Advanced ultrasound shear wave elastography methods

A thesis submitted for the
PhD Program in Technologies for Health and Well-being by

Enrique Manuel González Mateo

Instituto de Instrumentación para Imagen Molecular

Universitat Politécnica de Valencia



May 26, 2026

Supervisor

Noé Jimenez González Consejo Superior de Investigaciones Científicas, Spain

Science, my lad, is made up of mistakes, but they are mistakes which it is useful to make, because they lead little by little to the truth.

Jules Verne, Journey to the Center of the Earth

La ciencia, hijo mío, está hecha de errores; pero de errores que conviene cometer, porque conducen poco a poco a la verdad.

Julio Verne, Viaje al centro de la Tierra

Acknowledgements

When I finished the Telecommunications degree in Alicante, I was absolutely sure that I did not want to go into research, until one day my professor Alejandro Cebrecos suggested that I apply for some research fellowships at the Ultrasound Medical and Industrial Laboratory (UMIL). Thank you, Jandri, for noticing me and for seeing abilities in me that I was not even aware of myself. Once there, I later had the opportunity to start working and doing research with Noé Jiménez, a person I deeply admire both scientifically and personally. What struck me most was the way both Jandri and Noé, who at first may seem a little crazy, live science and manage to convey the passion and joy that come with devoting yourself to it. Had I not met them, I very much doubt that I would have started my path in research.

Shortly after joining UMIL, Noé came up with the crazy idea of starting a PhD, and, truth be told, he managed to convince me quite easily. It has been quite a challenge, because this line of research was built almost from scratch in the laboratory. I am deeply grateful to UMIL, and to Noé and Francisco Camarena, for the trust they placed in me and for always giving me the opportunity to continue at the research group and carry out this doctoral thesis.

This thesis has not been a solitary journey. Since I began it, our research group has grown a great deal, and I have been lucky enough to share this journey with many people. Thanks to all of them, this experience has been much easier to carry through. To Irene, for being my partner in academic delusions and the voice of reason at all times; this journey has been far more fun with you by my side. To Josep, for always listening to me. To Javi, for being my partner in musical nerdiness. To Alicia, for her wise advice. And to the daycare crew: Alba, Juanjo, Sergio, Víctor V., Víctor C., Nico, Samu, María, Suso, Adrián, and Jose Luis.

Chapters 5 and 6 of this thesis are the result of a research stay at Mayo Clinic in the United States. I would like to thank Matthew Urban for giving me the opportunity to work with him, a scientist I admire for his career in elastography. It was an incredible opportunity to work at a hospital of Mayo Clinic's calibre and to put into practice, in a clinical environment, many of the ideas and skills developed throughout this thesis. I would also like to thank Iris for making me feel at home during this stay, even while being more than 7,000 km away, and all the wonderful people I met along the way who

welcomed me as if I were one of the family: Claudia, Manolo, Kalvin, Alan, Gaby, and Andrés.

Comenzar esta tesis doctoral también supuso iniciar una nueva etapa en Valencia, lejos de mi ciudad, Murcia. Después de más de cinco años aquí, me llevo en el corazón a muchas personas que me han acompañado, directa o indirectamente, en los caminos valencianos: María, Jorge, Irene, Asun, Carlos, Alejandro, Paula, Marina, Paula G. y muchos más. Gracias a mi tío postizo Antonio Robles por ser mi faro en las ocasiones en las que he estado perdido. Además, uno de mis apoyos fundamentales durante todo este tiempo ha sido la música, y toda la gente maravillosa que he conocido gracias a ella, con la que he podido tocar en directo canciones que me encantan. Gracias Pedro, Mario, Nela y Álvaro.

Y por último, pero no menos importante, agradezco enormemente también a mi familia, y en especial a mis padres. Gracias por vuestra paciencia, por vuestra confianza y por estar siempre en cada etapa, celebrando lo bueno y sosteniéndome en lo difícil. Buena parte de lo que soy, y también de haber llegado hasta aquí, os lo debo a vosotros.

Abstract

Ultrasound shear-wave elastography (SWE) is a quantitative imaging modality for estimating the mechanical properties of a soft tissue by producing, observing and analyzing shear waves propagating through it. Since tissue stiffness is altered by many pathologies, SWE provides clinically relevant information beyond conventional B-mode ultrasound imaging. However, reliable mechanical estimation remains challenging because SWE performance depends on the combined effects of wave generation, motion tracking, and image reconstruction methods under attenuation, heterogeneity, noise, and acoustic output constraints. This thesis develops advanced SWE methods to improve image accuracy, computational efficiency, and shear-wave detectability under such practical limitations.

Its main contributions are fourfold. An efficient computational framework is first developed to model the complete SWE process from acoustic excitation to motion measurement, enabling controlled validation and method development. Second, a fast and robust spectral phase-velocity imaging method is developed by combining phase gradients with angular filtering, providing quantitative maps with real-time capability. Next, acoustic vortex elastography is proposed as a spatially structured excitation strategy for improving shear-wave detectability through torsional motion. Finally, temporally coded excitation based on maximum-length sequences (MLS) and pulse compression is introduced to enhance signal-to-noise ratio without increasing peak pressure, and is further assessed in challenging media such as heterogeneous, anisotropic, viscoelastic, noisy, moving, and *in vivo* tissues. Overall, the thesis shows that coordinated advances in modeling, excitation design, and reconstruction can extend the applicability of ultrasound SWE to challenging scenarios in which its use has traditionally been limited.

Resumen

La elastografía por ondas de cizallamiento (SWE, *shear-wave elastography*) es una técnica de imagen cuantitativa que permite estimar las propiedades mecánicas de un tejido blando mediante la generación, observación y análisis de las ondas de cizallamiento que se propagan a través de él. Dado que la rigidez de los tejidos se ve alterada por numerosas patologías, la SWE proporciona información clínicamente relevante que va más allá de las imágenes ecográficas convencionales en modo B. Sin embargo, la estimación mecánica fiable sigue siendo un reto, ya que el rendimiento de la SWE depende de los efectos combinados de la generación de ondas, el seguimiento del movimiento y los métodos de reconstrucción de imágenes bajo restricciones de atenuación, heterogeneidad, ruido y potencia acústica. Esta tesis desarrolla métodos avanzados de SWE para mejorar la precisión de la imagen, la eficiencia computacional y la detectabilidad de las ondas de cizallamiento bajo tales limitaciones prácticas.

Sus principales contribuciones son cuatro. En primer lugar, se desarrolla un marco computacional eficiente para modelar el proceso completo de la SWE, desde la excitación acústica hasta la medición del movimiento, lo que permite una validación controlada y el desarrollo de nuevos métodos. En segundo lugar, se desarrolla un método de imagen de velocidad de fase espectral rápido y robusto mediante la combinación de gradientes de fase con filtrado angular, lo que proporciona mapas cuantitativos con capacidad de implementación en tiempo real. A continuación, se propone un método de elastogradía basado en vórtices acústicos como estrategia de excitación espacialmente estructurada para mejorar la detectabilidad de las ondas de cizallamiento a través del movimiento torsional. Por último, se introduce una técnica basada en excitación codificada temporal mediante uso de secuencias de longitud máxima *(MLS, maximum-length sequence*) y compresión de pulsos para mejorar la relación señal-ruido sin aumentar la presión máxima. Además, el método de codificación por MLS es evaluado en medios relevantes, como tejidos heterogéneos, anisotrópicos, viscoelásticos, ruidosos, en movimiento y *in vivo*. En general, la tesis muestra que avances coordinados en modelización, diseño de excitación y reconstrucción pueden ampliar la aplicabilidad de la SWE por ultrasonidos a escenarios complejos en los que su uso ha sido tradicionalmente limitado.

Resum

L'elastografia per ones de cisalla per ultrasons (SWE, *shear-wave elastography*) és una modalitat d'imatge quantitativa per a estimar les propietats mecàniques d'un teixit tou mitjançant la generació, l'observació i l'anàlisi de les ones de cisalla que s'hi propaguen. Atés que la rigidesa dels teixits tous es veu alterada per nombroses patologies, la SWE proporciona informació clínicament rellevant més enllà de la que ofereix l'ecografia convencional en mode B. No obstant això, l'estimació fiable de paràmetres mecànics continua sent un repte, ja que el rendiment de la SWE depén dels efectes combinats de la generació d'ones, el seguiment del moviment i els mètodes de reconstrucció d'imatge en presència dels fenòmens d'atenuació, heterogeneïtat, soroll i restriccions en l'eixida acústica. Esta tesi desenvolupa mètodes avançats de SWE per a millorar la precisió de la reconstrucció d'imatge, l'eficiència computacional i la detectabilitat de les ones de cisalla sota estes limitacions pràctiques.

Les seues contribucions principals són quatre. En primer lloc, es desenvolupa un marc computacional eficient per a modelar el procés complet de SWE, des de l'excitació acústica fins a la mesura del moviment, la qual cosa permet una validació controlada i el desenvolupament de nous mètodes. En segon lloc, es desenvolupa un mètode ràpid i robust d'imatge espectral de velocitat de fase mitjançant la combinació de gradients de fase i filtratge angular, proporcionant mapes quantitatius amb capacitat de funcionament en temps real. A continuació, es proposa una nova tècnica d'elastografia per vòrtexs acústics com una estratègia d'excitació espacialment estructurada per a millorar la detectabilitat de les ones de cisalla mitjançant moviment torsional. Finalment, es proposa una estratègia d'excitació codificada temporalment basada en seqüències de longitud màxima (MLS, *maximum-length sequence*), combinada amb compressió de pols, per a millorar la relació senyal-soroll sense augmentar la pressió pic. Esta tècnica de codificació temporal MLS s'avalua en medis complexos com ara teixits heterogenis, anisòtrops, viscoelàstics, sorollosos, en moviment i *in vivo*. En conjunt, la tesi mostra que els avanços coordinats en modelatge, disseny d'excitació i reconstrucció poden ampliar l'aplicabilitat de la SWE per ultrasons a escenaris complexos en què el seu ús ha estat tradicionalment limitat.

Contents

1 Introduction **1**
1.1 Tissue stiffness as a clinically relevant biomarker 3
1.2 Ultrasound elastography 6
1.3 Motivation and objectives 10
1.4 Thesis overview 13

2 Modeling Shear Wave Propagation **21**
2.1 Introduction 23
2.2 Hybrid model for shear wave simulation 26
2.3 Application example 40
2.4 Shear-wave velocity characterization 52
2.5 Discussion 59
2.6 Conclusions 61

3 Phase Gradient Elastography **67**
3.1 Introduction 69
3.2 Phase gradient elastography method 70
3.3 Numerical and experimental validation 82
3.4 Computational complexity 92
3.5 Discussion 93
3.6 Conclusions 96

4 Acoustic Vortex Elastography **101**
4.1 Introduction 103
4.2 Acoustic vortex elastography 105
4.3 Generating acoustic vortices 107
4.4 Orbital angular momentum validation 112
4.5 Shear-wave generation efficiency 117
4.6 Elasticity imaging using vortex beams 120

4.7 Discussion . . . 127
4.8 Conclusions . . . 129

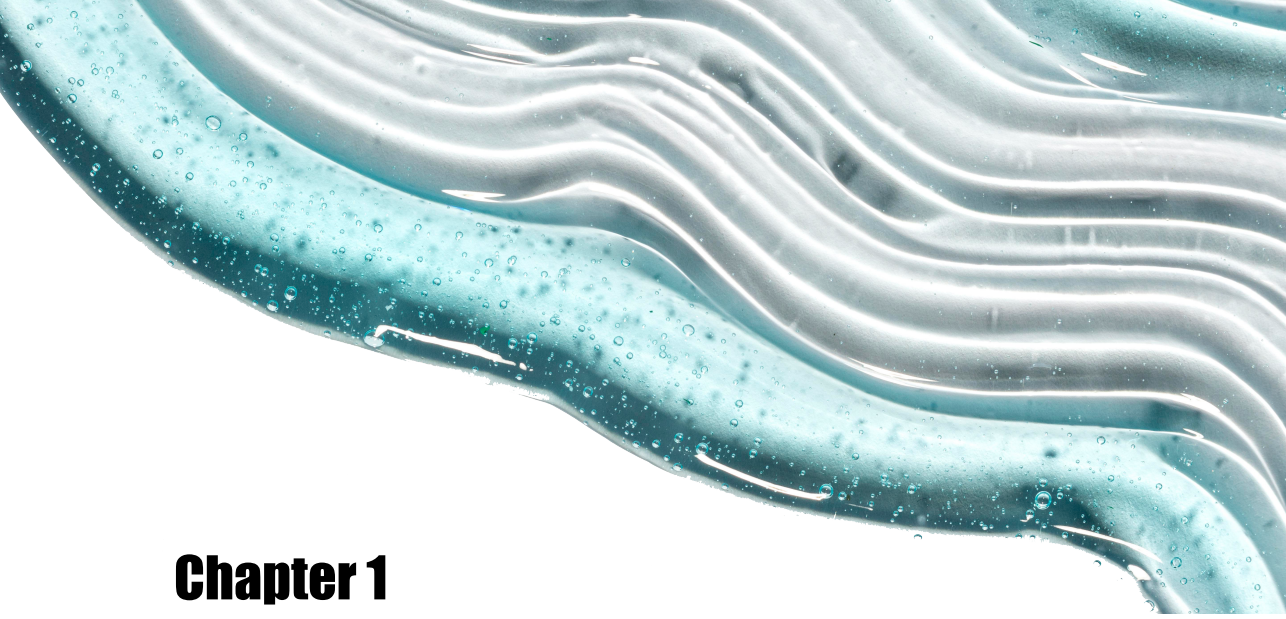
5 MLS-Coded Shear Wave Elastography 135
5.1 Introduction . . . 137
5.2 MLS-coded ARF pulse-compression method . . . 138
5.3 Numerical and experimental methods . . . 143
5.4 Results in homogeneous media . . . 146
5.5 Acoustic safety characterization . . . 154
5.6 Discussion . . . 159
5.7 Conclusions . . . 162

6 Coded-MLS ARF in Challenging Scenarios 167
6.1 Introduction . . . 169
6.2 Shared experimental methods . . . 170
6.3 Case studies . . . 172
6.4 Discussion . . . 196
6.5 Conclusions . . . 201

7 Conclusions 205
7.1 Conclusions . . . 207
7.2 Future perspectives . . . 210

A Appendix A - Kelvin–Voigt model 213

Research projects and funding 219



Chapter 1

Introduction

Abstract

This chapter presents elastography as a quantitative extension of palpation and introduces tissue stiffness as a clinically relevant biomarker. The text elaborates on the principle that mechanical imaging introduces complementary contrast by differentiating between the limited range of longitudinal wave speeds and the broader variability in shear-wave speeds. A generic elastography workflow is delineated in terms of excitation, motion detection, and quantification. The subsequent discussion focuses on ultrasound elastography, which is characterized as predominantly semi-quantitative. Shear-wave imaging is then outlined as a method for quantitative stiffness mapping through ultrafast tracking. The chapter concludes with a discussion of the practical challenges associated with isolating mechanical parameters in ultrasound SWE. Finally we present the thesis's motivation, objectives, and overview.

Outline

1.1 Tissue stiffness as a clinically relevant biomarker 3
1.2 Ultrasound elastography 6
1.2.1 Strain imaging 7
1.2.2 Shear-wave imaging 7
1.2.3 What can be measured with elastography 9
1.3 Motivation and objectives 10
1.3.1 From a powerful biomarker to a fragile measurement . . 10
1.3.2 Objectives 12
1.4 Thesis overview 13

1.1 Tissue stiffness as a clinically relevant biomarker

Medical imaging is built on contrast mechanisms: X-ray exploits attenuation linked to mass density [Röntgen, 1896], magnetic resonance leverages proton density and relaxation [Lauterbur, 1973], and ultrasound primarily reflects acoustic impedance and scattering [Griffin and Galambos, 1941]. Yet, many clinically meaningful changes are mechanical. For example, clinicians have implicitly used mechanical principles for centuries through palpation, in which pathological tissue commonly appears harder or softer than its surroundings. Palpation, however, is intrinsically subjective, depends on examiner's proficiency, has limited sensitivity to deep or small lesions, and is difficult to compare objectively across time. Elastography resolves these problems by translating "palpation" into a quantitative, image-based assessment of tissue mechanical response [Ophir et al., 1991].

A basic physical argument explains why elastography adds genuinely new information. In soft biological tissues, the speed of compressional or longitudinal waves is constrained by near-incompressibility and high water content and therefore varies within a relatively narrow range, typically around 1450–1600 m/s [Sarvazyan et al., 2011]. In contrast, shear wave propagation in soft tissues is mainly controlled by the shear modulus and can span orders of magnitude across tissue types and physiological states, from $\sim$1 m/s in very soft tissues to tens or even hundreds of m/s in stiffer structures [Ozturk et al., 2018]. This difference is illustrated in Fig. 1.1, where longitudinal speeds cluster in a narrow band around water-like values, whereas shear-wave speeds separate tissue classes over a much wider dynamic range. Elastography therefore adds a distinct source of contrast to conventional ultrasound, complementing the backscatter-based information of B-mode imaging with sensitivity to tissue stiffness.

From a clinical perspective, stiffness is relevant because it integrates multiple microstructural as well as physiological aspects that are altered by disease, including fibrosis, inflammation [Laharie et al., 2010], edema [Ophir et al., 1999] and changes in tissue morphology [Righetti et al., 1999]. These processes affect a broad range of organs and applications, making elastography attractive not as a niche tool but as a general-purpose biomarker for diagnosis, staging, and treatment monitoring [Kallel et al., 1999].

Elastography methods share a common workflow composed of three building blocks: (i) *mechanical excitation* to induce tissue deformation, (ii) *motion*

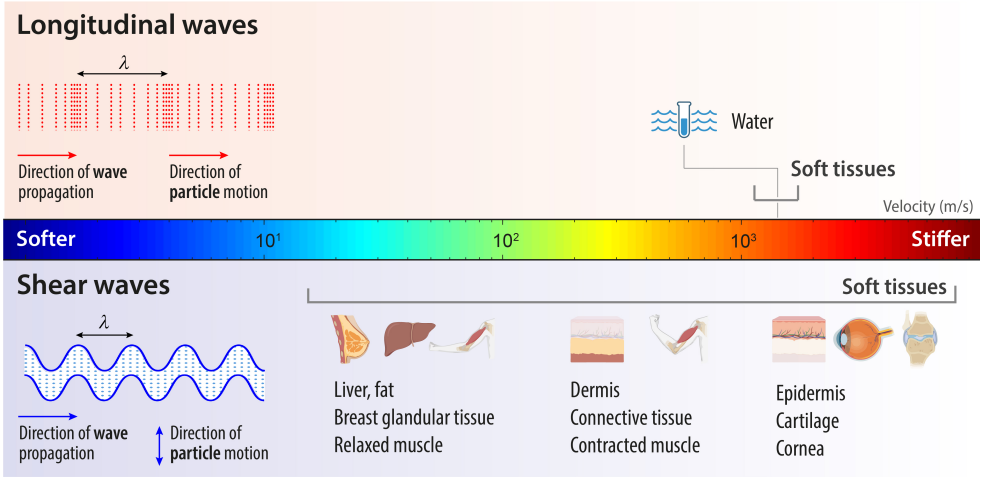


Figure 1.1: Illustrative comparison of characteristic longitudinal and shear wave speeds in soft biological materials. Longitudinal wave speeds cluster around 10^3 m/s (water-like values), while shear-wave speeds span a much broader range across tissues, enabling enhanced mechanical contrast for elastography. Figure created using Biorender icons.

detection to capture the resulting displacement field, and (iii) *quantification* to convert the measured deformation into a mechanical descriptor such as strain contrast, shear-wave speed, or viscoelastic parameters [Ozturk et al., 2018]. Excitation may be delivered externally through mechanical actuators or manual compression [Lerner et al., 1990], arise from endogenous motion such as respiration or blood flow [Gallot et al., 2011, Bae et al., 2007], or be generated internally through acoustic radiation force using a focused ultrasound beam [Sarvazyan et al., 2011]. Motion detection can then be performed with magnetic resonance imaging [Muthupillai et al., 1995], optical methods [Schmitt, 1998], or ultrasound [Ophir et al., 1999]. These modalities offer different trade-offs: magnetic resonance provides excellent volumetric characterization but with longer and more expensive acquisitions, optical methods achieve very high resolution at shallow depths, and ultrasound offers a favorable balance between penetration, temporal resolution, cost, and portability.

Quantification also spans a broad methodological spectrum. Some approaches estimate strain directly from the displacement field to obtain relative stiffness contrast [Ophir et al., 1991], whereas others exploit the physics of shear-wave propagation to recover quantitative mechanical parameters. These wave-based methods include propagation-speed estimators derived from

transient or time-of-flight analysis [Sandrin et al., 2003, Bercoff et al., 2004], inversion schemes based on full wave fields as in magnetic resonance elastography [Muthupillai et al., 1995], and dispersion-based approaches aimed at estimating viscoelastic parameters from the frequency dependence of shear-wave speed [Chen et al., 2004, Muller et al., 2009]. In general, strain-based methods are simpler and faster to implement but remain mostly qualitative, whereas shear wave-based techniques provide more quantitative and physically meaningful biomarkers at the cost of increased modeling and processing complexity.

Within this landscape, ultrasound has become a highly effective platform because it combines real-time operation, portability, wide availability, and cost effectiveness, enabling elastography to be deployed at scale in routine clinical workflows [DeWall, 2013]. This rapid translation is also motivated by a practical diagnostic gap: B-mode ultrasound primarily reports acoustic impedance and scattering, which may not uniquely reflect mechanical integrity. As a result, lesions or tissue states with comparable echogenicity can present markedly different stiffness, and stiffness-based imaging can therefore increase diagnostic specificity and improve longitudinal monitoring [Sigrist et al., 2017].

To emphasize how quickly elastography has grown from a research concept into a clinically used technology, Fig. 1.2 summarizes the annual number of publications indexed in the Scopus database between 2000 and 2025 whose titles explicitly include the term “elastography”. A steady increase in activity is observed, consistent with successful application across multiple diagnostic and treatment-monitoring scenarios. Notably, this acceleration coincides with the integration of elastography modes into clinical ultrasound scanners, enabling routine use throughout various patient populations worldwide.

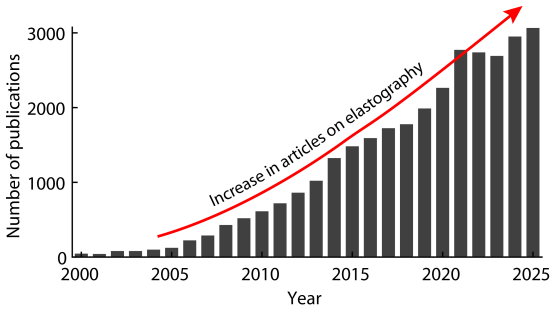
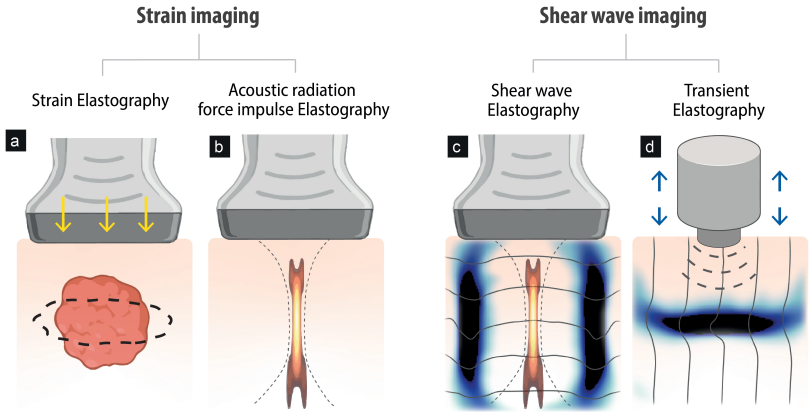


Figure 1.2: Annual number of publications indexed in the Scopus database between 2000 and 2025 whose titles explicitly include the term “elastography”.



Type of force	Quasi-static	Dynamic	Dynamic	Dynamic
Applied force	Manual compression	Acoustic radiation force	Acoustic radiation force	External vibration
Frequency range	10 Hz	Up to ~1000 Hz	Up to ~1000 Hz	10 - 100 Hz
Quantitative	No	No	Yes	Yes

Figure 1.3: Representative ultrasound elastography modalities grouped into strain imaging and shear-wave imaging. The excitation strategy determines the characteristic motion bandwidth and influences whether the output is qualitative or quantitative.

1.2 Ultrasound elastography

Building on the idea of stiffness as a broadly informative biomarker, this section narrows the focus to *ultrasound elastography*, the most widely deployed clinical implementation of elasticity imaging thanks to its real-time operation, portability, and incorporation within routine scanners [DeWall, 2013, Sigrist et al., 2017]. In ultrasound, elastography methods are commonly grouped into two main imaging modalities according to the physical quantity that is ultimately measured: *strain (compression) imaging* and *shear-wave imaging* [Sigrist et al., 2017, Ozturk et al., 2018]. Figure 1.3 summarizes representative techniques within each family.

All ultrasound elastography techniques share the same backbone: a mechanical stimulus produces deformation, ultrasound tracks motion, and a mechanical descriptor is estimated from the measured displacement or strain field [Ozturk et al., 2018]. The distinction between modalities derives from what is quantified and, crucially, from whether the applied stress is well enough

known to support absolute mechanical estimates.

1.2.1 Strain imaging

Strain elastography [Ophir et al., 1991] was the first ultrasound elastography approach to reach clinical practice [Konofagou and Ophir, 1998]. In its simplest form, a quasi-static stress is applied, and the resulting axial strain is estimated from pre- and post-compression ultrasound data. In principle, Hooke's law links both through Young's modulus. In practice, the applied stress is usually unknown or poorly controlled in freehand compression, so strain elastography is predominantly *semi-quantitative*, providing relative stiffness contrast rather than absolute modulus values.

Excitation in strain imaging can be delivered by manual probe compression, by external mechanical actuation, or by endogenous sources such as cardiac pulsation and respiration that naturally load tissue and generate measurable motion [Sarvazyan et al., 2011]. These methods commonly emphasize very low-frequency or quasi-static deformation content, from near-DC up to a few Hz for freehand compression, and up to tens of Hz when external vibration is used, depending on the implementation [Sigrist et al., 2017].

A key development in the 2000s was the use of *acoustic radiation force* to generate localized internal displacements without physically compressing the surface [Nightingale et al., 2002]. By focusing a moderate-intensity ultrasound beam, a short-duration internal "push" can be produced, yielding micrometer-scale displacements that are tracked with conventional ultrasound acquisitions. Because the stress distribution remains difficult to measure *in vivo*, these ARF-based strain methods generally preserve the relative nature of strain elastography while improving applicability to deeper targets [Konofagou et al., 2003].

1.2.2 Shear-wave imaging

Shear-wave imaging provides a clearer route to *quantitative* stiffness estimation by generating propagating shear waves and quantifying their speed [Bercoff et al., 2004]. The key observation is that shear waves in soft tissues propagate orders of magnitude more slowly than compressional waves, so their motion can be isolated and tracked by ultrasound despite being embedded in standard

pulse-echo acquisitions [Sarvazyan et al., 1998]. Once the shear-wave speed c_s is estimated, and under the common assumptions of an elastic, isotropic, nearly incompressible medium, Young's modulus can be approximated as $E \approx 3\rho c_s^2$, where ρ is the tissue density [Sandrin et al., 2003].

Two representative branches illustrate how shear waves are generated. In *transient elastography*, a low-frequency external actuator excites shear waves that are monitored ultrasonically along the excitation axis [Catheline et al., 1999, Bercoff et al., 2003]. In *ARF-based shear-wave imaging*, a focused ultrasound push produces an internal impulsive source that launches shear waves with higher effective bandwidth and enables a fully ultrasound-based implementation [Nightingale et al., 2002, González-Salido et al., 2016]. In both cases, ultrafast imaging is used to capture the transient wavefield, typically requiring kilohertz-range frame rates to resolve propagation over clinically relevant distances [Tanter et al., 2002, Montaldo et al., 2009].

From these basic actuation schemes, several families of shear-wave imaging techniques have emerged. Early implementations relied on one-dimensional tracking and pointwise speed estimation, whereas later developments incorporated two-dimensional imaging and more advanced reconstruction strategies [Sarvazyan et al., 1998, Nightingale et al., 2003]. A particularly influential milestone was *supersonic shear imaging* (SSI) in which a sequence of laterally shifted pushes generates a quasi-plane Mach cone and extends the spatial coverage of the induced wavefield, enabling real-time quantitative elasticity imaging [Bercoff et al., 2004]. Since then, numerous SWI variants have been proposed, including crawling waves [Manduca et al., 2003], shear-wave dispersion ultrasound vibrometry [Deffieux et al., 2009, Chen et al., 2009], comb-push excitation [Song et al., 2012], and fast shear compounding [Song et al., 2014], each reflecting different choices in excitation mechanism, dimensionality of tracking, reconstruction strategy, and compromise between speed, robustness, and implementation complexity.

Among these approaches, ARF-based shear-wave imaging has become especially attractive for clinical implementation because excitation, tracking, and quantification can be performed within the timescale of a standard ultrasound examination and, crucially, with the same device. Figure 1.4 illustrates this canonical push–track–quantify workflow. First, a focused push generates a localized internal displacement through acoustic radiation force [Nightingale et al., 2003]. Second, ultrafast acquisitions track the time-dependent propagation of shear waves, enabling micrometer-scale motion estimation

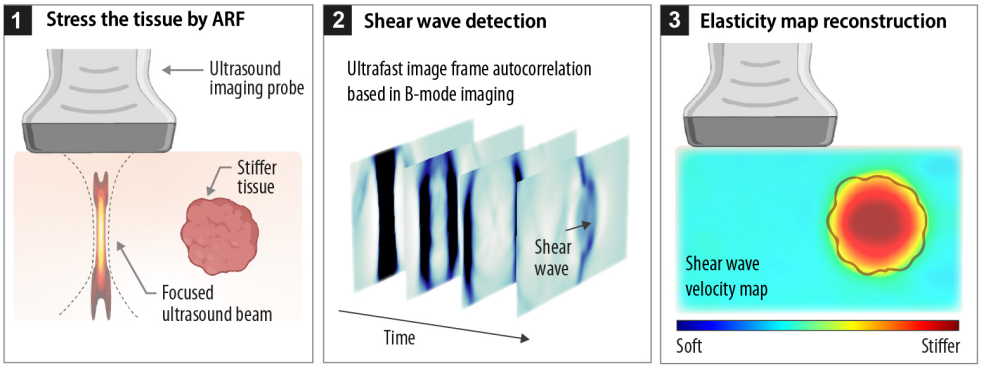


Figure 1.4: Canonical workflow of ARF-based shear-wave imaging. A localized push generates shear waves, ultrafast ultrasound tracks tissue motion over time, and a reconstruction stage estimates shear-wave speed to produce a quantitative stiffness map.

via correlation-based methods [Kasai et al., 1985, Loupas et al., 1995, Tanter et al., 2002]. Third, a reconstruction stage estimates local shear-wave speed and produces a quantitative stiffness map [Bercoff et al., 2004]. This basic procedure is not only a convenient description of clinical SWE, but also a useful scaffold for the contributions developed in this thesis, which address excitation design, motion tracking, and robust quantification under practical constraints.

1.2.3 What can be measured with elastography

Elastography is often described as a method to "measure stiffness," but the reported quantity depends on the excitation mechanism, the tracked motion component, and the bandwidth of the observed response. It therefore provides access to several mechanical descriptors rather than to a single universal parameter.

The most common clinical descriptor is *elastic stiffness*, usually expressed through Young's modulus, E, or shear modulus, μ. The basic strain-based elastography framework, together with quantitative imaging of strain and elastic modulus distributions under measured stress and appropriate corrections, was established in early seminal work [Ophir et al., 1991]. This also explains why strain elastography is mainly used for relative contrast: strain can be estimated from ultrasound tracking, whereas absolute modulus reconstruction requires knowledge of the applied stress field.

In ultrasound shear-wave elastography, however, the most direct observable is often the *shear-wave speed*. Under assumptions such as linearity, isotropy, and near incompressibility, shear-wave speed can be related to elastic stiffness, which is why it is widely used in practice as a surrogate mechanical descriptor [Sandrin et al., 2003, Bercoff et al., 2003, Hall et al., 2013]. Its practical value lies in the fact that it can be estimated directly from wave propagation without explicit knowledge of the applied stress field.

Soft tissues also exhibit *viscoelasticity*, so propagation is not fully described by a single constant speed. Viscosity introduces attenuation and dispersion [Sinkus et al., 2005], making the response frequency dependent. This motivates the distinction between *group velocity*, commonly estimated from time-of-flight measurements on broadband wave packets, and *phase velocity*, which characterizes the propagation of individual harmonic components and is central to spectral elastography [Parker, 2018]. In weakly dispersive media these two quantities may be similar, but in viscoelastic media they can differ substantially and provide complementary information [Gennisson et al., 2010].

Other properties can also be probed. For example, *nonlinear elasticity* has been assessed by combining static elastography with shear-wave elastography within an acoustoelasticity framework [Latorre-Ossa et al., 2012]. *Anisotropy* has been quantified *in vivo* from direction-dependent shear-wave propagation in muscle [Papazoglou et al., 2006, Brum et al., 2014]. *Heterogeneity* has likewise been exploited clinically, for example in breast elastography, where spatial stiffness patterns provide diagnostic information beyond mean values [Garra et al., 1997].

In principle, elastography can probe a broad range of mechanical behavior. In practice, however, interpretation is conditioned by departures from ideal assumptions and by experimental factors such as attenuation, boundaries, target geometry, region-of-interest (ROI) depth, and signal processing. These limitations motivate the methodological developments addressed in this thesis.

1.3 Motivation and objectives

1.3.1 From a powerful biomarker to a fragile measurement

The previous sections have shown that elastography can provide access to clinically relevant mechanical descriptors, and that ultrasound shear-wave

imaging offers a practical route to quantitative stiffness mapping [Sandrin et al., 2003, Bercoff et al., 2004]. Yet, the existence of a meaningful biomarker does not guarantee a reliable measurement. In ultrasound SWE, the quantity of interest is inferred from transient micrometer-scale displacements generated and tracked ultrasonically [Palmeri et al., 2006]. These waves propagate through heterogeneous, attenuating, dispersive, and often moving media, so quantitative estimation remains intrinsically sensitive to both physical and implementation-related limitations [Nightingale et al., 2003, Palmeri et al., 2008].

In practice, robustness depends on the full measurement chain: shear waves must be generated reproducibly, tracked with sufficient sensitivity and temporal resolution, and interpreted under assumptions that are only approximately satisfied *in vivo*. Early studies showed that measured shear-wave speed can depend on acquisition conditions such as transducer, depth, and lateral tracking range [Zhao et al., 2011], and that physiological motion, low displacement signal-to-noise ratio, and spatial inhomogeneity can corrupt time-of-flight estimation unless robust reconstruction strategies are used [Wang et al., 2010]. In addition, depth-dependent degradation, reflections, and acquisition or processing choices can further bias the reconstructed maps [Wang et al., 2014, Wood et al., 2016, Hollender et al., 2015]. This gap between biomarker relevance and measurement fragility motivates the present thesis, which focuses on methods to improve excitation, motion tracking, and quantitative reconstruction under realistic experimental constraints.

This thesis adopts the view that ultrasound shear-wave elastography can be most accurately described as a flowchart. A mechanical stimulus must be provided effectively and safely, the resulting motion must be captured with sufficient sensitivity and temporal sampling, and the mechanical descriptor must be reconstructed in a way that is robust to the complexity of real wavefields. Weakness in any of these steps can convert a strong biomarker into a noisy or biased estimate. The central motivation of this work is therefore to strengthen this flowchart end-to-end, with a focus on methods that improve dependability during realistic constraints rather than under idealized laboratory conditions.

1.3.2 Objectives

The thesis objectives are formulated to be measurable, technically grounded, and directly aligned with the chapter structure.

1. **Develop and validate an end-to-end forward model for ultrasound SWE.** An end-to-end simulation framework will be developed to predict the acoustic field, estimate radiation-force loading, simulate shear-wave propagation, and emulate ultrasound motion tracking. The model will be quantitatively benchmarked against established simulation tools and experimental measurements.

2. **Enable local, high-resolution, computationally efficient phase-velocity mapping.** A reconstruction method will be developed to estimate phase-velocity maps from complex wavefields using locally defined metrics. The objective is to support frequency-specific analysis while preserving spatial resolution and maintaining computational efficiency compatible with practical imaging.

3. **Increase shear-wave detectability under acoustic output restrictions via coded excitation.** A pulse-compression strategy based on maximum-length sequences will be implemented for ARF-based excitation. The resulting SNR gains will be quantified, and their impact on velocity estimation will be assessed together with the associated trade-offs in temporal resolution and exposure metrics.

4. **Demonstrate robustness in application-relevant, challenging conditions.** The proposed methods will be evaluated in situations that typically challenge shear-wave imaging, including heterogeneous and viscoelastic media, anisotropic tissue-like structures, safety-limited configurations, and active scenarios with external or physiological motion. A shared acquisition and processing method will be used to ensure equitable comparison across conditions.

5. **Explore structured actuation beyond conventional focusing for shear-motion generation.** Acoustic-vortex excitation will be investigated as an alternative strategy for generating shear motion. Its torque-driven response will be characterized and compared with standard techniques to assess its potential to improve motion generation under matched exposure constraints.

1.4 Thesis overview

This thesis is structured as a progressive account that strengthens ultrasound shear-wave elastography froms through three complementary dimensions: predicting the physics, quantifying motion robustly, and designing excitation strategies that improve detectability under feasibility considerations. Each chapter addresses a specific block of the elastography flowchart while providing the tools and evidence needed.

Chapter 2 presents an end-to-end simulation model for ultrasound SWE. The model will couple acoustic field prediction, radiation-force loading, shear-wave propagation, and ultrasound motion tracking. This chapter will provide the framework used to test parameter choices, interpret experimental observations, and benchmark reconstruction strategies against controlled references.

Chapter 3 lays out a reconstruction framework for local phase-velocity imaging based on phase-gradient and angular filter techniques. The chapter focuses on reliable estimation in complex wavefields, facilitating high-resolution maps while remaining computationally efficient. Frequency-specific formulations are also introduced to support viscoelastic analysis when required by the application.

Chapter 4 introduces acoustic vortex beams as an novel excitation mechanism for generating torsional shear motion in tissue. In particular, this chapter demonstrates the transfer of acoustic orbital angular momentum to a soft elastic medium using ultrasound vortex beams and characterizes the resulting torque-driven and torsional mechanical response by ultrasound tracking. The chapter shows how these singular beams can be used for shear wave elastography, comparing this novel technique with traditional methods, particularly in scenarios where conventional pushing excitations are inefficient or constrained.

Chapter 5 addresses the excitation block by proposing coded acoustic radiation force schemes based on maximum-length sequences and pulse compression. The chapter quantifies SNR gains and reviews their impact on phase-velocity estimation and measurement stability, with particular attention to regimes where typical single-push excitation is limited by exposure constraints or insufficient detectability.

Chapter 6 extend the coded excitation method to application-relevant scenarios that are known to challenge shear-wave imaging. The proposed coded excitation pipeline is evaluated in heterogeneous and viscoelastic media,

anisotropic tissues, and scenarios dominated by noise or motion. The objective of this chapter is to present an exposition of failure modes and to demonstrate the manner in which maximum-length sequences coded excitations and pulse compression approaches improve the reliability of shear wave elastography.

Finally, Chapter 7 offers a synopsis of the work, delineating the primary conclusions and proffering prospective avenues for future research.

The references cited in this work are listed at the end of each chapter to assist the reader in locating the sources.

References

Bae, U., Dighe, M., Dubinsky, T., Minoshima, S., Shamdasani, V., and Kim, Y. (2007). Ultrasound thyroid elastography using carotid artery pulsation: Preliminary study. *Journal of Ultrasound in Medicine*, 26(6):797–805.

Bercoff, J., Chaffai, S., Tanter, M., Sandrin, L., Catheline, S., Fink, M., Gennisson, J.-L., and Meunier, M. (2003). In vivo breast tumor detection using transient elastography. *Ultrasound in Medicine & Biology*, 29(10):1387–1396.

Bercoff, J., Tanter, M., and Fink, M. (2004). Supersonic shear imaging: A new technique for soft tissue elasticity mapping. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 51(4):396–409.

Brum, J., Bernal, M., Gennisson, J. L., and Tanter, M. (2014). In vivo evaluation of the elastic anisotropy of the human achilles tendon using shear wave dispersion analysis. *Physics in Medicine & Biology*, 59(3):505.

Catheline, S., Wu, F., and Fink, M. (1999). A solution to diffraction biases in sonoelasticity: The acoustic impulse technique. *The Journal of the Acoustical Society of America*, 105(5):2941–2950. _eprint: https://pubs.aip.org/asa/jasa/article-pdf/105/5/2941/8085677/2941_1_online.pdf.

Chen, S., Fatemi, M., and Greenleaf, J. F. (2004). Quantifying elasticity and viscosity from measurement of shear wave speed dispersion. *The Journal of the Acoustical Society of America*, 115(6):2781–2785.

Chen, S., Urban, M. W., and Greenleaf, J. F. (2009). Shearwave dispersion ultrasound vibrometry (sduv) for measuring tissue elasticity and viscosity. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 56(1):55–62.

Deffieux, T., Montaldo, G., Tanter, M., and Fink, M. (2009). Shear Wave Spectroscopy for In Vivo Quantification of Human Soft Tissues Visco-Elasticity. *IEEE Transactions on Medical Imaging*, 28(3):313–322. Conference Name: IEEE Transactions on Medical Imaging.

DeWall, R. J. (2013). Ultrasound Elastography: Principles, Techniques, and Clinical Applications. *Critical Reviews in Biomedical Engineering*, 41(1):1–19.

Gallot, T., Catheline, S., Roux, P., Brum, J., Benech, N., and Negreira, C. (2011). Passive elastography: shear-wave tomography from physiological-noise correlation in soft tissues. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 58(6):1122–1126.

Garra, B. S., Cespedes, E. I., Ophir, J., Spratt, S. R., Zuurbier, R. A., Magnant, C. M., and Pennanen, M. F. (1997). Elastography of breast lesions: initial clinical results. *Radiology*, 202(1):79–86.

Gennisson, J.-L., Deffieux, T., Macé, E., Montaldo, G., Fink, M., and Tanter, M. (2010). Viscoelastic and anisotropic mechanical properties of in vivo muscle tissue assessed by supersonic shear imaging. *Ultrasound in Medicine & Biology*, 36(5):789–801.

González-Salido, N., Medina, L., and Camacho, J. (2016). Full angle spatial compound of arfi images for breast cancer detection. *Ultrasonics*, 71:161–171.

Griffin, D. R. and Galambos, R. (1941). The sensory basis of obstacle avoidance by flying bats. *Journal of Experimental Zoology*, 86(3):481–506.

Hall, T. J., Milkowski, A., Garra, B., Carson, P., Palmeri, M., Nightingale, K., Lynch, T., Alturki, A., Andre, M., Audiere, S., Bamber, J., Barr, R., Bercoff, J., Bercoff, J., Bernal, M., Brum, J., Chan, H. W., Chen, S., Cohen-Bacrie, C., Couade, M., Daniels, A., DeWall, R., Dillman, J., Ehman, R., Franchi-Abella, S., Fromageau, J., Gennisson, J.-L., Henry, J. P., Ivancevich, N., Kalin, J., Kohn, S., Kugel, J., Lee, K., Liu, N., Loupas, T., Mazernik, J., McAleavey, S., Miette, V., Metz, S., Morel, B., Nelson, T., Nordberg, E., Oudry, J., Padwal, M., Rouze, N., Samir, A., Sandrin, L., Schaccitti, J., Schmitt, C., Shamdasani, V., Song, P., Switalski, P., Wang, M., Wear, K., Xie, H., and Zhao, H. (2013). Rsna/qiba: Shear wave speed as a biomarker for liver fibrosis staging. In *2013 IEEE International Ultrasonics Symposium (IUS)*, pages 397–400.

Hollender, P. J., Rosenzweig, S. J., Nightingale, K. R., and Trahey, G. E. (2015). Track location shear wave and acoustic radiation force impulse imaging: Initial in vivo results. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 62(3):565–579.

Kallel, F., Stafford, R. J., Price, R. E., Righetti, R., Ophir, J., and Hazle, J. D. (1999). The feasibility of elastographic visualization of hifu-induced thermal lesions in soft tissues. image-guided high-intensity focused ultrasound. *Ultrasound in Medicine & Biology*, 25(4):641–647.

Kasai, C., Namekawa, K., Koyano, A., and Omoto, R. (1985). Real-time two-dimensional blood flow imaging using an autocorrelation technique. *IEEE Transactions on Sonics and Ultrasonics*, 32(3):458–464.

Konofagou, E. and Ophir, J. (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 & Biology*, 24(8):1183–1199.

Konofagou, E. E., Thierman, J., Hynynen, K., and Ophir, J. (2003). Localized harmonic motion imaging: Theory, simulations and experiments. *Ultrasound in Medicine & Biology*, 29(10):1405–1413.

Laharie, D., Seneschal, J., Schaeverbeke, T., Doutre, M.-S., Longy-Boursier, M., Pellegrin, J.-L., Chabrun, E., Villars, S., Zerbib, F., and de Lédinghen, V. (2010). Assessment of liver fibrosis with transient elastography and FibroTest in patients treated with methotrexate for chronic inflammatory diseases: A case–control study. *Journal of Hepatology*, 53(6):1035–1040.

Latorre-Ossa, H., Gennisson, J.-L., De Brosses, E., and Tanter, M. (2012). Quantitative imaging of nonlinear shear modulus by combining static elastography and shear wave elastography. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 59(4):833–839.

Lauterbur, P. C. (1973). Image Formation by Induced Local Interactions: Examples Employing Nuclear Magnetic Resonance. *Nature*, 242(5394):190–191. Publisher: Nature Publishing Group.

Lerner, R. M., Huang, S. R., and Parker, K. J. (1990). "sonoelasticity" images derived from ultrasound signals in mechanically vibrated tissues. *Ultrasound in Medicine & Biology*, 16(3):231–239.

Loupas, T., Powers, J. T., and Gill, R. W. (1995). An axial velocity and displacement estimator for ultrasound imaging of fluid flow and tissue motion. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 42(4):672–688.

Manduca, A., Lake, D. S., Kruse, S. A., and Ehman, R. L. (2003). Sonoelastographic imaging of interference patterns for estimation of shear wave speed in biomaterials. *Physics in Medicine and Biology*, 48(10):1445–1460.

Montaldo, G., Tanter, M., Bercoff, J., Benech, N., and Fink, M. (2009). Coherent plane-wave compounding for very high frame rate ultrasonography and transient elastography. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 56(3):489–506.

Muller, M., Gennisson, J.-L., Deffieux, T., Tanter, M., and Fink, M. (2009). Quantitative viscoelasticity mapping of human liver using supersonic shear imaging: Preliminary in vivo feasibility study. *Ultrasound in Medicine & Biology*, 35(2):219–229.

Muthupillai, R., Lomas, D. J., Rossman, P. J., Greenleaf, J. F., Manduca, A., and Ehman, R. L. (1995). Magnetic Resonance Elastography by Direct Visualization of Propagating Acoustic Strain Waves. *Science*, 269(5232):1854–1857. Publisher: American Association for the Advancement of Science.

Nightingale, K., McAleavey, S., and Trahey, G. (2003). Shear-wave generation using acoustic radiation force: In vivo and ex vivo results. *Ultrasound in Medicine & Biology*, 29(12):1715–1723.

Nightingale, K., Soo, M. S., Nightingale, R., and Trahey, G. (2002). Acoustic radiation force impulse imaging: In vivo demonstration of clinical feasibility. *Ultrasound in Medicine & Biology*, 28(2):227–235.

Ophir, J., Alam, S. K., Garra, B., Kallel, F., Konofagou, E., Krouskop, T., and Varghese, T. (1999). Elastography: Ultrasonic estimation and imaging of the elastic properties of tissues. *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, 213(3):203–233. Publisher: IMECHE.

Ophir, J., Céspedes, I., Ponnekanti, H., Yazdi, Y., and Li, X. (1991). Elastography: a quantitative method for imaging the elasticity of biological tissues. *Ultrasonic Imaging*, 13(2):111–134.

Ozturk, A., Grajo, J. R., Dhyani, M., Anthony, B. W., and Samir, A. E. (2018). Principles of ultrasound elastography. *Abdominal Radiology*, 43(4):773–785.

Palmeri, M. L., McAleavey, S. A., Trahey, G. E., and Nightingale, K. R. (2006). Ultrasonic tracking of acoustic radiation force-induced displacements in homogeneous media. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 53(7):1300–1313.

Palmeri, M. L., Wang, M., Dahl, J. J., Frinkley, K. D., and Nightingale, K. R. (2008). Quantifying hepatic shear modulus in vivo using acoustic radiation force. *Ultrasound in Medicine & Biology*, 34(4):546–558.

Papazoglou, S., Rump, J., Braun, J., and Sack, I. (2006). Shear wave group velocity inversion in MR elastography of human skeletal muscle. *Magnetic Resonance in Medicine*, 56(3):489–497.

Parker, K. J. (2018). Group versus phase velocity of shear waves in soft tissues. *The Journal of the Acoustical Society of America*, 143(5):3222–3232.

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

Röntgen, W. C. (1896). On a New Kind of Rays. *Science*, 3(59):227–231. Publisher: American Association for the Advancement of Science.

Sandrin, L., Fourquet, B., Hasquenoph, J.-M., Yon, S., Fournier, C., Mal, F., Christidis, C., Ziol, M., Poulet, B., Kazemi, F., Beaugrand, M., and Palau, R. (2003). Transient elastography: A new noninvasive method for assessment of hepatic fibrosis. *Ultrasound in Medicine & Biology*, 29(12):1705–1713.

Sarvazyan, A., J. Hall, T., W. Urban, M., Fatemi, M., R. Aglyamov, S., and S. Garra, B. (2011). An Overview of Elastography-An Emerging Branch of Medical Imaging. *Current Medical Imaging Reviews*, 7(4):255–282.

Sarvazyan, A. P., Rudenko, O. V., Swanson, S. D., Fowlkes, J. B., and Emelianov, S. Y. (1998). Shear wave elasticity imaging: A new ultrasonic technology of medical diagnostics. *Ultrasound in Medicine & Biology*, 24(9):1419–1435.

Schmitt, J. M. (1998). OCT elastography: imaging microscopic deformation and strain of tissue. *Optics Express*, 3(6):199–211. Publisher: Optica Publishing Group.

Sigrist, R. M., Liau, J., Kaffas, A. E., Chammas, M. C., and Willmann, J. K. (2017). Ultrasound Elastography: Review of Techniques and Clinical Applications. *Theranostics*, 7(5):1303–1329.

Sinkus, R., Tanter, M., Xydeas, T., Catheline, S., Bercoff, J., and Fink, M. (2005). Viscoelastic shear properties of in vivo breast lesions measured by mr elastography. *Magnetic Resonance Imaging*, 23(2):159–165. Proceedings of the Seventh International Conference on Recent Advances in MR Applications to Porous Media.

Song, P., Macdonald, M., Behler, R., Lanning, J., Wang, M., Urban, M. W., Manduca, A., Zhao, H., Greenleaf, J. F., and Chen, S. (2014). Fast shear compounding using robust two-dimensional shear wave speed calculation and multi-directional filtering. *Ultrasound in Medicine & Biology*, 40(6):1343–1355.

Song, P., Zhao, H., Manduca, A., Urban, M. W., Greenleaf, J. F., and Chen, S. (2012). Comb-push ultrasound shear elastography (cuse): A novel method for two-dimensional shear elasticity imaging of soft tissues. *IEEE Transactions on Medical Imaging*, 31(9):1821–1832.

Tanter, M., Bercoff, J., Sandrin, L., and Fink, M. (2002). Ultrafast compound imaging for 2-d motion vector estimation: Application to transient elastography. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 49(10):1363–1374.

Wang, C. Z., Zheng, J., Huang, Z. P., Xiao, Y., Song, D., Zeng, J., and Zheng, H. R. (2014). Influence of measurement depth on the stiffness assessment of liver with real-time shear wave elastography. *Ultrasound in Medicine & Biology*, 40(3):461–469.

Wang, M. H., Palmeri, M. L., Rotemberg, V. M., Rouze, N. C., and Nightingale, K. R. (2010). Improving the robustness of time-of-flight based shear wave speed reconstruction methods using ransac in human liver in vivo. *Ultrasound in Medicine & Biology*, 36(5):802–813.

Wood, B. G., Barr, R. G., Kane, M., Rubens, D. J., and Parker, K. J. (2016). Evaluating the improvement in shear wave speed image quality using multi-dimensional directional filters in the presence of reflection artifacts. *Ultrasound in Medicine & Biology*, 42(8):1969–1985.

Zhao, H., Song, P., Urban, M. W., Kinnick, R. R., Yin, M., Greenleaf, J. F., and Chen, S. (2011). Bias observed in time-of-flight shear wave speed measurements using radiation force of a focused ultrasound beam. *Ultrasound in Medicine & Biology*, 37(11):1884–1892.



Chapter 2

Modeling Shear Wave Propagation for Ultrasound Elastography

Abstract

This chapter introduces an efficient numerical method for modeling shear-wave elastography in soft tissues. The proposed approach sequentially computes the acoustic field of a clinical linear array using a hybrid Rayleigh/angular-spectrum model. It then derives the resulting vectorial acoustic radiation force with a second-order formulation. Next, it drives shear-wave propagation in homogeneous Kelvin–Voigt media simulated with k–Wave. Finally, it embeds the computed motion into a speckle-generating medium to synthesize complex IQ frames that reproduce key features of experimental B-mode acquisitions. The acoustic pressure field was validated against two ultrasound simulation models and experimental measurements. On a representative 3D grid with $N \approx 5 \times 10^5$ points, the hybrid acoustic model generates the pressure field in 8.8 s, compared with 32.6 s for MUST and 144.1 s for k-Wave, demonstrating a substantial computational advantage without sacrificing quantitative agreement. The shear-wave dynamic response was assessed in homogeneous phantoms. Simulated shear wave dynamics agree with experiments in homogeneous phantoms, and shear-wave velocities agree within a few percent of manufacturer specifications. These findings demonstrate that the proposed approach is computationally efficient and quantitatively reliable for modeling shear wave elastography in soft, elastic, and viscoelastic solids.

Outline

2.1 Introduction 23
2.2 Hybrid model for shear wave simulation 26
2.2.1 Volumetric acoustic field 26
2.2.2 Acoustic radiation force 33
2.2.3 Shear wave simulation 34
2.2.4 Tracked shear wave motion generation 37
2.3 Application example 40
2.3.1 Validating the acoustic pressure field 41
2.3.2 Construction of volumetric ARF maps 47
2.3.3 Shear-wave motion validation 48
2.4 Shear-wave velocity characterization 52
2.4.1 Group and phase velocity metrics 53
2.4.2 Group velocity in homogeneous phantoms 54
2.4.3 Phase velocity in homogeneous phantoms 56
2.5 Discussion 59
2.6 Conclusions 61

2.1 Introduction

Clinical shear wave elastography (SWE) is still predominantly implemented with two-dimensional (2D) imaging, where the shear motion is tracked in the B-mode imaging plane, i.e., in the azimuth-axial (x-z) plane. This arrangement leaves the elevational (y) direction of the shear displacement unobserved, so that shear wave energy propagating out of imaging plane is not explicitly accounted for during stiffness reconstruction. As a result, conventional 2D processing can bias shear wave speed estimates [Wang et al., 2013, Lipman et al., 2016] when the wavefront exhibits substantial elevational spreading or anisotropic behavior, an issue that becomes more important for focused ARF excitations and heterogeneous media. In response, there is an active transition toward volumetric SWE using 2D matrix arrays or row–column addressed arrays [Gennisson et al., 2015, Huang et al., 2020, Bernal et al., 2021, Dong et al., 2024] that can generate and track shear waves in three dimensions (3D), enabling the estimation of fully 3D elastograms from volumetric motion fields.

The numerical reproduction of volumetric SWE experiments is challenging because the forward problem couples acoustic beam formation, acoustic radiation force generation, and transient elastodynamics across widely separated spatial and temporal scales. On the acoustic side, focused fields can be modeled through linear integral formulations based on Green's functions and spatial impulse responses [Jensen and Svendsen, 1992], through frequency-domain harmonic synthesis derived from the Rayleigh–Sommerfeld integral with far-field and paraxial approximations [Garcia, 2021], or through full-wave continuum models based on k-space pseudospectral solvers [Cox et al., 2007]. These formulations offer different trade-offs between fidelity and efficiency, depending on whether the priority is transducer realism, fast large-volume field synthesis, or the explicit representation of diffraction, heterogeneity, and absorption.

On the mechanical side, several numerical methods have been used to simulate ARF-driven shear-wave propagation. Finite-element formulations are well suited to complex geometries and boundary conditions [Palmeri et al., 2016, Merlini et al., 2025], whereas finite-differences [Qiang et al., 2015, Jiménez et al., 2012, Kijanka and Urban, 2019] and staggered-grid [Osika and Kijanka, 2024] approaches provide efficient time-domain modeling in viscoelastic tissue-like media. Pseudospectral stress–velocity solvers offer an efficient alternative for large domains and have already been adopted in SWE-oriented simulations

[Treeby et al., 2014, Prieur and Catheline, 2016]. However, fully coupled three-dimensional simulations remain computationally expensive for large parameter sweeps because the cost scales with both the volumetric grid size and the long time windows required to resolve millisecond shear-wave dynamics. This computational burden motivates the efficient stage-wise strategy adopted in this chapter.

In this chapter, an efficient bottom-up approach to simulate shear wave motion in soft solids for shear wave elastography is presented. Rather than introducing a new constituent numerical solver, the proposed framework integrates established and computationally efficient tools for each stage of the problem: acoustic field synthesis, ARF calculation, shear-wave propagation, and ultrasound motion tracking. Figure 2.1 represent the main steps of hybrid simulation method that links acoustic beam models, ARF computation and shear-wave elastic propagation.

Figure 2.1(a-b) illustrate the hybrid Rayleigh/Angular Spectrum (RAAS) strategy to calculate the volumetric acoustic pressure field. First, we use the first Rayleigh integral for projecting the normal velocity of an ultrasound source to a near-field plane [Bouwkamp, 1954, Sapozhnikov et al., 2015]. Then, we use the angular spectrum method to project this field to a dense 3D grid [Liu and Waag, 1997, Zeng and McGough, 2006]. We use this two-stage simulation because Rayleigh integral can model sources of arbitrary geometry, including narrow kernel spacing, and also arbitrary number of elements, including their temporal delay, apodization and elevation acoustic lens. As in shear wave elastography pushing ultrasound beams are usually long sinusoidal bursts, the Rayleigh integral is solved in the frequency-domain, therefore strongly reducing computational cost. Once the field is integrated on a transverse plane, the whole 3D harmonic field is calculated using the angular spectrum method, resulting in a faster approach for large volumetric domains.

In the subsequent step, depicted in Fig. 2.1(c), the ARF is then formed in second-order (SO) approximation from simulated pressure and particle velocity, yielding a time-averaged vector body force $\mathbf{F} = (F_x, F_y, F_z)$ that explicitly includes elevational forcing. The SO formulations account for beam structure and, for conventional probes and focused beams, capture the displacement-driving force more accurately than attenuated plane-wave surrogates; we therefore adopt SO by default [Prieur and Sapozhnikov, 2017, Prieur and Catheline, 2016].

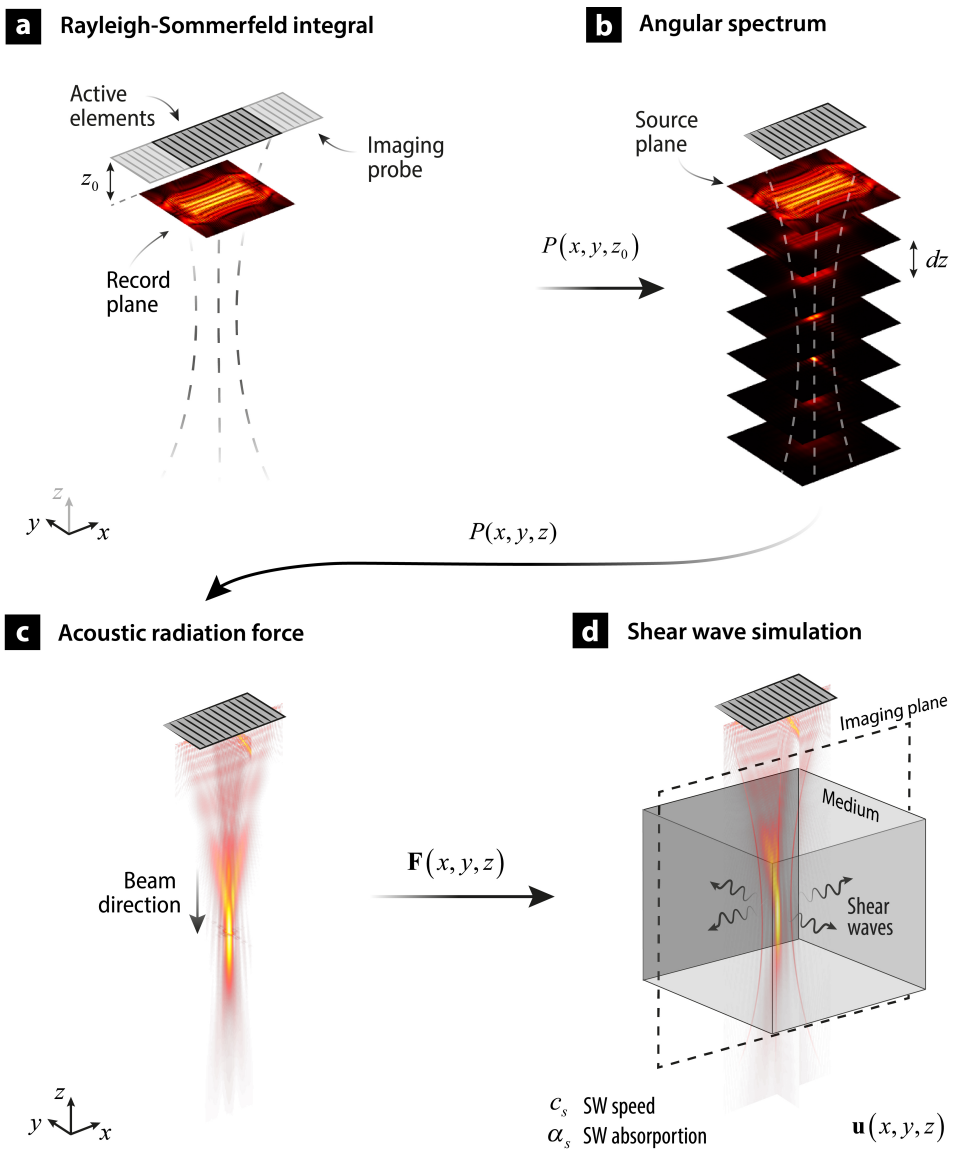


Figure 2.1: Hybrid acoustic modeling scheme. Schematic of the hybrid RAAS stage, where the (a) Rayleigh boundary synthesis on a reference plane is combined with (b) angular-spectrum propagation to build the three-dimensional pressure field that feeds the subsequent (c) acoustic radiation force and (d) shear-wave simulations.

The final stage of the elastic stage is resolved through the implementation of an elastodynamic solver, as illustrated in Fig. 2.1(d). The simulated soft tissue can be modeled as either elastic or viscoelastic; the latter utilizes a Kelvin-Voigt model[1]. This solver utilizes coupled first-order stress–velocity equations, which are discretized by a pseudospectral time-domain scheme with k-space corrector. This combination (i) provides a physically grounded loss model appropriate for soft tissue, (ii) supports arbitrary heterogeneous media via spatially dependent density and shear and compressional wave velocities [Firouzi et al., 2011], and (iii) attains high accuracy with fewer points per wavelength than low-order finite differences, reducing computational cost for large 3D volumes and parameter sweeps.

In order to make the simulation motion data more similar to an experimental case, the simulated shear wave displacements are embedded into synthetic ultrasound data. Volumetric shear wave displacement governs the motion dynamics of random clouds of point scatters. Then, we generate synthetic ultrasound RF data and complex IQ data for each time step through the MUST toolbox [Garcia, 2021], which generates RF and B-mode frames consistent with the imaging sequence used in experiments. The same beamforming and motion-estimation algorithms can then be applied to synthetic and experimental data, providing a direct and realistic route to validate reconstruction methods and shear wave velocity estimation techniques under controlled conditions.

In the following sections the accuracy and computational cost of the proposed method are assessed against established experimental measurements and other numerical tools.

2.2 Hybrid model for shear wave simulation

2.2.1 Volumetric acoustic field

The acoustic stage is deliberately split into two complementary operators to balance fidelity and efficiency in three dimensions. A Rayleigh–type boundary synthesis preserves near-field accuracy on a reference plane $z = z_0$, including the discrete element layout, kerf, elevation height and lens-induced phase. A subsequent plane-to-volume propagation by angular spectrum (AS) extrapolates this complex pressure map to the full axial extent of interest with

[1] Other viscoelastic models can also be used.

FFT-accelerated transforms.

Rayleigh-Sommerfeld integral

In the context of a field point denoted by $\mathbf{r}_0 = (x, y, z_0)$, where $z_0 > 0$, and operating within the time convention $e^{-\mathrm{i}\omega t}$, the Rayleigh–Sommerfeld representation of the second kind [Kirkup, 1994] yields a continuous expression.

$$P(\mathbf{r}_0, \omega) = \mathrm{i}\,\omega\,\rho_0 \int_S \frac{e^{-\mathrm{i}k(\omega)\|\mathbf{r}_0 - \mathbf{r}_s\|}}{2\pi\|\mathbf{r}_0 - \mathbf{r}_s\|}\, v_n(\mathbf{r}_s, \omega)\, \mathrm{d}S, \tag{2.1}$$

where $\mathbf{r}_s = (x_s, y_s, z_s)$ denotes a point on the radiating surface, ρ_0 is the medium density, ω is the angular frequency of the pushing burst, $k = \omega/c_0$ is the wavenumber and c_0 the sound speed in the propagation medium. The prescribed boundary quantity is the normal particle-velocity amplitude $v_n(\mathbf{r}_s, \omega)$ on S; outside S the rigid-baffle condition enforces $v_n = 0$. To account for compressional-wave attenuation, the wavenumber is taken as complex,

$$k(\omega) = k_r(\omega) - \mathrm{i}\,\alpha_p(\omega), \qquad k_r(\omega) = \frac{\omega}{c_0}, \tag{2.2}$$

so that the Rayleigh kernel in (2.1) includes an amplitude decay term $e^{-\alpha_p(\omega)\|\mathbf{r}_0 - \mathbf{r}_s\|}$ in addition to the usual phase accumulation.

For example, a linear array transducer is modeled as a finite set of rectangular piston elements mounted on a rigid baffle at $z = 0$. Let N_{el} be the number of elements along the lateral (azimuthal) direction x. Figure 2.2 represents the different parts that define each element: width w, elevation height L_e, and the inter-element gap (kerf) g. The pitch is therefore $p = w + g$. Denoting the element centers as

$$x_n = \Big(n - \frac{N_{el} + 1}{2}\Big)p, \qquad n = 1, \ldots, N_{el},$$

the active radiating surface can be written as

$$S = \bigcup_{n=1}^{N_{el}} \Big\{(x_s, y_s, z_s) :\ x_s \in \big[x_n - \tfrac{w}{2},\, x_n + \tfrac{w}{2}\big],\ y_s \in \big[-\tfrac{L_e}{2},\, \tfrac{L_e}{2}\big],\ z_s = 0\Big\}. \tag{2.3}$$

Each element is discretized into $N_{el_x} \times N_{el_y}$ patches with steps dx_e and dy_e in the lateral and elevation directions. This construction is illustrated in

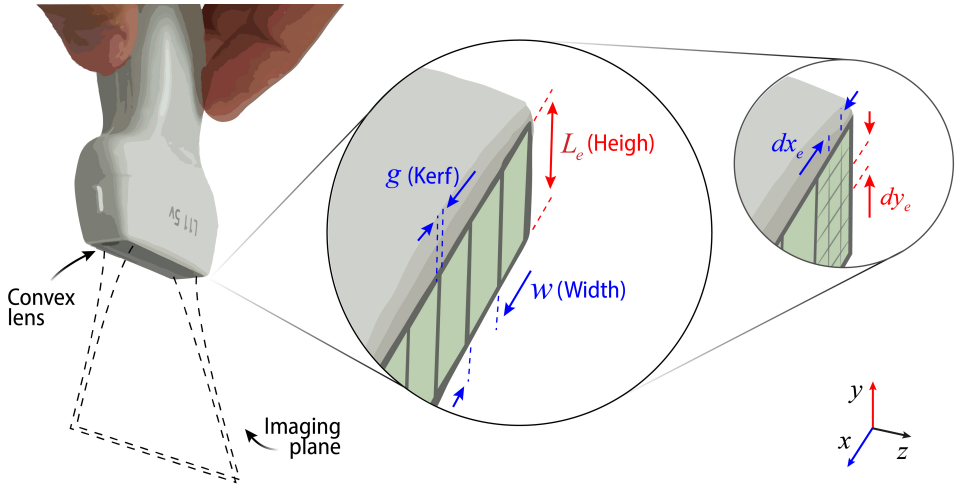


Figure 2.2: Linear-array probe geometry and element specifications.

Fig. 2.2, where the element width w, the kerf g, the element height L_e, and the local quadrature differential $\mathrm{d}S_e = dx_e\, dy_e$ are drawn explicitly. Writing the patch coordinates as $\mathbf{r}_{nq} = (x_{nq}, y_{nq})$ for $q = 1, \ldots, N_{el_x} N_{el_y}$ within the active rectangle $[x_n - w/2,\, x_n + w/2] \times [-L_e/2,\, L_e/2]$, the discretized Rayleigh sum reads

$$P(\mathbf{r}_0, \omega) \;\approx\; \frac{-\mathrm{i}\,\omega\,\rho_0}{2\pi} \sum_{n=1}^{N} \sum_{q=1}^{N_{el_x} N_{el_y}} \frac{e^{\mathrm{i}kR_{nq}(\mathbf{r}_0, \mathbf{r}_s)}}{R_{nq}(\mathbf{r}_0, \mathbf{r}_s)}\; v_n(\mathbf{r}_{nq}, \omega)\; dS_e, \tag{2.4}$$

with

$$R_{nq}(\mathbf{r}_0, \mathbf{r}_s) = \sqrt{(x - x_{nq})^2 + (y - y_{nq})^2 + z_0^2}. \tag{2.5}$$

The mesh is chosen sufficiently fine to resolve both the acoustic wavelength at least 2 samples per wavelength [Sapozhnikov et al., 2015] and any spatial variation of v_n due to apodization or element-wise phase programming.

Electronic steering and focusing in the imaging plane (x, z) are represented by complex element weights applied uniformly across each element face [Perrot et al., 2021]. If the transmit focus is prescribed at (x_F, z_F) in that plane, a standard choice is

$$v_n(\mathbf{r}_{nq}, \omega) = A_0\, a_n\, e^{-\mathrm{i}\omega\tau_n}, \tag{2.6}$$

where A_0 is a global amplitude, a_n encodes the element apodization, and τ_n

are the focusing delays given by

$$\tau_n = \frac{\sqrt{(x_n - x_F)^2 + (z_n - z_F)^2}}{c_0}. \tag{2.7}$$

In addition to electronic delays, most imaging arrays incorporate a convex refractive lens bonded on top of the matching layers that focuses acoustically in elevation (y-direction), concentrating the beam to a slice typically below 1 mm in thickness [Lee and Roh, 2017]. The lens introduces a deterministic elevation-dependent phase delay that can be modeled in the thin-lens approximation, by a multiplicative factor in v_n. Denoting by F_e the nominal elevation focal length of the lens and by y the elevation coordinate across the element, a convenient exact form is

$$\Phi_{\text{lens}}(y) = e^{\mathrm{i}k\sqrt{y^2+F_e^2}}, \tag{2.8}$$

which reduces, for $|y| \ll f_e$, to the quadratic thin-lens phase $\Phi_{\text{lens}}(y) \approx e^{\mathrm{i}\frac{k}{2F_e}y^2}$. Incorporating (2.8) into (2.6) yields

$$v_n(x_{nq}, y_{nq}, \omega) = A_0\, a_n\, e^{-\mathrm{i}\omega\tau_n}\; \Phi_{\text{lens}}(y_{nq}), \tag{2.9}$$

so that the elevation focusing provided by the convex lens is accounted for explicitly in the boundary condition that drives the Rayleigh integral.

Angular spectrum projection

The angular spectrum method represents the harmonic pressure field on a plane $z = z_0$ as a superposition of plane waves with transverse wavenumbers (k_x, k_y). The spatial Fourier transform of the complex pressure is

$$P(k_x, k_y, z_0, \omega) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} P(x, y, z_0, \omega)\, e^{-\mathrm{i}(k_x x + k_y y)}\, \mathrm{d}x\, \mathrm{d}y. \tag{2.10}$$

Under the same lossy compressional-wave model used in (2.2), the axial wavenumber associated with each spectral component becomes

$$k_z(k_x, k_y, \omega) = \operatorname{sgn}(z_i - z_0)\, \sqrt{k(\omega)^2 - k_x^2 - k_y^2}, \tag{2.11}$$

Here k_x and k_y are real transverse spatial frequencies, while attenuation is captured through the complex compressional wavenumber $k(\omega)$, yielding a generally complex k_z. Each plane-wave component undergoes a shift proportional to the propagation distance $|z_i - z_0|$, yielding

$$P(k_x, k_y, z_i, \omega) = P(k_x, k_y, z_0, \omega)\, e^{-\mathrm{i}\, k_z(k_x, k_y, \omega)\,(z_i - z_0)}. \tag{2.12}$$

The complex pressure field on the target plane is obtained by the inverse spatial Fourier transform,

$$P(x, y, z_i, \omega) = \frac{1}{(2\pi)^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} P(k_x, k_y, z_i, \omega)\, e^{\mathrm{i}(k_x x + k_y y)}\, \mathrm{d}k_x\, \mathrm{d}k_y, \tag{2.13}$$

which completes the plane-to-plane propagation used for volumetric reconstruction from a single reference plane.

Equations (2.10)–(2.13) implement the angular spectrum projection used for volumetric reconstruction from a single reference plane. For measurement-based back propagation in the frequency domain, time reversal is implemented by conjugating the measured spectrum on the input plane and applying the same forward propagation operator.

In addition, angular restriction should be applied for each plane to avoid nonphysical oscillations. These are caused by reflections on the boundaries due to the periodicity of the Fourier transforms. To mitigate them, the transverse spectrum is windowed with a low-pass circular mask $k_\perp \leq k_c$, where $k_\perp = \sqrt{k_x^2 + k_y^2}$. A practical choice that mitigates wrap-around over a finite lateral window of side D and propagation span $z = |z_i - z_0|$ is

$$k_c = k_r \sqrt{\frac{D^2/2}{D^2/2 + z^2}}, \tag{2.14}$$

commonly implemented [Zeng and McGough, 2008]. It is important to note that modes with sufficiently large $k_\perp$ yield a strongly attenuated axial component through (2.11), combining evanescent decay with material attenuation.

Positioning the RAAS hybrid against other simulators

This subsection situates the hybrid RAAS approach with respect to two open MATLAB toolboxes used for synthesizing acoustic pressure fields. For each

method, the main modeling assumptions and computational time complexity are outlined in variables that arise directly from the simulation, followed by a contrast with the proposed hybrid method.

MUST. The MUST Matlab UltraSound Toolbox[2] implements a linear, time-harmonic acoustic field model in which the emitted pressure is computed in the frequency domain under the assumption of small-signal propagation in a homogeneous background medium [Garcia, 2021, Cigier et al., 2022]. The broadband pressure is represented as a superposition of monochromatic components,

$$p(\mathbf{r}, t) = \Re\left\{\int_{-\infty}^{+\infty} P(\mathbf{r}, \omega)\, e^{-i\omega t}\, \mathrm{d}\omega\right\}, \tag{2.15}$$

so that, for each angular frequency ω, and the toolbox models the acoustic field as a linear time-harmonic pressure $P(\mathbf{r}, \omega)$ radiated by a prescribed aperture vibration [Cigier et al., 2022]. The field from the full array is obtained by computing the contribution of each element and summing them with the appropriate transmit delays and apodization weights [Garcia, 2022]. This formulation is mesh-free in the sense that the field is evaluated directly at user-defined points $\mathbf{r} = (x, y, z)$.

Since the field at each evaluation point is computed as a sum of effective sub-sources, N_s = elements × sub-elementation, the computational work is governed by the source-to-field superposition. A volumetric field built by stacking N_z planes costs

$$\mathrm{cost}_{\mathrm{MUST,\ volume}} = \mathcal{O}\big(N_s\, N^2\, N_z\big), \tag{2.16}$$

where $N = N_x = N_y$ to simplify the expression.

k-Wave. The k-Wave toolbox[3] solves the full wave equation on a discrete numerical grid using a pseudospectral time-domain method with a k-space dispersion corrector [Treeby and Cox, 2010]. In the acoustic configuration considered here, a homogeneous fluid is discretized on a Cartesian grid and the time evolution of the pressure and particle velocity fields is computed for N_t time steps. Each time step involves several three-dimensional FFTs, so

[2] https://www.biomecardio.com/MUST
[3] http://www.k-wave.org/

that the cost per step scales as $\mathcal{O}(N_z N^2 \log N_z N^2)$. The total computational cost is

$$\text{cost}_{\text{k-Wave}} = \underbrace{\mathcal{O}\big(N_t\big)}_{\text{time steps}} + \underbrace{\mathcal{O}\big(N_z\, N^2 \log(N_z\, N^2)\big)}_{\text{3D FFT}} \tag{2.17}$$

which can become substantial for the long bursts and dense three-dimensional grids required in shear-wave elastography studies.

RAAS. The hybrid RAAS route incurs the Rayleigh superposition cost only once on the reference plane, $\mathcal{O}(N_s N^2)$, and generates the entire axial stack by angular-spectrum propagation, i.e., a 2D FFT of the reference plane followed by pointwise spectral multiplication and one inverse Fourier transform per axial plane. The total cost of building a volume is

$$\text{cost}_{\text{RAAS, volume}} = \underbrace{\mathcal{O}\big(N_s\, N^2\big)}_{\text{Rayleigh plane}} + \underbrace{\mathcal{O}\big(N^2 \log N^2\big)}_{\text{2D FFT on } z_0} + \underbrace{\mathcal{O}\big(N_z\, N^2 \log N^2\big)}_{\text{2D IFFT per plane}}. \tag{2.18}$$

In typical imaging configurations N_s is orders of magnitude larger than $\log N^2$, so that, as soon as $N_z > 1$, the axial build in RAAS becomes substantially cheaper than re-summing all sub-sources plane by plane in MUST. MUST remains inexpensive per plane relative to full three-dimensional time-domain solvers, but for volumetric fields the FFT-based stacking of RAAS generally yields shorter runtimes while preserving near-field aperture fidelity on the reference plane. Furthermore, it should be noted that in the context of Rayleigh integral methods, a source grid spacing of two elements per wavelength is adequate for angular spectrum methods. This same value of two elements per wavelength is sufficient for each plane. Nonetheless, pseudospectral simulations mandate a grid comprising four to six elements per wavelength. The temporal sampling required for numerical stability is considerably larger than the Nyquist criterion.

In summary, RAAS is a hybrid Rayleigh-angular-spectrum strategy that preserves near-field source fidelity by imposing the aperture through a mesh-free Rayleigh superposition on a reference plane. It also obtains the full 3D field via FFT-based propagation. This design circumvents the necessity of repeated plane-by-plane source summation, thereby ensuring a substantial enhancement in efficiency for volumetric simulations. Notably, this design maintains a significant reduction in weight compared to full-wave time-domain solvers for the focused beams and three-dimensional pressure fields necessary for radiation-force and shear-wave simulations.

2.2.2 Acoustic radiation force

In the second step of the simulation flowchart, the time–averaged acoustic radiation force that drives shear-wave generation is derived from the harmonic pressure field.

The time-averaged radiation force that loads the elastic medium is constructed by building on the complex pressure field $P(\mathbf{r}, \omega_0)$. All beamforming phases, including electronic steering delays and the elevation–lens contribution, are already embedded in $P(\mathbf{r}, \omega_0)$. A second–order approximation of the ARF in a weakly viscous, weakly absorbing fluid is adopted, which yields a compact vector expression in terms of P and ∇P [Prieur and Sapozhnikov, 2017].

The first-order longitudinal particle velocity follows from linear momentum balance as

$$\mathbf{v}(\mathbf{r}, \omega) = \frac{1}{\mathrm{i}\,\omega\,\rho_0}\,\nabla P, \tag{2.19}$$

and the harmonic acoustic intensity is

$$\mathbf{I}(\mathbf{r}) = \frac{1}{2}\,\Re\{\, P\,\mathbf{v}^* \,\} = -\frac{1}{2\,\omega\,\rho_0}\,\Im\{\, P\,\nabla P^* \,\}, \tag{2.20}$$

where $\Re$ and $\Im$ denote the real and imaginary parts of a complex quantity, respectively.

Within the second-order approximation of the acoustic radiation force in a weakly viscous, weakly absorbing fluid, the body-force can be written as [Prieur and Sapozhnikov, 2017],

$$\mathbf{F}(\mathbf{r}) = \frac{\mathrm{i}\,\alpha_p(\omega)}{2\,\omega\,\rho_0\,c_0}\big(P\nabla P^* - P^*\nabla P\big). \tag{2.21}$$

where $\alpha_p(\omega)$ represents the losses for the longitudinal waves.

Combining Eqs. (2.19)–(2.20) with Eq. (2.21) recovers the compact intensity form

$$\mathbf{F}(\mathbf{r}) = \frac{2\,\alpha(\omega)}{c_0}\,\mathbf{I}(\mathbf{r}), \tag{2.22}$$

which holds without invoking a plane-wave approximation provided $\mathbf{I}$ is constructed from the local fields $(P, \mathbf{v})$ delivered by the RAAS method.

Usually longitudinal wave losses are specified in dB/(MHz$^\gamma$ cm) units, the absorption in Np/m entering Eqs. (2.21)–(2.22) is

$$\alpha_{\mathrm{Np/m}}(\omega) = 100 \times \frac{\ln 10}{20} \left(\frac{\omega}{2\pi\, 10^6}\right)^{\gamma} \times \alpha_{\mathrm{dB/(MHz^\gamma\, cm)}}\,, \tag{2.23}$$

with $\gamma \approx 1$ for many soft tissues [Chivers and Hill, 1975]. This conversion ensures that the experimentally tabulated attenuation coefficients can be used consistently in the ARF computation.

2.2.3 Shear wave simulation

In the third step of the simulation flowchart, the acoustic radiation force vector obtained in Sec. 2.2.2 is used as a 3D body-force source to trigger shear-wave propagation in a three-dimensional viscoelastic solid. This elastic response is simulated with the elastic solvers of the open–source k-Wave MATLAB toolbox using a pseudo–spectral time–domain strategy [Firouzi et al., 2011]. The ARF field is treated as a body–force density $\mathbf{F}(\mathbf{r}, t)$ applied to an isotropic Kelvin–Voigt medium, which is the constitutive model adopted in k-Wave to represent absorbing elastic solids [Treeby et al., 2014, Treeby and Cox, 2010].

The simulator evolves the particle motion through a first–order stress–velocity formulation. Let $\mathbf{u}(\mathbf{r}, t)$ be the displacement and $\mathbf{v} = \partial\mathbf{u}/\partial t$ the particle velocity. Linear momentum conservation with external loading is

$$\rho_0 \frac{\partial \mathbf{v}}{\partial t} = \nabla\cdot\, \boldsymbol{\sigma} + \mathbf{F}, \tag{2.24}$$

where $\boldsymbol{\sigma}$ is the Cauchy stress tensor and ρ_0 is the mass density. For infinitesimal deformations, the kinematics are described by the strain and strain-rate tensors

$$\boldsymbol{\varepsilon} = \mathrm{sym}(\nabla\mathbf{u})\,, \tag{2.25}$$

$$\dot{\boldsymbol{\varepsilon}} = \mathrm{sym}(\nabla\mathbf{v})\,, \tag{2.26}$$

where $\mathrm{sym}(\mathbf{A}) = (\mathbf{A} + \mathbf{A}^T)/2$ denotes the symmetric part of a second-order tensor. These tensors determine the internal stresses through the Kelvin–Voigt law

$$\boldsymbol{\sigma} = \lambda\, \mathrm{tr}(\boldsymbol{\varepsilon})\, \mathbf{I} + 2\mu\, \boldsymbol{\varepsilon} + \chi\, \mathrm{tr}(\dot{\boldsymbol{\varepsilon}})\, \mathbf{I} + 2\eta\, \dot{\boldsymbol{\varepsilon}}. \tag{2.27}$$

Here, λ and μ are the elastic Lamé parameters and χ and η are the Kelvin–Voigt viscosity coefficients. The trace terms $\mathrm{tr}(\boldsymbol{\varepsilon})$ and $\mathrm{tr}(\dot{\boldsymbol{\varepsilon}})$ describe the

volumetric response, whereas the terms proportional to $\boldsymbol{\varepsilon}$ and $\dot{\boldsymbol{\varepsilon}}$ describe the deviatoric response responsible for shear-wave propagation and dissipation. In k-Wave, (2.27) is advanced through the equivalent stress–rate form [Treeby et al., 2014]

$$\dot{\boldsymbol{\sigma}} = \lambda\, \nabla\cdot \mathbf{v}\, \mathbf{I} + 2\mu\, \mathrm{sym}(\nabla \mathbf{v}) + \chi\, \frac{\partial}{\partial t}(\nabla\cdot \mathbf{v})\, \mathbf{I} + 2\eta\, \frac{\partial}{\partial t}[\mathrm{sym}(\nabla \mathbf{v})]\,, \tag{2.28}$$

so that Eqs. (2.24) and (2.28) form a closed first–order system in $(\boldsymbol{\sigma}, \mathbf{v})$ discretized on a Cartesian grid and propagated using the k–space pseudospectral scheme.

The material parameters are specified through measurable wave speeds and attenuation coefficients. The low–frequency compressional and shear speeds, c_p and c_s, determine the elastic moduli via

$$\mu = \rho_0 c_s^2, \qquad \lambda + 2\mu = \rho_0 c_p^2, \tag{2.29}$$

while the quadratic absorption coefficients α_p and α_s determine the Kelvin–Voigt viscosity parameters,

$$\alpha_p = \frac{\chi + 2\eta}{2\rho_0 c_p^3}, \qquad \alpha_s = \frac{\eta}{2\rho_0 c_s^3}. \tag{2.30}$$

For the shear component, Appendix A derives the dispersion and attenuation relations associated with the divergence-free mode $\nabla\cdot \mathbf{v} = 0$, for which the volumetric terms vanish and the propagation is governed by the shear parameters (μ, η).

In order to reduce runtime, the simulations utilize a hypothetical compressional speed c_p, since a tissue-like value (~ 1540 m/s) would primarily constrict the CFL-limited time step without exerting an effect on the shear response of interest. The selection of c_p is constrained by compressibility, as demonstrated by equation Eq. (2.29), which establishes the relationship $\lambda/\mu = (c_p^2/c_s^2) - 2$. This means that c_p must remain sufficiently larger than c_s to keep the motion approximately divergence-free. Following [Prieur and Catheline, 2016], we set $c_p = 30$ m/s, a pragmatic choice for the soft-tissue shear-speed range considered here ($c_s \approx 0.5$–10 m/s), corresponding to a compressional-to-shear speed ratio $c_p/c_s \sim 3$–60, which keeps the response shear-dominated while relaxing the stable time step.

The ARF that drives the elastic simulation is computed from the monochromatic pressure field using the second–order expression in Eq. (2.21), yielding

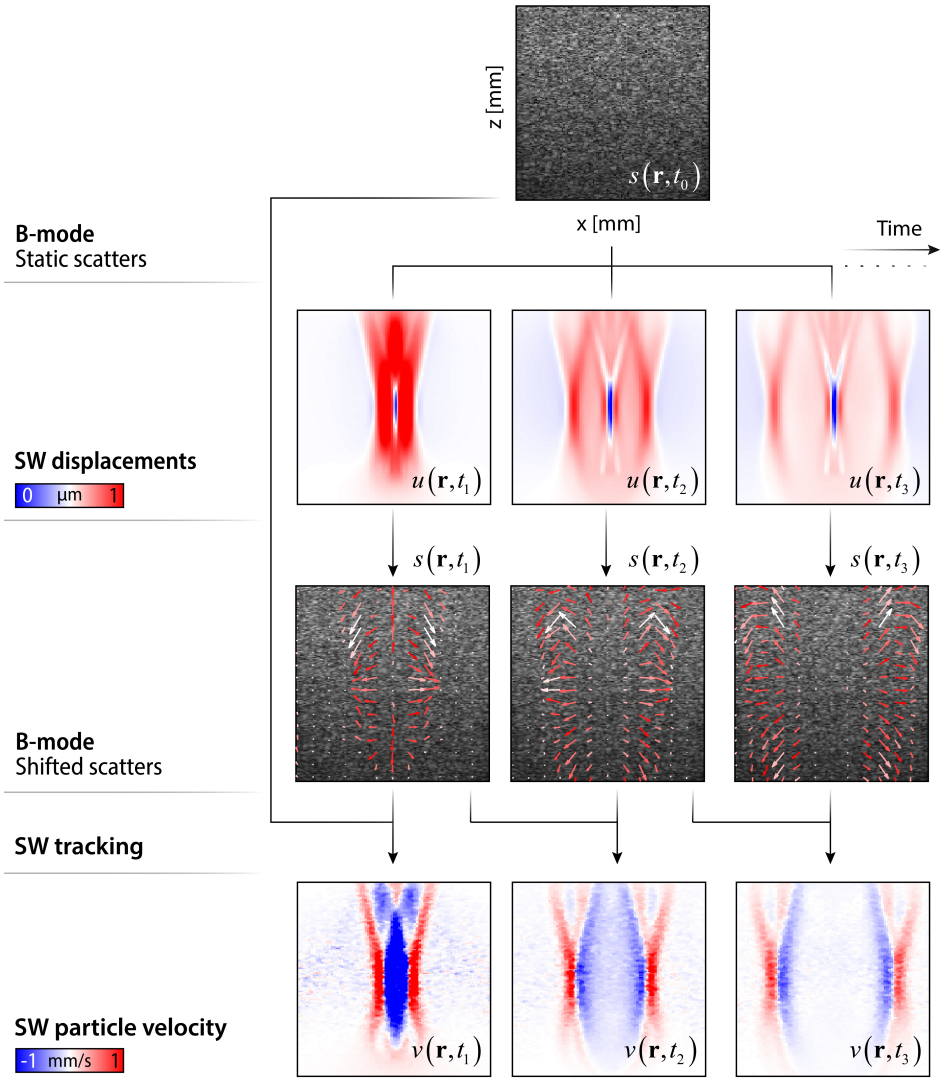


Figure 2.3: Flowchart for shear wave particle velocity calculation using Loupas autocorrelation from synthetic b-mode frames with embedded shear wave displacements

a vector-valued body-force density $\mathbf{F}(\mathbf{r}, \omega_0)$ on the solver grid. In the k-Wave injection scheme, volumetric body forces are applied as equivalent particle-velocity sources that can be written as [Prieur and Catheline, 2016]

$$\mathbf{v}_{\mathbf{F}}(\mathbf{r}, \omega_0) = \frac{1}{\rho_0} \mathbf{F}(\mathbf{r}, \omega_0) \odot \frac{d\mathbf{h}}{2\, c_0}, \tag{2.31}$$

where $d\mathbf{h} = (dx, dy, dz)^\top$ and $\odot$ denotes the Hadamard (element-wise) product.

Finally, the elastic solver returns the particle velocity as a vector field $\mathbf{v}(\mathbf{r}, t)$ over the full 3D domain. The corresponding displacement field $\mathbf{u}(\mathbf{r}, t)$ is obtained by temporal integration of each component,

$$u_i(\mathbf{r}, t) = \int_0^t v_i(\mathbf{r}, \tau)\, \mathrm{d}t, \qquad i \in \{x, y, z\}. \tag{2.32}$$

Subsequent processing focuses on the axial component z, since it is the component directly accessible in ultrasound experiments along the imaging beam.

2.2.4 Tracked shear wave motion generation

In order to mirror the ultrasound scanner acquisition used in realistic scenarios, we propose a synthetic imaging stage, providing a bridge between the simulated motion and the realistic measured complex IQ frames. Figure 2.3 illustrates the proposed method in its entirety. In this manner, the generation of synthetic ultrasound data is initiated through the embedding of simulated shear-wave displacements into a moving two-dimensional speckle-generating medium. Second, for each temporal frame, the ultrasound backscattering is simulated by calculating the RF signals using the MUST toolbox. Finally, a set of IQ frames are generated and local motion is estimated using an autocorrelation procedure.

Initialize reference frame

The speckle-generating medium is represented by a cloud of point scatters uniformly distributed over the same lateral and axial region as the elastic simulation field of view. Each scatter is assigned a random scattering amplitude drawn from a zero-mean Gaussian distribution, so that the resulting ultrasound IQ images exhibit fully developed speckle pattern. The initial scatterer

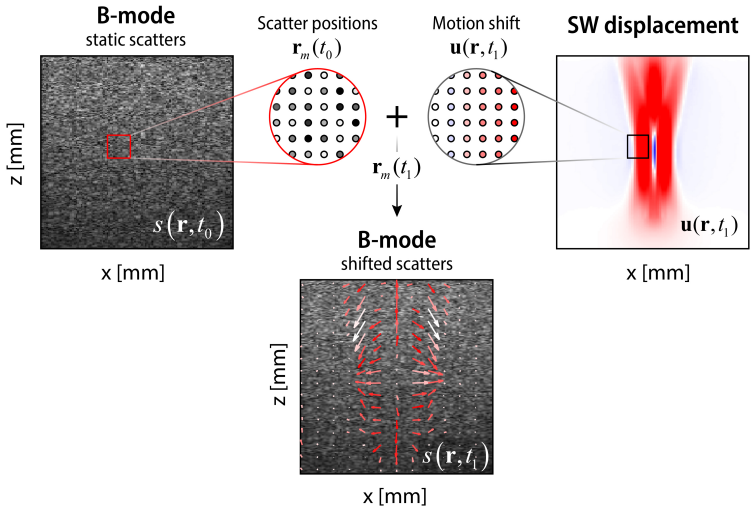


Figure 2.4: B-mode with embedded shear wave motion

positions, denoted by $\mathbf{r}_m(t_0) = (x_m, z_m)$, where t_0 denotes the reference initial time, are defined as continuous floating-point coordinates.

Shear wave displacement field

As illustrated in Fig. 2.4, shear-wave motion is embedded by advecting the scatterer cloud according to the simulated in-plane displacement field. At each synthetic frame time t_n, the displacement vector $\mathbf{u}(\mathbf{r}, t_n)$ is interpolated at the initial scatterer location $\mathbf{r}_m(t_0)$. The updated position of the m-th scatterer is then obtained as

$$\mathbf{r}_m(t_n) = \mathbf{r}_m(t_0) + \mathbf{u}\big(\mathbf{r}, t_n\big), \tag{2.33}$$

Update scattered medium

For each frame, the backscatterd RF signals are calculated using SIMUS algorithm [Garcia, 2022, Cigier et al., 2022] using the updated scatter distribution. The transmit configuration reproduces the ultrasound scanner tracking sequence. For example, in this case we use a set of compounding plane waves steered at small and we compound a final image by coherent plane wave

compounding [Montaldo et al., 2009] and DAS beamforming [Karaman et al., 1995, Thomenius, 1996]. The RF data are beamformed on a Cartesian image grid with lateral and axial sampling dx and dz. Finally, the beamformed RF signals are demodulated around the imaging center frequency, yielding complex IQ images $s(\mathbf{r}, t_n)$.

Shear wave tracking estimation

Subsequent to the acquisition of the set of complex image data $s(x, z, t_n)$, the shear-wave particle velocity is extracted directly from the complex ultrasound images.

$$s(\mathbf{r}, t_n) \;=\; I(\mathbf{r}, t_n) + \mathrm{i}\, Q(\mathbf{r}, t_n), \tag{2.34}$$

where I and Q are the in-phase and quadrature components provided either by the synthetic imaging chain (Sec.2.2.4) or by the ultrasound scanner in the experiments (Sec.2.3.3). The additional inputs to the estimator are the speed of sound used in beamforming, c_0, and the imaging central frequency f_0 and the pulse repetition period T_{PRF}.

The phase-based tracking follows the 2D autocorrelation approach of Loupas [Loupas et al., 1995], adapted here to tissue motion and shear-wave elastography. For each pair of consecutive frames at times t_n and $t_{n+1} = t_n + T_{\mathrm{PRF}}$, and for each image location $\mathbf{r} = (x, z)$, a local lag-one autocorrelation is computed over a small axial window centred at depth z,

$$R(\mathbf{r}, t_n) = \sum_{\tilde{z} \in \Omega_z} w(\tilde{z})\, s\big(\tilde{\mathbf{r}}, t_{n+1}\big)\; s^*\big(\tilde{\mathbf{r}}, t_n\big), \tag{2.35}$$

where $\tilde{\mathbf{r}} = (x, \tilde{z})$ denotes the positions inside the axial range gate, Ω_z denotes the set of axial samples inside the range gate around z, $w(\tilde{z})$ is a tapering window, and $(\cdot)^*$ denotes complex conjugation. This corresponds to the baseband version of the 2D autocorrelator, restricted here to a single slow-time lag and a short axial gate, consistent with Eq. (55) of [Loupas et al., 1995].

The argument of R gives the mean inter-pulse phase shift at that location,

$$\Delta\phi(\mathbf{r}, t_n) = \arg\big\{R(\mathbf{r}, t_n)\big\}, \tag{2.36}$$

which is related to the incremental axial displacement under a narrowband approximation by

$$\Delta u_z(\mathbf{r}, t_n) = \frac{c_0}{4\pi f_0}\, \Delta\phi(\mathbf{r}, t_n) = \frac{\lambda_0}{4\pi}\, \Delta\phi(\mathbf{r}, t_n). \tag{2.37}$$

The corresponding axial particle velocity between frames t_n and t_{n+1} is then approximated as

$$v_z\Big(\mathbf{r}, t_n + \tfrac{T_{\mathrm{PRF}}}{2}\Big) \approx \frac{\Delta u_z(\mathbf{r}, t_n)}{T_{\mathrm{PRF}}} = \frac{\lambda_0}{4\pi\, T_{\mathrm{PRF}}}\, \Delta\phi(\mathbf{r}, t_n). \tag{2.38}$$

Repeating this operation for all frame pairs yields a time-resolved estimate $v_z(\mathbf{r}, t)$ of the axial shear-wave motion.

2.3 Application example

This section presents an end-to-end application example of the proposed hybrid method, using a representative shear-wave elastography configuration driven by a linear-array transducer L11-5v (Verasonics Inc., Redmond, WA), whose main parameters are summarized in Table 2.1. First, the monochromatic acoustic pressure field generated by a focused pushing beam is computed with the RAAS front-end and benchmarked against two established simulation platforms (MUST and k-Wave) as well as experimental measurements. The validated harmonic pressure field is then used to derive the 3D acoustic radiation force vectorial field, which provide the mechanical source term for subsequent elastic simulations. Shear-wave propagation is finally evaluated consistently across three data modalities: full shear-wave elastic simulations, ultrasound-like synthetic IQ frames generated from the simulated motion, and experimental IQ acquisitions obtained with the same probe and sequence.

Parameter	Symbol	Value	Units
Probe model	-	L11-5v	-
Centre frequency	f_0	7.6	MHz
Number of elements	N_{el}	128	-
Pitch	p	0.30	mm
Kerf	g	0.03	mm
Elevation height	L_e	5	mm

Table 2.1: Main parameters of the L11-5v linear-array probe used in the acoustic simulations and hydrophone measurements.

2.3.1 Validating the acoustic pressure field

Numerical and experimental parameters

The hybrid acoustic stage was implemented and validated using a common transmit configuration in three numerical simulators and in a hydrophone measurement setup. The propagation medium was water, with density $\rho_0 = 1000$ kg/m^3 and sound speed $c_0 = 1540$ m/s, matching the experimental water-tank conditions.

A focused pushing beam was synthesized by driving a subset of 34 active elements with a burst at f_0, using an f-number of 2 and an electronic axial focus at $z_F = 20$ mm. This configuration produces an axially elongated focal region that favors the generation of an almost planar shear-wave front in depth, which is convenient for the plane-wave propagation assumptions of the velocity estimators used in later chapters. The same delay law, apodization and pulse length were imposed in all simulations and experiments.

For the acoustic simulations with the three analyzed simulators, a Cartesian grid was defined with uniform spacing = 0.05 mm in the three spatial directions, corresponding to four points per wavelength at f_0.

Acoustic simulations were conducted using the RAAS, MUST and k-Wave simulators. The simulation grids for the three simulators were found to be identical and exhibited uniform spacing, denoted as dh, with a value of = 0.05 mm in all three spatial directions. This configuration resulted in four points per wavelength at f_0. The grid extended 10.2 x 10.2 mm in the x/y direction, matching the active aperture of the array and covering the full elevation height. The axial direction z spanned from 0 to 30 mm, thus including the focal zone and the region where shear-wave generation is later analyzed.

In the proposed RAAS simulator, the array aperture was discretized into sub-elements in order to represent both the azimuthal apodization and the elevational focusing imposed by the convex acoustic lens of the L11-5v probe. Each physical element was divided into 5×76 patches, yielding a total of 380 sub-elements. The additional phase delays required to reproduce the elevational focus of the lens were applied at the sub-element level. The Rayleigh integral (2.4) was evaluated on a plane at $z_0 = 1$ mm from the exit plane. This initial complex pressure map $p(x, y, z_0; \omega_0)$ was then propagated

forward in steps of dh up to $z = 30$ mm by means of the angular-spectrum method described in Sec. 2.2.1, which provided the three-dimensional harmonic field $P_{\mathrm{RAAS}}(x, y, z; \omega_0)$ on the reference grid.

The MUST toolbox was configured through the SIMUS algorithm with the same transducer geometry, medium parameters and transmit sequence. The L11-5v array was modeled as a set of rectangular pistons with the same pitch and elevation height, subdivided into sub-elements in the azimuthal and elevational directions to mirror the discrete representation used in RAAS. SIMUS computes the complex pressure field over the frequency band of the transducer. In this work the complex harmonic component at the transmit frequency, $P_{\mathrm{MUST}}(x, y, z; \omega_0)$, was extracted from the SIMUS output. This choice ensures a consistent comparison with the monochromatic fields produced by RAAS and by the acoustic k-Wave model.

The k-Wave simulations were run in a homogeneous fluid domain with the same spatial grid and medium properties, using a Courant–Friedrichs–Lewy number of 0.3, 4 points per wavelength and a perfectly matched layer of 10 grid points on each side. The L11-5v array was represented as a distribution of normal surface velocity on the source plane, with the same sub-element partition and elevational focusing delays as in RAAS. The simulation was driven by a 5-cycle burst at f_0. The time-domain pressure waveforms at all grid points were then Fourier transformed and the complex amplitude at ω_0 was retained to form the harmonic field $P_{\mathrm{kW}}(x, y, z; \omega_0)$.

Experimental reference data were acquired in a water tank filled with degassed and deionised water maintained at 22–23 °C using a water conditioning unit (WCU Series, Sonic Concepts, USA). The linear-array probe was driven by a 50-cycle sinusoidal burst at f_0 using the Verasonics Vantage 256 research platform (Verasonics Inc., Redmond, WA). Acoustic signals were measured with a calibrated needle hydrophone (model HNA-0400, ONDA, Castellon), with a 400 μm sensitive element and a flat response over the frequency range of interest. The hydrophone was mounted on a three-axis motorized translation stage (5 μm resolution, PI Micos GmbH, Germany), and the pressure waveforms were acquire using a oscilloscope (14bits, 100 MS/s, model PXI5412, National Instruments, USA). A planar scan was performed at $z = 10$ mm, covering the interval -15 mm $< x <$ 15 mm and -15 mm $< y <$ 15 mm. Ten repetitions were averaged at each position to reduce noise. The recorded waveforms were band-pass filtered around f_0. The extraction of the complex pressure on the measurement plane was achieved through the implementation

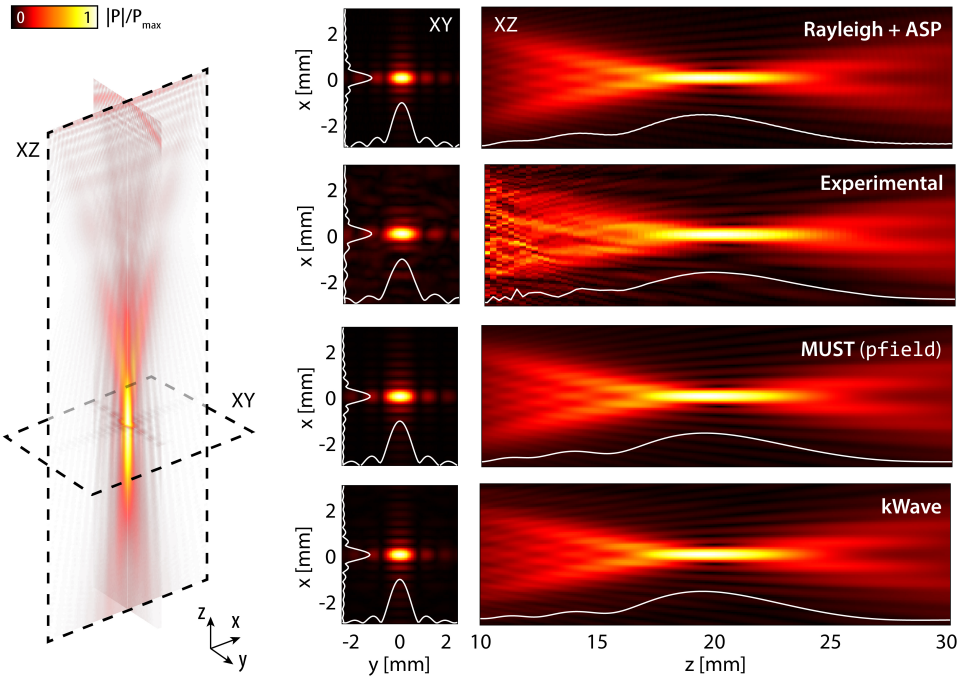


Figure 2.5: Comparison of acoustic pressure fields predicted by RAAS, MUST, k-Wave, and experiments.

of a temporal Fourier transform. This experimental plane was then forward-propagated to $z = 30$ mm using the same angular-spectrum method and spatial step dh as in the RAAS model, yielding an experimental three-dimensional harmonic field $P_{\text{exp}}(x, y, z; \omega_0)$ on the reference grid.

In all four cases the comparison is carried out in terms of the complex harmonic pressure at the transmit frequency.

Results

Figure 2.5 illustrates the overall structure of the focal beam. A volumetric rendering of the normalized pressure predicted by the hybrid model is shown together with the axial and elevational planes analyzed in this work. The corresponding xy and xz planes reveal that the main focal lobe, the axial beam convergence, and the pattern of grating and side lobes are visually

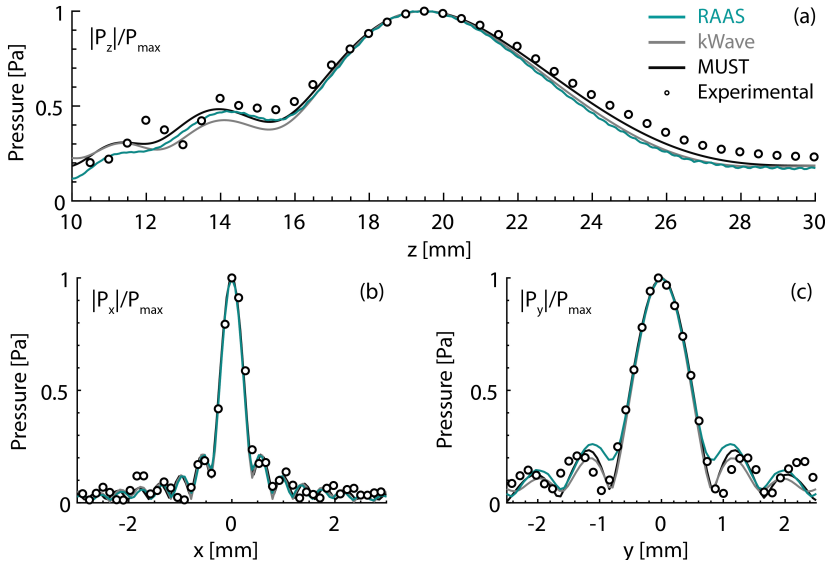


Figure 2.6: Acoustic beam pressure profile comparison for the L11-5v probe using the RAAS (green), k-Wave (gray), MUST (black) simulators and experimental measurements (dots). Plots represent the profiles for the axial (a), lateral (b) and elevational (c) directions.

consistent across all approaches. The experimental maps exhibit a slightly noisier background and mild asymmetries, whereas MUST and k-Wave produce smoother fields that closely resemble the hybrid solution.

A quantitative comparison in the focal region is summarized in Fig. 2.6. The transmit sequence was designed for a nominal focal depth of 20.04 mm. The experimental depth of maximum pressure is $z_{\text{focus}} = 19.48$ mm, whereas the hybrid model, MUST, and k-Wave predict 19.43, 19.47, and 19.47 mm, respectively. All values therefore lie within 0.5 mm of the design focus and within 0.05 mm of the experimental measurement. The full width at half maximum (FWHM) defined at -6 dB beamwidths in the lateral, elevational, and axial directions, derived from the same pressure maps, are reported in Table 2.2.

The agreement between simulated and measured profiles in the focal region is further quantified using the normalized root mean squared error (NRMSE) metric computed over a region of interest spanning from 10 to 30 mm in depth, from -3 to 3 mm laterally, and from -2.5 to 2.5 mm elevationally, with sampling steps identical to those used in the simulations and measurements.

FWHM [mm]	X	Y	Z
RAAS	0.49	1.00	7.48
MUST	0.51	1.02	7.96
k-Wave	0.49	0.98	7.56
Experimental	0.50	1.03	8.04

Table 2.2: FWHM of the normalized pressure field for the RAAS model, MUST, k-Wave, and the experimental measurement.

NRMSE	P_x	P_y	P_z
RAAS	0.0015	0.2344	0.0025
MUST	0.0017	0.1289	0.0017
k-Wave	0.0022	0.1546	0.0021

Table 2.3: NRMSE between simulated and experimental pressure profiles in the focal region, computed over a region of interest spanning from 10 to 30 mm in depth, from −3 to 3 mm laterally, and from −2.5 to 2.5 mm elevationally. Lower values indicate better agreement.

The NRMSE is an error norm normalized by the variability of the reference signal, so that NRMSE = 0 indicates a perfect match and larger values indicate increasing mismatch between x and x_{ref}. Table 2.3 reports the metric values for the axial, lateral, and elevational profiles. The hybrid model and the full-wave simulators achieve almost perfect agreement with the experimental reference for the axial and lateral profiles, while the elevational agreement is more modest but remains within the range of discrepancies observed between the different simulation platforms.

These results demonstrate that the RAAS model reproduces the focal position, beamwidth and on-axis pressure profile of the pushing beam with accuracy comparable to that of MUST and k-Wave, while matching closely the experimental hydrophone measurements. The good agreement in the axial and lateral profiles confirms that the discretized aperture and elevation lens are adequately captured in the hybrid model, providing a reliable acoustic input for the subsequent radiation-force and shear-wave simulations.

The computational performance of the hybrid acoustic model was evaluated by measuring the execution time required to generate volumetric pressure fields for different grid sizes and by comparing these measurements with the analytical cost expressions derived in Sec.2.2.1. The same transmit configuration and spatial sampling described in Sec.2.3.1 were used for all simulations, and the

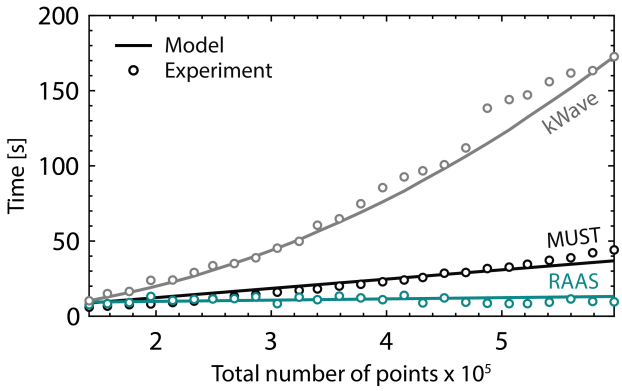


Figure 2.7: Execution time as a function of the total number of grid points N^2N_z for the RAAS model, MUST and k-Wave. Solid lines indicate least-squares fits to the analytical complexity laws derived for each simulator.

total number of grid points N^2N_z was varied by adjusting the number of axial planes while maintaining the fixed lateral (N_x) and elevational (N_y) dimensions.

Figure 2.7 reports the execution time as a function of the total number of grid points for the three simulators. The empirical trends follow closely the theoretical complexity laws: the MUST runtime grows approximately linearly with N_z for fixed N^2, consistent with the plane-by-plane source summation. The k-Wave runtime scales as Eq. (2.17), reflecting the cost of the three-dimensional FFTs performed at each time step. The hybrid RAAS model exhibits a markedly slower growth with N_z, since the Rayleigh summation is incurred only once on the reference plane and the axial build is dominated by FFTs on two-dimensional planes.

For the volumetric fields considered in this chapter, the hybrid model achieves a substantial reduction in execution time relative to MUST and k-Wave, particularly as the number of axial planes increases. This favourable scaling underpins the choice of RAAS as the acoustic stage in the proposed shear-wave simulation method, enabling the generation of dense three-dimensional pressure fields with a computational cost compatible with extensive parameter sweeps and three-dimensional elastography studies.

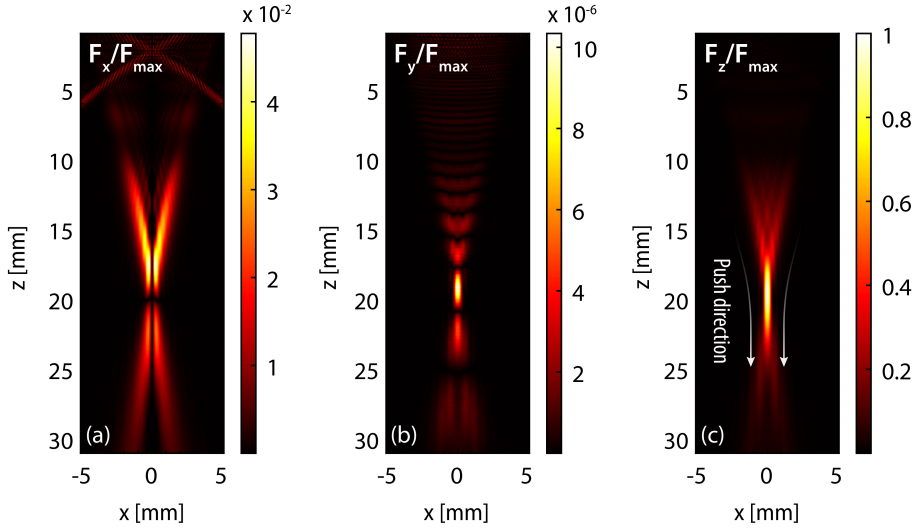


Figure 2.8: Components of the acoustic radiation force predicted by the hybrid Rayleigh and angular spectrum model. Maps show the xz-plane at $y = 0$ for $F_x/F_{\max}$, $F_y/F_{\max}$, and $F_z/F_{\max}$, with $F_{\max} = \max(F_z)$.

2.3.2 Construction of volumetric ARF maps

The complex pressure map $P(x, y, z; \omega_0)$ obtained on the reference grid in Sec. 2.2.1 was converted into a three-dimensional ARF distribution using Eq. (2.21). The ARF distribution analyzed in this step is that predicted by the hybrid RAAS model, using the monochromatic pressure at the transmit center frequency. Since Sec. 2.3.1 already established the agreement between simulated and measured pressure fields, only the force field generated by RAAS is examined here.

In this section, the simulated medium differ from those shown in Sec. 2.3.1, which assumed a propagation medium simulating water. Material properties studied in this study corresponded to the mechanical properties of the CIRS 039 ultrasound elasticity phantom (Computerized Imaging Reference Systems, Norfolk, USA) utilized in the subsequent shear wave experiments. The longitudinal propagation velocity and density values are maintained constant as in water like medium, but a typical soft tissue longitudinal absorption term, $\alpha_{\mathrm{L}} = 0.5$ dB/cm/MHz, is added.

Figure 2.8 shows the xz-plane of the three force components, F_x, F_y, and

F_z, normalized by the maximum axial force $F_{\max} = \max(F_z)$. The axial component F_z in Fig. 2.8(c) exhibits a single dominant focal lobe centered at a depth of $z = 19.5\,\text{mm}$, close to the nominal focal depth of 20.04 mm defined by the transmit sequence. The corresponding axial FWHM measured at $(x, y) = (0, 0)$ is 5.27 mm, indicating that the effective push region is extends along z direction. The force maps shown in Fig. 2.8 are restricted to a region of interest spanning from -5 to 5 mm in the lateral direction and from 0 to 30 mm in depth, encompassing the entire focal zone relevant for shear-wave generation.

The lateral components of the radiation force are substantially smaller than F_z. The F_x map in Fig. 2.8(a) presents a symmetric dipole-like pattern around the beam axis in the focal region, with a maximum amplitude $\max(F_x)/F_{\max} = 0.05$, indicating that the transverse force along x reaches only a few percent of the axial push. As illustrated in Fig. 2.8(b), the F_y component is several orders of magnitude weaker, with $\max(F_y)/F_{\max} = 1.05 \times 10^{-5}$, and displays a low-amplitude oscillatory pattern along z, consistent with the reduced aperture in the elevational direction.

Overall, the radiation force field is strongly dominated by the axial component, defining a well-localized push region aligned with the imaging axis that acts as the primary source term for the subsequent shear-wave propagation simulations.

2.3.3 Shear-wave motion validation

Simulation and experimental protocol

The elastic propagation of shear waves was modeled using four commercial liver fibrosis phantoms from the CIRS 039 Ultrasound Elasticity Set (Computerized Imaging Reference Systems, Norfolk, USA), labelled as Phantom 1 to Phantom 4. The manufacturer specifies central shear-wave speeds of 0.8, 1.6, 2.5 and 3.5 m/s, which were used as target shear-wave velocities in the elastic model. The methodology analyzed in the next subsection 2.3.3 combines three types of datasets for each phantom: simulated shear-wave motion data, synthetic motion from simulated acquisitions, and experimental measurements.

In the numerical simulations, each phantom is modelled as a homogeneous Kelvin–Voigt solid with density $\rho_0 = 1000\ \text{kg/m}^3$ and a hypotheti-

cal longitudinal wave speed $c_p = 30$ m/s. The target shear speeds c_s are imposed via Eq. (2.29), and the shear viscosity is set to $\eta = 0$, so that shear-wave attenuation is negligible over the distances of interest. Longitudinal wave losses are included through a constant attenuation coefficient $\alpha_{\mathrm{L}} = 0.5$ dB/cm/MHz. The elastic problem is solved on a Cartesian grid with uniform spacing $dx = dy = dz = 0.1$ mm, covering the full lateral aperture of the push and an axial range that includes both the focal region and the propagation path over which the shear wave is analysed. Perfectly matched layers of 10 grid points are placed on all faces to absorb outgoing waves. The time step is chosen assuming a CFL number of 0.3, resulting in $dt = 3.75$ μs. For each phantom, the total simulation time is set so that a shear wave travelling at the corresponding c_s crosses the region of interest once before reaching the absorbing boundaries.

The volumetric particle-velocity sources derived from the ARF field in Sec. 2.2.2 are mapped onto this elastic grid and applied with a rectangular temporal envelope of 100 μs, matching the duration of the experimental push. The same spatial source distribution is used in all four phantoms; only the KV parameters change between simulations. The axial particle-velocity component $v_z(x, y, z, t)$ is recorded. Output is low-pass filtered and downsampled to $dt = 25$ μs to reduce data volume while preserving the shear-wave bandwidth. For comparison with IQ-based reconstructions, the analysis is restricted to the central elevational plane $y_0 = 0$, yielding a two-dimensional motion field $v_z(x, y_0, z, t)$.

Synthetic IQ frames are generated using the same probe and field of view. For the motion-tracking imaging, a plane-wave compounding scheme with 3 steering angles $-3°$, $0°$ and $+3°$ are used, together with a pulse repetition period $T_{\mathrm{PRF}} = 100$ μs. The image grid is defined with $dx = dz = 0.144$ mm, and delay-and-sum beamforming with rectangular apodisation is applied on the full aperture. RF data are sampled at four times the imaging central frequency.

Finally, experimental shear-wave elastography measurements are performed in the CIRS phantoms. Each phantom is placed on a flat support and imaged from above with the L11-5v probe, ensuring good acoustic coupling. All acquisitions are carried out on a Verasonics Vantage 256 platform. The ARF excitation consists of a 100 μs burst at 5.5 MHz, driven at an input voltage of 50 V and electronically focused at $z_{\mathrm{foc}} = 20.04$ mm on the central axis with 34 active elements (f-number 2). Shear-wave propagation is tracked with a

high-frame-rate sequence: after the push, plane waves at $-3°$, $0°$ and $+3°$ are emitted and compounded in reception; the pulse repetition period between compounded frames is $T_{\mathrm{PRF}} = 100\ \mu$s, giving an effective frame rate of 10 kHz. The tracking window spans approximately 15 ms, long enough to capture the passage of the shear wave across the field of view in all four phantoms. The imaging region covers the full lateral aperture and an axial depth of about 35 mm, matching the synthetic datasets. The imaging system provides the complex phase and quadrature (IQ) data. For each phantom the resulting sequence $s(x, z, t)$ is stored with the same spatial sampling and frame interval as the synthetic IQ data.

Results

Figures 2.9(a-l) show the particle velocity obtained by the three approaches for Phantom 3, which has a nominal shear modulus of 16.04 kPa. For each dataset, four snapshots of $v_z(x, z, t)$ are shown at successive times $t = 1.35$, 2.25, 3.15, and 4.05 ms. The spatio-temporal maps correspond to the focal region of the push, obtained by averaging a 1 mm-thick window around the focal depth $z = 20.04$ mm, and cover a lateral range from -16.8 to 16.8 mm. The sequence of frames shows the lateral propagation of the shear wave away from the push location, with comparable wavefront curvature and symmetry in the three cases. For each time frame, lateral profiles at the focal depth are displayed in Figs. 2.9(m-p), allowing a direct comparison of the shear wave shape across simulation, synthetic, and experimental data. Despite the independent origins of the three datasets, the normalized profiles exhibit very similar temporal evolution and lateral structure.

The analysis was then extended to the four phantoms in the series. Figure 2.10 shows, for each phantom, waterfall representations of $v_z(x, t)$ at the focal depth. The maps reveal a systematic increase in the apparent propagation slope as phantom stiffness increases, consistent with higher shear wave speeds in the stiffer phantoms. This trend is reflected in the arrival time of the main wavefront at a fixed lateral position $x = -6.5$ mm, which decreases from approximately 6.8 ms in Phantom 1 to 1.7 ms in Phantom 4.

For each phantom, the bottom panels of Fig. 2.10 compare the temporal profiles of $v_z(t)$ at this fixed spatial location for the three datasets. The curves, normalized by their respective maxima, show a similar waveform morphology and timing in simulations, synthetic reconstructions, and experimental mea-

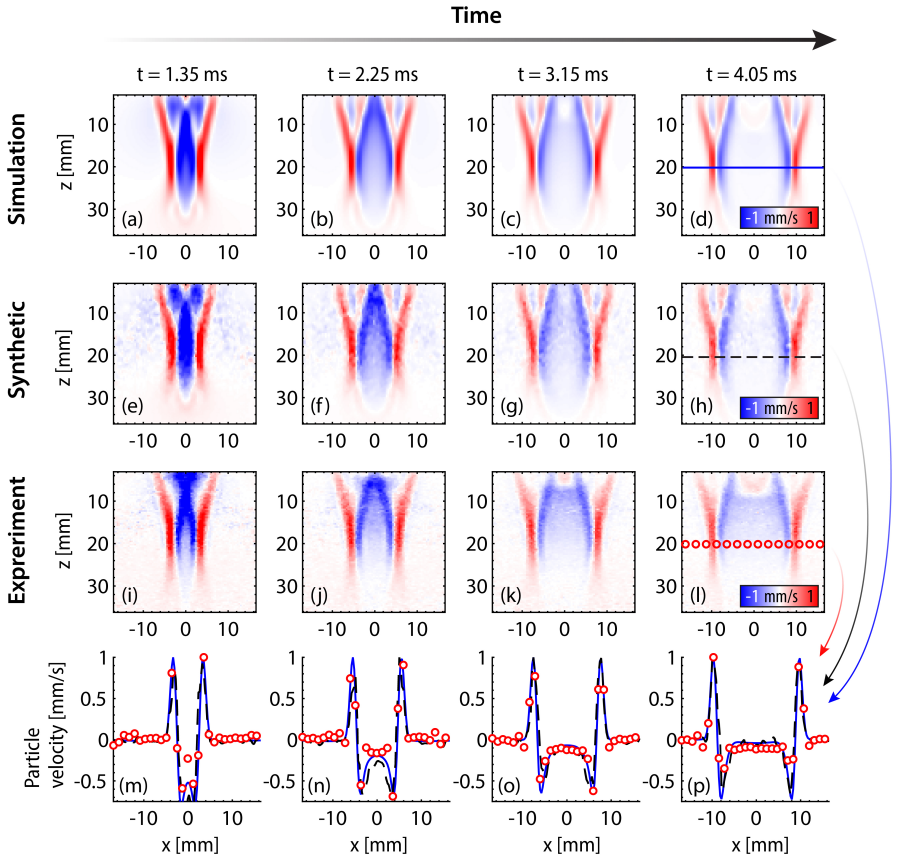


Figure 2.9: Comparison of normalized axial particle velocity $v_z(x, z, t)$ in Phantom 3 for full-wave simulations, synthetic reconstructions, and experimental measurements. Top panels show shear wave propagation in the XZ-plane at four time points, and the corresponding lateral profiles at the focal depth.

surements, indicating that the hybrid method preserves the main kinematic features of the shear wave across the range of stiffness considered. A quantitative analysis of the corresponding group and phase velocities will be presented in the following subsection.

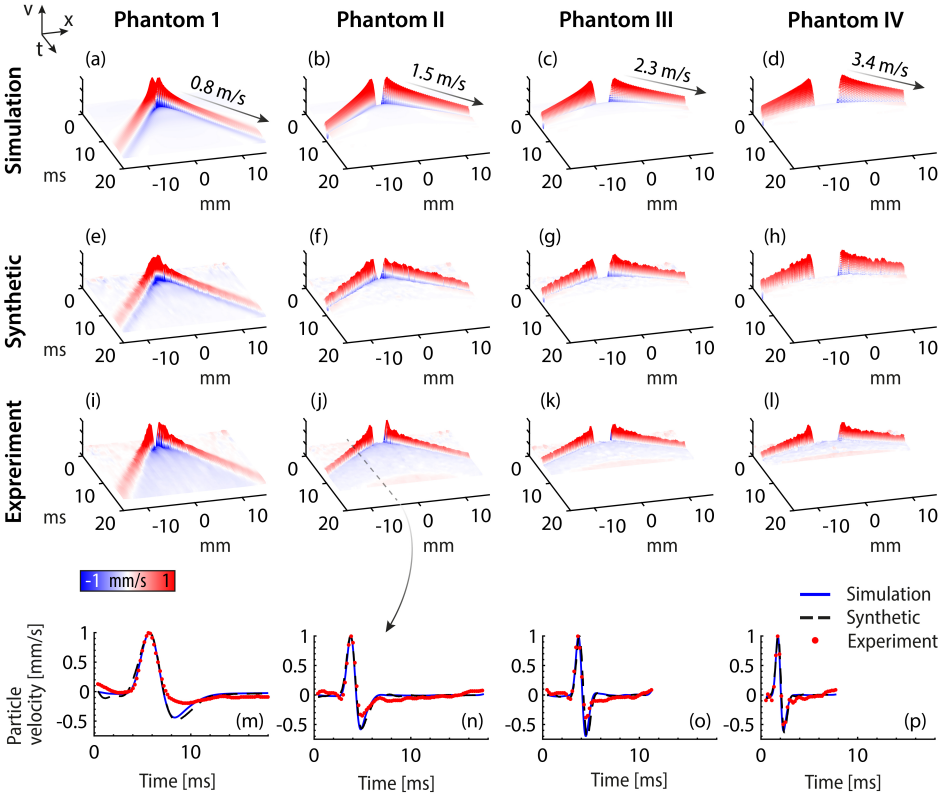


Figure 2.10: Waterfall representations of normalized axial particle velocity $v_z(x,t)$ at the focal depth for the four phantoms, comparing simulations, synthetic reconstructions, and experiments. The increasing slope of the shear wavefront with phantom stiffness and the temporal profiles at $x = -6.5$ mm illustrate the consistent propagation behaviour recovered by the three datasets.

2.4 Shear-wave velocity characterization

Finally, the shear-wave fields obtained from full-wave simulations, synthetic IQ reconstructions and experimental measurements are reduced to scalar metrics that can be compared across phantoms and datasets. Two complementary quantities are considered: the shear-wave group velocity c_g, which characterizes the propagation speed of the wave packet envelope, and the phase velocity $c_{\mathrm{ph}}(f)$, which describes the frequency-dependent speed of individual harmonic components. These metrics are extracted from the axial particle-velocity maps $v_z(x,z,t)$ estimated in Sec. 2.2.3 and are used to quantify stiffness and

viscoelastic dispersion in homogeneous and heterogeneous media.

2.4.1 Group and phase velocity metrics

For a shear disturbance propagating predominantly along the lateral direction x, the axial particle-velocity field at the focal depth can be approximated as a quasi-plane wave $v_z(x,t)$ with a slowly varying envelope. In this setting, two distinct propagation speeds are relevant. The group velocity $c_g(\omega) = d\omega/dk$ quantifies the frequency-dependent transport of the wave-packet envelope and, in practice, time-of-flight estimators yield an effective group-velocity value associated with a chosen feature of the band-limited packet (e.g., envelope maximum or correlation peak). The phase velocity $c_{\mathrm{ph}}(f)$ describes the speed of individual harmonic components at frequency f, and therefore directly captures the dispersive behaviour of viscoelastic media.

In non-dispersive elastic media, both velocities coincide and reduce to a single shear-wave speed c_s. In viscoelastic materials, however, the complex shear modulus becomes frequency dependent and induces dispersion. Under the Kelvin–Voigt rheology introduced in Sec.2.2.3, the shear response is governed by the elastic modulus μ and the viscosity η, which combine into a complex modulus $\mu^*(\omega) = \mu + \mathrm{i}\omega\eta$. Inserting this model into the shear-wave equation yields the phase-velocity dispersion relation

$$c_{\mathrm{ph}}(\omega) = \sqrt{\frac{2\left(\mu^2 + \omega^2\eta^2\right)}{\rho_0\left(\mu + \sqrt{\mu^2 + \omega^2\eta^2}\right)}}, \tag{2.39}$$

which monotonically increases with frequency when $\eta > 0$. In the purely elastic limit $\eta \to 0$, Eq. (2.39) reduces to the non-dispersive shear speed $c_{\mathrm{ph}}(\omega) = \sqrt{\mu/\rho_0} = c_s$. The Appendix A provides a detailed exposition of the derivation of the phase velocity for a plane shear wave propagating in one dimension.

Throughout the thesis, group-velocity estimates c_g provide a compact, scalar descriptor of the effective stiffness in a given region, while phase-velocity curves $c_{\mathrm{ph}}(\omega)$ and their Kelvin–Voigt fits are used to recover (μ, η) and to analyse dispersive effects.

A variety of estimators have been proposed to extract these quantities from spatio-temporal motion fields [Marmin et al., 2021, Rouze et al., 2010].

Time-of-flight methods for c_g measure the arrival time of a reference feature at different lateral positions and obtain the group velocity from the slope of the resulting space–time trajectory. Spectral approaches operate in the wavenumber–frequency domain [Kijanka and Urban, 2019] and estimate either the dominant slope of the propagation ridge or the frequency-dependent phase velocity from the dispersion relation. Phase velocity $c_{\mathrm{ph}}(\omega)$ can be obtained by tracking, at each frequency, the wavenumber of maximum spectral energy in the two-dimensional Fourier transform $V(k_x, \omega)$ of $v_z(x, t)$ [Alleyne and Cawley, 1991].

In this chapter, both metrics are extracted from the axial particle-velocity maps $v_z(x, t)$ at the focal depth, using the same processing for the three datasets (simulation, synthetic, and experiment). Group velocity is estimated from the trajectory of the wave-packet maximum in the (x, t) plane, while phase velocity is derived from the ridge of the k_x–ω spectrum and related to the Kelvin–Voigt parameters through Eq. (2.39). The corresponding procedures are detailed below in the context of the homogeneous CIRS liver phantoms.

2.4.2 Group velocity in homogeneous phantoms

Shear-wave group velocity was estimated from the spatio-temporal maps of axial particle velocity $v_z(x, t)$ extracted at the focal depth. For each phantom and each dataset, $v_z(x, t)$ was obtained by averaging the shear-wave propagation field $v_z(x, z, t)$ over a 1 mm-thick axial slab centred at $z = 20.04$ mm, corresponding to the nominal focus of the pushing beam. This averaging reduces local fluctuations and yields a one-dimensional representation of the dominant shear-wave front.

Within each time window, a reference motion trace was taken at a given lateral position and cross-correlated with traces at positions separated by a fixed lateral offset dx. The lag corresponding to the maximum of the normalised cross-correlation provided the local time-of-flight $dt(x)$ between those two positions. A linear regression was then fitted to the arrival times as a function of lateral position, so that the slope of this fit yielded a single group-velocity estimate $c_g = dx/dt$. This estimate should be interpreted as an effective group velocity associated with the dominant frequency content of the measured wave packet within the selected time window, rather than as $c_g(\omega) = d\omega/dk$ evaluated at a specific frequency. The coefficient of determination R^2 quantified how well a constant effective velocity described the observed propagation over

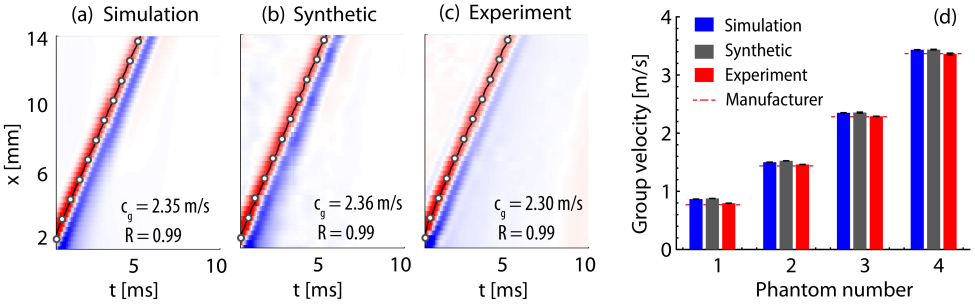


Figure 2.11: Shear wave group velocity estimation from spatio-temporal particle velocity maps. Particle velocity at the focal depth for simulation (a), synthetic (b), and experimental (c) data in Phantom 3, with peak trajectories and linear fits used to compute c_g. (d) Group velocity values for all four phantoms, comparing the three methods against the nominal manufacturer velocity.

c_g [m/s]	Phantom 1	Phantom 2	Phantom 3	Phantom 4
Manufacturer	0.810 ± 0.080	1.480 ± 0.060	2.310 ± 0.110	3.360 ± 0.070
Simulation	0.867 ± 0.001	1.501 ± 0.002	2.350 ± 0.003	3.434 ± 0.003
Synth	0.878 ± 0.001	1.524 ± 0.002	2.355 ± 0.008	3.436 ± 0.005
Experiment	0.795 ± 0.002	1.462 ± 0.002	2.290 ± 0.002	3.361 ± 0.016

Table 2.4: Effective group velocity $\bar{c_g}$ for the four phantoms, as provided by the manufacturer and estimated from simulations, synthetic reconstructions, and experimental measurements. Values are reported as mean ± standard deviation.

the analyzed range.

A representative example of this procedure is shown in Fig. 2.11 for Phantom 3. The three panels display the normalized particle velocity spatio-temporal maps for the simulation, synthetic, and experimental datasets, together with the discrete maxima and the corresponding regression line. In all three cases, the estimated group velocity is close to 2.3 m/s and the fits yield $R = 0.99$, indicating a highly coherent lateral propagation within the focal slab for the three datasets.

The effective group-velocity, $\bar{c_g}$, estimates for all phantoms are summarized in the bar plot of Fig. 2.11 and in Table 2.4. For each phantom, the three bars correspond to the simulation, synthetic, and experimental datasets, and the dashed horizontal line indicates the nominal value provided by the manufacturer. Error bars represent the standard deviation of the estimates

across all lateral positions and realizations considered in the regression. The measured velocities follow the expected increase with phantom stiffness, from approximately 0.8 m/s in Phantom 1 to around 3.4 m/s in Phantom 4, with small standard deviations for all three datasets.

2.4.3 Phase velocity in homogeneous phantoms

Phase-velocity dispersion was analyzed using the same spatio-temporal particle velocity maps employed for the group-velocity estimation. For each phantom and dataset, the axial particle velocity $v_z(x,t)$ was extracted by averaging the volumetric field over a 1 mm-thick slab centered at $z = 20.04$ mm. A two-dimensional Fourier transform was then applied to obtain the spectrum $V(k_x,\omega)$ in the wavenumber–frequency domain. The phase velocity $c_{\mathrm{ph}}(f)$ was computed by tracking, for each frequency, the wavenumber corresponding to the maximum spectral magnitude along the main propagation ridge,

$$c_{\mathrm{ph}}(\omega) = \frac{\omega}{\mathrm{Re}\{k_x(\omega)\}}, \tag{2.40}$$

where $\mathrm{Re}\{k_x(\omega)\}$ denotes the ridge wavenumber in the $k-\omega$ space. This procedure follows standard 2D Fourier-domain dispersion analysis used in transient elastography and shear-wave elastography studies [Catheline et al., 2004].

An example of this procedure is shown in Fig. 2.12 for Phantom 2. Figure. 2.12(a-c) displays the normalized $v_z(x,t)$ maps, Fig. 2.12(d-f) shows the corresponding $|V(k_x,\omega)|$ spectra, and Fig. 2.12(g-i) depicts the phase-velocity curves obtained from the ridge tracking. In each case, the black line represents the raw phase-velocity estimate, the circles indicate the frequency samples retained for the viscoelastic fitting, the red line shows the Kelvin–Voigt fit based on Eq. (2.39), and the blue dashed line marks the nominal shear-wave speed provided by the manufacturer. The curves exhibit a rapid increase of $c_{\mathrm{ph}}(\omega)$ at low frequencies, followed by a quasi-plateau in the frequency band where the spectrum concentrates most of its energy.

The frequency range used for Kelvin–Voigt fitting was chosen independently for each phantom based on the energy distribution in the k–ω spectrum of the motion wave. The lower bound was set at the frequency where the spectral amplitude first dropped by 3 dB below the maximum, while the upper bound was selected between the -6 dB and -12 dB levels depending on the

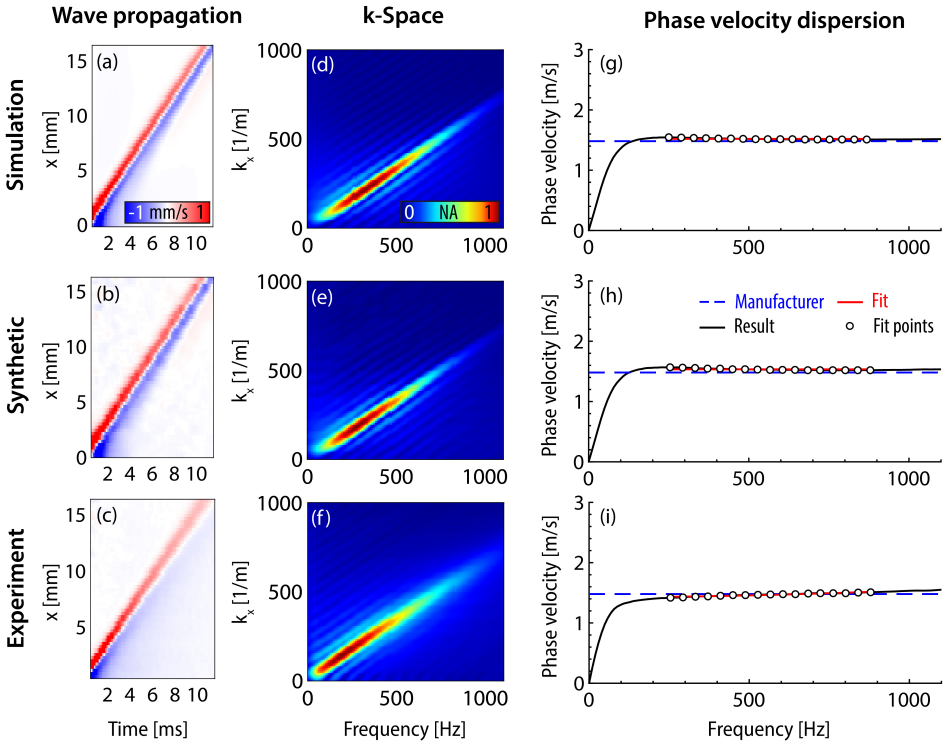


Figure 2.12: Example of phase-velocity estimation in Phantom 2. From left to right: shear-wave propagation map $v_z(x,t)$ at the focal depth, corresponding k–ω spectrum $|V(k_x,\omega)|$, and phase-velocity dispersion curves $c_{\mathrm{ph}}(\omega)$ for simulation, synthetic, and experimental data, together with the Kelvin–Voigt fit and the nominal manufacturer speed.

attenuation of the medium. This procedure resulted in the following frequency bands: 50–650 Hz for Phantom 1, 125–1100 Hz for Phantom 2, 150–1250 Hz for Phantom 3, and 250–1500 Hz for Phantom 4. Within each phantom, the same frequency band was used consistently for the simulation, synthetic, and experimental datasets.

Figure 2.13 summarises the phase-velocity dispersion curves for all four phantoms. For each case, the mean phase velocity, $\bar{c_{\mathrm{ph}}}$, within the selected frequency band was computed together with its standard deviation. The resulting values are reported in Table 2.5. Across the four phantoms, the mean phase velocities obtained from simulation, synthetic, and experimental data closely follow the increase in nominal shear-wave speed from 0.81 to 3.36 m/s.

Deviations from the manufacturer values remain below 8% for all phantoms and below 4% for Phantoms 2–4, with comparable agreement between the three datasets.

Finally, the Kelvin–Voigt model was used to convert the dispersion information into effective shear elastic and viscoelastic parameters for each phantom. In the manufacturer specifications and in the simulation and synthetic datasets, the media were modelled as purely elastic with $\eta = 0$, and the KV fits consistently recovered negligible viscosities while yielding shear moduli in close agreement with the prescribed values (0.66, 2.19, 5.34, and 11.29 kPa and their simulated counterparts). In contrast, the experimental dispersion curves resulted non-zero viscosities, with fitted shear moduli ranging from 0.59 to 11.05 kPa and viscosities between 0.12 and 0.27 Pa s for Phantoms 1–4. This slight viscoelastic response is also apparent in Fig. 2.12, where the phase-velocity curves for the simulated and synthetic data remain essentially flat within the analysed frequency bands, while the experimental curve exhibits a small but noticeable positive slope.

Together, the group-velocity and phase-velocity analyses provide a consistent set of stiffness and dispersion metrics across simulation, synthetic, and experimental datasets, and establish the notation and processing steps that will be used in the subsequent chapters of the thesis.

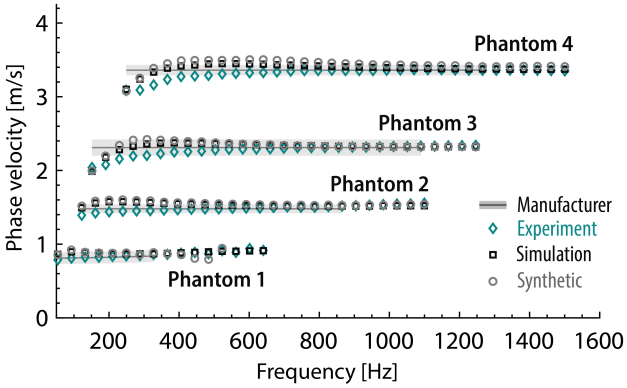


Figure 2.13: Phase-velocity dispersion curves $c_{\mathrm{ph}}(\omega)$ for the four phantoms, comparing simulation, synthetic reconstructions, and experimental measurements within the frequency bands of maximum spectral energy. The shaded regions indicate the manufacturer shear-wave speeds and their tolerance ranges.

$\bar{c}_{\mathrm{ph}}$ [m/s]	Phantom 1	Phantom 2	Phantom 3	Phantom 4
Manufacturer	0.81 ± 0.08	1.48 ± 0.06	2.31 ± 0.11	3.36 ± 0.07
Simulation	0.87 ± 0.03	1.51 ± 0.02	2.30 ± 0.08	3.36 ± 0.07
Synthetic	0.84 ± 0.03	1.53 ± 0.02	2.31 ± 0.09	3.40 ± 0.09
Experiment	0.84 ± 0.06	1.47 ± 0.05	2.25 ± 0.09	3.30 ± 0.09

Table 2.5: Mean phase velocity $\bar{c}_{\mathrm{ph}}$ and standard deviation within the selected frequency band for each phantom, as provided by the manufacturer and estimated from simulations, synthetic reconstructions, and experimental measurements.

2.5 Discussion

Beyond reproducing intermediate fields, the key outcome is that simulated shear-wave motion can be compared against synthetic and experimental IQ datasets using an identical tracking and velocity-analysis chain. Under this design, discrepancies between simulation and experiment are primarily attributable to modelling assumptions and uncertain material properties, rather than to mismatched reconstruction methods.

The acoustic stage is central to the practicality of the framework. By combining Rayleigh integration at the transducer surface with angular spectrum propagation in the bulk medium, the RAAS solver captures diffraction and interference effects that are essential for tightly focused beams, while remaining far less demanding than full three-dimensional time-domain solvers. The comparisons against MUST and k-Wave confirm that RAAS reproduces focal geometry and on-axis profiles with high fidelity. The computational study further supports the role of RAAS as a scalable solution for large parametric sweeps, where repeated field evaluations are required.

At the same time, the acoustic formulation adopts simplifying assumptions that define its range of validity. The present implementation assumes a homogeneous, weakly attenuating background with linear propagation at a single frequency. Consequently, nonlinear generation of higher harmonics and acoustic heterogeneities may become significant for very high-intensity pushes or strongly heterogeneous propagation paths. Within the voltage levels and burst durations considered in this chapter, the comparison with hydrophone measurements indicates that the linear approximation captures the dominant features of the driving field at the focus, which is the primary determinant of the ARF distribution used to excite shear waves.

The ARF computation constitutes the second pillar of the approach. Using the full vector second-order expression in Eq. (2.21), rather than an axial-only approximation, preserves the three-dimensional structure of the loading and enables the modelling of off-axis force components that can seed out-of-plane shear-wave energy. This is particularly relevant in the context of volumetric SWE, where energy leakage in the elevational direction has the potential to introduce bias in two-dimensional reconstructions. This phenomenon can also occur in alternative pushing strategies that will be developed during the course of this thesis. The remaining simplification is the use of constant sound speed and attenuation, which implies a linear scaling between pressure and ARF and ignores frequency-dependent absorption. Extending the ARF model to incorporate nonlinear propagation and frequency-dependent attenuation would increase physical realism. However, note that most acoustic pushing beams are quasi-monochromatic and narrowband around a central frequency, so frequency-dependent attenuation can be reasonably approximated by its value at ω_0, making the linear pressure-to-ARF scaling an adequate first-order assumption in the regime considered.

For shear-wave propagation, the elastic solver of k-Wave provides a full-wave solution of the viscoelastic equations in a three-dimensional domain, allowing shear reflections and out-of-plane spreading to emerge naturally. The use of homogeneous Kelvin–Voigt media matching the nominal CIRS 039 phantom specifications enables a controlled interpretation of the results, since the prescribed shear speeds and densities map directly to shear moduli and the zero-viscosity setting yields non-dispersive simulated waves. The agreement between simulated, synthetic, and experimental spatio-temporal patterns in Figs. 2.9 and 2.10 indicates that this idealised configuration is sufficient to reproduce the dominant propagation features in these phantoms. The limitations are clear in scope: real tissues, and often commercial phantoms, exhibit heterogeneity, anisotropy, and measurable viscoelasticity. These factors are intentionally excluded here to isolate the coherence of the coupling between acoustic, mechanical, and imaging stages, but they must be incorporated to address more complex media.

A distinctive strength of the proposed route is the explicit connection between mechanical simulations and ultrasound image formation. Synthetic IQ mirror the experimental acquisition in terms of probe geometry, sampling, and field of view. This design ensures that beamforming and demodulation effects are present in both synthetic and experimental data, reducing the

risk of attributing imaging-chain artefacts to shortcomings of the mechanical model.

The tracking stage builds on the 1D autocorrelation estimator applied directly to the complex IQ stream $s(x, z, t)$, converting inter-frame phase shifts into axial particle velocities. Applying the same estimator to synthetic and experimental data strengthens the validity of cross-dataset comparisons. The method nonetheless inherits the known constraints of phase-based tracking, including the need for small-displacement and narrowband conditions and the potential for decorrelation due to out-of-plane motion, multiple scattering, or low SNR. While these effects remain modest in the homogeneous phantom experiments presented here, they are expected to become more prominent in *in vivo* settings and should be considered when extending the framework.

The velocity analysis clarifies the complementary roles of group and phase velocity metrics. Velocity is estimated from the propagation of the wave-packet envelope and yields robust scalar stiffness trends that remain consistent across simulation, synthetic frames, and experiments. Phase velocity provides frequency-resolved information that is sensitive to weak dispersion. The flat phase-velocity curves in simulation and synthetic datasets are consistent with the elastic definition of the media, while the slight positive slope observed experimentally leads to non-zero effective viscosities and shear moduli below nominal values in Table 2.5. These results indicate that the physical phantoms exhibit mild viscoelasticity despite a nominally elastic specification, and they show that the proposed method is sensitive to such weak dispersive behaviour. The dispersion analysis also highlights limitations linked to empirical frequency-band selection and to the Kelvin–Voigt rheology itself, which represents only a single relaxation mechanism and should be interpreted as an effective description over the selected band.

2.6 Conclusions

This chapter describes a simulation method for shear wave elastography, linking RAAS-based acoustic beam modelling, second-order vector ARF estimation, full-wave three-dimensional shear-wave propagation, and ultrasound-like image formation and motion tracking. The acoustic stage reproduces the focal geometry and axial profiles obtained with MUST and k-Wave, as demonstrated in Figs. 2.5 and 2.6 and quantified in Tables 2.2 and 2.3, while offering a

substantially reduced computational burden that supports large parametric studies.

The framework facilitates comparable analysis of datasets by propagating shear waves driven by volumetric ARF fields and analyzing shear-wave motion through the same tracking and velocity-estimation chain for synthetic and experimental IQ data. In homogeneous commercial phantoms, group and phase velocity estimates remain consistent across simulation, synthetic data, and experiments, and the dispersion analysis reveals a mild viscoelastic signature in the experimental data that is absent in the purely elastic simulated media. The main limitations arise from linear single-frequency acoustic modelling, homogeneous Kelvin–Voigt mechanics, and the known constraints of phase-based tracking and empirical frequency-band selection. However, these assumptions are fulfilled for most elastography applications. In this way, the chapter provides a validated and computationally efficient approach for subsequent extensions to more complex media and acquisition conditions.

Publications

This section contains a compendium of the seminal scientific articles in which the results presented in this chapter have been published.

Conferences articles

- Jimenez, N., Carrión García, A., Andrés-Bautista, D., González-Mateo, E., Hernández-López, C., and Camarena, F., "Técnicas de holografía acústica para la validación rápida de hologramas acústicos," in *TECNIACÚSTICA 2022. 53º Congreso Español de Acústica. XII Congreso Ibérico de Acústica. European Symposium on NVH in Electric Vehicles*, Elche, España, 2022, pp. 924–929.

References

Alleyne, D. and Cawley, P. (1991). A two-dimensional fourier transform method for the measurement of propagating multimode signals. *The Journal of the Acoustical Society of America*, 89(3):1159–1168.

Bernal, M., Benech, N., Daigle, R., and Brum, J. (2021). Towards 3D passive shear elasticity imaging using row-columns arrays. In *2021 IEEE International Ultrasonics Symposium (IUS)*, pages 1–4. ISSN: 1948-5727.

Bouwkamp, C. (1954). Diffraction theory. *Reports on Progress in Physics*, 17(1):35–100.

Catheline, S., Gennisson, J.-L., Delon, G., Fink, M., Sinkus, R., Abouelkaram, S., and Culioli, J. (2004). Measurement of viscoelastic properties of homogeneous soft solid using transient elastography: An inverse problem approach. *The Journal of the Acoustical Society of America*, 116(6):3734–3741.

Chivers, R. C. and Hill, C. R. (1975). Ultrasonic attenuation in human tissue. *Ultrasound in Medicine & Biology*, 2(1):25–29.

Cigier, A., Varray, F., and Garcia, D. (2022). SIMUS: An open-source simulator for medical ultrasound imaging. Part II: Comparison with four simulators. *Computer Methods and Programs in Biomedicine*, 220:106774.

Cox, B. T., Kara, S., Arridge, S. R., and Beard, P. C. (2007). k-space propagation models for acoustically heterogeneous media: Application to biomedical photoacoustics. *The Journal of the Acoustical Society of America*, 121(6):3453–3464.

Dong, Z., Lok, U.-W., Lowerison, M. R., Huang, C., Chen, S., and Song, P. (2024). Three-Dimensional Shear Wave Elastography Using Acoustic Radiation Force and a 2-D Row-Column Addressing (RCA) Array. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 71(4):448–458.

Firouzi, K., Cox, B., Treeby, B., and Saffari, N. (2011). A first-order k-space model for elastic wave propagation in heterogeneous media. *The Journal of the Acoustical Society of America*, 129(4_Supplement):2611.

Garcia, D. (2021). Make the most of MUST, an open-source Matlab UltraSound Toolbox. In *2021 IEEE International Ultrasonics Symposium (IUS)*, pages 1–4. ISSN: 1948-5727.

Garcia, D. (2022). SIMUS: An open-source simulator for medical ultrasound imaging. Part I: Theory & examples. *Computer Methods and Programs in Biomedicine*, 218:106726.

Gennisson, J.-l., Provost, J., Deffieux, T., Papadacci, C., Imbault, M., Pernot, M., and Tanter, M. (2015). 4-D ultrafast shear-wave imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 62(6):1059–1065.

Huang, C., Song, P., Mellema, D. C., Gong, P., Lok, U.-W., Tang, S., Ling, W., Meixner, D. D., Urban, M. W., Manduca, A., Greenleaf, J. F., and Chen, S. (2020). Three-Dimensional Shear Wave Elastography on Conventional Ultrasound Scanners with External Vibration. *Physics in medicine and biology*, 65(21):215009.

Jensen, J. and Svendsen, N. (1992). Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 39(2):262–267.

Jiménez, N., Picó, R., Camarena, F., Redondo, J., and Roig, B. (2012). Ultrasonic evaluation of the hydration degree of the orange peel. *Postharvest Biology and Technology*, 67:130–137.

Karaman, M., Li, P.-C., and O'Donnell, M. (1995). Synthetic aperture imaging for small scale systems. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 42(3):429–442.

Kijanka, P. and Urban, M. W. (2019). Local Phase Velocity Based Imaging: A New Technique Used for Ultrasound Shear Wave Elastography. *IEEE Transactions on Medical Imaging*, 38(4):894–908. Conference Name: IEEE Transactions on Medical Imaging.

Kirkup, S. M. (1994). Computational solution of the acoustic field surrounding a baffled panel by the Rayleigh integral method. *Applied Mathematical Modelling*, 18(7):403–407.

Lee, W. and Roh, Y. (2017). Ultrasonic transducers for medical diagnostic imaging. *Biomedical Engineering Letters*, 7(2):91–97.

Lipman, S. L., Rouze, N. C., Palmeri, M. L., and Nightingale, K. R. (2016). Evaluating the Improvement in Shear Wave Speed Image Quality Using Multidimensional Directional Filters in the Presence of Reflection Artifacts. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 63(8):1049–1063.

Liu, D.-L. and Waag, R. C. (1997). Propagation and backpropagation for ultrasonic wavefront design. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 44(1):1–13.

Loupas, T., Powers, J., and Gill, R. (1995). An axial velocity estimator for ultrasound blood flow imaging, based on a full evaluation of the Doppler equation by means of a two-dimensional autocorrelation approach. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 42(4):672–688. Conference Name: IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control.

Marmin, A., Laloy-Borgna, G., Facca, S., Gioux, S., Catheline, S., and Nahas, A. (2021). Time-of-flight and noise-correlation-inspired algorithms for full-field shear-wave elastography using digital holography. *Journal of Biomedical Optics*, 26(8):086006.

Merlini, G., Imperiale, S., and Allain, J.-M. (2025). Finite element modelling of transient shear wave elastography for the cornea. *Journal of the Mechanics and Physics of Solids*, page 106363.

Montaldo, G., Tanter, M., Bercoff, J., Benech, N., and Fink, M. (2009). Coherent plane-wave compounding for very high frame rate ultrasonography and transient elastography. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 56(3):489–506.

Osika, M. and Kijanka, P. (2024). Ultrasound shear wave propagation modeling in general tissue–like viscoelastic materials. *Ultrasound in Medicine & Biology*, 50(4):627–638.

Palmeri, M. L., Qiang, B., Chen, S., and Urban, M. W. (2016). Guidelines for finite-element modeling of acoustic radiation force-induced shear wave propagation in tissue-mimicking media. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 64(1):78–92.

Perrot, V., Polichetti, M., Varray, F., and Garcia, D. (2021). So you think you can DAS? A viewpoint on delay-and-sum beamforming. *Ultrasonics*, 111:106309.

Prieur, F. and Catheline, S. (2016). Simulation of shear wave elastography imaging using the toolbox "k-Wave". *Proceedings of Meetings on Acoustics*, 29(1):020002. Publisher: Acoustical Society of America.

Prieur, F. and Sapozhnikov, O. A. (2017). Modeling of the acoustic radiation force in elastography. *The Journal of the Acoustical Society of America*, 142(2):947–961.

Qiang, B., Brigham, J. C., Aristizabal, S., Greenleaf, J. F., Zhang, X., and Urban, M. W. (2015). Modeling transversely isotropic, viscoelastic, incompressible tissue-like materials with application in ultrasound shear wave elastography. *Physics in Medicine & Biology*, 60(3):1289–1306.

Rouze, N. C., Wang, M. H., Palmeri, M. L., and Nightingale, K. R. (2010). Robust estimation of time-of-flight shear wave speed using a radon sum transformation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 57(12):2662–2670. Conference Name: IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control.

Sapozhnikov, O. A., Tsysar, S. A., Khokhlova, V. A., and Kreider, W. (2015). Acoustic holography as a metrological tool for characterizing medical ultrasound sources and fields. *The Journal of the Acoustical Society of America*, 138(3):1515–1532.

Thomenius, K. (1996). Evolution of ultrasound beamformers. In *1996 IEEE Ultrasonics Symposium. Proceedings*, volume 2, pages 1615–1622 vol.2. ISSN: 1051-0117.

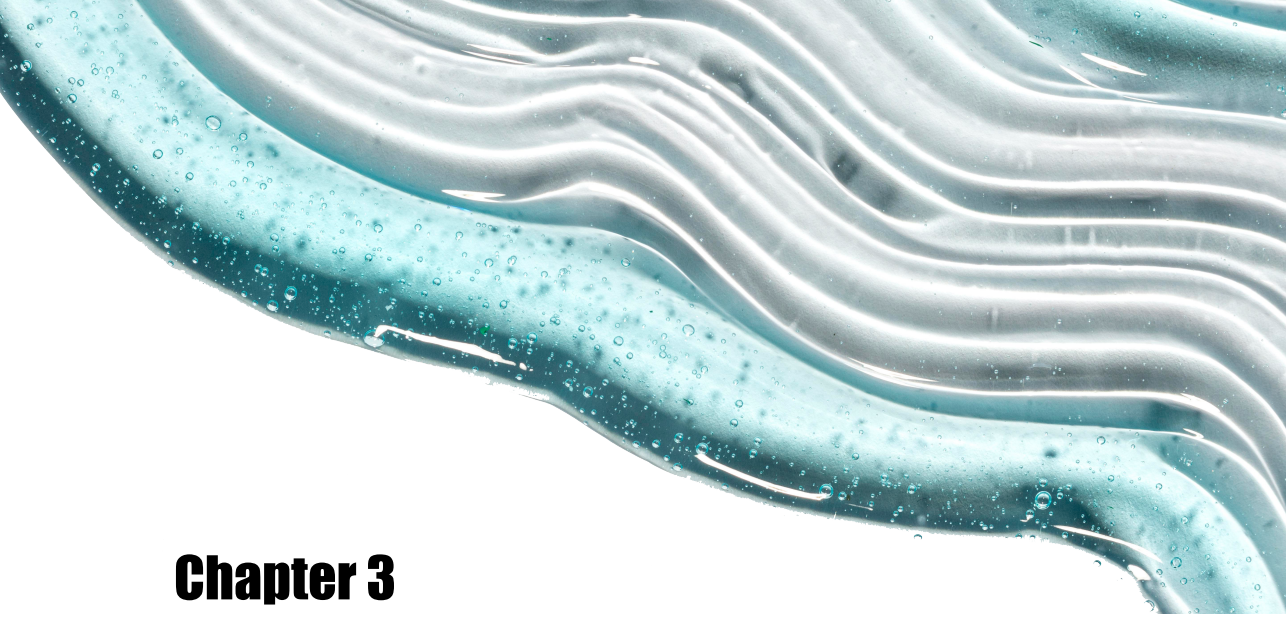
Treeby, B. E. and Cox, B. T. (2010). k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave fields. *Journal of Biomedical Optics*, 15(2):021314. Publisher: SPIE.

Treeby, B. E., Jaros, J., Rohrbach, D., and Cox, B. T. (2014). Modelling elastic wave propagation using the k-Wave MATLAB Toolbox. In *2014 IEEE International Ultrasonics Symposium*, pages 146–149.

Wang, M., Byram, B., Palmeri, M., Rouze, N., and Nightingale, K. (2013). Imaging Transverse Isotropic Properties of Muscle by Monitoring Acoustic Radiation Force Induced Shear Waves Using a 2-D Matrix Ultrasound Array. *IEEE Transactions on Medical Imaging*, 32(9):1671–1684.

Zeng, X. and McGough, R. J. (2006). Evaluation of Angular Spectrum Approach for Simulations of Spherically Focused Ultrasound Phased Arrays. In *2006 IEEE International Conference on Electro/Information Technology*, pages 66–71. ISSN: 2154-0373.

Zeng, X. and McGough, R. J. (2008). Evaluation of the angular spectrum approach for simulations of near-field pressures. *The Journal of the Acoustical Society of America*, 123(1):68–76.



Chapter 3

Real-time shear wave phase velocity mapping by phase gradients

Abstract

This chapter presents a phase-gradient elastography (PGE) method for mapping shear-wave phase velocity and viscoelastic properties with high spatial resolution. The method is conceived for scenarios where spectral information is required, but conventional time-of-flight or spectral techniques become limited by computational cost and spatial averaging. From acoustic radiation force excitations and ultrafast ultrasound motion tracking acquisitions, axial particle-velocity fields are reconstructed, filtered in the frequency–wavenumber domain to extract the shear-wave band of interest, and converted back to the spatial domain at selected frequencies. Local phase gradients are then used to estimate the phase-velocity maps at each frequency. By assembling these maps into local dispersion curves and fitting them to a Kelvin–Voigt model, spatially resolved estimates of shear modulus and viscosity are obtained. The chapter details the theoretical formulation, numerical implementation, and validation of PGE in simulated mediums, tissue-mimicking phantoms, and *ex vivo* liver, highlighting its potential for real-time spectral shear-wave elastography.

Outline

3.1 Introduction 69
3.2 Phase gradient elastography method 70
3.2.1 Overview and input data 70
3.2.2 Motion spectrum 72
3.2.3 Temporal, angular and bandpass filtering 73
3.2.4 Phase-gradient operator and local wavenumber 77
3.2.5 Phase-velocity maps and weighted averaging 80
3.3 Numerical and experimental validation 82
3.3.1 Numerical validation 82
3.3.2 Experimental phantom validation 88
3.3.3 *Ex vivo* bovine liver validation 90
3.4 Computational complexity 92
3.5 Discussion 93
3.6 Conclusions 96

3.1 Introduction

A fundamental challenge in the field of ultrasound shear-wave elastography is how to reconstruct local mechanical information from transient, spatially complex wavefields. In particular, when viscoelastic characterization is sought, the relevant quantity is not a single propagation speed but the phase velocity as a function of frequency.

From this perspective, shear-wave velocity estimators can be broadly divided into two groups: time-domain methods based on time-of-flight (TOF) measurements [Palmeri et al., 2008, Rouze et al., 2010, Zhao et al., 2011] and frequency-domain methods based on phase or spectral analysis [Manduca et al., 2001, Kijanka and Urban, 2019]. TOF approaches are simple and robust, which explains their adoption in many clinical systems, but they typically collapse the wavefield into a single effective speed and therefore overlook dispersion. Spectral methods, in contrast, estimate phase velocity as a function of frequency, enabling the recovery of dispersion curves and, under an appropriate rheological model, the estimation of elastic and viscous parameters. They have also been extended to characterize features such as anisotropy or nonlinearity [Li and Cao, 2017, Rizzuto et al., 2021].

Phase-gradient methods are not new in elastography. They have been used in harmonic MR elastography as local estimators of wavenumber, often combined with directional filtering to improve robustness in the presence of interfering waves [Manduca et al., 2003]. In ultrasound, phase-based dispersion measurements were also explored early on by tracking phase delays over propagation distance [Chen et al., 2004], and more recent work has revisited phase gradients as robust local estimators in complex shear-wave fields [Ormachea and Parker, 2021]. Accordingly, the contribution of this chapter is not the introduction of phase-gradient estimators themselves, but their integration into a transient ultrasound SWE framework for dense, frequency-resolved phase-velocity imaging.

The main limitation of spectral elastography is computational cost. Although real-time shear-wave elasticity imaging has been demonstrated [Bercoff et al., 2008], clinical implementations remain largely time-domain because full spatio-temporal spectral processing typically relies on large Fourier transforms and block-based strategies that scale poorly with image size [Kijanka and Urban, 2020]. This chapter addresses that limitation by adapting and evaluating a phase-gradient elastography (PGE) method that combines directional

filtering in the spatio-temporal frequency domain with a local estimator of the in-plane wavevector. The filtered harmonic wavefields are used to reconstruct local phase-velocity maps across frequency, enabling elasticity and viscoelastic parameter imaging from the same dataset while substantially reducing computational burden. The following sections present the formulation of the method and its numerical, phantom, *ex vivo*, and computational validation.

3.2 Phase gradient elastography method

3.2.1 Overview and input data

The starting point of the PGE method is a complex-valued representation of the shear-wave motion. This motion can be described either by the axial displacement $u(x, z, t)$ or by the particle velocity $v(x, z, t)$, where x and z denote the lateral and axial spatial coordinates, and t is the slow-time index of the ultrafast acquisition. In the frequency domain $V(x, z, \omega) = \mathrm{i}\,\omega\, U(x, z, \omega)$, so V is phase-shifted by $\pi/2$ relative to U. This constant offset is spatially uniform and cancels in the spatial phase gradients, yielding identical wavevector estimates in principle. In practice, working with v is preferable, since autocorrelation-based estimators are natively formulated as velocity estimators [Loupas et al., 1995, Pinton et al., 2006] and velocity suppresses slow drifts and quasi-static components that can contaminate the displacement field, yielding a more localized ridge in the k–ω domain. Throughout this chapter, PGE is therefore applied to the complex particle-velocity field reconstructed from beamformed RF data by converting to the IQ domain and applying a phase-shift estimator in slow time, as described in Sec. 2.2.4.

The overall objective of PGE is to recover, for each spatial location (x, z) and angular frequency ω within a suitable band, the local phase velocity

$$c_{\mathrm{ph}}(x, z, \omega) = \frac{\omega}{\|\Re\{\mathbf{k}(x, z, \omega)\}\|}, \tag{3.1}$$

where $\mathbf{k}(x, z, \omega)$ is the local wavevector of the shear wave. Once these local phase-velocity maps are available, dispersion curves $c_{\mathrm{ph}}(\omega)$ can be extracted in regions of interest (ROIs) and fitted to an appropriate viscoelastic model, enabling quantitative estimation of storage and loss moduli.

From a conceptual standpoint, PGE seeks to estimate $\Re\{\mathbf{k}(x, z, \omega)\}$ directly from the spatial variation of the phase of a suitably filtered complex motion

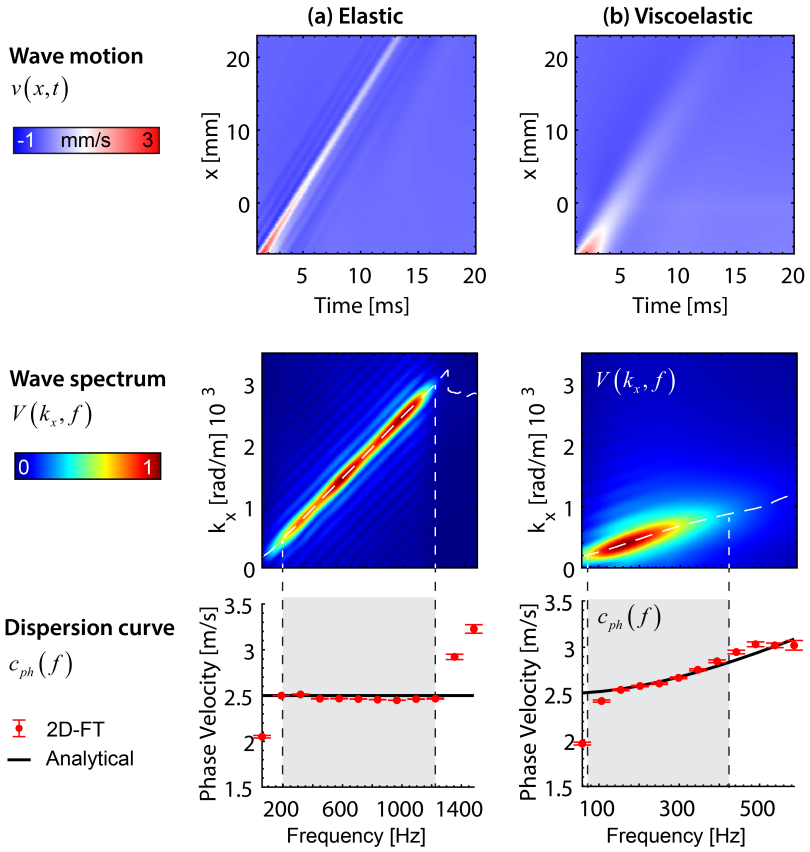


Figure 3.1: Shear-wave motion and spectral content in (a) an elastic and (b) a viscoelastic medium. Top row: wave motion in the space–time domain $v(x,t)$. Middle row: normalized spatio-temporal wave spectrum in the (k_x, f) domain. Bottom row: phase-velocity dispersion curves estimated with the 2D-FT method (red markers) and the analytical Kelvin–Voigt model (black line). The gray shaded region indicates the frequency band where the spectral amplitude is high and 2D-FT estimates are reliable.

field. By combining spectral filtering and a local phase-gradient operator, PGE delivers wavevector estimates at each pixel while preserving the native sampling of the motion data.

3.2.2 Motion spectrum

The first step in the PGE pipeline is to inspect the spectral content of the measured motion and to identify the frequency range where the propagating shear wave is dominant. To this end, a temporal Fourier transform is applied to $v_z(x, z, t)$ at each spatial position,

$$V(x, z, \omega) = \int_{-\infty}^{\infty} v_z(x, z, t)\, e^{-\mathrm{i}\omega t}\, \mathrm{d}t, \tag{3.2}$$

using an N_t-point fast Fourier transform (FFT). The modulus of $V(x, z, \omega)$ describes the local spectral content of the motion, whereas its phase encodes the propagation pattern of the shear wave.

The bandwidth of interest for shear-wave analysis is determined from the spatio-temporal spectrum of the motion. In practice, the (k_x, ω) spectrum is computed from $V(x, z, \omega)$ and the ridge associated with the dominant shear-wave component is identified (Fig. 3.1, middle row). The lower edge of the useful band is typically defined at the angular frequency $\omega_{-3\mathrm{dB}}$ where the spectral magnitude along this ridge has decreased by 3 dB with respect to its maximum, while the upper edge lies between $\omega_{-6\mathrm{dB}}$ and $\omega_{-12\mathrm{dB}}$, depending on the attenuation of the medium. Frequencies well below the spectral peak, associated with very long wavelengths and quasi-static motion, are not included in the analysis because they carry little spatial information and their spectral energy is low [Palmeri et al., 2014]. Within the resulting bandwidth, a discrete set of angular frequencies ω_0 is selected as working points for PGE, at which monochromatic phase-velocity maps will be reconstructed.

Figure 3.1(a) summarizes this analysis for a homogeneous elastic medium and Fig. 3.1(b) for a viscoelastic medium. The top row shows the space–time evolution of $v(x, t)$, with a nearly non-dispersive wavefront in the elastic case and a dispersive, attenuating wave in the viscoelastic case. The middle row displays the corresponding (k_x, f) spectra: in the elastic medium the energy concentrates along a straight ridge, consistent with a frequency-independent phase velocity, whereas in the viscoelastic medium the ridge bends and progressively vanishes at high frequencies, reflecting dispersion and viscous damping. The bottom row presents the phase-velocity curves estimated by the 2D-FT method (red markers) together with the analytical Kelvin–Voigt prediction (black line). The gray shaded region indicates the effective shear-wave bandwidth, defined as described above, where the spectral amplitude is high and 2D-FT estimates closely follow the analytical model.

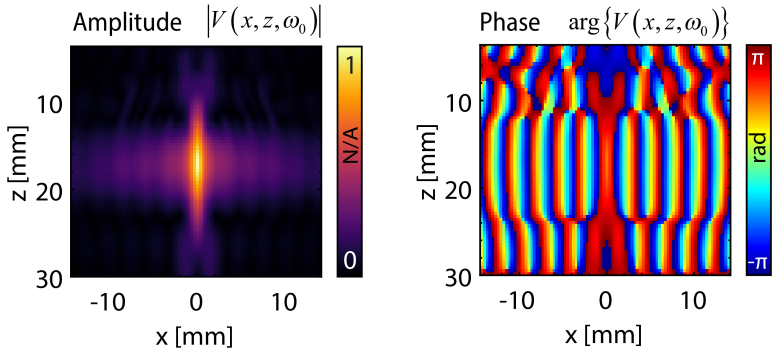


Figure 3.2: Spatial distribution of the complex motion spectrum at a single angular frequency ω_0 in a homogeneous elastic medium. Left: amplitude map $|V(x, z, \omega_0)|$. Right: phase map $\arg\{V(x, z, \omega_0)\}$. The amplitude is concentrated around the ARF focal region, whereas the phase exhibits nearly planar fringes along the lateral direction, indicative of a laterally propagating shear wave.

Figure 3.2 further illustrates the spatial structure of the complex spectrum at one of these angular frequencies $\omega_0 = 2pi\, f_0$ ($f_0 = 750$ Hz), selected within the shaded band of Fig. 3.1(a) for the homogeneous elastic case. The left panel shows the amplitude map $|V(x, z, \omega_0)|$, which is strongly concentrated around the ARF focal region and decays laterally and axially. The right panel displays the corresponding phase map $\arg\{V(x, z, \omega_0)\}$, where almost straight fringes are observed along x at a given depth, consistent with a locally planar shear wave propagating away from the push. The wavefront is not perfectly planar in the vicinity of the focus and near the probe surface, where effects such as near-field and boundary conditions are more pronounced. These phase patterns, obtained at carefully selected frequencies within the shear-wave band, are the basis for the local wavevector estimation performed in the subsequent phase-gradient step.

3.2.3 Temporal, angular and bandpass filtering

It is well established that reflections, interference between multiple wave components, and strongly curved shear-wave fronts can hinder a straightforward reconstruction of phase-velocity maps [Lipman et al., 2016]. Moreover, phase-gradient estimators are themselves sensitive to noise and motion artefacts, so that residual clutter or low-SNR regions can translate into local errors in the estimated phase-velocity field [Deffieux et al., 2009]. For this reason, the next

step in the processing is, at a given angular frequency ω_0, to extract a motion field that behaves as closely as possible to a single, locally planar travelling wave. To this end, the complex motion field $V(x, z, \omega_0)$ is first transformed into the spatial-frequency domain by means of a two-dimensional Fourier transform,

$$V(k_x, k_z, \omega_0) = \iint_{-\infty}^{\infty} V(x, z, \omega_0)\, e^{-i\,(k_x x + k_z z)}\, \mathrm{d}x\, \mathrm{d}z, \tag{3.3}$$

computed using an $N_x \times N_z$-point 2D-FFT, where k_x and k_z denote the lateral and axial wavenumber components, respectively. Filtering is then carried out directly on $V(k_x, k_z, \omega_0)$ using a combination of directional [Song et al., 2014, Zhao et al., 2014] radial and near-DC filters [Manduca et al., 2003], with the dual purpose of suppressing noise, reflections and quasi-static components.

Following the usual polar representation in the (k_x, k_z) plane, the radial wavenumber is defined as $k_r = \sqrt{k_x^2 + k_z^2}$ and the angular coordinate φ denotes the propagation angle of the local wavevector, measured counterclockwise from the positive k_x axis, i.e.,

$$\varphi = \mathrm{atan2}(k_z, k_x), \tag{3.4}$$

so that $k_x = k_r \cos\varphi$ and $k_z = k_r \sin\varphi$. Three separable filters are defined,

$$G_\varphi(\varphi) = \exp\left[-\tfrac{1}{2}\left(\frac{\varphi - \bar{\varphi}}{\Delta\varphi}\right)^2\right], \tag{3.5}$$

$$G_r(k_r) = \exp\left[-\tfrac{1}{2}\left(\frac{k_r - \bar{k}_r}{\Delta k_r}\right)^2\right], \tag{3.6}$$

$$G_0(k_r) = 1 - \exp\left[-\tfrac{1}{2}\left(\frac{k_r}{\Delta k_0}\right)^2\right], \tag{3.7}$$

where G_φ selects a cone of propagation directions around the mean angle $\bar{\varphi}$ with angular width $\Delta\varphi$, G_r confines the spectrum to a radial band centered at $\bar{k}_r$ with width Δk_r, and G_0 attenuates low-wavenumber components within a radius Δk_0 around the origin. As illustrated in Fig. 3.3, these filters can be tuned independently: $\bar{\varphi}$ controls the nominal direction of propagation (e.g., right- or left-going waves), $(\bar{k}_r, \Delta k_r)$ encode the expected range of shear-wave speeds $c_{\min} \leq \omega_0/k_r \leq c_{\max}$, and Δk_0 governs the rejection of quasi-static, slowly varying motion associated with the ARF push or physiological drift.

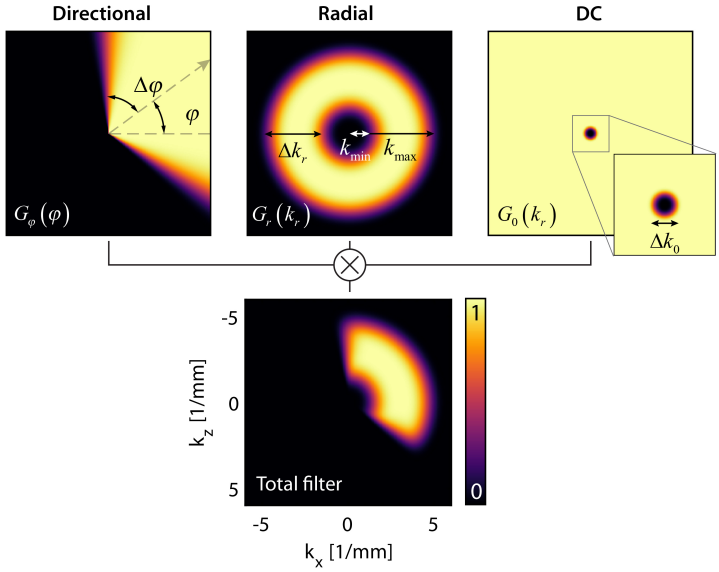


Figure 3.3: Directional, radial and near-DC filters in the (k_x, k_z) plane and their product. The angular filter $G_\varphi(\varphi)$ selects a cone of propagation directions around the mean angle $\bar{\varphi}$ with width $\Delta\varphi$. The radial filter $G_r(k_r)$ limits the spectrum to wavenumbers between $k_{\min}$ and $k_{\max}$, corresponding to a prescribed range of shear-wave speeds. The DC-rejection filter $G_0(k_r)$ suppresses low-wavenumber components within a radius Δk_0. Their product $G(k_x, k_z)$ defines the total filter used for PGE.

The total filter in the (k_x, k_z) plane is obtained as

$$G(k_x, k_z) = G_\varphi(\varphi)\, G_r(k_r)\, G_0(k_r), \tag{3.8}$$

and the filtered spectrum at ω_0 is computed by pointwise product in k-space,

$$V_{\text{filt}}(k_x, k_z, \omega_0) = V(k_x, k_z, \omega_0)\, G(k_x, k_z). \tag{3.9}$$

A product in the k-space domain corresponds to convolution in the spatial domain, so excessively narrow or sharply truncated filters would blur the wavefront and degrade spatial resolution. Conversely, overly broad filters may result in counter-propagating components that violate the local plane-wave assumption and destabilize the phase-gradient estimates. Figure 3.3 therefore serves not only to define the filter parameters, but also to emphasize this trade-off: the directional passband must be wide enough to accommodate modest deviations from the nominal propagation angle, while the radial band must track the main spectral ridge without cutting into its flanks.

The corresponding wavefield in the spatial domain is then recovered by an inverse 2D Fourier transform,

$$V_{\mathrm{filt}}(x, z, \omega_0) = \frac{1}{(2\pi)^2} \iint_{-\infty}^{\infty} V_{\mathrm{filt}}(k_x, k_z, \omega_0)\, e^{i\,(k_x x + k_z z)}\, \mathrm{d}k_x\, \mathrm{d}k_z, \tag{3.10}$$

which yields a complex-valued map in which the energy is concentrated in shear-wave components propagating within the selected angular sector and wavenumber band. Backward-traveling, low-wavenumber and out-of-band contributions are strongly attenuated, so that the residual phase varies smoothly along the main propagation direction and can be interpreted locally as that of a single dominant traveling plane wave for the purposes of the phase-gradient estimation.

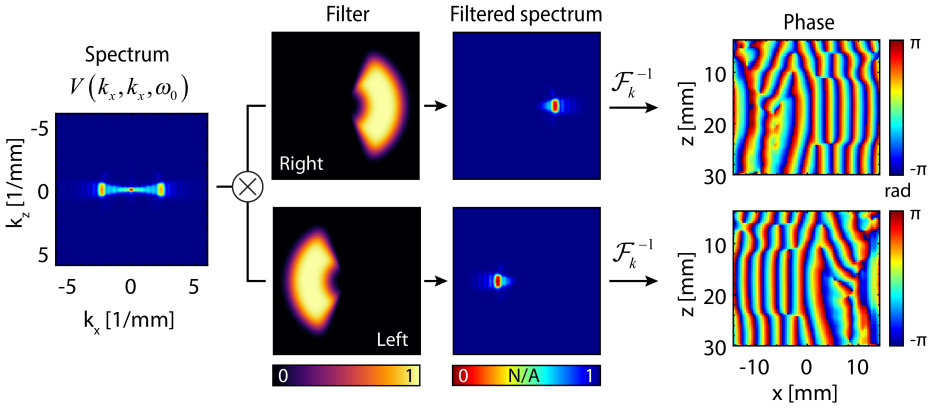


Figure 3.4: Example of directional filtering for right- and left-propagating shear waves at a fixed angular frequency ω_0. The unfiltered spectrum $V(k_x, k_z, \omega_0)$ (left) contains two symmetric lobes associated with waves traveling to the right and left. Directional filters centered on the positive and negative x-directions are applied in the (k_x, k_z) domain, yielding filtered spectra dominated by a single lobe. The inverse 2D Fourier transform produces spatial fields $V_{\mathrm{R}}(x, z, \omega_0)$ and $V_{\mathrm{L}}(x, z, \omega_0)$ whose phases (right) exhibit predominantly right- and left-going fringes, respectively, providing suitable inputs for local wavevector estimation by the phase-gradient operator.

The effect of directional filtering on a simulated shear-wave field is illustrated in Fig. 3.4, using the homogeneous elastic dataset introduced in Fig. 3.2. At ω_0, the unfiltered spectrum $V(k_x, k_z, \omega_0)$ exhibits two symmetric lobes centered near $(\pm k_s, 0)$, corresponding to waves propagating to the right and left of the ARF focus. Two complementary filters are then constructed by choosing $\bar{\varphi}$ around the positive and negative x-directions, respectively,

while keeping the same radial and DC parameters. When these filters are applied, each filtered spectrum contains primarily one of the lobes, effectively separating the right- and left-going components in the (k_x, k_z) domain.

Applying the inverse 2D Fourier transform to the filtered spectra yields two complex fields in the spatial domain, $V_{\mathrm{R}}(x, z, \omega_0)$ and $V_{\mathrm{L}}(x, z, \omega_0)$, whose phases are shown in the rightmost panels of Fig. 3.4. After directional filtering, the phase fringes become largely monotonic in the corresponding propagation direction, with interference between counter-propagating waves strongly reduced. This separation is crucial for PGE: the phase-gradient operator assumes that, within the filter passband, the local phase can be interpreted as that of a single dominant traveling plane wave. By removing backward-traveling components, near-field artifacts and low-wavenumber offsets, the directional and bandpass filters enforce this condition and yield phase maps in which the local wavevector is well defined and can be robustly estimated at each pixel in the subsequent step.

3.2.4 Phase-gradient operator and local wavenumber

After directional and bandpass filtering, the complex motion field at the working frequency ω_0,

$$V_{\mathrm{filt}}(x, z, \omega_0) = A(x, z, \omega_0)\, \exp\big[-\mathrm{i}\, \phi(x, z, \omega_0)\big], \tag{3.11}$$

is dominated by a single shear-wave component that is approximately planar within the filter defined in Sec. 3.2.3. In a general three-dimensional medium, a plane shear wave is described by a wavevector $\mathbf{k} = (k_x, k_y, k_z)$, whose magnitude $k = \|\mathbf{k}\|$ determines the phase velocity through $c_{\mathrm{ph}} = \omega_0/k$. In practice, the ultrasound acquisition only samples motion in the imaging plane $(y = 0)$, so PGE estimates the in-plane projection of the wavevector,

$$\mathbf{k}_{\mathrm{2D}} = k_x \hat{\mathbf{x}} + k_z \hat{\mathbf{z}}, \tag{3.12}$$

under the assumption that the out-of-plane component k_y is small compared with k_x and k_z in the region of interest. For an ideal plane wave propagating in the (x, z)-plane, the phase takes the form

$$\phi(x, z, \omega_0) = k_x x + k_z z + \phi_0, \tag{3.13}$$

so that the gradient of the phase is exactly equal to the in-plane wavevector,

$$\nabla\phi(x, z, \omega_0) = \frac{\partial \phi}{\partial x}\, \hat{\mathbf{x}} + \frac{\partial \phi}{\partial z}\, \hat{\mathbf{z}} = k_x \hat{\mathbf{x}} + k_z \hat{\mathbf{z}} = \mathbf{k}_{\mathrm{2D}}. \tag{3.14}$$

The purpose of the spectral filtering step is precisely to enforce conditions under which this local plane-wave approximation is valid: the field V_{filt} should be dominated by a single traveling wave whose phase varies smoothly along the main propagation direction, with backward-traveling and quasi-static components strongly suppressed.

In the discrete implementation, the wrapped phase of the filtered field is computed as

$$\phi_{i,j} = \arg\big\{V_{\mathrm{filt}}(x_i, z_j, \omega_0)\big\}, \tag{3.15}$$

where (x_i, z_j) denotes the sampling grid. A straightforward approach would be to unwrap $\phi_{i,j}$ and then estimate $\partial\phi/\partial x$ and $\partial\phi/\partial z$ using finite differences. However, numerical tests have shown that even small unwrapping errors can introduce spurious local gradients that severely affect the wavenumber calculation, especially near phase dislocations and in regions of low amplitude. To avoid these artifacts, PGE adopts a strategy that operates directly on wrapped phases and explicitly detects discontinuities.

The computation proceeds in three steps. First, centered finite differences of second-order accuracy [Fornberg, 1988] are applied to the wrapped phase map to obtain preliminary estimates of the wavevector components,

$$k_x^{(1)}(x_i, z_j, \omega_0) \approx \frac{\phi_{i+1,j} - \phi_{i-1,j}}{2\,\Delta x}, \tag{3.16}$$

$$k_z^{(1)}(x_i, z_j, \omega_0) \approx \frac{\phi_{i,j+1} - \phi_{i,j-1}}{2\,\Delta z}, \tag{3.17}$$

where Δx and Δz are the lateral and axial sampling steps. Since the phase is wrapped into $(-\pi, \pi]$, these estimates are contaminated near 2π jumps. In a second step, the same procedure is repeated on a π-shifted version of the wrapped phase,

$$\tilde{\phi}_{i,j} = \mathrm{wrap}\big(\phi_{i,j} + \pi\big), \tag{3.18}$$

yielding alternative gradients $k_x^{(2)}$ and $k_z^{(2)}$. The π-shift effectively moves the discontinuities, so locations that are close to a 2π jump in ϕ tend to be far from a jump in $\tilde{\phi}$, and vice versa.

In the third step, phase-discontinuity masks are constructed by thresholding the absolute differences between neighboring wrapped phases in both representations. Pixels where the local phase increment exceeds a prescribed threshold in one representation but not in the other are marked as unreliable for that representation and assigned to the complementary one. The final

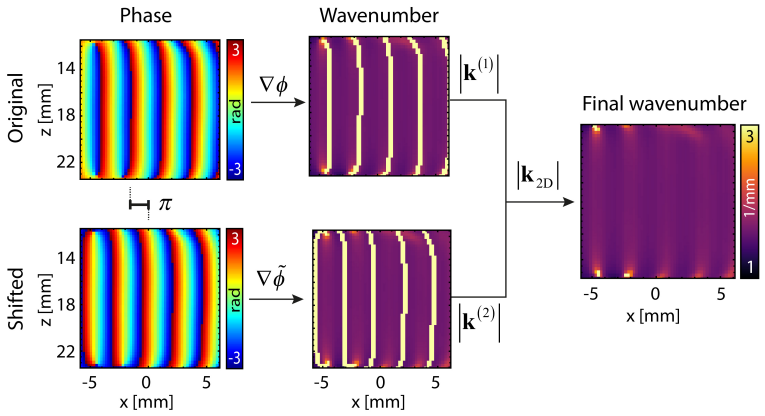


Figure 3.5: Illustration of the wavenumber computation using the phase-gradient approach for a homogeneous background with a stiffer circular inclusion. Left column: wrapped phase maps ϕ (top) and $\tilde{\phi} = \phi + \pi$ (bottom). Middle column: preliminary wavenumber magnitude maps $|\mathbf{k}^{(1)}|$ and $|\mathbf{k}^{(2)}|$ obtained by applying the phase gradient $\boldsymbol{\nabla}\phi$ and $\boldsymbol{\nabla}\tilde{\phi}$, respectively; large artifacts appear along 2π phase jumps. Right: final in-plane wavenumber magnitude $|\mathbf{k}_{2D}|$ after combining both estimates using discontinuity masks.

wavevector components are then obtained by combining the two gradient estimates using the corresponding masks,

$$k_x(x_i, z_j, \omega_0) = M_1(x_i, z_j)\, k_x^{(1)}(x_i, z_j, \omega_0) + M_2(x_i, z_j)\, k_x^{(2)}(x_i, z_j, \omega_0), \quad (3.19)$$

$$k_z(x_i, z_j, \omega_0) = M_1(x_i, z_j)\, k_z^{(1)}(x_i, z_j, \omega_0) + M_2(x_i, z_j)\, k_z^{(2)}(x_i, z_j, \omega_0), \quad (3.20)$$

where M_1 and M_2 are complementary binary masks ($M_1 + M_2 = 1$) that select, at each pixel, the most reliable gradient estimate. Pixels that lie in the vicinity of true phase dislocations may be flagged as unreliable in both representations; at these locations the gradient is not well defined for the current frequency, and the information is later recovered from neighboring frequencies when composing broadband phase-velocity maps or dispersion curves.

A convenient way to visualize this process is sketched in Fig. 3.5. The figure could display, for a given ω_0, (a) the wrapped phase map $\phi(x, z, \omega_0)$, (b) the π-shifted phase map $\tilde{\phi}(x, z, \omega_0)$, (c) the discontinuity masks derived from both representations, and (d) the resulting map of $\|\mathbf{k}_{2D}(x, z, \omega_0)\|$ after combining the two gradient estimates. Such a representation highlights how the directional filtering produces smooth, nearly planar phase fronts away from the push region.

3.2.5 Phase-velocity maps and weighted averaging

Once the in-plane wavevector has been estimated, the local phase velocity at a given angular frequency ω_0 follows directly from the dispersion relation. For each pixel (x, z), the magnitude of the 2D wavevector is

$$k_{2D}(x, z, \omega_0) = \sqrt{k_x^2(x, z, \omega_0) + k_z^2(x, z, \omega_0)}, \tag{3.21}$$

and the corresponding phase-velocity map is obtained as

$$c_{ph}(x, z, \omega_0) = \frac{\omega_0}{k_{2D}(x, z, \omega_0)}. \tag{3.22}$$

This map represents the shear-wave phase velocity associated with the dominant traveling wave selected by the spectral filters at ω_0.

In many situations, the same spatial location is probed multiple times under slightly different spectral conditions. This occurs, for instance, when several angular filters are applied around different propagation directions or when the PGE processing is repeated over a set of angular frequencies $\{\omega_m\}$ within the effective shear-wave bandwidth identified in Sec. 3.2.2. In all these cases, PGE produces a family of phase-velocity estimates

$$c_{ph}^{(r)}(x, z), \tag{3.23}$$

indexed by r, where each index can correspond to a particular frequency, a particular angular sector, or a frequency–angle pair. For each realization, a raw amplitude map

$$\tilde{A}^{(r)}(x, z) = \left|V_{filt}^{(r)}(x, z)\right| \tag{3.24}$$

is first computed from the filtered field and then converted into a binary confidence mask $A^{(r)}(x, z) \in \{0, 1\}$ by amplitude thresholding. Pixels whose amplitude exceeds 5% of the maximum spectral amplitude within the selected bandwidth are assigned $A^{(r)}(x, z) = 1$, whereas pixels below this threshold are set to $A^{(r)}(x, z) = 0$. In this way, only realizations and pixels where the shear-wave energy is sufficiently high are allowed to contribute to the final estimate. The top row of Fig. 3.6 illustrates this construction for two complementary angular filters ($\varphi = \pi$ and $\varphi = 0$), showing the directional phase-velocity maps $c_{ph}^{(1)}$ and $c_{ph}^{(2)}$ together with their associated logical masks $A^{(1)}$ and $A^{(2)}$.

Each directional map is then masked as $A^{(r)}(x, z)\, c_{ph}^{(r)}(x, z)$, removing unreliable regions such as zones where the shear-wave energy is negligible

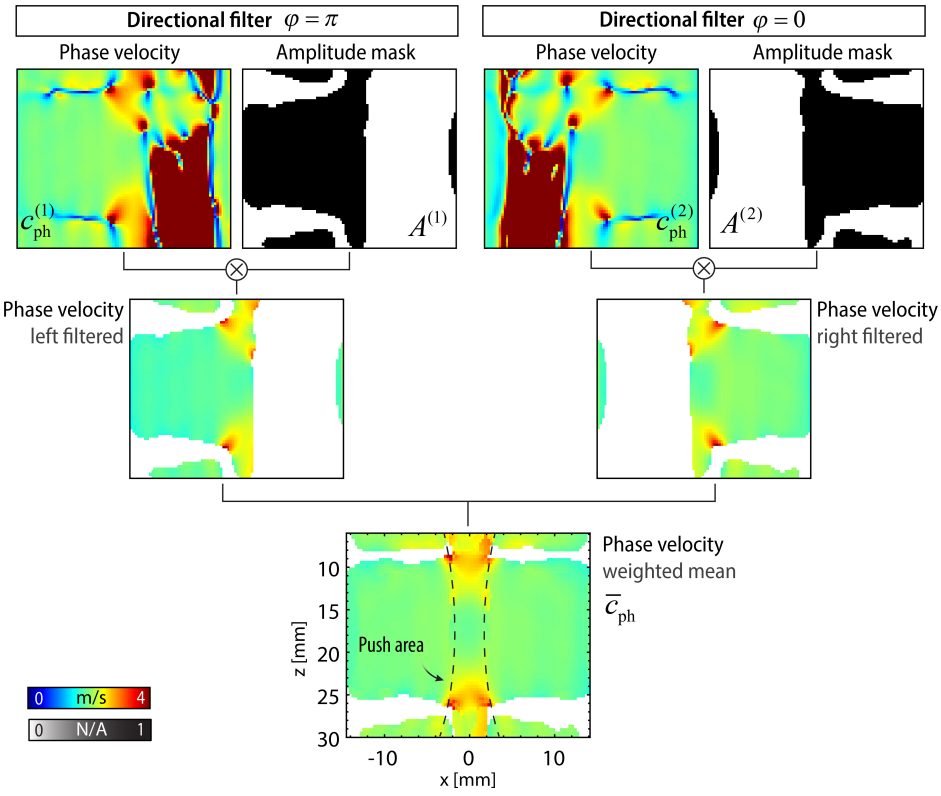


Figure 3.6: Phase-velocity reconstruction at a single angular frequency ω_0 using two directional filters ($\varphi = \pi$ and $\varphi = 0$). Top: directional phase-velocity maps $c_{\mathrm{ph}}^{(1)}$ and $c_{\mathrm{ph}}^{(2)}$ and their logical amplitude masks $A^{(1)}$ and $A^{(2)}$, obtained by thresholding the filtered-field magnitude at 5% of the maximum spectral amplitude. Middle: masked phase-velocity maps $A^{(r)} c_{\mathrm{ph}}^{(r)}$. Bottom: composite phase-velocity image $\bar{c}_{\mathrm{ph}}(x, z)$ after binary-mask averaging; the dashed lines indicate the ARF push region.

(middle row of Fig. 3.6). A final weighted phase velocity mean map is derived using [Hsiao et al., 2025],

$$\bar{c}_{\mathrm{ph}}(x, z) = \frac{\sum_r A^{(r)}(x, z)\, c_{\mathrm{ph}}^{(r)}(x, z)}{\sum_r A^{(r)}(x, z) + \varepsilon}, \tag{3.25}$$

where ε is a small regularization term that prevents division by zero in regions where the total amplitude is negligible. When the index r runs over a set of frequencies, Eq. (3.25) produces a broadband phase-velocity map in which

only frequency components with sufficient local shear-wave energy contribute. When r runs over several angular sectors, the same expression combines directional estimates into a single map while naturally excluding pixels where a given direction is poorly illuminated. The bottom panel of Fig. 3.6 shows the resulting composite phase-velocity image $\bar{c}_{\mathrm{ph}}(x, z)$ for the two-direction example, in which the left- and right-going reconstructions are merged into a single, spatially coherent representation of shear-wave speed.

3.3 Numerical and experimental validation

3.3.1 Numerical validation

The PGE method was first evaluated on simulated shear-wave datasets generated with the method described in Chapter 2. A realistic linear-array transmit/receive chain was modeled using a L11–5v probe (Verasonics Inc., Redmond, WA). The probe comprises 128 elements, each 0.27 mm wide and 5 mm high, separated by a 30 μm kerf, with an elevation focus of 18 mm and a central frequency of 7.6 MHz. Shear waves were induced by an ARF push synthesized with 32 active channels focused at 15 mm (effective $F\# = 1.74$) and a 200 μs pulse. The resulting intensity field was converted into a body-force distribution that drove the viscoelastic solver, yielding axial particle-velocity fields $v_z(x, z, t)$.

All simulated media shared the same longitudinal acoustic properties ($c =$ 1540 m/s, $\rho = 1000$ kg/m^3, absorption 0.5 dB/MHz/cm), while shear behavior followed a Kelvin–Voigt model. Three phantom families were considered: (i) purely elastic media with $\eta = 0$ and $\mu \in [2.25, 30.25]$ kPa to isolate stiffness effects without dispersion; (ii) viscoelastic media with $\mu = 6.25$ kPa and $\eta \in [0, 2]$ Pa·s to study increasing attenuation/dispersion; and (iii) a heterogeneous case with an 8 mm-diameter spherical inclusion ($\mu = 16$ kPa) in a background ($\mu = 4$ kPa), both with $\eta = 0$, to test stiffness-contrast imaging. Representative homogeneous viscoelastic cases are reported in Sec. 3.3.1, and the heterogeneous results in Sec. 3.3.1.

Across all simulations, PGE used a fixed processing configuration. In the k_x-k_z space, right-going components were selected with an angular filter centered at $\bar{\varphi} = 0$ rad and width $\Delta\varphi = \pi/4$. The radial band-pass was defined by a phase-velocity interval $[c_{\min}, c_{\max}]$. For homogeneous media,

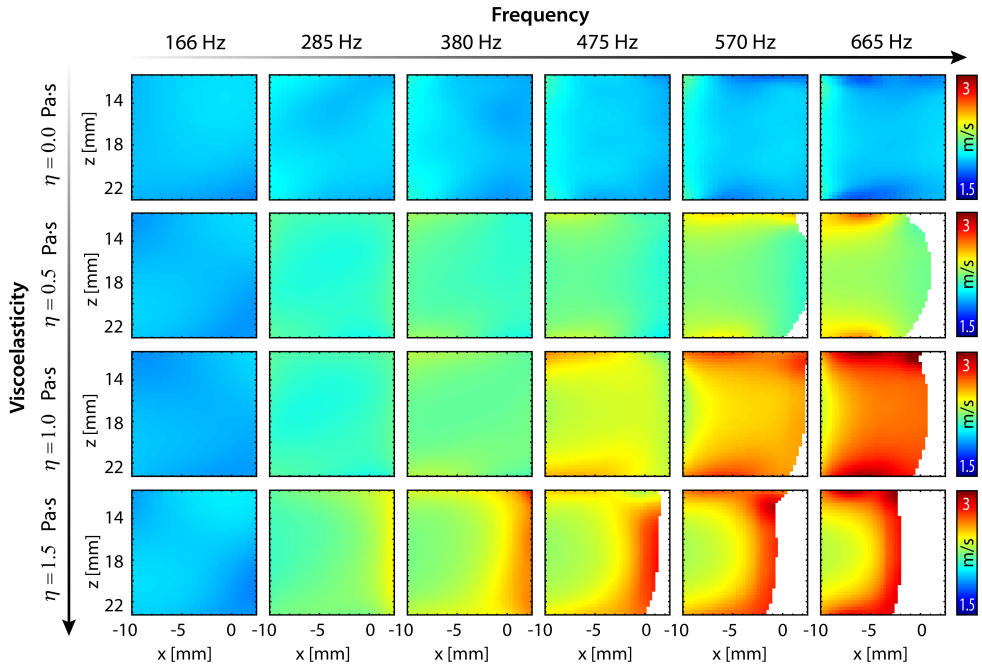


Figure 3.7: Phase-velocity maps obtained with PGE in four homogeneous viscoelastic media with fixed shear modulus $\mu = 4$ kPa and viscosities $\eta = 0.0,\ 0.5,\ 1.0,\ 1.5$ Pa · s. Rows correspond to increasing viscosity and columns to six representative frequencies within the shear-wave band.

$[c_{\min}, c_{\max}]$ was centered on the mean phase velocity $c_{\mathcal{F}_{2D}}$ obtained with the 2D-FT method [Catheline et al., 2004] at the central shear-wave frequency, using $c_{\min} = c_{\mathcal{F}_{2D}} - c_k$ and $c_{\max} = c_{\mathcal{F}_{2D}} + c_k$ with $c_k = 0.5$ m/s [Kijanka and Urban, 2019]. For heterogeneous media, a fixed range $0.5 \leq c_{\mathrm{ph}} \leq 10$ m/s was used. Quasi-static components were suppressed with a DC-rejection radius $\Delta k_0 = k_{\min}/8$, where $k_{\min} = \omega_0/c_{\max}$.

For global spectral reference, all datasets were also analyzed with the conventional 2D-FT approach [Catheline et al., 2004]. In heterogeneous media, a local phase-velocity imaging (LPVI) method was additionally implemented for comparison [Kijanka and Urban, 2019, Kijanka and Urban, 2020]. For each medium and ROI, dispersion curves $c_{\mathrm{ph}}(f)$ were obtained by ROI-averaging PGE maps or by applying the reference estimators over depth-resolved windows, and then fitted to the Kelvin–Voigt dispersion relation to estimate μ and η. The same fit was applied pixel-wise to produce spatial maps $\mu(x, z)$ and $\eta(x, z)$.

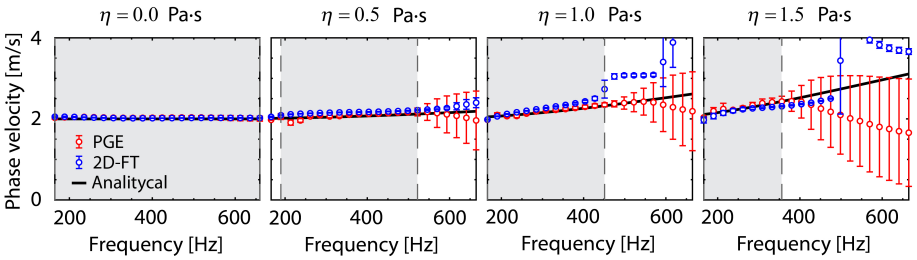


Figure 3.8: Dispersion curves in the four homogeneous phantoms of Fig. 3.7. For each medium, the mean phase velocity was computed over the spatial ROI shown in the corresponding phase-velocity maps of Fig. 3.7. Red markers show PGE estimates and blue markers show 2D-FT estimates, together with the analytical Kelvin–Voigt prediction (black line). Shaded regions denote the frequency bands used for Kelvin–Voigt fitting.

Homogeneous elastic and viscoelastic phantoms

PGE was first validated in homogeneous Kelvin–Voigt media with fixed shear modulus $\mu = 4$ kPa and viscosities $\eta = 0.0,\ 0.5,\ 1.0,\ 1.5$ Pa $\cdot$ s. For each dataset, the common processing and filter settings of Sec. 3.3.1 were applied to obtain monochromatic phase-velocity maps $c_{\mathrm{ph}}(x, z, f_i)$ over the shear-wave bandwidth. Figure 3.7 illustrates the resulting maps across representative frequencies. In the elastic case ($\eta = 0$), c_{ph} remains essentially uniform and frequency-independent, matching the value expected from $\mu = 4$ kPa. As viscosity increases, the maps exhibit the expected positive dispersion (higher c_{ph} at higher frequency) while progressively losing spatial uniformity at the highest frequencies, consistent with increased attenuation and reduced SNR away from the focal region; in the most viscous media this also leads to masked pixels where the spectral amplitude is insufficient.

The corresponding dispersion curves are summarized in Fig. 3.8, where $c_{\mathrm{ph}}(f)$ was extracted by ROI-averaging the PGE maps and compared against a conventional 2D-FT estimator and the analytical Kelvin–Voigt prediction. Both methods yield nearly flat curves for $\eta = 0$ and follow the theoretical dispersion for $\eta = 0.5$ Pa $\cdot$ s. For $\eta = 1.0$ and 1.5 Pa $\cdot$ s, the usable frequency band narrows as high-frequency components are rapidly attenuated, and deviations and uncertainty increase near the band edge, in agreement with the spatial degradation observed in Fig. 3.7. The shaded regions indicate the frequency ranges retained for Kelvin–Voigt fitting, selected from the motion spectrum to exclude noise-dominated bins.

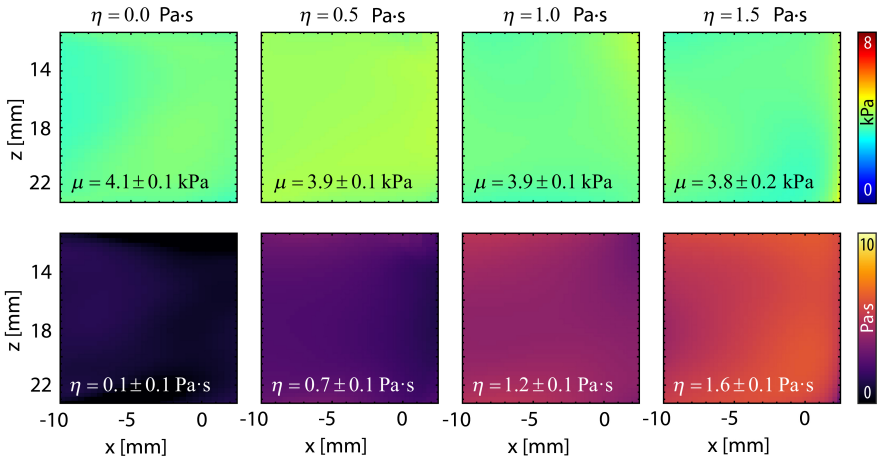


Figure 3.9: Pixel-wise Kelvin–Voigt fitting of the phase-velocity maps in homogeneous viscoelastic phantoms. Each column corresponds to a given viscosity $\eta = 0.0,\ 0.5,\ 1.0,\ 1.5$ Pa·s, and the two rows show the recovered shear elastic modulus $\mu(x, z)$ and shear viscosity $\eta(x, z)$, respectively.

Beyond ROI-based analysis, pixel-wise Kelvin–Voigt fitting was performed to recover spatial maps $\mu(x, z)$ and $\eta(x, z)$ (Fig. 3.9). The estimated μ remains close to the nominal 4 kPa and largely uniform across viscosities, with minor edge degradation when the available fitting bandwidth is reduced. The recovered η maps are near zero for the elastic phantom and remain approximately constant at the imposed levels for viscoelastic cases, while naturally deteriorating in low-amplitude peripheral regions.

Finally, Fig. 3.10 condenses the broader parameter sweep introduced in Sec. 3.3.1. In purely elastic media, recovered μ closely tracks the theoretical values with only small relative errors. In the viscoelastic sweep at fixed μ, μ remains stable while the fitted η follows the ground truth up to $\eta \approx 1.5$–2 Pa·s, beyond which attenuation and bandwidth reduction increasingly limit fitting accuracy.

Heterogeneous phantoms

The spatial performance of the PGE method in the presence of stiffness contrast was evaluated using the heterogeneous phantom introduced in Sec. 3.3.1. The medium consisted of a purely elastic background with shear modulus $\mu = 4$ kPa

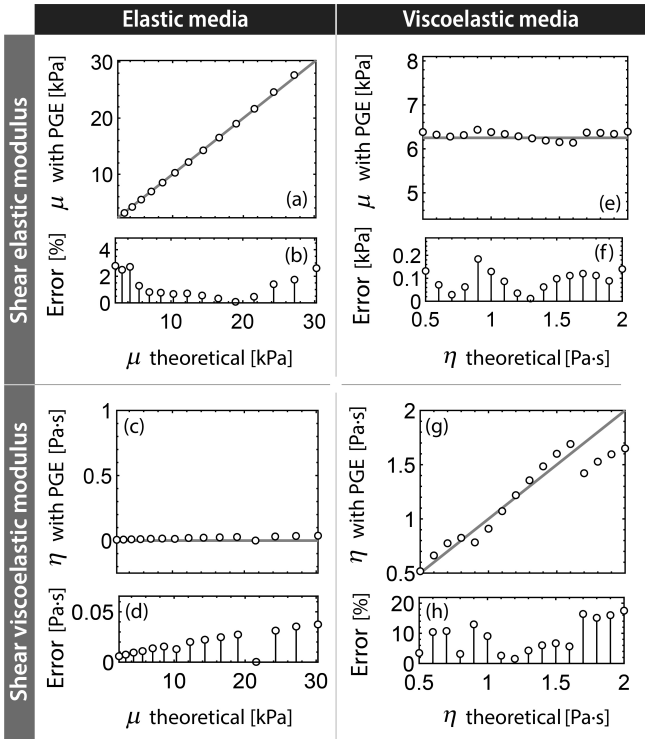


Figure 3.10: Systematic assessment of PGE-based Kelvin–Voigt parameter estimation in simulated homogeneous media. Recovered shear elastic (a) and viscosity (c) modulus as a function of theoretical μ for purely elastic media; (b) anc (d) represent the corresponding relative error. Recovered shear elastic (e) and viscosity (g) modulus versus theoretical η; (f) and (h) present the relative error. Viscosity is accurately retrieved up to $\eta \approx 1.5$–2 Pa $\cdot$ s, beyond which strong attenuation and a reduced useful frequency band degrade the estimation.

and viscosity $\eta = 0$ Pa $\cdot$ s, containing a single spherical inclusion of diameter 8 mm with $\mu = 16$ kPa and the same viscosity. For comparison, a local phase-velocity imaging (LPVI) algorithm based on block-wise wavenumber estimation was implemented following [Kijanka and Urban, 2019, Kijanka and Urban, 2020], using 2×2 mm windows to balance spatial resolution and robustness.

Figure 3.11 compares phase-velocity maps reconstructed with PGE (top row) and LPVI (bottom row) at three frequencies (800, 1000 and 1200 Hz), together with the mean maps obtained by averaging over the 800–1200 Hz band. In all cases, the stiff inclusion appears as a nearly circular region of higher shear-wave speed surrounded by a slower background. The reconstructed

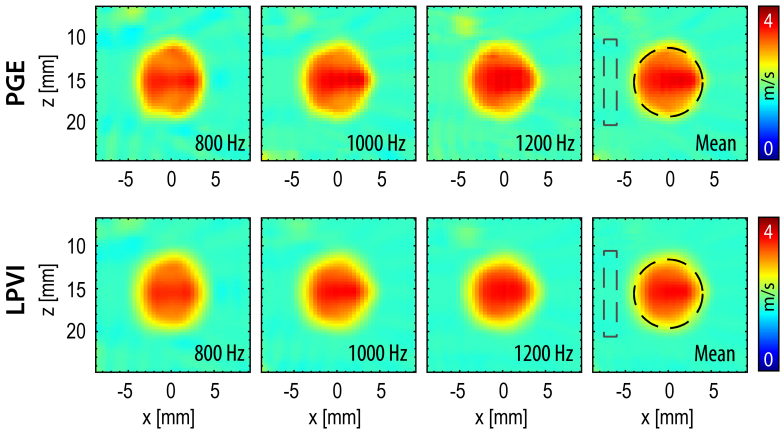


Figure 3.11: Phase-velocity maps in the heterogeneous simulated phantom at 800, 1000 and 1200 Hz (columns 1–3) and for the frequency-averaged image (column 4). Top row: PGE; bottom row: LPVI. The circular high-velocity region corresponds to the stiff inclusion, and the dashed circle and rectangle mark the inclusion and background ROIs used for quantitative analysis.

velocity levels are consistent with the nominal contrast between $\mu = 4$ and 16 kPa. At individual frequencies, both methods recover a smooth, frequency-independent inclusion whose apparent size agrees well with the true geometry. The mean maps further reduce residual speckle and highlight the inclusion and background regions used for quantitative analysis: a circular ROI centred on the inclusion (dashed circle) and a lateral rectangular ROI in the background (dashed box). Visual inspection shows that PGE and LPVI produce very similar mean maps, with comparable edge sharpness and contrast, despite relying on different local estimators (pixel-wise phase gradient versus block-based wavenumber).

The agreement between both methods is examined more closely in Fig. 6.3. Panels 6.3(a) and 6.3(b) display lateral and axial phase-velocity profiles extracted through the centre of the inclusion from the mean maps of Fig. 3.11. In both directions, the profiles show a clear step in phase velocity between background and inclusion, and the positions of the left/right and lower/upper edges (vertical dashed lines) coincide with the known inclusion boundaries. PGE and LPVI curves nearly overlap, indicating that both methods offer similar spatial resolution at the chosen frequency band. Small differences in the slope of the transition regions can be attributed to the intrinsic smoothing

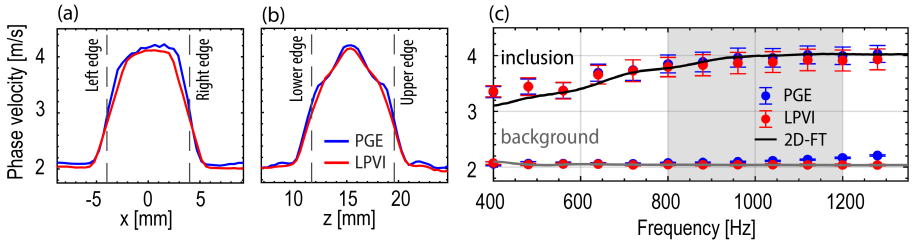


Figure 3.12: Quantitative comparison between PGE and LPVI in the heterogeneous phantom. (a) Lateral and (b) axial phase-velocity profiles across the centre of the inclusion, with vertical dashed lines indicating the estimated edges. (c) Phase-velocity dispersion curves in inclusion and background for PGE and LPVI, with 2D-FT providing a global spectral reference and error bars denoting the standard deviation within each ROI.

introduced by the LPVI block size, whereas PGE operates at the pixel scale; however, neither method exhibits spurious oscillations or ringing artefacts around the interface.

Finally, Fig. 6.3(c) reports phase-velocity dispersion curves extracted from the inclusion and background ROIs. For each frequency, PGE and LPVI estimates are averaged within the corresponding ROI, and the resulting curves $c_{\mathrm{ph}}(f)$ are compared with reference values obtained from a 2D-FT analysis performed over extended lateral–temporal windows spanning the same regions. As expected for a purely elastic medium, the dispersion curves for both background and inclusion are nearly flat over the analysed band, with a constant velocity offset reflecting the stiffness contrast. PGE, LPVI and 2D-FT provide very similar phase-velocity levels, and the spread of the PGE and LPVI data around the 2D-FT reference remains small. These results indicate that, in the presence of moderate stiffness contrast, PGE yields local phase-velocity maps and dispersion curves that are consistent with established spectral approaches, while providing full-field images at each frequency without the need to select an explicit block size.

3.3.2 Experimental phantom validation

Experimental validation was conducted on a CIRS 049 elasticity phantom (Computerized Imaging Reference Systems, Norfolk, USA) featuring a homogeneous background and four spherical inclusions (Type I–IV) with increasing stiffness and comparable echogenicity. From the manufacturer Young's moduli

(background $E = 18$ kPa; inclusions $E = 6, 9, 36, 70$ kPa) and assuming an incompressible isotropic elastic solid,

$$E \approx 3\mu = 3\rho_0 c_s^2, \tag{3.26}$$

reference shear-wave speeds were $c_s \approx 2.4$ m/s for the background and $\approx$ 1.4, 1.7, 3.5, 4.8 m/s for Type I–IV.

Data were acquired with the same L11–5v probe of Sec. 3.3.1. Plane-wave tracking was performed at 10 kHz PRF and axial particle velocity $v_z(x, z, t)$ was estimated from IQ data using the 2D Loupas autocorrelator [Loupas et al., 1995]. For each inclusion, the wavefields were cropped and processed with PGE using the same directional and radial filtering as in simulations; LPVI was also applied for comparison.

Phase-velocity maps $c_{\mathrm{ph}}(x, z, f)$ were computed over the shear-wave band and frequency-compounded over high-SNR intervals: 350–550 Hz for the softer inclusions (Type I–II) and 415–810 Hz for the stiffer ones (Type III–IV). The resulting mean maps (Fig. 3.13) delineate all inclusions with the expected ordering relative to the background (Type I–II slower; Type III–IV faster) while preserving an approximately uniform background. ROIs within each sphere and in the surrounding material (Fig. 3.13) were used for quantitative analysis. Local edge artefacts are more apparent in the softer inclusions, consistent with wavefront refraction and out-of-plane contributions in a 2D acquisition, whereas the stiffer inclusions yield smoother reconstructions.

Dispersion curves were obtained by ROI-averaging PGE and LPVI estimates at each frequency and compared against a 2D-FT reference computed over extended lateral segments (Fig. 3.13, bottom row). For all regions, PGE and LPVI closely follow the 2D-FT curves, which are nearly flat over the analysed bands, as expected for a predominantly elastic phantom.

Table 3.1 reports mean phase velocities and standard deviations in the background and inclusions. The background is consistently recovered near 2.4 m/s, and both PGE and LPVI provide similar estimates within the spheres: agreement is closest for Type I, with moderate underestimation for the stiffer inclusions (Type II–IV). Overall, these measurements show that PGE yields phase-velocity maps and averaged values consistent with LPVI and manufacturer-derived references while operating at full image-grid resolution.

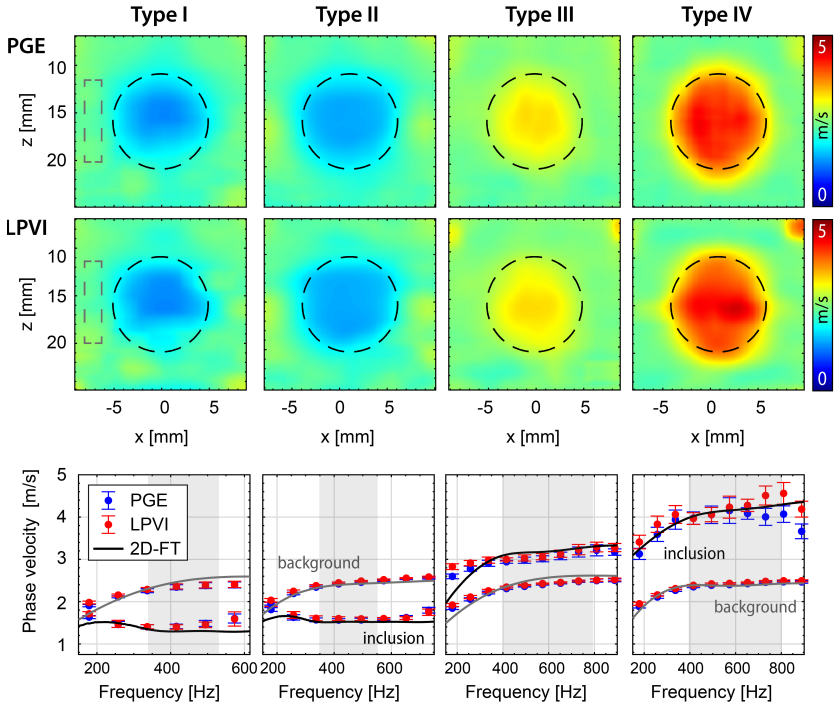


Figure 3.13: Phase-velocity imaging in the experimental elasticity phantom. Top row: mean phase-velocity maps reconstructed with PGE for the four spherical inclusions (Type I–IV, from softest to stiffest) and the surrounding background. Middle row: corresponding LPVI maps; circular and rectangular ROIs indicate inclusion and background regions used for quantitative analysis. Bottom row: phase-velocity dispersion curves in inclusion and background for PGE and LPVI, with 2D-FT providing a global spectral reference and error bars denoting the standard deviation within each ROI.

Mean c_{ph} (m/s)	Type I		Type II		Type III		Type IV	
	Mean ± Std	Ref.	Mean ± Std	Ref.	Mean ± Std	Ref.	Mean ± Std	Ref.
Inclusion PGE	1.42 ± 0.06	1.41	1.52 ± 0.04	1.73	3.17 ± 0.08	3.46	4.05 ± 0.18	4.83
Inclusion LPVI	1.40 ± 0.07		1.55 ± 0.04		3.23 ± 0.08		4.26 ± 0.23	
Background PGE	2.30 ± 0.06	2.45	2.39 ± 0.02	2.45	2.51 ± 0.02	2.45	2.42 ± 0.02	2.45
Background LPVI	2.32 ± 0.06		2.41 ± 0.02		2.54 ± 0.02		2.44 ± 0.02	

Table 3.1: Mean shear wave velocity and standard deviation were calculated for the four spherical inclusions and background of the phantom using the phase gradient elastography (PGE), local wavenumber approach (LPVI), and manufacturer (Ref.) methods.

3.3.3 *Ex vivo* bovine liver validation

Validation was extended to real tissue using an *ex vivo* bovine liver sample. A 10 mm-diameter cylindrical core was extracted, immersed in a 10% formalin

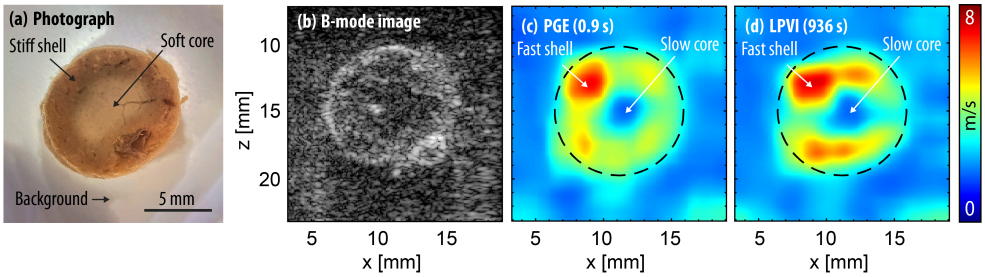


Figure 3.14: Experimental validation in an ex vivo bovine liver sample. (a) Photograph of a sliced sample, where a darker, formalin-stiffened shell surrounds a lighter core. (b) B-mode ultrasound image showing the cylindrical liver inclusion embedded in a gelatin phantom. (c) Mean phase-velocity map between 350 and 550 Hz reconstructed with the PGE method. (d) Corresponding LPVI map.

bath, and placed in vacuum for 30 min to remove trapped air and promote homogeneous fixation. This protocol produced a mechanically heterogeneous cylinder, with formalin predominantly stiffening the outer layers while leaving a comparatively softer inner core. The treated sample was embedded in a tissue-mimicking gelatin phantom tuned to a shear-wave speed of ~ 2.5 m/s following [Cook et al., 2011]. Measurements used the same Verasonics L11–5v configuration, ARF push sequence, and processing parameters as in Sec. 3.3.2.

Figure 3.14 reports the outcomes. The inclusion is only weakly contrasted in B-mode, Fig. 3.14(b), whereas the photograph of the sliced sample, Fig. 3.14(a) shows a darker peripheral ring around a lighter core, suggesting a fixation-induced stiffness gradient not directly visible in ultrasound intensity.

Phase-velocity maps were reconstructed with PGE and LPVI over six discrete frequencies spanning 350–550 Hz, selected from the motion spectrum to balance SNR and spatial resolution. Figures. 3.14(c,d) illustrate the frequency-averaged phase velocity maps across 350–550 Hz. Both methods consistently reveal three regions: (i) an approximately uniform background, (ii) a high-velocity annulus aligned with the liver periphery (stiffened shell), and (iii) a lower-velocity central core, consistent with weaker cross-linking in the interior.

The reconstructed field of view contains 155×315 pixels with $\Delta x = 0.10$ mm and $\Delta z = 0.05$ mm. LPVI used a 2.5 mm block size to resolve the central region while maintaining robust spectral estimation, yielding slightly smoother interfaces due to spatial averaging, whereas PGE, computed at pixel scale, provided sharper delineation of the shell–core boundaries.

Computation time strongly differed between methods. For the 155×315 grid, reconstructing a single-frequency map required ~ 156.7 s with LPVI-2D versus 0.16 s with PGE. Over the six frequencies, total reconstruction time was 935.5 s (≈ 15.6 min) for LPVI-2D compared with 0.9 s for PGE. Despite this $\sim 10^3$ speedup, both approaches produced consistent elastographic patterns, supporting the use of PGE for rapid characterization of stiffness variations in real tissue.

3.4 Computational complexity

Local phase velocity imaging (LPVI) was used in this work as a spectral benchmark for PGE. In LPVI-2D [Kijanka and Urban, 2019], local phase velocity is obtained by tiling the shear-wave motion field with overlapping lateral–axial blocks, computing a 2D Fourier transform per block, and extracting the dominant wavevector from the spectrum peak. A faster variant, LPVI-1D [Kijanka and Urban, 2020], reduces cost by performing 1D transforms along the propagation direction under an approximately planar-wave assumption. In both cases, block-wise processing yields a phase-velocity estimate at each block centre but requires many local Fourier transforms per image.

In contrast, PGE estimates $c_{\mathrm{ph}}(x, z, f)$ from the spatial phase gradient of a globally filtered monochromatic field. For each frequency, the data are transformed to wavenumber space once, directional and radial filters isolate the dominant component, and the filtered field is transformed back to the spatial domain, where phase gradients are computed by finite differences at every pixel. This replaces repeated local transforms with a small number of global FFTs and point-wise operations, without requiring explicit block extraction.

The relationship between computation time and image size was quantified in Fig. 3.15. Images of varying size, ranging from 2.5×10^3 to 120×10^3 pixels, were processed to determine the time required to reconstruct a single-frequency phase-velocity map. A linear model $t = a\,n + b$ was fitted to the data, where n is the number of pixels and t the computation time. The fitted parameters confirm that all three methods exhibit an approximately linear algorithmic complexity $\mathcal{O}(n)$, as expected for FFT-based implementations with fixed filter settings. However, the slopes of the fits differ by several orders of magnitude: the LPVI-1D slope is roughly 7×10^2 times larger than that of PGE, and the

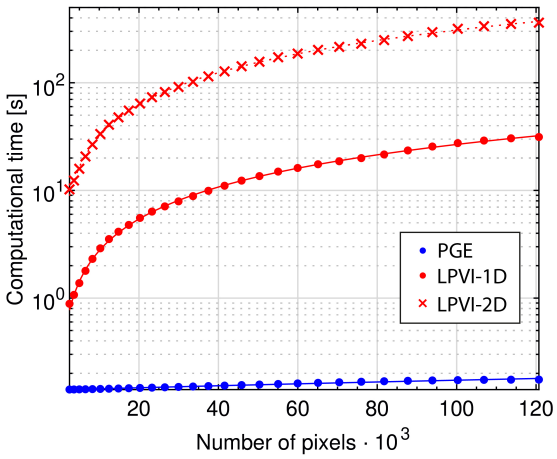


Figure 3.15: Computation time as a function of the number of pixels n for the proposed PGE method (blue) and the LPVI comparison method (red). LPVI is evaluated in both its one-dimensional (LPVI-1D, red circles) and two-dimensional (LPVI-2D, red crosses) implementations. Solid lines show linear fits $t = a\,n + b$, indicating approximately linear complexity for all methods, with markedly steeper slopes for the LPVI variants.

LPVI-2D slope is close to 8×10^3 times larger.

These trends match the experimental timings: a 150×150-pixel phase-velocity map is reconstructed in ~ 0.14 s with PGE versus ~ 72.8 s with a parallelized LPVI-2D implementation on the same desktop workstation. Notably, LPVI was explicitly optimized and parallelized, whereas PGE achieved these runtimes without comparable optimization. Overall, LPVI retains advantages in block-design flexibility and prior validation, while PGE provides similar elastographic information at a substantially lower computational cost, enabling real-time or near real-time workflows.

3.5 Discussion

This chapter evaluated phase-gradient elastography (PGE) as a spectral shear-wave imaging strategy that estimates the local in-plane wavevector from the phase gradients of a globally filtered monochromatic field, rather than from repeated local spectral searches. This design choice is central to the observed behaviour of the method: once the displacement/velocity field has been narrowed to a single dominant travelling component at a given frequency,

local phase gradients provide dense pixel-wise estimates of (k_x, k_z), and thus of $c_{\mathrm{ph}}(x, z, f)$, without the block-size trade-off intrinsic to other spectral approaches. In practice, this turns a traditionally heavy spectral inversion problem into an operation dominated by a small number of global FFTs plus lightweight point-wise arithmetic, i.e., a workflow that is compatible with clinical timing constraints rather than merely aspirational about them.

A first implication is that the accuracy of PGE is primarily governed by how well the preprocessing step enforces the "single locally planar travelling wave" condition. In shear-wave elastography, this condition is challenged by reflections, multi-path interference, and wavefront curvature, all of which can compromise the phase structure and destabilize gradient-based wavenumber estimates [Lipman et al., 2016]. Moreover, phase-gradient estimators are intrinsically sensitive to noise and motion artefacts, so that low-SNR regions can translate into locally noisy phase-velocity values [Deffieux et al., 2009]. The combined angular, radial and near-DC filtering used in this chapter therefore plays a dual role: (i) it rejects counter-propagating and quasi-static components that violate the travelling-wave model, and (ii) it regularizes the phase sufficiently for finite-difference gradients to remain stable. The results in homogeneous media are consistent with this interpretation: when the shear-wave spectrum is sufficiently broad and energetic, PGE retrieves dispersion curves and Kelvin–Voigt parameters with low error; when viscosity becomes large, attenuation narrows the usable band and the estimator becomes progressively more noise-limited, delimiting the practical operating range.

In heterogeneous media, the experiments and simulations indicate that PGE preserves stiffness contrast and achieves boundaries that are effectively resolved at the scale of approximately one shear wavelength, in agreement with the behaviour expected from a frequency-resolved wave imaging method. Importantly, PGE reaches this spatial detail without explicitly specifying a block size, whereas LPVI achieves robustness by spatial averaging and therefore produces smoother interfaces by construction. These complementary behaviours suggest a useful interpretation: PGE tends to be resolution-forward (pixel-scale estimates that depend on spectral fidelity), whereas LPVI is robustness-forward (local spectral estimates stabilized by averaging). In practice, the relative advantage depends on the wavefield complexity and on the signal bandwidth available at the selected frequency.

The experimental phantom results further highlight a fundamental limitation shared by all 2D spectral estimators applied to projected displacement

fields: only the in-plane projection of the true 3D wavevector is accessible. In the softer inclusions, refraction can effectively focus energy and increase the contribution of out-of-plane components, leading to local artefacts and distortions near interfaces; for stiffer inclusions, refraction tends to steer energy away from the imaging plane, and the reconstructed maps typically appear smoother. This behaviour, observed in this chapter, reinforces that local biases can arise near strong heterogeneities when wave propagation departs from an in-plane model. From a methodological standpoint, this motivates two practical recommendations: (i) interpret interface-adjacent pixels with caution in strongly refractive settings, and (ii) favour frequency bands where the wavefront remains coherent in-plane and the directional spectrum is well separated.

The *ex vivo* liver experiment illustrates the most clinically relevant outcome of PGE: stiffness patterns comparable to LPVI can be recovered while reducing reconstruction time by orders of magnitude. This is not merely a performance metric; it changes how spectral elastography can be used. Frequency compounding, multi-frequency dispersion analysis, and repeated acquisitions (e.g., for averaging or motion rejection) become feasible without incurring prohibitive latency. In other words, PGE behaves like a spectral method that can keep up with the scanner instead of asking the scanner to slow down.

Finally, several implementation choices are worth making explicit because they condition reproducibility and robustness. First, the selection of the frequency band from the k–ω spectrum is not cosmetic: it determines the available wavelength content and the SNR regime in which phase gradients are computed. Second, the directional and radial filter widths must balance rejection of interfering components against spatial blurring introduced by overly narrow passbands. Third, upstream displacement/velocity estimation sets the noise floor; in low-motion regions, improving tracking SNR can be as important as refining the spectral inversion. These considerations collectively suggest that the next technical step for broad deployment is not a more complex phase-gradient operator, but rather more systematic, data-driven strategies to adapt frequency selection and filter parameters to the measured wavefield.

3.6 Conclusions

This chapter presented and validated phase-gradient elastography as a computationally efficient spectral method to reconstruct frequency-dependent phase-velocity maps from ARF-induced shear waves. By estimating the local in-plane wavevector from the spatial gradients of the phase of a filtered monochromatic field, PGE provides dense pixel-wise maps $c_{\mathrm{ph}}(x, z, f)$ without explicit block processing.

Across numerical validations in homogeneous Kelvin–Voigt media, PGE reproduced the expected phase-velocity dispersion and enabled accurate recovery of shear modulus and viscosity within the regime where the shear-wave spectrum maintains sufficient energy and adequate resolution. In heterogeneous configurations, PGE recovered stiffness contrasts consistent with established spectral references while preserving spatial detail at the scale expected from wavelength-limited imaging.

Experimental results in elasticity phantoms and an *ex vivo* bovine liver sample confirmed that PGE translates to real acquisitions, delineating inclusions and layered stiffness patterns that are not evident in B-mode. Critically, the chapter quantified that this elastographic information can be obtained at a computational cost that is orders of magnitude lower than other benchmark methods, supporting real-time or near real-time spectral shear-wave imaging.

The main limitations identified are linked to wavefield complexity and data quality: reduced bandwidth in highly attenuative media, sensitivity to residual reflections/noise when filtering is insufficient, and intrinsic effects in the presence of out-of-plane propagation near strong heterogeneities. Within these constraints, PGE emerges as a practical route to bring frequency-resolved shear-wave elastography closer to routine, low-latency imaging.

Publications

This section contains a compendium of the seminal scientific articles in which the results presented in this chapter have been published.

Journal articles

- González-Mateo, E., Camarena, F., and Jimenez, N., "Real-time ultrasound shear wave elastography using a local phase gradient," *Computer Methods and Programs in Biomedicine*, vol. 260, 2025.

Conferences articles

- González-Mateo, E., Camarena, F., and Jimenez, N., "Acoustic vortices and phase gradient techniques for ultrasound elastography," in *International Tissue Elasticity Conference (ITEC 2024)*, Lyon, France, 2024.
- González-Mateo, E., Camarena, F., and Jimenez, N., "Enhanced Real-Time Ultrasound Elastography with Phase Gradient," in *IEEE Ultrasonics, Ferroelectrics, and Frequency Control Joint Symposium (UFFC-JS 2024)*, Taipei, Taiwan, 2024.
- González-Mateo, E., Camarena, F., and Jimenez, N., "Fast shear wavevector estimator for real-time elasticity map reconstruction," in *2023 IEEE International Ultrasonics Symposium (IUS)*, Montreal, Canada, 2023.
- González-Mateo, E., Jimenez, N., and Camarena, F., "Mapeado de la elasticidad de tejidos mediante un método rápido basado en gradiente de fase," in *TECNIACÚSTICA 2022. 53º Congreso Español de Acústica. XII Congreso Ibérico de Acústica. European Symposium on NVH in Electric Vehicles*, Elche, España, 2022, pp. 1127–1132.

References

Bercoff, J., Criton, A., Bacrie, C., Souquet, J., Tanter, M., Gennisson, J., Deffieux, T., Fink, M., Juhan, V., Colavolpe, A., Amy, D., and Athanasiou, A. (2008). ShearWave™ Elastography A new real time imaging mode for assessing quantitatively soft tissue viscoelasticity. In *2008 IEEE Ultrasonics Symposium*, pages 321–324.

Catheline, S., Gennisson, J.-L., Delon, G., Fink, M., Sinkus, R., Abouelkaram, S., and Culioli, J. (2004). Measurement of viscoelastic properties of homogeneous soft solid using transient elastography: An inverse problem approach. *The Journal of the Acoustical Society of America*, 116(6):3734–3741.

Chen, S., Fatemi, M., and Greenleaf, J. F. (2004). Quantifying elasticity and viscosity from measurement of shear wave speed dispersion. *The Journal of the Acoustical Society of America*, 115(6):2781–2785.

Cook, J. R., Bouchard, R. R., and Emelianov, S. Y. (2011). Tissue-mimicking phantoms for photoacoustic and ultrasonic imaging. *Biomedical Optics Express*, 2(11):3193–3206.

Deffieux, T., Montaldo, G., Tanter, M., and Fink, M. (2009). Shear Wave Spectroscopy for In Vivo Quantification of Human Soft Tissues Visco-Elasticity. *IEEE Transactions on Medical Imaging*, 28(3):313–322.

Fornberg, B. (1988). Generation of finite difference formulas on arbitrarily spaced grids. *Mathematics of Computation*, 51(184):699–706.

Hsiao, W.-C., Yeh, C.-L., Kao, Y.-C. J., Chu, S.-Y., and Hsieh, B.-Y. (2025). Spatiotemporal monitoring of tissue coagulation treated by continuous high-intensity focused ultrasound with pulse-inversion shear wave elastography. *Ultrasonics*, 156:107766.

Kijanka, P. and Urban, M. W. (2019). Local Phase Velocity Based Imaging: A New Technique Used for Ultrasound Shear Wave Elastography. *IEEE Transactions on Medical Imaging*, 38(4):894–908.

Kijanka, P. and Urban, M. W. (2020). Fast Local Phase Velocity-Based Imaging: Shear Wave Particle Velocity and Displacement Motion Study. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 67(3):526–537.

Li, G.-Y. and Cao, Y. (2017). Assessing the mechanical properties of anisotropic soft tissues using guided wave elastography: Inverse method and numerical experiments. *The Journal of the Acoustical Society of America*, 142(3):1526–1536.

Lipman, S. L., Rouze, N. C., Palmeri, M. L., and Nightingale, K. R. (2016). Evaluating the Improvement in Shear Wave Speed Image Quality Using Multidimensional Directional Filters in the Presence of Reflection Artifacts. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 63(8):1049–1063.

Loupas, T., Powers, J., and Gill, R. (1995). An axial velocity estimator for ultrasound blood flow imaging, based on a full evaluation of the Doppler equation by means of a two-dimensional autocorrelation approach. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 42(4):672–688.

Manduca, A., Lake, D. S., Kruse, S. A., and Ehman, R. L. (2003). Spatio-temporal directional filtering for improved inversion of MR elastography images. *Medical Image Analysis*, 7(4):465–473.

Manduca, A., Oliphant, T. E., Dresner, M. A., Mahowald, J. L., Kruse, S. A., Amromin, E., Felmlee, J. P., Greenleaf, J. F., and Ehman, R. L. (2001). Magnetic resonance elastography: non-invasive mapping of tissue elasticity. *Medical Image Analysis*, 5(4):237–254.

Ormachea, J. and Parker, K. J. (2021). Reverberant shear wave phase gradients for elastography. *Physics in Medicine and Biology*, 66(17):175001.

Palmeri, M., Wang, M., Dahl, J., Frinkley, K., and Nightingale, K. (2008). Quantifying hepatic shear modulus in vivo using acoustic radiation force. *Ultrasound in Medicine & Biology*, 34(4):546–558.

Palmeri, M. L., Deng, Y., Rouze, N. C., and Nightingale, K. R. (2014). Dependence of shear wave spectral content on acoustic radiation force excitation duration and spatial beamwidth. In *2014 IEEE International Ultrasonics Symposium*, pages 1105–1108.

Pinton, G., Dahl, J., and Trahey, G. (2006). Rapid tracking of small displacements with ultrasound. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 53(6):1103–1117.

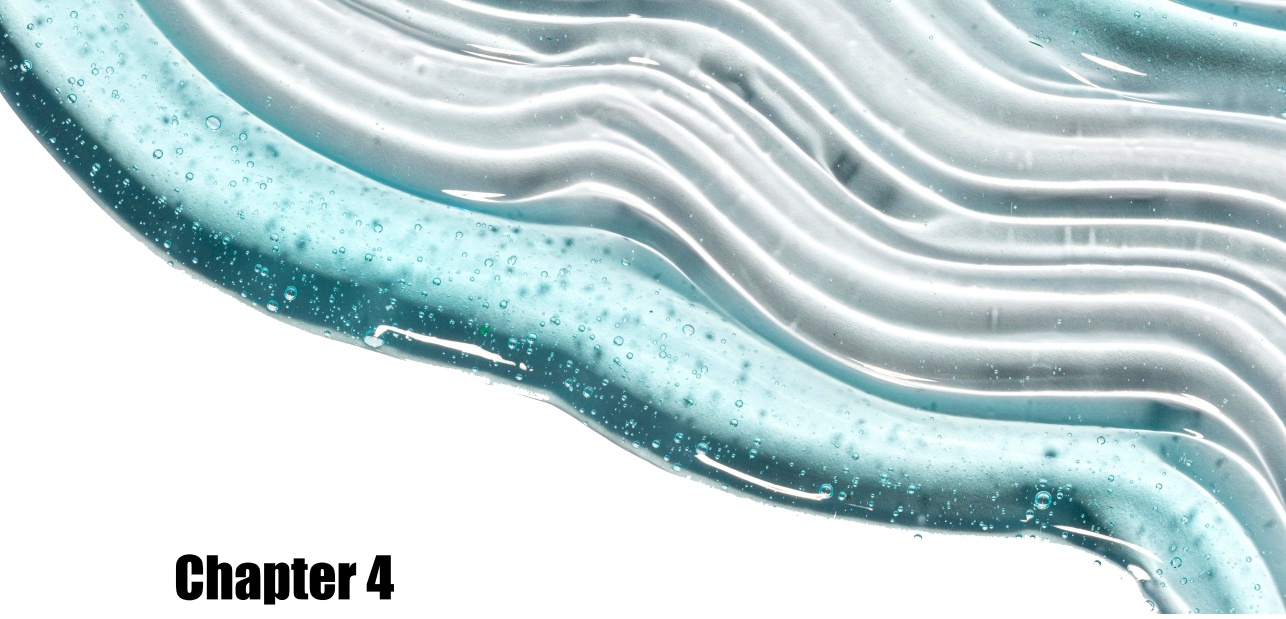
Rizzuto, E., De Luca, R., Musarò, A., and Del Prete, Z. (2021). Measuring and Modelling Nonlinear Elasticity of Ex Vivo Mouse Muscles. *Journal of Healthcare Engineering*, 2021:5579232.

Rouze, N. C., Wang, M. H., Palmeri, M. L., and Nightingale, K. R. (2010). Robust estimation of time-of-flight shear wave speed using a radon sum transformation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 57(12):2662–2670.

Song, P., Manduca, A., Zhao, H., Urban, M. W., Greenleaf, J. F., and Chen, S. (2014). Fast Shear Compounding Using Robust 2-D Shear Wave Speed Calculation and Multi-directional Filtering. *Ultrasound in Medicine & Biology*, 40(6):1343–1355.

Zhao, H., Song, P., Meixner, D. D., Kinnick, R. R., Callstrom, M. R., Sanchez, W., Urban, M. W., Manduca, A., Greenleaf, J. F., and Chen, S. (2014). External Vibration Multi-Directional Ultrasound Shearwave Elastography (EVMUSE): Application in Liver Fibrosis Staging. *IEEE Transactions on Medical Imaging*, 33(11):2140–2148. Conference Name: IEEE Transactions on Medical Imaging.

Zhao, H., Song, P., Urban, M. W., Kinnick, R. R., Yin, M., Greenleaf, J. F., and Chen, S. (2011). Bias Observed in Time-of-Flight Shear Wave Speed Measurements Using Radiation Force of a Focused Ultrasound Beam. *Ultrasound in Medicine & Biology*, 37(11):1884–1892.



Chapter 4

Acoustic Vortex Elastography

Abstract

This chapter introduces a novel ultrasound elastography technique based on acoustic vortices. By focusing acoustic vortices carrying orbital angular pseudo-momentum, a torque is exherted at the focal spot and, therefore, the soft tissue is twisted generating shear rotational waves in a efficient way. Then, by tracking and processing this motion, we image the shear wave velocity distribution of the tissue. In this chapter we validate the method using a particular realization based on acoustic lenses to generate the vortex and an ultrafast ultrasound imaging system to detect the shear wave motion. Vortex-fields generated with 3D-printed lenses are validated by hydrophone mapping measurements, confirming the expected annular pressure distribution and helical phase winding. The observed tissue response under the action of the vortex shows a strong rotational component, providing direct evidence of torque-driven deformation and showing that vortex excitation enhances torsional components relative to conventional focusing. For acoustic beam of same mechanical index, increasing topological charge improves shear-wave motion amplitude and SNR, which translates into more robust phase-velocity reconstructions in a heterogeneous phantom and *ex vivo* kidney tissue. In this way, vortices might improve shear wave elastography.

Outline

4.1 Introduction . . . 103
4.2 Acoustic vortex elastography . . . 105
4.3 Generating acoustic vortices . . . 107
4.3.1 Vortex lens design . . . 107
4.3.2 Acoustic field validation protocol . . . 109
4.3.3 Metrics . . . 111
4.3.4 Results . . . 111
4.4 Orbital angular momentum validation . . . 112
4.4.1 Simulation and experimental protocol . . . 113
4.4.2 Metrics . . . 115
4.4.3 Results . . . 115
4.5 Shear-wave generation efficiency . . . 117
4.5.1 Simulation and experimental protocol . . . 117
4.5.2 Metrics . . . 118
4.5.3 Results . . . 118
4.6 Elasticity imaging using vortex beams . . . 120
4.6.1 Heterogeneous phantom . . . 120
4.6.2 *Ex vivo* kidney . . . 123
4.7 Discussion . . . 127
4.8 Conclusions . . . 129

4.1 Introduction

Shear Wave Elastography (SWE) is contingent upon the effectiveness of ultrasound excitation in inducing a quantifiable shear-wave field within a concise, well-defined time window [Sarvazyan et al., 1998, Nightingale et al., 2003]. In practice, the excitation stage determines whether the subsequent tracking and inversion operate with a high-SNR wavefield or with motion that is barely distinguishable from decorrelation and physiological clutter. In this chapter, the emphasis is therefore placed on the excitation itself: how to reshape ARF actuation so that shear waves are generated more reliably and with a richer spatial radiation pattern, while remaining compatible with conservative acoustic output.

In conventional SWE, ARF is dominated by linear momentum transfer along the beam axis. This typically produces a localized push and a shear-wave motion field that is larger in a limited angular sector around the focal region. When operating at reduced output—for instance to comply with safety constraints or to improve robustness in fragile or highly attenuating scenarios like neonates [Nemes et al., 2024] or lung tissue [Child et al., 1990]—the induced displacements can become marginal, degrading motion tracking and ultimately the precision of stiffness estimates. Importantly, peak output cannot be increased arbitrarily: mechanical index (MI) and thermal index (TI) considerations place explicit constraints on diagnostic ultrasound exposure [Abbott, 1999]. Moreover, transient temperature rises during extended or repeated pulse bursts can become a practical limiting factor [Herman and Harris, 2002]. These constraints motivate excitation strategies that introduce new degrees of freedom beyond the standard axial push.

Acoustic vortex (AV) beams provide such a degree of freedom. In the harmonic regime, using the time convention $e^{-i\omega_0 t}$, the complex pressure field of an idealized acoustic vortex can be written as

$$p(r, \theta, z, t) = P(r, z)\, e^{i\ell\theta} e^{-i\omega_0 t}, \tag{4.1}$$

where $P(r, z)$ is the complex spatial envelope and $\ell \in \mathbb{Z}$ is the topological charge, i.e., the number of 2π phase wraps accumulated over one azimuthal turn [Hefner and Marston, 1999, Thomas and Marchiano, 2003]. Unlike a conventional focused beam, a focused AV beam exhibits azimuthal phase winding, a phase singularity on axis, and an annular focal intensity distribution. These helical-phase and doughnut-like amplitude signatures are the defining

field features of vortex excitation.

This helical structure does not only reshape the focal spot. It also introduces an azimuthal phase gradient and therefore a circulating component of energy flow, which enables orbital angular pseudo-momentum (OAM) transport. In a lossy medium, this angular momentum can be converted into acoustic radiation torque [Anhäuser et al., 2012]. For a single-mode vortex field, the axial OAM density m_z and the beam energy density w satisfy [Volke-Sepúlveda et al., 2008]

$$\frac{m_z}{w} = \frac{\ell}{\omega_0}, \tag{4.2}$$

so that increasing $|\ell|$ increases OAM per unit power, while changing the sign of ℓ reverses the direction of angular momentum transport [Thomas and Marchiano, 2003]. Transfer of angular momentum from acoustic vortices to matter has already been demonstrated in free space and in liquids, establishing torque as a measurable and controllable outcome of vortex excitation [Volke-Sepúlveda et al., 2008, Anhäuser et al., 2012, Baresch et al., 2018]. However, to the best of our knowledge, an experimental demonstration of OAM transfer from an acoustic vortex to soft tissue for elastography has not been reported.

The technological ecosystem around vortices has matured rapidly. Vortex fields can be produced by dedicated transducer geometries [Hefner and Marston, 1999, Jiménez et al., 2021], by passive diffractive devices such as spiral gratings [Jiménez et al., 2016, Jiménez et al., 2018], or by holographic and phase-engineered acoustic elements that shape the field with high flexibility [Marzo et al., 2015, Jiménez-Gambín et al., 2019, Jiménez-Gambín et al., 2020]. In parallel, vortex-based manipulation has expanded into ultrasonic trapping and control of particles and microbubbles, illustrating that OAM can be exploited under practical acoustic intensities and imaging-compatible conditions [Baresch et al., 2016, Marzo et al., 2018, Lo et al., 2021]. These developments suggest that vortex beams can be engineered with sufficient stability and reproducibility to be considered as a realistic actuation modality in elastography.

From an elastography perspective, the motivation is straightforward: twist rather than push. A vortex excitation distributes ARF over an annular region and introduces a rotational component of motion associated with torque. This geometry is expected to (i) improve displacement detectability at a given MI by activating rotational shear components that are weak or absent in conventional pushes, and (ii) enrich the angular content of the launched shear-wave field,

which is beneficial whenever reconstruction quality is limited by a narrow radiation pattern, interference, or heterogeneous propagation. Moreover, in the context of a conventional focused beam, it is not possible to pull the tissue; however, by employing a vortex, and modulating the topological charge accordingly, it is possible to reverse the direction of the torque. Building on these ideas, a vortex-based elastography concept was previously introduced as a route with controllable waveform design [Jiménez et al., 2020], and it has been formalized in an intellectual-property framework [Jiménez et al., 2021].

This chapter presents the first experimental realization of ultrasound elastography driven by focused acoustic vortices. The following sections describe the design of passive vortex lenses compatible with medical ultrasound hardware, the acoustic characterization of the resulting vortex fields, and their validation for elastography in tissue-mimicking and *ex vivo* configurations. Because vortex elastography relies on precise control of phase and pressure distributions, acoustic field mapping and reporting follow established best practices for hydrophone measurements and exposimetry [Harris et al., 2023, Canney et al., 2008]. The overarching goal is to establish angular momentum transfer as a practical and safe actuation mechanism that complements conventional ARF excitation, opening a path toward more efficient shear-wave imaging when acoustic output is constrained [Herman and Harris, 2002].

4.2 Acoustic vortex elastography

Acoustic vortex elastography was implemented using a dedicated dual-purpose device that combines vortex-based mechanical actuation with ultrafast ultrasound motion tracking. The system was built on a programmable research ultrasound scanner (Vantage 256, Verasonics Inc., Kirkland, WA, USA) and integrates three main components: a focused annular transducer for actuation, a removable 3D-printed vortex lens for wavefront shaping, and a concentric imaging probe for motion detection. Its operating principle is illustrated in Fig. 4.1.

The procedure is sequential. First, the annular source emits a focused ultrasound burst shaped by the vortex lens, generating an acoustic vortex beam at the target region. As a result, the tissue is not only pushed but also twisted, producing a localized torsional mechanical response. This excitation can launch shear waves with a marked rotational component, which constitutes

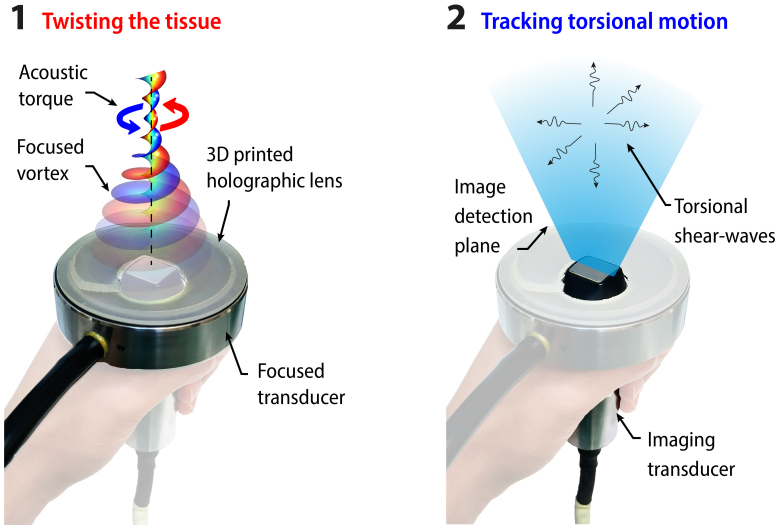


Figure 4.1: Vortex elastography device concept. A focused annular transducer fitted with a 3D-printed vortex lens generates a vortex beam capable of imparting acoustic torque to the medium in order to twist the tissue. A concentric imaging probe acquires ultrafast data to track the resulting torsional shear-wave motion within a defined detection plane.

the basis of the proposed elastographic approach.

After the excitation burst, the same device switches to imaging mode. A concentric phased-array probe acquires ultrafast ultrasound data in a defined detection plane intersecting the actuation region. These acquisitions capture the shear-wave motion induced by the vortex beam. As illustrated in Figure 4.1, the elastographic measurement process comprises two consecutive stages. The first method involves the induction of vortex-induced twisting of the tissue, and the second method involves ultrafast tracking of the resulting motion field.

Vortex actuation was produced with a 1.1 MHz four-element ring-shaped focused source (H-101, Sonic Concepts, WA, USA) with a nominal focal distance of $F = 64$ mm and outer/inner apertures of 64/37 mm. To generate the different vortex states, removable soft silicone lenses (Elite Double 8, Zhermack SpA, Badia Polesine, Italy; longitudinal wave speed $c_L = 1000$ m/s; attenuation coefficient $\alpha = 0.3$ dB cm^{-1} MHz^{-1}) were mounted in front of the source. In addition to encoding the vortex phase profile, the lenses shifted the effective focus from 64 to 40 mm, bringing the region of maximum acoustic stress closer to the imaging plane. Five lenses with topological charges

$\ell \in \{0, 1, 2, 3, 4\}$ were fabricated, all designed for the same target focal distance.

Motion tracking was performed with a concentric 5.2 MHz 64-element phased-array probe (IP-105, Sonic Concepts, WA, USA). Ultrafast imaging was carried out using plane-wave transmissions with three-angle coherent compounding at $-15°$, $0°$, and $+15°$. Consecutive compounded frames were acquired every 100 μs, corresponding to an effective frame rate of 3.3 kHz. This configuration provided high temporal resolution for resolving the rapidly evolving displacement field generated after vortex excitation.

Overall, the device was designed to spatially align actuation and detection within a compact coaxial geometry. This arrangement ensures that the region where acoustic torque is delivered is directly accessible to ultrafast imaging, enabling the experimental observation of vortex-induced torsional motion and its associated shear-wave dynamics.

4.3 Generating acoustic vortices

4.3.1 Vortex lens design

In this work, vortex beams were generated using custom 3D-printed refractive lenses mounted coaxially on a spherically curved annular transducer of radius of curvature ROC. The lens is parameterized by the integer charge ℓ and designed by enforcing local phase matching on the spherical aperture. The coordinate R denotes the radial distance from the center of curvature of the transducer, so that $R = \mathrm{ROC}$ corresponds to the bare spherical aperture. The parameter d_0 defines the prescribed minimum lens thickness, measured along the local radial direction.

At the excitation frequency ω_0, the wavenumbers in the surrounding medium and in the lens are

$$k_0 = \frac{\omega_0}{c_0}, \qquad k_L = \frac{\omega_0}{c_L}, \tag{4.3}$$

with sound speeds c_0 and c_L. The aperture is discretized in spherical coordinates (θ, ϕ), and candidate surface points are explored along the local radial coordinate R within the admissible interval

$$R \in [\mathrm{ROC} - d_0 - h_{\max},\ \mathrm{ROC} - d_0], \tag{4.4}$$

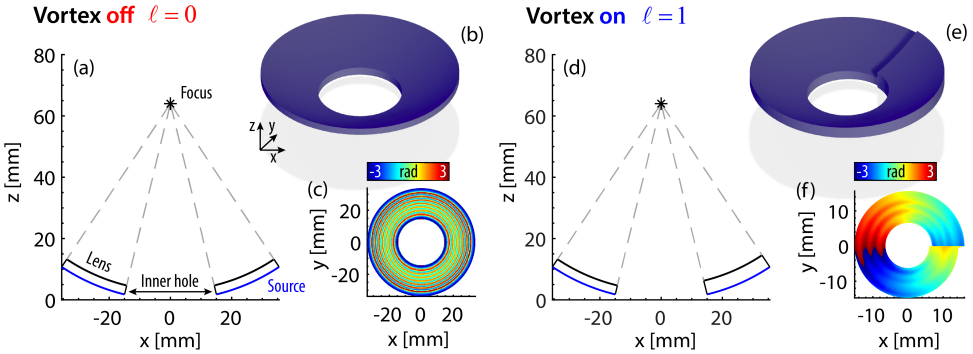


Figure 4.2: Representative design outputs for $\ell = 0$ and $\ell = 1$ showing a lateral cross-section of the lens profile, a 3D rendering of the fabricated geometry, and the phase map at the lens exit surface. For $\ell = 0$, the phase is axisymmetric and encodes only the focusing term. For $\ell = 1$, an additional azimuthal phase ramp is imposed, producing a helical phase profile over the annular aperture.

where $h_{\max}$ is the maximum allowed lens relief. Thus, $R = \mathrm{ROC} - d_0$ represents the thinnest admissible surface, whereas smaller values of R correspond to locally thicker lens regions.

A candidate point is written in Cartesian form as

$$\mathbf{r}(\theta, \phi; R) = \begin{bmatrix} R\cos\phi\cos\theta \\ R\cos\phi\sin\theta \\ R\sin\phi \end{bmatrix}, \tag{4.5}$$

and the target focus is $\mathbf{r}_F = (0, 0, F)^\top$, where F is the prescribed focal distance. The desired phase in the surrounding medium corresponds to a converging spherical wave,

$$\phi_{\mathrm{src}}(\theta, \phi; R) = -k_0 \left\| \mathbf{r}(\theta, \phi; R) - \mathbf{r}_F \right\|. \tag{4.6}$$

The lens-imparted phase is modeled as propagation through the lens plus the helical term,

$$\phi_{\mathrm{lens}}(\theta, \phi; R) = -k_L \, L(\theta, \phi; R) \; - \; \ell\,\theta, \tag{4.7}$$

where $L(\theta, \phi; R)$ is the path length in the lens.

Both phases are referenced to a boundary point (θ_0, ϕ_0),

$$\begin{aligned} \Delta\phi(\theta, \phi; R) = {} & \phi_{\mathrm{lens}}(\theta, \phi; R) - \phi_{\mathrm{src}}(\theta, \phi; R) \\ & - \Big[\phi_{\mathrm{lens}}(\theta_0, \phi_0; R_0) - \phi_{\mathrm{src}}(\theta_0, \phi_0; R_0) \Big]. \end{aligned} \tag{4.8}$$

A continuous branch is enforced across the aperture. Among the admissible solutions, the largest radial solution $R_h(\theta, \phi)$ is selected at each aperture point, where R_h denotes the radial coordinate of the final lens exit surface. This choice minimizes the local lens thickness while satisfying

$$t(\theta, \phi) = \mathrm{ROC} - R_h(\theta, \phi) \geq d_0. \tag{4.9}$$

Figure 4.2 illustrates representative outputs for $\ell = 0$ and $\ell = 1$. Finally, the surface mesh was exported as STL for fabrication.

4.3.2 Acoustic field validation protocol

In order to validate the designed lenses, hydrophone scans were conducted in degassed deionized water at 22–23 °C using a fiber-optic reflectance hydrophone (HFO-690, ONDA; 100 μm tip; 3 kHz–150 MHz bandwidth) on a three-axis stage (M-403/C-863, PI Micos; 5 μm resolution), following exposimetry best practices [Canney et al., 2008, Harris et al., 2023]. Signals were digitized at 125 MS/s (PicoScope 5444D, 12-bit). The source was driven at 1.1 MHz with 50-cycle bursts. For the transversal and axial planes measurements 30 V were used. In addition, voltages ranging from 1.6 to 50 V were utilized to characterize the mechanical index of the designed lenses. For each topological charge, pressure was mapped in a transverse plane at $z = 40$ mm and in an axial plane at $y = 0$ over ± 6 mm laterally. At each point, 10 waveforms were averaged.

Numerical acoustic fields were computed with the hybrid Rayleigh–angular-spectrum solver presented in Chap. 2, where the vortex lens was modeled as an ideal spatial phase delay applied over the spherical aperture. The transducer field was first synthesized with the Rayleigh formulation and then propagated volumetrically using the angular spectrum method. For each topological charge, the simulated pressure field was evaluated in the same transverse and axial planes used experimentally. The complex pressure at the excitation frequency was extracted to reconstruct magnitude and phase maps, which were directly compared with the hydrophone measurements to verify the expected annular focal distribution and azimuthal phase winding.

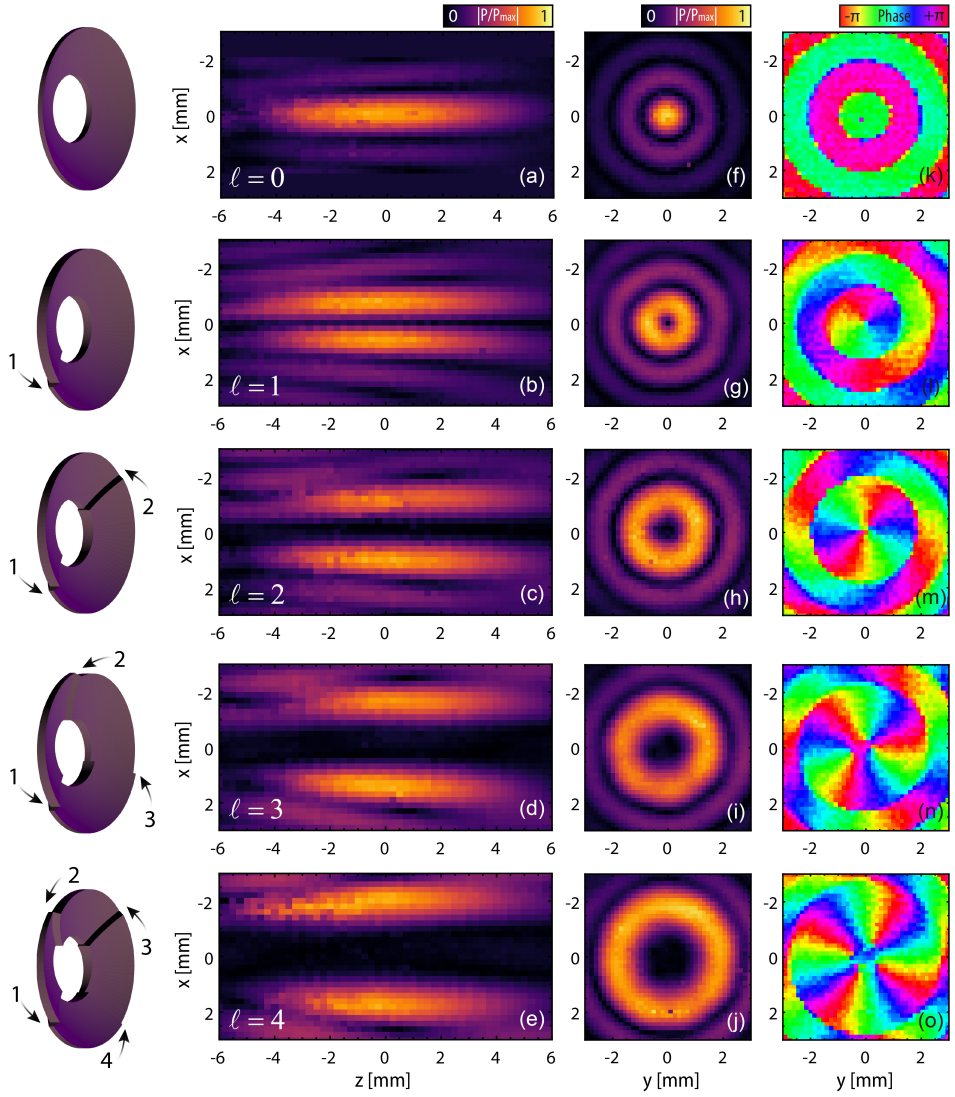


Figure 4.3: Hydrophone-based characterization of the acoustic field generated by the annular source combined with vortex lenses. The figure reports axial pressure magnitude maps in the x–z plane, transverse pressure magnitude maps in the x–y plane at $z = 40$ mm, and corresponding transverse phase maps showing the azimuthal winding that emerges for $\ell \geq 1$.

4.3.3 Metrics

Mechanical index (MI) was estimated for each lens configuration to support exposure-matched comparisons [Abbott, 1999,Herman and Harris, 2002]. Peak rarefactional pressure at the focus was obtained over an input-voltage sweep. Soft-tissue derating used $\alpha = 0.3$ dB cm^{-1} MHz^{-1} and $z_F = 40$ mm:

$$p_{r,0.3}(\ell, V_{\text{in}}) = p_r(\ell, V_{\text{in}})\, 10^{-\alpha\, f_0\, z_F/20}. \tag{4.10}$$

MI was computed as

$$\text{MI}(\ell, V_{\text{in}}) = \frac{p_{r,0.3}(\ell, V_{\text{in}})}{\sqrt{f_0}}, \tag{4.11}$$

with $p_{r,0.3}$ in MPa and f_0 in MHz [Abbott, 1999].

The lateral beamwidth was reported as the -6 dB full width at half maximum (FWHM). The degree of agreement between the simulated and measured lateral profiles was quantified using a goodness-of-fit parameter calculated using the normalized mean squared error (NRMSE). A perfect fit between the simulated and measured profiles is indicated by a 0 value. In order to validate the helical phase structure, the unwrapped acoustic phase was sampled along a circular trajectory of radius 1 millimeters and compared against the expected linear phase ramp as a function of azimuth.

4.3.4 Results

Figure 4.3 summarizes the measured focal structure of the focused vortex beams. For $\ell = 0$, the field exhibited a conventional focused pattern with a single dominant lobe in the axial map and an approximately axisymmetric focal spot in the transverse plane. When the vortex lenses were introduced, the focal intensity reorganized into a toroidal distribution with a pressure null on axis, and the transverse phase maps displayed the characteristic azimuthal winding. The number of phase wraps increased with ℓ, consistent with the intended helical wavefront encoded by the lens design.

As illustrated in Fig. 4.4(a), the measured lateral pressure profiles at a distance of 40 millimeters from the origin exhibited a high degree of correlation with the numerical predictions for $\ell = 0$, 1, and 3. The NRMSE between numerical and experimental data were determined to be 0.03, 0.04, and 0.02, respectively. The experimentally measured FWHM increased with topological

charge, from 1.2 mm for $\ell = 0$ to 2.4 mm for $\ell = 1$, 3.5 mm for $\ell = 2$, 4.5 mm for $\ell = 3$, and 5.7 mm for $\ell = 4$. This trend reflects the redistribution of acoustic energy into an annular focus as the helical phase gradient increases.

Phase measurements are presented in Fig. 4.4(b). These measurements provide further confirmation of the expected azimuthal structure. The unwrapped phase evolves linearly with angle for $\ell \geq 1$.

Figure 4.5 summarizes the estimated derated mechanical index as a function of the driving voltage for evaluated topological charges. As expected, MI increased with input voltage for all configurations. For a fixed voltage, MI decreased monotonically with increasing ℓ, reflecting the redistribution of acoustic energy into a wider annular focal structure for higher-charge vortices, which reduces the peak rarefactional pressure that governs MI. For example, at $V_{\text{in}} = 20$ V, the estimated MI ranged from 8.2 for $\ell = 0$ down to 3.6 for $\ell = 4$, corresponding to a reduction of approximately 56% between the non-vortex and highest-charge conditions.

4.4 Orbital angular momentum validation

This section tests whether focused vortices generate a measurable torsional response in a bulk soft solid, consistent with OAM transfer. Torsional motion was evaluated in a homogeneous gelatin phantom and benchmarked against

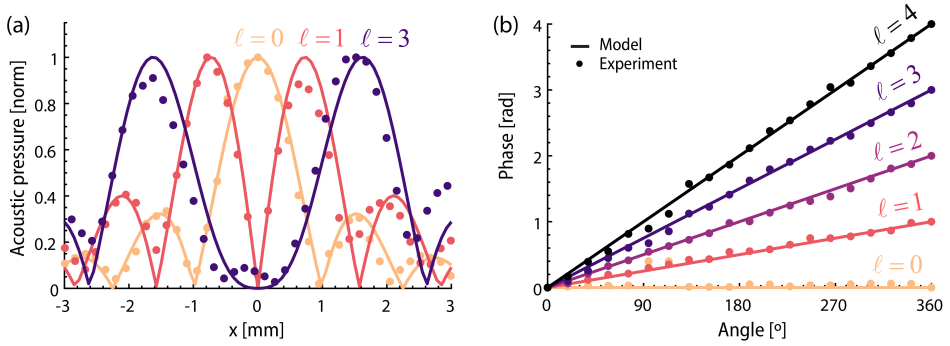


Figure 4.4: Quantitative validation of the vortex fields against numerical predictions. (a) Comparison normalized lateral pressure profiles at $z = 40$ mm for selected topological charges and (b) unwrapped phase as a function of azimuth along a circle of radius 5 mm at $z = 40$ mm.

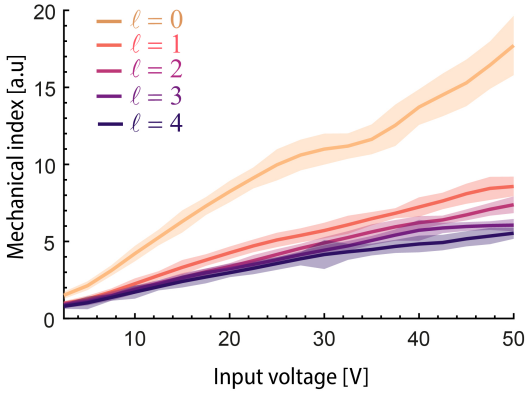


Figure 4.5: Estimated mechanical index as a function of input driving voltage for vortex lenses with topological charge $\ell = 0$ to $\ell = 4$. Increasing ℓ systematically reduced MI at a given voltage, consistent with a lower peak rarefactional pressure due to the annular redistribution of the focal energy.

shear-wave simulations.

4.4.1 Simulation and experimental protocol

Experiments were conducted in a custom homogeneous gelatin phantom with Young's modulus of 45 kPa [Cook et al., 2011]. The push transmission is characterized by a burst tone at 1.1 MHz of 100 μs of duration. The reference case transmission $\ell = 0$ and the vortex one $\ell = 1$ were evaluated at the maximum driving voltage (50 V), corresponding to mechanical indices of 15 and 8, respectively.

To reconstruct the in-plane motion within the transverse plane intersecting the focal region (z_f), two orthogonal acquisitions were performed with a 7.6 MHz 128-element linear-array probe (L11-5v, Verasonics, Kirkland, WA), as illustrated in Fig. 4.6. The probe was repositioned between acquisitions so that its axial tracking direction was aligned first with the lateral axis to measure $v_x(x, y, z_f, t)$ and then with the elevational axis to measure $v_y(x, y, z_f, t)$ under matched actuation conditions. Ultrafast imaging was performed using three-angle compounded plane-wave transmissions at steering angles $-3°$, $0°$, and $+3°$, yielding an effective compounded frame rate of 3.33 kHz. The two datasets

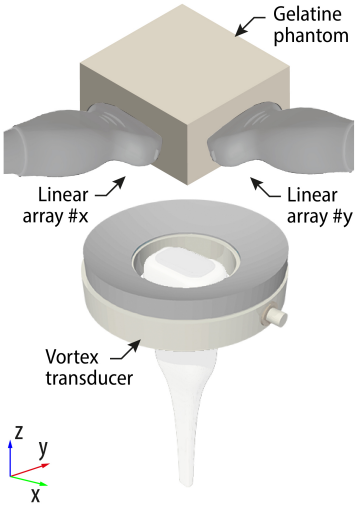


Figure 4.6: Experimental configurations for the torsional-response experiment validation. This experiment uses two orthogonal linear probes immersed in a water tank to capture the particle velocity for the x and y component to compound the v_θ field.

were then registered in space and time to form the in-plane velocity field

$$\mathbf{v}(x, y, z_f, t) = \begin{bmatrix} v_x(x, y, z_f, t) \\ v_y(x, y, z_f, t) \end{bmatrix}. \tag{4.12}$$

Displacements were estimated from IQ data using a two-dimensional autocorrelation method [Loupas et al., 1995], and the first 0.3 ms after the push were excluded from the analysis to avoid contamination from the excitation burst.

Numerical simulations were configured to mirror this acquisition and provide a reference. The steady-state complex pressure field for each topological charge was computed with the hybrid Rayleigh–angular-spectrum solver (Sec. 2.2.1), assuming linear propagation in water and incorporating the vortex lens as an ideal spatial phase delay. The corresponding 3D acoustic radiation force distribution was obtained from the harmonic pressure field using the second-order formulation (Sec. 2.2.2), with a temporal envelope matched to the 100 μs push. Elastic-wave propagation was simulated in k-Wave (Sec. 2.2.3) using an isotropic elastic medium with $\rho = 1000\ \mathrm{kg\,m^{-3}}$, shear modulus $\mu = 45$ kPa, and attenuation $0.5\ \mathrm{dB\,cm^{-1}\,MHz^{-1}}$, on a 3D Cartesian grid with $dx = dy = dz = 0.25$ mm over a $15 \times 15 \times 30$ mm domain. The simulated vectorial particle velocity field, v_x and v_y, was sampled in the same

transverse observation plane to match the two orthogonal probe orientations that were used in the experiment. This enabled a consistent computation of the rotational component and derived metrics.

4.4.2 Metrics

The rotational displacement component was computed by projecting $\mathbf{u}(x, y, t)$ onto the local azimuthal direction around the beam axis. Defining the polar angle

$$tan[\Theta(x, y)] = \frac{y - y_0}{x - x_0} \tag{4.13}$$

where (x_0, y_0) denotes the beam axis location in the measurement plane, the rotational displacement was defined as

$$\bar{v}_\theta(x, y, t) = v_x(x, y, t)\ \sin\Theta(x, y)\ +\ v_y(x, y, t)\ \cos\Theta(x, y). \tag{4.14}$$

Prior to projection, v_x and v_y were temporally band-pass filtered over the shear-wave band identified by short-time spectral analysis, with a typical passband of 100–1000 Hz, and spatially low-pass filtered to suppress speckle-scale fluctuations using a lateral cutoff of approximately $0.5\,\lambda$.

To summarize torsional response around the vortex ring, a fixed annular region of interest was used with inner radius $r_1 = 2.5$ mm and outer radius $r_2 = 5.5$ mm, centered at (x_0, y_0). The root-mean-square rotational displacement over a window $\Delta t = 5$ ms starting 0.3 ms after the push was computed as

$$v_{\theta,\mathrm{RMS}} = \frac{1}{N} \sum_{(x,y)\in \mathrm{ROI}} \sqrt{\frac{1}{T} \sum_{t\in\Delta t} v_\theta^2(x, y, t)}, \tag{4.15}$$

where N is the number of pixels inside the annulus and T is the number of temporal samples inside Δt. In addition, representative time traces $v_\theta(t)$ were extracted at fixed points on the annulus to compare waveform structure between excitations.

4.4.3 Results

Figure 4.7(a-b) summarizes the in-plane displacement components measured under conventional focusing. The corresponding v_x and v_y maps exhibit

predominantly radial patterns, and the projected rotational component remains close to baseline, as shown in Figure 4.7(c). By contrast, vortex excitation generates a clear rotational signature. Figure 4.7(e-f) shows the measured in-plane components for $\ell = 1$, and the reconstructed rotational field concentrates around the vortex ring, as shown in Figure 4.7(g).

Quantitatively, the RMS rotational displacement increased by approximately a factor of 15 for vortex excitation compared with conventional focusing. The mean value of $v_{\theta,\mathrm{RMS}}$ was 0.34 mm for $\ell = 1$ and 0.02 mm for $\ell = 0$, computed over the same annular ROI and temporal window. This difference occurred despite the $\ell = 1$ condition being delivered at a lower mechanical index, with MI of 8 for $\ell = 1$ versus MI of 14 for $\ell = 0$.

Time traces at a representative location on the ring further highlighted the contrast between conditions. The $\ell = 1$ excitation, Figure 4.7(h), produced a clear oscillatory rotational response, while the $\ell = 0$, Figure 4.7(d), trace remained close to baseline. The simulated $v_\theta(t)$ closely matched the experimental waveform at the same location, with a correlation coefficient of 0.92 and a mean absolute error of 0.12. These results provide direct experimental evidence that focused acoustic vortices can deliver a measurable torsional component in soft tissue surrogates, consistent with prior theoretical and experimental

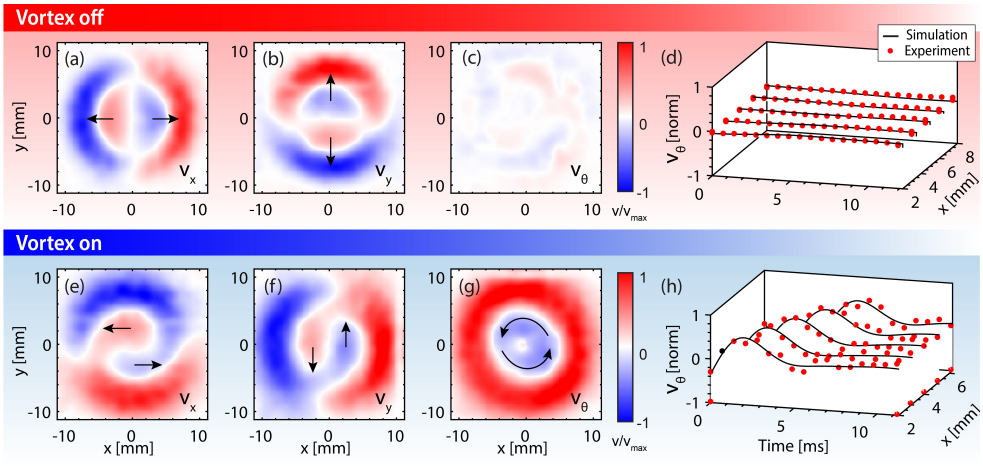


Figure 4.7: Direct evidence of orbital angular momentum transfer in a homogeneous gelatin phantom. In-plane displacement components v_x and v_y were acquired in a transverse plane and combined to compute the rotational component v_θ for conventional focusing and vortex excitation.

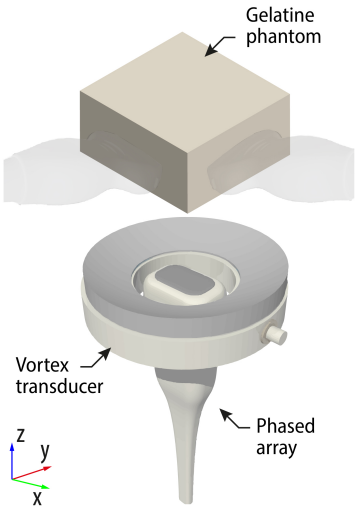


Figure 4.8: Experimental configurations for the shear-wave efficiency experiment validation. This experiment uses a concentric probe to measure axial particle velocity along the focal depth under controlled MI and topological charge.

studies of acoustic torque mediated by orbital angular momentum [Guo et al., 2022].

4.5 Shear-wave generation efficiency

This experiment quantifies how efficiently vortex beams generate detectable shear waves depending on the topological charge under controlled acoustic exposure.

4.5.1 Simulation and experimental protocol

The shear-wave generation efficiency was evaluated in the same 45-kPa gelatin phantom for topological charges $\ell = 0$ to $\ell = 4$, using the vortex transmission parameters described in Sec. 4.2. Shear-wave motion was tracked with the concentric IP-105 imaging probe, as illustrated in Fig. 4.8. Frames were acquired using ultrafast plane-wave imaging with by coherently compounding three steered transmissions at $-15°$, $0°$, and $+15°$ per frame resulting in a

PRF of 3.3 kHz. Acoustic exposure was controlled by sweeping the driving voltage over the available range to vary the mechanical index.

Numerical simulations were used to benchmark the measured trends using the same acoustic field computation, force deposition, and elastic method previously described for the OAM experiment (Sec. 4.4.1).

4.5.2 Metrics

Axial particle velocity $v_z(x, z, t)$ was reduced to lateral–time maps $v_z(x, t)$ by axial averaging over 1 mm centered at $z_F = 40$ mm. To probe axial repeatability, the window was shifted in 0.2 mm steps over $z = 38.0$ to 41.8 mm (20 realizations). Peak amplitude was

$$v_{z,\mathrm{pk}} = \max_{x,t} |v_z(x, t)|. \tag{4.16}$$

Signal-to-noise ratio was computed from the axial velocity maps by comparing a signal window around the propagating wavefront to a noise window away from the wavefront, using an RMS-based definition in decibels:

$$\mathrm{SNR} = 20 \log_{10} \left(\frac{\mathrm{RMS}\{v_z\}_{\mathrm{signal}}}{\mathrm{RMS}\{v_z\}_{\mathrm{noise}}} \right). \tag{4.17}$$

4.5.3 Results

The vortex-induced shear-wave response at a moderated mechanical index is first illustrated through spatiotemporal particle-velocity maps. Figure 4.9(a) shows that case $\ell = 0$ produced a weak and poorly contrasted propagating feature within the same acquisition window. As the topological charge increased, the propagating wavefront became progressively stronger and more coherent, as shown in Fig. 4.9(b-e). This trend is also represented Fig. 4.9(f), where a representative waveform extracted at $x = 10$ mm and $z = 40$ mm confirms that the peak particle velocity increased monotonically with ℓ. The measured peak values at this spatial location were 0.42, 0.86, 1.15, 1.69, and 2.05 mm/s for $\ell = 0$ to $\ell = 4$, respectively.

To quantify this trend SNR metric was calculated from spatio-temporal maps represented in Figs. 4.9(a-e). The dashed lines indicate the signal region of interest for SNR calculation, the rest of the map is the noise area. The

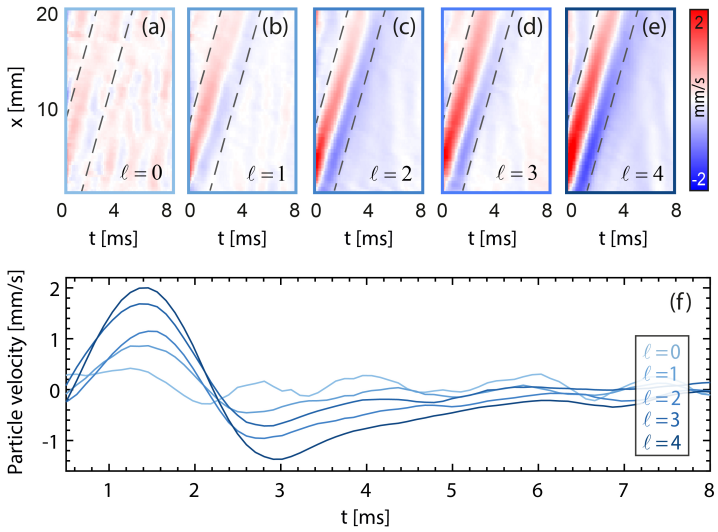


Figure 4.9: Spatiotemporal signatures of shear-wave generation across topological charge. (a-e) Spatially averaged lateral–time particle-velocity maps $v_z(x,t)$ for $\ell = 0$ to $\ell = 4$. The dashed lines indicate the signal ROI used for SNR calculation. (f) Representative waveforms extracted at $x = 10$ mm and $z = 40$ mm.

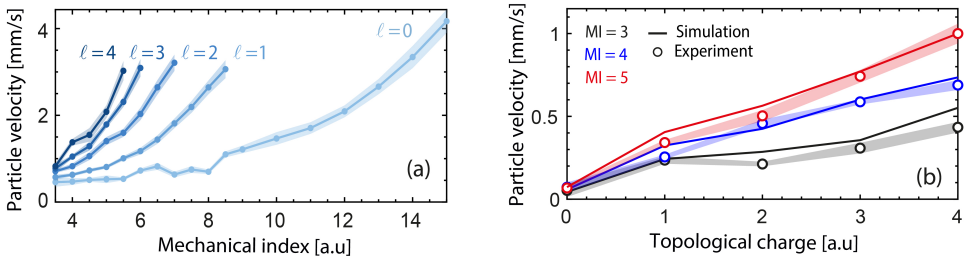


Figure 4.10: (a) Peak axial particle velocity $v_{z,\mathrm{pk}}$ as a function of mechanical index for $\ell = 0$ to $\ell = 4$. (b) Simulated and experimental peak velocities as a function of topological charge at matched MI values of 3, 4, and 5.

SNR of the axial velocity field also increased with ℓ: 1.3, 8.5, 11.4, 12.7, and 13.6 dB for $\ell = 0$ to $\ell = 4$, respectively.

Peak axial particle velocity was then quantified over a mechanical-index sweep for each topological charge. Figure 4.10(a) reports $v_{z,\mathrm{pk}}$ as a function of MI for $\ell = 0$ to 4. For conventional focusing, $v_{z,\mathrm{pk}}$ increased with MI from 0.53 ± 0.17 (MI = 5) to 6.09 ± 0.07 mm/s (MI = 16). Vortex excitation

yielded substantially larger amplitudes at moderate MI. At MI = 5, the $\ell = 4$ condition reached 2.08 ± 0.12 mm/s, nearly four times the $\ell = 0$ value at the same MI. The maximum MI reachable for each lens decreased with increasing ℓ, due to the redistribution of acoustic energy and the system voltage limit. In this configuration, $\ell = 4$ reached MI = 5.5 at the maximum 50 V drive.

To enable a direct comparison between experiment and simulation at matched exposure, peak velocities were extracted for MI values of 3, 4, and 5 and plotted as a function of topological charge. Figure 4.10(b) shows that the peak particle velocity increased monotonically with ℓ for all three MI levels, and that the experimental trend closely followed the simulated prediction. This agreement supports the interpretation that the observed efficiency gains originate from the vortex-driven redistribution of force deposition in the focal region.

4.6 Elasticity imaging using vortex beams

This section evaluates whether vortex-based actuation can improve phase-velocity elastography under constrained acoustic output. The assessment of performance is conducted in a controlled heterogeneous phantom and in an *ex vivo* rabbit kidney.

4.6.1 Heterogeneous phantom

Experiment protocol

Measurements were performed in a gelatin phantom consisting of a stiff cylindrical inclusion embedded in a softer background. The inclusion diameter was 10 mm and its center was located at 25 mm depth. The inclusion was prepared following the tissue-mimicking phantom recipe reported in [Cook et al., 2011], which provides an expected Young's modulus of approximately 40 kPa for this composition. Figure 4.11(p) illustrates a schematic of the experiment.

The vortex beam and imaging transmission parameters used were the same as the ones described in Sec. 4.2. The acoustic output was varied by sweeping the driving voltage and reporting exposure in terms of MI, with acquisitions

performed over the MI range 5–8 whenever available for each ℓ. A conventional high-output reference was additionally acquired using $\ell = 0$ at MI = 12.

Phase-velocity maps $c_{\mathrm{ph}}(x, z)$ were reconstructed with the phase gradient elastography method described in Chap. 3 [González-Mateo et al., 2025]. In brief, the motion field was transformed to the frequency domain, filtered in the (k_x, k_z) domain to enforce physically plausible phase-velocity ranges and suppress reflected components, and mapped back to the space–frequency domain to estimate the local wavevector from phase gradients. Final maps were obtained by averaging phase-velocity estimates over a frequency window centered at the spectral peak of the measured shear-wave motion.

Metrics

Quantitative assessment relied on two delineated regions of interest. Inclusion and background ROIs were defined from B-mode anatomy. In Figure 4.11, these ROIs are indicated by the dashed and dash-dotted circles, respectively. For each ROI, the mean phase velocity μ and its standard deviation σ were computed. Inclusion–background separation was summarized through a contrast-to-noise ratio reported in dB as

$$\mathrm{CNR} = 20 \log_{10} \left(\frac{\left| \mu_{\mathrm{inc}} - \mu_{\mathrm{bg}} \right|}{\sigma_{\mathrm{bg}}} \right). \tag{4.18}$$

To capture variability, repeated acquisitions ($n = 3$) were performed for each configuration.

Results

Representative phase-velocity maps are shown in Figure 4.11. For $\ell = 0$, increasing MI from 5 to 12 progressively enhanced inclusion visibility and reduced spatial fluctuations, as seen by comparing Figs. 4.11(a-e). When vortex lenses were used, inclusion delineation remained observable at lower MI, with examples shown for $\ell = 1$ in Figs. 4.11(f-i). Some ℓ–MI combinations were not acquired due to the limited achievable output for higher ℓ at the same driving constraints.

At MI = 5, the inclusion ROI statistics clearly depended on the topological charge, as shown in Figure 4.12(a). Conventional focusing at low MI underes-

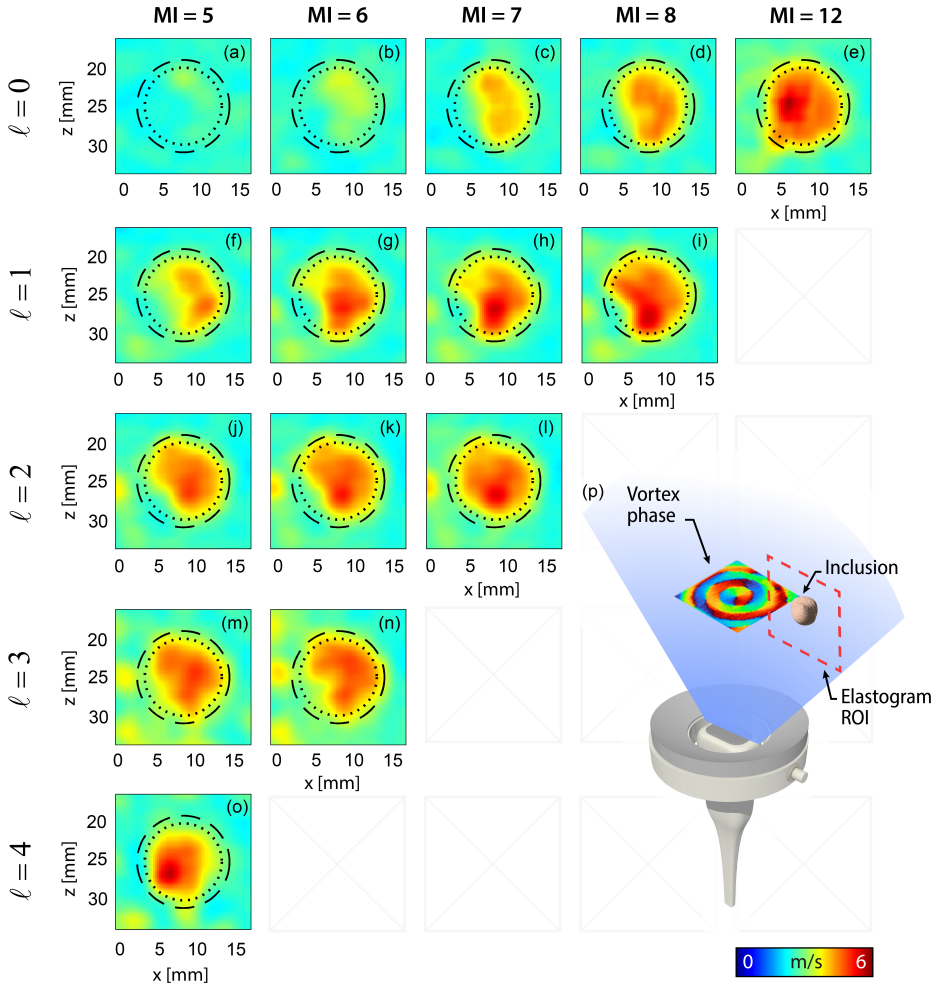


Figure 4.11: Phase-velocity imaging in a heterogeneous phantom with a stiff inclusion embedded in a softer background. (a)–(e) $\ell = 0$ at MI = 5, 6, 7, 8, and 12 reference. (f)–(i) $\ell = 1$ at MI = 5–8. (j)–(l) $\ell = 2$ at MI = 5–7. (m)–(n) $\ell = 3$ at MI = 5–6. (o) $\ell = 4$ at MI = 5. The dashed circles indicate the inclusion margin, and the dash-dotted circles indicate the background ROI used for statistical analysis.

timated the inclusion stiffness, whereas vortex actuation recovered inclusion phase-velocity values comparable to the high-MI reference. In particular, the $\ell = 0$ reference acquired at MI = 12 yielded a mean inclusion phase velocity of 3.67 ± 0.22 m/s, and vortex configurations at MI = 5 produced closely matching central values. By contrast, background estimates remained

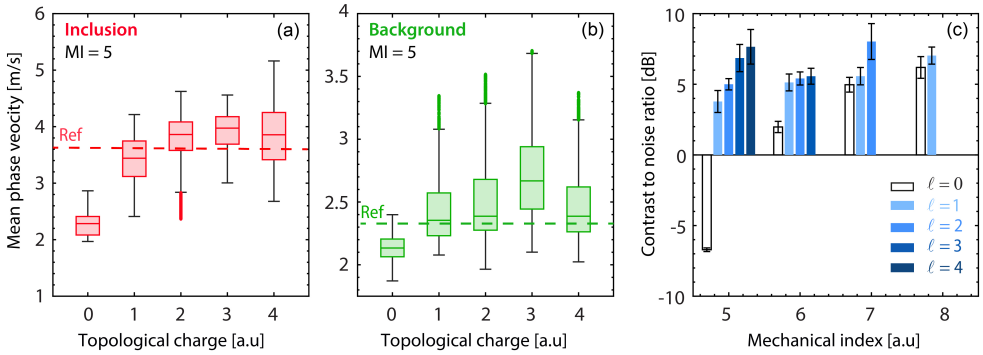


Figure 4.12: Statistical summary of heterogeneous phantom reconstructions. (a) Boxplots of the mean phase velocity within the inclusion ROI at MI = 5 for all ℓ, with the MI = 12 reference for $\ell = 0$ shown as a dashed red line. (b) Corresponding background ROI values under the same conditions, with the reference background mean shown as a dashed green line. (c) Contrast-to-noise ratio as a function of MI for each ℓ over MI = 5–8.

relatively stable across configurations and close to the high-MI reference level, Figure 4.12(b).

CNR trends are summarized in Figure 4.12(c). At low-MI, conventional focusing provided limited inclusion–background separability, whereas vortex actuation produced a higher contrast at the same MI. The effect was strongest at MI = 5, where CNR exhibited an increase from −6.87 dB for $\ell = 0$ to 7.65 dB for $\ell = 4$. The minus value of CNR for the $\ell = 0$ case indicates the absence of inclusion detection in the phase velocity map, as illustrated in Fig. 4.11(a). As the MI increased, the topological charge yielding the highest CNR shifted toward lower values of ℓ. This shift is attributable to the limitation of driving output voltages in accordance with ℓ.

4.6.2 *Ex vivo* kidney

Experiment protocol

Ex vivo experiments were conducted on rabbit kidneys Figure 4.13(a). Kidneys were obtained from a local supplier after animal sacrifice for food purposes; no animals were sacrificed for this study. Samples were stored in degassed water, and embedded in a homogeneous gelatin phantom ($c_s \approx 1.5$ m/s at room temperature) to provide stable coupling [Cook et al., 2011]. The organ was

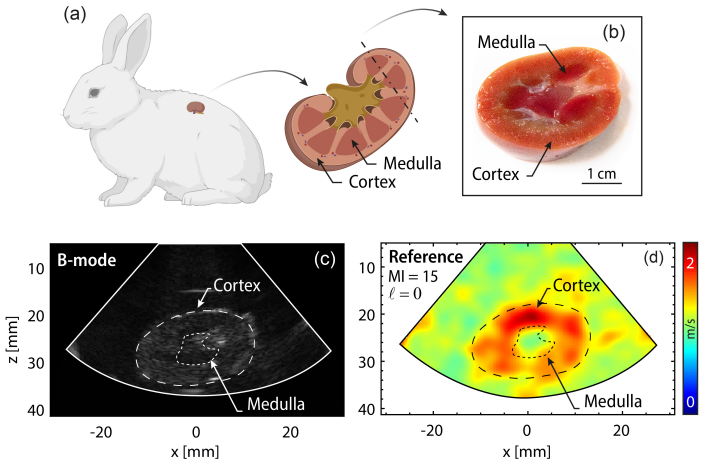


Figure 4.13: *Ex vivo* rabbit kidney experiment and reference elastogram. (a) Schematic representation of rabbit kidney anatomy highlighting cortex and medulla. (b) Photograph of the longitudinal kidney section used for imaging. (c) B-mode image with cortex and medulla ROIs defined from anatomical landmarks. (d) Reference phase-velocity map acquired with $\ell = 0$ at MI = 15, showing the expected cortex–medulla pattern, with cortex faster than medulla.

oriented longitudinally to delineate cortex and medulla in B-mode as shown in Fig. 4.13(b-c).

Excitation and tracking followed Sec. 4.2. Moderate-output acquisitions used $\ell = 0$ or $\ell = 1$ at MI=6–8. Figure 4.13(d) shows a reference phase velocity map using a high-output push $\ell = 0$ at MI=15. Cortex and medulla ROIs were delineated on B-mode images based on anatomical landmarks as illustrated in Fig. 4.13(c).

Metrics

Phase-velocity maps $c_{\mathrm{ph}}(x, z)$ were reconstructed using the phase gradient elastography method [González-Mateo et al., 2025]. Phase-velocity maps were obtained by averaging single-frequency estimates over a window centered at the spectral peak, 150-350 Hz, improving robustness while preserving spatial resolution.

For each spatial region, the mean phase velocity μ and standard deviation

σ were computed within the corresponding ROI. The signed mean contrast was defined as

$$\Delta\mu = \mu_{\text{cortex}} - \mu_{\text{medulla}}. \tag{4.19}$$

To capture cortex–medulla relation, contrast-to-noise ratio was reported in decibels as

$$\text{CNR} = 20\log_{10}\left(\frac{\Delta\mu}{\sqrt{\sigma^2_{\text{cortex}} + \sigma^2_{\text{medulla}}}}\right). \tag{4.20}$$

Finally, spatial agreement with the anatomical reference segmentation was quantified using the Dice similarity coefficient

$$\text{Dice} = \frac{2|A \cap B|}{|A| + |B|}, \tag{4.21}$$

where B denotes the binary cortex mask manually delineated on the B-mode image using visible anatomical boundaries between cortex and medulla, and A denotes the binary cortex mask manually delineated on the reconstructed phase-velocity map as the peripheral region exhibiting higher phase velocity than the inner medulla. The Dice coefficient ranges from 0 (no overlap) to 1 (perfect overlap), with higher values indicating more accurate spatial recovery of the cortex region.

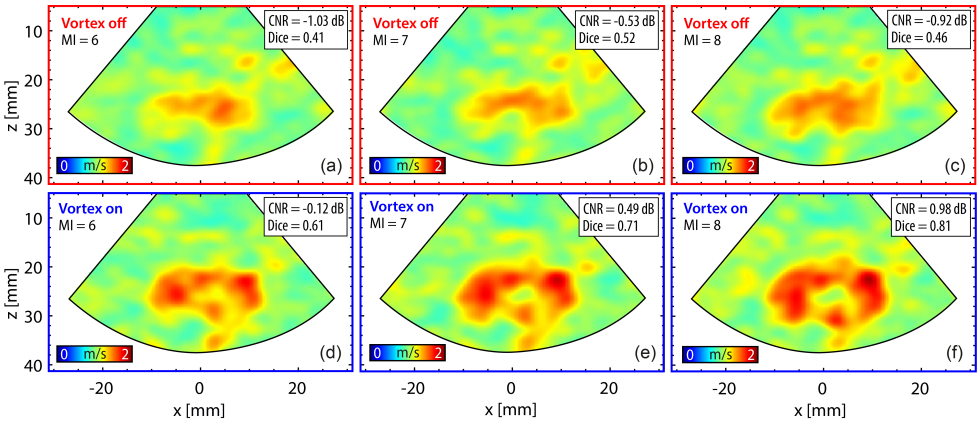


Figure 4.14: Phase-velocity maps of *ex vivo* kidney at moderate output. (a–c) Conventional excitation ($\ell = 0$, vortex off) at MI = 6, 7, and 8. (d–f) Vortex excitation ($\ell = 1$, vortex on) at MI = 6, 7, and 8. Insets report the signed $\text{CNR}_{\pm}$ and Dice coefficient for each reconstruction.

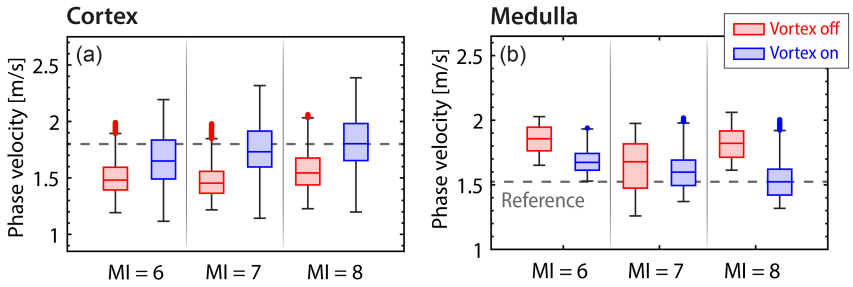


Figure 4.15: Statistical summary of cortex–medulla phase-velocity values in *ex vivo* kidney. (a) Boxplots of mean cortex phase velocity for vortex off and vortex on reconstructions at MI = 6–8, with the MI = 15 reference indicated by the dashed line. (b) Corresponding medulla statistics.

Results

Figure 4.13(d) illustrates the result for the high-MI reference condition ($\ell = 0$, MI = 15). Here the phase-velocity map exhibited the expected cortex–medulla structure, with a stiffer outer ring (cortex) and a softer inner region (medulla). Across repeated acquisitions, mean phase velocities were 1.84 ± 0.07 m/s in cortex and 1.42 ± 0.06 m/s in medulla, consistent with reported speed values in kidney [Liu et al., 2017, Singh et al., 2017, Vasconcelos et al., 2024].

At moderate output, representative maps are shown in Fig. 4.14. With conventional excitation ($\ell = 0$), cortex–medulla contrast remained weak over MI = 6–8, yielding negative CNR values in all cases and Dice coefficients below 0.6, as observed in Figs. 4.14(a–c). By contrast, vortex excitation ($\ell = 1$) improved both contrast and spatial agreement under matched output conditions, as shown in Figs. 4.14(d–f). In this case, CNR was close to zero at MI = 6, became positive at MI = 7, and reached nearly 1 dB at MI = 8, while Dice increased from 0.61 at MI = 6 to 0.81 at MI = 8. These quantitative results are listed in Table 4.1. The same trend is summarized in Fig. 4.15, which shows that vortex excitation brought both cortex and medulla phase-velocity estimates closer to the reference condition than conventional excitation.

	Vortex off, $\ell = 0$		**Vortex on**, $\ell = 1$	
MI	CNR [dB]	Dice	CNR [dB]	Dice
6	−1.03	0.41	−0.12	0.61
7	−0.53	0.52	0.49	0.71
8	−0.92	0.46	0.98	0.81

Table 4.1: CNR and Dice coefficient obtained in the *ex vivo* kidney experiment for conventional excitation (vortex off) and vortex excitation (vortex on) at different MI values.

4.7 Discussion

This chapter shows that focused acoustic vortices provide an elastography actuation channel that is qualitatively different from conventional ARF pushing: orbital angular momentum transfer enables radiation torque and, consequently, measurable rotational motion. Hydrophone measurements verified the two field features required for torque generation—annular field distribution and azimuthal phase winding (Fig. 4.3)—and the gelatin-phantom experiment provided direct mechanical evidence through a strong torsional response that was absent for $\ell = 0$ (Fig. 4.7). This behavior is consistent with OAM transfer rather than a trivial exposure increase, since $v_{\theta,\mathrm{RMS}}$ was about fifteen times higher for $\ell = 1$ than for $\ell = 0$, and the measurements closely matched simulations driven by the measured source fields.

Beyond torsion, vortex structuring improved shear-wave detectability under exposure normalization. At MI= 5, peak axial velocity increased from $\sim$0.42 mm/s ($\ell = 0$) to $\sim$2.05 mm/s ($\ell = 4$), accompanied by an $\sim$11 dB SNR gain (Fig. 4.9). This benefit is most relevant when MI is the binding constraint, but the hydrophone measurements also highlight a practical trade-off: at fixed drive, MI decreases with $|\ell|$ because peak rarefactional pressure is redistributed over a wider annulus (Fig. 4.5). For instance, at 40 V, MI decreases from 5.8 for $\ell = 0$ to 2.7 for $\ell = 4$. Thus, the best-performing charge depends on the available output level, since higher ℓ values broaden the field but reduce peak pressure. In practice, these results indicate that the topological charge cannot be selected independently of the acoustic output. A higher ℓ can improve shear-wave visibility under matched MI, but it also lowers the peak pressure at a given driving voltage. The choice of ℓ should therefore balance field structuring and available output for the specific imaging condition.

These excitation-level gains translated into more reliable phase-velocity map reconstructions in challenging settings. In the heterogeneous phantom, vortex actuation recovered inclusion contrast at MI= 5, where conventional focusing failed to delineate the target (CNR improvement > 6 dB; Fig. 4.12); inclusion phase velocity increased from $\sim$2.5 m/s ($\ell = 0$) to $\sim$3.8 m/s ($\ell = 3$), approaching the 3.6 m/s reference. In *ex vivo* kidney, vortex excitation stabilized cortex–medulla contrast at MI= 6–8, yielding positive signed CNR and Dice values exceeding 0.8 at MI= 8 (Fig. 4.14). Together, these results indicate that structured actuation is particularly advantageous in output-limited or weak-contrast regimes where conventional excitation is more prone to ambiguous reconstructions.

The main barrier to clinical translation of the present realization is acoustic exposure. The current prototype operates at 1.1 MHz, below the frequencies commonly used for SWE pushing (typically 3–7 MHz) [Maksuti et al., 2016, Tanter et al., 2008], and the achieved MI exceeds the diagnostic guidance value of 1.9 [Food and Drug Administration, 2021]. These experiments should therefore be interpreted primarily as a proof of physical feasibility enabled by the available hardware. Using only the $1/\sqrt{f}$ dependence of MI (Eq. 4.11) as a first-order estimate, increasing the frequency from 1.1 to 4 MHz would reduce MI by a factor of about 0.52. Under this approximation, a condition producing MI= 3 at 1.1 MHz would correspond to MI= 1.57 at 4 MHz. However, this is only a first-order estimate, since it neglects frequency-dependent changes in gain, focal field formation, attenuation, and absorption. Quantitative prediction at higher frequencies would require dedicated simulation and acoustic re-characterization.

Finally, several practical aspects could improve the method. The *ex vivo* kidney study included only a limited number of specimens, so larger datasets will be needed to strengthen the contrast and segmentation analysis. From a system-design perspective, clinical implementation would benefit from more integrated architectures, such as matrix arrays, which could combine vortex generation and ultrafast tracking in a single platform, improve alignment robustness, and allow more flexible control of ℓ and focal geometry under safety constraints.

4.8 Conclusions

This chapter established focused acoustic vortices as an elastography actuation modality enabled by a momentum channel that is absent in conventional ARF: the transfer of orbital angular momentum to soft media. This concept was realized experimentally using a focused annular transducer, passive vortex lenses, and a concentric ultrasound imaging system for ultrafast motion tracking. Acoustic characterization confirmed that the fabricated lenses generated the expected annular focal structure and azimuthal phase winding, and the torsional-validation experiment provided direct evidence of torque-driven motion through a pronounced increase in the rotational component under vortex excitation.

At matched MI, increasing topological charge improved SNR, thereby leading to an increase in shear-wave detectability. This demonstrate that structured excitation can enhance motion observables under exposure-normalized conditions. These gains translated into more robust elasticity imaging, enabling recovery of inclusion contrast in a heterogeneous phantom at reduced output and restoring cortex–medulla differentiation in *ex vivo* kidney where conventional excitation failed across the tested conditions.

The present implementation is not yet suitable for clinical use, since it operates at 1.1 MHz, which exceeds diagnostic MI limits. However, the results point to a clear route toward clinical implementation based on higher-frequency operation and integrated array architectures. In summary, acoustic vortices have been demonstrated to be a novel mechanism for interacting with soft tissues during elastography and as a useful tool for efficiently generating tissue motion. Consequently, vortex elastography has emerged as an alternative to conventional pushing elastography in scenarios where output is limited or weak contrasts are present.

Publications

This section contains a compendium of the seminal scientific articles in which the results presented in this chapter have been published.

Journal articles

- González-Mateo, E., Rodríguez-Sendra, J., Camarena, F., and Jimenez, N., "Midiendo la elasticidad de tejidos biológicos empleando ultrasonidos," *Revista de Acústica (Online)*, vol. 54, no. 3–4, pp. 12–19, 2023.

Conferences articles

- González-Mateo, E., Camarena, F., and Jimenez, N., "Acoustic vortices and phase gradient techniques for ultrasound elastography," in *International Tissue Elasticity Conference (ITEC 2024)*, Lyon, France, 2024.
- González-Mateo, E., Camarena, F., and Jimenez, N., "Vórtices acústicos para medir la elasticidad de tejidos biológicos," in *TECNIACÚSTICA 2024. 55º Congreso Español de Acústica. XIII Encuentro Ibérico de Acústica*, Faro, Portugal, 2024, pp. 656–660.
- González-Mateo, E., Jimenez, N., and Camarena, F., "Quasi-omnidirectional shear wave generation using acoustic vortices for elastography," in *2022 IEEE International Ultrasonics Symposium (IUS)*, Venice, Italy, 2022, pp. 1–4.
- Jimenez, N., González-Mateo, E., Camarena, F., and Staliunas, K., "Rotating acoustic drills by the interference of detuned vortices," in *2022 IEEE International Ultrasonics Symposium (IUS)*, Venice, Italy, 2022, pp. 1–3.
- Jimenez, N., González-Mateo, E., Camarena, F., and Staliunas, K., "Acoustic drill beams: twisting wave intensity by the superposition of detuned vortices," in *TECNIACÚSTICA 2022. 53º Congreso Español de Acústica. XII Congreso Ibérico de Acústica. European Symposium on NVH in Electric Vehicles*, Elche, España, 2022, pp. 1141–1144.
- Jimenez, N., González-Mateo, E., Camarena, F., and Staliunas, K., "Synthesizing acoustic drill beams by the superposition of detuned vortices," in *182nd Meeting of the Acoustical Society of America*, Denver, EEUU, 2022, p. 121.
- Jimenez, N., González-Mateo, E., Benlloch Baviera, J. M., and Camarena, F., "Reflectors for Therapeutic Vortex Beam Generation," in *21th International Symposium for Therapeutic Ultrasound (ISTU 2022)*, Toronto, Canadá, 2022, pp. 342–343.

References

Abbott, J. G. (1999). Rationale and derivation of MI and TI—a review. *Ultrasound in Medicine & Biology*, 25(3):431–441.

Aleef, T. A., Lobo, J., Baghani, A., Mohammed, S., Eskandari, H., Moradi, H., Rohling, R., Goldenberg, S. L., Morris, W. J., Mahdavi, S. S., and Salcudean, S. E. (2023). Multi-Frequency 3D Shear Wave Absolute Vibro-Elastography (S-WAVE) System for the Prostate. *IEEE Transactions on Medical Imaging*, 42(11):3436–3450.

Anhäuser, A., Wunenburger, R., and Brasselet, E. (2012). Acoustic Rotational Manipulation Using Orbital Angular Momentum Transfer. *Physical Review Letters*, 109(3):034301. Publisher: American Physical Society.

Ashkin, A., Dziedzic, J. M., Bjorkholm, J. E., and Chu, S. (1986). Observation of a single-beam gradient force optical trap for dielectric particles. *Optics Letters*, 11(5):288–290. Publisher: Optica Publishing Group.

Ashkin, A., Dziedzic, J. M., and Yamane, T. (1987). Optical trapping and manipulation of single cells using infrared laser beams. *Nature*, 330(6150):769–771. Number: 6150 Publisher: Nature Publishing Group.

Baresch, D., Thomas, J.-L., and Marchiano, R. (2016). Observation of a Single-Beam Gradient Force Acoustical Trap for Elastic Particles: Acoustical Tweezers. *Physical Review Letters*, 116(2):024301. Publisher: American Physical Society.

Baresch, D., Thomas, J.-L., and Marchiano, R. (2018). Orbital Angular Momentum Transfer to Stably Trapped Elastic Particles in Acoustical Vortex Beams. *Physical Review Letters*, 121(7):074301. Publisher: American Physical Society.

Canney, M. S., Bailey, M. R., Crum, L. A., Khokhlova, V. A., and Sapozhnikov, O. A. (2008). Acoustic characterization of high intensity focused ultrasound fields: A combined measurement and modeling approach. *The Journal of the Acoustical Society of America*, 124(4):2406–2420.

Child, S. Z., Hartman, C. L., Schery, L. A., and Carstensen, E. L. (1990). Lung damage from exposure to pulsed ultrasound. *Ultrasound in Medicine & Biology*, 16(8):817–825.

Cook, J. R., Bouchard, R. R., and Emelianov, S. Y. (2011). Tissue-mimicking phantoms for photoacoustic and ultrasonic imaging. *Biomedical Optics Express*, 2(11):3193–3206.

Démoré, C. E., Dahl, P. M., Yang, Z., Glynne-Jones, P., Melzer, A., Cochran, S., MacDonald, M., and Spalding, G. C. (2014). Acoustic Tractor Beam. *Physical Review Letters*, 112(17):174302. Publisher: American Physical Society.

Fan, D. L., Zhu, F. Q., Cammarata, R. C., and Chien, C. L. (2011). Electric tweezers. *Nano Today*, 6(4):339–354.

Food and Drug Administration (2021). Marketing clearance of diagnostic ultrasound systems and transducers. U.S. Department of Health and Human Services, Center for Devices and Radiological Health.

Fu, S. and Gao, C. (2023). *Optical Vortex Beams: Fundamentals and Techniques*. Advances in Optics and Optoelectronics. Springer Nature, Singapore.

González-Mateo, E., Camarena, F., and Jiménez, N. (2025). Real-time ultrasound shear wave elastography using a local phase gradient. *Computer Methods and Programs in Biomedicine*, 260:108529.

Guo, S., Ya, Z., Wu, P., and Wan, M. (2022). A review on acoustic vortices: Generation, characterization, applications and perspectives. *Journal of Applied Physics*, 132(21):210701.

Harris, G. R., Howard, S. M., Hurrell, A. M., Lewin, P. A., Schafer, M. E., Wear, K. A., Wilkens, V., and Zeqiri, B. (2023). Hydrophone Measurements for Biomedical Ultrasound Applications: A Review. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 70(2):85–100. Conference Name: IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control.

Hefner, B. T. and Marston, P. L. (1999). An acoustical helicoidal wave transducer with applications for the alignment of ultrasonic and underwater systems. *The Journal of the Acoustical Society of America*, 106(6):3313–3316.

Herman, B. A. and Harris, G. R. (2002). Models and regulatory considerations for transient temperature rise during diagnostic ultrasound pulses. *Ultrasound in Medicine & Biology*, 28(9):1217–1224.

Jiménez, N., Ealo, J. L., Muelas-Hurtado, R. D., Duclos, A., and Romero-García, V. (2021). Subwavelength acoustic vortex beams using self-demodulation. *Physical review applied*, 15.

Jiménez, N., Benlloch, J. M., and Camarena, F. (2020). A new elastographic technique using acoustic vortices. In *2020 IEEE International Ultrasonics Symposium (IUS)*, pages 1–4. ISSN: 1948-5727.

Jiménez, N., Benlloch Baviera, J. M., and Camarena, F. (2021). Procedimiento de obtención de propiedades elásticas de un sólido blando que hace uso de vórtices acústicos.

Jiménez, N., Picó, R., Sánchez-Morcillo, V., Romero-García, V., García-Raffi, L. M., and Staliunas, K. (2016). Formation of high-order acoustic Bessel beams by spiral diffraction gratings. *Physical Review E*, 94(5):053004. Publisher: American Physical Society.

Jiménez, N., Romero-García, V., García-Raffi, L. M., Camarena, F., and Staliunas, K. (2018). Sharp acoustic vortex focusing by Fresnel-spiral zone plates. *Applied Physics Letters*, 112(20):204101.

Jiménez-Gambín, S., Jiménez, N., Benlloch, J. M., and Camarena, F. (2019). Generating Bessel beams with broad depth-of-field by using phase-only acoustic holograms. *Scientific Reports*, 9(1):20104.

Jiménez-Gambín, S., Jiménez, N., and Camarena, F. (2020). Transcranial Focusing of Ultrasonic Vortices by Acoustic Holograms. *Physical Review Applied*, 14(5):054070.

Kijanka, P. and Urban, M. W. (2024). Ultrasound Shear Elastography with Expanded Bandwidth (USEWEB): A Novel Method for 2D Shear Phase Velocity Imaging of Soft Tissues. *IEEE Transactions on Medical Imaging*, pages 1–1. Conference Name: IEEE Transactions on Medical Imaging.

Kreider, W., Yuldashev, P. V., Sapozhnikov, O. A., Farr, N., Partanen, A., Bailey, M. R., and Khokhlova, V. A. (2013). Characterization of a Multi-element Clinical HIFU System Using Acoustic Holography and Nonlinear Modeling. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 60(8):1683–1698.

Li, P., Mao, Z., Peng, Z., Zhou, L., Chen, Y., Huang, P.-H., Truica, C. I., Drabick, J. J., El-Deiry, W. S., Dao, M., Suresh, S., and Huang, T. J. (2015). Acoustic separation of circulating tumor cells. *Proceedings of the National Academy of Sciences*, 112(16):4970–4975. Publisher: Proceedings of the National Academy of Sciences.

Liu, H.-C., Zeng, Y., Gong, C., Chen, X., Kijanka, P., Zhang, J., Genyk, Y., Tchelepi, H., Wang, C., Zhou, Q., and Zhao, X. (2024). Wearable bioadhesive ultrasound shear wave elastography. *Science Advances*, 10(6):eadk8426. Publisher: American Association for the Advancement of Science.

Liu, X., Li, N., Xu, T., Sun, F., Li, R., Gao, Q., Chen, L., and Wen, C. (2017). Effect of renal perfusion and structural heterogeneity on shear wave elastography of the kidney: an in vivo and ex vivo study. *BMC nephrology*, 18(1):265.

Lo, W.-C., Huang, Y.-L., Fan, C.-H., and Yeh, C.-K. (2021). 3-D Ultrafast Ultrasound Imaging of Microbubbles Trapped Using an Acoustic Vortex. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 68(12):3507–3514. Conference Name: IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control.

Loupas, T., Powers, J., and Gill, R. (1995). An axial velocity estimator for ultrasound blood flow imaging, based on a full evaluation of the Doppler equation by means of a two-dimensional autocorrelation approach. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 42(4):672–688.

Maksuti, E., Widman, E., Larsson, D., Urban, M. W., Larsson, M., and Bjällmark, A. (2016). Arterial Stiffness Estimation by Shear Wave Elastography: Validation in Phantoms with Mechanical Testing. *Ultrasound in Medicine & Biology*, 42(1):308–321.

Marzo, A., Caleap, M., and Drinkwater, B. W. (2018). Acoustic Virtual Vortices with Tunable Orbital Angular Momentum for Trapping of Mie Particles. *Physical Review Letters*, 120(4):044301. Publisher: American Physical Society.

Marzo, A., Seah, S. A., Drinkwater, B. W., Sahoo, D. R., Long, B., and Subramanian, S. (2015). Holographic acoustic elements for manipulation of levitated objects. *Nature Communications*, 6(1):8661. Number: 1 Publisher: Nature Publishing Group.

Nemes, A. F., Toma, A. I., Dima, V., Serboiu, S. C., Necula, A. I., Stoiciu, R., Ulmeanu, A. I., Marinescu, A., and Ulmeanu, C. (2024). Use of Lung Ultrasound in Reducing Radiation Exposure in Neonates with Respiratory Distress: A Quality Management Project. *Medicina*, 60(2).

Nightingale, K., McAleavey, S., and Trahey, G. (2003). Shear-wave generation using acoustic radiation force: in vivo and ex vivo results. *Ultrasound in Medicine and Biology*, 29(12):1715–1723.

Owen, J., Pankhurst, Q., and Stride, E. (2012). Magnetic targeting and ultrasound mediated drug delivery: Benefits, limitations and combination. *International Journal of Hyperthermia*, 28(4):362–373. Publisher: Taylor & Francis

Qiang, B., Xu, Q., Pan, Y., Wang, J., Shen, C., Peng, X., Shen, W., Zhang, Y., and Zhu, X. (2024). Shear wave elastography: A noninvasive approach for assessing acute kidney injury in critically ill patients. *PloS One*, 19(1):e0296411.

Sarvazyan, A. P., Rudenko, O. V., Swanson, S. D., Fowlkes, J. B., and Emelianov, S. Y. (1998). Shear wave elasticity imaging: a new ultrasonic technology of medical diagnostics. *Ultrasound in Medicine & Biology*, 24(9):1419–1435.

Seliverstova, E., Caenen, A., Bézy, S., Nooijens, S., Voigt, J.-U., and D'hooge, J. (2022). Comparing Myocardial Shear Wave Propagation Velocity Estimation Methods Based on Tissue Displacement, Velocity and Acceleration Data. *Ultrasound in Medicine & Biology*.

Shi, C., Dubois, M., Wang, Y., and Zhang, X. (2017). High-speed acoustic communication by multiplexing orbital angular momentum. *Proceedings of the National Academy of Sciences*, 114(28):7250–7253. Publisher: Proceedings of the National Academy of Sciences.

Singh, H., Panta, O. B., Khanal, U., and Ghimire, R. K. (2017). Renal Cortical Elastography: Normal Values and Variations. *Journal of Medical Ultrasound*, 25(4):215–220.

Tanter, M., Bercoff, J., Athanasiou, A., Deffieux, T., Gennisson, J.-L., Montaldo, G., Muller, M., Tardivon, A., and Fink, M. (2008). Quantitative Assessment of Breast Lesion Viscoelasticity: Initial Clinical Results Using Supersonic Shear Imaging. *Ultrasound in Medicine and Biology*, 34(9):1373–1386.

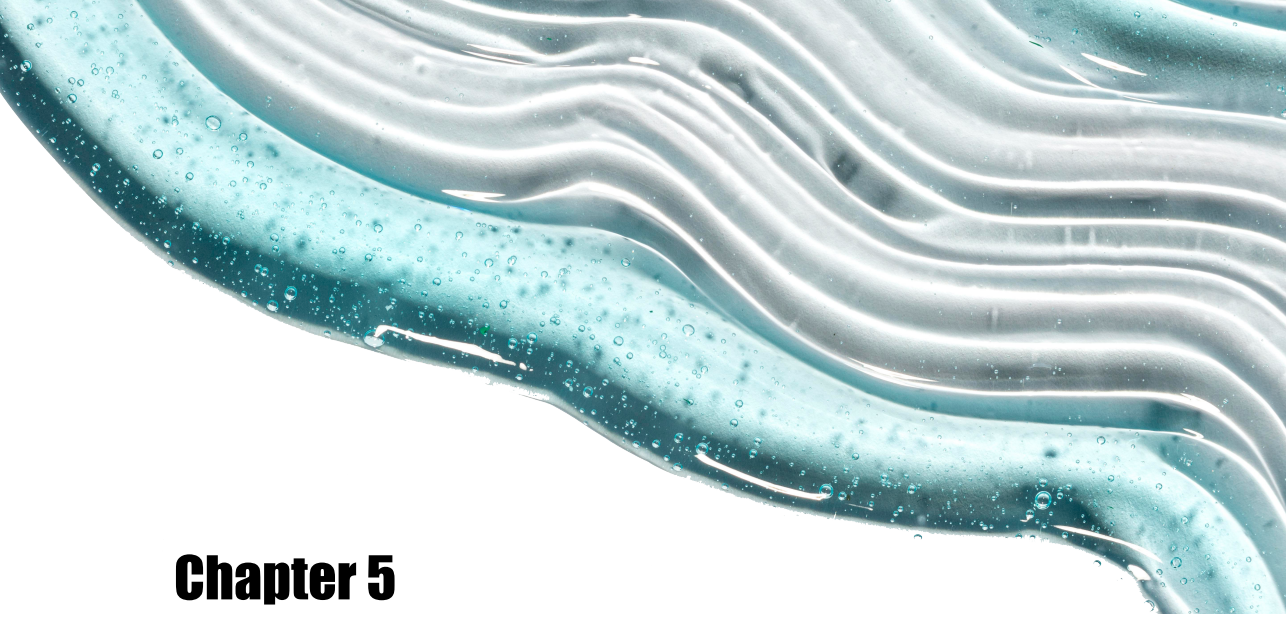
Thanaboonnipat, C., Sutayatram, S., Buranakarl, C., and Choisunirachon, N. (2019). Renal shear wave elastography and urinary procollagen type III amino-terminal propeptide (uPIIINP) in feline chronic kidney disease. *BMC Veterinary Research*, 15(1):54.

Thomas, J.-L. and Marchiano, R. (2003). Pseudo angular momentum and topological charge conservation for nonlinear acoustical vortices. *Physical Review Letters*, 91(24):244302.

Treeby, B. E., Jaros, J., Rohrbach, D., and Cox, B. T. (2014). Modelling elastic wave propagation using the k-Wave MATLAB Toolbox. In *2014 IEEE International Ultrasonics Symposium*, pages 146–149.

Vasconcelos, L., Kijanka, P., Grande, J. P., Oliveira, R., Amador, C., Aristizabal, S., Sanger, N. M., Rule, A. D., Atwell, T. D., and Urban, M. W. (2024). Kidney cortex shear wave motion simulations based on segmented biopsy histology. *Computer Methods and Programs in Biomedicine*, 245:108035.

Volke-Sepúlveda, K., Santillán, A. O., and Boullosa, R. R. (2008). Transfer of Angular Momentum to Matter from Acoustical Vortices in Free Space. *Physical Review Letters*, 100(2):024302. Publisher: American Physical Society.



Chapter 5

MLS-coded shear wave elastography

Abstract

Shear-wave elastography (SWE) often operates close to the motion noise floor when acoustic output is limited. This chapter evaluates a pulse-compression strategy in which acoustic radiation force (ARF) is modulated by Maximum Length Sequences (MLS) and the resulting coded shear-wave response is compressed by frequency-domain deconvolution. Matched simulations and homogeneous phantom experiments show that MLS-coded pushes preserve the underlying shear-wave dynamics while providing monotonic SNR gains, more stable group-velocity estimates in the time domain, and smoother phase-velocity mapping in the frequency domain. Because long temporal sequences increase time-average exposure and limit the achievable repetition rate, the method is characterized in terms of mechanical index (MI) and intensity, highlighting an explicit trade-off between SNR, temporal resolution, and thermal dose. Overall, long MLS codes at low MI can approach the velocity-estimation quality obtained with higher-MI single-push excitations while keeping peak pressure unchanged. Taken together, these results establish MLS-coded ARF as a practical strategy to improve SWE robustness under low-SNR and safety-constrained operating conditions.

Outline

5.1 Introduction 137
5.2 MLS-coded ARF pulse-compression method 138
5.2.1 Maximum-Length Sequences and binary ARF coding . . 138
5.2.2 Pulse-compression method 140
5.2.3 Bit-level timing and implementation 141
5.3 Numerical and experimental methods 143
5.3.1 Homogeneous phantom experiments 143
5.3.2 Shear-wave simulations 144
5.3.3 Motion tracking and post-processing 145
5.3.4 Performance metrics 145
5.4 Results in homogeneous media 146
5.4.1 Numerical and experimental validation 146
5.4.2 SNR enhancement 148
5.4.3 Group-velocity estimation 151
5.4.4 Phase-velocity mapping 152
5.5 Acoustic safety characterization 154
5.5.1 Mechanical index 155
5.5.2 Spatial-peak pulse-average intensity 156
5.5.3 Spatial-peak time-average intensity 157
5.6 Discussion 159
5.7 Conclusions 162

5.1 Introduction

Clinical implementations of shear wave elastography (SWE) rely on acoustic radiation force (ARF) to generate localized shear waves whose propagation is tracked to estimate tissue mechanical properties [Sandrin et al., 2003, Deffieux et al., 2009, Gennisson, 2018, Bercoff et al., 2004]. In practice, the amplitude of ARF-induced displacements is intrinsically low, and shear waves undergo rapid attenuation in biological media through absorption, geometric spreading, scattering, and out-of-plane leakage [Deng et al., 2015, Bruce et al., 2017, Zhou et al., 2020]. As a result, the signal-to-noise ratio (SNR) of the tracked motion deteriorates with depth and in highly attenuating tissues. Simply increasing peak acoustic intensity can partially compensate for this decay, but such strategies are fundamentally constrained by safety regulations on acoustic output and by hardware limits of the transmit beam.

Pulse-compression and coded-excitation techniques were originally developed in radar and sonar [Klauder et al., 1960, Anstey, 1964], later adopted in system identification [Nowak and Veen, 1994], and have been successfully translated to medical ultrasound imaging [Chiao and Hao, 2005]. In this context, coded excitation distributes acoustic energy over a longer duration without increasing peak pressure, thereby improving SNR at the receiver [Nowak and Veen, 1994, Yücel et al., 2015]. This longer, lower-peak-amplitude transmission strategy also helps reduce cavitation and propagation nonlinearities near the transducer.

In ultrasound elastography, chirps and Barker or Golay codes have been used to generate high-SNR B-mode images and to enhance tissue displacement tracking [Liu and Insana, 2005, Song et al., 2015]. In addition chirp-coded ARF pushes with short durations (on the order of 100 μs) have been proposed to boost displacement amplitudes [Ding et al., 2025]. However, these strategies generally do not encode the shear wave itself and therefore do not fully exploit the coding structure for deconvolution-based pulse compression.

Works on amplitude-modulated ARF have explored Barker and chirp codes to enhance shear-wave signals [Nguyen et al., 2015]. However, chirp excitations are generally more demanding to implement in transmission, often requiring AWG-based waveform synthesis. By contrast, systematic evaluation of long-duration binary sequences such as maximum length sequences (MLS) for direct ARF excitation remains limited. MLS provide near-ideal two-level

autocorrelation, an approximately flat spectral density, and robustness to noise, which are advantageous for both excitation and deconvolution-based pulse compression [Bleakley and Scaife, 1995, White and Challis, 1994]. Moreover, unlike many coded-elastography approaches that rely on separate actuation and detection systems [He et al., 2017, Nguyen et al., 2014], MLS-based ARF can be implemented within a single ultrasound platform, thereby reducing system complexity and improving prospects for clinical translation.

To address these limitations, this chapter investigates a coded-excitation strategy for ARF-based SWE based on long-duration maximum length sequences (MLS). The proposed approach modulates the ARF transmission with an MLS, uses the same clinical ultrasound probe for both actuation and motion tracking, and subsequently applies frequency-domain deconvolution to recover a time-localized, high-SNR shear-wave response. In this way, acoustic energy and displacement detectability can be increased without raising peak pressure, while preserving compatibility with standard ultrasound hardware. The method is validated through matched phantom experiments and full shear-wave simulations, and its impact is quantified in terms of SNR improvements and the stability of group- and phase-velocity estimates across different stiffness levels and excitation parameters. Throughout the chapter, these results are interpreted in the broader context of the simulation method and phase-velocity reconstruction tools developed earlier in the thesis, positioning MLS-coded excitation as a practical route to enhance the robustness and clinical viability of SWE.

5.2 MLS-coded ARF pulse-compression method

5.2.1 Maximum-Length Sequences and binary ARF coding

Maximum-length sequences, also known as m-sequences, are pseudorandom sequences generated using linear feedback shift registers (LFSR) defined over Galois fields [Golomb, 1982]. For a register of order N, the corresponding MLS has a period $L = 2^N - 1$ and visits every nonzero N-bit state exactly once per cycle [Gemba et al., 2021]. MLS exhibit a two-level periodic autocorrelation, with a single prominent peak at zero lag and nearly constant side-lobes elsewhere, which makes them particularly suitable for impulse-response measurement and deconvolution [Rife and Vanderkooy, 1989]. In the frequency

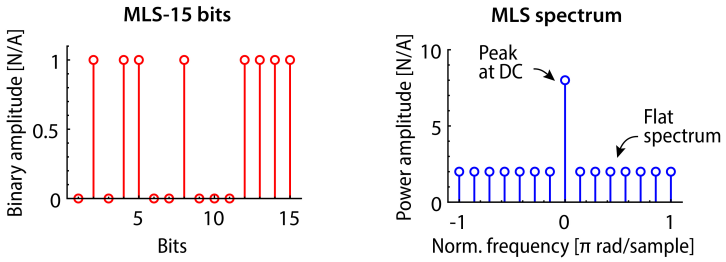


Figure 5.1: Discrete binary MLS signal of 15 bits length, $L = 2^4 - 1$, and the corresponding flat spectrum behavior of this type of pseudorandom sequences.

domain, MLS display an approximately flat power spectrum across their support, providing broadband excitation and facilitating stable deconvolution in noisy conditions [White and Challis, 1994]. The MLS used in this work were generated using the equations and primitive polynomials described in [Buracas and Boynton, 2002], ensuring maximal period and the desired autocorrelation properties for the selected register orders.

In the present method, these properties are exploited to encode the ARF excitation itself. A binary sequence $s(t)$ is constructed from the MLS by assigning each bit to a fixed rectangular pulse of duration T_b, such that $s(t) \in \{0, 1\}$ indicates whether an ARF push is present ($s = 1$) or absent ($s = 0$) within that slot. The continuous-time ARF excitation waveform used for the pushing beam is then defined as

$$w(t) = s(t)\, u(t), \tag{5.1}$$

where $u(t) = \sin(\omega_0 t)$ denotes a sinusoidal signal and ω_0 is the angular frequency of the pushing beam with a prescribed number of cycles. When $s(t) = 1$, a short high-frequency burst is transmitted and generates an ARF-induced shear disturbance; when $s(t) = 0$, no ARF push is emitted and the tissue/soft-solid is left to evolve freely. In this way, the MLS determines the temporal pattern of pushes and silent intervals, and the total excitation time is given by $T_\mathrm{seq} = L\, T_\mathrm{b}$.

As in other coded-actuation approaches, the measured motion can be interpreted as the tissue response to a long, binary-modulated excitation whose temporal structure is known *a priori*. The distinctive feature of the present scheme is that the code is embedded directly into the ARF push generated by the imaging probe, rather than being delivered by a separate external

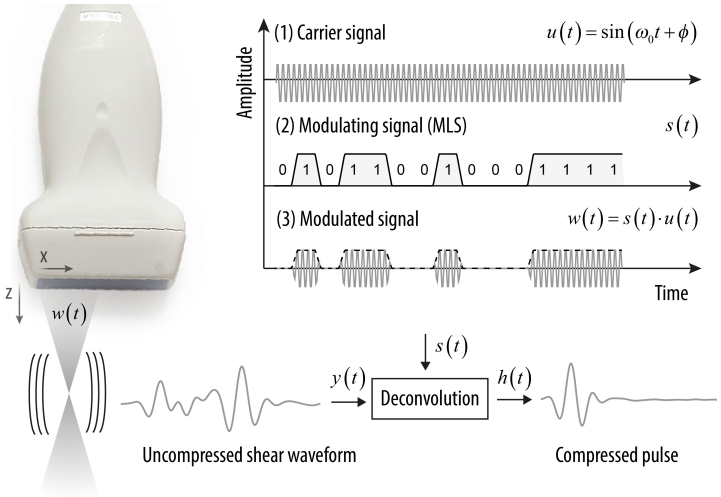


Figure 5.2: MLS-coded ARF pulse compression method. (a) Schematic flowchart of the MLS excitation procedure. The carrier ultrasonic signal, $u(t)$, is modulated by the MLS signal, $s(t)$, and the modulated signal, $w(t)$, is applied to the transducer. The ultrasound beam generates an ARF whose envelope is proportional to the sequence $s(t)$. The resulting shear waves, $y(x,t)$, are tracked by the ultrasound probe. Their frequency-domain deconvolution with $s(t)$ provides the shear impulse response $h(t,x)$.

actuator. This single-probe configuration not only simplifies the experimental setup, but also ensures that shear-wave generation remains localized within the imaging plane, thereby reducing out-of-plane wave components. This point is particularly relevant for the reconstruction methods used in this thesis, whose accuracy depends on tracking predominantly in-plane wave propagation, and therefore contributes directly to improved elastographic image quality.

5.2.2 Pulse-compression method

The interaction between the MLS-coded ARF excitation and the tissue is modeled, to first order, as a linear time-invariant (LTI) system. Under this assumption the measured tissue motion $y(t)$ at a given spatial location can be expressed as the convolution between the input signal that modulates the force $s(t)$ and the tissue impulse response $h(t)$, plus additive noise $n(t)$:

$$y(t) = s(t) * h(t) + n(t), \tag{5.2}$$

where the operator $*$ denotes convolution in time. In the context of SWE, $y(t)$ represents the tracked shear-wave displacement or particle velocity at a fixed point, while $h(t)$ can be interpreted as the shear-wave response that would be produced ideally by a band-limited, time-localized ARF push at that location.

The goal of pulse compression is to recover an estimate of $h(t)$ from the recorded signal $y(t)$ and the known modulation signal $s(t)$. To this end, a frequency-domain deconvolution is applied:

$$h(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \frac{Y(\omega)}{S(\omega)} \, e^{i\omega t} \, d\omega \tag{5.3}$$

where $Y(\omega)$ and $S(\omega)$ are the Fourier transforms of $y(t)$ and $s(t)$, respectively. Owing to the near-ideal autocorrelation and approximately flat spectral magnitude of MLS, $S(\omega)$ remains close to constant over the relevant bandwidth and does not introduce strong spectral nulls. Consequently, the deconvolution can effectively invert the code without markedly amplifying noise in specific frequency bands, yielding a time-compressed response with a narrow dominant peak and reduced side-lobes in the recovered impulse response [Rife and Vanderkooy, 1989, Gemba et al., 2021, White and Challis, 1994]. In practice, the deconvolution is implemented in discrete time using fast Fourier transforms (FFT), and the same operation can be applied independently at each spatial location to produce a time-compressed shear-wave field $h(x, z, t)$. This time-compressed pulses will be later used for group velocity estimation and phase-velocity mapping.

5.2.3 Bit-level timing and implementation

The MLS-coded excitation was implemented by assigning each bit a fixed time slot of duration T_b, composed of an ARF push, a ring-down delay, and a tracking window,

$$T_\mathrm{b} = T_\mathrm{push} + T_\mathrm{rd} + T_\mathrm{trk}. \tag{5.4}$$

Figure 5.3 illustrate the timing diagram. For bits with $s = 1$, an ARF burst is transmitted during T_push, followed by a delay T_rd that allows residual transducer ringing and near-field reverberations to decay before motion tracking begins. The subsequent interval T_trk is then used to acquire tracking frames. For bits with $s = 0$, no ARF burst is emitted, but the same delay and tracking intervals are preserved. As a result, all bits have identical duration, which ensures a well-defined motion PRF and a uniformly sampled temporal signal

suitable for stable deconvolution. Similar push–track delays are routinely used in SWE and ultrasound monitoring of HIFU therapy to prevent corruption of motion estimates by the strong driving pulse [Jeong et al., 2010, Song et al., 2013].

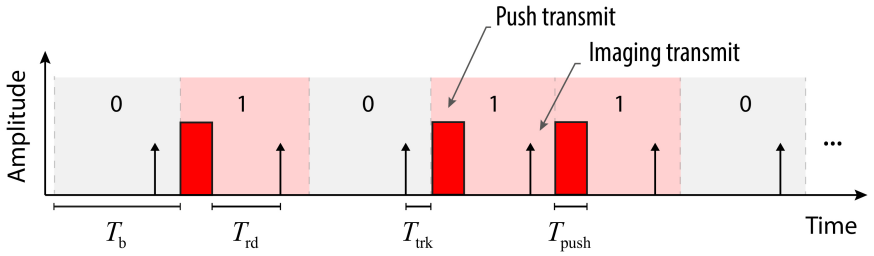


Figure 5.3: Timing diagram for the event sequence programming, illustrating how the total MLS bit time T_b is composed of an ARF push T_{push}, a delay between push and tracking pulses T_{rd}, and a tracking image temporal window T_{trk}.

In practice, the temporal envelope differs slightly from the idealized description in Eq. (5.1), because the ARF push occupies only the initial portion of each bit. Accordingly, the transmitted waveform is more precisely written as

$$w(t) = m(t)\, u(t), \tag{5.5}$$

where $m(t)$ denotes the practical binary modulation actually implemented by the sequence. Let $s[k] \in \{0, 1\}$, with $k = 0, \ldots, L-1$, denote the discrete MLS of length L. The corresponding continuous-time modulation is

$$m(t) = \sum_{k=0}^{L-1} s[k]\, \Pi\left(\frac{t - kT_b}{T_{push}}\right), \tag{5.6}$$

where $\Pi(\cdot)$ is a rectangular window equal to one over the push interval and zero otherwise. Thus, bits with $s[k] = 1$ activate the ARF field only during the first T_{push} of the k-th bit, whereas bits with $s[k] = 0$ correspond to silent intervals. This formulation captures the practical sequence timing while preserving the known binary structure required for subsequent deconvolution.

The total duration of the coded excitation is $T_{seq} = L\,T_b$, where L is the code length. In this work, three MLS lengths were considered, $L = 31$, 63, and 127, yielding excitation durations of a few tens of milliseconds for the selected T_b. The effective duty cycle, defined as the fraction of time during

which ARF pushes are active, is given by

$$D = \frac{N_1 \, T_{\text{push}}}{T_{\text{seq}}}, \tag{5.7}$$

where $N_1 = (L+1)/2$ is the number of ones in the MLS. This duty-cycle control allows the total deposited energy to increase with code length while keeping the peak intensity of each individual push unchanged.

5.3 Numerical and experimental methods

This section describes the experimental implementation of MLS-coded ARF in calibrated elasticity phantoms, the corresponding full shear-wave numerical simulations, the displacement-estimation and preprocessing, and the metrics used to quantify performance and acoustic safety.

5.3.1 Homogeneous phantom experiments

Experiments were conducted in four homogeneous elasticity phantoms (Model 039, CIRS Inc., Norfolk, VA, USA) with nominal Young's moduli of 10, 25, 45, and 80 kPa. These values span a range representative of clinically relevant soft-tissue stiffness. Assuming isotropic, nearly incompressible media, the expected shear-wave velocities were estimated from the relation $E \approx 3\mu_1 = 3\rho_0 c_s^2$, yielding approximate speeds of 1.8, 2.9, 3.9, and 5.2 m/s for the four phantoms, respectively.

Shear-wave generation and tracking were performed on a Verasonics V1 research scanner (Verasonics Inc., Kirkland, WA, USA) using a 128-element L11-4v linear array transducer (Verasonics Inc., Redmond, WA, USA). The transducer was operated at a center frequency of 7.24 MHz with an effective fractional bandwidth of approximately 80%. ARF pushes were transmitted using 32 central elements, focused at a lateral position $x = 0$ mm and a depth $z = 20$ mm (f-number = 2). Each push consisted of a 450 cycles burst at 4.5 MHz. Matching the configuration adopted in the MLS timing scheme described in Fig. 5.3, each bit in the MLS-coded excitation is assigned a fixed time slot of $T_{\text{b}} = 350$ μs, divided into a $T_{\text{push}} = 100$ μs ARF push, a $T_{\text{rd}} = 170$ μs delay for reverberation decay, and an $T_{\text{trk}} = 80$ μs acquisition window for shear wave tracking.

Four excitation modes were tested: a conventional single-push (SP) configuration and MLS-coded excitations with lengths $L = 31$, 63, and 127 bits. For each mode, input voltages of 20, 30, 40, 50, and 60 V were applied at the Verasonics transmit stage, covering a range of MI values compatible with diagnostic imaging. Each condition—phantom stiffness, excitation type, and voltage—was repeated three times to assess repeatability.

5.3.2 Shear-wave simulations

Three-dimensional full-wave simulations were carried out using the method introduced in Chapter 2 to replicate shear-wave generation and propagation under conditions matched to the phantom experiments. The acoustic source was designed to emulate the L11-4v array geometry and the same transmission parameters and MLS excitations as in the experiments.

The background medium was modeled as a homogeneous, isotropic elastic solid with parameters corresponding to the 25 kPa phantom: hypothetical compressional speed $c_\ell = 1540$ m/s, shear wave speed $c_s = 2.9$ m/s, density $\rho_0 = 1000$ kg/m^3, and frequency-dependent ultrasound attenuation of 0.5 dB/MHz/cm. Simulations were run on a 0.3 CFL and 4 points per wavelength at the maximum shear wave frequency ($\sim 1500\ Hz$), which ensured numerical stability and limited numerical dispersion over the frequency range of interest.

The resulting body-force source applied in elastic simulation can thus be written as

$$\mathbf{f}(\mathbf{r}, t) = F(\mathbf{r})\, m(t), \tag{5.8}$$

with $m(t)$ defined as above for SP and for each MLS length. This construction ensures that all excitation schemes share the same spatial ARF pattern and differ only in their temporal coding, so that any observed differences in shear-wave response arise solely from the MLS modulation and subsequent pulse compression.

The simulated particle-velocity field $\mathbf{v}(\mathbf{r}, t)$ was recorded over a region of interest encompassing the ARF focus and the main propagation path. For consistency with the experimental processing, the axial component $v_z(x, z, t)$ was extracted at the central elevational plane and stored on the same spatiotemporal grid used to analyze the experimental phantom data.

5.3.3 Motion tracking and post-processing

In the experimental results the displacement estimation was performed from raw IQ data acquired by the Verasonics system. The axial component of shear-wave particle velocity was estimated using the Loupas two-dimensional autocorrelation method [Loupas et al., 1995], with an axial kernel of 1 mm and a three-sample temporal sliding window. IQ data were sampled at 28.96 MHz with demodulation centered at 7.24 MHz.

For SNR and group-velocity analyses, the full axial particle-velocity field $v_z(x, z, t)$ was reduced to a lateral–time map $v(x, t)$ by averaging over a 1 mm axial window centered at the ARF focal depth of 20 mm. To probe spatial repeatability, this axial window was shifted in 0.2 mm steps over a 4 mm range around the focus (from $z = 18$ to 22 mm), yielding 20 realizations per configuration. Each realization thus corresponded to a local shear-wave propagation map $v(x, t)$ extracted under nominally identical excitation conditions but from a different axial slice.

Before metric computation, all datasets underwent directional and wavenumber filtering to isolate the main propagating shear-wave component (see Sec. 3.2.3). A directional filter to select rightward-propagating waves ($\bar{\varphi} = 0$ and $\Delta\varphi = \pi/4$ rad) was used. In addition, suppressing counter-propagating and quasi-static components, while a radial wavenumber bandpass filter was applied to retain phase velocities between 1 and 10 m/s.

The assessment of phase velocity maps was conducted by employing the PGE method [González-Mateo et al., 2025], which was introduced in Chapter 3. The particle velocity motion $v_z(x, z, t)$ after MLS deconvolution functions as the input field for the PGE algorithm. Directional and bandpass filters were employed to extract both the rightward- and leftward-propagating shear waves ($\bar{\varphi} = [0, \pi]$ rad) within the effective shear-wave band, at frequencies that fall within the range of the shear-wave band.

5.3.4 Performance metrics

First, the performance of MLS-coded excitation was quantified primarily in terms of shear-wave SNR and group velocity. Both metrics were derived from the lateral–time particle velocity maps $v(x, t)$ obtained as described above.

SNR was defined as

$$\mathrm{SNR} = 20 \log_{10} \left(\frac{\mathrm{RMS}_{\mathrm{signal}}}{\mathrm{RMS}_{\mathrm{noise}}} \right), \tag{5.9}$$

where $\mathrm{RMS}_{\mathrm{signal}}$ denotes the root-mean-square amplitude of the shear-wave displacement within a time window encompassing the main propagating wavefront, and $\mathrm{RMS}_{\mathrm{noise}}$ is the RMS of pre-push baseline data taken before ARF excitation. SNR-gain for MLS-coded excitations was computed as the difference in SNR (in dB) between each MLS configuration and the corresponding SP reference at the same input voltage.

Group velocity c_g was estimated from the arrival-time progression of the shear wave across lateral positions (see Sec. 2.4.1). For each axial realization, time-domain displacement traces were extracted at lateral positions separated by a fixed step Δx. The time-of-flight between neighboring traces was obtained from the lag maximizing their cross-correlation, and a global group velocity estimate was then computed from the slope of a linear regression of arrival time versus lateral position. The quality of this fit was quantified by the coefficient of determination R^2. For each excitation condition, the distribution of c_g and R^2 values over the 20 axial windows was used to summarize robustness and repeatability.

Finally, single-frequency phase-velocity maps were reconstructed over the 300–400 Hz band, and broadband maps were obtained by averaging the spatial phase velocity map $c_{\mathrm{ph}}(x, z, \omega_0)$ across this band. The resulting maps were used to derive summary statistics of phase velocity in the homogeneous phantom background, excluding the push-beam region where phase-velocity estimation is known to be unreliable [Kijanka and Urban, 2021].

5.4 Results in homogeneous media

5.4.1 Numerical and experimental validation

Figures 5.4(a–d) summarize the temporal behavior of the excitation and response for each configuration. For each excitation type, the transmitted code $s(t)$, the measured shear-wave response $y(t)$, and the recovered impulse response $h(t)$ are shown at a lateral position 15 mm away from the ARF focus. Solid lines correspond to the simulated data and dotted lines to the experimental measurements. Visual inspection reveals a close agreement between

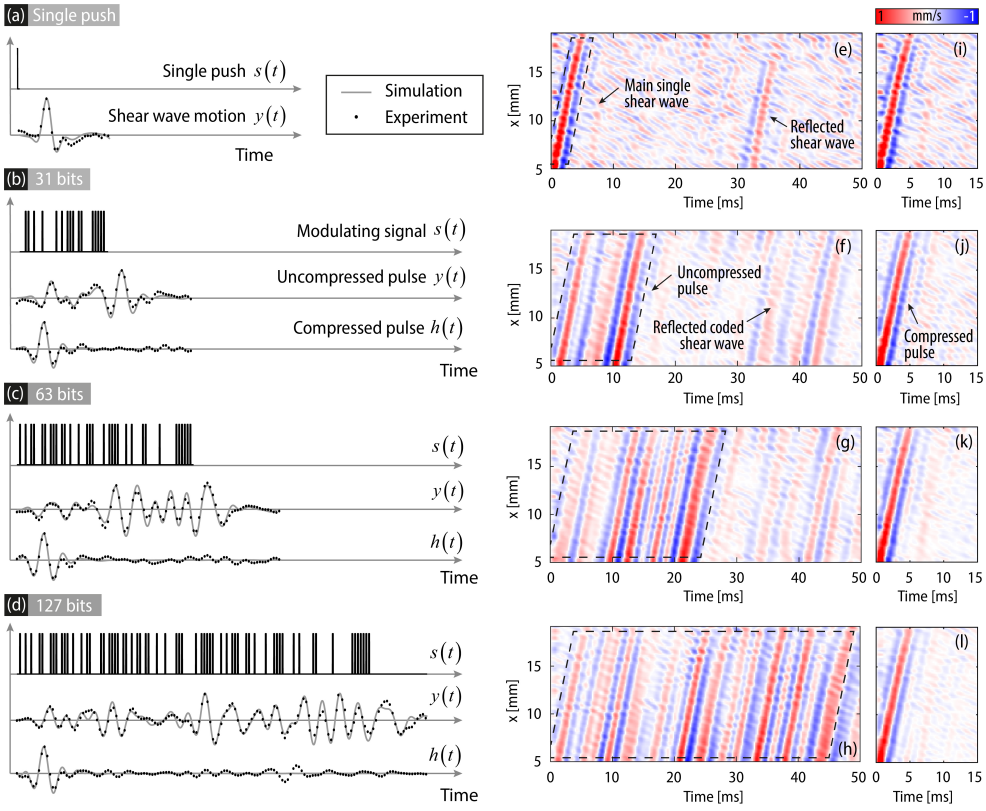


Figure 5.4: Comparison of single-push and MLS-coded ARF excitations in homogeneous phantoms and matched k-Wave simulations. (a) Single-push excitation: input waveform $s(t)$ and resulting shear-wave motion $y(t)$ at a lateral position 15 mm from the focus. (b–d) MLS-coded excitations with 31, 63, and 127 bits: modulating code $s(t)$, uncompressed shear-wave response $y(t)$, and compressed impulse response $h(t)$. Solid lines correspond to simulations and dots to experimental data. (e–h) Lateral–time maps $y(x,t)$ before compression, showing the extended coded shear-wave motion and the main travelling shear wave (dashed boxes). (i–l) Lateral–time motion $h(x,t)$ after frequency-domain deconvolution.

simulation and experiment for both the coded responses and their compressed counterparts. This agreement was quantified using the normalized root-mean-square error (NRMSE) computed on amplitude-normalized waveforms, with 0 indicating a perfect match. The resulting NRMSE values were 0.09 for SP, 0.11 for 31 bits, 0.15 for 63 bits, and 0.23 for 127 bits. All configurations thus exhibit error norms well below 0.25, confirming that the numerical model captures the experimental waveforms with high fidelity. The slightly larger

discrepancies observed for the longest code are consistent with the longer propagation distance and increased sensitivity to attenuation and experimental imperfections, but remain within acceptable bounds for quantitative SWE.

The effect of MLS coding on the duration and morphology of the propagating shear waves is illustrated in the lateral–time maps of Figs. 5.4(e–h). Compared with the SP case, MLS excitations generate elongated wavefields whose temporal extent scales with code length: approximately 11 ms for 31 bits, 22 ms for 63 bits, and 45 ms for 127 bits. With $T_{\mathrm{push}} = 100\ \mu$s, the duty cycle defined in (5.7) is 14.75%, 14.51%, and 14.40% for the 31-, 63-, and 127-bit sequences, respectively, indicating that the fraction of time during which ARF is active changes only marginally with code length.

After frequency-domain deconvolution, the reconstructed impulse responses $h(t)$ [Figs. 5.4(i–l)] are temporally compact and closely match the SP reference [Fig. 5.4(i)] despite the very different coded durations. In practice, the long MLS excitations are collapsed into short, high-SNR shear-wave impulses, so that the effective push remains narrow in time. In the SP case, no deconvolution is applied and the measured response matches the impulse response, $y(t) = h(t)$, providing a direct benchmark for the compressed waveforms.

5.4.2 SNR enhancement

The effect of MLS coding on shear-wave visibility was first examined qualitatively from space–time maps of the compressed particle velocity field. Figure 5.5 shows representative examples for the 45 kPa phantom, for four excitation voltages (30, 40, 50 and 60 V) and four excitation schemes (SP and MLS with 31, 63, and 127 bits). Each panel displays the compressed shear-wave motion in the (x, t) plane, together with the SNR value computed in the analysis window indicated by the dashed line. At 30 V, the SP configuration yields a weak and highly corrupted wavefront (SNR ≈ -1.5 dB), whereas MLS-coded excitations progressively suppress background noise and enhances the propagating wavefront signal, with SNR increasing to 2.1, 3.3, and 4.1 dB for 31-, 63-, and 127-bit codes, respectively. As the voltage is raised to 40–60 V, the SP wavefront becomes visible, but MLS still provides a consistent improvement of 1–3 dB, with the 127-bit configuration reaching SNR values around 6.5 dB. These maps illustrate the two main effects of coding: at low voltages, MLS converts barely detectable shear waves into clearly discernible fronts; at higher voltages, it further enhances contrast and suppresses residual clutter around

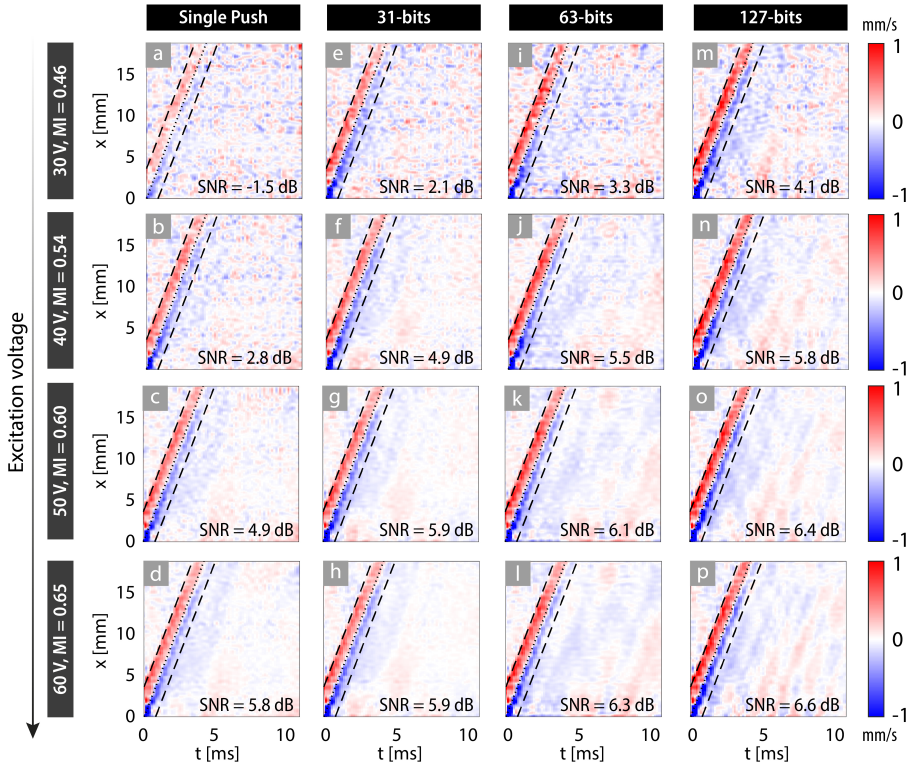


Figure 5.5: Effect of MLS coding on shear-wave SNR in the 45 kPa phantom. Lateral–time maps of the compressed shear-wave velocity are shown for SP and MLS excitations with 31, 63, and 127 bits at four excitation voltages (30, 40, 50, and 60 V). Dashed lines indicate the analysis window used for SNR computation, and the corresponding SNR value is reported in each panel.

the main wave.

The statistical behavior of this SNR gain across phantoms, code lengths, and voltages is summarized in Fig. 5.6. For each phantom stiffness (10, 25, 45, and 80 kPa), violin plots show the distribution of SNR gain (in dB) relative to the SP reference at the same voltage, pooling the 20 axial windows and repeated acquisitions. Across all phantoms and operating conditions, the median SNR gain is positive for every MLS configuration, confirming that coding does not merely re-shape the wavefield but systematically increases displacement SNR. The largest gains are observed at low voltages (20–30 V), where SP is closest to the noise floor: in the softest phantom (10 kPa), the

127-bit code reaches median gains of 6–7 dB, while 31- and 63-bit codes already achieve improvements of 3–5 dB. As stiffness increases, the absolute gain tends to decrease slightly, reflecting stronger attenuation and higher propagation speeds, but the qualitative trend remains the same.

A noteworthy feature of Fig. 5.6 is the behavior of intermediate code lengths. In all phantoms, 31 and 63 bits recover a substantial fraction of the SNR improvement obtained with 127 bits, typically between one half and two thirds of the maximum gain, while requiring much shorter excitation durations. This indicates a regime of diminishing returns with increasing MLS length: longer codes always provide additional SNR, but with progressively smaller incremental benefit. From a practical viewpoint, the code length can therefore be selected as a design parameter that trades acquisition time for SNR. In shallow or moderately attenuating regions, 31- or 63-bit excitations may provide sufficient enhancement with minimal temporal overhead, whereas deeper targets or highly attenuating tissues may justify the use of 127-bit codes to maximize robustness. Overall, these results show that MLS-coded ARF offers a flexible and monotonic mechanism to boost shear-wave SNR without modifying peak acoustic output, which is later controlled independently through the safety constraints discussed in Section 5.5.

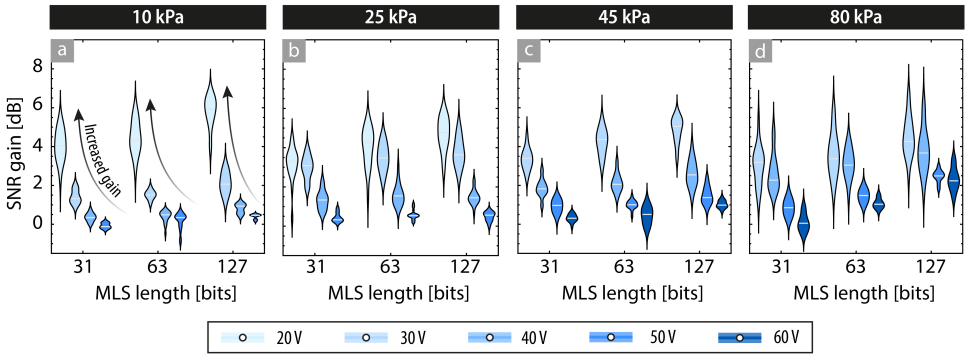


Figure 5.6: SNR gain (in dB) achieved by MLS-coded excitations relative to the SP reference for four homogeneous phantoms (10, 25, 45, and 80 kPa). For each phantom, violin plots show the distribution of SNR gain across axial windows and repeated acquisitions as a function of MLS length (31, 63, and 127 bits) and selected excitation voltages within the 20–60 V range. Only voltage conditions that produced a detectable shear-wave response and allowed reliable SNR estimation were included. Consequently, not all phantoms contain the five possible voltage levels: low-voltage acquisitions were excluded in the stiffer phantoms when the propagated shear wave was too close to the noise floor.

5.4.3 Group-velocity estimation

The influence of MLS-coded excitation on shear-wave group velocity was assessed in all four homogeneous phantoms using the cross-correlation and linear-regression procedure described in Section 5.3.3. For each configuration, arrival times were estimated from lateral profiles of the compressed displacement field and group velocity c_g was obtained from the slope of a linear fit, with the coefficient of determination R^2 used as an indicator of propagation consistency.

Figures 5.7(a–d) show representative arrival-time curves for SP excitation in the 10, 25, 45, and 80 kPa phantoms. The scatter of the points and the relatively low R^2 values, particularly in the stiffer phantoms, reflect the difficulty of tracking the wavefront under low-SNR conditions. Figures 5.7(e–h) display the corresponding fits for 127-bit MLS-coded excitation. In all cases, the arrival times align much more closely with a straight line and R^2 approaches unity. For example, in the 25 kPa phantom the coefficient of determination increased from 0.50 (SP) to 1.00 (MLS-127), and in the 80 kPa phantom from 0.01 to 0.97, indicating that coding greatly stabilizes the cross-correlation peak and yields a more coherent apparent wavefront.

The statistical behavior of the group-velocity estimates is summarized in Figs. 5.7(i–l), which present violin plots of c_g over the 20 axial windows for each phantom and excitation scheme. In the 45 kPa phantom, SP produced a mean velocity of 3.93 m/s but with very large dispersion (standard deviation 24.89 m/s, IQR 3.84 m/s), whereas MLS-coded excitations yielded tightly clustered distributions around the expected shear-wave speed: 3.92 m/s (σ = 0.12 m/s, IQR 0.30 m/s) for 31 bits, 3.91 m/s (σ = 0.08 m/s, IQR 0.12 m/s) for 63 bits, and 3.90 m/s (σ = 0.06 m/s, IQR 0.07 m/s) for 127 bits. In the stiffest phantom (80 kPa), the interquartile range was reduced from 15.16 m/s for SP to 0.14 m/s for MLS-127, illustrating how coding converts a practically unusable estimate into a narrow, well-defined distribution.

Across all stiffness levels, MLS-coded excitation preserves the expected mean group velocity while markedly reducing variability. Intermediate code lengths already provide substantial improvements: 31- and 63-bit excitations recover most of the precision gain of 127 bits, with progressively narrower distributions as the code length increases. Thus, MLS coding not only boosts displacement SNR but also translates this gain into more reliable and repeatable group-velocity estimates, particularly in the most challenging, high-stiffness

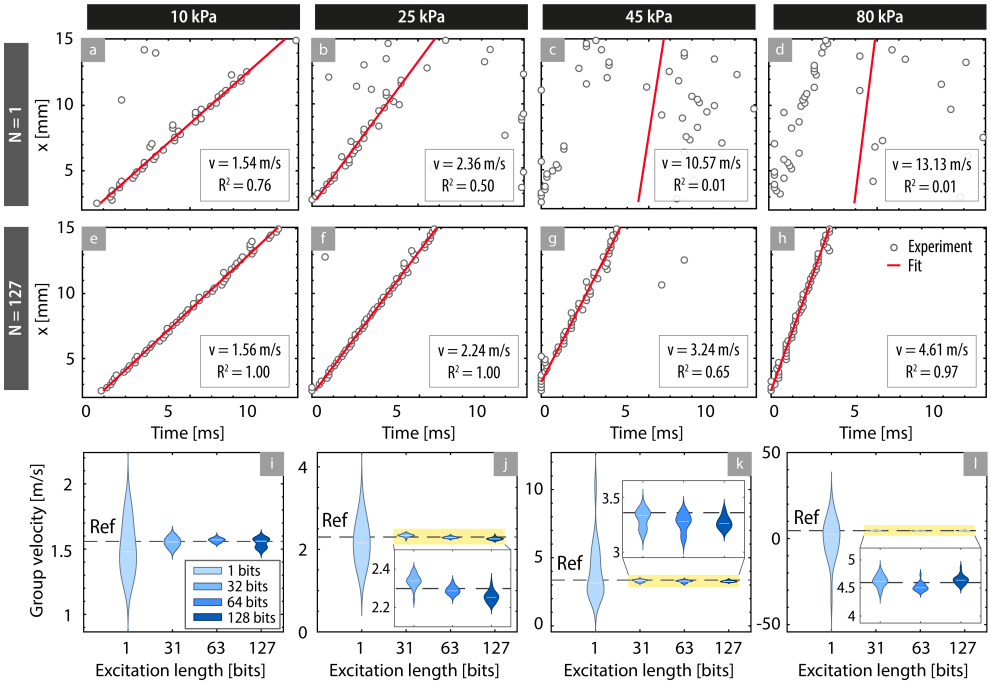


Figure 5.7: Group-velocity estimation across homogeneous phantoms for SP and MLS-coded excitations. (a–d) Example arrival-time curves $t_{\mathrm{arr}}(x)$ and linear fits for SP excitation in the 10, 25, 45, and 80 kPa phantoms, with the corresponding coefficient of determination R^2. (e–h) Equivalent results for 127-bit MLS-coded excitation. (i–l) Violin plots of group-velocity estimates c_g obtained from 20 axial windows for each phantom, comparing SP with MLS codes of 31, 63, and 127 bits.

phantoms where SP alone often fails to provide robust measurements due to rapid shear-wave attenuation.

5.4.4 Phase-velocity mapping

The impact of MLS-coded excitation on phase-velocity estimation was assessed in the homogeneous 45 kPa phantom. Phase-velocity maps were reconstructed over the 300–400 Hz band for a conventional single-push (1-bit) excitation and for a 127-bit MLS sequence operated at the same mechanical index, and compared with a high mechanical index single-push reference.

Figures 5.8(a) and 5.8(b) show broadband phase-velocity maps for the

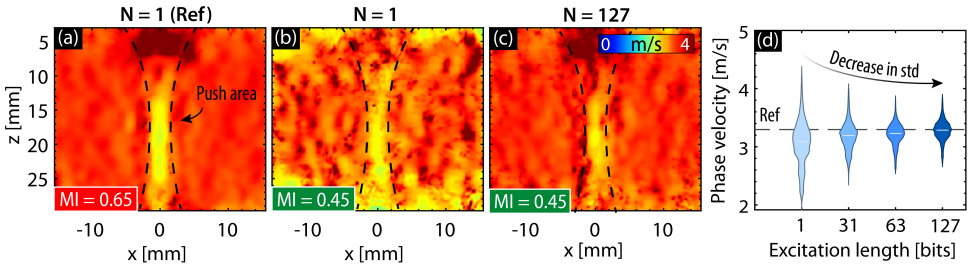


Figure 5.8: PGE-based phase-velocity estimation of the homogeneous 45 kPa phantom. (a) Phase-velocity map for a 1-bit SP excitation at MI = 0.45 (30 V). (b) Phase-velocity map for a 127-bit MLS excitation at the same MI. (c) SP reference acquired at higher drive level (MI = 0.65, 60 V), used as a benchmark; the dashed contour in all maps indicates the push-beam region excluded from the analysis. (d) Violin plots of the mean background phase velocity for low-MI SP and MLS excitations of different lengths. The dashed line marks the mean phase velocity of the high-MI reference.

1-bit SP and 127-bit MLS excitations at MI = 0.45 (30 V). In the SP case, the background exhibits pronounced spatial fluctuations and a mottled appearance, indicative of limited SNR in the underlying displacement field. When the excitation is encoded with a 127-bit MLS at the same MI, the phase-velocity map becomes visibly smoother and more homogeneous, with reduced small-scale variability. Figure 5.8(c) displays the SP reference acquired at higher drive level (MI = 0.65, 60 V), which yields a similarly uniform background. Thus, for the same MI, the long MLS code recovers a phase-velocity map whose visual quality approaches that of a conventional single push operated at substantially higher MI.

The corresponding quantitative behavior is depicted in Fig. 5.8(d), which presents violin plots of background phase velocity for the low-MI SP and for MLS excitations of 31, 63, and 127 bits. All values are averaged over the 300–400 Hz band, and the dashed horizontal line marks the mean phase velocity obtained with the high-MI SP reference. The analysis focused on the background region, excluding the push-beam area where the local wavefield is strongly disturbed and PGE estimates are less reliable [Kijanka and Urban, 2021]. At MI = 0.45, the single push excitation produces a broad distribution with noticeable scatter around the reference value. As code length increases, the distributions progressively narrow and their medians converge toward the reference mean, with the 127-bit configuration showing the smallest spread and the closest agreement with the high-MI case. In other words, MLS coding at fixed MI primarily acts to reduce variability and residual bias in phase

speed.

Overall, these results indicate that MLS-coded ARF and PGE form an effective combination for quantitative phase-velocity mapping in homogeneous media. By increasing displacement SNR at a given MI, long MLS codes enable phase-velocity maps that are comparable in uniformity and stability to those obtained with higher-MI single pushes, while keeping peak mechanical output unchanged. This property is particularly relevant in contexts where acoustic safety limits constrain MI, but robust and spatially homogeneous phase-velocity estimates are still required for stiffness quantification or viscoelastic model fitting.

5.5 Acoustic safety characterization

The use of long-duration coded excitations in MLS-based ARF naturally raises questions regarding acoustic safety. This section quantifies the mechanical index (MI), the spatial-peak pulse-average intensity (I_{SPPA}), and the spatial-peak time-average intensity (I_{SPTA}) for the configurations used in this work, and compares them against the regulatory thresholds defined for general diagnostic imaging by the U.S. Food and Drug Administration (FDA) [Food and Drug Administration, 2021]. Unless otherwise stated, the reference limits considered are MI ≤ 1.9, $I_{SPPA} \leq 190$ W/cm^2, and $I_{SPTA} \leq 720$ mW/cm^2.

Free-field pressure measurements were acquired in a degassed water tank using a calibrated membrane hydrophone (HMB-0500, Onda Inc., Sunnyvale, CA, USA; 500 µm sensitive element) positioned at different focus positions in depth (10, 20, 30 mm) of the L11-4v probe with an active aperture equal to a f-number of 2. For each input voltage, the hydrophone recorded the focal pressure waveform produced by a representative 32-cycle burst at the carrier frequency $f_0 = 4.5$ MHz. Short bursts were used instead of full-duration MLS pushes to avoid hydrophone damage and excessive thermal loading, in line with standard practice for characterizing high-duty-cycle or therapeutic-like excitations [Herman and Harris, 2002]. The acoustic output characterization and hydrophone measurement procedure followed the Acoustic Output Measurement Standard for Diagnostic Ultrasound Equipment [AIUM-NEMA, 2007].

Figure 5.9 provides a schematic view of the temporal intensity pattern associated with an MLS-coded ARF excitation. Each blue bar represents an

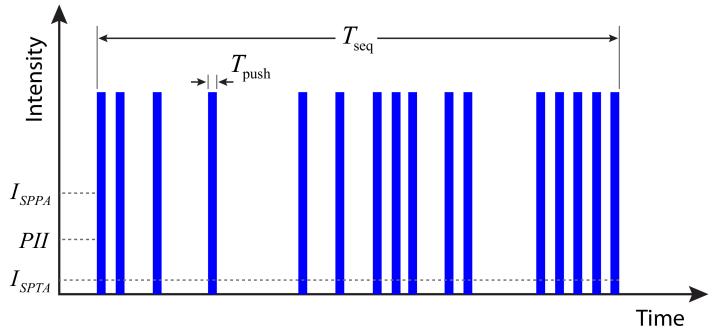


Figure 5.9: Schematic representation of the temporal intensity pattern for an MLS-coded ARF excitation. Each blue bar denotes a high-intensity push of duration T_{push} within a sequence of total duration T_{seq}. The $\mathrm{I_{SPPA}}$ corresponds to the average intensity during one push, the PII is proportional to the area under a single bar, and the $\mathrm{I_{SPTA}}$ represents the average intensity over the entire sequence, including silent intervals.

intensity push of duration T_{push} embedded within the full sequence of length T_{seq}. The spatial-peak pulse-average intensity $\mathrm{I_{SPPA}}$ corresponds to the average intensity during one push, the pulse intensity integral (PII) is proportional to the energy deposited per push and equals the area under one bar, and $\mathrm{I_{SPTA}}$ represents the long-term average over the entire sequence, including both active and silent intervals. The diagram also clarifies the different roles of peak pressure and temporal coding: $\mathrm{I_{SPPA}}$ is governed by the pulse amplitude, MI by the peak rarefactional pressure, and $\mathrm{I_{SPTA}}$ by the combined effect of pulse intensity, duty cycle, and sequence repetition rate.

5.5.1 Mechanical index

The mechanical index was computed as

$$\mathrm{MI} = \frac{|p_{\mathrm{r}}|}{\sqrt{f_0}}, \tag{5.10}$$

where p_{r} is the derated peak rarefactional pressure in MPa and f_0 is the center frequency in MHz [Abbott, 1999]. For each input voltage, MI was evaluated at three focal depths (10, 20, and 30 mm) by axially scanning the hydrophone around the nominal focus.

Figure 5.10(a) shows MI as a function of input voltage for the three focal depths. As expected, MI increases monotonically with voltage and is higher

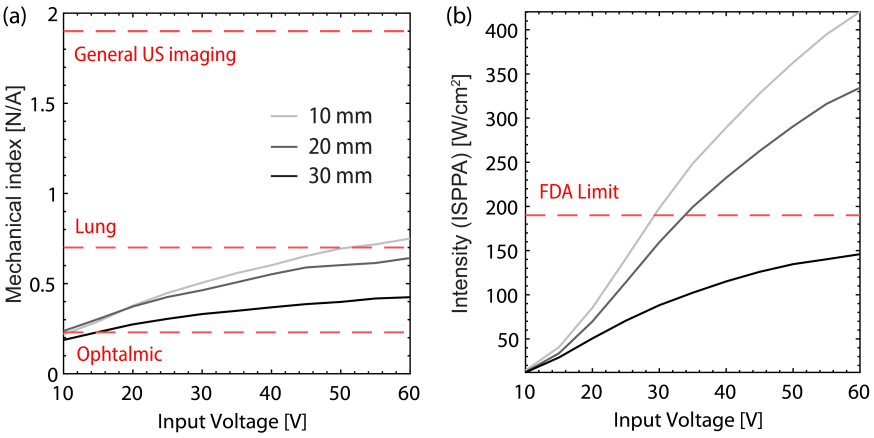


Figure 5.10: (a) Mechanical index (MI) as a function of input voltage for three focal depths (10, 20, and 30 mm). Horizontal dashed lines indicate FDA thresholds for general ultrasound imaging (MI = 1.9), lung imaging (MI = 0.7), and ophthalmic imaging (MI = 0.23). (b) Spatial-peak pulse-average intensity as a function of input voltage for the same focal depths. The horizontal dashed line indicates the FDA limit for derated I_{SPPA} (190 W/cm^2).

at shallower depths, where the focal pressure is larger. The maximum MI measured across all configurations was 0.75 at 60 V and 10 mm focal depth, well below the general FDA limit of 1.9 for diagnostic imaging. The figure also displays the modality-specific thresholds for lung elastography (MI = 0.7) and ophthalmic imaging (MI = 0.23). Depending on focus depth, the measured MI values remain within these stricter limits over the lower voltage range, indicating that the proposed MLS-coded ARF protocol can be operated safely even in more restrictive clinical scenarios, provided that appropriate voltage levels and focal configurations are selected.

5.5.2 Spatial-peak pulse-average intensity

To connect the short-burst measurements with the energy delivered by the longer ARF pushes embedded in the MLS, the pulse intensity integral (PII) was first evaluated for the 32-cycle reference burst according to the standard [AIUM-NEMA, 2007]:

$$\mathrm{PII} = \frac{1}{\rho_0 c_0} \int_0^{T_{\mathrm{pulse}}} p^2(t)\,\mathrm{d}t, \tag{5.11}$$

where ρ_0 and c_0 are the density and sound speed in water, $p(t)$ is the recorded focal pressure waveform, and T_{pulse} is the duration of the 32-cycle burst ($T_{\text{pulse}} = 32/f_0 \approx 7.11$ μs at 4.5 MHz). This quantity has units of energy per unit area (J/cm^2). The corresponding spatial-peak pulse-average intensity was then obtained as

$$I_{\text{SPPA}} = \frac{\text{PII}}{T_{\text{pulse}}}. \tag{5.12}$$

Under the assumption that the ARF pushes used within the MLS sequence are quasi-stationary sinusoidal bursts with the same focal amplitude as the 32-cycle reference, I_{SPPA} depends only on the measured PII and is independent of the actual push duration T_{push}, as illustrated schematically in Fig. 5.9.

Figure 5.10(b) depicts I_{SPPA} as a function of input voltage for different focal depths. As expected, I_{SPPA} grows approximately quadratically with voltage, reflecting the dependence of acoustic intensity on pressure amplitude. For input voltages above 30 V, I_{SPPA} exceeds the FDA upper limit of 190 W/cm^2. However, as discussed by Herman and Harris [Herman and Harris, 2002], regulatory compliance can be established by meeting either the MI or the I_{SPPA} limit, but not necessarily both simultaneously. Given that MI remains below its threshold for all investigated configurations (see Fig. 5.10(a)), the MLS-coded ARF method remains compliant with FDA regulations even in regimes where I_{SPPA} exceeds 190 W/cm^2. In this sense, MI becomes the primary safety constraint, while I_{SPPA} provides complementary information about local peak power levels.

5.5.3 Spatial-peak time-average intensity

Whereas MI and I_{SPPA} are governed by peak pressure and pulse-averaged intensity, respectively, the spatial-peak time-average intensity I_{SPTA} quantifies the average acoustic power delivered over the full repetition period and is therefore the most sensitive metric to long coded excitations and high repetition rates. For the MLS-coded ARF protocol, I_{SPTA} can be expressed as

$$I_{\text{SPTA}} = I_{\text{SPPA}}\, D\, T_{\text{seq}}\, \text{PRF}, \tag{5.13}$$

where D is the duty cycle of the ARF pushes within a single MLS sequence, defined in Eq. (5.7), T_{seq} is the total sequence duration, and PRF is the repetition frequency of the full SWE acquisition. Equivalently, since $DT_{\text{seq}} =$

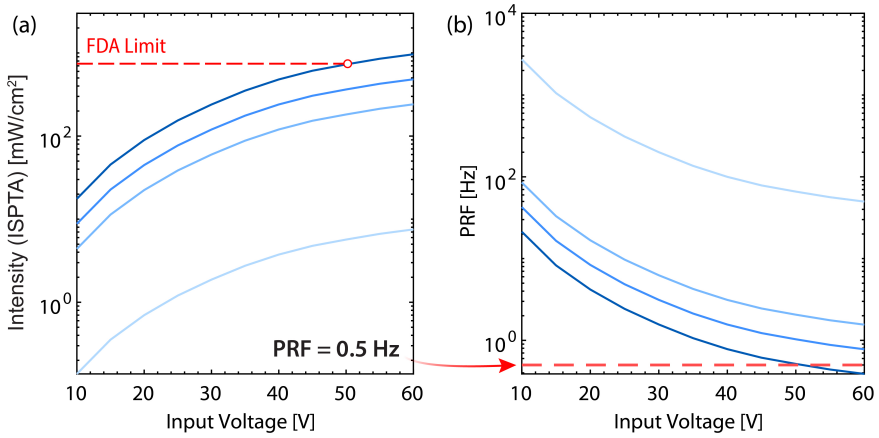


Figure 5.11: Time-averaged safety metrics for MLS-coded ARF excitations. (a) Spatial-peak time-average intensity (I_{SPTA}) as a function of input voltage for SP and MLS-coded excitations at a fixed PRF of 0.5 Hz. The horizontal dashed line indicates the FDA limit of 720 mW/cm^2. (b) Maximum allowable repetition frequency PRF_{max} that ensures I_{SPTA} remains below the FDA threshold, plotted for different MLS lengths and input voltages.

$N_1 T_{push}$, the above expression can be written as

$$I_{SPTA} = I_{SPPA}\, N_1 T_{push}\, PRF. \tag{5.14}$$

Figure 5.11(a) shows I_{SPTA} as a function of input voltage for SP and MLS-coded excitations at a fixed PRF of 0.5 Hz. Due to the longer excitation duration, MLS configurations exhibit higher I_{SPTA} than the SP reference at the same voltage, and I_{SPTA} increases with both voltage and sequence length. For the longest code (127 bits), the 0.5 Hz protocol exceeds the 720 mW/cm^2 threshold at 50 V, indicating that time-averaged exposure becomes the most restrictive constraint when simultaneously using long codes and high drive levels.

To translate these results into practical operating guidelines, the maximum allowable PRF ensuring I_{SPTA} compliance was also computed. Solving (5.14) for PRF yields

$$PRF_{max} = \frac{{I_{SPTA}}^{(lim)}}{I_{SPPA}\, N_1 T_{push}}. \tag{5.15}$$

where ${I_{SPTA}}^{(lim)} = 720$ mW/cm^2 is the FDA limit for general imaging. Figure 5.11(b) displays PRF_{max} as a function of input voltage and MLS length.

As expected, higher voltages and longer codes require reduced PRF to keep I_{SPTA} below the regulatory threshold. For example, at 60 V and 127 bits, PRF_{max} must be limited to approximately 0.4 Hz to satisfy the I_{SPTA} constraint, whereas shorter codes can operate at higher repetition rates under the same voltage.

Overall, the safety analysis indicates that MLS-coded excitations can be operated within standard diagnostic limits, but it also highlights the trade-off between SNR gains and acoustic exposure. MI remains comfortably below the general and modality-specific thresholds across all tested configurations, and I_{SPPA} exceeds its nominal limit only at higher voltages where MI still guarantees regulatory compliance. In contrast, I_{SPTA} becomes the dominant constraint for long MLS codes and high voltages, requiring a reduction in PRF to maintain safe time-averaged exposure. These considerations are essential when translating the proposed pulse-compression strategy to clinical protocols that demand both high SNR and strict control of acoustic output.

5.6 Discussion

The results in this chapter indicate that embedding an MLS directly into the ARF excitation and recovering the impulse response by pulse compression is an effective strategy to enhance shear-wave elastography. Matched simulations and phantom experiments in homogeneous media confirmed that the coding–compression method behaves as intended: long coded pushes generate extended shear-wave fields that can be collapsed into compact, SP-like impulse responses with high fidelity (Fig. 5.4). The close agreement between simulated and experimental waveforms, together with the small NRMSE values obtained for all code lengths, supports both the validity of the MLS-coded ARF formulation and the consistency of the underlying shear-wave modeling method introduced in Chapter 2.

A central outcome is the monotonic increase in displacement SNR with MLS length, together with the emergence of a regime of diminishing returns. Lateral–time maps and violin plots of SNR gain in homogeneous phantoms (Figs. 5.5 and 5.6) show that long codes such as 127 bits can recover shear waves that are barely visible under low-voltage SP excitation, especially in the softest phantom where median SNR gains reach 6–7 dB at 20–30 V. At the same time, intermediate code lengths (31 and 63 bits) systematically recover

one half to two thirds of the maximum gain while requiring substantially shorter excitation durations. This behavior is consistent with the near-ideal autocorrelation of MLS and with the expectation that longer codes distribute energy over a larger number of pushes. From a design perspective, code length therefore acts as a tunable parameter that trades acquisition time and motion stationarity for SNR. In practice, this temporal penalty is not a major limitation for SWE, because complete shear-wave acquisitions are typically repeated at low rates, often from sub-hertz to a few hertz, whereas only the motion tracking within each acquisition operates at kilohertz frame rates. Under these conditions, the longer duration of MLS coding remains compatible with usual elastographic operation.

The analysis of group and phase velocities shows how these SNR gains translate into quantitative stiffness estimates. In all homogeneous phantoms, MLS coding preserved the expected mean group velocity while markedly increasing the coherence of arrival-time curves and the coefficient of determination of linear fits (Fig. 5.7). Similarly, in the 45 kPa phantom the combination of MLS compression and PGE produced phase-velocity maps whose spatial homogeneity and statistical dispersion approached those of a high-MI SP reference, even when operating at substantially lower MI (Fig. 5.8). These observations indicate that MLS-coded excitation primarily improves precision and robustness, rather than altering the underlying mechanical response. In practical terms, long MLS codes operated at low MI can reproduce the group- and phase-velocity estimates obtained with high-voltage single pushes, effectively exchanging excitation duration for SNR without increasing peak mechanical output.

Acoustic safety introduces an additional design dimension that is particularly relevant for clinical translation. The hydrophone-based characterization performed in Section 5.5 showed that, for the explored configurations, MI remains below the general diagnostic limit across voltages and code lengths, and that I_{SPPA} only approaches its regulatory bound at the highest drive levels. In contrast, I_{SPTA} emerges as the dominant constraint when long MLS codes are combined with high repetition frequencies. The analysis of I_{SPTA} and of the corresponding maximum allowable PRF (Fig. 5.11) makes explicit how increasing voltage or code length requires a proportional reduction in PRF to maintain safe time-averaged exposure. Importantly, the phase-velocity results demonstrate that comparable quantitative performance can be achieved either with high-MI SP or with long MLS codes at lower MI and PRF. This property

is especially attractive in scenarios where MI is tightly constrained—such as ophthalmic imaging, lung elastography, or neonatal applications—because it enables robust shear-wave measurements while respecting strict peak-pressure limits.

From an implementation standpoint, the MLS-coded ARF method can be interpreted as an excitation-layer modification that is largely transparent to downstream SWE processing. On the hardware side, the method relies on a single linear array operating in focused-transmit and plane-wave receive modes, with the only additional requirement being a bit-synchronous scheduling of ARF pushes and tracking events, as described in Section 5.2.3. Such timing control is readily available on research platforms and could, in principle, be embedded into clinical scanners as a dedicated SWE mode. On the processing side, the incremental computational cost is modest: frequency-domain deconvolution is dominated by FFTs, which are highly parallelizable and compatible with real-time implementation on modern CPUs/GPUs. Once the coded signals have been compressed, the same group-velocity, dispersion, and phase-velocity approaches used in conventional SWE can be applied without modification.

Two practical limitations of the current implementation are worth noting. First, the fixed bit period of 350 µs, chosen to accommodate push–track ring-down and avoid contamination of the tracking pulses by ARF reverberations, constrains the effective shear-wave tracking frame rate to below 3 kHz. This ceiling may be restrictive in applications that require very high temporal resolution, such as rapid dynamic imaging of transient deformations in arterial elastography [Roy et al., 2025], or tendon and other musculoskeletal applications where shear-wave propagation can be fast and strongly dependent on loading, contraction state, or joint position [Shinohara et al., 2010]. Future developments could explore more aggressive clutter- and noise-suppression strategies [Takagi et al., 2021] that preserve phase integrity while shortening the inter-bit delay, thereby increasing the usable tracking frame rate without compromising pulse compression.

Second, long MLS codes increase time-average acoustic output and load the transmit chain. Although peak pressure and MI are preserved, the extended ARF "on-time" raises I_{SPTA}, so that, for example, 127-bit sequences at 50 V can exceed the 720 mW/cm^2 FDA limit unless PRF is reduced, exposing a direct trade-off between SNR gains and permissible thermal dose. At the hardware level, the higher duty factor demands sufficient RMS current and

thermal headroom in the amplifier and probe; otherwise, protection circuitry or waveform distortion may erode the benefits of pulse compression. In practice, safe deployment of MLS-coded SWE requires verifying these margins for the specific system and, if needed, jointly derating PRF, drive voltage, and code length while keeping MI and I_{SPPA} within regulatory bounds.

5.7 Conclusions

This chapter has presented an MLS-coded ARF method for shear-wave elastography in which a binary pseudo-random sequence is embedded directly into the ARF excitation and later used for by frequency-domain pulse compression. Using matched phantom experiments and full-wave shear-wave simulations, it has been shown that long MLS-modulated pushes can be collapsed into compact, single-push like impulse responses that preserve the underlying shear-wave dynamics. In homogeneous phantoms spanning 10–80 kPa, increasing the MLS length led to monotonic SNR gains, more coherent group-velocity estimates with reduced dispersion, and PGE-based phase-velocity maps that converged towards high-MI reference values even under low-drive conditions.

A dedicated safety analysis demonstrated that these benefits can be obtained within standard diagnostic-output limits, provided that repetition frequency is adjusted to account for sequence length and drive voltage. At fixed MI, long MLS codes were able to reproduce the quantitative velocities of high-voltage SP excitations, thereby enabling high-SNR, quantitative SWE in regimes where MI must remain low. Overall, MLS-coded ARF emerges from this chapter as a practical and scalable strategy to boost the reliability of shear-wave measurements using a single clinical probe and standard SWE processing. The next chapter builds on this foundation to examine how the method performs in heterogeneous, viscoelastic, anisotropic, and dynamically perturbed media, with the aim of assessing its impact on more clinically realistic elastography scenarios.

Publications

This section contains a compendium of the seminal scientific articles in which the results presented in this chapter have been published.

Conferences articles

- González-Mateo, E., Rodríguez-Sendra, J., Camarena, F., and Jimenez, N., "Shear-Wave Pulse-Compression Elastography using a Modulated Acoustic Radiation Force Excitation," in *IEEE Ultrasonics, Ferroelectrics, and Frequency Control Joint Symposium (UFFC-JS 2024)*, Taipei, Taiwan, 2024.
- González-Mateo, E., Camarena, F., and Jimenez, N., "Improving SNR and bandwidth of shear-wave displacements via modulated force excitation," in *TECNIACÚSTICA 2023. 54º Congreso Español de Acústica. XIII Encuentro Ibérico de Acústica y el International Symposium on Acoustics in Biomedical Engineering*, Cuenca, España, 2023, pp. 896–899.

Awards

- Premio Andrés Lara para jóvenes investigadores, *Sociedad Española de Acústica*, 2023

References

Abbott, J. G. (1999). Rationale and derivation of mi and ti—a review. *Ultrasound in Medicine & Biology*, 25(3):431–441.

AIUM-NEMA (2007). *Acoustic Output Measurement Standard for Diagnostic Ultrasound Equipment*. American Institute of Ultrasound in Medicine and National Electrical Manufacturers Association.

Anstey, N. A. (1964). Correlation techniques -a review. *Geophysical Prospecting*, 12(4):355–382.

Bercoff, J., Tanter, M., and Fink, M. (2004). Supersonic shear imaging: a new technique for soft tissue elasticity mapping. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 51(4):396–409.

Bleakley, C. and Scaife, R. (1995). New formulas for predicting the accuracy of acoustical measurements made in noisy environments using the averaged m-sequence correlation technique. *J. Acoustical Society of America*, 97(2):1329–1332.

Bruce, M., Kolokythas, O., Ferraioli, G., Filice, C., and O'Donnell, M. (2017). Limitations and artifacts in shear-wave elastography of the liver. *Biomedical Engineering Letters*, 7(2):81–89.

Buracas, G. T. and Boynton, G. M. (2002). Efficient design of event-related fmri experiments using m-sequences. *NeuroImage*, 16(3 Pt 1):801–813.

Chiao, R. Y. and Hao, X. (2005). Coded excitation for diagnostic ultrasound: a system developer's perspective. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 52(2):160–170.

Deffieux, T., Montaldo, G., Tanter, M., and Fink, M. (2009). Shear wave spectroscopy for in vivo quantification of human soft tissues visco-elasticity. *IEEE Transactions on Medical Imaging*, 28(3):313–322.

Deng, Y., Palmeri, M. L., Rouze, N. C., Rosenzweig, S. J., Abdelmalek, M. F., and Nightingale, K. R. (2015). Analyzing the impact of increasing mechanical index and energy deposition on shear wave speed reconstruction in human liver. *Ultrasound Medicine & Biology*, 41(7):1948–1957.

Ding, J., Li, Y., Jiao, Y., Wang, N., and Cui, Y. (2025). A method of enhanced shear wave elastography based on chirp coded excitation. *Ultrasonics*, 149:107569.

Food and Drug Administration (2021). Marketing clearance of diagnostic ultrasound systems and transducers. U.S. Department of Health and Human Services, Center for Devices and Radiological Health.

Gemba, K. L., Vazquez, H. J., Fialkowski, J., Edelmann, G. F., Dzieciuch, M. A., and Hodgkiss, W. S. (2021). A performance comparison between m-sequences and linear frequency-modulated sweeps for the estimation of travel-time with a moving source. *The Journal of the Acoustical Society of America*, 150(4):2613–2623.

Gennisson, J.-L. (2018). Musculoskeletal applications of supersonic shear imaging. In *Ultrasound Elastography for Biomedical Applications and Medicine*, pages 534–544. John Wiley & Sons, Ltd.

Golomb, S. (1982). *Shift Register Sequences*. Aegean Park Press.

González-Mateo, E., Camarena, F., and Jiménez, N. (2025). Real-time ultrasound shear wave elastography using a local phase gradient. *Computer Methods and Programs in Biomedicine*, 260:108529.

He, X.-N. et al. (2017). Improved shear wave motion detection using coded excitation for transient elastography. *Scientific Reports*, 7(1):44483.

Herman, B. A. and Harris, G. R. (2002). Models and regulatory considerations for transient temperature rise during diagnostic ultrasound pulses. *Ultrasound Medicine & Biology*, 28(9):1217–1224.

Jeong, J. S., Cannata, J. M., and Shung, K. K. (2010). Adaptive HIFU noise cancellation for simultaneous therapy and imaging using an integrated HIFU/imaging transducer. *Physics in Medicine & Biology*, 55(7):1889.

Kijanka, P. and Urban, M. W. (2021). Local phase velocity based imaging of viscoelastic phantoms and tissues. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 68(3):389–405.

Klauder, J. R., Price, A. C., Darlington, S., and Albersheim, W. J. (1960). The theory and design of chirp radars. *Bell System Technical Journal*, 39(4):745–808.

Liu, J. and Insana, M. F. (2005). Coded pulse excitation for ultrasonic strain imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 52(2):231–240.

Loupas, T., Powers, J., and Gill, R. (1995). An axial velocity estimator for ultrasound blood flow imaging, based on a full evaluation of the doppler equation by means of a two-dimensional autocorrelation approach. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 42(4):672–688.

Nguyen, T.-M., Arnal, B., Song, S., Huang, Z., Wang, R. K., and O'Donnell, M. (2015). Shear wave elastography using amplitude-modulated acoustic radiation force and phase-sensitive optical coherence tomography. *J. Biomedical Optics*, 20(1):016001.

Nguyen, T.-M., Song, S., Arnal, B., Wong, E. Y., Huang, Z., Wang, R. K., and O'Donnell, M. (2014). Shear wave pulse compression for dynamic elastography using phase-sensitive optical coherence tomography. *Journal of Biomedical Optics*, 19(1):016013.

Nowak, R. D. and Veen, B. D. V. (1994). Random and pseudorandom inputs for volterra filter identification. *IEEE Transactions on Signal Processing*, 42(8):2124–2135.

Rife, D. D. and Vanderkooy, J. (1989). Transfer-function measurement with maximum-length sequences. *Journal of the Audio Engineering Society*, 37:419–444.

Roy, T., Lee, H.-K., Capron, C. B., Lopez-Jimenez, F., Hesley, G. K., Greenleaf, J. F., Urban, M. W., and Guddati, M. N. (2025). Estimation of in vivo human carotid artery elasticity using arterial dispersion ultrasound vibrometry. *Ultrasound in Medicine & Biology*, 51(2):250–261.

Sandrin, L., Tanter, M., Gennisson, J., Catheline, M., and Fink, M. (2003). Transient elastography: a new noninvasive method for assessment of hepatic fibrosis. *Ultrasound Medicine & Biology*, 29(12):1705–1713.

Shinohara, M., Sabra, K. G., Gennisson, J.-L., Fink, M., and Tanter, M. (2010). Real-time visualization of muscle stiffness distribution with ultrasound shear wave imaging during muscle contraction. *Muscle & Nerve*, 42(3):438–441.

Song, J. H., Yoo, Y., Song, T.-K., and Chang, J. H. (2013). Real-time monitoring of HIFU treatment using pulse inversion. *Physics in Medicine & Biology*, 58(15):5333. Publisher: IOP Publishing.

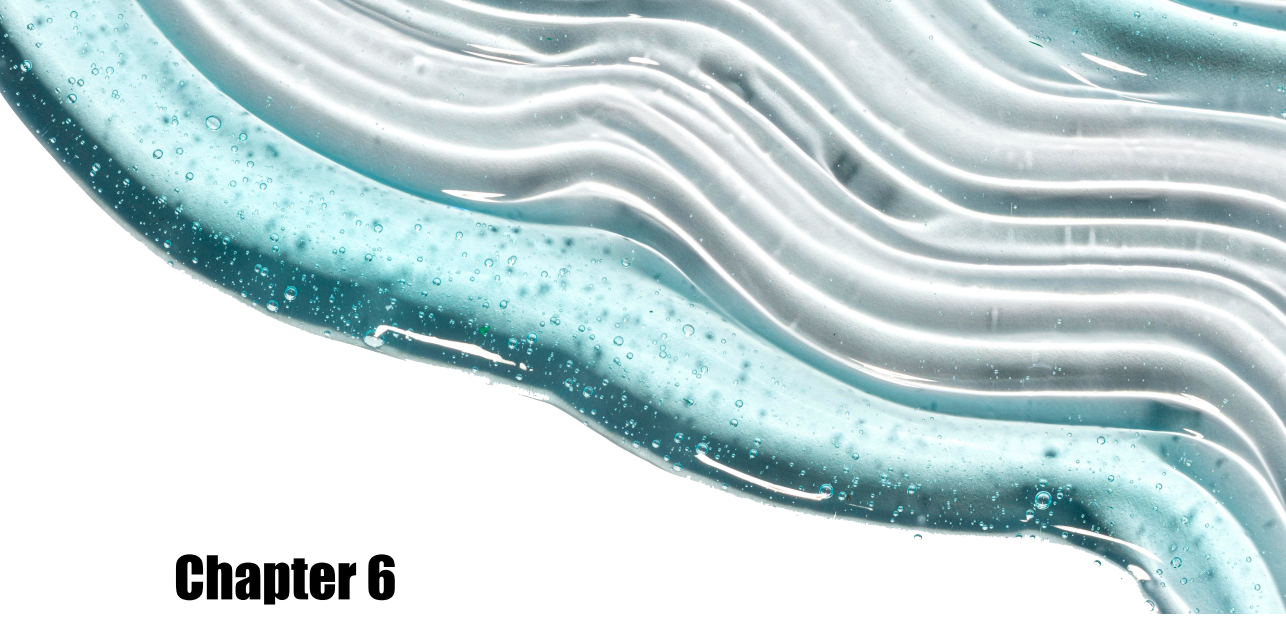
Song, P., Urban, M. W., Manduca, A., Greenleaf, J. F., and Chen, S. (2015). Coded excitation plane wave imaging for shear wave motion detection. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 62(7):1356–1372.

Takagi, R., Koseki, Y., Yoshizawa, S., and Umemura, S. I. (2021). Investigation of feasibility of noise suppression method for cavitation-enhanced high-intensity focused ultrasound treatment. *Ultrasonics*, 114:106394.

White, J. D. H. and Challis, R. E. (1994). Pseudo random binary sequence techniques for ultrasonic nde of high loss materials. In *IEEE Colloquium on Advanced Techniques for Collection and Interpretation of NDT Data*, pages 5/1–5/4.

Yücel, M. K. et al. (2015). Pulse-compression based iterative time-of-flight extraction of dispersed ultrasonic guided waves. In *2015 IEEE 13th International Conference on Industrial Informatics (INDIN)*, pages 809–815.

Zhou, B., Yang, X., Zhang, X., Curran, W. J., and Liu, T. (2020). Ultrasound elastography for lung disease assessment. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 67(11):2249–2257.



Chapter 6

Coded-MLS ARF in Challenging Scenarios

Abstract

This chapter evaluates MLS-Coded ARF pulse compression in application-relevant media selected to shear wave elastography under physically challenging conditions. The considered scenarios span heterogeneous propagation, anisotropic tissue, viscoelastic dispersion, pleural-adjacent imaging under mechanical-index constraints, and dynamic settings dominated by external or physiological motion. Using a shared acquisition and processing method, performance is assessed through shear-wave velocity metrics and phase-velocity mapping, focusing on robustness and sensitivity when conventional single-push excitations become SNR-limited or safety-limited. The results delineate operating regimes where MLS-Coded ARF improves measurement stability across realistic constraints encountered in practice.

Outline

6.1 Introduction 169
6.2 Shared experimental methods 170
6.2.1 MLS-Coded excitation settings 170
6.2.2 Reference acquisition 171
6.2.3 Motion estimation and post-processing 171
6.3 Case studies 172
6.3.1 Heterogeneous media mapping 172
6.3.2 Anisotropy assessment in skeletal muscle 175
6.3.3 Viscoelastic media analysis 180
6.3.4 Pleural interface tissue 183
6.3.5 Moving and vibrating media: robustness to external motion 188
6.3.6 Physiologically moving media: carotid artery 191
6.4 Discussion 196
6.5 Conclusions 201

6.1 Introduction

A central practical constraint motivating coded excitations is that improving shear-wave motion SNR by simply increasing peak acoustic output is not always desirable or feasible. In clinical SWE, shear-wave speed reconstruction can benefit from increased acoustic output; however, this strategy is intrinsically bounded by acoustic exposure limits and safety considerations and may be further restricted in sensitive or motion-prone examinations [Deng et al., 2015,Schoen et al., 2025]. In vascular SWE, additional constraints arise because arterial pulsatility and wall dynamics introduce endogenous motion that mixes with the ARF-induced shear-wave response, complicating the separation of both contributions during motion tracking and velocity estimation [Nordenfur et al., 2025, Maksuti et al., 2016]. Within this context, pulse-compression concepts provide an alternative route to increasing effective shear-wave motion SNR by distributing excitation energy over time and recovering an impulse-like response through matched processing [Nguyen et al., 2014, Song et al., 2015]. For Maximum Length Sequences (MLS), their correlation properties and practical implementation are well established in the broader coded-excitation literature [Golomb, 2017].

The application studies in this chapter are designed to probe complementary failure modes of conventional processing. First, heterogeneous phantoms with stiffness inclusions are used to assess phase-velocity mapping under spatially varying wavefields, where directional filtering and local spectral selection are typically required to suppress interference and improve estimation stability [Manduca et al., 2003, Kijanka and Urban, 2019]. Secondly, anisotropy is evaluated in both *ex vivo* pork loin and *in vivo* human biceps through along-fiber and across-fiber acquisition geometries, leveraging the sensitivity of shear-wave propagation to tissue microstructure and directionality [Deffieux et al., 2009]. Third, viscoelastic behavior is analyzed via frequency–phase-velocity spectra and dispersion-curve fitting in a viscoelastic phantom and *ex vivo* liver, where dispersion has been used to characterize soft-tissue rheology [Muller et al., 2009]. Fourth, robustness to mechanically induced motion is investigated using a controlled loudspeaker-based paradigm, motivated by the strong dependence of quantitative SWE variance on noise, decorrelation, and spatiotemporal clutter [Deffieux et al., 2012]. Finally, an ECG-triggered *in vivo* human carotid experiment is included to evaluate performance under physiologically gated conditions in which arterial dispersion and cycle-to-cycle variability can affect shear-wave velocity estimation [Roy et al., 2025, Maksuti

et al., 2016].

This chapter extends the MLS-coded ARF pulse-compression framework introduced in Chapter 5 toward application-oriented scenarios that more closely reflect the challenges encountered in practice. Whereas the previous chapter established the method under controlled validation conditions, the focus here shifts to situations in which SWE performance is compromised by structural heterogeneity, anisotropic propagation, viscoelastic dispersion, and contamination from externally induced or physiological motion. Throughout the chapter, acquisitions and processing follow the same core implementation, including motion estimation, deconvolution-based pulse compression, and subsequent shear-wave velocity analysis.

6.2 Shared experimental methods

This section summarizes the experimental instrumentation and the processing steps that were shared across all studies reported in this chapter. Experiment-specific details such as phantom composition, anatomical target, motion paradigm, and acquisition geometry, are provided in the corresponding sections.

6.2.1 MLS-Coded excitation settings

All experiments were performed using a Verasonics V1 research ultrasound system (Verasonics Inc., Kirkland, WA, USA) coupled to a linear-array transducer (L11-4v, Verasonics Inc., Kirkland, WA, USA) with 128 elements and a nominal center frequency of 7.24 MHz. ARF excitation was delivered using a focused transmit with an active aperture of 32 elements and a focal depth of 20 mm, resulting in an f-number of 2. Unless otherwise stated, the ARF bursts were transmitted at 4.5 MHz.

Coded excitations were implemented using maximum-length sequences (MLS), where each sequence bit had a duration of 350 µs (corresponding to a bit rate of ∼2.8 kHz). Positive bits ($s_i = 1$) triggered a 100 µs ARF burst, while zero bits ($s_i = 0$) produced no acoustic emission. MLS lengths of 31, 63, and 127 bits were evaluated, corresponding to excitation durations of approximately 11 ms, 22 ms, and 45 ms, respectively. For comparison, a 1-bit excitation (single push, SP) employing the same focal configuration was also acquired. The transmit voltage was adjusted in each experiment

to comply with standard safety metrics, including mechanical index and spatial-peak temporal-average intensity. A dedicated safety characterization of the MLS-coded ARF implementation with these parameters was reported in Section 5.5.

Shear-wave motion was tracked using ultrafast plane-wave imaging at 7.24 MHz. The acquisition pulse repetition frequency (PRF) was selected to match the MLS bit rate (2.8 kHz) so that displacement/velocity estimates were available at each bit interval. This timing ensured that the temporally extended coded response was sampled with uniform cadence during both the excitation and the subsequent propagation window.

6.2.2 Reference acquisition

A high-voltage SP acquisition was used as a reference condition to provide a robust benchmark wavefield for velocity estimation across all experimental scenarios. The reference excitation employed the same focal configuration (depth and aperture) as the coded excitations but with an increased transmit voltage (60 V, MI = 0.65). Only for the reference acquisition, shear-wave tracking was performed with Time-Aligned Plane Wave Compounding (TAPWC) [Capriotti et al., 2020] using three steering angles ($-3°$, $0°$, $+3°$), yielding an effective frame rate of 12.5 kHz. TAPWC was not used in the MLS-coded case because of the limited frame-rate acquisition due to the push beam noise. This high-output, high-frame-rate condition served as the baseline for assessing the fidelity of the estimated wavefronts, group-velocity consistency, and dispersion-related analyses.

6.2.3 Motion estimation and post-processing

Beamformed in-phase and quadrature (IQ) data were exported from the Verasonics system and processed offline. Particle-velocity fields were estimated using an autocorrelation-based approach [Kasai et al., 1985]. For MLS-coded acquisitions, the measured particle-velocity field corresponds to the convolution between the MLS-coded excitation and the impulse response. To recover an estimate of the impulse response, pulse compression was performed by deconvolving the motion field with the known MLS sequence, following the framework introduced in Chapter 5. Specifically, for each spatial location, the temporal waveform was matched-processed with the MLS code, producing a

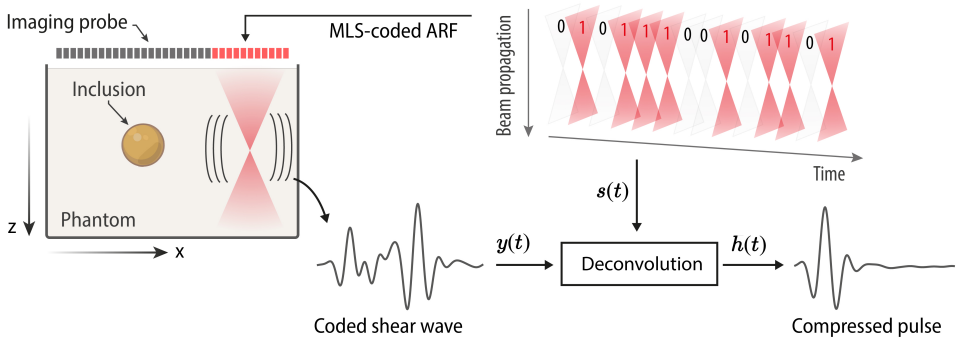


Figure 6.1: Heterogeneous phantom experiment and pulse-compression workflow. A circular inclusion was embedded in a homogeneous background and interrogated with MLS-coded ARF. The recorded coded shear-wave response $y(t)$ was deconvolved with the excitation sequence $s(t)$ to obtain a compressed impulse-like response $h(t)$ for subsequent velocity estimation.

compressed waveform that approximates the response to an impulse-like ARF push. Figure 6.1 represents the flowchart of the post-processing method.

To improve robustness in the presence of interference and reflections, a directional filter was first applied to isolate predominantly rightward-propagating components [Song et al., 2014]. Subsequently, a wavenumber-domain bandpass filter was used to retain velocity content consistent with typical shear-wave propagation in soft solids (1–10 m/s) and to suppress out-of-range components [Manduca et al., 2003]. The filtered spatiotemporal fields served as input for the experiment-specific estimators reported in the following sections, including time-of-flight group-velocity analysis, phase-velocity mapping, and frequency-dependent dispersion characterization.

6.3 Case studies

6.3.1 Heterogeneous media mapping

Spatial heterogeneity challenges shear-wave processing because stiffness gradients promote scattering, interference, and locally nonplanar wavefronts. Under these conditions, phase-based estimators can become unstable unless the wavefield is carefully isolated in the spectral domain by directional filtering and frequency selection [Kijanka and Urban, 2019]. This section evaluates the

performance of MLS-coded ARF combined with pulse compression for phase-velocity imaging in a standard inclusion-phantom paradigm (Figs. 6.1–6.3), while keeping the acoustic output at a moderate level.

Experiment protocol

A calibrated heterogeneous phantom containing circular inclusions of different stiffness embedded in a softer background was imaged. Three inclusion targets were considered (9, 36, and 70 kPa), selected to span low-, mid-, and high-contrast regimes relative to the background (18 kPa). For each target, acquisitions were performed using a conventional SP excitation and three MLS lengths (31, 63, and 127 bits) at the same acoustic output (30 V, MI = 0.45). In addition, a high-output and high-frame-rate SP acquisition was recorded as a reference condition (see Section 6.2.2). All excitations were placed laterally offset ($x = 12$ mm) from the inclusion to minimize direct push artifacts inside the inclusion region.

Metrics

The particle-velocity wavefield was first estimated and pulse compressed by deconvolution with the matched MLS code (Fig. 6.1). Phase-velocity maps were then reconstructed using the phase-gradient framework described in Chapter 3. Inclusion detectability in these maps was quantified through a contrast-to-noise ratio (CNR) computed between the inclusion ROI (dashed circle in Fig. 6.2) and a surrounding background ROI. This metric summarizes, in a single scalar value, the degree of inclusion–background separation while accounting for the spatial variability of the background.

$$\mathrm{CNR} = 20 \log_{10} \left(\frac{\left| \tilde{c}_{\mathrm{in}} - \tilde{c}_{\mathrm{bg}} \right|}{\sigma_{\mathrm{bg}}} \right), \tag{6.1}$$

where $\tilde{c}_{\mathrm{in}}$ and $\tilde{c}_{\mathrm{bg}}$ are the median phase velocities within the inclusion and background ROIs, respectively, and σ_{bg} is the standard deviation of phase velocity within the background ROI. Median estimates were used to mitigate the influence of local outliers and residual reconstruction artifacts, so that the resulting CNR (in dB) reflects both the effective contrast of the inclusion and the granularity of the surrounding map.

Results

Representative phase-velocity maps are shown in Fig. 6.2. Across all stiffness targets, increasing MLS length yielded visibly smoother maps, with a more spatially uniform inclusion region and reduced background granularity. These qualitative trends were accompanied by systematic CNR increases with code length for all inclusions.

For the 9 kPa inclusion, the SP reconstruction exhibited pronounced background fluctuations and a less uniform inclusion interior (CNR = 4.2 dB). With MLS excitation, CNR increased to 5.6 dB (31-bit), 7.2 dB (63-bit), and 7.8 dB (127-bit), approaching the 8.14 dB of the reference case. The corresponding maps show progressively cleaner inclusion boundaries and a more homogeneous inclusion ROI as MLS length increases.

For the 36 kPa inclusion, the inclusion was visible under all excitation types,

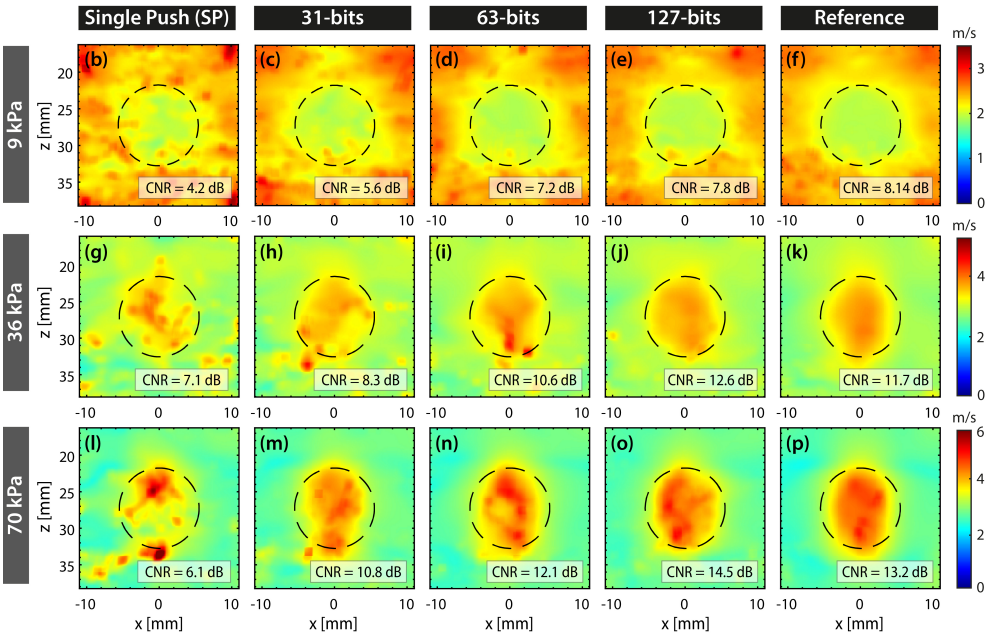


Figure 6.2: Phase-velocity maps in heterogeneous phantoms with inclusions of 9, 36, and 70 kPa. Columns compare SP, 31-, 63-, and 127-bit MLS excitations against a high-SNR reference. The dashed circle denotes the inclusion ROI used for CNR computation, reported within each panel. Color scales are adapted per stiffness level to preserve dynamic range.

although the reconstructed phase velocity maps differed in spatial smoothness and stability. The CNR increased from 7.1 dB for SP to 8.3 dB, 10.6 dB, and 12.6 dB for the 31-, 63-, and 127-bit MLS excitations, respectively, whereas the reference condition yielded a CNR of 11.7 dB.

For the 70 kPa inclusion, SP reconstructions presented stronger spatial variability and localized hot spots compared to longer coded excitations. CNR rose from 6.1 dB (SP) to 10.8 dB (31-bit), 12.1 dB (63-bit), and 14.5 dB (127-bit), while the reference condition provided 13.2 dB. In this stiff-inclusion case, longer MLS codes consistently yielded clearer inclusion boundaries and a more stable background.

The lateral phase-velocity profiles in Fig. 6.3 corroborate the map-level observations. In the 9 kPa case, the inclusion produced a velocity trough, and the profiles became progressively smoother and more symmetric as MLS length increased, converging toward the reference. In the 36 and 70 kPa cases, the inclusion produced a central velocity peak; increasing MLS length improved the profile smoothness and sharpened the transition at the inclusion boundaries, with the 127-bit MLS condition showing the closest agreement to the reference profile across the three stiffness targets.

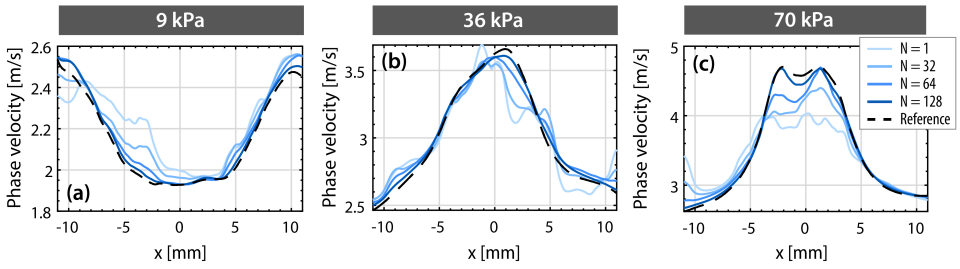


Figure 6.3: Lateral phase-velocity profiles across heterogeneous inclusions (9, 36, and 70 kPa). Profiles were extracted through the inclusion center and compared across excitation lengths and the reference measurement (dashed).

6.3.2 Anisotropy assessment in skeletal muscle

This section evaluates the MLS-coded ARF pulse-compression framework under direction-dependent propagation in skeletal muscle, where the measured shear-wave motion varies with the relative orientation between the imaging plane and the muscle fiber direction [Deffieux et al., 2009]. Two acquisition

geometries were considered: (i) probe alignment *along* the dominant fiber direction and (ii) probe alignment *across* the fibers. The experiment was performed in *ex vivo* porcine muscle and *in vivo* human biceps to assess repeatability across controlled and physiological conditions.

Experiment protocol

Figure 6.4 summarizes the acquisition scheme for both *ex vivo* and *in vivo* measurements. The *ex vivo* experiments were conducted on porcine skeletal muscle samples obtained from the loin region of a commercial food supplier. These samples were maintained at ambient temperature during the imaging measurements. The *in vivo* measurements were performed in a healthy 44-year-old male. In both cases, B-mode imaging was used to identify the dominant fiber orientation and to guide probe placement either parallel or perpendicular to the fibers. For each orientation, shear-wave datasets were acquired using SP and MLS-coded excitations (31-, 63-, and 127-bit), following the shared experimental platform and post-processing method described in Section 6.2. The SP reference datasets shown in Figs. 6.5–6.6 were obtained using the high-voltage reference method with high-frame-rate tracking.

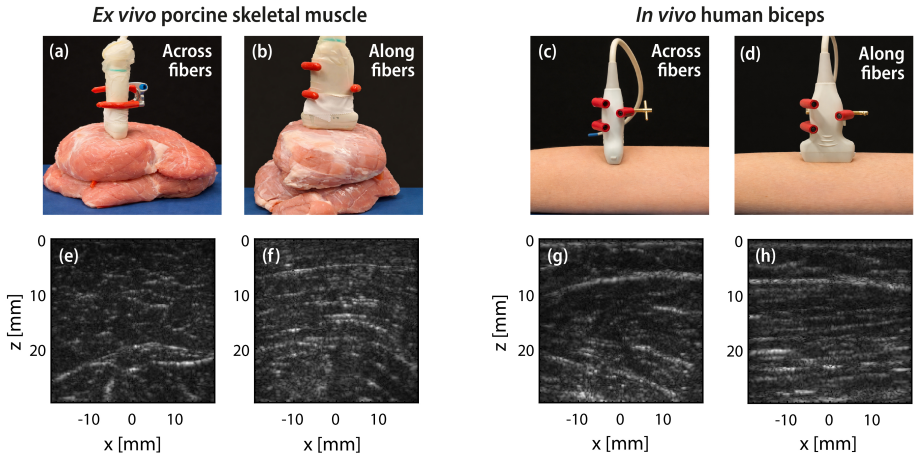


Figure 6.4: Schematic overview of the anisotropy experiments in skeletal muscle. B-mode imaging was used to identify the dominant fiber direction and guide acquisitions with the probe aligned along-fiber and across-fiber for both *ex vivo* and *in vivo* cases.

Metrics

For visualization and group-velocity estimation, spatiotemporal maps were computed from the decoded filtered shear-wave particle velocity by averaging within a 1-mm axial window centered at the ARF focus ($z = 20$ mm). Group velocity estimates were obtained from time-of-flight (ToF) tracking across a 15 mm lateral span, with the fitted ToF trend overlaid as a dashed line in the representative maps. For each excitation configuration, statistical group-velocity estimation was repeated across axial windows centered from -5 to $+5$ mm relative to the focal depth, using a 0.5 mm step. The resulting distributions were summarized in the corresponding violin plots in the rightmost panels of Figs. 6.5 and 6.6.

Results

Ex vivo porcine skeletal muscle

Representative spatiotemporal maps for the *ex vivo* experiment are reported in Fig. 6.5 for both along-fiber and across-fiber probe orientations. The reference acquisitions exhibited well-defined traveling-wave signatures with high regression quality in both geometries. In the along-fiber configuration, the reference yielded $c_g = 9.98$ m/s with $R^2 = 0.97$, as shown in Fig. 6.5(a). In the across-fiber configuration, the reference decreased to $c_g = 5.07$ m/s with $R^2 = 0.98$, as shown in Fig. 6.5(e).

At the lower-voltage SP settings used for the application experiments, the representative maps showed reduced wavefront definition and the ToF regression produced unstable estimates. In the along-fiber case, SP produced a sign-inverted estimate of $c_g = -2.63$ m/s with low fit quality ($R^2 = 0.16$), as shown in Fig. 6.5(b). In the along-fiber case, SP yielded an implausible velocity of $c_g = 33.18$ m/s with a failed regression ($R^2 = 0.00$), as shown in Fig. 6.5(f). After MLS decoding and pulse compression, the 127-bit datasets recovered coherent propagation patterns with ToF trends closely matching the reference. The along-fiber acquisition yielded $c_g = 5.15$ m/s with $R^2 = 0.99$ in Fig. 6.5(g), while the across-fiber acquisition yielded $c_g = 11.29$ m/s with $R^2 = 0.97$ in Fig. 6.5(c).

The embedded violin plots summarize the group-velocity distributions across axial windows for SP and the three MLS lengths. In the along-fiber

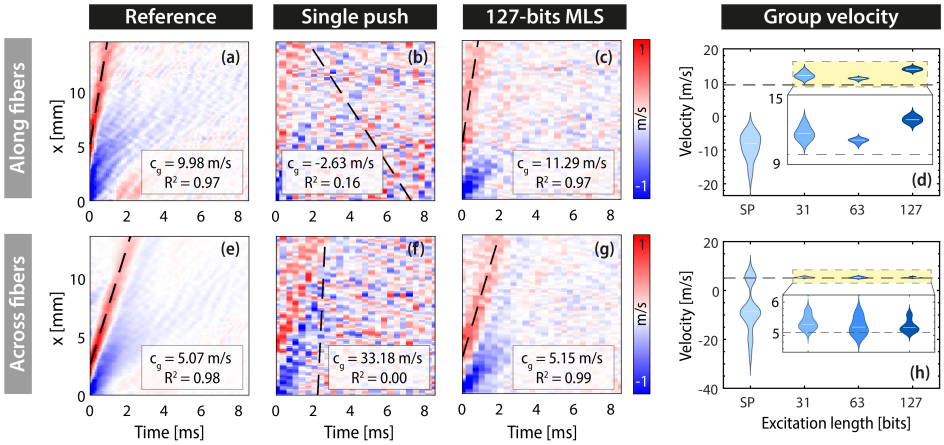


Figure 6.5: *Ex vivo* skeletal muscle anisotropy. Left: representative spatiotemporal particle-velocity maps for the reference, SP, and decoded 127-bit MLS acquisitions, shown for along-fiber (top row) and across-fiber (bottom row) probe orientations. Dashed lines indicate the ToF trend used for group-velocity estimation. Right: group-velocity distributions across axial windows for SP and MLS lengths (31/63/127 bits), with the reference range indicated by dashed horizontal lines.

geometry, the MLS-coded cases concentrated around the reference range, whereas SP produced a broad spread with prominent outliers in Fig. 6.5(d). In the across-fiber geometry, the MLS-coded distributions remained compact and close to the reference, while SP exhibited increased dispersion and sign inversions under the adopted propagation-direction convention in Fig. 6.5(h).

In vivo human biceps

The *in vivo* biceps results are reported in Fig. 6.6 using the same layout as the *ex vivo* figure. The reference acquisitions again displayed consistent traveling-wave trajectories in both geometries, with distinct group velocities between the along-fiber and across-fiber orientations. Along the fibers, Fig. 6.6(a), the reference yielded $c_g = 1.91$ m/s with $R^2 = 0.99$. Across the fibers, Fig. 6.6(e), the reference increased to $c_g = 1.41$ m/s with $R^2 = 0.87$.

For SP, the representative along-fiber map yielded an inconsistent ToF trend and an unreliable group-velocity estimate. In the along-fiber configuration the SP case produced a plausible velocity but with increased variability in the motion field, as shown in Fig. 6.6(b). In contrast, the across-fiber SP case

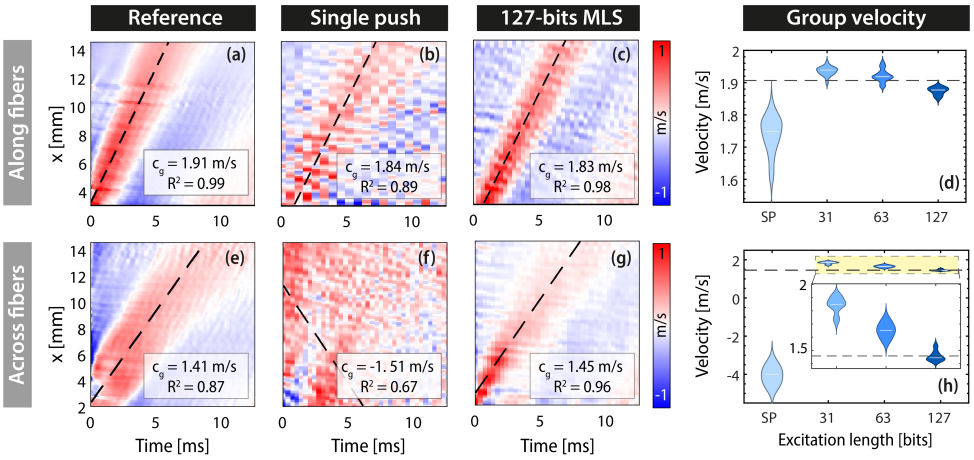


Figure 6.6: *In vivo* skeletal muscle anisotropy (human biceps). Left: representative spatiotemporal particle-velocity maps for the reference, SP, and decoded 127-bit MLS acquisitions, shown for along-fiber (top row) and across-fiber (bottom row) probe orientations. Dashed lines indicate the ToF trend used for group-velocity estimation. Right: group-velocity distributions across axial windows for SP and MLS lengths (31/63/127 bits), with the reference range indicated by dashed horizontal lines.

yielded an inconsistent ToF trend and an unreliable group-velocity estimate. In the along-fiber configuration, the estimate inverted sign ($c_g = -1.51$ m/s), as shown in Fig. 6.6(f). After MLS decoding and pulse compression, the 127-bit acquisitions produced clear propagation trajectories in both orientations. The along-fiber case yielded $c_g = 1.83$ m/s in Fig. 6.6(c), while the across-fiber case yielded $c_g = 1.45$ m/s in Fig. 6.6(g), both closely aligned with the corresponding reference examples.

The embedded violin plots summarize the group-velocity distributions across axial windows for SP and the three MLS lengths. In the along-fiber geometry, the MLS-coded distributions remained narrow and close to the reference, while SP showed increased dispersion relative to the coded cases in Fig. 6.6(d). In the across-fiber geometry, the MLS-coded cases yielded compact distributions centered near the reference range, whereas SP exhibited a broader spread and intermittent sign inversions in Fig. 6.6(h).

6.3.3 Viscoelastic media analysis

Frequency-dependent shear-wave propagation is a recurrent practical challenge when the goal is to extract quantitative viscoelastic indicators from SWE data. This section evaluates the MLS-coded pulse-compression framework under conditions where shear-wave dispersion is expected to be measurable and clinically relevant. The objective is to quantify how coded excitations affect (i) the effective SNR of the tracked motion in attenuating viscoelastic media and (ii) the reliability of dispersion-curve estimation at the reduced acoustic-output settings.

Experiment protocol

Two distinct viscoelastic media were examined: a custom-made viscoelastic phantom (CIRS Inc., Norfolk, VA, USA) and an *ex vivo* porcine liver. The phantom exhibited viscoelastic behavior with frequency-dependent shear-wave propagation previously documented, although no explicit reference elasticity or viscosity curves were available. The *ex vivo* porcine liver samples were obtained at Mayo Clinic from pigs euthanized for other educational research experiments conducted under approved institutional protocols. After collection, the liver tissue was kept refrigerated in degassed water and allowed to reach room temperature before measurement.

For this application study, the SP and MLS-coded acquisitions were performed at the reduced transmit voltage of 25 V (MI = 0.40). For each medium, four excitation configurations were evaluated: SP and MLS lengths of 31, 63, and 127 bits. Also a high-voltage (60 V, MI = 0.65) reference single push was used. Both datasets were acquired with the shared Verasonics-based platform, including the same imaging sequence and motion-estimation method described in Section 6.2. Finally, the coded motion was matched-processed using the known MLS sequence and directional filtering to isolate the dominant propagation direction and wavenumber-band selection to suppress out-of-range velocity content.

Metrics

For each excitation configuration and axial window, a *signal-to-noise ratio* was computed from the processed spatiotemporal particle-velocity map $v(x, t)$.

The motion SNR was defined in decibels as

$$\mathrm{SNR} = 20 \log_{10}\left(\frac{A_\mathrm{s}}{A_\mathrm{n}}\right), \tag{6.2}$$

where A_s and A_n are the RMS particle-velocity amplitudes computed over a wave-dominated spatiotemporal ROI, Ω_s, and a noise ROI, Ω_n, respectively:

$$A_\mathrm{s} = \sqrt{\frac{1}{N_\mathrm{s}} \sum_{(x,t)\in\Omega_\mathrm{s}} |v(x,t)|^2}, \qquad A_\mathrm{n} = \sqrt{\frac{1}{N_\mathrm{n}} \sum_{(x,t)\in\Omega_\mathrm{n}} |v(x,t)|^2}, \tag{6.3}$$

with N_s and N_n denoting the number of samples in the corresponding ROIs. Both regions were selected on the same $v(x,t)$ map, as illustrated in Figs. 6.7(a) and 6.8(a), with Ω_s aligned with the dominant propagation ridge and Ω_n placed in regions devoid of visible traveling-wave energy. The resulting SNR values are reported on the representative maps shown in Figs. 6.7(b–e) and 6.8(b–e), and their distributions across axial windows are summarized in Fig. 6.9(a,c).

Dispersion was quantified by estimating phase velocity as a function of frequency from the processed shear-wave spatiotemporal field. For each axial window, the motion field was mapped to the (k,ω) domain and the dominant ridge was tracked over frequency to obtain discrete phase-velocity samples $c_\mathrm{ph}(f)$. These samples (white markers in Figs. 6.7 and 6.8) were compared with the reference dispersion curve measured from the high-voltage SP acquisition (solid curve). This reference curve was used as an empirical baseline for comparison, and the agreement between both curves was quantified using a MATLAB goodness-of-fit metric reported as a percentage in Figs. 6.7(g–j) and 6.8(g–j). In addition, a band-averaged phase velocity was computed per window by averaging $c_\mathrm{ph}(f)$ over the frequency range where the ridge energy was consistently detectable; the resulting phase velocities at 450 Hz are reported in Fig. 6.9(b,d).

All windowed statistics were computed over a set of axial windows centered at the push focus. Windows were placed at axial offsets spanning $[-5, 5]$ mm relative to the focal depth, in increments of 0.5 mm, and each metric was computed independently per window before aggregating into violin distributions.

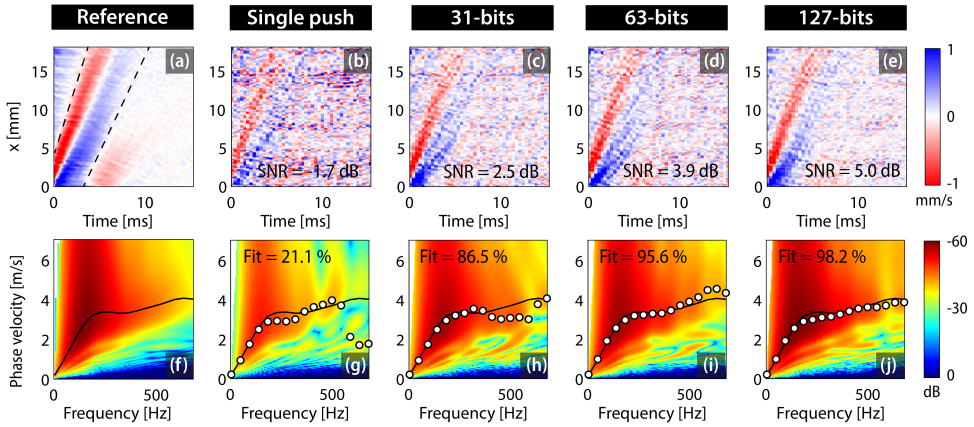


Figure 6.7: Viscoelastic phantom: (top) spatiotemporal wavefield maps and SNR evaluation windows for the reference, SP, and MLS excitations; (bottom) frequency–phase-velocity spectra showing the extracted dispersion samples (markers) and the reference dispersion curve obtained from the high-voltage SP acquisition (solid line). The reported fit metric quantifies the agreement between the extracted samples and this reference curve.

Results

Viscoelastic phantom

Representative spatiotemporal maps for the viscoelastic phantom are shown in Fig. 6.7. The high-voltage reference acquisition in Fig. 6.7(a) exhibits a clear traveling-wave pattern with limited background contamination. At the reduced-voltage SP setting, the map in Fig. 6.7(b) shows substantially increased spatiotemporal clutter, and the reported SNR is negative (SNR = −1.7 dB). After MLS decoding, progressively cleaner trajectories are observed as the code length increases, with SNR values of 2.5 dB for 31 bits in Fig. 6.7(c), 3.9 dB for 63 bits in Fig. 6.7(d), and 5.0 dB for 127 bits in Fig. 6.7(e).

The corresponding dispersion representations are shown in Fig. 6.7(f–j). The reference map in Fig. 6.7(f) displays a well-defined ridge across frequency, consistent with a dispersive response. In the reduced-voltage SP case, the ridge becomes fragmented and the extracted phase-velocity samples show reduced coherence, which is reflected in a low fit value of 21.1% in Fig. 6.7(g). In contrast, the MLS-coded cases produce consistent ridge structure and more stable phase-velocity samples, with fit values of 86.5% for 31 bits in Fig. 6.7(h), 95.6% for 63 bits in Fig. 6.7(i), and 98.2% for 127 bits in Fig. 6.7(j).

Windowed statistics for this phantom are summarized in Fig. 6.9(a,b). The SNR distributions in Fig. 6.9(a) show a shift from a negative-centered SP distribution to positive-centered MLS distributions, with a monotonic increase in central tendency from 31 to 127 bits. The phase-velocity distributions in Fig. 6.9(b) similarly indicate reduced spread for MLS-coded cases relative to SP, together with a stronger concentration around the reference range indicated by the dashed line.

Ex vivo porcine liver

Results for *ex vivo* porcine liver are reported in Fig. 6.8 using the same layout. The reference spatiotemporal map in Fig. 6.8(a) shows a coherent traveling-wave footprint, whereas the reduced-voltage SP map in Fig. 6.8(b) exhibits elevated background fluctuations, with SNR $= -1.4$ dB. After MLS decoding, clearer propagation patterns are observed for all code lengths, with SNR values of 2.6 dB in Fig. 6.8(c), 3.2 dB in Fig. 6.8(d), and 4.1 dB in Fig. 6.8(e).

Dispersion maps and the corresponding reference dispersion curves are shown in Fig. 6.8(f–j). The SP case in Fig. 6.8(g) yields a moderate fit value of 50.6%, with extracted samples showing visible scatter around the reference curve. In contrast, the MLS-coded cases show more consistent detection of the dominant ridge and closer agreement with the reference dispersion curve from the high-voltage SP acquisition, yielding fit values of 88.9% for 31 bits in Fig. 6.8(h), 92.2% for 63 bits in Fig. 6.8(i), and 97.9% for 127 bits in Fig. 6.8(j).

The SNR distributions in Fig. 6.9(c) again shift upward for MLS-coded acquisitions, with increasing central tendency from 31 to 127 bits. The phase-velocity distributions in Fig. 6.9(d) show that MLS-coded cases yield tighter distributions than SP and cluster closer to the dashed reference level, while SP exhibits broader spread consistent with higher variability in the estimated dispersion curves.

6.3.4 Pleural interface tissue

Ultrasound elastography in lung-adjacent conditions is constrained by acoustic-output safety, as MI limits can be more restrictive than in standard abdominal or musculoskeletal examinations. This limitation stems from the tissue–air

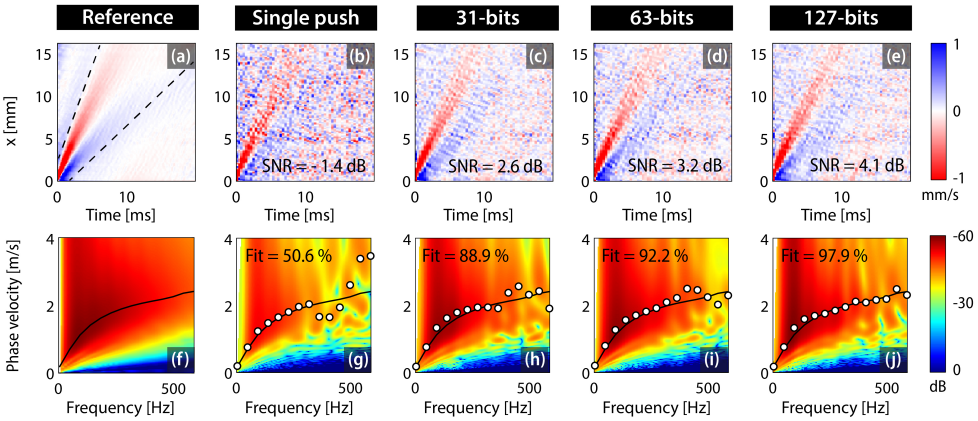


Figure 6.8: Ex vivo porcine liver: (top) spatiotemporal wavefield maps and SNR evaluation windows for the reference, SP, and MLS excitations; (bottom) frequency–phase-velocity spectra showing the extracted dispersion samples (markers) and the reference dispersion curve obtained from the high-voltage SP acquisition (solid line). The reported fit metric quantifies the agreement between the extracted samples and this reference curve..

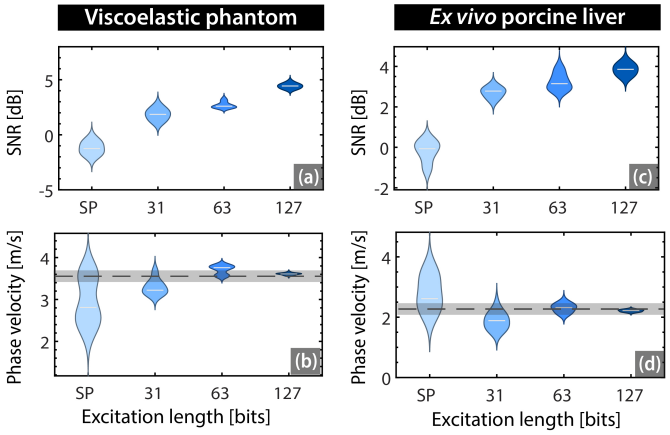


Figure 6.9: Summary statistics for viscoelastic phantom and *ex vivo* liver experiments. Violin plots report SNR and phase velocity at 450 Hz as a function of excitation length; dashed horizontal lines indicate the reference value.

interface, which increases susceptibility to mechanically driven bioeffects; in ARF-based approaches, high pressures in gas-containing structures have been linked to capillary hemorrhage in small-animal models [Miller et al., 2019]. Within this safety-limited regime, pulse-compression with MLS-coded ARF

provides a means to distribute excitation energy over time while operating at reduced MI.

Pulmonary edema [Wiley et al., 2021] was selected as a clinically motivated perturbation because fluid accumulation within the lung parenchyma alters the effective mechanical loading at the pleura and, consequently, the boundary condition experienced by shear waves propagating in the superficial soft-tissue layer. From an elastographic standpoint, increased fluid content can modify the apparent stiffness-related response measured above the pleura primarily by changing mechanical coupling across the pleural interface.

Experiment protocol

A layered, custom lung-mimicking phantom was assembled to emulate a superficial soft-tissue layer above a pleural interface and an underlying porous lung compartment, with a controllable degree of fluid loading. The phantom consisted of three stacked layers: (i) a 14-mm-thick chicken-breast slab, (ii) a 2-mm nitrile sheet acting as a pleura phantom, and (iii) a 14-mm-thick porous sponge as the inner lung tisuee phantom. Each layer is coupled with ultrasound coupling gel to ensure adhesion. The lateral extent of the phantom was approximately 100 mm. Acoustic coupling between layers was ensured using ultrasound coupling gel. For the sake of simplicity of this proof-of-concept, no ribs were included in this phantom.

To emulate a pulmonary edema condition in a controlled and repeatable manner, edema was simulated by progressively injecting water into the porous sponge layer (lung surrogate) using a syringe and needle, with injected volumes ranging from 0 to 30 mL in 3 mL steps. After each injection, a 2-min equilibration period was observed before acquisition, while keeping the probe position fixed to avoid geometric variability across conditions. Figure 6.10 illustrates the experimental concept (a) and the corresponding B-mode appearance (b), where the pleura line and typical A-line artifacts are visible.

All acquisitions were performed with an L11-5v linear array probe (Verasonics Inc., Redmond, WA), comprising 128 elements (element width 0.27 mm, height 5 mm, kerf 30 µm), an elevational focus of 18 mm, and a nominal center frequency of 7.6 MHz. The ARF push beam was transmitted using 24 active channels and focused at 0 mm in lateral direction and at a depth of 13 mm proximal to the pleura line, resulting in an f-number of 2. The ARF

burst duration was 65 µs at a center frequency of 6 MHz. Three excitation conditions were compared: a high-output single push at 50 V (MI = 1.4), a reduced-voltage SP at 15 V (MI = 0.7) and a 63-bit MLS excitation at 15 V (MI = 0.7).

For the MLS-coded acquisition, the bit interval was set to 250 µs, corresponding to the sum of the ARF burst duration (65 µs), an inter-burst waiting time (105 µs), and the acquisition time allocated per bit (80 µs). The total coded-excitation duration was therefore 15.75 ms.

Shear-wave tracking and motion estimation followed the same method described in Sections 6.2.1–6.2.3, with the only modification being the tracking transmit center frequency (7.6 MHz) to match the L11-5v configuration.

Metrics

Performance was quantified through phase-velocity mapping in the superficial layer above the pleura line. Phase velocity was estimated using the PGE method applied to the particle-velocity field, following pulse compression for the MLS-coded acquisitions. Phase-velocity maps were reconstructed as a

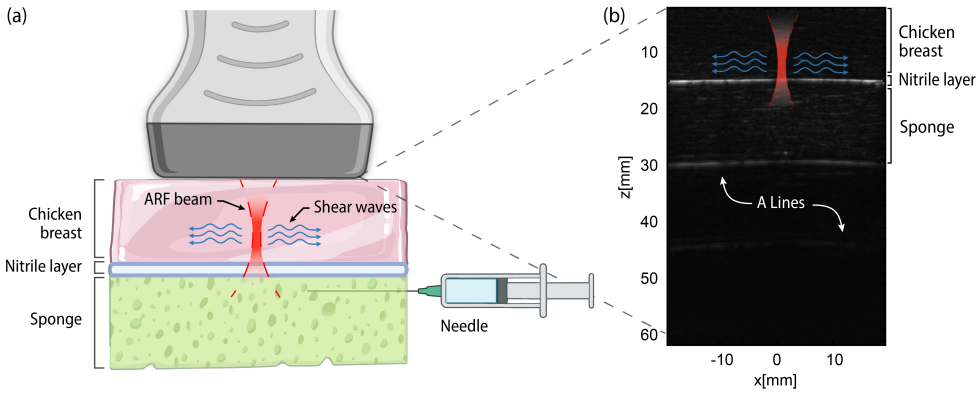


Figure 6.10: Pulmonary-edema surrogate experiment. (a) Schematic of the layered phantom and acquisition geometry. A chicken-breast layer overlies a nitrile pleura surrogate and a porous sponge acting as the lung compartment. Water is injected into the sponge to emulate progressive edema. (b) Representative B-mode image showing the chicken-breast layer, the pleura line (nitrile interface), and the sponge compartment, along with typical A-line artifacts. The ARF focus is positioned proximal to the pleura line, and shear-wave propagation is analyzed in the superficial layer.

weighted average over the frequency range 200–400 Hz. Directional filtering was restricted to rightward-propagating components, and maps were reconstructed only on the right-hand side of the ARF push.

For each injected volume, a rectangular ROI of 16 mm × 7 mm was defined above the pleura line on the right side of the push. Two summary metrics were extracted from this ROI, the mean phase velocity $\bar{c}_{\mathrm{ph}}$ and the spatial standard deviation $\sigma_{c_{\mathrm{ph}}}$ within the ROI. This standard deviation is reported as the shaded band in Figs. 6.11(d,e).

To quantify agreement with the reference condition across the complete 0–30 mL range, the mean absolute error (MAE), mean bias, and maximum absolute deviation of $\bar{c}_{\mathrm{ph}}$ were also computed for the low-MI 1- and 63-bits conditions relative to the reference.

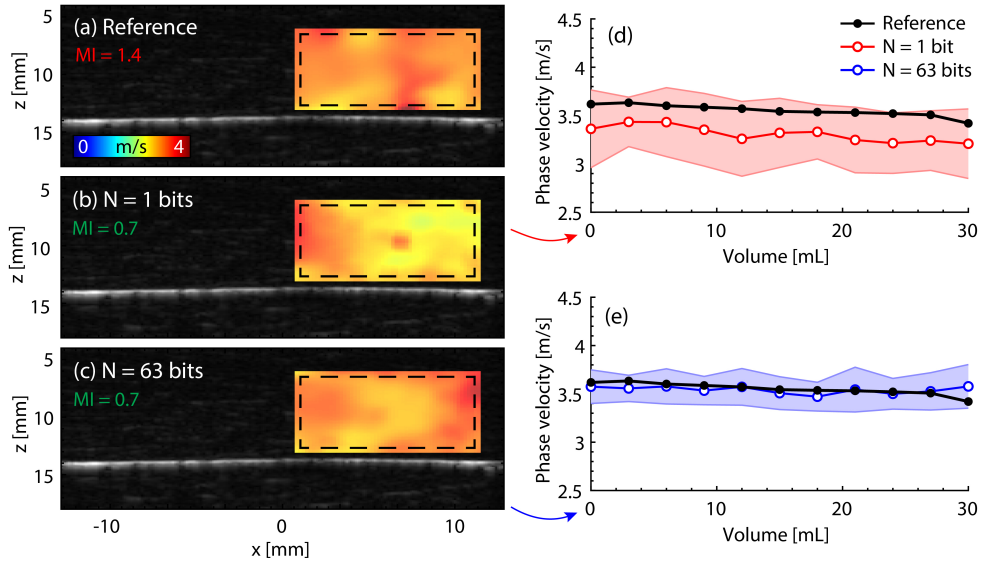


Figure 6.11: Pulmonary-edema surrogate results. (a–c) Phase-velocity maps overlaid on B-mode for a representative fluid-loading state (12 mL), comparing (a) reference SP (MI = 1.4), (b) low-MI SP (N = 1, MI = 0.7), and (c) low-MI MLS-coded ARF (N = 63, MI = 0.7). The dashed rectangle indicates the ROI used for summary statistics. (d-e) Mean phase velocity within the ROI as a function of injected volume for the reference, N = 1 and N = 63 conditions.

Results

Representative phase-velocity maps at an injected volume of 12 mL are shown in Fig. 6.11(a-c) for the three excitation conditions. Since the superficial chicken-breast layer is macroscopically homogeneous a spatially smooth reconstruction is expected within the ROI. The reference, Fig. 6.11(a), and the low-MI 63-bit condition, Fig. 6.11(c), exhibit comparatively uniform phase-velocity patterns, whereas the low-MI 1-bit case, Fig. 6.11(b), shows larger spatial fluctuations and localized heterogeneities within the ROI.

The evolution of the ROI-averaged phase velocity, $\bar{c}_{\mathrm{ph}}$, over the injected-volume range is summarized in Fig. 6.11(d-e). Across 0–30 mL, the reference mean phase velocity decreased from 3.62 to 3.42 m/s. Under low-MI single-push excitation (1-bit), $\bar{c}_{\mathrm{ph}}$ remained consistently lower than the reference at all volume steps, with MAE = 0.24 m/s and a mean bias of −0.24 m/s. Spatial variability was also higher in this condition, with a mean ROI standard deviation of 0.27 m/s, reported as the shaded band.

With low-MI MLS-coded excitation (63-bit), $\bar{c}_{\mathrm{ph}}$ more closely followed the reference trend, yielding MAE = 0.05 m/s and a mean bias of −0.01 m/s. The corresponding mean spatial standard deviation decreased to 0.10 m/s. The maximum absolute deviation relative to the reference was 0.16 m/s, observed at 30 mL.

6.3.5 Moving and vibrating media: robustness to external motion

Mechanically induced motion unrelated to the ARF excitation, such as physiological motion or exogenous vibrations due to probe or patient motion, can contaminate the measured shear-wave wavefield and degrade the detectability of the main propagating component. To emulate this scenario in a controlled manner, an experiment was designed where an additional background shear-like motion was generated independently of the ARF timing, i.e., with no temporal correlation to the excitation code. The resulting datasets enabled a systematic evaluation under progressively stronger uncorrelated disturbances.

Experiment protocol

A homogeneous elastic phantom (Model 039, CIRS Inc., Norfolk, VA) with a nominal elasticity of 25 kPa was used. Background motion was generated by coupling a loudspeaker equipped with a central stem to the lateral base of the phantom. The loudspeaker was driven continuously with a band-limited random waveform generated by a function generator, with spectral content centered around 1 kHz. Since the loudspeaker drive was applied independently of the ARF transmission, and the AWG was randomly triggered before each acquisition, the induced motion did not share a timing reference with either SP or MLS excitations. Therefore, both wave fields were considered temporally incoherent. The experimental setup is illustrated in Fig. 6.12.

Two baseline datasets were acquired and subsequently combined to build noisy conditions with controlled amplitude. The Gaussian noisy motion dataset was first recorded with the loudspeaker active and ARF transmission disabled, as depicted in Fig. 6.12(a) and illustrated by the spatiotemporal example

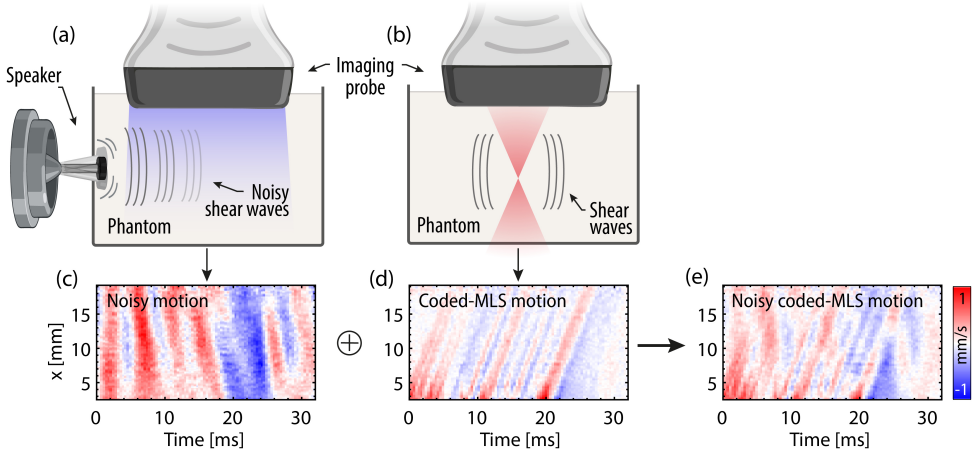


Figure 6.12: Experimental setup to emulate mechanically induced motion during SWE acquisitions. (a) A loudspeaker coupled to the phantom generates background shear-wave motion, which is recorded by the imaging probe as a noisy motion field. (b) A clean MLS-coded ARF excitation generates the target shear-wave response in the same phantom. (c) Representative spatiotemporal map of speaker-induced motion. (d) Representative spatiotemporal map of the clean coded-MLS shear-wave motion. (e) The final dataset used for processing is obtained by superposition of (c) and (d), producing noisy coded-MLS motion.

in Fig. 6.12(c). For this acquisition the same tracking configuration as in Section 6.2.1 were used. A clean motion dataset was then acquired in the same phantom with the loudspeaker inactive while transmitting either SP or MLS-coded ARF excitations, following the configuration shown in Fig. 6.12(b) and exemplified in Fig. 6.12d. Finally, noisy datasets were constructed by digitally scaling the noise-only data and superimposing it onto the signal-only response, yielding the combined noisy coded-MLS motion field shown in Fig. 6.12(e). This strategy enabled the same background-motion realization to be applied across SP and all MLS lengths, ensuring fair comparisons under matched noise conditions.

The noise level was parameterized by the peak particle-velocity amplitude of the noise-only component, spanning 0.2 to 8.2 mm/s. For reference, the peak particle velocities measured in the corresponding noise-free signal-only datasets were 2.7 mm/s for SP and 5.3 mm/s for the uncompressed 127-bit MLS dataset. Following the shared processing method described in Sec. 6.2, $v(x,t)$ were reconstructed at the focal depth by averaging axial particle velocity over a 1.0 mm depth range centered at the ARF focus.

Metrics

Robustness to uncorrelated background motion was quantified using the SNR metric defined previously in Eq. (6.2). SNR was computed for each excitation configuration (SP, and 31-, 63-, and 127-bit MLS) at each imposed noise amplitude. The signal and noise regions employed in the SNR calculation are indicated by dashed boundaries in Fig. 6.13.

Results

The global behavior as noise strength increases is summarized in Fig. 6.14, where SNR is reported as a function of the imposed background-motion amplitude for SP and for 31-, 63-, and 127-bit MLS excitations. In the SP condition, the map is dominated by the mechanically induced background motion, and the main shear-wave trajectory is weakly distinguishable within the selected time window, as illustrated in Fig. 6.13(a). The SNR reported for this representative SP case equals -9.5 dB.

For MLS-coded excitations, the measured maps in Fig. 6.13(b) through

Fig. 6.13(d) exhibit a temporally extended coded response, where the desired shear-wave signature coexists with the same background motion. After deconvolution, the compressed responses in Fig. 6.13(f) through Fig. 6.13(h) recover a more localized propagation ridge within the analysis window. The SNR values obtained after pulse compression for this representative realization were -2.9 dB, -2.0 dB and 1.2 dB for each MLS length respectively. Relative to SP, these correspond to SNR gains of 6.6 dB, 7.5 dB, and 10.7 dB, respectively, for the same noise level and evaluation window.

The dependence of robustness on the imposed noise amplitude is summarized in Fig. 6.14, where SNR is plotted as a function of mechanically induced noise for SP and MLS-coded excitations. Across the explored range, SNR decreases as noise amplitude increases for all methods. However, MLS-coded excitations remain consistently above SP, and longer sequences preserve higher SNR at progressively stronger background motion, reflecting a more stable separation between the dominant propagation ridge and the incoherent motion field under the same additive-noise construction.

6.3.6 Physiologically moving media: carotid artery

This experiment investigated the temporal consistency of group-velocity estimation in the common carotid artery under *in vivo* conditions, where arterial wall motion can introduce low-frequency components that degrade the interpretability of shear-wave propagation maps and increase variability in velocity estimates. To probe this effect in a controlled manner, acquisitions were synchronized to the cardiac cycle using ECG triggering and repeated across multiple cardiac stages.

Experiment protocol

Measurements were performed in a healthy volunteer, a 44-year-old male. According to the study framework used for this pilot acquisition, Institutional Review Board approval was not required. The subject was positioned supine and the ultrasound probe was placed over the left common carotid artery while ECG was recorded to provide a trigger signal for gated acquisitions.

Figure 6.15 summarizes the experimental configuration, including the ECG trigger path to the Verasonics system and the probe positioning over the neck.

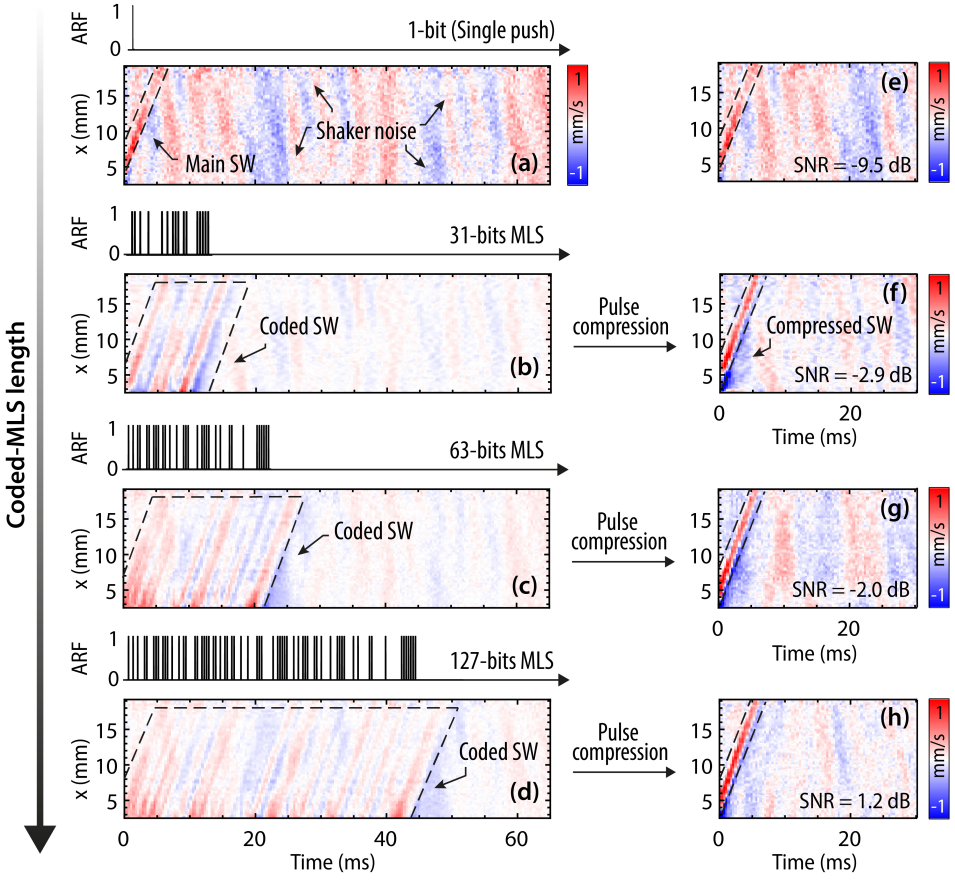


Figure 6.13: Representative noisy-motion experiment for a fixed noise level, showing spatiotemporal particle-velocity maps before and after pulse compression for SP and increasing MLS lengths. (a) SP case in the presence of shaker-induced background motion. (b-d) Measured noisy coded motion for 31-bit, 63-bit, and 127-bit MLS, respectively, with the applied ARF code shown above each map. (e-h) show the corresponding outputs after pulse compression. Dashed lines indicate the main propagation ridge and the temporal window used for SNR evaluation.

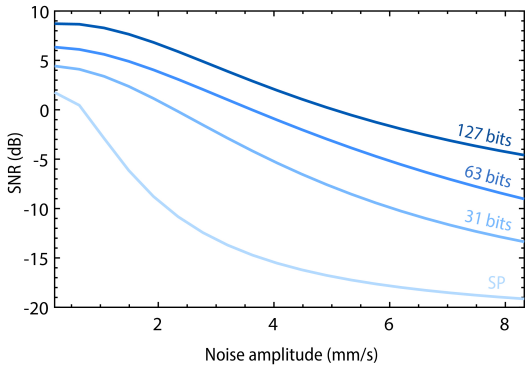


Figure 6.14: SNR as a function of mechanically induced noise amplitude for SP and MLS-coded excitations. Longer MLS sequences exhibit increased robustness, maintaining higher SNR under progressively stronger background motion.

The artery was first localized in B-mode and the ARF focus was positioned on the anterior wall of the artery, which provided higher motion SNR than the posterior wall in this acquisition geometry. The B-mode panel in Fig. 6.15 indicates the focal location and the expected region of shear-wave propagation along the anterior wall segment.

ECG gating was then used to trigger acquisitions at ten uniformly spaced stages along the cardiac cycle as shown in Fig. 6.16(a). For each stage, the same acquisition protocol was repeated five times per excitation configuration to quantify stage-wise repeatability.

The tested excitations comprised SP, 31-bit MLS, and 63-bit MLS, all transmitted at an input voltage of 25 V (MI = 0.40) to intentionally operate under motion-challenged conditions. A high-voltage 60 V (MI = 0.65) SP acquisition was additionally collected as a reference condition. The 127-bit MLS configuration was not used because its acoustic output can exceed the FDA-recommended spatial-peak temporal-average intensity limit of 720 mW/cm^2, as previously characterized for the same implementation in the previous Section 5.5.3.

For completeness of the exposure conditions at 25 V, the measured MI was 0.45. Assuming a heart rate of 80 BPM for time-averaging, the resulting I_{SPTA} values were 3.6 mW/cm^2 for SP, 115.7 mW/cm^2 for 31-bit MLS, and 231.4 mW/cm^2 for 63-bit MLS.

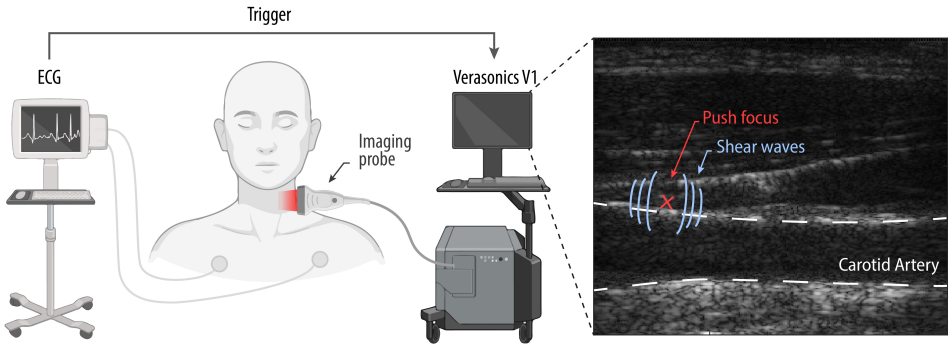


Figure 6.15: ECG-triggered *in vivo* carotid experiment. The imaging probe was positioned over the common carotid artery while ECG triggering synchronized acquisitions to the cardiac cycle. The B-mode image indicates the ARF focal location at the anterior wall and the expected shear-wave propagation region along the superficial wall segment.

Metrics

For each cardiac stage and excitation configuration, the axial particle-velocity field was processed to form a spatiotemporal particle velocity map $v(x,t)$ along the superficial carotid wall at the ARF focal depth. Stage-wise group velocity, denoted c_g, was computed using the same time-of-flight procedure. In brief, a traveling-wave trajectory was identified over the lateral observation window corresponding to the upper wall and the time-of-flight trend was fitted to yield a single c_g estimate per stage and repetition.

To assess temporal consistency across the cardiac cycle, the five repeated measurements acquired at each stage were summarized by their mean and standard deviation. The mean described the stage-dependent variation in c_g, while the standard deviation reflected repeatability under physiological motion and ECG gating. In the plots reported in Fig. 6.16, the experimental estimates are shown as solid markers, and their variability across repetitions is represented by a shaded band. A high-voltage SP curve is included as a reference trajectory for comparison across stages.

Results

Figure 6.16 reports ECG-gated results across the ten cardiac stages for SP, 31-bit MLS, and 63-bit MLS. The ECG trace in Fig. 6.16(a), illustrates the

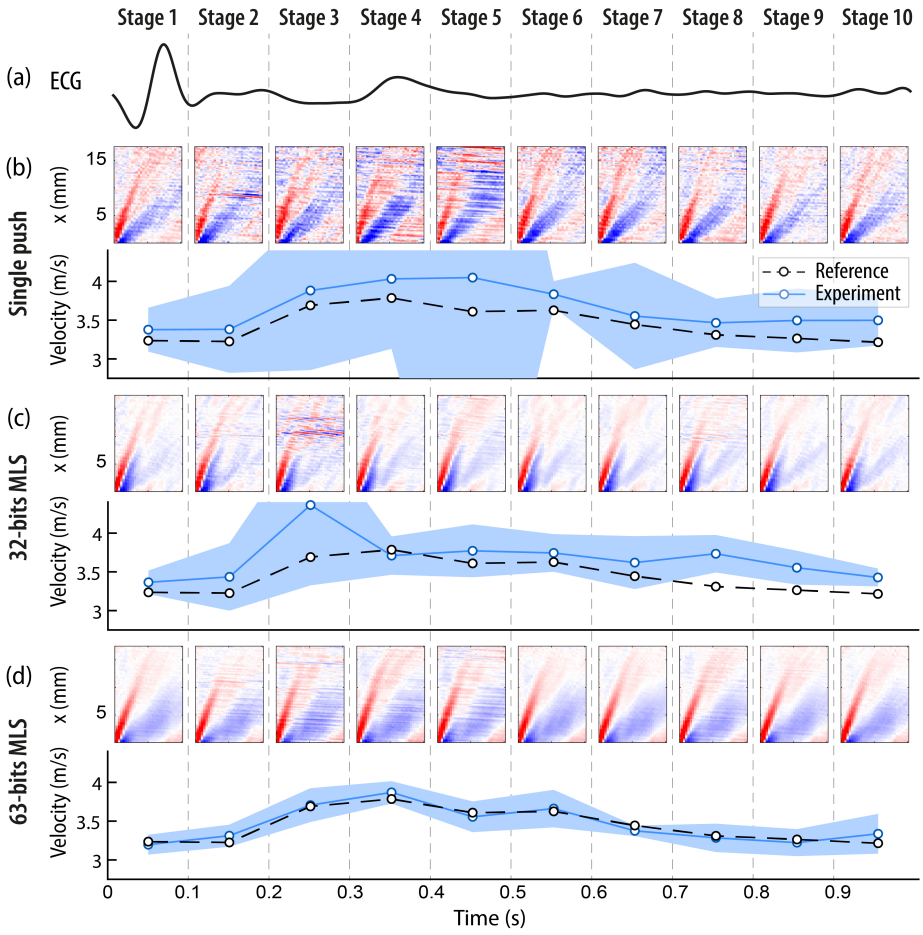


Figure 6.16: ECG-gated carotid results across ten cardiac stages for SP and MLS-coded excitations. (a) ECG trace and the ten uniformly spaced trigger stages. (b) SP results, including representative spatiotemporal $v(x,t)$ maps for each stage and the corresponding stage-wise group-velocity estimates compared with the high-voltage SP reference and (c) and (d) represents the sane outputs for 31 and 63-bit MLS. In the group-velocity plots, solid markers denote the mean across five repetitions per stage and the shaded band indicates the associated variability, while the dashed curve denotes the reference measurement.

division of the cardiac cycle into ten uniformly spaced stages used to trigger the acquisitions. For each excitation configuration, representative $v(x,t)$ maps are displayed above the corresponding group-velocity curve, enabling a direct visual association between wavefront regularity and the resulting stage-wise

c_g estimates.

For SP acquisitions, the representative maps in Fig. 6.16(b), show that several early-to-mid stages exhibit irregular or distorted wavefronts, particularly in stages 2 through 5. In the corresponding group-velocity curve this variability is reflected by large stage-to-stage excursions relative to the high-voltage reference. A prominent example occurs at stage 5, where the SP estimate reached 6.17 m/s while the reference at the same stage was 3.83 m/s. In contrast, later stages can align more closely with the reference, as observed at stage 10 where SP yielded 3.31 m/s compared with 3.42 m/s for the reference, together with a reduced spread indicated by a standard deviation of 0.39 m/s.

For 31-bit MLS, the representative maps in Fig. 6.16(c), present a more regular traveling-wave appearance across stages, while still showing visible differences in wavefront definition between early and late phases of the cycle. The resulting stage-wise c_g estimates ranged from 3.44 to 3.91 m/s across the cardiac stages, with standard deviations generally below 0.7 m/s. Agreement with the reference was stage-dependent. For example, at stage 3 the estimate was 3.62 m/s compared with 3.88 m/s for the reference, whereas at stage 9 the estimate was 3.73 m/s, closely matching the reference value of 3.71 m/s.

For 63-bit MLS, the maps in Fig. 6.16(d), show consistently well-defined propagation patterns across the full set of stages, and the corresponding group-velocity curve remains confined to a narrower range. Across all stages, c_g remained within 3.44 to 3.76 m/s and the standard deviation stayed below 0.35 m/s. Stage-wise agreement with the reference was close throughout the cycle, including stage 2 where 63-bit MLS yielded 3.68 m/s compared with 3.64 m/s for the reference. At stage 5, the 63-bit MLS estimate was 3.67 m/s while the reference was 3.81 m/s, reducing the deviation observed for SP at the same stage.

6.4 Discussion

This chapter evaluated MLS-coded ARF pulse compression in four application-driven scenarios that deliberately stress conventional SWE processing: spatially complex propagation in heterogeneous phantoms, direction-dependent propagation in anisotropic muscle, frequency-dependent behavior in viscoelastic media, and motion contamination from external or physiological sources. Across these settings, the same operating principle was tested under reduced-voltage

conditions: ARF energy is distributed over time using a binary pseudo-random sequence and subsequently collapsed into an impulse-like response by matched processing, thereby improving the effective motion SNR without increasing peak pressure.

In heterogeneous phantoms, the key practical challenge is not only a low displacement amplitude, but also the impact of the spatial complexity of the wavefield in the reconstructed shear wave velocity map. Reflections, interference, and inclusion-induced wavefront distortions can increase local variability and reduce the separability between inclusion and background in phase-velocity reconstructions. Under these conditions, the progressive CNR increase with MLS length shows that pulse compression improves inclusion detectability by increasing inclusion–background separation and lowering background variability. For example, in the 36 kPa inclusion, CNR increased from 7.1 dB with SP to 8.3 dB, 10.6 dB, and 12.6 dB for 31-, 63-, and 127-bit MLS, respectively, while the background standard deviation decreased from 0.19 m/s to 0.15 m/s, 0.14 m/s, and 0.12 m/s. These trends are consistent with the interpretation that, when phase-gradient estimators operate near their noise floor, phase noise and local artifacts are translated into spatially structured errors in the final map, which are then suppressed as the coherent, code-correlated component becomes dominant after pulse compression.

Skeletal muscle provides a stringent test because the expected along-fiber propagation is both faster and typically more coherent than across-fiber propagation, while the across-fiber case is more attenuated and therefore more prone to tracking noise. The results show that MLS-coded excitation restores the directional dependence of group velocity that is otherwise obscured when conventional SP acquisitions become dominated by variability. This effect is particularly relevant in the stiff and high-speed *ex vivo* measurements, where SP estimates can deviate dramatically from physically plausible trajectories, with reported deviations exceeding 20 m/s in the most challenging case. In contrast, longer MLS sequences recover a clear propagation ridge after decoding, supporting a more stable time-of-flight trend extraction. Importantly, this improvement is achieved despite the intrinsic limitation imposed by the fixed MLS bit timing, which constrains the motion sampling rate to the kHz range and can reduce the representation of very high-frequency wave components in fast media. Practically, this defines a regime where longer MLS sequences can compensate for reduced temporal sampling by increasing shear-wave energy, but only up to the point where PRF-limited bandwidth becomes the dominant

bottleneck.

In viscoelastic characterization, the objective shifts from estimating a single speed to recovering a frequency-dependent curve with sufficient fidelity to support a reference curve. In this setting, improvements in motion SNR translate directly into dispersion-curve stability because the spectral peak in the k–ω domain becomes better localized, reducing the susceptibility of the estimator to spurious maxima and residual interference. The reported goodness-of-fit increases relative to the high-voltage reference support this mechanism: in the viscoelastic phantom, the fit value increased from 21.1% with SP to 79.4% for 127-bit MLS, respectively, while the *ex vivo* liver results improved from 50.6% with SP to 97.9% 127-bit MLS. Notably, these gains are most consequential at higher frequencies, where shear-wave attenuation is strongest and where dispersion-induced differences are typically more informative, but also more difficult to measure reliably. The results therefore suggest that coded excitation can expand the practically usable frequency range for shear-wave spectroscopy under voltage-limited conditions, while still requiring careful consideration of stationarity assumptions when long sequences are used in media whose mechanical state may evolve within tens of milliseconds.

The pleura-adjacent pulmonary-edema application illustrates the value of MLS-coded ARF in a safety-limited regime where MI constraints restrict peak pressure in the presence of a tissue–air interface. In the layered, pulmonary-edema-mimicking phantom, progressive fluid loading in the porous sponge layer provided a controlled modulation of the effective boundary condition across the pleura line, and the superficial phase velocity estimated above the pleura exhibited a clear negative dependence on injected volume (3.62 m/s at 0 mL to 3.42 m/s at 30 mL for the reference). This trend is consistent with reports showing that increased lung water can reduce measured shear wave speeds and apparent elasticity under lung-surface and parenchymal elastography paradigms, plausibly through changes in effective loading, mass density, and damping as fluid replaces air and alters the mechanical contribution of air-filled microstructure [McGee et al., 2012, Zhang et al., 2019]. Under low-MI single-push excitation, the trend was obscured by systematic underestimation and elevated spatial variability, Fig. 6.11(d), consistent with insufficient shear-wave SNR for stable directional filtering and PGE reconstruction. By contrast, low-MI MLS-coded excitation preserved the decreasing trend and remained closely aligned with the high-output reference, as shown in Fig. 6.11(e). It also markedly reduced spatial dispersion within the ROI, lowering the MAE from

0.24 to 0.05 m/s and the mean ROI standard deviation from 0.27 to 0.10 m/s. These results suggest that pulse compression can retain quantitative sensitivity to edema-driven boundary-condition changes even when peak acoustic output must remain small due to safety constraints.

The experiments on moving and shaking media isolates a failure mode that often appears in practice in less controlled forms: background motion that is not temporally correlated with the ARF excitation. In this regime, pulse compression provides a structural advantage because the shear-wave response generated by the MLS-coded excitation follows the temporal pattern of the code, whereas the mechanically induced background motion does not. As a result, after decoding, the ARF-induced response is compressed into a more localized propagation ridge, while the unrelated background motion remains temporally spread and contributes less strongly to the final map. This behavior is visible both qualitatively in the recovered propagation ridge and quantitatively in the representative SNR values measured after deconvolution. For the fixed-noise example, SNR improved with gains of 6.6 dB, 7.5 dB, and 10.7 dB for 31-, 63-, and 127-bit MLS compared with the SP case. The summary curve across noise amplitudes further shows that, while all methods degrade as mechanically induced motion increases, longer MLS sequences remain systematically more robust, indicating that the separation between the dominant shear-wave ridge and background motion is preserved over a wider disturbance range.

The carotid experiment adds two clinically relevant constraints: physiological motion that varies over the cardiac cycle and regulatory limits that restrict the use of long, high-energy excitations. The ECG-gated results show that conventional SP acquisitions can exhibit stage-dependent distortions of the propagation map that translate into large excursions in the estimated c_g, including a peak of 6.17 m/s at stage 5 compared with 3.83 m/s in the high-voltage reference. Notably, the largest deviations occur in stages 2 through 5, which coincide with the phase of stronger intrinsic arterial wall motion driven by pulsatile blood flow. In this interval, the vessel undergoes rapid diameter changes and elevated out-of-plane and axial components that can decorrelate motion tracking and bias the apparent propagation ridge in $v(x, t)$, thereby perturbing the time-of-flight trend used for c_g estimation. In contrast, MLS-coded acquisitions yield tighter stage-wise ranges and improved agreement with the reference trajectory, with the 31-bit MLS spanning 3.44–3.91 m/s and typically remaining within sub-m/s deviations in later stages. In contrast, MLS-coded

acquisitions yield tighter stage-wise ranges and improved agreement with the reference trajectory, with the 31-bit MLS spanning 3.44–3.91 m/s and typically remaining within sub-m/s deviations in later stages. This stabilization is achieved while operating at 25 V and maintaining MI and I_{SPTA} within the safety limits. Rather than reflecting only an increase in displacement amplitude, this behavior suggests that deconvolution attenuates motion components that are not temporally correlated with the excitation sequence, so that the resulting motion data are more strongly dominated by the ARF-induced shear waves. At the same time, the need to exclude the 127-bit sequence in this *in vivo* setting emphasizes that MLS length selection is inherently application dependent: longer codes improve robustness, but they increase exposure and extend the stationarity window, which can be incompatible with dynamic targets such as the arterial wall.

Three constraints emerge as central for translating MLS-coded SWE beyond controlled experiments. First, sequence duration imposes an implicit stationarity assumption: as the code length increases into the tens-of-milliseconds range, tissue properties and boundary conditions must remain sufficiently stable during the coded excitation for the decoded impulse response to remain interpretable. Second, the fixed MLS bit duration limits the effective motion sampling rate and therefore the accessible shear-wave bandwidth, which can become limiting in stiff media or along fast propagation directions. Third, time-averaged exposure can become the dominant safety bottleneck for long codes, as illustrated by the exclusion of the 127-bit configuration in the carotid experiment due to I_{SPTA} constraints. These limitations do not negate the benefits of coded excitation, but they define a practical design space where MLS length, bit timing, transmit voltage, and repetition strategy must be co-optimized for the target anatomy, expected dynamics, and regulatory limits.

The results motivate several directions that are technically plausible within the same framework. Adaptive selection of MLS length and transmit settings could be driven by real-time motion SNR estimates, selecting short codes when temporal fidelity dominates and longer codes when coherence is the limiting factor. In parallel, strategies that relax the fixed-bit constraint, such as application-specific timing or hybrid coding across multiple shorter sequences, could mitigate PRF limitations while preserving the pulse compression advantage. Finally, broader *in vivo* validation across subjects and anatomies remains essential to quantify robustness under realistic variability in scattering, attenuation, and motion, especially in safety-restricted environments where

the ability to increase peak acoustic output is fundamentally limited.

6.5 Conclusions

MLS-coded ARF pulse compression was shown to improve SWE robustness under reduced-voltage conditions across heterogeneous, anisotropic, viscoelastic, and moving and vibrating tissues. Quantitatively, longer MLS sequences increased contrast in heterogeneous reconstructions while reducing background variability, stabilized group-velocity estimation in anisotropic muscle where single-push became unreliable, improved dispersion-curve fidelity in viscoelastic media as reflected by large gains in goodness-of-fit, and enhanced resilience to uncorrelated motion by preserving a decodable, code-correlated propagation ridge. The *in vivo* carotid experiment further demonstrated that coded excitation can reduce cardiac stage-dependent variability without increasing peak acoustic pressure.Collectively, these findings support MLS-coded excitation as a scalable mechanism to trade excitation duration for a more robustly recoverable shear-wave response, thereby expanding the operating envelope of SWE in applications where conventional strategies are constrained by attenuation, complexity, or safety limits.

Publications

This section contains a compendium of the seminal scientific articles in which the results presented in this chapter have been published.

Conferences articles

- González-Mateo, E., Urban, M. W., and Jimenez, N., "MLS-Coded ARF Excitation for Noise-Resilient Shear Wave Elastography," in *2025 IEEE International Ultrasonics Symposium (IUS 2025)*, Utrecht, Netherlands, 2025.
- González-Mateo, E., Urban, M. W., and Jimenez, N., "Improving Shear Wave Elastography by coded Acoustic Radiation Force," in *International Tissue Elasticity Conference (ITEC 2025)*, Seattle, USA, 2025.
- González-Mateo, E., Rodríguez-Sendra, J., Urban, M. W., and Jimenez, N., "Pulse-compression for lung elastography using low mechanical index transmissions," in *International Lung Ultrasound Symposium (ILUS 2025)*, Trento, Italy, 2025, p. 10.
- Rodríguez-Sendra, J., González-Mateo, E., Jimenez, N., and Camarena, F., "Singular value decomposition beamforming and pulse-coded excitation for lung ultrasound aberration correction," in *TECNIACÚSTICA 2022. 53º Congreso Español de Acústica. XII Congreso Ibérico de Acústica. European Symposium on NVH in Electric Vehicles*, Elche, España, 2022, pp. 1151–1156.

References

Capriotti, M., Greenleaf, J. F., and Urban, M. W. (2020). Time-aligned plane wave compounding methods for high-frame-rate shear wave elastography: Experimental validation and performance assessment on tissue phantoms. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 67(12):2580–2592.

Deffieux, T., Gennisson, J.-L., Larrat, B., Fink, M., and Tanter, M. (2012). The variance of quantitative estimates in shear wave imaging: Theory and experiments. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 59(11):2390–2410.

Deffieux, T., Montaldo, G., Tanter, M., and Fink, M. (2009). Shear wave spectroscopy for in vivo quantification of human soft tissues visco-elasticity. *IEEE Transactions on Medical Imaging*, 28(3):313–322.

Deng, Y., Palmeri, M. L., Rouze, N. C., Rosenzweig, S. J., Abdelmalek, M. F., and Nightingale, K. R. (2015). Analyzing the impact of increasing mechanical index and energy deposition on shear wave speed reconstruction in human liver. *Ultrasound in Medicine & Biology*, 41(7):1948–1957.

Golomb, S. W. (2017). *Shift register sequences: secure and limited-access code generators, efficiency code generators, prescribed property generators, mathematical models*. World Scientific.

Kasai, C., Namekawa, K., Koyano, A., and Omoto, R. (1985). Real-time two-dimensional blood flow imaging using an autocorrelation technique. *IEEE Transactions on Sonics and Ultrasonics*, 32(3):458–464.

Kijanka, P. and Urban, M. W. (2019). Local phase velocity based imaging: A new technique used for ultrasound shear wave elastography. *IEEE Transactions on Medical Imaging*, 38(4):894–908.

Maksuti, E., Widman, E., Larsson, D., Urban, M. W., Larsson, M., and Bjällmark, A. (2016). Arterial stiffness estimation by shear wave elastography: Validation in phantoms with mechanical testing. *Ultrasound in Medicine & Biology*, 42(1):308–321.

Manduca, A., Lake, D., Kruse, S., and Ehman, R. (2003). Spatio-temporal directional filtering for improved inversion of mr elastography images. *Medical Image Analysis*, 7(4):465–473. Medical Image Computing and Computer Assisted Intervention.

McGee, K. P., Mariappan, Y. K., Hubmayr, R. D., Carter, R. E., Bao, Z., Levin, D. L., Manduca, A., and Ehman, R. L. (2012). Magnetic resonance assessment of parenchymal elasticity in normal and edematous, ventilator-injured lung. *Journal of Applied Physiology (Bethesda, Md.: 1985)*, 113(4):666–676.

Miller, D. L., Dong, Z., Dou, C., Patterson, B., and Raghavendran, K. (2019). Pulmonary capillary hemorrhage induced by super sonic shear wave elastography in rats. *Ultrasound in Medicine & Biology*, 45(11):2993–3004.

Muller, M., Gennisson, J.-L., Deffieux, T., Tanter, M., and Fink, M. (2009). Quantitative viscoelasticity mapping of human liver using supersonic shear imaging: Preliminary in vivo feasability study. *Ultrasound in Medicine & Biology*, 35(2):219–229.

Nguyen, T.-M., Song, S., Arnal, B., Wong, E. Y., Huang, Z., Wang, R. K., and O'Donnell, M. (2014). Shear wave pulse compression for dynamic elastography using phase-sensitive optical coherence tomography. *Journal of Biomedical Optics*, 19(1):016013.

Nordenfur, T., Caidahl, K., Lindberg, L., Urban, M. W., and Larsson, M. (2025). Safety of shear wave elastography as evidenced from carotid artery strain and strain rate induced by acoustic radiation force impulse and arterial pulsations. *Ultrasound in Medicine & Biology*, 51(5):742–750.

Roy, T., Lee, H.-K., Capron, C. B., Lopez-Jimenez, F., Hesley, G. K., Greenleaf, J. F., Urban, M. W., and Guddati, M. N. (2025). Estimation of *In Vivo* Human Carotid Artery Elasticity Using Arterial Dispersion Ultrasound Vibrometry. *Ultrasound in Medicine & Biology*, 51(2):250–261.

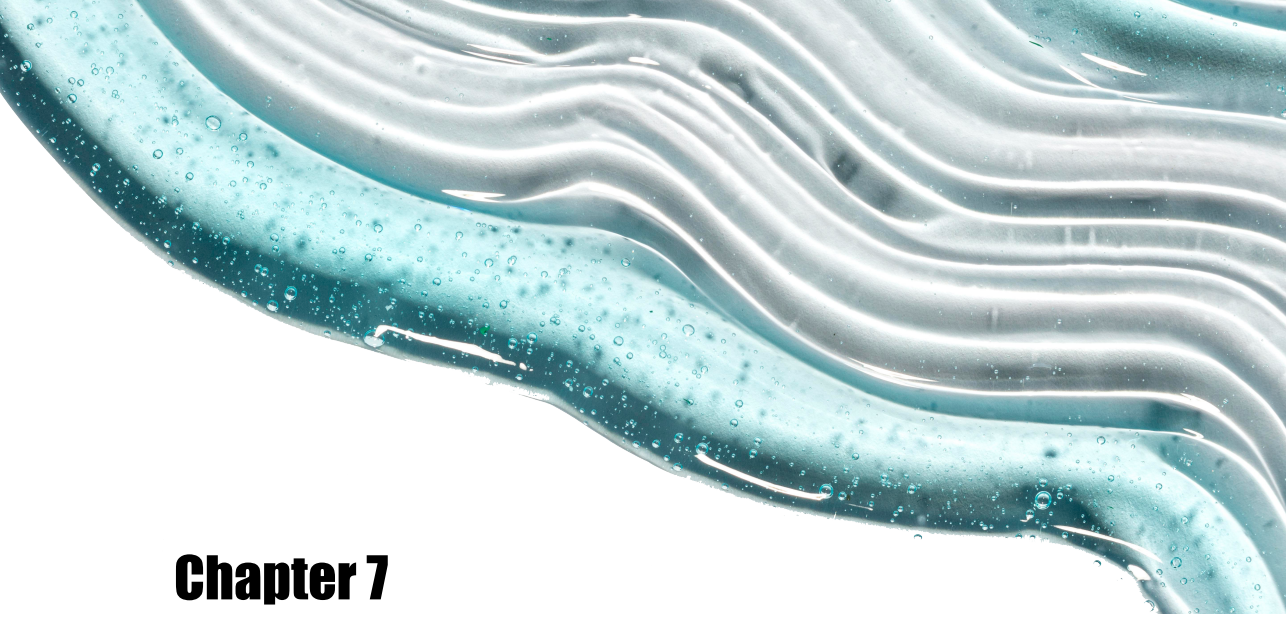
Schoen, S., Wang, M., Dayavansha, S., Naja, K., Kumar, V., Tadross, R., Pope, K., Ling, L., Hunt, D., Peters, M. K., Iafrate, A., Mercaldo, N. D., Sandstrom, K., Kim, T., Washburn, M., Pierce, T. T., and Samir, A. E. (2025). Increased mechanical index improves shear wave elastography: Pilot study of signal enhancement. *Ultrasound in Medicine & Biology*.

Song, P., Manduca, A., Zhao, H., Urban, M. W., Greenleaf, J. F., and Chen, S. (2014). Fast shear compounding using robust 2-D shear wave speed calculation and multi-directional filtering. *Ultrasound in medicine & biology*, 40(6):1343–1355. Publisher: Elsevier.

Song, P., Urban, M. W., Manduca, A., Greenleaf, J. F., and Chen, S. (2015). Coded excitation plane wave imaging for shear wave motion detection. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 62(7):1356–1372.

Wiley, B. M., Zhou, B., Pandompatam, G., Zhou, J., Kucuk, H. O., and Zhang, X. (2021). Lung ultrasound surface wave elastography for assessing patients with pulmonary edema. *IEEE Transactions on Biomedical Engineering*, 68(11):3417–3423.

Zhang, X., Zhou, B., and Zhang, A. X. (2019). A Pilot Study of Wet Lung Using Lung Ultrasound Surface Wave Elastography in an Ex Vivo Swine Lung Model. *Applied Sciences (Basel, Switzerland)*, 9(18):3923.



Chapter 7

Concluding remarks

Abstract

This chapter synthesizes the conclusions of the thesis under a pipeline perspective of ultrasound shear-wave elastography, demonstrating that reliability is enhanced when excitation, motion estimation, and reconstruction are concurrently designed under realistic constraints. The primary contributions are encapsulated within three key elements: a computationally efficient end-to-end simulator that facilitates optimized three-dimensional acoustic field computation and synthetic ultrasound frames, a local and efficient phase velocity spectral estimator derived from a phase-gradient estimator, and excitation strategies that enhance detectability by redistributing ARF energy over time through MLS-coded excitations and in the spatial domain through acoustic vortex beams. Future perspectives in this field are expected to place significant emphasis on three-dimensional reconstruction and spatiotemporal ARF coding.

Outline

7.1 Conclusions 207
7.2 Future perspectives 210

7.1 Conclusions

This thesis argues that improving the clinical dependability of ultrasound shear-wave elastography (SWE) is primarily achieved through coherent flowchart design. In particular, reliability increases when excitation, motion estimation, and reconstruction are treated as coupled stages and are validated under realistic constraints.

Desiging and validating shear wave elastography techniques

A first substantive contribution concerns a efficient end-to-end simulator for SWE. Although full-wave modeling can, in principle, connect acoustic field generation, radiation-force loading, shear-wave propagation, and reconstruction, its routine use is frequently limited by the different scales of shear and longitudinal wavelengths and different time scales of the acoustic and elastic problems, leading to a high computational cost. The framework developed in this thesis addresses this bottleneck by implementing an optimized strategy under the assumptions of shear wave elastography in soft tissues, thereby enabling protocol-level sensitivity analyses over excitation geometries and acquisition settings within practical computational budgets.

A second key novelty is the integration of synthetic ultrasound data that emulate the experimental measurement process. Specifically, the simulator generates synthetic B-mode/IQ frames in which shear-wave–induced deformations are embedded, allowing motion tracking and reconstruction to be applied to realistic image-derived observables rather than to idealized displacement fields alone. This completes the simulation flowchart by aligning the forward model with the same data pathway used in experiments, and thus enables validation of not only the elstodynamic model but also the measurement and processing stages that ultimately determine the performance of a given shear wave elasticity technique.

In this thesis, this efficient method was used to test and validate the proposed techniques, explore and tune the algorithm parameters for image reconstruction methods, and design the geometry and parameters of experimental devices prior to their manufacturing.

Efficient spectral image generation

A secondary conclusion of this thesis is that practical spectral SWE benefits from reconstruction approaches that are both local and computationally efficient. Phase-gradient elastography (PGE) contributes to this objective by providing frequency-resolved phase-velocity estimates without reliance on heavy block-based processing. While related phase-based approaches were introduced early in ultrasound vibration imaging and later formalized in elastographic reconstruction and shear-wave spectroscopy [Manduca et al., 2001, Manduca et al., 2003, Chen et al., 2004, Deffieux et al., 2009], the present work combines this local phase principle with modern angular and spectral filtering to achieve a robust and efficient reconstruction method. This enables dense pixel mapping while retaining spatial specificity. Consequently, multi-frequency characterization becomes a viable option within processing budgets compatible with iterative analysis and, potentially, real-time or near-real-time workflows.

From a methodological perspective, PGE reinterprets spectral reconstruction as a local estimation problem, wherein directional filtering and phase-gradient operators are employed to isolate wave components that are consistent with a traveling-wave interpretation at a specific angular frequency. This formulation is particularly advantageous when the wavefield exhibits spatially varying behavior, as it supports localized estimation without imposing global homogeneity assumptions over large windows.

The performance of the method depends strongly on the effectiveness of the wavefield conditioning stage. Directional and spectral filtering are essential to suppress interfering or reflected components and isolate the wave content relevant for local phase-velocity estimation. Under these conditions, PGE can operate robustly even in wavefields that are not globally planar. Its main limitation arises when when the wavefield cannot be conditioned to support a stable local phase interpretation within the region of interest. A key conclusion is that the performance of PGE depends not only on the local estimator, but also on the ability of the preceding filtering stage to isolate the wave components of interest.

Efficient motion generation strategies

A third set of conclusions concerns excitation design as a key lever for improving SWE robustness when increasing peak output is neither desirable nor feasible. Chapters 4 and 5 show that, despite their different physical implementations, acoustic vortex actuation and MLS-coded ARF excitation follow a common principle: both reshape the spatiotemporal distribution of the applied force to improve shear-wave generation without increasing peak pressure. In this thesis, MLS coding redistributes excitation in time through a pseudo-random sequence, whereas acoustic vortices redistribute it in space through a toroidal forcing pattern. In both cases, the practical outcome is improved shear-wave detectability under lower-voltage and lower-mechanical-index operation.

Chapter 4 shows that acoustic vortices provide a complementary actuation mechanism by transferring both linear and angular momentum to soft tissue. This enables torque-driven actuation and generates measurable rotational shear motion, whose direction can be reversed by changing the sign of the topological charge. This additional degree of freedom can improve detectability, but it also requires acquisition and reconstruction strategies adapted to the resulting wavefield. The results therefore suggest that successful translation will depend on actuation efficiency and on the co-design of vortex parameters with motion estimation and reconstruction methods.

The results of Chapters 5–6 support the conclusion that MLS-coded ARF pulse compression provides a practical route to enhance shear-wave detectability under output constraints. By spreading energy over time and coherently recovering it through pulse compression, the method increases effective displacement SNR and stabilizes phase-velocity estimation when conventional single-push acquisitions become detectability-limited. Importantly, this benefit extends beyond idealized conditions, remaining effective in media and scenarios affected by attenuation, heterogeneity, dispersion, and motion. A further conclusion is that code length and timing are meaningful design variables, enabling an explicit trade-off between detectability gains and temporal footprint while maintaining low peak output.

Elastography under challenging scenarios

The value of a SWE method is best judged not only in homogeneous or high-SNR conditions, but in scenarios where wave propagation, motion tracking, or

reconstruction become intrinsically difficult. Several chapters show that the proposed methods remain informative under such conditions and, in different ways, extend the range of situations in which quantitative elastography can be applied reliably.

In Chapter 3, the *ex vivo* liver experiment shows that PGE can recover layered stiffness patterns comparable to established spectral references while reducing reconstruction time by orders of magnitude. This is important not only from a computational perspective, but also because it makes repeated acquisitions, frequency compounding, and multi-frequency analysis feasible in practice. In Chapter 4, the *ex vivo* kidney study shows that vortex excitation can restore cortex–medulla differentiation under moderate-output conditions where conventional focusing remains ambiguous, thereby improving contrast recovery in weak-detectability regimes. Finally, Chapter 6 demonstrates the efficacy of MLS-coded excitation under a range of conditions, including heterogeneity, anisotropy, viscoelasticity, noise, movement, and *in vivo* settings. This excitation method preserves measurable wave content and enhances the robustness of reconstruction, particularly in scenarios where conventional excitation becomes unreliable. Taken together, these results show that the proposed methods extend the reliable use of SWE beyond ideal conditions.

7.2 Future perspectives

Two directions appear particularly promising for extending the scope of this thesis while preserving its central design philosophy: advancing toward volumetric reconstruction and integrating temporal and spatial excitation coding under practical exposure constraints.

A natural continuation of the phase-gradient framework is its extension to three-dimensional elastography. Current implementations based on conventional linear probes are inherently affected by out-of-plane motion, which can bias local phase-velocity estimates when the wavefield is not well described by an in-plane model. At the same time, volumetric ultrasound imaging is becoming increasingly practical through matrix arrays and row-column architectures. In this context, the low computational cost of the proposed reconstruction method becomes a major advantage: a 3D implementation could provide volumetric elasticity maps at clinically relevant rates while improving accuracy by incorporating the full three-component shear-wave motion. Beyond

improved precision, volumetric mapping would also enable a more faithful characterization of complex targets, where geometry and boundary conditions generate inherently three-dimensional wavefields. This extension would also open a natural path toward magnetic resonance elastography (MRE), where wave propagation and inversion are intrinsically three-dimensional and could benefit from similar local phase-based strategies.

The excitation results of this thesis also motivate a unified view of ARF design in which detectability is improved by redistributing the excitation either temporally (MLS-coded) or spatially (vortex-coded). A particularly attractive future direction is to combine both concepts into a single hybrid strategy. In this setting, the sign of the vortex topological charge could be incorporated into a binary temporal sequence by assigning the ± 1 states of the code to opposite vortex charges on a per-bit basis, thereby defining a form of spatiotemporal ARF coding. Such a strategy could further improve detectability without increasing peak output, while also providing additional control over the directional content of the induced shear-wave field.

A more ambitious but coherent continuation of this thesis is to bring these developments together within a single system: a two-dimensional imaging array capable of generating vortex-coded ARF excitations, implementing temporal MLS coding, and acquiring volumetric motion data for 3D PGE reconstruction. This would integrate the main threads of the thesis into a unified framework combining time coding, space coding, and volumetric reconstruction in a single platform. Future work should therefore focus on defining the appropriate encoding and pulse-compression strategy, optimizing the actuation field for efficient and MI-compliant operation, and assessing the compatibility of this combined approach with volumetric motion tracking and 3D reconstruction assumptions. In this way, the long-term perspective is not only to improve each stage of SWE separately, but to merge them into an integrated elastography system designed from the outset for robustness under realistic experimental and clinical conditions.

References

Chen, S., Fatemi, M., and Greenleaf, J. F. (2004). Quantifying elasticity and viscosity from measurement of shear wave speed dispersion. *The Journal of the Acoustical Society of America*, 115(6):2781–2785.

Deffieux, T., Montaldo, G., Tanter, M., and Fink, M. (2009). Shear wave spectroscopy for in vivo quantification of human soft tissues visco-elasticity. *IEEE Transactions on Medical Imaging*, 28(3):313–322.

Manduca, A., Lake, D. S., Kruse, S. A., and Ehman, R. L. (2003). Spatio-temporal directional filtering for improved inversion of mr elastography images. *Medical Image Analysis*, 7(4):465–473.

Manduca, A., Oliphant, T. E., Dresner, M. A., Mahowald, J. L., Kruse, S. A., Amromin, E., Felmlee, J. P., Greenleaf, J. F., and Ehman, R. L. (2001). Magnetic resonance elastography: non-invasive mapping of tissue elasticity. *Medical Image Analysis*, 5(4):237–254.

Yamakoshi, Y., Sato, J., and Sato, T. (1990). Ultrasonic imaging of internal vibration of soft tissue under forced vibration. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 37(2):45–53.

Appendix A

Kelvin–Voigt model

This appendix summarizes the Kelvin–Voigt viscoelastic model used by the k-Wave elastic solver in § 2.2.3 and derives the dispersion relations. We obtain the frequency dispersion and absorption of a plane wave in a Kelvin-Voigt medium. The derivation follows a wavenumber-based approach, in which the complex wavenumber is decomposed into propagating and absorbing contributions.

Complex wavenumber

For a monochromatic plane wave with angular frequency ω, the spatial dependence can be written in terms of a complex wavenumber $k(\omega)$ as

$$u(x,t) = U\, e^{i\omega t - ik(\omega)x}. \tag{A.1}$$

The complex wavenumber is decomposed as

$$k(\omega) = k_r(\omega) + i\,\alpha(\omega), \tag{A.2}$$

so that $u(x,t) = U\, e^{i\omega t - ik_r x} e^{-\alpha x}$. The real part $k_r(\omega)$ governs the propagating oscillatory component and the imaginary part $\alpha(\omega) \geq 0$ defines the absorption coefficient in Np/m. The phase velocity is defined by

$$c_p(\omega) = \frac{\omega}{k_r(\omega)}. \tag{A.3}$$

The group velocity is defined as the derivative of the angular frequency with respect to the real wavenumber,

$$c_g(\omega) = \frac{d\omega}{dk_r(\omega)}. \tag{A.4}$$

Momentum balance

Let $\mathbf{u}(\mathbf{r}, t)$ be the displacement, $\mathbf{v} = \partial\mathbf{u}/\partial t$ the particle velocity, ρ_0 the density, and $\boldsymbol{\sigma}(\mathbf{r}, t)$ the symmetric Cauchy stress tensor. Linear momentum conservation can be written compactly

$$\rho_0 \frac{\partial \mathbf{v}}{\partial t} = \nabla \cdot \boldsymbol{\sigma}. \tag{A.5}$$

Constitutive relations for a Kelvin–Voigt medium

Under infinitesimal deformations, the strain tensor and its time derivative are written in tensor form as

$$\boldsymbol{\varepsilon} = \mathrm{sym}(\nabla \mathbf{u}), \tag{A.6}$$

$$\frac{\partial \boldsymbol{\varepsilon}}{\partial t} = \mathrm{sym}(\nabla \mathbf{v}). \tag{A.7}$$

where $\mathrm{sym}(\mathbf{A}) = \frac{1}{2}\left(\mathbf{A} + \mathbf{A}^{\mathbf{T}}\right)$ denotes the symmetric part of a second-order tensor.

The Kelvin–Voigt model assumes that the stress is the sum of an elastic contribution proportional to strain and a viscous contribution proportional to strain rate. For an isotropic viscoelastic solid, the constitutive law is

$$\boldsymbol{\sigma} = \lambda\, \mathrm{tr}(\boldsymbol{\varepsilon})\, \mathbf{I} + 2\mu\, \boldsymbol{\varepsilon} + \chi\, \mathrm{tr}\left(\frac{\partial \boldsymbol{\varepsilon}}{\partial t}\right) \mathbf{I} + 2\eta\, \frac{\partial \boldsymbol{\varepsilon}}{\partial t}, \tag{A.8}$$

where λ and μ are the elastic Lamé parameters and χ and η are the Kelvin–Voigt viscosity coefficients. The trace terms $\mathrm{tr}(\boldsymbol{\varepsilon})$ and $\mathrm{tr}(\dot{\boldsymbol{\varepsilon}})$ describe the volumetric response, whereas the terms proportional to $\boldsymbol{\varepsilon}$ and $\dot{\boldsymbol{\varepsilon}}$ describe the deviatoric response responsible for shear-wave propagation and dissipation.

Differentiating (A.8) and using (A.7) yields the stress–rate form used by the k-Wave elastic solver,

$$\dot{\boldsymbol{\sigma}} = \lambda\, \nabla\cdot \mathbf{v}\, \mathbf{I} + 2\mu\, \mathrm{sym}(\nabla \mathbf{v}) + \chi\, \frac{\partial}{\partial t}(\nabla\cdot \mathbf{v})\, \mathbf{I} + 2\eta\, \frac{\partial}{\partial t}[\mathrm{sym}(\nabla \mathbf{v})]\,. \tag{A.9}$$

Equations (A.5) and (A.9) define a closed first-order system in $(\boldsymbol{\sigma}, \mathbf{v})$. The dispersion relations derived below correspond to the divergence-free shear mode $\nabla\cdot \mathbf{v} = 0$, for which the volumetric terms proportional to (λ, χ) do not contribute and the response is governed solely by the shear parameters (μ, η), consistently with the discussion in § 2.2.3.

Plane shear wave in Kelvin–Voigt media

A plane shear wave propagating along x is introduced as transverse particle motion polarized along z direction with no dependence on y or z directions,

$$\mathbf{u}(x,t) = u(x,t)\, \mathbf{e}_z, \qquad \mathbf{v}(x,t) = v(x,t)\, \mathbf{e}_z, \qquad \frac{\partial(\cdot)}{\partial y} = 0, \qquad \frac{\partial(\cdot)}{\partial z} = 0, \tag{A.10}$$

which is a specialization of the kinematic definitions (A.6)–(A.7). Under (A.10), $\nabla\cdot \mathbf{v} = 0$, so the volumetric terms proportional to λ and χ in (A.8) and (A.9) do not contribute to this shear-wave mode. Moreover, the only stress component involved in the z-polarized propagation along x is σ_{xz}, so that the momentum balance (A.5) reduces to

$$\rho_0\, \frac{\partial v}{\partial t} = \frac{\partial \sigma_{xz}}{\partial x}, \tag{A.11}$$

and the Kelvin–Voigt stress–rate relation (A.9) reduces to the shear component

$$\frac{\partial \sigma_{xz}}{\partial t} = \mu\, \frac{\partial v}{\partial x} + \eta\, \frac{\partial^2 v}{\partial x\, \partial t}. \tag{A.12}$$

Combining (A.11) and (A.12) yields the Kelvin–Voigt shear-wave equation in velocity form

$$\rho_0\, \frac{\partial^2 v}{\partial t^2} = \mu\, \frac{\partial^2 v}{\partial x^2} + \eta\, \frac{\partial^3 v}{\partial x^2\, \partial t}. \tag{A.13}$$

Derivation of the dispersion relation

A monochromatic plane-wave solution as in (A.1) is introduced in (A.13), yielding

$$\rho_0\,\omega^2 = (\mu + i\omega\eta)\,k^2, \qquad k^2(\omega) = \frac{\rho_0\,\omega^2}{\mu + i\omega\eta}. \tag{A.14}$$

Note the Kelvin–Voigt complex shear modulus is usually defined as $\mu^*(\omega) = \mu + i\omega\eta$, so (A.14) can be written as $k^2 = \rho_0\omega^2/\mu^*(\omega)$. Multiplying by the complex conjugate gives

$$k^2(\omega) = \frac{\rho_0\,\omega^2}{\mu^2 + \omega^2\eta^2}\,(\mu - i\omega\eta)\,. \tag{A.15}$$

The separation of $k(\omega)$ into propagation and attenuation parts follows from the standard square-root identities applied to $k^2(\omega)$,

$$k_r^2(\omega) = \frac{\left|k^2(\omega)\right| + \Re\{k^2(\omega)\}}{2}, \qquad \alpha^2(\omega) = \frac{\left|k^2(\omega)\right| - \Re\{k^2(\omega)\}}{2}, \tag{A.16}$$

where

$$\Re\{k^2(\omega)\} = \frac{\rho_0\,\omega^2\,\mu}{D(\omega)}, \qquad \left|k^2(\omega)\right| = \frac{\rho_0\,\omega^2}{\sqrt{D(\omega)}}, \tag{A.17}$$

and

$$D(\omega) = \mu^2 + \omega^2\eta^2. \tag{A.18}$$

Substituting (A.17) into (A.16) yields

$$k_r^2(\omega) = \frac{\rho_0\,\omega^2}{2D(\omega)}\left(\mu + \sqrt{D(\omega)}\right), \tag{A.19}$$

$$\alpha^2(\omega) = \frac{\rho_0\,\omega^2}{2D(\omega)}\left(\sqrt{D(\omega)} - \mu\right). \tag{A.20}$$

Phase velocity and attenuation

The Kelvin–Voigt shear-wave dispersion is fully characterized by the frequency dependence of the phase velocity $c_p(\omega)$ and the attenuation coefficient $\alpha(\omega)$. From (A.3) and (A.19), the phase velocity is

$$c_p(\omega) = \sqrt{\frac{2\,(\mu^2 + \omega^2\eta^2)}{\rho_0\left(\mu + \sqrt{\mu^2 + \omega^2\eta^2}\right)}}, \tag{A.21}$$

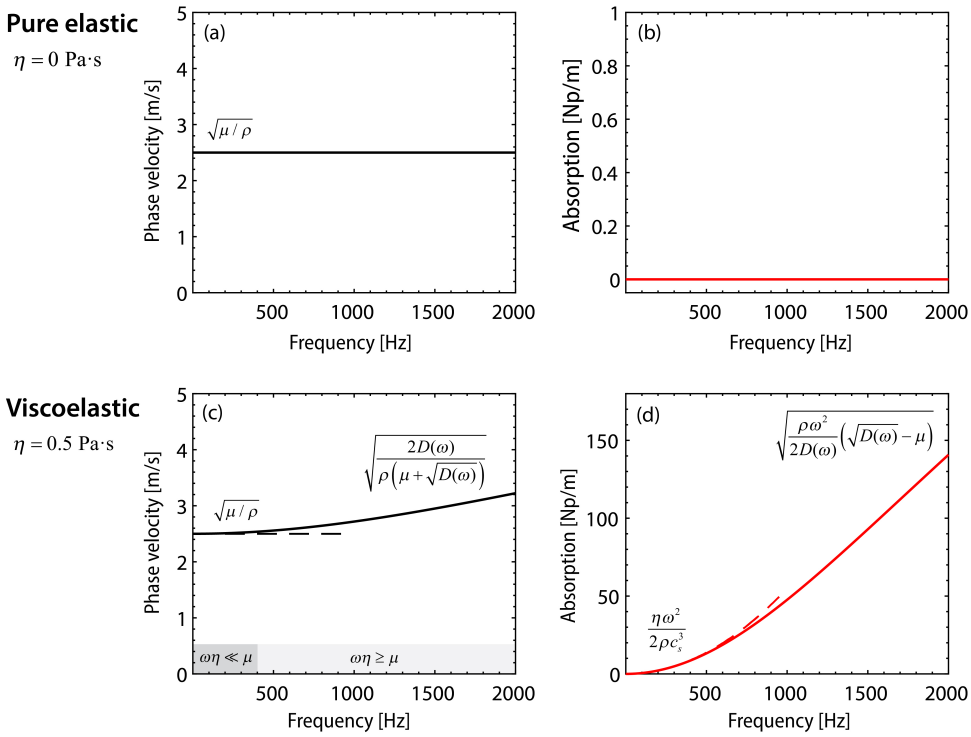


Figure A.1: Kelvin–Voigt shear-wave dispersion and absorption for elastic and viscoelastic media. (a)–(b) Purely elastic case ($\eta = 0$), showing a frequency-independent phase velocity $c_p = \sqrt{\mu/\rho_0} = c_s$ and zero attenuation $\alpha = 0$. (c)–(d) Viscoelastic Kelvin–Voigt case ($\eta \neq 0$), where (A.21) and (A.22) predict a monotonic increase of phase velocity with frequency and a frequency-dependent attenuation. The low-frequency regime $\omega\eta \ll \mu$ corresponds to the initial plateau $c_p \approx c_s$ in (c), while the leading-order approximation $\alpha(\omega) \approx \eta\omega^2/(2\rho_0 c_s^3)$ describes the initial quadratic growth of absorption in (d) before higher-frequency deviations appear.

and the corresponding attenuation coefficient follows from (A.20) as

$$\alpha(\omega) = \omega \sqrt{\frac{\rho_0}{2\left(\mu^2 + \omega^2\eta^2\right)}\left(\sqrt{\mu^2 + \omega^2\eta^2} - \mu\right)}. \tag{A.22}$$

Figure A.1 summarizes the qualitative trends implied by (A.21)–(A.22). In the purely elastic limit $\eta = 0$, the wavenumber becomes purely real and proportional to frequency, so the medium is non-dispersive. The phase velocity is constant and reduces to the shear speed $c_s = \sqrt{\mu/\rho_0}$, while the attenuation vanishes, $\alpha(\omega) = 0$, as shown in Fig. A.1(a)–(b). In this same limit, the

linear relation $k_r(\omega) = \omega/c_s$ also implies a constant group velocity, hence $c_g(\omega) = c_s$ and $c_p(\omega) = c_g(\omega) = c_s$. When $\eta \neq 0$, the viscous term introduces dispersion and absorption: the phase velocity increases with frequency and the attenuation rises from zero, as illustrated in Fig. A.1(c)–(d). These effects become more pronounced as the measurement bandwidth extends toward higher frequencies, because the ratio $\omega\eta/\mu$ increases across the spectrum.

The low-frequency regime is obtained under the asymptotic condition $\omega\eta \ll \mu$. Introducing the dimensionless parameter $\varepsilon(\omega) = \omega\eta/\mu$ and factoring out μ, the Taylor approximation $\sqrt{1+\varepsilon^2} \approx 1 + \frac{1}{2}\varepsilon^2$ yields the leading-order behaviour. For the phase velocity, the purely elastic shear speed is recovered,

$$c_p(\omega) \approx \sqrt{\frac{\mu}{\rho_0}} = c_s, \tag{A.23}$$

which explains the low-frequency plateau in Fig. A.1(c). In this regime, the associated wavenumber is approximately linear in frequency and therefore the group velocity also approaches the same constant value, $c_g(\omega) \approx c_s$.

In contrast, the attenuation in (A.22) involves the term $\sqrt{\mu^2 + \omega^2\eta^2} - \mu$, which is identically zero in the purely elastic limit; consequently, the leading-order behaviour must be obtained from the first non-vanishing term in its Taylor expansion. Using $\sqrt{1+\varepsilon^2} - 1 \approx \frac{1}{2}\varepsilon^2$ yields the quadratic dependence on frequency

$$\alpha(\omega) \approx \frac{\eta\,\omega^2}{2\rho_0\,c_s^3}, \tag{A.24}$$

consistent with the near-zero absorption and the initial curvature observed in Fig. A.1(d). As the bandwidth extends toward frequencies where $\omega\eta \gtrsim \mu$, higher-order terms in ε become non-negligible and the full expressions (A.21)–(A.22) are required to describe the stronger dispersion and absorption at higher frequencies.

Research projects and funding

This doctoral thesis was supported by the *Agencia Valenciana de la Innovación* through the transfer projects *Dispositivo de imagen elastográfica cuantitativa empleando vórtices acústicos* (INNVA1/2020/92) and *Prototipo de litotricia extracorpórea por vórtices acústicos* (INNVA1/2022/37); by the UPV R&D line *Línea I+D Imagen, Caracterización y Terapia con Ultrasonidos* (C17217); by the UPV R&D line *Línea I+D Imagen, Diseño y Construcción de un Litotriptor Ultrasónico basado en Vórtices Acústicos* (C17889); by the *Agencia Estatal de Investigación* through the R&D line *Línea I+D Hologramas caleidoscópicos para imagen y terapia en el cerebro empleando ultrasonidos de bajo coste* (20236CNS293); and by the *Ministerio de Ciencia, Innovación y Universidades* through the R&D grant *Procesos de calibración para detectores de imagen molecular* (PTA2024-024684-I).

Part of this thesis was also developed during a predoctoral research stay at *Mayo Clinic* in 2024, supported by the competitive mobility grant *iMOVE* (IMOVE24009), funded by the *Consejo Superior de Investigaciones Científicas*. This stay was particularly relevant for the work on MLS-coded excitation and its evaluation in challenging scenarios presented in Chapters 5 and 6.