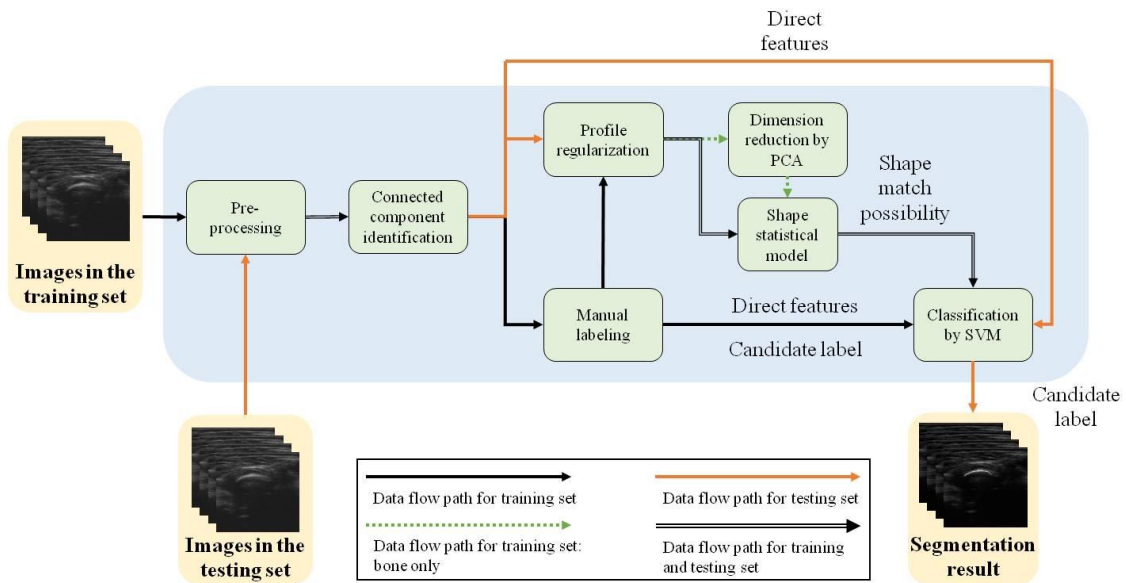


Figure 25 Illustration of complete segment steps.

Create Segmentation Candidates and shape model

The image after the pre-processing is transformed to black and white, and morphological thin operation is used to obtain segmentation candidates, which contains mainly the bone surface and connective tissue interfaces. In addition, bone fracture and some high intensity artifacts may also exist in the current image.

Segmentation candidates are drawn from the candidate library with each connected component manually labeled as the bone surface or otherwise. When a bone

surface is marked, the shape is resampled using 10 points evenly spaced along the horizontal direction by linear interpolation, and then its size is rescaled to normalize the horizontal coordinate range. In this way, variability in the horizontal coordinates is completely controlled for, and, as a result, only the 10 vertical coordinates for each surface candidate are used to represent the shape. The mean shape can be produced by averaging the entire swarm of regularized shapes, and the deviations of each shape from the mean shape are combined to construct the sampled covariance matrix. By performing Principal Component Analysis (PCA) we can describe a curve by:

$$S \approx \tilde{S} = A(\bar{S} + Pb) \quad (5)$$

where $S = (x_1, x_2 \dots x_{10})$, with each component being one of the resampled vertical coordinates, A is a size factor used to restore the original scale of the bone, \bar{S} is the mean shape, $P = (p_1 | p_2 \dots | p_n)$ contains n eigenvectors of the difference matrix, and b is a t -dimensional vector defined by

$$b = P^T(S - \bar{S}) \quad (6)$$

For free-hand ultrasound scans, the bone shape exhibits a lot of variations due to the different incident angles and positions. For one type of imaging target, we create a shape model so that 90% energy. For magnetic resonance or computed tomography images, 3 to 4 eigenvectors could produce a P matrix that contains 90% of energy, but for free-hand ultrasound scans, more eigenvectors need to be used because of more variations caused by imaging position difference. In this case, feature vector will have higher dimensions which can increase computation cost dramatically. To reduce the feature vector dimension, we can use the statistical model to fit a candidate's natural

shape and give a score for the shape match possibility, indicative of the likelihood that one shape can be represented by the model. The score is calculated by the normalized Mean Square Error (MSE) of S and \tilde{S} :

$$P = \frac{1}{10} \sum_{i=1}^{10} \frac{(x_i - \tilde{x}_i)^2}{x_i^2} \quad (7)$$

Intensity and Shape Features

For each bone surface segmentation candidate, we extract a set of local features from the points defining the surface based on intensity, shadow and shape features. A major reason for choosing these features is to make sure that the classifier can be robust to different bone surface conditions (i.e., intact or fractured). The intensity- and shadow-based features can help determine small bone pieces, and shape-based features can improve bone surface segmentation accuracy, especially by ruling out connective tissues. The pooled features are: the average intensity of the sampled segmentation points, the average shadow intensity and its standard deviation, the width and height of the segmented hyperechoic regions, shape match possibility and size factor. These features are combined and used for Support Vector Machine (SVM) (Cortes and Vapnik 1995) training. We use sigmoid kernel for SVM and perform parameter grid search to find the best parameters for each classifier (Claesen and Moor 2015).

Experimental testing

Controlled ex-vivo and in-vivo experiments were used to assess the performance of the proposed technique. All experiments were performed using a Sonix RP diagnostic ultrasound imaging system (Ultrasonix Medical Corp., Richmond, BC, Canada), which has been described previously (Parmar, Yang et al. 2015, Shajudeen and Righetti 2017).

Images were acquired with the center frequency set to 10 MHz or 6.6 MHz depending on the depth of the samples.

Controlled experiment

For the controlled experiments, we used data previously obtained from samples with controlled defects. The overall procedure about the samples and how to obtain the controlled defects has been described previously (Parmar, Longsine et al. 2010). Since in these samples the size of the defects was known from the optical measurements, these experimental data could be used to assess the accuracy of the proposed segmentation method.

Ex-vivo experiment

Ex vivo data were used to validate the proposed method in realistic but more controlled scenarios than in in vivo experiments. Ultrasound data were acquired from the tibia in the lower limb of sheep animals using the system with a center frequency of 10 MHz or 6.6 MHz depending on the depth. Ten sheep tibia experiments were performed. Two-three data sets were acquired from the same leg at different locations. During the acquisition, the transducer was moved free-hand along the direction of the tibia long axis with the imaging plane perpendicular to the bone axis (as typically required by 3D reconstructions). Each volume scan took 2 to 4 seconds, covered a distance of 4-10 cm and contained 100-300 frames. The angle between bone surface and axis of insonification and the transducer-specimen contact varied during the free-hand scan. The presence of connective tissue and soft tissue heterogeneity generated a number of hyperechoic regions in these images in more superficial areas than the bone surface.

In-vivo experiment

In vivo data were collected from 7 live sheep tibias (intact) in free hand mode. In vivo data acquisition was approved by the Houston Methodist Research Institute, Institutional Animal Care and Use Committee (IACUC) (ARO #60598-MS-DRP, Award W911NF-11-1-0266). The same scan protocol was used as in ex vivo experiments. The acquisition depth was set to 4cm, and three focal zones were set with the bone surface situated approximately in the center. B-mode images were saved as AVI videos and then individually extracted for segmentation purposes.

Performance Analysis

The accuracy of the proposed segmentation method was assessed in terms of the pixel error. The average pixel error is defined as the average difference between locations of the segmentation curve and manual delineation curve on the same A-line:

$$E = \frac{1}{N} \sum \|x_i - x'_i\|, (8)$$

where N is the A-line number while x_i and x'_i are the locations of the auto-segmented bone curve and manually segmented bone curve on the i^{th} column of the image. To calculate the average pixel error, 205 images were randomly selected from the image set and manually delineated by an independent ultrasound expert by following the procedures in Foroughi et al. Foroughi et al. (2007). About 70% of the selected images were used for modeling and training while auto-segmentation was applied to the other 30%. The auto-segmentation results were then compared to the ones obtained by the ultrasound expert to calculate the error. The number of false positive and false negative was also evaluated as delineated in Shajudeen and Righetti (2017).

RESULTS

Segmentation of Bone Surfaces from Ex Vivo Samples

The goal of the ex vivo study is to test the performance of the proposed segmentation method in B-mode images with varying image quality but in conditions more controllable by the operator than in in vivo experiments. We note that the signal to noise ratio (SNR) and contrast in these sheep animal data is generally low presumably due to ultrasonic attenuation Parmar et al. (2010). In Figure 3 of this section, we report selected representative ex vivo results from the sheep animal data.




Figure 26 Segmentation results of ex-vivo experiments

In the B-mode images (first column of Figure 3), skin, subcutaneous tissue, muscle and bone are visible. After the segmentation (third column of Figure 3), only the bone surface is extracted, and the boundary between the subcutaneous tissue and muscle, which has similar intensity to that of the boundary of the bone surface, is not present in general. Note that in the second row of Figure 3, the connective tissue above the bone has higher intensity than the bone surface in the original B-mode image. In the segmented image, however, it is not present as the algorithm correctly identifies it as a non-bone surface. From these preliminary results, it appears that the proposed method can differentiate high intensity connective tissues from high intensity bone surfaces, which is a major source of false positives in ultrasonic bone segmentation applications. Also, the method still performs well even when the shadow area underneath the bone surface may not be clearly defined. These experiments also demonstrate that the method can be used in free-hand ultrasound scan mode, which typically results in limited image quality due to the probe's freedom of motion but is essential in applications in vivo.

Segmentation of Bone Surfaces from In Vivo Samples

In this section, we evaluate our method using in vivo data sets. In vivo experiments would allow us to test the performance of the method on a variety of bone shapes. Note that the in vivo images were acquired from a larger scan area than the ex vivo images, and the acquisition was also under more noisy conditions presumably due to the motion of the sheep. Consequently, the bone sonographic profile can have more variations. The display resolution was around 400 as acquired from saved AVI files. As ultrasound acquisition speed was very high, adjacent frames are very similar in contact

to each other. Hence, it is not necessary to incorporate all of them in the analysis. Under this premise, 105 frames were selected to construct the dataset for the training and testing.




Figure 27 Different bone surface segmented using our algorithm.

Figure 27 shows segmentation results obtained from those in vivo sheep experiments. The results demonstrate that, although the shape of the bone varies at different scan locations and the illumination conditions change among animal subjects, the algorithm can adapt to these changes and detect the bone surface with fairly high accuracy.

Statistical results

The results of our statistical analysis are shown in Table 2. From Table 2, we observe that the segmentation results from the ex vivo study is more accurate than the segmentation results from the in vivo study, as expected. In addition, segmentation accuracy is improved in the presence of the shape information. In particular, false positives decreased greatly, in the presence of the shape information.

Table 2 Algorithm detection rate for the two studies

	Avg. Error(pixel)	False Positive	False Negative
Ex-vivo	3.23	17.2%	3.9%
In-vivo	5.16	16.8%	7.8%
Ex-vivo without SSM	5.53	27.2%	5.3%
In-vivo without SSM	7.22	28.4%	8.7%

DISCUSSION

In this paper, we present a method to segment the bone surface in noisy ultrasound images based on SSM. While SSM has been extensively investigated for MRI and CT applications, its suitability for ultrasound bone imaging has not been evaluated yet. The *ex vivo* and *in vivo* results included in this paper show that, when combined with phase symmetry and shadow information, the SSM can significantly improve accuracy of bone surface detection and segmentation.

The SSM-based segmentation method is particularly suitable for freehand ultrasound scans. Common errors arising in parameter estimation-based methods due to variations occurring in a volumetric free-hand scan do not affect this method since the shape information input to the classifier is entirely obtained from a data-driven statistical model. Thus, the proposed segmentation method allows us to automatically segment a large number of images once the classifier has been trained. The method is sensitive to the shadow underneath the bone and performs well even when the intensity contrast between the bone surface and other tissue regions is not high (as in the sheep data used in this study). In those cases, pure intensity-based methods would likely fail. As demonstrated by our results, most soft tissue hyperechoic areas are suppressed in the processed images. In applications aiming at 3D bone reconstructions based on volume rendering, the false positives caused by the soft tissue may be further reduced using thresholding. This method is currently implemented using MATLAB and takes about 2.3 seconds to process each frame with 400×450 resolution on an Intel Core i7-4700 MQ laptop. In order to accelerate computation, GPGPU might be used to parallelize all

computations Yang et al. (2011). By using more efficient computer languages such as C and optimizing the algorithm properly, we would expect this algorithm to achieve at least semi real-time performance on ultrasound imaging systems.

As mentioned previously, the SSM can be used in many ways. One popular method is marginal space learning, which has been applied to many applications for MR and CT imagingZheng et al. (2008); Kelm et al. (2013). However, the complexity of marginal space learning increases exponentially with the number of eigenvectors. Our method does not search for eigenvalue directly to reconstruct curved bone surface patterns in the B-mode image. This obviates the need of a large number of eigenvectors to describe bone shape variations, which is a typical problem that affects free-hand ultrasound imaging. Therefore, in general, the proposed method is time-efficient and adaptive to a variety of experimental conditions.

There are a number of factors that could affect the performance of the proposed method. The first is the quality of the ultrasound images. As demonstrated by our results, the higher is the bone/soft tissue contrast in the ultrasound images and, in general, the better is the performance of the segmentation method. This is true, in general, for most ultrasound bone segmentation methods proposed in the literature. However, a strength of the proposed method with respect to previously proposed intensity-based and phased-based algorithms, is that a good performance (error rate at around 16.8% or below) is in general maintained also in the cases where the quality of the B-mode images is fairly poor. Therefore, in these very difficult cases, the availability of the proposed segmentation methods could be helpful. The second is the features used for the

classification. For example, radiofrequency (RF) data have the potential to provide information regarding reflected power Wen and Salcudean (2007). Ultrasound elastographic features could further help identifying soft tissue structures Parmar et al. (2015); Tang et al. (2017), and reduce influence of connective tissues, which create high reflection in B-mode images. However acquiring 3D elastography data is challenging and not easy for wide area free hand scans using regular probes Gennisson et al. (2013); Lindop et al. (2006). In addition, an increase in the number of features used for the classification will, in general, lead to more complex and less computationally efficient algorithms.

In this paper, we have studied the performance of the algorithm only using long bones from a large animal model (sheep). Naturally, it is expected that the performance of the method varies depending on the bone that is imaged. Based on our past experience, bone surface segmentation from humans typically yield better results than those obtained from large animal models.

Furthermore, we have not studied the performance of the algorithm when a bone fracture is present. In this case, the method may need to be modified. It may require generation of a cascading classifier. This classifier would first determine if there are fractures in an image, and then apply models of the corresponding category (i.e., intact or fractured) to segment the bone surface. The assessment of the proposed method in the presence of fractures is left for future work.

CONCLUSION

In this paper a new bone segmentation method is proposed. This method is designed to work robustly for in vivo free hand scans. Shadow effects are used firstly to enhance contrast between soft tissue and bone surface. Then, phase property is used to extract patterns in the images. Finally, intensity property and SSM are combined to describe each segmentation candidate. A trained SVM classifier is then used to classify the bone surface according to the features extracted. Based on ex vivo and in vivo experimental data, the proposed method may be used in clinical bone imaging applications.

5 CONCLUSIONS

SUMMARY

Ultrasound image has a lot of potentials in orthopedics applications. This dissertation studies some fundamental aspects in 3D Ultrasound reconstruction in orthopedics. Real-time high quality strain elastography opens possibility to acquire 3D US elastography. New way to process Doppler signal can improve contrast between bone surface and connective tissue, and can eventually reduce the computation cost for bone segmentation. Shape model based segment method offers an alternative before previous image enhancement method can be implemented in real-time.

Section 2 discussed keys and limitations in using CUDA for strain elastography computation, and proposed a hybrid method which take advantage of both CPU and GPGPU to achieve the goal of success. Simulation data and real data are used for performance analysis. This implementation is finally integrated with Sonix US system to display real-time strain elastography on screen.

The method used in section 3 is inspired by strain elastography but applied in US Doppler image. Using in-vivo and in-vitro experiment, the feasibility of using pre- and post-compression Doppler signal to increase contrast between tissue-bone surface and connective tissue surface is verified. The enhancement can eventually simplify bone segmentation method used for surface reconstruction and give faster reconstruction results.

In section 4 a segment method is discussed using regular US image. Images are manually segmented and a statistical shape model is generated from these models. For

surface detection, all curves in US image is extracted and the coefficients fit to this model is calculated. If the coefficients are outliers, the curve is considered not a bone surface. The method is tested using in-vivo data from long bones as they are most concerned in many clinical settings.

FUTURE WORK

Based on the studies in this dissertation, a series of future research plan can be considered.

3D Elastography

3D elastography is a very important research as the possibility to have elastography of a entire volume can greatly extend its usability in clinical investigation. In section 2 the software to generate real-time 2D strain elastography is ready to use and integrated with commercial US imaging system. The follow up study can include the following part:

1. Use electro-magnetic position sensor to record position of each image. Electro-magnetic position sensor has been used on commercial US 3D acquisition systems to generate 3D tissue volumes. For this study, it would be interesting to see if it is possible to record actual strain as well for pre- and post-compression RF frames. As strain elastography is qualitative, recording the actual strain might be very important to set stretch factor in elastography algorithm and make images acquired from different plans more comparable. If actual strain is too small to be detected, adaptive stretch method should be used to estimate real strain and get best image contrast.

2. Develop a proper normalization algorithm to combine strain elastograph acquired at different positions. Although strain elastograph cannot be compared in quantitative way, adjacent images may be able to normalized properly to form a 3D volume considering the fact that tissues are continuous. Elastography acquired from nearby positions with similar strain should be comparable and normalized.

3. To get best result, mechanical apparatus should be developed to control strain applied. It can also get highly precise image position information as well.

4. Choose an organ to have human studies. In such study, the key is to compare elastography detected suspicion region and other modality detected suspicion region. This study would provide more comprehensive understanding about the correlation between elastography and CT/MRI. It may help decide the scope of application for strain elastography in clinic, and provide patients with more affordable and accessible service in the future.

As mentioned before, current shear wave based quantitative elastography needs a couple seconds to acquire one frame. This makes continuous acquisition very hard for 3D elastography. If 3D strain elastography can be achieved with reasonable position attachments, it would give strain elastography another important reason to stay on stage.

Current study in section 2 offered only axial strain image. It would be easy to modify the algorithm to get axial shear strain image as well for other related research projects.

3D Bone surface reconstruction

As described before, this dissertation is aimed to create 3D reconstruction of bone surface in multiple US image modalities. Though the purpose is not achieved, the tools to achieve the purpose are prepared.

With position sensor attached to US probe, regular 2D US scan can be combined to form a 3D volume. Bone surface could be extracted from the volume. Possible research topic include:

1. Repeatability study. With freehand scan, it is important to know how repeatable the 3D reconstructions are. Between different scans probe can have different compression on tissue and different velocity at each position. How to eliminate the effect of these factors and get repeatable result is very important for follow up studies. If there are inevitable reconstruction errors, it is also important to know the range.

2. 3D volume registration. Unlike CT or MR, US images object from one direction. This means if a comprehensive 3D volume is required, different US scan volumes from different directions need to be registered and combined together. While soft tissue can deform during scan, bone surface is a better reference for volume registration.

3. 3D bone fracture identification. As 2D US image is just an image of one cross section of tissue volume, in many cases it is hard to determine if one surface is fracture or not in one image or one direction. With 3D technique, algorithm can have more characteristics to determine if one object is fracture or not.

REFERENCES

- Agarwal, R., P. Gosain, J. N. Kirkpatrick, T. Alyousef, R. Doukky, G. Singh and C. A. Umscheid (2012). "Tissue Doppler imaging for diagnosis of coronary artery disease: a systematic review and meta-analysis." Cardiovascular Ultrasound **10**(47).
- Al Naimi, A., M. Fittschen and F. Bahlmann (2014). "Measuring cervical strain with tissue Doppler imaging depending on the shape and placement of the region of interest and its correlation with cervical consistency index." European Journal of Obstetrics & Gynecology and Reproductive Biology **179**: 246-250.
- Alfiansyah, A., R. Streichenberger, M. Bellemare and O. Coulon (2006). Automatic segmentation of hip bone surface in ultrasound images using an active contour. IFBME.
- Allen, G. M. and D. J. Wilson (1999). "Ultrasound and the diagnosis of orthopaedic disorders." J Bone Joint Surg Br **81**(6): 944-951.
- ATI (2010). "ATI Stream Computing OpenCL Programming Guide."
- Barratt, D. C., A. H. Davies, A. D. Hughes, S. A. Thom and K. N. Humphries (2001). "Optimisation and evaluation of an electromagnetic tracking device for high-accuracy three-dimensional ultrasound imaging of the carotid arteries." Ultrasound Med Biol **27**(7): 957-968.
- Barratt, D. C., G. P. Penney, C. S. Chan, M. Slomczykowski, T. J. Carter, P. J. Edwards and D. J. Hawkes (2006). "Self-calibrating 3D-ultrasound-based bone registration for minimally invasive orthopedic surgery." IEEE transactions on medical imaging **25**(3): 312-323.
- Beek, M., P. Abolmaesumi, S. Luenam, R. Ellis, R. Sellens and D. Pichora (2008). "Validation of a new surgical procedure for percutaneous scaphoid fixation using intra-operative ultrasound." Medical image analysis **12**(2): 152-162.
- Blankstein, A. (2011). "Ultrasound in the diagnosis of clinical orthopedics: The orthopedic stethoscope." World J Orthop **2**(2): 13-24.
- Bude, R. O. and J. M. Rubin (1996). "Power Doppler sonography." Radiology **200**(1): 21-23.
- Carson, P. L., T. V. Oughton, W. R. Hendee and A. S. Ahuja (1977). "Imaging soft tissue through bone with ultrasound transmission tomography by reconstruction." Medical physics **4**(4): 302-309.
- Cespedes, I., Y. Huand, J. Ophir and S. Spratt (1995). "Methods for Estimation of Subsample Time Delays of Digitized Echo Signals." Ultrasonic Imaging **17**: 142-171.

- Chen, L., Y. Kim and C. Moore (2007). "Diagnosis and guided reduction of forearm fractures in children using bedside ultrasound." Pediatric emergency care **23**(8): 528-531.
- Chen, Y., Y. Lin, S. Wang, S. Lin, K. Shung and C. Wu (2014). "Monitoring tissue inflammation and responses to drug treatments in early stages of mice bone fracture using 50 MHz ultrasound." Ultrasonics **54**(1): 177-186.
- Cho, S. H., J. Y. Lee, J. K. Han and B. I. Choi (2010). "Acoustic radiation force impulse elastography for the evaluation of focal solid hepatic lesions: preliminary findings." Ultrasound Med Biol **36**(2): 202-208.
- Claesen, M. and B. D. Moor (2015). Hyperparameter Search in Machine Learning. 11th Metaheuristics International Conference.
- Cooley, J. W. and J. W. Tukey (1965). "An algorithm for the machine calculation of complex Fourier series." Math. Comput. **19**: 297-301.
- Cootes, T. F., C. J. Taylor, D. H. Cooper and J. Graham (1995). "Active Shape Models-Their Training and Application." Computer Vision and Image Understanding **61**(1): 38-59.
- Cortes, C. and V. N. Vapnik (1995). "Support-vector networks." Machine Learning **20**(3): 273-297.
- Culjat, M. O., M. Choi, R. S. Singh, W. S. Grundfest, E. R. Brown and S. N. White (2008). "Ultrasound detection of submerged dental implants through soft tissue in a porcine model." The Journal of prosthetic dentistry **99**(3): 218-224.
- Dai, Y. (2010). "Real-time visualized freehand 3D ultrasound reconstruction based on GPU." IEEE Transaction on Information Technology in Biomedicine **14**(6): 1338-1345.
- de Jesus, J. O., L. Parker, A. J. Frangos and L. N. Nazarian (2009). "Accuracy of MRI, MR arthrography, and ultrasound in the diagnosis of rotator cuff tears: a meta-analysis." AJR Am J Roentgenol **192**(6): 1701-1707.
- Deka, S., X. Yang and R. Righetti (2009). Elastography on GPGPU: A feasibility study. The 26th Annual Houston Conference on Biomedical Engineering Research (HSEMB), Houston, TX.
- Desai, R. R., T. A. Krouskop and R. Righetti (2010). "Elastography using harmonic ultrasonic imaging." Ultrasonic Imaging **32**: 103-117.

Doctor, A., B. Vondenbusch and J. Kozak (2011). "Bone segmentation applying rigid bone position and triple shadow check method based on RF data." Acta of Bioengineering and Biomechanics **13**(2): 3-11.

E. E. Konofagou, T. V., J. Ophir, and S. K. Alam (September 1999). "Power spectral strain estimators in elastography." **25**(7): 1115-1129.

Fayaz, H. C., P. V. Giannoudi, M. S. Vrahas, R. M. Smith, C. Moran, H. C. Pape, C. Krettek and J. B. Jupiter (2011). "The role of stem cells in fracture healing and nonunion." Int Orthop. **35**(11): 1587-1597.

Foroughi, P., E. Boctor*, M. J. Swartz, R. H. Taylor* and G. Fichtinger (2007). Ultrasound Bone Segmentation Using Dynamic Programming. 2007 IEEE Ultrasonics Symposium: 2523-2526.

Franz, A. M., K. Marz, A. Seitel, H. G. Kenngott, M. Wagner, A. Preukschas, H. P. Meinzer, I. Wolf and L. Maier-Hein (2013). "Combined Modality for Ultrasound Imaging and Electromagnetic Tracking." Biomed Tech (Berl).

Garra, B., I. Cespedes, J. Ophir, R. Zuurbier, C. Magnant and M. Pennanen (1997). "Analysis of Breast Lesions Using Elastography: Initial Clinical Results." Radiology **202**(79-86).

Gennisson, J.-L., T. Deffieux, M. Fink and M. Tanter (2013). "Ultrasound elastography: Principles and techniques." Diagnostic and Interventional Imaging **94**(5): 487-495.

Giannetti, R., A. Petrella, J. Bach, A. K. Silverman and M. A. Sáenz-Nuño (2017). "In vivo Bone Position Measurement Using High-Frequency Ultrasound Validated with 3-D Optical Motion Capture Systems: A Feasibility Study." Journal of Medical and Biological Engineering **37**(4).

Gluer, C. C. and R. Barkmann (2003). "Quantitative ultrasound: use in the detection of fractures and in the assessment of bone composition." Curr Osteoporos Rep **1**(3): 98-104.

Gnudi, S., N. Malavolta, P. Calderoni, G. Bettelli and G. Gualtieri (1996). "Ultrasound in the evaluation of bone fragility caused by osteoporosis: a comparison between different sites of measurement." Chir Organi Mov **81**(4): 383-387.

Gorcsan, J. r. and H. Tanaka (2011). "Echocardiographic assessment of myocardial strain." J Am Coll Cardiol **2011** **58**(14): 1401-1403.

Hacihaliloglu, I., R. Abugharbieh, A. Hodgson and R. Rohling (2008). "Bone segmentation and fracture detection in ultrasound using 3D local phase features." Med Image Comput Assist Interv **11**: 287-295.

Hacihaliloglu, I., R. Abugharbieh, A. J. Hodgson and R. N. Rohling (2009). "Bone surface localization in ultrasound using image phase based features." Ultrasound in Medicine and Biology **35**(9): 1475-1487.

Hacihaliloglu, I., R. Abugharbieh, A. J. Hodgson and R. N. Rohling (2011). "Automatic adaptive parameterization in local phase feature-based bone segmentation in ultrasound." Ultrasound in Medicine and Biology **37**(10): 1689-1703.

Hacihaliloglu, I., A. Rasoulian, R. Rohling and P. Abolmaesumi (2014). "Local phase tensor features for 3-D ultrasound to statistical shape+pose spine model registration." IEEE Transactions on Medical Imaging **33**(11): 2167-2179.

Hagiwara, Y., Y. Saijo, A. Ando, K. Kobayashi, A. Tanaka, N. Hozumi, K. Hatori and E. Itoi (2011). "High Frequency Ultrasound Imaging of Cartilage-Bone Complex." Acoustical Imaging **30**: 119-124.

Hall, T., Y. Zhu and C. Spalding (2003). "In vivo real-time freehand palpation imaging." Ultrasound in Medicine and Biology **427-435**.

Harris, M. (2007) "Optimizing parallel reduction in cuda."

He, P. and J. Zheng (2001). Segmentation of tibia bone in ultrasound images using active shape models. IEEE/EMBS, Istanbul, Turkey.

Heaney, R. P., L. V. Avioli, C. H. Chesnut, 3rd, J. Lappe, R. R. Recker and G. H. Brandenburger (1989). "Osteoporotic bone fragility. Detection by ultrasound transmission velocity." JAMA **261**(20): 2986-2990.

Hijazy, A., H. Al-Smoudi, M. Swedan, N. Qaddoum, H. Al-Nashash and K. Ramesh (2006). "Quantitative monitoring of bone healing process using ultrasound." Journal of the Franklin Institute **343**(4): 495-500.

Hill, C. and G. terHaar (1995). "Review article: High intensity focused ultrasound-potential for cancer treatment " British Journal of Radiology **68** (816): 1296-1303

Hoyt, K., F. Forsberg and J. Ophir (2006). "Comparison of shift estimation strategies in spectral elastography." Ultrasonics **44**: 99-108.

Jensen, J. A. (1996). Estimation of Blood Velocities Using Ultrasound, Cambridge University Press, Cambridge, UK.

- Jia, X., Y. Lou and R. Li (2010). "GPU-based fast cone beam CT reconstruction from undersampled and noisy projection data via total variation " Medical Physics **37**(4): 1757-1760.
- Kallel, F., R. J. Stafford, R. E. Price, R. Righetti, J. Ophir and J. D. Hazle (1999). "The Feasibility of Elastographic Visualization of HIFU-Induced Thermal Lesions in Soft-Tissue." Ultrasound in Medicine and Biology **25**: 641-647.
- Kasai, C., K. Namekawa, A. Koyano and R. Omoto (1985). "Real-time two-dimensional blood flow imaging using an autocorrelation technique." IEEE Trans Sonics Ultrason. **32**: 458-464.
- Kelm, B., M. Wels, S. Zhou, S. Seifert, M. Suehling, Y. Zheng and D. Comaniciu (2013). "Spine detection in CT and MR using iterated marginal space learning." Medical image analysis **17**(8): 1283-1292.
- Konofagou, E. E., T. P. Harrigan, J. Ophir and T. A. Krouskop (2001). "Poroelastography: imaging the poroelastic properties of tissues." Ultrasound in Medicine and Biology **27** (10): 1387-1397.
- Konofagou, E. E. and J. Ophir (1998). "A new elastographic method for estimation and imaging of lateral displacements, lateral strains, corrected axial strains and Poisson's ratios in tissues. ." Ultrasound in Medicine and Biology **24**(8): 1183-1199.
- Konofagou, E. E., T. Varghese, J. Ophir and S. K. Alam (September 1999). "Power spectral strain estimators in elastography." Ultrasound in Medicine and Biology **25**(7): 1115-1129.
- Kovesi, P. (1997). Symmetry and Asymmetry from Local Phase. AI'97, Tenth Australian Joint Conference on Artificial Intelligence.
- Kowal, J., C. Amstutz, F. Langlotz, H. Talib and M. G. Ballester (2007). "Automated bone contour detection in ultrasound B-mode images for minimally invasive registration in computer-assisted surgery—an in vitro evaluation." The International Journal of Medical Robotics and Computer Assisted Surgery **3**(4): 341–348.
- Kruskal, J. B., P. A. Newman, L. G. Sammons and R. A. Kane (May 2004). "Optimizing Doppler and Color Flow US: Application to Hepatic Sonography." Radio Graphics **24**(3).

Lasaygues, P., E. Ouedraogo, J.-P. Lefebvre, M. Gindre, M. Talmant and P. Laugier (2005). "Progress towards in vitro quantitative imaging of human femur using compound quantitative ultrasonic tomography." Physics in medicine and biology **50**(11): 2633.

Laugier, P. and G. Haët (2011). Bone quantitative ultrasound, Springer.

Li, H., L. Le, M. Sacchi and E. Lou (2013). "Ultrasound imaging of long bone fractures and healing with the split-step fourier imaging method." Ultrasound in Medicine and Biology **39**(8): 1482-1490.

Lindop, J. E., G. M. Treece, A. H. Gee and R. W. Prager (2006). "3D Elastography using freehand ultrasound." **32**: 529-545.

Liu, D. (2010). "Real-Time 2-D Temperature Imaging Using Ultrasound." IEEE Trans on Biomedical Engineering **57**(1): 12-16.

Luebke, D. (2008). CUDA: Scalable parallel programming for high-performance scientific computing. 5th IEEE International Symposium on Biomedical Imaging - From Nano to Macro. Paris, FRANCE. **1-4**: 836-838.

Madersbacher, S., C. Kratzik and M. Susani (1995). "High-intensity focused ultrasound as a minimally invasive treatment for benign prostatic hyperplasia." Urologe-Ausgabe A **34**(2): 98-104.

Namekawa, K., C. Kasai, M. Tsukamoto and A. Koyano (1983). "Realtime bloodflow imaging system utilizing auto-correlation techniques." Ultrasound in Medicine and Biology **2**: 203-208.

NVIDIA (2010) "NVIDIA CUDA Compute Unified Device Architecture Programming Guide, Version 3.0."

Ophir, J. (1991). "Elastography: a quantitative method for imaging the elasticity of biological tissues." Ultrasonic Imaging **13**(2):111-34.

Owens, J. D., D. Luebke, N. Govindaraju, M. Harris, J. Kruger, A. E. Lefohn and T. J. Purcell (2007). "A Survey of General-Purpose Computation on Graphics Hardware." Computer Graphics **26**: 80–113.

Parmar, B. J., W. Longsine, E. P. Sabonghy, A. Han, E. Tasciotti, B. K. Weiner, M. Ferrari and R. Righetti (2010). "Characterization of controlled bone defects using 2D and 3D ultrasound imaging techniques." Physics in medicine and biology **55**: 4839-4859.

Parmar, B. J., X. Yang, A. Chaudhry, P. S. Shajudeen, S. P. Nair, B. K. Weiner, E. Tasciotti, T. A. Krouskop and R. Righetti (2015). "Ultrasound elastography assessment of bone/soft tissue interface." Physics in medicine and biology **61**(1): 131.

Poonai, N., F. Myslik, G. Joubert, J. Fan, A. Misir, V. Istasy, M. Columbus, R. Soegtrop, A. Goldfarb and D. Thompson (2017). "Point-of-care ultrasound for nonangulated distal forearm fractures in children: Test performance characteristics and patient-centered outcomes." Academic Emergency Medicine **24**(5): 607-617.

Protopappas, V., M. Vavva, D. Fotiadis and K. Malizos (2008). "Ultrasonic monitoring of bone fracture healing." IEEE Trans Ultrason. Ferroelectr. Freq. Control **55**(6): 1243-1255.

Ricciardi, L., A. Perissinotto and M. Dabala (1993). "Mechanical monitoring of fracture healing using ultrasound imaging " Clinical orthopaedics and related research **293**: 71-76.

Righetti, R., F. Kallel, R. J. Stafford, R. E. Price, T. A. Krouskop, H. J.D. and J. Ophir (1999). "Elastographic Characterization of HIFU-induced lesions in canine livers." Ultrasound in Medicine and Biology **25**(7): 1099-1113.

Righetti, R., J. Ophir and T. Krouskop (2005). "A method for generating permeability elastograms and poisson's ratio time-constant elastograms." Ultrasound in Medicine and Biology **31**: 803-816.

Righetti, R., J. Ophir and P. Ktonas (2002). "Axial Resolution in Elastography." Ultrasound in Medicine and Biology **28**: 101-113.

Rosenzweig, S. (2011). "GPU-based real-time small displacement estimation with ultrasound." IEEE Trans Ultrason. Ferroelectr. Freq. Control **58**(2): 399-405.

Rubens, D. J., S. Bhatt, S. Nedelka and J. Cullinan (Jan. 2006). "Doppler Artifacts and Pitfalls." Ultrasound Clinics **1**(1).

Samant, S., J. Xia and P. Muyan-Ozcelilk (2008). "High performance computing for deformable image registration: Towards a new paradigm in adaptive radiotherapy " Medical Physics **35**(8).

Sambasubramanian, S. (August 2010). Comparison of the performance of different time delay estimation techniques for ultrasound elastography. MS, Texas A&M University.

Schaar, J. A., C. L. d. Korte, F. Mastik, S. Chaylendra, G. Pasterkamp, E. Boersma, P. W. Serruys and A. F. W. v. d. Steen (2003). "Characterizing vulnerable plaque features with intravascular elastography." Circulation **108**: 2636-2641.

- Sebag, F., J. Vaillant-Lombard, J. Berbis, V. Griset, J. F. Henry, P. Petit and C. Oliver (2010). "Shear wave elastography: a new ultrasound imaging mode for the differential diagnosis of benign and malignant thyroid nodules." J Clin Endocrinol Metab **95**(12): 5281-5288.
- Sehgal, C., D. Lewallen, J. Nicholson, R. A. Robb and J. F. Greenleaf (1988). Ultrasound transmission and reflection computerized tomography for imaging bones and adjoining soft tissues. Ultrasonics Symposium, 1988. Proceedings., IEEE 1988, IEEE.
- Shajudeen, P. M. S. and R. Righetti (2017). "Spine surface detection from local phase-symmetry enhanced ridges in ultrasound images." Medical Physics **44**(11): 5755-5767.
- Slane, L. C., J. A. Slane and L. Scheys (2017). "The measurement of medial knee gap width using ultrasound." Arch Orthop Trauma Surg **137**(8): 1121-1128.
- Smiseth, O., A. Stoylen and H. Ihlen (2004). "Tissue Doppler imaging for the diagnosis of coronary artery disease." Curr Opin Cardiol **19**(5): 421-429.
- Souchon, R., O. Rouviere, A. Gelet, V. Detti, S. Srinivasan, J. Ophir and J. Y. Chapelon (2003). "Visualisation of HIFU lesions using elastography of human prostate in vivo: preliminary results." Ultrasound in Medicine and Biology **29**: 1007-1015.
- Srinivasan, S., T. Krouskop and J. Ophir (2004). "Comparing elastographic strain images with modulus images obtained using nanoindentation: preliminary results using phantoms and tissue samples." Ultrasound in Medicine and Biology **30**(3): 329-343.
- Srinivasan, S. and J. Ophir (2003). "A zero-crossing strain estimator in elastography." Ultrasound in Medicine and Biology **29**: 227-238.
- Srinivasan, S., J. Ophir and S. K. Alam (2002). "Elastographic Imaging Using Staggered Strain Estimates." Ultrasonic Imaging **25**: 229-245.
- T, V. and O. J (1997). "A theoretical framework for performance characterization of elastography: the strain filter." IEEE Trans Ultrason. Ferroelectr. Freq. Control **44**(1): 164-172.
- Tang, S., A. Chaudhry, N. Kim, J. Reddy and R. Righetti (2017). "Effect of bone-soft tissue friction on ultrasound axial shear strain elastography." Physics in medicine and biology **62**(15): 6074.
- Teefey, S. A., B. Petersen and H. Prather (2009). "Shoulder Ultrasound vs MRI for rotator cuff pathology." PM&R **1**(5): 490-495.

- Thomas, J., R. Peters and P. Jeanty (1991). "Automatic segmentation of ultrasound images using morphological operators." IEEE Trans Med Imaging **10**(2): 180-186.
- Varghese, T. and J. Ophir (1996). "Noise reduction in elastograms using temporal stretching with multicompression averaging." Ultrasound in Medicine and Biology **22**: 1043-1052.
- Varghese, T. and J. Ophir (1998). "An analysis of elastographic contrast-to-noise ratio performance." Ultrasound in Medicine and Biology **24**: 915-924.
- Varghese, T., J. Ophir, E. E. Konofagou, F. Kallel and R. Righetti (2003). "Trade-offs between the axial resolution and the signal-to-noise ratio in elastography." Ultrasound in Medicine and Biology **29**: 847-866.
- Vincent, D., J. Tonetti and J. Troccaz (2004). Automatic delineation of the osseous interface in ultrasound images by information fusion. Proceedings of the Seventh International Conference on Information Fusion.
- Viola, F. and W. Walker (2003). "A comparison of the performance of time-delay estimators in medical ultrasound." IEEE Trans Ultrason. Ferroelectr. Freq. Control **50**: 392-401.
- Wang, M. Y., X. B. Wang, X. H. Sun, F. L. Liu and S. C. Huang (2016). "Diagnostic value of high-frequency ultrasound and magnetic resonance imaging in early rheumatoid arthritis." Exp Ther Med **12**(5): 3035-3040.
- Weiss, D. B., J. A. Jacobson and M. A. Karunakar (2005). "The use of ultrasound in evaluating orthopaedic trauma patients." J Am Acad Orthop Surg **13**(8): 525-533.
- Wen, X. and S. E. Salcudean (2007). Enhancement of Bone Surface Visualization Using Ultrasound Radio-frequency Signals. IEEE Ultrasonics Symposium: 2535-2538.
- Xu, F. and K. Mueller (2005). "Accelerating popular tomographic reconstruction algorithms on commodity PC graphics hardware." IEEE Trans Nuclear Science **52**(3): 654-663
- Yu, S., K. K. Tan, B. L. Sng, S. Li and A. T. Sia (2014). "Feature extraction and classification for ultrasound images of lumbar spine with support vector machine." Conf Proc IEEE Eng Med Biol Soc **2014**: 4659-4662.
- Zahiri-Azar, R. and S. Salcudean (2008). Time Delay Estimation of Discrete Samples in Ultrasound Echo Signals. 2008 30TH Annual International Conference of the IEEE Engineering in Medicine and Biology Society. **1-8**: 2976-2979.

Zahiri-Azar, R. and S. E. Salcudean (2006). "Motion Estimation in Ultrasound Images Using Time Domain Cross Correlation With Prior Estimates." IEEE Transactions on Biomedical Engineering **53**: 1990 - 2000

Zheng, R., L. H. Le, M. D. Sacchi, D. Ta and E. Lou (2007). "Spectral ratio method to estimate broadband ultrasound attenuation of cortical bones in vitro using multiple reflections." Physics in medicine and biology **52**(19): 5855.

Zheng, Y., A. Barbu, B. Georgescu, M. Scheuering and D. Comaniciu (2008). "Four-chamber heart modeling and automatic segmentation for 3-D cardiac CT volumes using marginal space learning and steerable features." IEEE Trans Med Imaging **27**(11): 1668-1681.