Visual Surface Wave Elastography: Revealing Subsurface Physical Properties via Visible Surface Waves

Alexander C. Ogren* Berthy T. Feng* Jihoon Ahn Katherine L. Bouman Chiara Daraio
California Institute of Technology

Abstract

Wave propagation on the surface of a material contains information about physical properties beneath its surface. We propose a method for inferring the thickness and stiffness of a structure from just a video of waves on its surface. Our method works by extracting a dispersion relation from the video and then solving a physics-based optimization problem to find the best-fitting thickness and stiffness parameters. We validate our method on both simulated and real data, in both cases showing strong agreement with ground-truth measurements. Our technique provides a proof-of-concept for at-home health monitoring of medically-informative tissue properties, and it is further applicable to fields such as human-computer interaction.

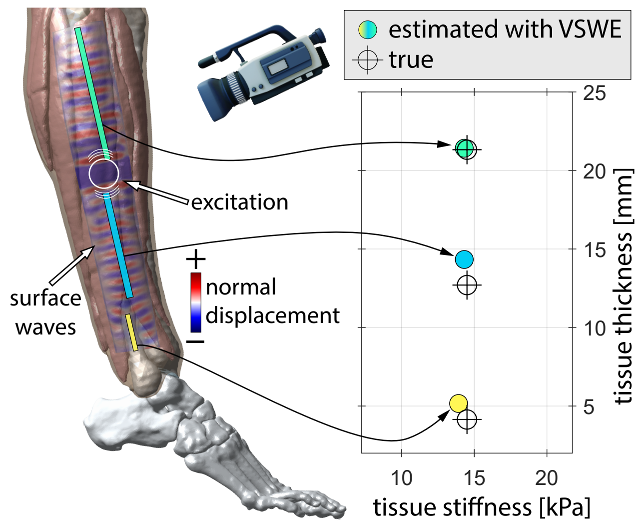

Figure 1. Estimating subsurface tissue properties of a human leg with VSWE. From a video of surface wave motion, VSWE estimates the thickness and stiffness of the soft tissue layer. In this example, we are able to recover the three different thicknesses of the three regions highlighted on the leg. We also recover the stiffness of the leg tissue, which does not change across the leg.

1. Introduction

How waves propagate on the surface of a medium reveals information about its subsurface properties. For example, by watching how ocean waves evolve and break as they near the shore, one can infer the rises and dips of the seafloor below [49]. Waves can also be observed on the surfaces of biological systems. Imagine applying a massage gun to your calf. The ripples on your skin convey information about the underlying layers of fat, muscle, and bone. In fact there is a well-defined relationship linking the thickness and stiffness of each layer to the wave propagation behavior. We propose *Visual Surface Wave Elastography* (VSWE), a physics-based method to estimate the thickness and stiffness of a medium from a video of waves on its surface.

We are primarily motivated by the task of biological tissue characterization, which has broad applications including at-home health monitoring. For example, tumors [26, 30], musculoskeletal degeneration [31, 44, 45], and liver disease [23, 42] often lead to changes in tissue thickness or stiffness. However, existing techniques for elastography (i.e., measuring tissue stiffness) require specialized ultrasound [43, 50] and sometimes magnetic resonance [34] devices, along with medical experts to operate the equipment [7, 12]. We show that it is possible to obtain coarse estimates of thickness and stiffness given surface waves measured with just a video camera. Another application is human computer interaction (HCI), where visually inferring subsurface changes in muscle stiffness may unlock new modes of gesture recognition. In biomechanics, many regions of the body are often modeled as a series of layers of biological tissue, such as skin, muscle, or bone [4, 22]. Similarly, in this work, we model the medium as a layer of soft tissue with unknown thickness and stiffness atop a layer of bone. We leverage the mathematical relationship between the geometric and mechanical properties of the soft layer and the wave propagation behavior of the medium.

The wave propagation behavior of a material is compactly described by a mathematical object known as a *dispersion relation*. Under some common biomechanical as-

*equal contribution

sumptions in our layer model (discussed further in Sec. 3.1), the thickness and stiffness of the soft layer fully determine the dispersion relation. The main idea of our method is to find the thickness and stiffness values that lead to a dispersion relation that matches the dispersion relation extracted from surface waves observed in the video.

In this paper, we first discuss related work on video-based and wave-based material characterization and provide background on wave mechanics. We then describe our method, which consists of taking a video of surface waves, extracting the dispersion relation from the video, and then solving an optimization problem to find the best-fitting thickness and stiffness parameters. We validate our method on real and simulated data, including real videos of gelatin-based phantoms that mimic biological tissue and realistic simulations of a human leg.

2. Related work

2.1. Video-based material characterization

Videos encode rich information about the environment and physical objects. For example, previous work has leveraged surface vibrations [14] or sound [13, 19] in video to recover unknown information. For material characterization tasks, the way an object moves in a video provides useful information about its physical properties. Videos have been used to estimate material parameters of specific object classes, such as fabrics [8, 11, 35], rods [15, 16], and trees [47], whose dynamics are entirely visible. Going a step further, Visual Vibration Tomography (VVT) [18] infers subsurface material properties. Specifically, it recovers the spatially-varying stiffness and density throughout a 3D object with known geometry by analyzing its global vibrational modes.

In contrast to vibrational modes that exist at discrete frequencies, surface waves contain information over a continuous range of frequencies and wavevectors. Surface waves have been under-utilized for video-based material characterization, even though they can also be observed in video and contain useful information about underlying physical properties [51]. Our work shows how to leverage surface waves in video to estimate geometric and mechanical properties. A surface-wave-based approach is especially beneficial when one wishes to analyze local structure without having to model the global structure (as would be required by VVT), which may be complicated or have unknown geometry. In the case of biological tissue characterization, for example, it would be impractical to model the 3D vibrational modes of the entire human anatomy in order to estimate local tissue properties. Our surface-wave-based approach circumvents the need to model a complex geometry and solve for its global modes by targeting local regions where we can analyze wave modes of a simpler geometry.

2.2. Wave-based material characterization

In general, wave-based imaging relies on an understanding of how physical characteristics of interest affect wave propagation. For example, magnetic resonance imaging [25] uses knowledge of how radio-frequency waves are absorbed and re-emitted by different types of tissues in the body. Optical tomography methods leverage near-infrared light measurements at multiple positions to recover a 3D map of tissue absorption [5, 33]. Ultrasound techniques, which are pervasive in non-destructive testing [29], medical imaging [6], and wearable technology [27], use high-frequency mechanical bulk waves to characterize tissue and locate features. Among tissue characterization techniques, transient elastography [50], shear wave elastography [43], and magnetic resonance elastography [34] leverage knowledge of bulk wave physics to estimate the elasticity of tissues and organs. However, such methods often require not only high-end, expensive equipment, but also trained medical specialists [7, 12], making regular screening infeasible. In contrast, our work leverages surface waves, which are generally less expensive to observe than their bulk counterparts.

Although very different fields from our application of interest, geophysics and seismic imaging provide insights into the possibility of using surface waves to infer subsurface features [24, 32, 36, 37, 40, 52]. Inspired by the success of surface-wave methods in seismology, we propose to harness surface waves observed in video to infer subsurface tissue properties. Whereas previous work [9] suggested an approach using sensors sparsely placed on the skin, we leverage dense visual data. Vision-based tissue characterization would enable the next generation of health monitoring systems that take advantage of the ubiquity of visual sensors.

3. Background

3.1. Dispersion relations

Just as sound waves can be expressed as the combination of simple harmonic modes, waves traveling through any medium can be expressed as the combination of wave modes of different spatial and temporal frequencies. A *dispersion relation* compactly describes wave propagation by defining all the possible wave modes of a medium. Specifically, it defines the *wavevector* and *frequency* of each possible wave mode, where the wavevector and frequency indicate the spatial and temporal rates of oscillation, respectively.

Mathematically, the dispersion relation can be determined by solving the harmonic elastic wave equation (a PDE) subject to phase-shifted periodic boundary conditions. Let $\mathbf{x} = [x, y, z]^\intercal$ denote the spatial location and $\mathbf{u}(\mathbf{x}) = [u(\mathbf{x}), v(\mathbf{x}), w(\mathbf{x})]^\intercal$ denote the displacement at $\mathbf{x}$. Assuming an isotropic linear-elastic material, the harmonic elastic wave equation can be written as

$$-\omega^2 \mathcal{M}(\mathbf{u})(\mathbf{x}) = \mathcal{K}(\mathbf{u})(\mathbf{x}) \qquad \forall \mathbf{x} \in \Omega, \tag{1}$$

where $\mathcal{M}$ and $\mathcal{K}$ denote linear operators representing the mass and stiffness of the medium. The mass operator $\mathcal{M}$ depends on the density field $\rho(\mathbf{x})$. The stiffness operator $\mathcal{K}$ depends on the elastic modulus $E(\mathbf{x})$ and Poisson's ratio $\nu(\mathbf{x})$. The eigenvector solution $\mathbf{u}(\mathbf{x})$ describes the shape of the wave mode, and the eigenfrequency ω defines the temporal frequency of the wave. Note that the solutions also depend on the domain Ω, which in our case is parameterized by the thickness and length of the tissue layer.

When solving Eq. (1), boundary conditions are necessary to impose real-life assumptions of the physical system. In our setting, we assume that Ω is a finite subregion of the full wave medium (e.g., a section of a leg). Furthermore, to simplify our analysis, we target waves traveling in one direction (i.e., along the x direction). This leads us to apply 1D periodic boundary conditions on the domain $x \in [0, a]$ to impose the assumption that the wave medium continues past the boundary of Ω.

Specifically, we apply the Bloch-Floquet [10, 21] (a.k.a. phase-shifted) periodic boundary conditions, which define a boundary condition for every wavenumber $\gamma \in [0, \pi/a]$. We note that in general γ is a wavevector, but in the case of 1D wave analysis, it is a wavenumber that indicates the number of wavelengths per unit length. For a given γ, the boundary condition is defined as

$$\mathbf{u}(x = a, y, z) = \mathbf{u}(x = 0, y, z)e^{\mathbf{i}\gamma a}. \tag{2}$$

In words, the solution at one end ($x = 0$) only differs from the solution at the other end ($x = a$) by the phase shift γa.

To compute a dispersion relation, we solve Eq. (1) subject to Eq. (2) for $\gamma \in [0, \pi/a]$. For each γ, we solve a generalized eigenvalue problem and obtain multiple $(\omega, \mathbf{u}(\mathbf{x}))$ solution pairs. The dispersion relation is exactly the set of all ω solutions for every γ. Supplementary details on solid mechanics may be found in Appendix D.

3.2. Assumptions for tissue characterization

For biological tissue characterization, we can model the medium as a layer of soft tissue atop a hard bone layer, where the top layer has uniform thickness T. We note that throughout this work, we refer to the elastic modulus of the tissue as its stiffness. We make several assumptions to reduce the complexity of the problem:

1. The soft layer has a uniform stiffness $E(\mathbf{x}) = E$.
2. The density and Poisson's ratio of the soft tissue are known. We set $\rho = 1\,\mathrm{g\,cm^{-3}}$ [39] and $\nu = 0.45$ [28].
3. The bone is much stiffer than the soft tissue. This means neither the thickness nor the exact stiffness of the bone layer matters, and we can model it as motionless.

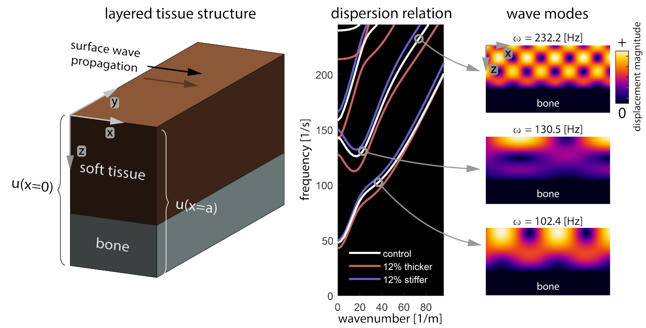

Figure 2. The dispersion relation depends on both the thickness and the stiffness of the soft tissue. We show how the dispersion relation changes when the soft tissue is made slightly thicker or slightly stiffer, with examples of the wave modes that are embedded in each dispersion relation. Wave mode dynamics occurring beneath the surface affect their expression on the surface.

These assumptions are common in biomechanical analysis [4, 22]. Under these assumptions, the stiffness operator $\mathcal{K}$ in Eq. (1) only depends on the elastic modulus E (a.k.a. stiffness) of the soft tissue, and the spatial domain Ω on which we solve Eq. (1) only depends on the thickness T. This mean we can fully determine the dispersion relation based on T and E, so we denote the dispersion relation as

$$\mathfrak{D}(T, E) = \{\omega_i(\gamma)\}_{i=1}^{N} \quad \forall \gamma \in [0, \pi/a], \tag{3}$$

where N is the number of branches, or the number of eigenvalue solutions computed for Eq. (1).

The dispersion relation is only concerned with the spatial (γ) and temporal (ω) frequencies of the mode shapes rather than the mode shapes $\mathbf{u}(\mathbf{x})$ themselves. From waves that express themselves, even in part, on the surface, we can extract spatial and temporal frequency information to obtain a dispersion relation. Fig. 2 shows how the dispersion relation changes with the thickness and stiffness of the soft tissue layer. *Our method leverages the sensitivity of the dispersion relation to small changes in these parameters.*

4. Methods

This section describes the computational method of VSWE. The input is a video of motion on the surface of the medium of interest, and the output is the estimated thickness and stiffness of the medium. In between there are two broad stages: (1) extracting a dispersion relation from the video and (2) solving for the tissue properties that best agree with the observed dispersion relation. Fig. 3 presents an overview of the pipeline.

4.1. Extracting a dispersion relation from video

4.1.1. Motion extraction

The first step is to quantify the image-space displacements in the video. Since surface wave motion tends to be small, we use phase-based motion processing [46], which is sensitive to sub-pixel displacements. Specifically, we compute

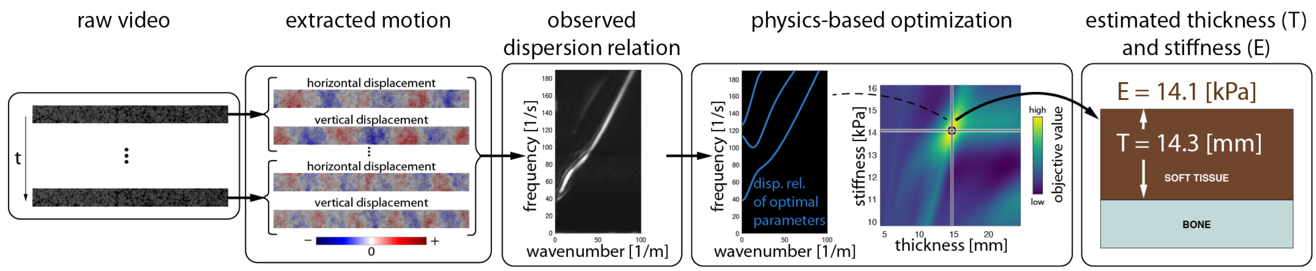

Figure 3. Method overview. Given a video of the surface of the medium of interest, we first extract motion fields in image-space. From these we extract the observed dispersion relation, which dictates the spatial and temporal frequencies of waves that travel through the medium. We solve a physics-based optimization problem to estimate the thickness and stiffness values that lead to a theoretical dispersion relation that best agrees with the observed dispersion relation.

local phase shifts in a complex steerable pyramid [41] and then convert these to pixel displacements [20]. The result is the image-space horizontal displacement $\tilde{u}(\tilde{x}, \tilde{y}, t)$ and vertical displacement $\tilde{v}(\tilde{x}, \tilde{y}, t)$ at each pixel $(\tilde{x}, \tilde{y})$ and video frame t, where the displacements are relative to the first frame. (Image-space coordinates are notated with a tilde to avoid confusion with world-space coordinates.)

4.1.2. Estimating the dispersion relation

We process the image-space surface displacements into a dispersion relation via the fast Fourier transform (FFT). We assume that the waves are traveling in the horizontal direction in image-space. For each row of pixels in the $\tilde{u}(\tilde{x}, \tilde{y}, t)$ video, we take a 2D FFT, transforming the space dimension $\tilde{x}$ to the wavenumber dimension γ and the time dimension t to the frequency dimension ω. That is, for the row where $\tilde{y} = \tilde{y}_i$, we convert the $\tilde{u}(\tilde{x}, \tilde{y} = \tilde{y}_i, t)$ signal to the complex-valued signal $\widehat{\tilde{u}}(i)(\gamma, \omega)$, whose magnitude gives the dispersion relation. To improve the signal-to-noise ratio, we average across all the rows and both displacement directions, resulting in the observed dispersion relation $\mathbf{D}_{\text{obs}}$:

$$\mathbf{D}_{\text{obs}} := \frac{1}{2H} \sum_{i=1}^{H} \left(\left| \widehat{\tilde{u}}(i) \right| + \left| \widehat{\tilde{v}}(i) \right| \right), \tag{4}$$

where H is the number of rows in the image. Fig. 4 illustrates this process, showing how the 2D FFT pulls out spatial and temporal frequency content.

The main challenge is that the observed dispersion relation likely does not fully agree with the true theoretical dispersion relation of the medium. First of all, it is incomplete: wave modes that do not prominently express on the surface or require large amounts of energy to propagate may not appear strongly in the video, and frequencies above the Nyquist sampling rate of the video cannot be captured. Second of all, the extracted motion may be noisy due to camera noise or spurious motion (from, e.g., lights flickering or rigid-body motion). VSWE has some margin of error due to the discrepancy between the observed and true dispersion relations, although we show in our experiments in Sec. 5.1 it is sensitive to 5% changes in the true parameters.

4.2. Estimating the thickness and stiffness

The next stage is to find the geometric and mechanical parameters that best explain the observed dispersion relation. We do this via physical simulation and optimization. For a given hypothesized thickness T and stiffness E, we employ the finite element method (FEM) to numerically compute the dispersion relation. We wrote specialized FEM code to do so efficiently, which can be accessed at github/aco8ogren/tissue-dispersion. The goal is to maximize an objective function that rewards similarity between $\mathfrak{D}(T, E)$ and $\mathbf{D}_{\text{obs}}$.

After testing various objective functions, we found that treating dispersion relations as images and using the structural similarity (SSIM) index [48] to work best. That is, we aim to maximize the SSIM between the images of the observed and proposed dispersion relations. The observed dispersion relation $\mathbf{D}_{\text{obs}}$, since it is derived via a 2D FFT, is already represented as an image. However, the physics-based dispersion relation $\mathfrak{D}(T, E)$ is computed as a set of curves. We transform $\mathfrak{D}(T, E)$ into an image $\mathbf{D}_{\text{hyp}}(T, E)$ by assigning intensities with a Gaussian kernel based on the distance from each (γ, ω) pixel to the curves. Then we solve the following optimization problem:

$$T^*, E^* = \underset{T,E}{\arg\max}\, \text{SSIM}\left(\mathbf{D}_{\text{hyp}}(T, E), \mathbf{D}_{\text{obs}}\right). \tag{5}$$

We solve Eq. (5) with a grid search over possible thickness and stiffness values, although more efficient approaches can be taken for computationally-constrained settings.

5. Results

We validate our approach on real and simulated data. We demonstrate remarkable accuracy on real gelatin samples, and we demonstrate recovering spatially-varying thickness and stiffness across a simulated 3D human leg.

5.1. Simulated plane strain

We created a two-layer tissue model in COMSOL [1] that would allow us to simulate samples with ground-truth geometric and mechanical parameters. We set stiffness values

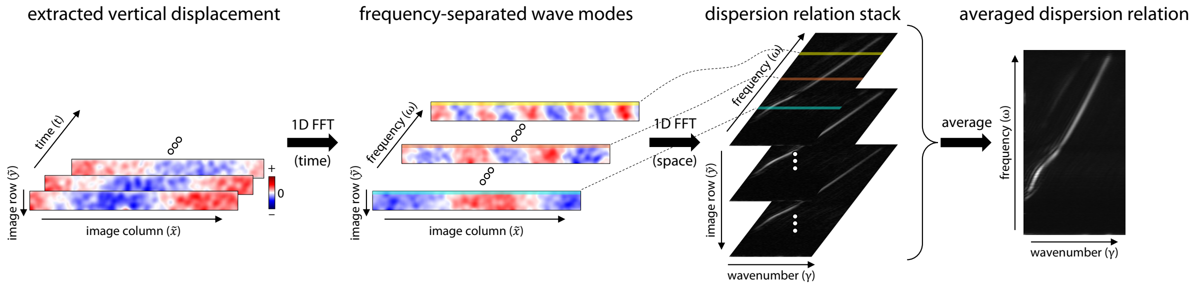

Figure 4. Obtaining a dispersion relation from image-space motion. Here we demonstrate with the vertical displacements taken from a real video of a gelatin sample. Taking a 1D FFT across time for every pixel allows us to separate wave modes by frequency (ω). Taking another 1D FFT across space (in the $\tilde{x}$ direction) decomposes the wave modes into spatial frequencies (γ). The result is a 2D FFT representing the dispersion relation. We average the dispersion relations across image rows and both displacement directions ($\tilde{u}$ and $\tilde{v}$).

on the order of $10\ \mathrm{kPa}$, which is a reasonable range for soft tissue [17]. We simulated the model's response to a chirp excitation signal applied to the leftmost side of the surface of the sample. The simulations were done with the plane-strain assumption, which allows for modeling 3D dynamics with only two dimensions by assuming that strain is zero in the z direction (more details in Appendix D). We used the simulated horizontal and vertical displacements, sampled at 600 Hz and $\sim 4300\ \mathrm{samp/m}$ over 1 s and $29.5\mathrm{cm}$, to obtain dispersion relations.

We tested the sensitivity of our method by applying slight changes to the true thickness and stiffness of the simulated sample. That is, for a certain thickness and stiffness, we can simulate samples assuming those parameters as well as perturbations of those parameters and assess whether VSWE picks up on these slight changes. As an experiment, we simulated samples at 5 distinct thicknesses and 4 distinct moduli. For each thickness-stiffness pair, we looked at 5% and 10% perturbations of either parameter. This led to 9 clusters with 9 simulated samples each. As Fig. 5 shows, within each cluster, the estimated parameters track with changes in the true parameters, even when the true parameters have changed by only 5%.

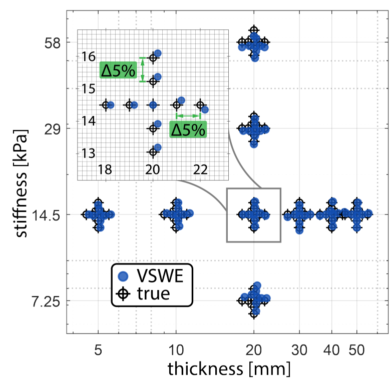

Figure 5. Sensitivity analysis. We applied VSWE to plane-strain simulations with a range of true thickness and stiffness values, along with slight perturbations of those values. In each cluster, we perturbed the true parameters by $\pm 5\%$ and $\pm 10\%$ from the central values and found that VSWE was sensitive to these changes.

5.2. Real videos of gelatin-based phantoms

To test our method on real data, we created gelatin-based phantoms (i.e., samples mimicking biological tissue) of varying thicknesses. We poured varying amounts of gelatin (1000, 1100, and $1500\ \mathrm{mL}$) into the same-sized container, leading to three different thicknesses. We set the samples in the refrigerator for about 24 hours. Once they were set, we sprinkled garlic powder onto the samples to create texture for motion extraction.

For each set sample, we measured the thickness with calipers, and we used rheometry to obtain ground-truth stiffness values. To obtain videos, we applied a shaker at one end of the sample to excite waves with a chirp signal and recorded the surface of the sample with a high-speed camera at 600 FPS at a resolution of 96×1280 pixels (each video was about four seconds long, with an observation extent of about $30\mathrm{cm}$). Fig. 6 shows pictures of the experiment setup, and more details can be found in Appendix F.

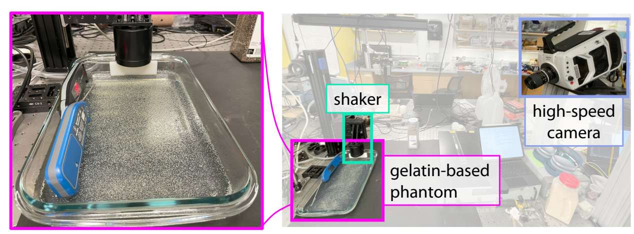

Figure 6. Experiment setup. A shaker was applied on one side of the gelatin-based phantom to excite waves in the medium. A high-speed camera, zoomed in on the phantom's surface, captured videos of the response to the shaker. As the left picture shows, we took temperature recordings with a thermometer.

For each of the three samples, we took many rheometry measurements and videos over the course of about one hour after removing the sample from the refrigerator. In total we obtained about 60 videos per sample over time, each with corresponding ground-truth parameters. Fig. 7 shows the inferred parameters for each sample and each point in time. Fig. 7(a) shows that our method clearly identifies the three

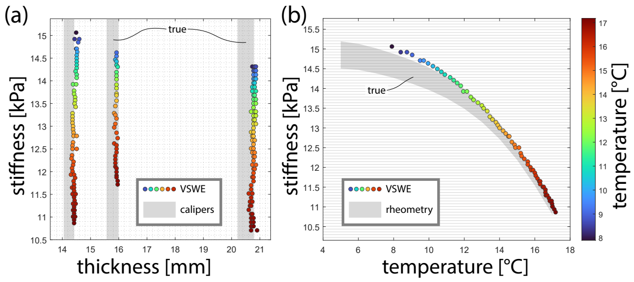

Figure 7. **(a) Estimated thickness and stiffness for three phantoms over a range of temperatures.** The phantoms were produced with three different volumes (1000, 1100, 1500 mL) of gelatin but set in the same-sized container, thus generating three different thicknesses. The thickness of each sample was measured with calipers. We produced a confidence interval of the ground-truth thickness by taking the 0.2 and 0.8 quantiles of multiple calipers measurements. The VSWE-estimated thickness consistently falls within the confidence interval. Note that while each phantom was polymerized according to the same approximate recipe, slight variations in the recipe, along with slightly different polymerization durations, led to some variation in the stiffness across samples. The temperature changes occurred as the samples spent more time at room temperature after being removed from the fridge. **(b) Estimated stiffness values for the thinnest sample over the same range of temperatures.** The ground-truth elastic modulus was measured with rheometry over a similar temperature range. Because the modulus of hydrogels can depend on both temperature and frequency [53], the rheometry measurements show a different stiffness range at each temperature. Note that the rheometry measurements were only taken between 10–100 Hz due to instrument limitations, while the excitation signal ranged from 40–200 Hz. This may explain some of the discrepancy between VSWE and rheometry, especially for the colder (and thus stiffer) samples which exhibit a stronger expression of higher-frequency wave dynamics. Regardless, in each case, VSWE estimates the stiffness extremely well, within 1.2% error of the rheometry range.

different thicknesses of the samples. As a sample spends more time out of the refrigerator, its temperature increases and its stiffness decreases. Fig. 7(b) shows that the estimated stiffness decreases accordingly with temperature.

5.3. 3D human leg with spatially-varying thickness

As another step towards realism, we tested VSWE on a simulated female human leg. The anatomical geometry of the leg was obtained as an STL file from the dataset of Andreassen et al. [3], who created 3D models from the National Library of Medicine's Visible Human Project [2]. We simulated the leg's response to a chirp excitation applied on the leg in COMSOL, running a full 3D physics simulation without any assumptions besides linear elasticity. For the sake of computational feasibility, we ran the simulation on the lower half of the leg, which we refer to as the calf region. To obtain dispersion relations, we considered the simulated displacements in two directions: one tangent to the leg's surface and one normal to the leg's surface.

Displacements were sampled at $800\mathrm{Hz}$ for $1\mathrm{s}$. The observation windows varied per region of the leg, with sizes between $\sim 3\mathrm{cm}$ and $\sim 11.5\mathrm{cm}$. See Appendix E for more details.

We applied a sweeping window across the upper calf region, estimating the thickness and modulus at each window location. Fig. 8 shows how, as the window sweeps around the leg, the inferred thickness changes. We computed the true thickness by taking the distance from the point on the skin to the nearest point on the bone. This results in a thickness distribution within the window, since the thickness changes slightly in the lengthwise direction of the leg, as well. Fig. 8 shows the inferred thickness and the distribution of ground-truth thicknesses for each location of the sweeping window. We find that the estimated thickness roughly agrees with the true thickness distribution. There appears to be slightly more error near the edges of the simulated domain, which may be due to boundary effects. While the physics-based simulation in COMSOL is well-modeled in the interior of the domain, there may be artifacts near the boundary of the domain, making it harder to infer the correct tissue parameters near the boundary.

We also applied VSWE to three different windows in the lower calf region. Rather than perform a sweep in this region, we chose to focus on three distinct windows since there are three subregions in the lower calf area that have notably different thicknesses, due to the tissue structure changing a lot near the ankle. As Fig. 8 shows, the inferred thickness agrees with the true thickness distribution for each subregion. Our method also recovered the constant stiffness throughout the calf, as shown in Fig. 1.

5.3.1. Objective function ablation

As mentioned in Sec. 4.2, we find that SSIM performs best for the optimization objective function. This choice was crucial for obtaining acceptable results with the 3D leg. Fig. 9 shows the optimization landscape for different choices of the objective function, demonstrated on both the upper calf simulation and a real video of gelatin. We compare SSIM to the following objective functions: a curve-based objective function, the negative mean squared error (MSE) between the images, and the peak signal-to-noise ratio (PSNR) between the images. The curve-based loss function goes through all the curves (i.e., $\{\omega_i(\gamma)\}_{i=1}^{N} \forall \gamma \in [0, \pi/a]$) in $\mathfrak{D}(T, E)$ and integrates the values of the observed dispersion relation $\mathbf{D}_{\text{obs}}$ along the curves. Fig. 9 shows that SSIM leads to the sharpest loss landscape and the most accurate estimated parameters. More details can be found in Appendix B.

6. Discussion

6.1. Practical considerations

VSWE's performance depends on many factors, including the types of waves supported by the tissue (e.g., surface

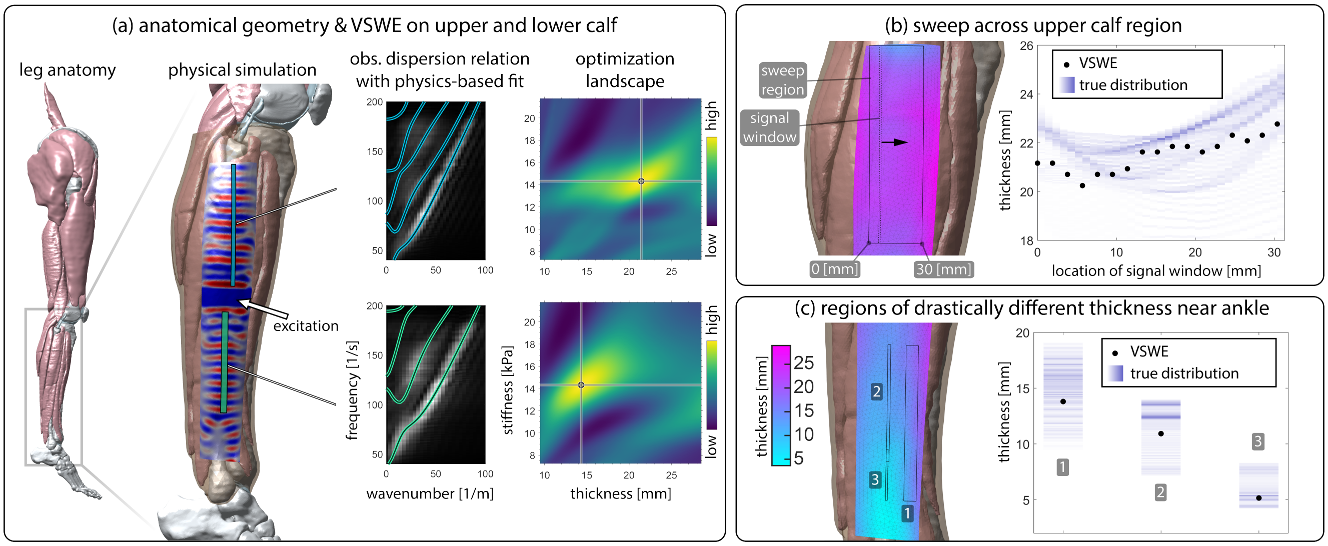

Figure 8. 3D anatomical inference. Panel **(a)** shows the realistic 3D leg anatomy (taken from the Visible Human Project [2]). We ran a full 3D simulation of the leg's response to a chirp excitation applied on the skin. The observed dispersion relation, along with the optimization landscape and hypothesized dispersion relation of the optimal parameters, is shown for two regions: one on the upper calf and one on the lower calf near the ankle. Panel **(b)** shows results of sweeping an observation window across a region of the upper calf. A ground-truth thickness distribution was obtained by computing the distance from the skin to the nearest point on the bone across the observation window. The VSWE-estimated thickness as the observation window slides from left to right tracks with the changing thickness of the leg. Panel **(c)** shows results for three distinct regions near the ankle. The VSWE estimations reflect the drastically different thicknesses of these regions.

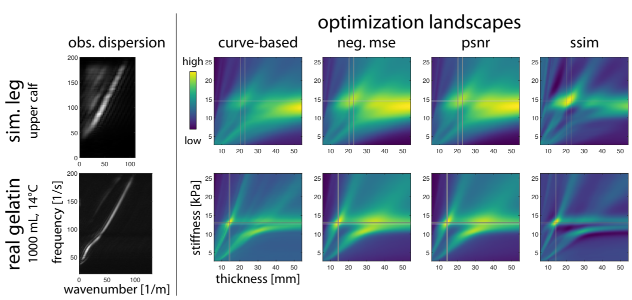

Figure 9. Ablation of optimization objective function. The optimization landscape is shown for various objective functions, applied in two different scenarios. Each landscape is visualized with its own color map limits, but bright yellow and dark blue indicate high (good) and low (bad) objective values, respectively. SSIM gives the sharpest optimization landscape. For the simulated upper calf, all the objective functions besides SSIM lead to the wrong optimal parameters (cross hairs indicate target parameters). Additionally, note how the simulated upper calf leads to a blurrier dispersion relation and hence blurrier optimization landscapes than the real gelatin sample, which has a simple geometry.

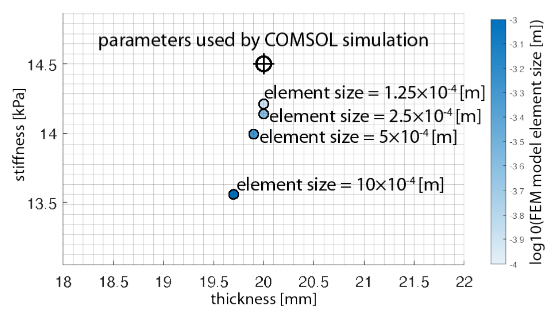

Figure 10. Ablation of FEM mesh resolution. As the mesh becomes finer, the VSWE-estimated parameters approach the true parameters used for the COMSOL plane-strain simulation.

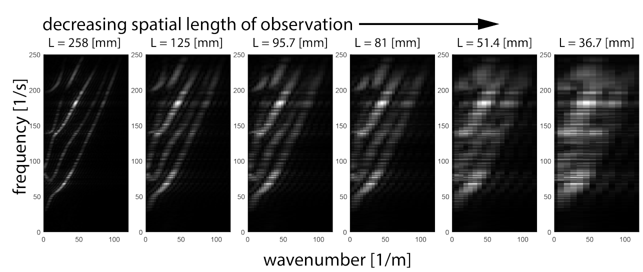

Figure 11. Ablation of spatial extent of observation. Keeping all else fixed, decreasing the spatial length L of the observation window degrades the quality of the observed dispersion relation. A degraded dispersion relation in turn leads to a fuzzy optimization landscape (see Fig. 9 for an example).

waves must be observable), the quality and extent of the observation (e.g., sufficient texture for motion extraction), and the quality of the physics model.

One practical consideration is the resolution of the FEM mesh used to model physics, as the element size determines the modeling accuracy and the smallest modeled wavelength. Fig. 10 shows that as the element size decreases, the estimated parameters approach the true parameters.

Another factor is that the spatial length L of the obser-

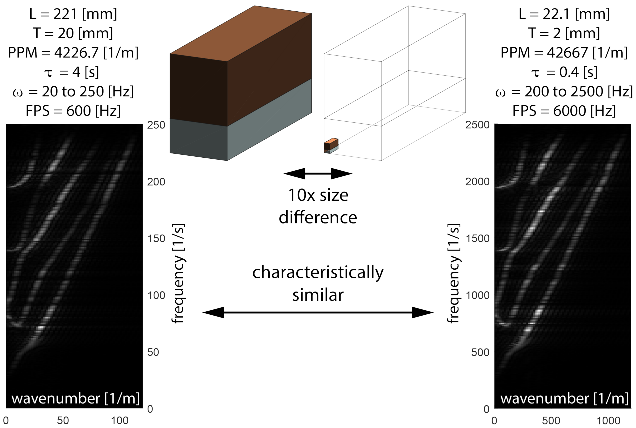

Figure 12. Leveraging similitude of characteristic numbers enables the application of VSWE to extremely different parameter ranges. In this simulated example, we show that by preserving the values of π_{1-6}, we see a similar observed signal from a system with parameters $10\times$ different than the primary ranges studied in this paper. In particular, we rescaled the size of the observation window, the thickness of the tissue, the spatial sampling rate, the total observation time, the frequency range of the excitation, and the temporal sampling rate. Note that the observation window on the right is smaller than that of the most degraded FFT in Fig. 11, yet the FFT here has not degraded.

vation window determines the largest wavelength and the number of wavelengths that can be observed. Fig. 11 shows how, keeping all else fixed, the FFT degrades as the observed domain shrinks, which makes inference more challenging.

6.2. Characteristic numbers

Characteristic numbers (a.k.a. dimensionless numbers or $\pi-$groups) allow us to roughly quantify limiting effects like those in the ablation studies of Fig. 10 and Fig. 11. They can be used to understand working limits for VSWE, and they give guidance for applying VSWE to scenarios beyond those studied in this work.

Some characteristic numbers reflect the observability of wave modes, such as $\pi_1 := \gamma L$, $\pi_2 := \text{PPM}/\gamma$, $\pi_3 := \omega\tau$, and $\pi_4 := \text{FPS}/\omega$, where PPM (pixels per meter) is the spatial sampling rate, FPS (frames per second) is the temporal sampling rate, and τ is the total observation time. Increasing π_{1-4} improves VSWE's performance by improving the quality of the FFT dispersion. The effect of varying π_1 can be observed in Fig. 11.

Other characteristic numbers reflect the numerical accuracy of the FEM model used to fit the observation, such as $\pi_5 := \frac{1}{\gamma e}$, where e is the FEM element size. Increasing π_5 improves performance by ensuring that physics is well-modeled in the parameter-fitting process (Fig. 10).

Yet other characteristic numbers reflect the actual physics that are possible in the system. For example, $\pi_6 := \gamma T$ conveys the "shallowness" of the wave. If π_6 is too large, then waves behave similarly to how they would in an infinitely-thick medium, making it difficult to pinpoint the thickness. If π_6 is too small, waves above a certain wavelength are not permitted to exist, making it necessary to capture smaller wavelengths.[2]

Across our experiments, we found that π_1, π_4, and π_6 were often the most relevant to practical working limits and helped us to identify and tune the minimum L, minimum FPS, and valid T ranges, respectively.[3] See Appendix C for examples of π values in our experiments.

Fig. 12 provides an example of how to take advantage of characteristic numbers. Specifically, we overcome the challenge of a smaller observed domain (as illustrated in Fig. 11) by rescaling other parameters to preserve the relevant characteristic numbers and thus maintain the quality of the observed dispersion relation. In general, one should look to characteristic numbers as guidelines for adjusting system and/or inference parameters in order to maintain the performance of VSWE across a variety of scenarios.

7. Conclusion

We have introduced Visual Surface Wave Elastography for estimating subsurface geometric and mechanical properties from visible surface waves. Our method works by extracting spatial and temporal frequency information from surface waves observed in video. By representing this information in the form of a dispersion relation, we can use a fully physics-based approach to solve for the properties that explain the wave behavior observed in the video. We validated our method on both simulated and real data, showing that our method is sensitive to small changes in the properties of tissue-like media. We demonstrated the promise of applying VSWE in the real world by testing it on real videos of gelatin-based phantoms and on a simulated realistic 3D human leg. Our work shows the potential of everyday visual data to reveal critical information deep beneath the skin.

Acknowledgments

ACO and CD are supported by DOE award no. DE-SC0021253 and the New Frontiers of Sound (NewFoS) NSF center. BTF and KLB acknowledge support from the NSF GRFP, the Amazon AI4Science Partnership Discovery Grant, and the Heritage Medical Research Institute (HMRI) Fellowship. JA and CD are supported by the HMRI and the NSF Center to Stream Healthcare in Place (C2SHIP) award no. 2052827.

[2]A small π_6 explains why ocean waves break near the shore [49]. It is also the reason that a higher-frequency excitation was helpful for the very thin tissue near the ankle in Fig. 8.

[3]Other characteristic numbers may prove to be the limiting factor in other scenarios or with different hardware.

References

[1] COMSOL AB. Comsol multiphysics® v5.5, 2019. 4

[2] M.J. Ackerman. The visible human project. *Proceedings of the IEEE*, 86(3):504–511, 1998. 6, 7

[3] Thor E Andreassen, Donald R Hume, Landon D Hamilton, Karen E Walker, Sean E Higinbotham, and Kevin B Shelburne. Three dimensional lower extremity musculoskeletal geometry of the visible human female and male. *Scientific Data*, 10(1):34, 2023. 6

[4] J. W. Andrews, M. J. Adams, and T. D. Montenegro-Johnson. A universal scaling law of mammalian touch. *Science Advances*, 6(41):eabb6912, 2020. 1, 3

[5] S R Arridge. Optical tomography in medical imaging. *Inverse Problems*, 15(2):R41, 1999. 2

[6] Danilo Avola, Luigi Cinque, Alessio Fagioli, Gianluca Foresti, and Alessio Mecca. Ultrasound medical imaging techniques: a survey. *ACM Computing Surveys (CSUR)*, 54 (3):1–38, 2021. 2

[7] R. A. Bell. Economics of MRI technology. *Journal of magnetic resonance imaging: JMRI*, 6(1):10–25, 1996. 1, 2

[8] Kiran S. Bhat, Christopher D. Twigg, Jessica K. Hodgins, Pradeep K. Khosla, Zoran Popović, and Steven M. Seitz. Estimating cloth simulation parameters from video. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, page 37–51, Goslar, DEU, 2003. Eurographics Association. 2

[9] Marcin Białas and Bojan B. Guzina. On the viscoelastic characterization of thin tissues via surface-wave sensing. *International Journal of Solids and Structures*, 48(14):2209–2217, 2011. 2

[10] Felix Bloch. Über die quantenmechanik der elektronen in kristallgittern. *Zeitschrift für physik*, 52(7):555–600, 1929. 3

[11] Katherine L Bouman, Bei Xiao, Peter Battaglia, and William T Freeman. Estimating the material properties of fabric from video. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1984–1991, Sydney, Australia, 2013. IEEE. 2

[12] Nathalie J Bureau and Daniela Ziegler. Economics of musculoskeletal ultrasound. *Current radiology reports*, 4:1–7, 2016. 1, 2

[13] Ziyang Chen, Shengyi Qian, and Andrew Owens. Sound localization from motion: Jointly learning sound direction and camera rotation. *International Conference on Computer Vision (ICCV)*, 2023. 2

[14] Abe Davis, Michael Rubinstein, Neal Wadhwa, Gautham Mysore, Fredo Durand, and William T. Freeman. The visual microphone: Passive recovery of sound from video. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 33(4):79:1–79:10, 2014. 2

[15] Abe Davis, Katherine L Bouman, Justin G Chen, Michael Rubinstein, Fredo Durand, and William T Freeman. Visual vibrometry: Estimating material properties from small motion in video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5335–5343, Boston, MA, USA, 2015. IEEE. 2

[16] Abe Davis, Katherine L. Bouman, Justin G. Chen, Michael Rubinstein, Oral Büyüköztürk, Frédo Durand, and William T. Freeman. Visual vibrometry: Estimating material properties from small motions in video. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(4):732–745, 2017. 2

[17] V Egorov, S Tsyuryupa, S Kanilo, M Kogit, and AJME Sarvazyan. Soft tissue elastometer. *Medical engineering & physics*, 30(2):206–212, 2008. 5

[18] Berthy T Feng, Alexander C Ogren, Chiara Daraio, and Katherine L Bouman. Visual vibration tomography: Estimating interior material properties from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16231–16240, 2022. 2

[19] Chao Feng, Ziyang Chen, and Andrew Owens. Self-supervised video forensics by audio-visual anomaly detection. *Computer Vision and Pattern Recognition (CVPR)*, 2023. 2

[20] David J Fleet and Allan D Jepson. Computation of component image velocity from local phase information. *International journal of computer vision*, 5(1):77–104, 1990. 4

[21] Gaston Floquet. Sur les équations différentielles linéaires à coefficients périodiques. In *Annales scientifiques de l'École normale supérieure*, pages 47–88, 1883. 3

[22] Freija Geldof, Behdad Dashtbozorg, Benno HW Hendriks, Henricus JCM Sterenborg, and Theo JM Ruers. Layer thickness prediction and tissue classification in two-layered tissue structures using diffuse reflectance spectroscopy. *Scientific reports*, 12(1):1698, 2022. 1, 3

[23] Penelope C. Georges, Jia-Ji Hui, Zoltan Gombos, Margaret E. McCormick, Andrew Y. Wang, Masayuki Uemura, Rosemarie Mick, Paul A. Janmey, Emma E. Furth, and Rebecca G. Wells. Increased stiffness of the rat liver precedes matrix deposition: Implications for fibrosis. *American Journal of Physiology-Gastrointestinal and Liver Physiology*, 293(6):G1147–G1154, 2007. 1

[24] N. Gucunski and R. D. Woods. Numerical simulation of the SASW test. *Soil Dynamics and Earthquake Engineering*, 11 (4):213–227, 1992. 2

[25] Ray Hashman Hashemi, William G Bradley, and Christopher J Lisanti. *MRI: the basics: The Basics*. Lippincott Williams & Wilkins, 2012. 2

[26] Kenneth Hoyt, Benjamin Castaneda, Man Zhang, Priya Nigwekar, P Anthony di Sant'Agnese, Jean V Joseph, John Strang, Deborah J Rubens, and Kevin J Parker. Tissue elasticity properties as biomarkers for prostate cancer. *Cancer Biomarkers*, 4(4-5):213–225, 2008. 1

[27] Hao Huang, Ray S Wu, Muyang Lin, and Sheng Xu. Emerging wearable ultrasound technology. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 71(7): 713–729, 2023. 2

[28] Yousef Javanmardi, Huw Colin-York, Nicolas Szita, Marco Fritzsche, and Emad Moeendarbary. Quantifying cell-generated forces: Poisson's ratio matters. *Communications physics*, 4(1):237, 2021. 3

[29] Jitesh Jodhani, Abhay Handa, Anandita Gautam, Ramakant Rana, et al. Ultrasonic non-destructive evaluation of compos-

ites: A review. *Materials Today: Proceedings*, 78:627–632, 2023. 2

[30] Sangwoo Kwon, Woochul Yang, Donggerami Moon, and Kyung Sook Kim. Comparison of Cancer Cell Elasticity by Cell Type. *Journal of Cancer*, 11(18):5403–5412, 2020. 1

[31] Lilian Lacourpaille, François Hug, Arnaud Guével, Yann Péréon, Armelle Magot, Jean-Yves Hogrel, and Antoine Nordez. Non-invasive assessment of muscle stiffness in patients with duchenne muscular dystrophy. *Muscle & Nerve*, 51(2):284–286, 2015. 1

[32] Chih-Ping Lin, Chun-Hung Lin, and Chih-Jung Chien. Dispersion analysis of surface wave testing – SASW vs. MASW. *Journal of Applied Geophysics*, 143:223–230, 2017. 2

[33] Chao Liu, Akash K. Maity, Artur W. Dubrawski, Ashutosh Sabharwal, and Srinivasa G. Narasimhan. High resolution diffuse optical tomography using short range indirect subsurface imaging. In *2020 IEEE International Conference on Computational Photography (ICCP)*, pages 1–12, 2020. 2

[34] Yogesh K. Mariappan, Kevin J. Glaser, and Richard L. Ehman. Magnetic resonance elastography: A review. *Clinical Anatomy*, 23(5):497–511, 2010. 1, 2

[35] Eder Miguel, Derek Bradley, Bernhard Thomaszewski, Bernd Bickel, Wojciech Matusik, Miguel A Otaduy, and Steve Marschner. Data-driven estimation of cloth simulation models. In *Computer Graphics Forum*, pages 519–528, Hoboken, NJ, USA, 2012. Wiley Online Library, Wiley Online Library. 2

[36] Soheil Nazarian and Kenneth Stokoe. Evaluation of Moduli and Thicknesses of Pavement Systems by Spectral-Analysis-of-Surface-Waves Method. *The Study of New Technologies for Pavement Evaluation Research Project 3-8-80-256, as conducted for Texas State Department of Highways and Public Transportation, by the Center for Transportation Research Bureau of Engineering Research - The University of Texas at Austin*, 1983. 2

[37] Choon Park. MASW for geotechnical site investigation. *The Leading Edge*, 32(6):656–662, 2013. 2

[38] Lord Rayleigh. On Waves Propagated along the Plane Surface of an Elastic Solid. s1-17(1):4–11. 12

[39] Babette S. Zemel. Chapter 18 - body composition during growth and development. In *Human Growth and Development (Second Edition)*, pages 461–486. Academic Press, Boston, second edition edition, 2012. 3

[40] Nikolai M. Shapiro, Michel Campillo, Laurent Stehly, and Michael H. Ritzwoller. High-Resolution Surface-Wave Tomography from Ambient Seismic Noise. *Science*, 307 (5715):1615–1618, 2005. 2

[41] Eero P Simoncelli and William T Freeman. The steerable pyramid: A flexible architecture for multi-scale derivative computation. In *Proceedings., International Conference on Image Processing*, pages 444–447, Washington, D.C., USA, 1995. IEEE, IEEE. 4

[42] Shan Sun, Zhenyuan Song, Scott J. Cotler, and Michael Cho. Biomechanics and functionality of hepatocytes in liver cirrhosis. *Journal of Biomechanics*, 47(9):2205–2210, 2014. 1

[43] Mihra S. Taljanovic, Lana H. Gimber, Giles W. Becker, L. Daniel Latt, Andrea S. Klauser, David M. Melville, Liang Gao, and Russell S. Witte. Shear-Wave Elastography: Basic Physics and Musculoskeletal Applications. *RadioGraphics*, 37(3):855–870, 2017. 1, 2

[44] N. H. Varady and A. J. Grodzinsky. Osteoarthritis year in review 2015: Mechanics. *Osteoarthritis and Cartilage*, 24 (1):27–35, 2016. 1

[45] Kelley M. Virgilio, Kyle S. Martin, Shayn M. Peirce, and Silvia S. Blemker. Multiscale models of skeletal muscle reveal the complex effects of muscular dystrophy on tissue mechanics and damage susceptibility. *Interface Focus*, 5(2): 20140080, 2015. 1

[46] Neal Wadhwa, Michael Rubinstein, Frédo Durand, and William T Freeman. Phase-based video motion processing. *ACM Transactions on Graphics (TOG)*, 32(4):1–10, 2013. 3

[47] Bohan Wang, Yili Zhao, and Jernej Barbič. Botanical materials based on biomechanics. *ACM Transactions on Graphics (TOG)*, 36(4):1–13, 2017. 2

[48] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004. 4, 11

[49] Paul Webb. *Introduction to oceanography*. Roger Williams University, 2021. 1, 8, 12

[50] Vincent Wai-Sun Wong and Henry Lik-Yuen Chan. Transient elastography. *Journal of Gastroenterology and Hepatology*, 25(11):1726–1731, 2010. 1, 2

[51] Tianyuan Zhang, Mark Sheinin, Dorian Chan, Mark Rau, Matthew O'Toole, and Srinivasa G. Narasimhan. Analyzing physical impacts using transient surface wave imaging. In *Proc. IEEE CVPR*, 2023. 2

[52] Zhendong Zhang, Tariq Alkhalifah, and Yike Liu. Full dispersion-spectrum inversion of surface waves. *JGR Solid Earth*, 2024. 2

[53] S Zörner, M Kaltenbacher, R Lerch, A Sutor, and M Döllinger. Measurement of the elasticity modulus of soft tissues. *Journal of Biomechanics*, 43(8):1540–1545, 2010. 6