

Analyzing Physical Impacts using Transient Surface Wave Imaging

Tianyuan Zhang¹ Mark Sheinin¹ Dorian Chan¹ Mark Rau²
 Matthew O'Toole¹ Srinivasa G. Narasimhan¹
¹Carnegie Mellon University ²Stanford University

Abstract

The subtle vibrations on an object's surface contain information about the object's physical properties and its interaction with the environment. Prior works imaged surface vibration to recover the object's material properties via modal analysis, which discards the transient vibrations propagating immediately after the object is disturbed. Conversely, prior works that captured transient vibrations focused on recovering localized signals (e.g., recording nearby sound sources), neglecting the spatiotemporal relationship between vibrations at different object points. In this paper, we extract information from the transient surface vibrations simultaneously measured at a sparse set of object points using the dual-shutter camera described by Sheinin et al. [37]. We model the geometry of an elastic wave generated at the moment an object's surface is disturbed (e.g., a knock or a footstep) and use the model to localize the disturbance source for various materials (e.g., wood, plastic, tile). We also show that transient object vibrations contain additional cues about the impact force and the impacting object's material properties. We demonstrate our approach in applications like localizing the strikes of a ping-pong ball on a table mid-play and recovering the footsteps' locations by imaging the floor vibrations they create.

1. Introduction

Our environment is teeming with vibrations created by the interaction of physical objects. Some vibrations, like a knock on the door or the sound of a ball bouncing off the ground, can be perceived by humans because they are transmitted from the vibrating object's surface via the air. However, many vibrations that fill our world are too subtle for auditory-based remote sensing. Moreover, much like ripples in a pond, the transient spatial shapes such vibrations create on object surfaces are a *visual* cue that can disclose the source of the disturbance and other object properties.

Object vibrations can be divided into two main types: transient and modal. For example, consider the vibrations of a tuning fork. When struck, the impulse creates transient waves propagating from the impact source until they




Figure 1. When physical objects interact, like a ping pong ball bouncing off the table, they create minute vibrations that propagate through the objects' surfaces and interiors. The transient vibrations that occur immediately on impact, exaggerated here for visualization, carry information about the impact source location. We image the surface vibrations at a sparse set of locations using the imaging system of Sheinin *et al.* [37]. We model the elastic wave propagation and recover the impact source locations without a direct line-of-sight on the impacted surface. Visit the project page for videos of results [1].

reach and vibrate the fork's entire body. After a short time interval, the transient vibrations die down, leaving the fork to vibrate at its resonant modal frequencies. Modal analysis, which aims to measure these resonant frequencies [11, 13, 42], can reveal the tuning fork's designed tone (e.g., 440 Hz for the A tone) and can also be used to analyze the fork's material properties [9, 14, 18].

While extremely useful, modal analysis ignores the *transient* vibrations that occur at the moment of impact. Such transient vibrations contain valuable cues about the disturbance's origin, its magnitude, and the properties of the object causing the disturbance (e.g., a falling basketball vs. a falling rock). Prior works that did sense transient vibrations primarily focused on localized low-dimensional signals such as heartbeats [42, 44, 48], music and speech [8, 15, 37, 45, 46, 48], and musical instruments [37]. These works disregard the spatiotemporal relationship between transient vibrations at different object points.

This paper focuses on recovering the physical location of an impacting object from transient surface vibra-




Figure 2. Elastic wave propagation in isotropic objects. (a) An electronic knocker creates repeated short knocks on a whiteboard. For each knock, a laser Doppler vibrometer (LDV) sensor is used to optically measure the temporal vertical displacement at a single point. Aggregating and synchronizing measurements from multiple board points generates a video showing the surface displacement with time. (b) Displacement 1ms after impact. Observe the circular shape of the outgoing wave. (c) Displacement 3.1ms after impact. Here, the outgoing wave has reflected from the board’s boundaries.

tions measured simultaneously at multiple surface points using the dual-shutter camera of Sheinin *et al.* [37]. This task opens the door to potential applications like localizing sound sources in walls (*e.g.*, pipe bursting), localizing bullet or bird impacts on airplanes mid-flight, or impacts on ship hulls from dockside, tugs, or other debris, localizing shell-ground impacts on battlefields, localizing people in building fires or hostage situations by observing external vibrations on ceilings or side walls, and more.

While, in general, object shape and material determine its vibration profile, we show that immediately after impact, there exists a short time interval (~ 1.5 ms long) where the surface vibrations can be modeled as an outwardly propagating elastic wave. We derive an approximate model of the wave’s geometry for both isotropic and anisotropic materials and develop a backprojection-based algorithm to localize the impact sources using the vibrations within this time interval. Unlike prior works that merely visualize acoustic wave propagation [36], we explicitly model its transient behavior and show that only a sparse set of points is required to determine the wave’s source.

We verified our approach on various materials, including wood, plastic, glass, porcelain, and gypsum. In our experiments, we localized impact sources with an average error between 1.1 cm and 2.9 cm for 40 cm \times 40 cm and 90 cm \times 90 cm surfaces, respectively. We also show applications like localizing ping-pong ball strikes on the table mid-play and localizing footsteps through floor vibrations beyond a camera’s line of sight.

Beyond impact localization, we show that the transient surface vibrations can convey more information about the impacting object and the impacted surface. For surfaces of unknown material, we estimate the material anisotropy by measuring vibrations at known surface points and fitting a material-specific wave propagation model parameter. Our preliminary experiments suggest that the transient vi-

brations’ amplitudes relate to the force applied to disturb the object [20, 28, 31], and that the vibrations’ frequency content depends on the stiffness and shape on the impacting object. We thus believe our work can inspire a new class of *transient* vibration imaging approaches that opens the door for novel vision tasks.

2. Related works

Non-line-of-sight imaging Our method can analyze object interactions beyond the camera’s line of sight. This task relates to optical non-light-of-sight (NLOS) methods that capture light scattering from LOS surfaces to form images of objects around corners [10, 27, 30, 41, 47]. However, unlike optical NLOS, relying on vibrations does not presuppose the existence of a light path between sensor and object, but only the visibility of the *impacted* surface. Through-wall NLOS methods were also explored since longer wavelengths (wifi) can penetrate walls [3, 49]. However, these methods require specialized antenna arrays and can not operate for materials that RF signals can not penetrate (*e.g.*, metals). Our method also relates to seismic imaging, where geophones measure seismic waves at multiple earth points to recover below-ground geological structures [26].

Capturing object vibrations Piezoelectric transducers embedded within the object can provide 1D pressure readings for multiple object points [43]. However, such contact-based sensing limits impact localization to specialized scenarios. Laser Doppler Vibrometers (LDV) can sense vibrations remotely [34], but most LDVs yield a 1D signal (*e.g.*, transverse surface velocity) and are constrained to measure a single point at a time. In this work, we capture the surface vibration using the dual-shutter camera as described in the paper of Sheinin *et al.* [37]. The dual-shutter camera relies on speckle-based vibrometry, which amplifies minute surface vibrations by illuminating the object’s surface with coherent light and imaging interference that the reflected light creates. [4–8, 19, 24, 37–40, 48]. As such, the dual-shutter camera combines three key desirable properties for impact localization: it provides *non-contact sensing* (a) of *2D surface tilts* (b) at *multiple surface points* (c). Our localization method relies on the 2D *geometric* signal stemming from combining (b) and (c).

3. Background

3.1. Elastic waves propagation

Consider a planar object as shown in Fig. 2(a). Let $\mathbf{x} \equiv (x, y)$ denote the spatial coordinates coincident with the object’s surface, and z denote the coordinate perpendicular to the object’s surface plane. The object’s surface is located at $z = 0$. Now, consider a short impulse of force applied to the object surface at position $\mathbf{x} = (0, 0)$. The impact cre-




Figure 3. Speckle-based vibrometry. An object’s surface is illuminated by a laser. A camera is focused on a plane located some distance away from the object’s surface and images the resulting interference pattern (*i.e.*, speckle). In this configuration, the focus-plane speckle is highly sensitive to minute surface tilts, causing the image-plane speckle to shift in relation to the surface tilts.

ates vibrations along the object’s surface and interior. When the object is made of isotropic homogeneous elastic material, the object’s vibration can be described using the elastic wave equation [2, 25]. Namely, the impact creates a *wave* that propagates from the impact location outward. On the object’s surface, the wave creates minute vertical displacements (*i.e.*, along the z axis), which can be sensed remotely using interferometry or speckle-based vibrometry.

Now consider a surface point at $\mathbf{x}_n = (x_n, y_n)$. Let $u_z(\mathbf{x}_n, t)$ denote the vertical displacement (*i.e.*, surface height) at \mathbf{x}_n as a function of time t , where $t=0$ marks the moment of impact. Roughly speaking, the surface height at \mathbf{x}_n is perturbed by several wavefronts. The first wavefront to disturb the point is the P-wave (or pressure wave). The P-wave is a longitudinal wave and therefore yields little vertical displacement. The P-wave is followed by the S- (shear) and R- (Rayleigh) waves which arrive later. The S- and R-waves contain a transverse motion component that causes vertical displacement (see Fig. 2(b)) [32]. Finally, the reflected waves from the object’s edges yield complicated surface vibration patterns due to interference (Fig. 2(c)).

3.2. Speckle-based vibrometry

Speckle-based vibrometry relies on illuminating an object’s surface point with a coherent light source (*e.g.*, a laser) and imaging the resulting speckle-pattern formed on a plane *away* from the object surface (see Fig. 3). The focus-plane speckle pattern is created by the random interference of light reflected from the surface’s microscopic structure. The interference can be constructive or destructive, yielding an image with randomly distributed bright and dark patches.

Minute surface vibrations cause the imaged speckle pattern to shift in the image plane. Specifically, under the configuration illustrated in Fig. 3, the speckle-pattern image shifts (d_x, d_y) are mostly caused by surface tilts angles $\boldsymbol{\theta} \equiv (\theta_x, \theta_y)$ [48]. The measured image-domain shifts can be converted into surface tilts via a linear per-axis factor:

$$(\theta_x, \theta_y) = (h_x, h_y) \odot (d_x, d_y), \quad (1)$$

where \odot is an element-wise product. The scaling factors (h_x, h_y) depend on various factors including the camera optics and focus setting, as well as camera-object distance. See supplementary for details on how to calibrate (h_x, h_y) .

4. Transient Circular Wavefronts

Fig. 2 shows that a short impact on the object’s surface creates elastic waves propagating outward from the impact location. For an infinitely wide isotropic homogeneous surface, symmetry dictates that the resulting displacement is circularly symmetric for $\forall t > 0$.¹ Thus, assuming the impact occurred at point $(0, 0)$, and substitution x, y by $r = \sqrt{x^2 + y^2}$, the surface height at each point can be expressed as $u_z(r, t)$. Therefore, the surface gradient is

$$\begin{aligned} \mathbf{g}(\mathbf{x}, t) &= \nabla u_z = \left(\frac{\partial u_z(r, t)}{\partial x}, \frac{\partial u_z(r, t)}{\partial y} \right) = \\ &= \left(\frac{\partial u_z(r, t)}{\partial r} \frac{\partial r}{\partial x}, \frac{\partial u_z(r, t)}{\partial r} \frac{\partial r}{\partial y} \right) = \\ &= \frac{1}{r} \frac{\partial u_z(r, t)}{\partial r} (x, y) = \frac{\partial u_z(r, t)}{\partial r} \hat{\mathbf{x}}, \end{aligned} \quad (2)$$

where $\hat{\mathbf{x}}$ is a unit vector pointing to \mathbf{x} .

Eq. (2) shows that once the first transverse wave reaches a point at radius r , the surface gradient at that point always points towards or away from the impact source, depending on the sign of $\frac{\partial u_z(r, t)}{\partial r}$. However, Eq. (2) does not hold for finite objects. In finite objects, when the outgoing elastic waves reach the boundaries, reflected waves appear and interfere across the object’s surface, causing $\mathbf{g}(\mathbf{x}, t)$ to point in arbitrary directions.

Our key observation is that for every surface point \mathbf{x} , there may exist a short interval of time after an impact

$$T^c(\mathbf{x}) \equiv [t^{\text{start}}, t^{\text{end}}] \quad (3)$$

during which the outgoing vertical elastic waves (S- and R-waves) generated by the impact have reached the point without strong interference from the object boundaries. Therefore, during $T^c(\mathbf{x})$, point \mathbf{x} is displaced by an approximately circular wavefront for which the gradient at \mathbf{x} indicates the impact source location:

$$\mathbf{g}(\mathbf{x}, t) = A(r, t) \hat{\mathbf{x}} \quad \forall t \in T^c(\mathbf{x}), \quad (4)$$

where $A(r, t)$ is a time-dependent displacement amplitude. Note again that, depending on the sign of $A(r, t)$, the gradient can either point toward or away from the impact source. In this paper, we will refer to $T^c(\mathbf{x})$ as the *stable time interval*, since during this interval, the surface gradient at point \mathbf{x} consistently points at or away from the impact source’s location (see Fig. 4(b)). Conversely, at other times $t \notin T^c(\mathbf{x})$, $\mathbf{g}(\mathbf{x}, t)$ may behave erratically (see Fig. 4(c)).

¹This assumes that the impulse is a delta function in space.

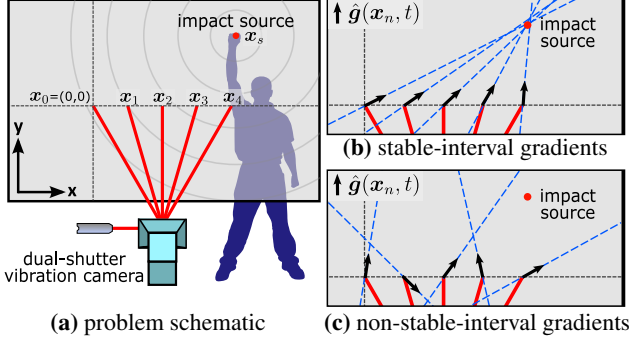


Figure 4. Transient vibration imaging. **(a)** A dual-shutter vibration camera simultaneously captures 2D vibration at $N=5$ surface points. A short impulse of force is applied to the surface at \mathbf{x}_s . **(b)** For a short time interval, defined as the *stable time interval*, the impact generates elastic waves having circular wavefronts. Upon reaching the measured points, the wavefronts create a vertical displacement whose surface gradient points towards or away from the impact source. **(c)** Outside the stable time period, the surface gradients may point in arbitrary directions.

We experimentally show that $T^c(\mathbf{x})$ exists for a variety of materials. In Section 5, we show that measuring multiple surface points having $T^c(\mathbf{x})$ can help localize an unknown impact source. In Section 6, we extend source localization to objects made of a non-isotropic material.

5. Transient-based Source Localization

We capture vibrations using a dual-shutter speckle-based vibration camera [37]. The camera measures the vibrations at N locations on a planar elastic isotropic surface. Let \mathbf{x}_n denote the board measurements locations, where $n = [0, 1, \dots, N-1]$. For convenience, from this point onward, we set the axes origin on the $x-y$ plane to coincide with \mathbf{x}_0 , namely $\mathbf{x}_0 = (0, 0)$.

A short force impulse is applied to the object at an unknown position \mathbf{x}_s . We assume that points \mathbf{x}_n and \mathbf{x}_s are located on the object surface facing the camera (*i.e.*, $z=0$). Moreover, for thin planar objects, we assume that an impact at the object’s back side (as illustrated in Fig 4(a)) is approximately equivalent to an impact at $z=0$.

Our camera measures the surface’s instantaneous tilts for each point n in both axes $\boldsymbol{\theta}_n(t)$, which relate to the surface gradient by an element-wise tangent:

$$\mathbf{g}(\mathbf{x}_n, t) = \tan(\boldsymbol{\theta}_n(t)). \quad (5)$$

Let $t=0$ coincide with the camera’s first vibration measurement. Let $T^c(\mathbf{x}_n)$ denote the stable time interval of point n . Each point n may have a separate stable time interval due to its unique distance to the impact source. Moreover, some measurement configurations may yield measurement points having no stable interval (*i.e.*, $T^c(\mathbf{x}_n) \in \emptyset$). For example, a measurement point located too close to the object boundary

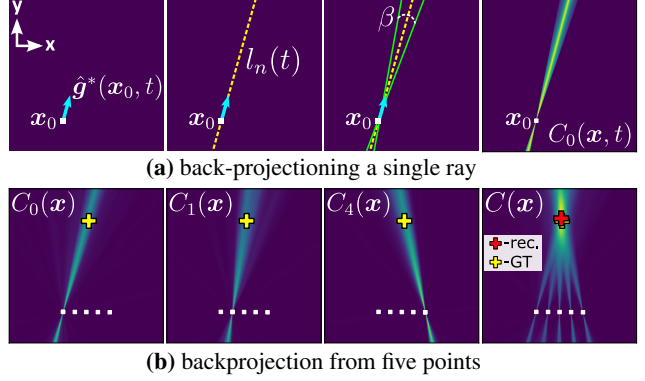


Figure 5. Impact source localization using backprojection. **(a)** The surface gradient during the stable time interval defines a line l that intersects the impact source position. At each time step within the stable interval, we cast a cone of rays centered at $l_n(t)$. **(b)** Per point, we integrate the cones across all times within their corresponding stable time intervals. Finally, we sum the votes from all N points to yield the final backprojection voting map $C(\mathbf{x})$. The impact source is the point \mathbf{x} that maximizes $C(\mathbf{x})$.

may incur wave reflections almost instantaneously with the arrival of the S- and R-waves.

After an impact, the gradient $\mathbf{g}(\mathbf{x}_n, t)$, $t \in T^c(\mathbf{x}_n)$ at each point n defines a line on the $x-y$ plane that approximately intersects the impact source position (see Fig 4(b)). Therefore, to recover \mathbf{x}_s , we must complete two tasks: (a) find $T^c(\mathbf{x}_n)$ for each point and (b) compute the intersecting lines per point and find where all N lines intersect to recover \mathbf{x}_s .

Source localization using backprojection We frame the impact source localization problem as searching for a point that maximizes the agreement between the measured directions for all N points. Inspired by prior works [16, 17, 30], we devise a voting-based method. As illustrated in Fig. 5(a), for each point n , we cast a cone of rays along the line dictated by the surface gradient direction during $T^c(\mathbf{x}_n)$. The cone has an angle of β to take into account possible errors in gradient angles. A map $C(\mathbf{x})$ accumulates the votes from all points n during their individual stable time intervals (Fig. 5(b)). Finally we recover \mathbf{x}_s as the position having the highest value in $C(\mathbf{x})$:

$$\mathbf{x}_s^* = \arg \max_{\mathbf{x}} C(\mathbf{x}). \quad (6)$$

To compute $C(\mathbf{x})$, we first initialize an accumulator array to zero, *i.e.*, $C(\mathbf{x})=0, \forall \mathbf{x}$. Then, for every point n , we apply the following procedure. First, we recover the gradient direction from Eq. (5):

$$\hat{\mathbf{g}}^*(\mathbf{x}_n, t) = \frac{\tan(\boldsymbol{\theta}_n(t))}{\|\tan(\boldsymbol{\theta}_n(t))\|_2} \approx \frac{\boldsymbol{\theta}_n(t)}{\|\boldsymbol{\theta}_n(t)\|_2}, \quad (7)$$

where the second transition in Eq. (7) is due to the small surface vibration displacement angles. Recall that we do not know whether the gradient points to or away from the impact source. Therefore, we define a line that originates at \mathbf{x}_n and runs along the direction dictated by the gradient (yellow dotted line in Fig. 5(a)):

$$l_n(t) = \mathbf{x}_n + s \hat{\mathbf{g}}^*(\mathbf{x}_n, t), \quad s \in \mathbb{R}. \quad (8)$$

At each time step, we compute a weighted 2D cone that follows the bisector line $l_n(t)$. The contribution of point n to the backprojection voting map at time $t \in T^c(\mathbf{x}_n)$ is:

$$C_n(\mathbf{x}, t) = \begin{cases} 0 & \phi_n(\mathbf{x}, t) > \cos \frac{\beta}{2} \\ \exp \left[-\frac{1}{\sigma^2} d(\mathbf{x}, l_n(t))^2 \right] & \text{otherwise,} \end{cases} \quad (9)$$

where

$$\phi_n(\mathbf{x}, t) = \hat{\mathbf{g}}^*(\mathbf{x}_n, t)^T (\mathbf{x} - \mathbf{x}_n) / \|\mathbf{x} - \mathbf{x}_n\|_2 \quad (10)$$

is the cosine of the angle between line $l_n(t)$ and vector $\mathbf{x} - \mathbf{x}_n$, d is the perpendicular distance between $l_n(t)$ and \mathbf{x} , and $\sigma=5$. We integrate over all times $t \in T^c(\mathbf{x}_n)$,

$$C_n(\mathbf{x}) = \sum_{t \in T^c(\mathbf{x}_n)} C_n(\mathbf{x}, t), \quad (11)$$

to get the contribution per point, and over all N points to get the final backprojection voting map:

$$C(\mathbf{x}) = \sum_n C_n(\mathbf{x}). \quad (12)$$

Please see the supplementary materials for a summary of the backprojection algorithm.

Estimating the stable time intervals Following the discussion in Section 4, the stable time interval per point n starts with the arrival of the first transverse wavefront. Therefore, the stable interval start time t_s is determined by the time at which the vibration magnitude surpasses a pre-defined threshold P :

$$t_n^{\text{start}} = \arg \min_t (\|\boldsymbol{\theta}_n(t)\|_2 > P). \quad (13)$$

We apply a high-pass filter to $\boldsymbol{\theta}_n(t)$ before Eq. (13) to increase robustness to ambient low-frequency vibrations.

A short time after the arrival of the first transverse wavefront, reflections from the object's boundaries interfere at point n causing the surface gradient to point in arbitrary directions. The stable interval's duration depends on various factors, including the material of the object, and the distance of the impact and measurement points from the object's boundary. Nevertheless, we experimentally found that a duration of $t^{\text{end}} - t^{\text{start}} = 1.5$ ms holds in most cases.

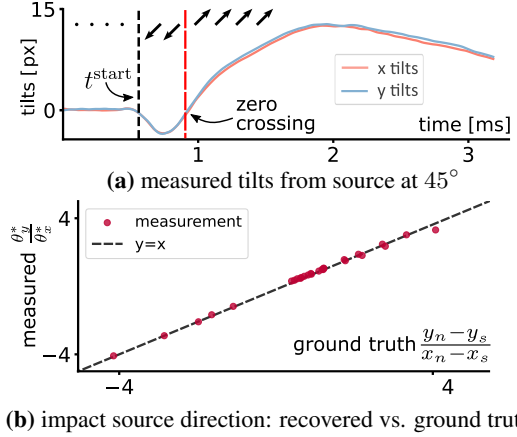


Figure 6. Gradient stable time interval. (a) The measured surface tilts correspond to the instantaneous gradient direction. For each measurement point, we detect the start of the stable interval when the tilts magnitude crosses a pre-defined threshold. In (a), the impact source is located at 45° with respect to the measured point, yielding a ratio of $\theta_y/\theta_x \approx 1$. (b) Experimental validation of the stable time interval hypothesis. The plot shows the recovered gradient direction vs. the ground truth direction computed by knocking at various known \mathbf{x}_s .

Fig. 6(a) shows an example of the measured tilts for a point located at a 45° angle from the impact source. Observe that the tilts in both axes are almost identical as the vibrations begin. Around 1 ms, the gradient direction flips from -135° to 45° . Yet, we do not rely on the gradient's sign, but only on the line it draws on the x - y plane.

Fig. 6(b) shows the matching between the recovered and ground truth gradient directions for a plurality of impact-source locations. The experiment involved knocking at various known plane positions \mathbf{x}_s and comparing the expected gradient direction, up to sign, expressed as $(y_n - y_s)/(x_n - x_s)$. The plot shows good correspondence for various directions, verifying the stable interval assumption.

6. Modeling Anisotropic Materials

In Sections 4 and 5, we assumed an elastic isotropic material. While many materials can be treated as isotropic, there are some notable ubiquitous exceptions like wood (technically considered orthotropic), Polyvinyl chloride (PVC), and porcelain. These materials consist of micro-structures oriented in a preferred direction. For example, Fig. 7(a) shows a thin slab of Engelmann spruce where the fiber direction is along the image's horizontal axis.

In anisotropic materials, the speed of sound varies with the relative angle to the micro-structure direction [12, 23]. This means that a surface impact creates non-circular wavefronts (Fig. 7(a)) [21]. Nevertheless, our experiments show that the surface gradient directions induced by the impact on anisotropic elastic materials have an approximately el-




Figure 7. Anisotropic wave propagation can be approximated by elliptical level-sets. (a) LDV vibration measurements in Engelmann Spruce. The red curves mark the true level sets (75% percentile), while the blue curve marks a fitted ellipse having $m^2 = 3.5$. (b) For anisotropic materials, the measured surface gradient relates to the impact source location via a scalar factor m^2 .

lptical shape. Specifically, the elastic wave level-sets, assuming w.l.o.g $x_s = (0, 0)$, can be approximated in the x - y plane using

$$x^2/m^2 + y^2 = R^2, \quad (14)$$

where m is the aspect ratio of the ellipse level-set and R is a constant (see Fig. 7(a)). Thus, the surface gradients can be approximated using:

$$\frac{g_y}{g_x} \approx \frac{\theta_y}{\theta_x} \approx m^2 \frac{y}{x}. \quad (15)$$

Therefore, given m , we can use Eqs. (8)-(13) to locate the impacts on anisotropic surfaces by replacing $\hat{g}^*(x_n, t)$ in Eq. (8) with

$$\hat{h}^*(x_n, t) \approx \frac{\theta_n(t) \odot (1, m^2)}{\|\theta_n(t) \odot (1, m^2)\|_2}. \quad (16)$$

As further described in the supplementary, we find m per material by capturing knocks at known x_s .

7. Experimental Evaluation

Using a system as described in Sheinin *et al.*'s paper [37], we generated five laser dots using an Edmund Optics 80 grooves/mm transmission grating beamsplitter. Like Sheinin *et al.*, we illuminated the scene using a low-power Thorlabs 4.5 mW 532 nm laser (Thorlabs CPS532) and boosted the signal by placing retro-reflective tape at the measured points. The camera sampled the scene vibrations at 63 kHz.

To validated our models, we generate dense vibration measurements for various surfaces using a laser Doppler vibrometer (LDV) (Polytec PDV-100) and mirror galvanometer synchronized with an impact hammer (PCB 086E80).

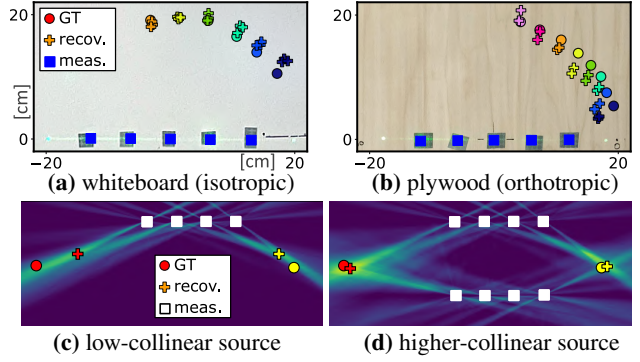


Figure 8. Source localization on isotropic and anisotropic materials. (a) Impact localization on a whiteboard using five measurement points. (b) Localization on a slab of plywood. Average localization error was 1.1 cm and 2.1 cm for the whiteboard and plywood, respectively. (c) Collinear measurement points have high uncertainty at grazing angles. (d) Adding non-collinear measurement points reduces the triangulation uncertainty.

The hammer measured the instantaneous force applied during the strike [33]. The dense vibration videos were created by repeatedly knocking the board at the same location while moving the LDV's measurement position at each repetition.

7.1. Impact-source localization

We tested impact localization on various isotropic and anisotropic materials. Fig. 8 shows example localization for an isotropic whiteboard and an orthotropic sheet of plywood. The whiteboard shown in Fig. 8(a) has dimensions of 110 cm by 290 cm, while the plywood slab in Fig. 8(b) was 85 cm by 65 cm. Our imaging system was set around two meters from the boards in both experiments. While the whiteboard exhibited isotropic behavior, the plywood had an elliptical coefficient of $m^2 = 0.67$, which we calibrated by knocking at a set of known points.

On each board, we knocked at a set of points having a radius of about 20 cm from the middle measurement point. We repeat each point a few times. Fig. 8 shows that our method can recover the knock locations accurately. The orthotropic plywood shows a more considerable variance in estimation accuracy. This could result from the heterogeneous planar grain arrangement and the plywood being constructed by gluing several layers of wood (in the z -axis) with different fiber orientations. Please see the supplementary materials for experiments on additional materials, including fiberboard, particle board glass, PVC panels, porcelain, and gypsum.

The dual-shutter camera used in this work was limited to producing sets of collinear points. As seen in Fig. 8, the accuracy of triangulation via backprojection depends on the angle between the source point and the line defined by the measurement points. This behavior agrees with prior analyses of the triangulation error with respect to angles to tar-




Figure 9. Localizing ping-pong ball strikes mid-play. Our camera measures five points on the table’s *bottom* surface (see Fig. 1). The five markers on the top side visualize the bottom locations. We visualize the recovered ball strike locations using two concentric red circles. We also super-impose the backprojection voting map per strike (bright yellow). The motion blur trajectories help infer the ball’s “real” impact locations. See the project page for videos of results [1].

get [35]. In the extreme case, the location of a source that is collinear to the measurement point positions can not be recovered using a line configuration. Fig. 8(c)-(d) shows that non-collinear point measurement arrangements can yield results with superior accuracy.

Localizing ping-pong ball strikes mid-game Fig. 9 shows the application of our method to localize the strikes of a ping-pong ball on a table during play. Here our camera measures the ping-pong table vibrations from below. A side RGB camera is used to super-impose the recovered impact locations on the table surface. The table surface is a wood and aluminum composite and is approximately isotropic with respect to the generated elastic waves. We recorded ten video clips lasting eight seconds each. Each clip contains between two to six hits. As can be seen, *our method correctly recovers the ball landing positions, mid-play, without a line of sight to the ball*. Our average error for these ten video clips is around 2.9 cm. Please see the supplementary materials for videos and more results.

Localizing footsteps using vibrations When walking, the foot creates floor vibrations that originate at the stepping location. Therefore, our system can localize the step locations by observing the floor vibrations. Fig. 10 shows footstep localization on a hardwood floor. It is noteworthy that the experimental conditions in Fig. 10 deviate from the assumptions of Sections 4-5 in several regards. First, the force profile exerted on the floor by the stepping leg is not a short and localized impulse, but a prolonged pressure having a spatially wide footprint (foot’s stepping area). Secondly, the floor is a heterogeneous medium since it is constructed by stacking up many bamboo planks.




Figure 10. Localizing footsteps using vibrations. A footstep creates vibrations that propagate from the step location through the floor medium. **(Top-row)** Our camera recovers the footstep locations by measuring the floor vibrations, without requiring line of sight. **(Middle-row)** Recovery using five floor points. **(Bottom-row)** Synthesized recovery using ten floor points.

Still, our method could accurately infer the footstep direction using five measurement points (Fig. 10(Middle-row)) and the footstep location using ten non-collinear measurement points (Fig. 10(Bottom-row)). Since our camera can currently only support five measurement points, the experiment results in (Fig. 10(Bottom-row)) were synthesized using two pairs of non-simultaneous measurements having five points each. Please see supplementary for details on how we create synthesized results.

7.2. Measuring material anisotropy

So far, we have concerned ourselves with recovering unknown impact sources given a known (or calibrated) material. However, measuring vibrations for *known* impulse locations allows inferring the material anisotropy factor m by finding the value which best fits the observed vibrations. We tested our method on various materials (see Fig. 11). While all the isotropic materials have $m \approx 1$, the anisotropic materials displayed a variety of m values, suggesting that recovering m and comparing it to a pre-collected dataset may help classify the material remotely.

8. Towards inferring force & object shape

Based on the experiments described in this section, we postulate that the transient vibrations contain additional cues about the impact force and impacting material properties. The impact force relates to the vibrations’ magnitude. We demonstrate this relation in Fig. 12(a), where we drop

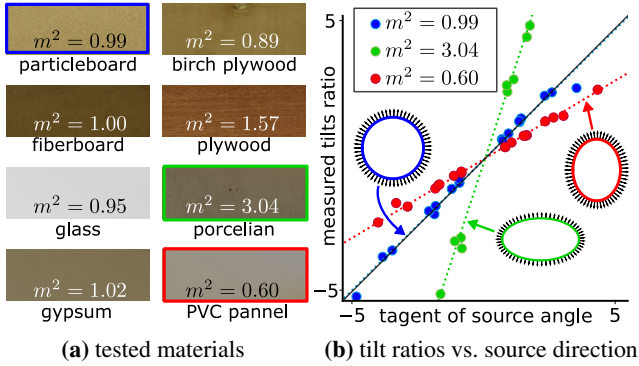


Figure 11. Transient vibration analysis for different materials. (a) We calibrate the wave propagation model for various materials. Calibration consists for knocking on the surface at several known locations. (b) For each material, we fit the anisotropy factor m . Isotropic materials yield $m = 1$, representing circular wavefronts (blue circle). Anisotropic-material wavefronts are approximately elliptical (red and blue ellipses). Once m is known, we can apply our method to localize impulses at unknown surface locations.




Figure 12. Vibrations contain cues about force & impacting object. (a) We dropped the ball from different heights and measured the peak vibration magnitude. (b) The object’s shape and stiffness affect the resulting vibrations’ spectral composition.

a ping-pong ball from varying known heights at the same point. The plot in Fig. 12(a) shows a square root relation between peak vibration magnitude and the drop height. Since in free fall, the velocity has the same square root relation to height; the vibration magnitude is linear to the ball’s velocity, which is linear to the ball’s peak generated force [22]. Our experiment measured the vibration magnitudes when dropping the ball at a single point. However, recovering the peak force at every surface point (using the same five measurement points) requires accounting for more factors, like the distance between the impact and measurement points.

We also posit that the vibrations’ spectral decomposition holds cues about the shape and stiffness of the impacting object. Our experiments show that sharp and hard objects produce impacts that resemble a spatiotemporal delta function and yield high-frequency vibrations. Conversely, soft objects having a larger spatial footprint upon impact yielded smoother measured vibrations (see Fig. 12(b)). Our preliminary analysis leads us to believe that the shape of the impact object could be reasoned about from the vibrations.

9. Discussion

Beyond homogeneous planar surfaces We assumed a simple propagation model that applies to homogeneous planar surfaces. Our model does not extend easily for more complex structures such as curved objects, surfaces having multiple concatenated materials (*e.g.*, floor-to-wall vibrations), and highly heterogeneous materials. In these cases, elastic wave propagation strongly depends on the object’s unknown spatially-varying geometry and material. Nevertheless, the vibrations generated in such complex structures are fertile ground for future works to extract more novel scene information. Another exciting research avenue is to localize sound sources in 3D, like the locations of survivors trapped beneath the rubble of fallen buildings.

Limitations and opportunities of speckle-vibrometry

Speckle-vibrometry is an active approach in which the optical signal depends on the amount of light reflected from the target. Thus, performance degrades for low-albedo materials, specular materials, and far-away objects. Our camera sampled the scene at 63 kHz, which is too slow to reliably measure time-of-arrival delays between our measurement points for most materials.² However, faster future systems might also use time-of-arrival to improve localization.

While our camera only measured surface tilts, prior works used speckle to sense additional motions, including rotation, translation, and axial shift [24, 29, 50]. Thus, future works could incorporate these degrees of freedom to model elastic wave propagation more accurately and produce better source localizations.

10. Conclusion

This paper explores the mostly untapped potential of imaging and extracting information from transient surface vibrations. We presented an approach to localize impact sources by measuring transient surface vibrations at multiple surface points. We also showed that these vibrations might contain more information about the physical interactions like the impact force and properties of impacting object. We believe that our method paves the way for future research to develop new types of scene understanding based on the subtle vibration cues people and objects create, which are invisible to human eyes and conventional cameras.

Acknowledgements: This work was supported in parts by NSF Grants IIS-1900821 and CCF-1730147. We thank B. Li, H. Yu, J. Chen for help with the experiments.

²*E.g.*, the speed of sound in wood is generally between 3300-5000 m/s, which yields about one sample delay between two measurement points.

References

- [1] Analyzing physical impacts using transient surface wave imaging: Project webpage. https://imaging.cs.cmu.edu/transient_vibrations/, 2023. 1, 7
- [2] Jan Achenbach. *Wave propagation in elastic solids*. Elsevier, 2012. 3
- [3] Fadel Adib and Dina Katabi. See through walls with wifi! In *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*, pages 75–86, 2013. 2
- [4] Marina Alterman, Chen Bar, Ioannis Gkioulekas, and Anat Levin. Imaging with local speckle intensity correlations: theory and practice. *ACM Transactions on Graphics (TOG)*, 40(3):1–22, 2021. 2
- [5] Marina Alterman, Chen Bar, Ioannis Gkioulekas, and Anat Levin. Near-field imaging inside scattering layers. In *Computational Optical Sensing and Imaging*, pages CW3B–2. Optica Publishing Group, 2021. 2
- [6] Chen Bar, Marina Alterman, Ioannis Gkioulekas, and Anat Levin. A monte carlo framework for rendering speckle statistics in scattering media. *ACM Transactions on Graphics (TOG)*, 38(4):1–22, 2019. 2
- [7] Chen Bar, Marina Alterman, Ioannis Gkioulekas, and Anat Levin. Single scattering modeling of speckle correlation. In *2021 IEEE International Conference on Computational Photography (ICCP)*, pages 1–16. IEEE, 2021. 2
- [8] Silvio Bianchi and Emanuele Giacomozzi. Long-range detection of acoustic vibrations by speckle tracking. *Applied optics*, 58(28):7805–7809, 2019. 1, 2
- [9] 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, 2013. 1
- [10] Katherine L Bouman, Vickie Ye, Adam B Yedidia, Frédo Durand, Gregory W Wornell, Antonio Torralba, and William T Freeman. Turning corners into cameras: Principles and methods. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2270–2278, 2017. 2
- [11] Oral Buyukozturk, Justin G Chen, Neal Wadhwa, Abe Davis, Frédo Durand, and William T Freeman. Smaller than the eye can see: Vibration analysis with video cameras. In *19th World Conference on Non-Destructive Testing 2016 (WC-NDT)*, 2016. 1
- [12] P Chadwick and GD Smith. Foundations of the theory of surface waves in anisotropic elastic materials. *Advances in applied mechanics*, 17:303–376, 1977. 5
- [13] Justin G Chen, Abe Davis, Neal Wadhwa, Frédo Durand, William T Freeman, and Oral Büyüköztürk. Video camera-based vibration measurement for civil infrastructure applications. *Journal of Infrastructure Systems*, 23(3):B4016013, 2017. 1
- [14] 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, 2015. 1
- [15] Abe Davis, Michael Rubinstein, Neal Wadhwa, Gautham J Mysore, Fredo Durand, and William T Freeman. The visual microphone: Passive recovery of sound from video. 2014. 1
- [16] Richard O Duda and Peter E Hart. Use of the hough transformation to detect lines and curves in pictures. *Communications of the ACM*, 15(1):11–15, 1972. 4
- [17] Lee A Feldkamp, Lloyd C Davis, and James W Kress. Practical cone-beam algorithm. *Josa a*, 1(6):612–619, 1984. 4
- [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. 1
- [19] DA Gregory. Basic physical principles of defocused speckle photography: a tilt topology inspection technique. *Optics & Laser Technology*, 8(5):201–213, 1976. 2
- [20] Michael A Greminger and Bradley J Nelson. Vision-based force measurement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(3):290–298, 2004. 2
- [21] Klaus Helbig and Leon Thomsen. 75-plus years of anisotropy in exploration and reservoir seismics: A historical review of concepts and methods. *Geophysics*, 70(6):9ND–23ND, 2005. 5
- [22] Mont Hubbard and WJ Stronge. Bounce of hollow balls on flat surfaces. *Sports Engineering*, 4(2):49–61, 2001. 8
- [23] KA Ingebrigtsen and A Tonning. Elastic surface waves in crystals. *Physical Review*, 184(3):942, 1969. 5
- [24] Kensei Jo, Mohit Gupta, and Shree K Nayar. Spedo: 6 dof ego-motion sensor using speckle defocus imaging. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4319–4327, 2015. 2, 8
- [25] Eduardo Kausel. Lamb’s problem at its simplest. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 469(2149):20120462, 2013. 3
- [26] Q Liu and YJ Gu. Seismic imaging: From classical to adjoint tomography. *Tectonophysics*, 566:31–66, 2012. 2
- [27] Xiaochun Liu, Sebastian Bauer, and Andreas Velten. Phasor field diffraction based reconstruction for fast non-line-of-sight imaging systems. *Nature communications*, 11(1):1645, 2020. 2
- [28] MT Martin and JF Doyle. Impact force identification from wave propagation responses. *International journal of impact engineering*, 18(1):65–77, 1996. 2
- [29] Alex Olwal, Andrew Bardagjy, Jan Zizka, and Ramesh Raskar. Speckleeye: gestural interaction for embedded electronics in ubiquitous computing. In *CHI’12 Extended Abstracts on Human Factors in Computing Systems*, pages 2237–2242. 2012. 8
- [30] Matthew O’Toole, David B Lindell, and Gordon Wetzstein. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature*, 555(7696):338–341, 2018. 2, 4
- [31] Siyou Pei, Pradyumna Chari, Xue Wang, Xiaoying Yang, Achuta Kadambi, and Yang Zhang. Forcesight: Non-contact force sensing with laser speckle imaging. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, pages 1–11, 2022. 2

- [32] Ante Qu and Doug L James. On the impact of ground sound. *arXiv preprint arXiv:1909.09235*, 2019. 3
- [33] Mark Rau, Julius O Smith, and Doug L James. Augmenting a single-point laser doppler vibrometer to perform scanning measurements. *The Journal of the Acoustical Society of America*, 151(4):A157–A157, 2022. 6
- [34] Steve Rothberg, JR Baker, and Neil A Halliwell. Laser vibrometry: pseudo-vibrations. 1989. 2
- [35] Markus Rumpler, Arnold Irschara, and Horst Bischof. Multi-view stereo: Redundancy benefits for 3d reconstruction. In *35th Workshop of the Austrian Association for Pattern Recognition*, volume 4, page 25. OAGM, 2011. 7
- [36] Ryusuke Sagawa, Yusuke Higuchi, Ryo Furukawa, and Hiroshi Kawasaki. Acquisition and visualization of microvibration of a sound wave in 3d space. *Journal of Robotics and Mechatronics*, 34(5):1024–1032, 2022. 2
- [37] Mark Sheinin, Dorian Chan, Matthew O’Toole, and Srinivasa G Narasimhan. Dual-shutter optical vibration sensing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16324–16333, 2022. 1, 2, 4, 6
- [38] Yi Chang Shih, Abe Davis, Samuel W Hasinoff, Frédo Durand, and William T Freeman. Laser speckle photography for surface tampering detection. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 33–40. IEEE, 2012. 2
- [39] Brandon M Smith, Pratham Desai, Vishal Agarwal, and Mohit Gupta. Colux: Multi-object 3d micro-motion analysis using speckle imaging. *ACM Transactions on Graphics (TOG)*, 36(4):1–12, 2017. 2
- [40] Brandon M Smith, Matthew O’Toole, and Mohit Gupta. Tracking multiple objects outside the line of sight using speckle imaging. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6258–6266, 2018. 2
- [41] Andreas Velten, Thomas Willwacher, Otkrist Gupta, Ashok Veeraraghavan, Mounsi G Bawendi, and Ramesh Raskar. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature communications*, 3(1):745, 2012. 2
- [42] 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. 1
- [43] Wikipedia. Piezoelectric accelerometer — Wikipedia, the free encyclopedia, 2023. [Online; accessed on March 21, 2023]. 2
- [44] Hao-Yu Wu, Michael Rubinstein, Eugene Shih, John Guttag, Frédo Durand, and William Freeman. Eulerian video magnification for revealing subtle changes in the world. *ACM transactions on graphics (TOG)*, 31(4):1–8, 2012. 1
- [45] Nan Wu and Shinichiro Haruyama. Fast motion estimation of one-dimensional laser speckle image and its application on real-time audio signal acquisition. In *2020 the 6th International Conference on Communication and Information Processing*, pages 128–134, 2020. 1
- [46] Nan Wu and Shinichiro Haruyama. The 20k samples-per-second real time detection of acoustic vibration based on displacement estimation of one-dimensional laser speckle images. *Sensors*, 21(9):2938, 2021. 1
- [47] Shumian Xin, Sotiris Nousias, Kiriakos N Kutulakos, Aswin C Sankaranarayanan, Srinivasa G Narasimhan, and Ioannis Gkioulekas. A theory of fermat paths for non-line-of-sight shape reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6800–6809, 2019. 2
- [48] Zeev Zalevsky, Yevgeny Beiderman, Israel Margalit, Shimshon Gingold, Mina Teicher, Vicente Mico, and Javier Garcia. Simultaneous remote extraction of multiple speech sources and heart beats from secondary speckles pattern. *Optics express*, 17(24):21566–21580, 2009. 1, 2, 3
- [49] Mingmin Zhao, Tianhong Li, Mohammad Abu Alsheikh, Yonglong Tian, Hang Zhao, Antonio Torralba, and Dina Katabi. Through-wall human pose estimation using radio signals. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7356–7365, 2018. 2
- [50] Jan Zizka, Alex Olwal, and Ramesh Raskar. Specklesense: fast, precise, low-cost and compact motion sensing using laser speckle. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, pages 489–498, 2011. 8