Supplementary Material for Visual Surface Wave Elastography

A. Runtime analysis

The primary computational cost of VSWE stems from calculating dispersion relations to fit observed data. To address this efficiently, we precomputed dispersion relation datasets over grids of stiffness and thickness values suitable for soft tissue characterization. These datasets can be reused for multiple fitting scenarios, eliminating the need for repeated dispersion calculations. Although employing more efficient optimization methods than grid search would substantially reduce computational cost, we provide representative runtimes below to illustrate the expense associated with computing dispersion relations in this study.

In each dispersion dataset, there are several parameters that have bearing on the cost to compute the dataset. For the computation of a dispersion dataset containing 12 eigenvalue branches and 60 wavenumber values, on a grid of 41 stiffness (E) values $\times$ 41 thickness (T) values, the runtimes for various numbers of elements corresponding to different element sizes e are tabulated below. Computations were performed on an Intel Xeon CPU E5-2663 v3 with 20 cores, and all computations used less than 64 GB of memory. While we found $e = 0.5, N_{\text{ele}} = 240$ to be sufficient for reasonable performance in our settings, using smaller e with higher N_{ele} may be desirable if modeling higher frequency wave modes.

- $e = 0.5$ mm, $N_{\text{ele}} = 240$ had runtime of 157 s
- $e = 0.25$ mm, $N_{\text{ele}} = 960$ had runtime of 470 s
- $e = 0.125$ mm, $N_{\text{ele}} = 3840$ had runtime of 3169 s

B. Optimization objective functions

Here we define the objective functions we considered when developing our approach. Recall that $\mathbf{D}_{\text{obs}}$ is the observed FFT-derived dispersion relation. For a given hypothesized (T, E) pair, $\mathfrak{D}(T, E)$ is the physics-based dispersion relation, and $\mathbf{D}_{\text{hyp}}(T, E)$ is the image version of $\mathfrak{D}(T, E)$.

Curve-based objective function We define a simple and intuitive curve-based objective function f that assigns points for the observed dispersion relation having high magnitude at the points that exist in the hypothesized physics-based dispersion relation $\mathfrak{D}(T, E) = \{\omega(\gamma)\}_{i=1}^{N} \forall \gamma \in [0, \pi/a]$:

$$f\left(\mathfrak{D}(T, E), \mathbf{D}_{\text{obs}}\right) = \sum_{i=1}^{N} \int_{0}^{\pi/a} \mathbf{D}_{\text{obs}}\left(\gamma, \omega_i(\gamma)\right) \mathrm{d}\gamma. \quad (6)$$

Image-based objective functions To convert the dispersion relations from curve format to image format, we assign pixel values with a Gaussian kernel based on the distance from the pixel to any point on any of the curves. More precisely, the value of a pixel located at point $(\gamma_{\text{im}}, \omega_{\text{im}})$ is given by

$$v(\omega_{\text{im}}, \gamma_{\text{im}}) = \exp\left(\frac{-d_{\text{min}}^2}{2\sigma^2}\right), \quad (7)$$

where

$$d_{\text{min}}(\omega_{\text{im}}, \gamma_{\text{im}}) = \min_{\gamma_c, i}\left(\sqrt{(\omega_{\text{im}} - \omega_c^i(\gamma_c))^2 + (\gamma_{\text{im}} - \gamma_c)^2}\right), \quad (8)$$

and $\omega_c^i(\gamma_c)$ denotes the dispersion relation in curve form on any band i, and σ is a parameter that controls the width of the curve once converted to image form.

We use standard definitions of MSE, PSNR, and SSIM [48].

C. Characteristic numbers

Here we provide examples of characteristic number values from our experiments (both real and synthetic). While we have listed several of what we think are the most practically relevant characteristic numbers in our analysis, others can surely be defined. We emphasize that the ranges of characteristic numbers here do not comprehensively represent *all* systems where VSWE will function smoothly, but rather they represent *a few examples* of systems where we have observed VSWE to function smoothly. It would be a worthwhile exercise to determine more carefully curated bounds for these characteristic numbers in order to more precisely define where the method can be used.

As a reminder for the reader, the definition for each characteristic number is given in Tab. 1 along with a brief intuitive description of what it conveys.

Characteristic numbers with known working limits Some characteristic numbers have well-known limits from previously existing scientific knowledge. For example, Nyquist limits (and practical version of them) for $\pi_2 :=$ PPM$/\gamma$ and $\pi_4 :=$ FPS$/\omega$ give well-defined guidelines for spatial and temporal sample rates if one knows approximately what frequency and wavenumber ranges will be observed. These can help to determine what hardware would be required to apply VSWE to a particular system. As demonstrated by our synthetic experiments, for an upper

Characteristic number definitions

π	Mathematical definition	Intuitive description
π_1	γL	number of waves that fit spatially in the observation
π_2	PPM/γ	number of spatial sample points per wave
π_3	$\omega\tau$	number of waves that fit temporally in the observation
π_4	FPS/ω	number of temporal sample points per wave
π_5	$1/(\gamma e)$	number of FEM mesh elements per wave
π_6	γT	ratio of thickness divided by wavelength

Table 1. Characteristic number definitions and intuitive descriptions.

Characteristic number values

π	Real	Sim. leg (lower, thin)	Sim. leg (lower)	Sim. leg (upper)
π_1	$[2.1, 27]$	$[0.7, 4.6]$	$[0.9, 8]$	$[1.7, 9.6]$
	number of waves that fit spatially in the observation			
π_2	$[43, 548]$	$[51, 799]$	$[46, 400]$	$[47, 266]$
	number of spatial sample points per wave			
π_3	$[160, 800]$	$[58, 175]$	$[120, 360]$	$[100, 360]$
	number of waves that fit temporally in the observation			
π_4	$[3, 15]$	$[2.3, 7]$	$[4.4, 13.3]$	$[4.4, 16]$
	number of temporal sample points per wave			
π_5	$[22, 286]$	$[26, 400]$	$[23, 200]$	$[23, 133]$
	number of FEM mesh elements per wave			
π_6	$[0.1, 1.9]$	$[0.05, 0.78]$	$[0.12, 1.5]$	$[0.3, 2]$
	ratio of thickness divided by wavelength			

Table 2. Characteristic number ranges for the highest and lowest frequencies and wavenumbers observed in each experiment. For the simulated experiments in this table, thickness ranges were also considered for the computation of π_6.

calf about 20 mm thick, one would need to observe frequencies between $[50, 180]$ Hz and wavenumbers between $[15, 85]$ m^{-1}. Applying Nyquist limits to π_2 and π_4 then tells us that we need at minimum 360 FPS and 170 px/m — well within the capacities of consumer cameras.

Leveraging characteristic similarity Other characteristic numbers do not yet have clearly defined bounds. For example, while characteristic shallowness ($\pi_6 := \gamma T$) has some rule-of-thumb bounds for mechanical analysis [38, 49], they carry heavy assumptions and do not have direct relevance for inference problems. However, even without firm bounds, characteristic numbers can prove useful as invariants. If one finds a set of characteristic numbers that are confirmed to be VSWE-compatible in one scenario (e.g., by experiment), they can expect to preserve the performance of VSWE by preserving the values of the characteristic numbers when transferring to a new scenario. This is known as leveraging *similitude*: when the relevant characteristic numbers of two different systems have the same values, the systems are called *similar*.[4]

For example, when shifting from inference over a thicker region of tissue to a thinner region of tissue, we can adjust other parameters (i.e., ones that we can control) to preserve the characteristic shallowness of the waves ($\pi_6 := \gamma T$). In the VSWE pipeline, this helps to maintain the quality and scope of observation of the dispersion relation. Because T is decreasing, we would need to increase γ somehow. To do this, we would increase the frequencies ω (a common heuristic from elastic wave mechanics says that wave speed $c \sim \sqrt{E/\rho} \sim \omega/\gamma$, and here, the material properties E and ρ are fixed). These increases in γ and ω demand commensurate increases in the minimum acceptable FPS and PPM (i.e., to preserve $\pi_2 := \mathrm{PPM}/\gamma$ and $\pi_4 := \mathrm{FPS}/\omega$), but allow us to reduce the spatial and temporal extents of our observation (τ and L) while still preserving $\pi_1 := \gamma L$ and $\pi_3 := \omega\tau$.

Characteristic numbers are often coupled As demonstrated by the previous example, it is often the case that changing one system parameter affects several characteristic numbers in different ways, and controlling them individually is not always possible. As a more hardware-related example, increasing the FPS of the video to increase $\pi_4 := \mathrm{FPS}/\omega$ may require us to lower the resolution, having the side-effect of lowering $\pi_2 := \mathrm{PPM}/\gamma$ (assuming focus distance and field of view are fixed).

D. Mechanics

D.1. Linear elasticity & the elastic wave equation

Linear elasticity is the theory that dictates how a solid moves when subjected to external forces — in our case, vibrations that generate measurable surface waves.

Momentum balance Within a solid body, material moves and displaces as a result of stresses exerted by neighboring pieces of material and external forces applied to the solid body. We can express this with Cauchy's momentum balance

$$\rho\,\ddot{\mathbf{u}} = \nabla\cdot\boldsymbol{\sigma} + \mathbf{f}, \tag{9}$$

where ρ is density, and vectors $\mathbf{u}$, $\boldsymbol{\sigma}$, and $\mathbf{f}$ are displacement, stress, and body forces (respectively) in each of the three spatial directions. Note that each of ρ, $\mathbf{u}$, $\boldsymbol{\sigma}$, and $\mathbf{f}$ are taken to be functions of space ($\mathbf{x}$). Cauchy's momentum balance is essentially the continuum mechanics version of Newton's second law.

[4] Beyond elastodynamics, similitude is also often used in fluid mechanics. Instead of constructing an expensive full-sized airplane or wing for an aerodynamic test, engineers often opt to run a characteristically similar experiment on a scaled-down version of the object, preserving the characteristic number known as the *Reynold's number* by modifying fluid density, velocity, and viscosity to balance out the size scale difference.

Linear-elastic constitutive law In general, if a body of material displaces non-uniformly (such that $\mathbf{u}$ has a non-zero spatial gradient), stress ($\boldsymbol{\sigma}$) often develops within the material (and in many cases acts as a restorative force to return the body to some equilibrium configuration). The exact relationship between the spatial gradient of displacement (often written compactly as strain, ϵ) and stress σ is called *a constitutive law*, which is different for different classes of materials. In our work, we restrict our studies to a *linear-elastic* constitutive law.

Linear elasticity proposes a constitutive law where the stress, $\boldsymbol{\sigma}$, (like a pressure, or area-normalized force) that a piece of material feels is *linearly* proportional to the strain, $\boldsymbol{\varepsilon}$, (the symmetrized spatial gradient of displacement) that the piece of material is experiencing. It can be written more concisely if instead of elastic modulus E and Poisson's ratio ν, we use the Lamé parameters $\mu := \frac{E}{2(1+\nu)}$ and $\lambda := \frac{E\nu}{(1+\nu)(1-2\nu)}$:

$$\boldsymbol{\sigma} = \lambda\,(\nabla\cdot\mathbf{u})\,\mathbf{I} + 2\mu\,\boldsymbol{\varepsilon}, \tag{10}$$

where $\boldsymbol{\varepsilon} := \frac{1}{2}\big(\nabla\mathbf{u} + \nabla\mathbf{u}^\top\big)$.

Elastic wave equation Substituting the linear-elastic constitutive law into Cauchy's momentum balance yields the elastic wave equation (a linear PDE problem where one traditionally solves for $\mathbf{u}$):

$$\rho\,\ddot{\mathbf{u}} = (\lambda+\mu)\,\nabla(\nabla\cdot\mathbf{u}) + \mu\,\nabla^2\mathbf{u} + \mathbf{f}, \tag{11}$$

which describes how elastic waves travel through linear elastic materials.

D.2. Transient analysis

Transient analysis means that time is being explicitly modeled. In transient simulations, we solve Eq. (11) for $\mathbf{u}(\mathbf{x},t)$ by some choice of time-stepping scheme. Although these simulations are often more costly than their assumption-leveraging counterparts (e.g., harmonic analysis, described in Appendix D.3.1), they are able to model very general situations with few restrictions and with high fidelity.

In our work, *all* of the simulated experiments used transient analysis to represent the observed systems because we wanted the simulated material behavior to involve as few simplifying assumptions as possible and match as closely as possible with the behavior of real-world materials.

D.3. Bloch-Floquet analysis

Bloch-Floquet (BF) analysis aims to reveal the natural wave modes of a system comprehensively, finding waves of every possible (γ,ω) frequency-wavevector pair. BF analysis applies a *time-harmonic assumption* and *BF periodic boundary conditions* (a spatially quasi-periodic assumption) to Eq. (11). Each of these assumptions is discussed in this section. In our work, *all* of the dispersion relations computed to fit the observed dispersion relations used BF analysis.

D.3.1. Time-harmonic assumption

BF analysis is a subcategory of *harmonic analysis*, and as such it employs the same time-harmonic assumption from harmonic analysis. The time-harmonic assumption says the solution is harmonic (sinusoidal) in time, written as

$$\mathbf{u}(\mathbf{x},t) = \mathbf{u}_0(\mathbf{x})e^{\mathbf{i}\omega t}, \tag{12}$$

where ω represents the time frequency of the harmonic solution, and $\mathbf{u_0}(\mathbf{x})$ is explicitly *not* a function of time. In general, assuming the displacement field is a harmonic function at frequency ω only makes sense when the external forces applied to the system are also harmonic functions at the same frequency. That is,

$$\mathbf{f}(\mathbf{x},t) = \mathbf{f}_0(x)e^{\mathbf{i}\omega t}. \tag{13}$$

In cases where we are looking for natural modes (e.g., in our work we look for natural wave modes), we set external forces $\mathbf{f}_0 = 0$. However, there are other situations (e.g., in *frequency-domain* analysis, which is *not* used in this work) where non-zero forcing would be appropriate. Applying the harmonic assumption in Eq. (12) and assuming $\mathbf{f_0} = 0$, Eq. (11) simplifies to

$$\mu\nabla^2\mathbf{u}_0 \;+\; (\lambda+\mu)\,\nabla\big(\nabla\cdot\mathbf{u}_0\big) \;+\; \rho\,\omega^2\,\mathbf{u}_0 \;=\; \mathbf{0}, \tag{14}$$

which is a linear eigenvalue PDE, where $\mathbf{u_0}(\mathbf{x})$ is the eigenfunction (representing the displacement field of the wave mode), ω^2 is the eigenvalue, and ω is the frequency (a.k.a. *eigenfrequency*) of the wave mode.

D.3.2. Bloch-Floquet boundary conditions

Alongside the time-harmonic assumption, Bloch analysis also adopts a *spatially* quasi-periodic assumption: that the solution field repeats from unit cell to unit cell, but with a phase shift. This is imposed through *BF periodic boundary conditions*:

$$\mathbf{u}(\mathbf{x}+\mathbf{a}_n) = \mathbf{u}(\mathbf{x})\,e^{\mathbf{i}\,\boldsymbol{\gamma}\cdot\mathbf{a}_n}, \qquad n = 1,\ldots,D, \tag{15}$$

where D is the spatial dimensionality of the problem ($D = 1,2,3$ is common for real-world problems), $\mathbf{a}_n$ is the size of the modeled material domain in each direction (commonly known as a *lattice vector*), and γ is the wavevector (as usual) which defines the lengthscale of Bloch-Floquet periodicity in each direction. An intuitive way to think about these boundary conditions is "the solution at one boundary is the same as the solution at the opposite boundary, but with a phase shift prescribed by the wavevector."

D.4. Plane-strain assumption

When most of the significant mechanics of a three-dimensional system are happening in only two of the dimensions, it may be reasonable to only model those two dimensions. One standard assumption is called *plane strain*,

which assumes zero strain in the direction not explicitly modeled. This is commonly employed in situations where a material body is very thick in one direction or has fixed boundary conditions preventing strain in the out-of-plane direction. Letting y be the direction of zero strain, this assumption can be written as $u_y = 0$ and $\partial_y(\,\cdot\,) = 0$. Plane strain simplifies Eq. (11) considerably and significantly reduces the amount of computation required for simulations.

The transient models for the simulated experiments from the sensitivity demonstration in Fig. 5 and ablation studies in Fig. 10 and Fig. 11 use this assumption for computational efficiency since we ran such a large number of them. However, the transient models for the simulated experiments from Sec. 5.3 use full 3D physics (no plane-strain assumption) to maximize real-world fidelity.

All dispersion datasets used to perform fits use the plane-strain assumption (see Appendix D.3). This is exactly ideal in our simulated plane-strain studies (as the assumption is identical to the observed physics) and close-to-ideal in our real-world gelatin studies (because the samples were thick *and* had fixed out-of-plane boundary conditions discouraging out-of-plane strain). However, it is a less-than-perfect assumption for the simulated anatomical geometry studies. It is promising that the plane-strain assumption works for inference problems on the anatomical geometry, despite the fact that it is less than ideal, demonstrating a measure of robustness for VSWE.

E. Simulated experiment details

In this section we give supplementary details surrounding the simulated experiments.

E.1. Sensitivity experiments (Fig. 5)

The simulated experiments in Fig. 5 were simulated on a 30 cm domain and used an observation window of 29.5 cm (the remaining 0.5 cm region was where the input excitation was applied). The simulated duration was 1 s and was equivalent to the observed duration. Displacements were sampled at 1280 positions and 600 times, equating to sample rates of ~ 4336 $\mathrm{px/m}$ and 600 FPS.

E.2. Mesh resolution ablation experiments (Fig. 10)

The simulated experiments for the mesh size ablation study in Fig. 10 were simulated on a 60 cm domain (unrealistic, but we chose this to be large so that domain size was not a confounding factor) and used an observation window of 59.5 cm (the remaining 0.5 cm region was where the input excitation was applied). The simulated duration was 1 s and was equivalent to the observed duration. Displacements were sampled at 1280 positions and 600 times, equating to sample rates of ~ 2150 $\mathrm{px/m}$ and 600 FPS.

E.3. Observation length ablation experiments (Fig. 11)

The simulated experiments for the ablation study where the spatial length of the observed surface was ablated in Fig. 11 were simulated on a 30 cm domain and used observation windows of varying spatial extents. Even though the spatial extent of the observation changed, the spatial distance between sample points was maintained constant at ~ 4336 $\mathrm{px/m}$. The simulated duration was 1 s and was equivalent to the observed duration. Displacements were sampled at 600 times, equating to a sample rate of 600 FPS.

E.4. Anatomical experiments (Fig. 8)

The operational parameters of the simulated experiments in the anatomical section were not as uniform as those from the plane-strain section, mainly due to the more complicated nature of the geometry and the fact that simulation of the thin region near the ankle required its own separate simulation due to computational cost reasons. As such, details will be given for each scenario presented in the main body of the paper, as well as for each of the two transient simulations.

E.5. Anatomical simulations

Here we give the details of the two anatomical simulations.

Common between the two Many of the parameters are shared between the two simulations. Both models prescribed $E = 14.5$ kPa, $\rho = 1000$ $\mathrm{kg/m^3}$, and $\nu = 0.45$ everywhere in the soft tissue, and they treated the bone interface as a fixed boundary. Both models used low-reflecting boundary conditions on a segmented subsection of the calf to avoid modeling the entire domain. Both models used a quadratic tetrahedral finite element discretization, though with different element sizes. Both models used an excitation region that spanned laterally across the outer surface of the leg, with a length of 1.5 cm in the direction of wave propagation.

Simulation 1 simulated both the upper-calf and lower-calf regions to demonstrate simultaneous inferences of these significantly different regions from a single excitation. The prescribed chirp excitation was a smoothed chirp sweeping from 25 to 200 Hz. This simulation used a target finite element (spatial) discretization size of 3 mm. We used a physical timestep of 800 Hz and solved for a duration of 2 s. The length of the simulated domain was ~ 27 cm with approximate cross-section dimensions of $\sim 2-5$ cm $\times$ $\sim 2-5$ cm with significant variance along the length. The model contained $\sim 180,000$ elements and took 42 h 45 min to solve on an Intel Xeon CPU E5-2663 with 20 cores.

Simulation 2 specifically simulated the lower-calf region near the ankle to demonstrate inference in a region of the body where the soft tissue is much thinner and varies significantly. The prescribed chirp excitation was a smoothed chirp sweeping from 40 to 400 Hz. This simulation used a target finite element (spatial) discretization size of 2 mm. We used a physical timestep of 800 Hz and solved for a duration of 0.5 s. This timestep was likely a bit too coarse for the upper end of the excitation frequency (400 Hz) in this simulation. The model used low-reflecting boundary conditions on a segmented subsection of the lower-calf to avoid modeling the entire domain. The length of the simulated domain was ~ 11 cm with approximate cross-section dimensions of $\sim 1.5-4$ cm $\times \sim 1.5-4$ cm with significant variance along the length. The model contained $\sim 130,000$ elements and took 5 h to solve on an Intel Xeon CPU E5-2663 with 20 cores.

The main difference between simulation 1 and simulation 2 is that simulation 2 modeled a thinner region of the body than simulation 1, so the waves in this region were smaller in both space and time (i.e., higher wavenumbers and frequencies). Thus finer spatial and temporal discretizations were needed, but over shorter spatial and temporal observation windows. Additionally, the excitation frequency range was approximately doubled.

E.6. Anatomical VSWE inferences

Fig. 8(a) In this panel, two inferences are presented — one inference in the upper calf and one in the lower calf. Both used simulation 1.

In Fig. 8(a), the upper calf inference used the following parameters:

- ~ 11.3 cm spatial observation window
- 452 spatial sample points
- ~ 4000 px/m spatial sample rate
- 2 s temporal observation window
- 1601 temporal sample points
- 800 FPS temporal sample rate

In Fig. 8(a), the lower calf inference used the following parameters:

- ~ 9.3 cm spatial observation window
- 373 spatial sample points
- ~ 4000 px/m spatial sample rate
- 2 s temporal observation window
- 1601 temporal sample points
- 800 FPS temporal sample rate

Fig. 8(b) This panel presents a series of inferences as an observation window sweeps laterally across the upper calf, making an inference at each location in the sweep. All inferences in this sweep used simulation 1. In Fig. 8(b), the inferences in the upper calf sweep used the following parameters:

- ~ 11.6 cm spatial observation window
- 463 spatial sample points
- ~ 4000 px/m spatial sample rate
- 2 s temporal observation window
- 1601 temporal sample points
- 800 FPS temporal sample rate

Fig. 8(c) In this panel, three inferences are presented in three different regions near the ankle of significantly different thickness. All inferences in this panel used simulation 2. In Fig. 8(c), the inference on region **1** used the following parameters:

- ~ 9.1 cm spatial observation window
- 730 spatial sample points
- ~ 8000 px/m spatial sample rate
- 0.5 s temporal observation window
- 401 temporal sample points
- 800 FPS temporal sample rate

In Fig. 8(c), the inference on region **2** used the following parameters:

- ~ 6.9 cm spatial observation window
- 549 spatial sample points
- ~ 8000 px/m spatial sample rate
- 0.5 s temporal observation window
- 401 temporal sample points
- 800 FPS temporal sample rate

In Fig. 8(c), the inference on region **3** used the following parameters:

- ~ 3 cm spatial observation window
- 240 spatial sample points
- ~ 8000 px/m spatial sample rate
- 0.5 s temporal observation window
- 401 temporal sample points
- 800 FPS temporal sample rate

F. Experiment details

In this section we give supplementary details surrounding the real-world experiments.

F.1. Real-world VSWE gelatin experiments

Camera We used a Phantom V1610 high-speed camera. We operated it at 600 FPS and a resolution of 96×1280 pixels (cropping was later applied before inputting the videos to the VSWE pipeline). For our setup, this spatial resolution equates to a spatial sampling rate of ~ 3840 px/m. Each video was taken in grayscale at a bit-depth of 16 and then downsampled to a bit-depth of 8 before entering the VSWE pipeline (for computational efficiency). Each video was originally ~ 8 seconds long but trimmed to ~ 4 seconds before entering the VSWE pipeline.

Geometry The gelatin was polymerized in a glass baking dish with an interior dimension of ~ 30 cm in length and 19 cm in width. Waves traveled in the 30 cm direction, originating near one end of the glass baking dish, and propagating length-wise toward the other end. The excitation head attached to the shaker was ~ 10 cm in width.

The camera was positioned ~ 1.2 m from the gelatin surface, and the viewing angle of the camera was at a $\sim 30^\circ$ angle with the surface.

Temperature measurements The gelatin was polymerized in a refrigerator $\sim 8\,^\circ$C for $\sim 24 \pm 3$ h. After being removed from the refrigerator, the gelatin was immediately textured with garlic powder and subsequently positioned in the camera's line of sight with the shaker applied to its surface. A Thermoworks Thermapen One thermometer was used to measure temperature. The thermomenter was positioned such that the probe was approximately halfway between the upper and lower surface of the gelatin, a few centimeters from the edge of the glass baking dish. As often as the hardware would allow (some amount of time is needed to write the video to disk), a video was taken. For each video, the current time and temperature were recorded.

F.2. Rheometry

Rheometry was used as the ground-truth measurement for the gelatin stiffness. Here, we give more details on the rheometry setup. A rheometer is an expensive piece of bench-top equipment that is meant to measure the stiffnesses and viscosities of soft solids and fluids. These measured values are often dependent on temperature, excitation frequency, humidity, and a variety of other factors, so high-end rheometers often have built-in systems to control these variables.

Rheometers use material samples that are shaped like disks. Between this fact and the price tag of a rheometer, they are unfit for the in-vivo measurement tasks that VSWE aims to deliver, but are very helpful for benchmarking against a well-established material measurement method in a controlled environment.

Polymerization All gelatin in this work was polymerized at a weight concentration of 4%. We polymerized our gelatin samples directly on the rheometer with gelatin taken directly from the gelatin samples prepared for the real-world VSWE gelatin experiments. The gelatin samples were taken when it was in liquid state, just before setting the remainder to polymerize in the refrigerator for the VSWE experiment. To replicate the polymerization environment as exactly as possible, we set the rheometer gelatin and the VSWE gelatin to polymerize at the same time (give or take a few minutes) and took the VSWE videos concurrently with the rheometer measurements. During polymerization, we set the rheometer stage to the same temperature as the refrigerator. This way, we could ensure that both samples had the exact same amount of time to polymerize, and at the same temperature.

Stiffness measurements The rheometer measures the stiffness of the gelatin by twisting the disk-shaped sample in a shearing motion, back and forth, at a prescribed frequency. The deformation applied is small, so as not to damage the sample. In our experiments we applied 7% shear strain. This measurement gives the *shear modulus* G of the gelatin, which can then be converted to an elastic modulus E by assuming a value for the Poisson's ratio ν. We assume $\nu = 0.45$, which is reasonable for a hydrogel material like gelatin.

Sweep over frequency and temperature This stiffness measurement is repeated over a grid of frequencies and temperatures, since frequency and temperature both affect how stiffly the gelatin behaves. For each gelatin sample, the stiffness was measured over the ranges of $[10, 100]$ Hz and $[5, 17]\,^\circ$C. These ranges were chosen to emulate the frequencies and temperatures observed in some preliminary VSWE experiments; however, we should note that the rheometer was limited to not exceed 100 Hz, whereas the VSWE videos observed surface waves up to at least ~ 200 Hz. The rheometry measurements indicated that the gelatin stiffens as excitation frequency increases, though we have no ground-truth stiffness measurements from rheometry in the range of $[100, 200]$ Hz. This mismatch between frequency ranges may explain why the VSWE inferences appear to be at the upper bound, or slightly above, the rheometry measurements.

F.3. Excitation system

To excite the surface of the gelatin, we uploaded chirp waveforms to an Agilent signal generator. The signal generator passed this signal to a B&K amplifier, which then passed the signal to a B&K shaker. Finally, the shaker mechanically applied the excitation to the surface of the gelatin.