



Motion microscopy for visualizing and quantifying small motions

Neal Wadhwa^{a,1}, Justin G. Chen^{a,b}, Jonathan B. Sellon^{c,d}, Donglai Wei^a, Michael Rubinstein^e, Roozbeh Ghaffari^d, Dennis M. Freeman^{c,d,f}, Oral Büyüköztürk^b, Pai Wang^g, Sijie Sun^g, Sung Hoon Kang^{g,h,i}, Katia Bertoldi^g, Frédo Durand^{a,f}, and William T. Freeman^{a,e,f,2}

^aComputer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139; ^bDepartment of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139; ^cHarvard-MIT Program in Health Sciences and Technology, Cambridge, MA 02139; ^dResearch Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139; ^eGoogle Research, Google Inc. Cambridge, MA 02139; ^fDepartment of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139; ^gSchool of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138; ^hDepartment of Mechanical Engineering, Johns Hopkins University, Baltimore, MD 21218; and ⁱHopkins Extreme Materials Institute, Johns Hopkins University, Baltimore, MD 21218

Edited by William H. Press, University of Texas at Austin, Austin, TX, and approved August 22, 2017 (received for review March 5, 2017)

Although the human visual system is remarkable at perceiving and interpreting motions, it has limited sensitivity, and we cannot see motions that are smaller than some threshold. Although difficult to visualize, tiny motions below this threshold are important and can reveal physical mechanisms, or be precursors to large motions in the case of mechanical failure. Here, we present a “motion microscope,” a computational tool that quantifies tiny motions in videos and then visualizes them by producing a new video in which the motions are made large enough to see. Three scientific visualizations are shown, spanning macroscopic to nanoscopic length scales. They are the resonant vibrations of a bridge demonstrating simultaneous spatial and temporal modal analysis, micrometer vibrations of a metamaterial demonstrating wave propagation through an elastic matrix with embedded resonating units, and nanometer motions of an extracellular tissue found in the inner ear demonstrating a mechanism of frequency separation in hearing. In these instances, the motion microscope uncovers hidden dynamics over a variety of length scales, leading to the discovery of previously unknown phenomena.

visualization | motion | image processing

Motion microscopy is a computational technique to visualize and analyze meaningful but small motions. The motion microscope enables the inspection of tiny motions as optical microscopy enables the inspection of tiny forms. We demonstrate its utility in three disparate problems from biology and engineering: visualizing motions used in mammalian hearing, showing vibration modes of structures, and verifying the effectiveness of designed metamaterials.

The motion microscope is based on video magnification (1–4), which processes videos to amplify small motions of any kind in a specified temporal frequency band. We extend the visualization produced by video magnification to scientific and engineering analysis. In addition to visualizing tiny motions, we quantify both the object’s subpixel motions and the errors introduced by camera sensor noise (5). Thus, the user can see the magnified motions and obtain their values, with variances, allowing for both qualitative and quantitative analyses.

The motion microscope characterizes and amplifies tiny local displacements in a video by using spatial local phase. It does this by transforming the captured intensities of each frame’s pixels into a wavelet-like representation where displacements are represented by phase shifts of windowed complex sine waves. The representation is the complex steerable pyramid (6), an over-complete linear wavelet transform, similar to a spatially localized Fourier transform. The transformed image is a sum of basis functions, approximated by windowed sinusoids (Fig. S1), that are simultaneously localized in spatial location (x, y), scale r , and orientation θ . Each basis function coefficient gives spatially

local frequency information and has an amplitude $A_{r,\theta}(x, y)$ and a phase $\phi_{r,\theta}(x, y)$.

To amplify motions, we compute the unwrapped phase difference of each coefficient of the transformed image at time t from its corresponding value in the first frame,

$$\Delta\phi_{r,\theta}(x, y, t) := \phi_{r,\theta}(x, y, t) - \phi_{r,\theta}(x, y, 0). \quad [1]$$

We isolate motions of interest and remove components due to noise by temporally and spatially filtering $\Delta\phi_{r,\theta}$. We amplify the filtered phase shifts by the desired motion magnification factor to obtain modified phases for each basis function at each time t . We then transform back each frame’s steerable pyramid to produce the motion-magnified output video (Fig. S2) (3).

We estimate motions under the assumption that there is a single, small motion at each spatial location. In this case, each coefficient’s phase difference, $\Delta\phi_{r,\theta}$, is approximately equal to the dot product of the corresponding basis function’s orientation and the 2D motion (7) (*Relation Between Local Phase Differences and Motions*). The reliability of spatial local phase varies across scale and orientations, in direct proportion to the coefficient’s amplitude (e.g., coefficients for basis functions orthogonal to an edge are more reliable than those along it) (Fig. S3 and

Significance

Humans have difficulty seeing small motions with amplitudes below a threshold. Although there are optical techniques to visualize small static physical features (e.g., microscopes), visualization of small dynamic motions is extremely difficult. Here, we introduce a visualization tool, the motion microscope, that makes it possible to see and understand important biological and physical modes of motion. The motion microscope amplifies motions in a captured video sequence by re-rendering small motions to make them large enough to see and quantifies those motions for analysis. Amplification of these tiny motions involves careful noise analysis to avoid the amplification of spurious signals. In the representative examples presented in this study, the visualizations reveal important motions that are invisible to the naked eye.

Author contributions: N.W., J.G.C., J.B.S., D.W., M.R., R.G., D.M.F., O.B., S.H.K., K.B., F.D., and W.T.F. designed research; N.W., J.G.C., J.B.S., D.W., R.G., P.W., S.S., S.H.K., and W.T.F. performed research; N.W., J.G.C., J.B.S., and D.W. analyzed data; and N.W., J.G.C., J.B.S., D.W., R.G., D.M.F., O.B., P.W., S.S., S.H.K., K.B., F.D., and W.T.F. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

¹Present address: Google Research, Google Inc. Mountain View, CA 94043.

²To whom correspondence should be addressed. Email: billf@mit.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1703715114/-DCSupplemental.

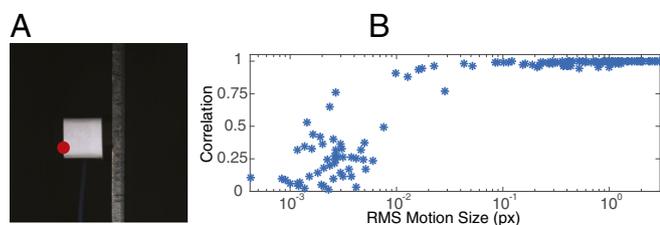


Fig. 1. A comparison of our quantitative motion estimation vs. a laser vibrometer. Several videos of a cantilevered beam excited by a shaker were taken with varying focal length, exposure times, and excitation magnitude. The horizontal, lateral motion of the red point was also measured with a laser vibrometer. (A) A frame from one video. (B) The correlation between the two signals across the videos vs. root mean square (RMS) motion size in pixels (px). Only motions at the red point in A were used in our analysis. More results are in Fig. S4.

Low-Amplitude Coefficients Have Noisy Phase). We combine information about the motion from multiple orientations by solving a weighted least squares problem with weights equal to the amplitude squared. The result is a 2D motion field. This processing is accurate, and we provide comparisons to other algorithms and sensors (Fig. 1, *Synthetic Validation*, and Figs. S4 and S5).

For a still camera, the sensitivity of the motion microscope is mostly limited by local contrast and camera noise—fluctuations of pixel intensities present in all videos (5). When the video is motion-magnified, this noise can lead to spurious motions, especially at low-contrast edges and textures (Fig. S6). We measure motion noise level by computing the covariance matrix of each estimated motion vector. Estimating this directly from the input video is usually impossible, because it requires observing the motions without noise. We solve this by creating a simulated noisy video with zero motion, replicating a static frame of the input video and adding realistic, independent noise to

each frame. We compute the sample covariance of the estimated motion vectors in this simulated video (Fig. S7 and *Noise Model and Creating Synthetic Video*). We show analytically, and via experiments in which the motions in a temporal band are known to be zero, that these covariance matrices are accurate for real videos (*Analytic Justification of Noise Analysis* and Figs. S8 and S9). We also analyze the limits of our technique by comparing to a laser vibrometer and show that, with a Phantom V-10 camera, at a high-contrast edge, the smallest motion we can detect is on the order of 1/100th of a pixel (Fig. 1 and Fig. S4).

Results and Discussion

We applied the motion microscope to several problems in biology and engineering. First, we used it to reveal one component of the mechanics of hearing. The mammalian cochlea is a remarkable sensor that can perform high-quality spectral analysis to discriminate as many as 30 frequencies in the interval of a semitone (8). These extraordinary properties of the hearing organ depend on traveling waves of motion that propagate along the cochlear spiral. These wave motions are coupled to the extremely sensitive sensory receptor cells via the tectorial membrane, a gelatinous structure that is 97% water (9).

To better understand the functional role of the tectorial membrane in hearing, we excised segments of the tectorial membrane from a mouse cochlea and stimulated it with audio frequency vibrations (Movie S1 and Fig. 2A). Prior work suggested that motions of the tectorial membrane would rapidly decay with distance from the point of stimulation (10). The unprocessed video of the tectorial membrane appeared static, making it difficult to verify this. However, when the motions were amplified 20 times, waves that persisted over hundreds of micrometers were revealed (Movie S1 and Fig. 2B–E).

Subpixel motion analysis suggests that these waves play a prominent role in determining the sensitivity and frequency selectivity of hearing (11–14). Magnifying motions has provided new insights into the underlying physical mechanisms of hearing.

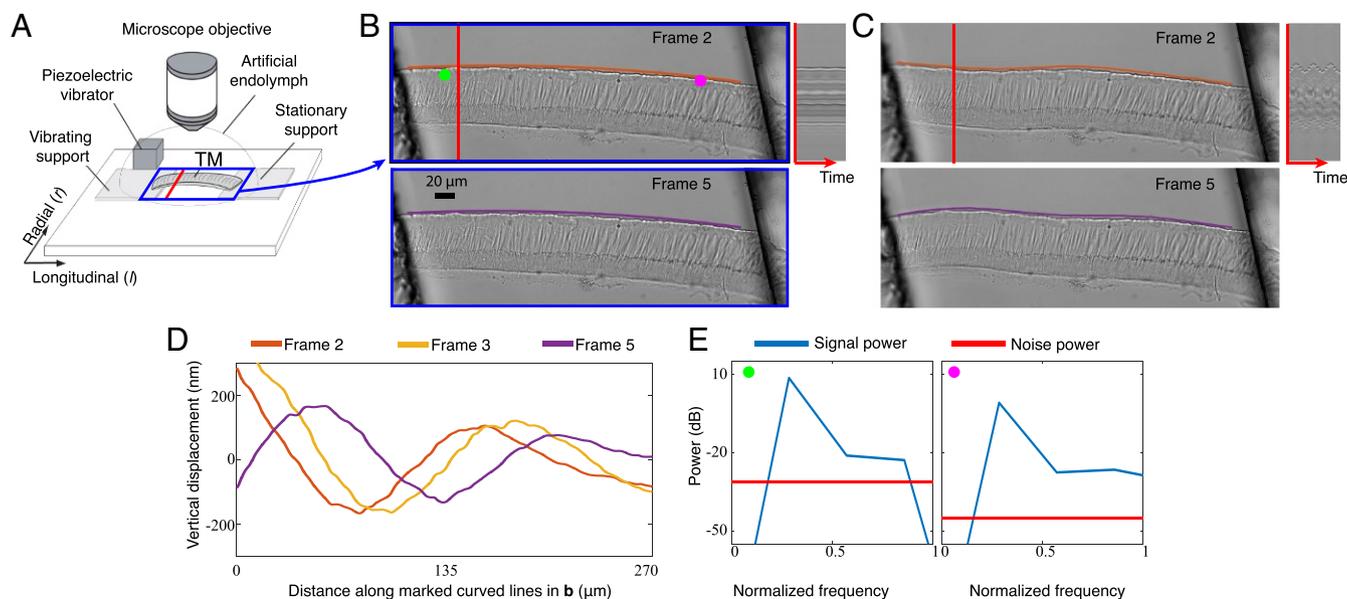


Fig. 2. Exploring the mechanical properties of a mammalian tectorial membrane with the motion microscope. (A) The experimental setup used to stroboscopically film a stimulated mammalian tectorial membrane (*Tecta*^{Y1870C/+}). Subfigure Copyright (2007) National Academy of Sciences of the United States of America. Reproduced from ref. 12. (B) Two of the eight captured frames. (Movie S1, data previously published in ref. 13). (C) Corresponding frames from the motion-magnified video in which displacement from the mean was magnified 20 \times . The orange and purple lines on top of the tectorial membrane in B are warped according to magnified motion vectors to produce the orange and purple lines in C. (D) The vertical displacement along the orange and purple lines in B is shown for three frames. (E) The power spectrum of the motion signal and noise power is shown in the direction of least variance at the magenta and green points in B.

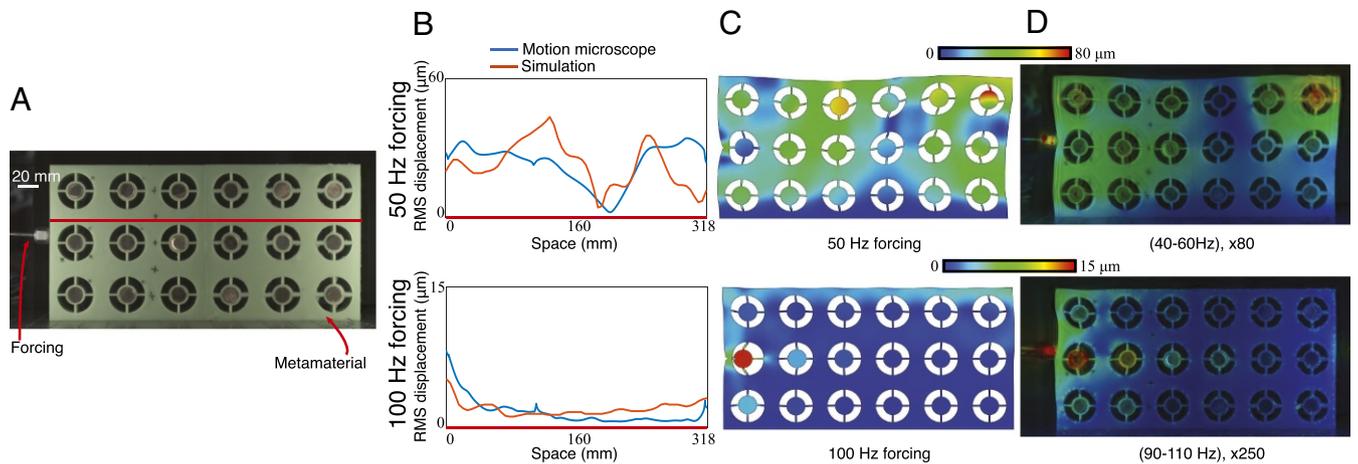


Fig. 4. The motion microscope is used to investigate properties of a designed metamaterial. (A) The metamaterial is forced at 50 Hz and 100 Hz in two experiments, and a frame from the 50-Hz video is shown. (B) One-dimensional slices of the displacement amplitude along the red line in A are shown for both a finite element analysis simulation and the motion microscope. (C) A finite element analysis simulation of the displacement of the metamaterial. Color corresponds to displacement amplitude, and the material is warped according to magnified simulated displacement vectors. (D) Results from the motion microscope are shown. Displacement magnitudes are shown in color at every point on the metamaterial, overlaid on frames from the motion-magnified videos (Movies S4 and S5).

The elastic metamaterial was forced at two frequencies, 50 Hz and 100 Hz, and, in each case, it was filmed at 500 frames per second (FPS) (Fig. 4A). The motions in 20-Hz bands around the forcing frequencies were amplified, revealing that the metamaterial functions as expected (24), passing 50-Hz waves and rapidly attenuating 100-Hz waves (Movies S4 and S5). We also compared our results with predictions from a finite element analysis simulation (Fig. 4B and C). In Fig. 4D, we show heatmaps of the estimated displacement amplitudes overlaid on the motion-magnified frames. We interpolated displacements into textureless regions, which had noisy motion estimates. The agreement between the simulation (Fig. 4C) and the motion microscope (Fig. 4D) demonstrates the motion microscope's usefulness in verifying the correct function of the metamaterial.

Conclusion

Small motions can reveal important dynamics in a system under study, or can foreshadow large-scale motions to come. Motion microscopy facilitates their visualization, and has been demonstrated here for motion amplification factors from $20\times$ to $400\times$ across length scales ranging from 100 nm to 0.3 mm.

Materials and Methods

Quantitative Motion Estimation. For every pixel at location (x, y) and time t , we combine spatial local phase information in different subbands of the frames of the input video using the least squares objective function,

$$\arg \min_{u, v} \sum_i A_{r_i, \theta_i}^2 \left[\left(\frac{\partial \phi_{r_i, \theta_i}}{\partial x}, \frac{\partial \phi_{r_i, \theta_i}}{\partial y} \right) \cdot (u, v) - \Delta \phi_{r_i, \theta_i} \right]^2. \quad [2]$$

Arguments have been suppressed for readability; $A_{r_i, \theta_i}(x, y, t)$ and $\phi_{r_i, \theta_i}(x, y, t)$ are the spatial local amplitude and phase of a steerable pyramid representation of the image, and $u(x, y, t)$ and $v(x, y, t)$ are the horizontal and vertical motions, respectively, at every pixel. The solution ($\mathbf{V} = (u, v)$) is our motion estimate and is equal to

$$\mathbf{V} = (\mathbf{X}^T \mathbf{W} \mathbf{X})^{-1} (\mathbf{X}^T \mathbf{W} \mathbf{Y}), \quad [3]$$

where \mathbf{X} is $N \times 2$ with i th row $(\frac{\partial \phi_{r_i, \theta_i}}{\partial x}, \frac{\partial \phi_{r_i, \theta_i}}{\partial y})$, \mathbf{Y} is $N \times 1$ with i th row $\Delta \phi_{r_i, \theta_i}$, and \mathbf{W} is a diagonal $N \times N$ matrix with i th diagonal element A_{r_i, θ_i}^2 .

To increase the signal-to-noise ratio, we assume the motion field is constant in a small window around each pixel. This gives additional constraints from neighboring pixels, weighted by both their amplitude squared and the corresponding value in a smoothing kernel K , to the objective described in Eq. 3. To handle temporal filtering, we replace the local phase variations $\Delta \phi_{r, \theta}(x, y, t)$ with temporally filtered local phase variations.

We use a four-orientation complex steerable pyramid specified by Portilla and Simoncelli (25). We use only the two highest-frequency scales of the complex steerable pyramid, for a total of eight subbands. We use a Gaussian spatial smoothing kernel with a SD of 3 pixels and a support of 19×19 pixels. The temporal filter depends on the application.

Noise Model and Creating Synthetic Video. We estimate the noise level function (26) of a video. We apply derivative of Gaussian filters to the image in the x and y directions and use them to compute the gradient magnitude. We exclude pixels where the gradient magnitude is above 0.05 on a 0 to 1 intensity scale. At the remaining pixels, we take the temporal variance and mean of the image. We divide the intensity range into 64 equally sized bins. For each bin, we take all pixels with mean inside that bin and take the mean of the corresponding temporal variances of I to form 64 points that are linearly interpolated to estimate the noise level function f .

Estimating Covariance Matrices of Motion Vectors. For an input video $I(x, y, t)$, we use the noise level function f to create a synthetic video

$$I_s(x, y, t) = I_0(x, y, 0) + I_n(x, y, t) \sqrt{f(I_0(x, y, 0))} \quad [4]$$

that is N frames long. We estimate the covariance matrices of the motion vectors by taking the temporal sample covariance of I_s ,

$$\Sigma_v = \frac{1}{N-1} \sum_t (\mathbf{V}_s(x, y, t) - \bar{\mathbf{V}}_s(x, y)) (\mathbf{V}_s(x, y, t) - \bar{\mathbf{V}}_s(x, y))^T, \quad [5]$$

where $\bar{\mathbf{V}}_s(x, y)$ is the mean over t of the motion vectors.

The temporal filter reduces noise and decreases the covariance matrix. Oppenheim and Schaffer (27) show that a signal with independent and identically distributed (IID) noise of variance σ^2 , when filtered with a filter with impulse response $T(t)$, has variance $\sum_t T(t)^2 \sigma^2$. Therefore, when a temporal filter is used, we multiply the covariance matrix by $\sum_t T(t)^2$.

Comparison of Our Motion Estimation to a Laser Vibrometer. We compare the results of our motion estimation algorithm to that of a laser vibrometer, which measures velocity using Doppler shift (28). In the first experiment, a cantilevered beam was shaken by a mechanical shaker at 7.3 Hz, 58.3 Hz, 128 Hz, and 264 Hz, the measured modal frequencies of the beam. The relative amplitude of the shaking signal was varied between a factor of 5 and 25 in 2.5 increments. We simultaneously recorded a 2,000 FPS video of the beam with a high-speed camera (VisionResearch Phantom V-10) and measured its horizontal velocity with a laser vibrometer (Polytec PDV100). We repeated this experiment for nine different excitation magnitudes, three focal lengths (24 mm, 50 mm, 85 mm) and eight exposure times (12.5 μ s, 25 μ s, 50 μ s, 100 μ s, 200 μ s, 300 μ s, 400 μ s, 490 μ s), for a total of 20 high-speed videos. The beam had an accelerometer mounted on it (white object in Fig. 1A), but we did not use it in this experiment.

We used our motion estimation method to compute the horizontal displacement of the marked, red point on the left side of the accelerometer from the video (Fig. 1A). We applied a temporal band-stop filter to remove motions between 67 Hz and 80 Hz that corresponded to camera motions caused by its cooling fan's rotation. The laser vibrometer signal was integrated using discrete, trapezoidal integration. Before integration, both signals were high-passed above 2.5 Hz to reduce low-frequency noise in the integrated vibrometer signal. The motion signals from each video were manually aligned. For one video (exposure, 490 μ s; excitation, 25; and focal length, 85 mm), we plot the two motion signals (Fig. S4 B–D). They agree remarkably well, with higher modes well aligned and a correlation of 0.997.

To show the sensitivity of the motion microscope, we plot the correlation of our motion estimate and the integrated velocities from the laser vibrometer vs. motion size (RMS displacement). Because the motion's average size varies over time, we divide each video's motion signal into eight equal pieces and plot the correlations of each piece in each video in Fig. S4 E and F. For RMS displacements on the order of 1/100th of a pixel, the correlation between the two signals varies between 0.87 and 0.94. For motions larger than 1/20th of a pixel, the correlation is between 0.95 and 0.999. Possible sources of discrepancy are noise in the motion microscope signal, integrated low-frequency noise in the vibrometer signal, and slight misalignment between the signals. Displacements with RMS smaller than 1/100th of a pixel were noisier and had lower correlations, indicating that noise in the video prevents the two signals from matching.

As expected, correlation increases with focal length and excitation magnitude, two things that positively correlate with motion size (in pixels) (Fig. S4 G and H). The correlation also increases with exposure, because videos with lower exposure times are noisier (Fig. S4I).

Filming Bridge Sequence. The bridge was filmed with a monochrome Point Gray Grasshopper3 camera (model GS3-U3-2356M-C) at 30 FPS with a resolution of 800 \times 600. The central span of the bridge lifted to accommodate marine traffic. Filming was started about 5 s before the central span was lowered to its lowest point.

The accelerometer data were doubly integrated using trapezoidal integration to displacement. In Fig. 3 C and D, both the motion microscope displacement and the doubly integrated acceleration were band-passed with a first-order band-pass Butterworth filter with the specified parameters.

Motion Field Interpolation. In textureless regions, it may not be possible to estimate the motion at all, and, at one-dimensional structures like edges, the motion field will only be accurate in the direction perpendicular to the edge. These inaccuracies are reflected in the motion covariance matrix. We show how to interpolate the motion field from accurate regions to inaccurate regions, assuming that adjacent pixels have similar motions.

We minimize the following objective function:

$$\sum_x (\mathbf{V}_S(\mathbf{x}) - \mathbf{V}(\mathbf{x})) \Sigma_V^{-1}(\mathbf{x}) (\mathbf{V}_S(\mathbf{x}) - \mathbf{V}(\mathbf{x}))^T + \lambda_S \sum_{y \in \mathcal{N}(\mathbf{x})} (\mathbf{V}_S(\mathbf{x}) - \mathbf{V}_S(\mathbf{y})) (\mathbf{V}_S(\mathbf{x}) - \mathbf{V}_S(\mathbf{y}))^T \quad [6]$$

- Liu C, Torralba A, Freeman WT, Durand F, Adelson EH (2005) Motion magnification. *ACM Trans Graph* 24:519–526.
- Wu HY, et al. (2012) Eulerian video magnification for revealing subtle changes in the world. *ACM Trans Graph* 31:1–8.
- Wadhwa N, Rubinstein M, Durand F, Freeman WT (2013) Phase-based video motion processing. *ACM Trans Graph* 32:80.
- Wadhwa N, Rubinstein M, Durand F, Freeman WT (2014) Riesz pyramid for fast phase-based video magnification. *IEEE International Conference on Computational Photography* (Inst Electr Electron Eng, New York), pp 1–10.
- Nakamura J (2005) *Image Sensors and Signal Processing for Digital Still Cameras* (CRC, Boca Raton, FL).
- Simoncelli EP, Freeman WT (1995) The steerable pyramid: A flexible architecture for multi-scale derivative computation. *Int J Image Proc* 3:444–447.
- Fleet DJ, Jepson AD (1990) Computation of component image velocity from local phase information. *Int J Comput Vis* 5:77–104.
- Dallos P, Fay RR (2012) *The Cochlea*, Springer Handbook of Auditory Research (Springer Science, New York), Vol 8.
- Thalmann I (1993) Collagen of accessory structures of organ of Corti. *Connect Tissue Res* 29:191–201.
- Zwislocki JJ (1980) Five decades of research on cochlear mechanics. *J Acoust Soc Am* 67:1679–1685.

where \mathbf{V}_S is the desired interpolated field, \mathbf{V} is the estimated motion field, Σ_V is its covariance, $\mathcal{N}(\mathbf{x})$ is the four-pixel neighborhood of \mathbf{x} , and λ_S is a user-specified constant that specifies the relative importance of matching the estimated motion field vs. making adjacent pixels have similar motion fields. The first term seeks to ensure that \mathbf{V}_S is close to \mathbf{V} , weighted by the expected amount of noise at each pixel. The second term seeks to ensure that adjacent pixels have similar motion fields.

In Fig. 4D, we produce the color overlays by applying the above processing to the estimated motion field with $\lambda_S = 300$ and then taking the amplitude of each motion vector. We also set components of the covariance matrix that were larger than 0.1 square pixels to be an arbitrarily large number (we used 10,000 square pixels).

Finite Element Analysis of Acoustic Metamaterial. We use Abaqus/Standard (29), a commercial finite-element analyzer, to simulate the metamaterial's response to forcing. We constructed a 2D model with 37,660 nodes and 11,809 eight-node plane strain quadrilateral elements (Abaqus element type CPE8H). We modeled the rubber as Neo-Hookean, with shear modulus 443.4 kPa, bulk modulus 7.39×10^5 kPa, and density 1,050 kg·m³ (Abaqus parameters C10 = 221.7 kPa, D1 = 2.71×10^{-9} Pa⁻¹). We modeled the copper core with shear modulus 4.78×10^7 kPa, bulk modulus 1.33×10^8 kPa, and density 8,960 kg·m³ (Abaqus parameters C10 = 2.39×10^7 kPa, D1 = 1.5×10^{-11} Pa⁻¹). Geometry and material properties are specified in Wang et al. (24). The bottom of the metamaterial was given a zero-displacement boundary condition. A sinusoidal displacement loading condition at the forcing frequency was applied to a node located halfway between the top and bottom of the metamaterial.

Validation of Noise Analysis with Real Video Data. We took a video of an accelerometer attached to a beam (Fig. S9A). We used the accelerometer to verify that the beam had no motions between 600 Hz and 700 Hz (Fig. S9B). We then estimated the in-band motions from a video of the beam. Because the beam is stationary in this band, these motions are entirely due to noise, and their temporal sample covariance gives us a ground-truth measure of the noise level (Fig. S9C). We used our simulation with a signal-dependent noise model to estimate the covariance matrix from the first frame of the video, the specific parameters of which are shown in Fig. S9D. The resulting covariance matrices closely match the ground truth (Fig. S9 E and F), showing that our simulation can accurately estimate noise level and error bars.

We also verify that the signal-dependent noise model performs better than the simpler constant variance noise model, in which noise is IID. The result of the constant noise model simulation produced results that are much less accurate than the signal-dependent noise model (Fig. S9 G and H).

In Fig. S9, we only show the component of the covariance matrix corresponding to the direction of least variance, and only at points corresponding to edges or corners.

ACKNOWLEDGMENTS. We thank Professor Erin Bell and Travis Adams at University of New Hampshire and New Hampshire Department of Transportation for their assistance with filming the Portsmouth lift bridge. This work was supported, in part, by Shell Research, Quanta Computer, National Science Foundation Grants CGV-1111415 and CGV-1122374, and National Institutes of Health Grant R01-DC00238.

- Sellon JB, Farrahi S, Ghaffari R, Freeman DM (2015) Longitudinal spread of mechanical excitation through tectorial membrane traveling waves. *Proc Natl Acad Sci USA* 112:12968–12973.
- Ghaffari R, Aranyosi AJ, Freeman DM (2007) Longitudinally propagating traveling waves of the mammalian tectorial membrane. *Proc Natl Acad Sci USA* 104:16510–16515.
- Sellon JB, Ghaffari R, Farrahi S, Richardson GP, Freeman DM (2014) Porosity controls spread of excitation in tectorial membrane traveling waves. *J Biophys* 106:1406–1413.
- Ghaffari R, Aranyosi AJ, Richardson GP, Freeman DM (2010) Tectorial membrane traveling waves underlie abnormal hearing in tectb mutants. *Nat Commun* 1:96.
- Ewins DJ (1995) *Modal Testing: Theory and Practice*, Engineering Dynamics Series (Res Stud, Baldock, UK) Vol 6.
- Salawu O (1997) Detection of structural damage through changes in frequency: A review. *Eng Struct* 19:718–723.
- Hermans L, van der Auweraer H (1999) Modal testing and analysis of structures under operational conditions: Industrial applications. *Mech Sys Signal Process* 13:193–216.
- Hussein MI, Leamy MJ, Ruzzene M (2014) Dynamics of phononic materials and structures: Historical origins, recent progress, and future outlook. *Appl Mech Rev* 66:040802.
- Khelif A, Choujaa A, Benchabane S, Djafari-Rouhani B, Laude V (2004) Guiding and bending of acoustic waves in highly confined phononic crystal waveguides. *Appl Phys Lett* 84:4400–4402.

Supporting Information

Wadhwa et al. 10.1073/pnas.1703715114

Modal Shapes of a Pipe

We made a measurement of a pipe being struck by a hammer, viewed end on by a camera, to capture its radial–circumferential vibration modes. A standard 4" schedule 40 PVC pipe was recorded with a high-speed camera at 24,000 FPS, at a resolution of 192×192 (Fig. S10A). Fig. S10 B–F shows frames from the motion-magnified videos for different resonant frequencies showing the mode shapes, a comparison of the quantitatively measured mode shapes with the theoretically derived mode shapes, and the displacement vs. time of the specific frequency band and the estimated noise SD. The tiny modal motions are seen clearly. Obtaining vibration data with traditional sensors with the same spatial density would be extremely difficult, and accelerometers placed on the pipe would alter its resonant frequencies.

This sequence also demonstrates the accuracy of our noise analysis. The noise standard deviations show that the detected motions before impact, when the pipe is stationary, are likely spurious.

Synthetic Validation

We validate the accuracy of our motion estimation on a synthetic dataset and compare its accuracy to Ncorr, a digital image correlation technique (30) used by mechanical engineers (31). In this experiment, we did not use temporal filtering.

We created a synthetic dataset of frame pairs with known ground-truth motions between them. We took natural images from the frames of real videos (Fig. S5A) and warped them according to known motion fields using cubic b-spline interpolation (32). Sample motions fields, shown in Fig. S5B, were produced by Gaussian-blurring IID Gaussian random variables. We used Gaussian blurs with SDs, ranging from zero (no filtering) to infinite (a constant motion field). We also varied the RMS amplitude of the motion fields from 0.001 pixels to three pixels. For each set of motion field parameters, we sampled five different motion fields to produce a total of 155 motion fields with different amplitudes and spatial coherence. To test the accuracy of the algorithms rather than their sensitivity to noise, no noise was added to the image pairs.

We ran our motion estimation technique and Ncorr on each image pair. We then computed the mean absolute difference between the estimated and ground-truth motion fields. Then, for each set of motion field parameters, we averaged the mean absolute differences across image pairs and divided the result by the RMS motion amplitude to make the errors comparable over motion sizes. The result is the average relative error as a percentage of RMS motion amplitude (Fig. S5C).

Both Ncorr and our method perform best when the motions are spatially coherent (filter standard deviations greater than 10 pixels) with relative errors under 10%. This reflects the fact that both methods assume the motion field is spatially smooth. Across motion sizes, our method performs best for subpixel motions (5% relative error). This is probably because we assume that the motions are small when we linearize the phase constancy equation (Eq. S3). Ncorr has twice the relative error (10%) for the same motion fields.

The relative errors reported in Fig. S5C are computed over all pixels, including those that are in smooth, textureless regions where it is difficult to estimate the motions. If we restrict the error metric to only take into account pixels at edges and corners, the average relative errors for small (< 1 pixel RMS), spatially

coherent (filter SD > 10 pixels) motions drops by a factor of 2.5 for both methods.

We generated synthetic images that are slight translations of each other and added Gaussian noise to the frames (Fig. S8A). For each translation amount, we compute the motion between the two frames over 4,000 runs. We compute the sample covariance matrix over the runs as a measure of the ground-truth noise level. We also used our noise analysis to estimate the covariance matrix at the points denoted in red.

The off-diagonal term of the covariance matrix should be zero for the synthetic frames in Fig. S8A. For both examples, it is within 10^{-5} square pixels of zero for all translation amounts (Fig. S8B).

The relative errors of the horizontal and vertical variances vs. translation (Fig. S8 C and D) are less than 5% for subpixel motions. This is likely due to the random nature of the simulation. For motions greater than one pixel, the covariance matrix has a relative error of less than 25%.

Relation Between Local Phase Differences and Motions

Fleet and Jepson have shown that contours of constant phase in image subbands such as those in the complex steerable pyramid approximately track the motion of objects in a video (7). We make a similar phase constancy assumption, in which the following equation relates the phase of the frame at time 0 to the phase of future frames:

$$\phi_{r,\theta}(x, y, 0) = \phi_{r,\theta}(x - u(x, y, t), y - v(x, y, t), t), \quad [\text{S1}]$$

where $\mathbf{V}(x, y, t) := (u(x, y, t), v(x, y, t))$ is the motion we seek to compute. We Taylor-expand the right-hand side around (x, y) to get

$$\Delta\phi_{r,\theta} = \left(\frac{\partial\phi_{r,\theta}}{\partial x}, \frac{\partial\phi_{r,\theta}}{\partial y} \right) \cdot (u, v) + O(u^2, v^2), \quad [\text{S2}]$$

where $\Delta\phi_{r,\theta}(x, y, t) := \phi_{r,\theta}(x, y, t) - \phi_{r,\theta}(x, y, 0)$, arguments have been suppressed and $O(u^2, v^2)$ represents higher order terms in the Taylor expansion. Because we assume the motions are small, higher order terms are negligible and the local phase variations are approximately equal to only the linear term,

$$\Delta\phi_{r,\theta} = \left(\frac{\partial\phi_{r,\theta}}{\partial x}, \frac{\partial\phi_{r,\theta}}{\partial y} \right) \cdot (u, v). \quad [\text{S3}]$$

Fleet has shown that the spatial gradients of the local phase, $\left(\frac{\partial\phi_{r,\theta}}{\partial x}, \frac{\partial\phi_{r,\theta}}{\partial y} \right)$, are roughly constant within a subband and that they are approximately equal to the peak tuning frequency of the corresponding subband's filter (33). This frequency is a 2D vector oriented orthogonal to the direction the subband selects for, which means that the local phase changes only provide information about the motions perpendicular to this direction.

Low-Amplitude Coefficients Have Noisy Phase

Each frame of the input video $I(x, y, t)$ is transformed to the complex steerable pyramid representation by being spatially band-passed by a bank of quadrature pairs of filters $g_{r,\theta}$ and $h_{r,\theta}$, where r corresponds to different spatial scales of the pyramid and θ corresponds to different orientations. We use the filters of Portilla and Simoncelli, which are specified and applied in the frequency domain (25). For one such filter pair, the result is a set of complex coefficients $S_{r,\theta} + iT_{r,\theta}$ whose real and imaginary part are given by

$$S_{r,\theta} = g_{r,\theta} * I \text{ and } T_{r,\theta} = h_{r,\theta} * I, \quad [\text{S4}]$$

where the convolution is applied spatially at each time instant t . This filter pair is converted to amplitude $A_{r,\theta}$ and phase $\phi_{r,\theta}$ by the operations

$$A_{r,\theta} = \sqrt{S_{r,\theta}^2 + T_{r,\theta}^2} \text{ and } \phi_{r,\theta} = \tan^{-1}(T_{r,\theta}/S_{r,\theta}). \quad [\text{S5}]$$

Filters $g_{r,\theta}$ and $h_{r,\theta}$ are in quadrature relationship, which means that they select for the same frequencies, but are 90 degrees out of phase like \sin and \cos . A consequence is that they are uncorrelated and have equal RMS values. Complex coefficients at antipodal orientations are conjugate symmetric and contain redundant information. Therefore, we only use a half-circle of orientations.

Suppose the observed video $I(x, y, t)$ is contaminated with IID noise $I_n(x, y, t)$ of variance σ^2 ,

$$I(x, y, t) = I_0(x, y, t) + I_n(x, y, t), \quad [\text{S6}]$$

where $I_0(x, y, t)$ is the underlying noiseless video. This noise causes the complex steerable pyramid coefficients to be noisy, which causes the local phase to be noisy. We show that the local phase at a point has an approximate Gaussian distribution when the amplitude is high and is approximately uniformly distributed when the amplitude is low.

The transformed representation has response

$$g_{r,\theta} * I_0 + g_{r,\theta} * I_n \text{ and } h_{r,\theta} * I_0 + h_{r,\theta} * I_n. \quad [\text{S7}]$$

The first term in each expression is the noiseless filter response, which we denote $S_{0,r,\theta} = g_{r,\theta} * I_0$ for the real part and $T_{0,r,\theta} = h_{r,\theta} * I_0$ for the imaginary part. The second term in each expression is filtered noise, which we denote as $S_{n,r,\theta}$ and $T_{n,r,\theta}$. At a single point, $S_{n,r,\theta}$ and $T_{n,r,\theta}$ are Gaussian random variables with covariance matrix equal to

$$\sigma^2 \begin{pmatrix} \sum_{x,y} g_{r,\theta}(x, y)^2 & \sum_{x,y} g_{r,\theta}(x, y)h_{r,\theta}(x, y) \\ \sum_{x,y} g_{r,\theta}(x, y)h_{r,\theta}(x, y) & \sum_{x,y} h_{r,\theta}(x, y)^2 \end{pmatrix} \\ = \sigma^2 \sum_{x,y} g_{r,\theta}(x, y)^2 I, \quad [\text{S8}]$$

where I is the identity matrix and equality follows from the fact that $g_{r,\theta}$ and $h_{r,\theta}$ are quadrature pairs.

We suppress the indices r, θ in this section for readability. From Eq. S5, the noiseless and noisy phases are given by

$$\phi_0 = \tan^{-1}(T_0/S_0) \text{ and } \phi = \tan^{-1}((T_0 + T_n)/(S_0 + S_n)). \quad [\text{S9}]$$

Their difference linearized around (S_0, T_0) is

$$\tan^{-1}\left(\frac{T_0 + T_n}{S_0 + S_n}\right) - \tan^{-1}\left(\frac{T_0}{S_0}\right) \\ = \frac{S_n S_0 - T_n T_0}{A_0^2} + O\left(\frac{S_n^2, S_n T_n, T_n^2}{A_0^4}\right). \quad [\text{S10}]$$

The terms S_n^2 and T_n^2 are expected to be equal to their variance $\sigma^2 \sum g_{r,\theta}(x, y)^2$. Therefore, if $A_0^2 \gg \sigma^2 \sum g_{r,\theta}(x, y)^2$, higher order terms are negligible. In this case, we see that the phase is approximately a linear combination of Gaussian random variables and is therefore Gaussian. This is illustrated empirically by local phase histograms of the green and blue points in Fig. S3A–E.

For these high-amplitude points, we compute the variance of the phase of a coefficient,

$$E \left[\left(\tan^{-1}\left(\frac{T_0 + T_n}{S_0 + S_n}\right) - \tan^{-1}\left(\frac{T_0}{S_0}\right) \right)^2 \right] \quad [\text{S11}]$$

$$\approx E \left[\left(\frac{T_0 S_n - S_0 T_n}{A_0^2} \right)^2 \right] \quad [\text{S12}]$$

$$= E \left[\frac{T_0^2 S_n^2 - 2 T_0 S_0 S_n T_n + S_0^2 T_n^2}{A_0^4} \right] \quad [\text{S13}]$$

$$= \frac{\sigma^2 \sum g_{r,\theta}^2 (T_0^2 + S_0^2)}{A_0^4} \quad [\text{S14}]$$

$$= \frac{\sigma^2 \sum g_{r,\theta}^2}{A_0^2}. \quad [\text{S15}]$$

The first approximation follows from the linearization of Eq. S10.

When the amplitude is low compared with the noise level ($A_0^2 \ll \sigma^2 \sum g_{r,\theta}(x, y)^2$), the linearization of Eq. S10 is not accurate. In this case, $S_0 \approx 0$ and $T_0 \approx 0$, and phase is given by

$$\tan^{-1}\left(\frac{T_n}{S_n}\right). \quad [\text{S16}]$$

T_n and S_n are uncorrelated Gaussian random variables with equal variance, which means that the phase is a uniformly random number. The phase at such points contains no information and intuitively corresponds to places where there is no image content in a given pyramid level (Fig. S3E, red point).

Noise Model and Creating Synthetic Video

We adopt a signal-dependent noise model, in which each pixel is contaminated with spatially independent Gaussian noise with variance $f(I)$, where I is the pixel's mean intensity (26, 34). Liu et al. (26) refer to this function f as a “noise level function,” and we do the same. This reflects that sensor noise is well modeled by the sum of zero-mean Gaussian noise sources, some of which have variances that depend on intensity (5). We show that this noise model is an improvement over a constant variance noise model in Fig. S9.

The noise level function f is estimated from temporal variations in the input video, with observed intensities $I(x, y, t)$. Assuming that I is the sum of noiseless intensity I_0 and a zero-mean Gaussian noise term I_n with variance $f(I_0)$, the temporal variations are given by the following Taylor expansion:

$$I(x, y, t) = I_0(x, y, t) + I_n(x, y, t) \quad [\text{S17}]$$

$$= I_0(x - u(x, y, t), y - v(x, y, t), 0) + I_n(x, y, t) \quad [\text{S18}]$$

$$\approx I_0(x, y, 0) - \frac{\partial I_0}{\partial x} u(x, y, t) - \frac{\partial I_0}{\partial y} v(x, y, t) \\ + I_n(x, y, t). \quad [\text{S19}]$$

The second equality is the brightness constancy assumption of optical flow (35, 36). We exclude pixels where the spatial gradient $\left(\frac{\partial I_0}{\partial x}, \frac{\partial I_0}{\partial y}\right)$ has high magnitude from our analysis. At the remaining pixel, temporal variations in I are mostly due to noise,

$$I(x, y, t) \approx I_0(x, y, 0) + I_n(x, y, t). \quad [\text{S20}]$$

At these pixels, we take the temporal variance and mean of I , which, in expectation, are $f(I_0)$ and I_0 , respectively. To increase robustness, we divide the intensity range into 64 equally sized bins. For each bin, we take all those pixels with mean inside that bin and take the mean of the corresponding temporal variances of I to estimate the noise level function f .

With f in hand, we can take frames from existing videos and use them to create simulated videos with realistic noise, but with known, zero motion. In Fig. S6A, we take a frame $I_0(x, y, 0)$ from a video of the metamaterial, filmed with a Phantom V-10, and add noise to it via the equation

$$I_S(x, y, t) = I_0(x, y, 0) + I_n(x, y, t) \sqrt{f(I_0(x, y, 0))}, \quad [\text{S21}]$$

where I_n now is Gaussian noise with unit variance. We motion-magnify the resulting video 600 times in a 20-Hz band centered

at 50 Hz to show that motion-magnified noise can cause spurious motions (Fig. S6B).

We use the same simulation to create synthetic videos with which to estimate the covariance matrix of the motion vectors.

We quantify the noise in the motion vectors by estimating their covariance matrices $\Sigma_V(x, y)$. These matrices reflect variations in the motion caused by noise. It is not usually possible to directly estimate them from the input video, because both motions and noise vary across frames and the true motions are unknown. Therefore, we create a noisy, synthetic video $I_S(x, y, t)$ with known zero true motion (Eq. S21 and Fig. S7A–C).

We estimate the motions in I_S (Fig. S7D) using our technique with spatial smoothing, but without temporal filtering, which we handle in a later step. This results in a set of 2D motion vectors $\mathbf{V}_S(x, y, t)$, in which all temporal variations in \mathbf{V}_S are due to noise. The sample covariance matrix over the time dimension is

$$\Sigma_V = \frac{1}{N-1} \sum_t (\mathbf{V}_S(x, y, t) - \bar{\mathbf{V}}_S(x, y)) (\mathbf{V}_S(x, y, t) - \bar{\mathbf{V}}_S(x, y))^T, \quad [\text{S22}]$$

where $\bar{\mathbf{V}}_S(x, y)$ is the mean over t of the motion vectors. Σ_V is a 2×2 symmetric matrix, defined at every pixel, with only three unique components. In Fig. S7E, we show these components, the variances of the horizontal and vertical components of the motion, and their covariance.

The motion \mathbf{V} projected onto a direction vector $\mathbf{d}_\theta := (\cos(\theta), \sin(\theta))$ is $\mathbf{V} \cdot \mathbf{d}_\theta$ and has variance $\sigma_V^2(\theta) = \mathbf{d}_\theta^T \Sigma_V \mathbf{d}_\theta$. Of particular interest is the direction θ of least variance that minimizes $\sigma_V^2(\theta)$. In the case of an edge in the image, the direction of least variance is usually normal to the edge.

Analytic Justification of Noise Analysis

We analyze only the case when the amplitudes at a pixel in all subbands are large ($A_{r,\theta} \gg \sigma^2 \sum g^2$), because the local phases have a Gaussian distribution in this case. Such points intuitively correspond to places where there is image content in at least two directions. In this case, we show that the sample covariance matrix computed using a simulated video with no motions is accurate for videos with subpixel small motions.

We reproduce the linearization of phase constancy equation (Eq. S3) with noise terms added to the phase variations (n_t) and phase gradient (n_x, n_y),

$$\Delta\phi_{r,\theta} + n_t = (u, v) \cdot \left(\frac{\partial\phi_{r,\theta}}{\partial x} + n_x, \frac{\partial\phi_{r,\theta}}{\partial y} + n_y \right). \quad [\text{S23}]$$

The total noise term in this equation is $n_t + un_x + vn_y$. The noise terms n_t , n_x and n_y are of the same order of magnitude. Since u and v are much less than 1 pixel, the predominant source of noise is from n_t , the effects of n_x and n_y are negligible, and we can ignore them, allowing us to write the noisy version of the equation as

$$\Delta\phi_{r,\theta} + n_t = (u, v) \cdot \left(\frac{\partial\phi_{r,\theta}}{\partial x}, \frac{\partial\phi_{r,\theta}}{\partial y} \right). \quad [\text{S24}]$$

The motion estimate \mathbf{V} is the solution to a weighted least squares problem, $\mathbf{V} = (\mathbf{X}^T \mathbf{W} \mathbf{X})^{-1} \mathbf{X}^T \mathbf{W} \mathbf{Y}$. To simplify notation, let $\mathbf{B} = (\mathbf{X}^T \mathbf{W} \mathbf{X})^{-1} \mathbf{X}^T \mathbf{W}$, the parts of the equation that don't depend on time. Then, the flow estimate is

$$\mathbf{V} = \mathbf{B} \mathbf{Y}. \quad [\text{S25}]$$

where the elements of \mathbf{Y} are the local phase variations over time. \mathbf{X} and \mathbf{W} contain the spatial gradients of phase and amplitude, respectively. We have demonstrated that \mathbf{X} is close to noiseless

(Eq. S24), and our assumption about the amplitudes being large means \mathbf{W} is also approximately noiseless, which means that \mathbf{B} is noiseless.

We split \mathbf{Y} into the sum of its mean \mathbf{Y}_0 and variance, a multivariate Gaussian random variable, denoted as \mathbf{Y}_n , that has zero mean and variance that depends only on image noise and local image content. Then, the flow estimate is

$$\mathbf{V} = \underbrace{\mathbf{B} \mathbf{Y}_0}_{\text{True flow}} + \underbrace{\mathbf{B} \mathbf{Y}_n}_{\text{Noise Term (Covariance Matrix)}}. \quad [\text{S26}]$$

The noise term doesn't depend on the value of the true flow $\mathbf{B} \mathbf{Y}_0$. Therefore, the estimated covariance matrix is valid even when the motions are nonzero but small.

From Eq. S15, we know that \mathbf{Y}_n is proportional to the amount of sensor noise. Therefore, an interesting consequence of Eq. S26 is that the motion covariance matrix will be linearly proportional to the variance of the sensor noise.

Using the Motion Microscope and Limitations

The motion microscope works best when the camera is stable, such as when mounted on a tripod. If the camera is handheld or the tripod is not sturdy, the motion microscope may only detect camera motions instead of subject motions. The easiest way to solve this problem is to stabilize the camera. If the camera motions are caused by some periodic source, such as camera fans, that occur at a temporal frequency that the subject is not moving at, they can be removed with a band-stop filter.

If adjacent individual pixels have different motions, it is unlikely that the motion microscope will be able to discern or differentiate them. The motion microscope assumes motions are locally constant, and it won't be able to properly process videos with high spatial frequency motions. The best way to solve this problem is to increase the resolving power or optical zoom of the imaging system. If noise is not a concern, reducing the width of the spatial Gaussian used to smooth the motions may also help.

It is not possible for the motion microscope to quantify the movement of textureless regions or the full 2D movement of 1D edges. However, the motion covariance matrices produced by the motion microscope will alert the user as to when this is happening. The visualization will still be reasonable, as it is not possible to tell if a textureless region was motion-amplified in the wrong way.

Saturated, pure white pixels may pose a problem for the motion microscope. Slight color changes are the underlying signal used to detect tiny motions. For pixels with clipped intensities, this signal will be missing, and it may not be possible to detect motion. The best ways to mitigate this problem are to reduce the exposure time, change the position of lights, or use a camera with a higher dynamic range.

The motion microscope only quantifies lateral, 2D motions in the imaging plane. Motions of the subject toward or away from the camera will not be properly reported. For a reasonable camera, these 2D motions can be multiplied by a constant to convert from pixels to millimeters or other units of interest. However, if the lens has severe distortion, the constant will vary across the scene. If quantification of the motions is important, the images will need to be lens distortion-corrected before the motion microscope is applied.

The motion microscope visualization is unable to push pixels beyond the support of the basis functions of the steerable pyramid. In practice, this results in a maximum amplified motion of around four pixels. If the amplification is high enough to push an image feature past this, ringing artifacts may occur.

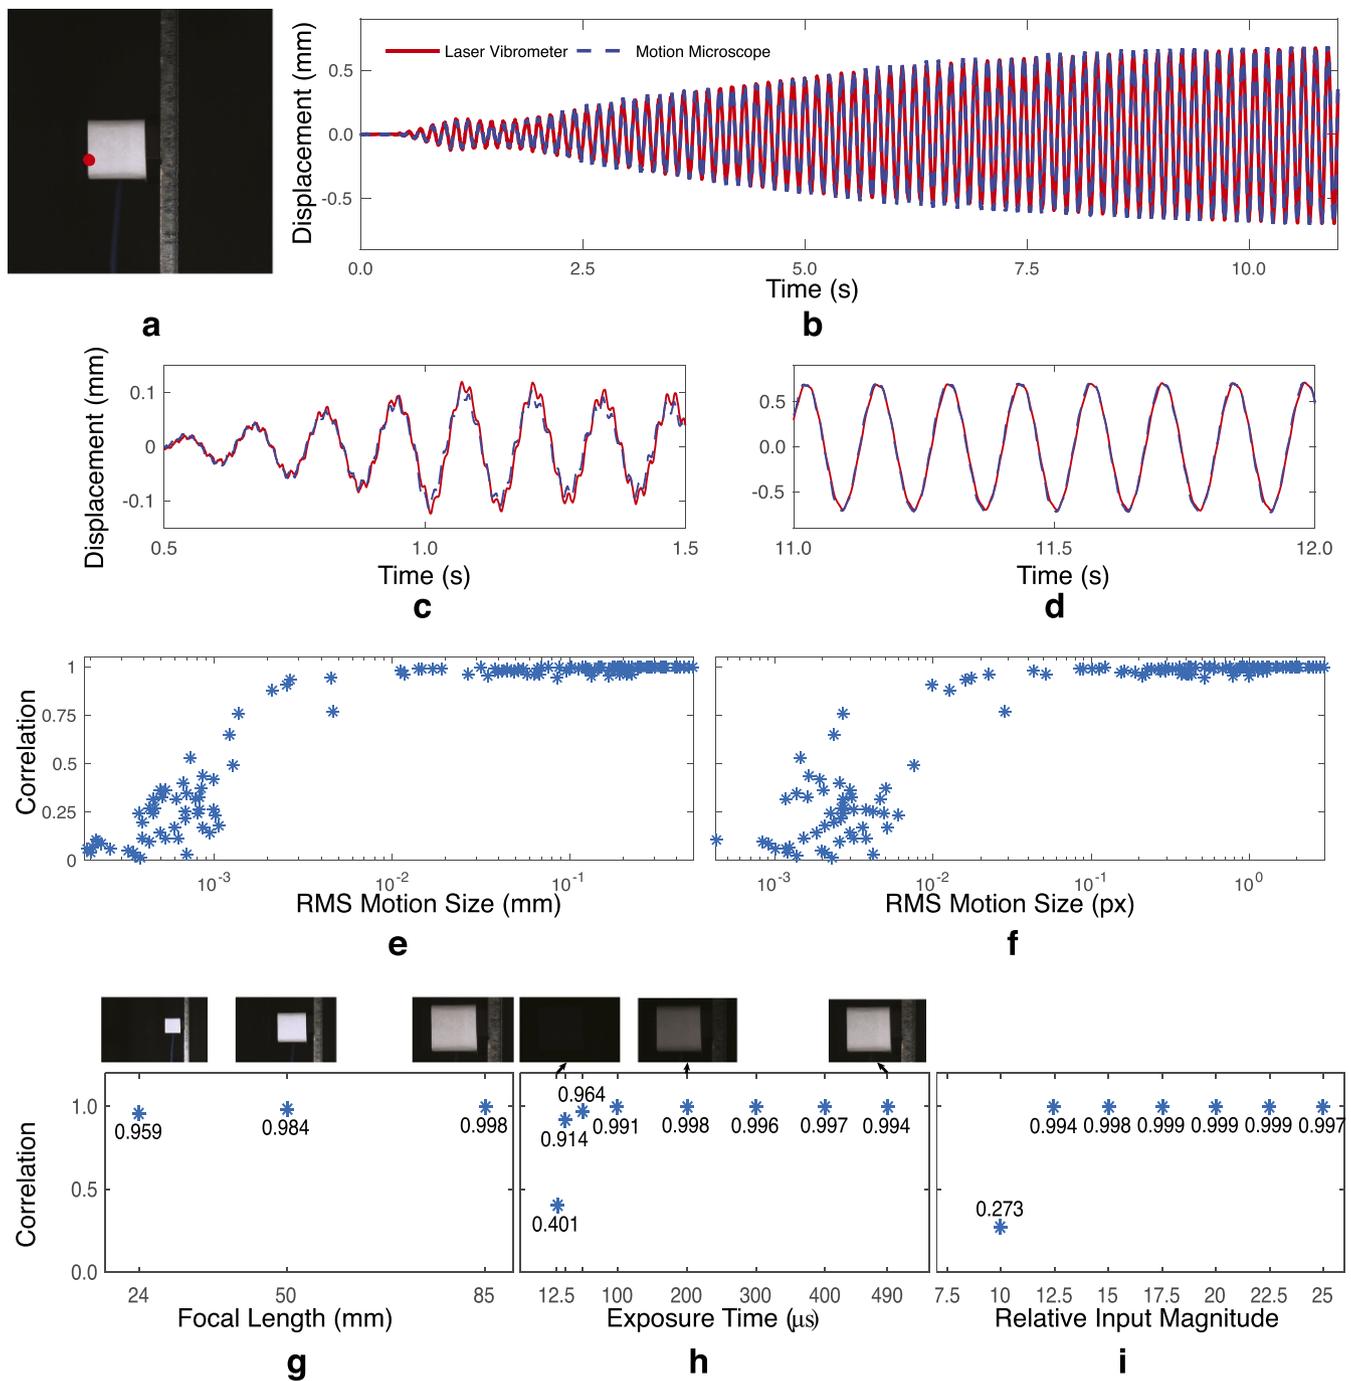


Fig. 54. A comparison of our quantitative motion estimation vs. a laser vibrometer. Several videos of a cantilevered beam excited by a shaker were taken, with varying focal length, exposure times, and excitation magnitude. (A) A frame from one video. (B) Motions from the motion microscope at the red point are compared with the integrated velocities from a laser vibrometer. (C) B from 0.5 s to 1.5 s. (D) B from 11 s to 12 s. (E) The correlation between the two signals across the videos vs. the RMS motion size in pixels, measured by the laser vibrometer. (F) The correlation between the two signals across the videos vs. the RMS motion size in pixels measured by the laser vibrometer. (G) The correlation between the signals vs. focal length (exposure time, 490 μ s; excitation magnitude, 15). (H) Correlation vs. exposure time (focal length, 85 mm; excitation magnitude, 15). Cropped frames from the corresponding videos are shown above. (I) Correlation vs. relative excitation magnitude (focal length, 85 mm; exposure time, 490 μ s). Only motions at the red point in A were used in our analysis.

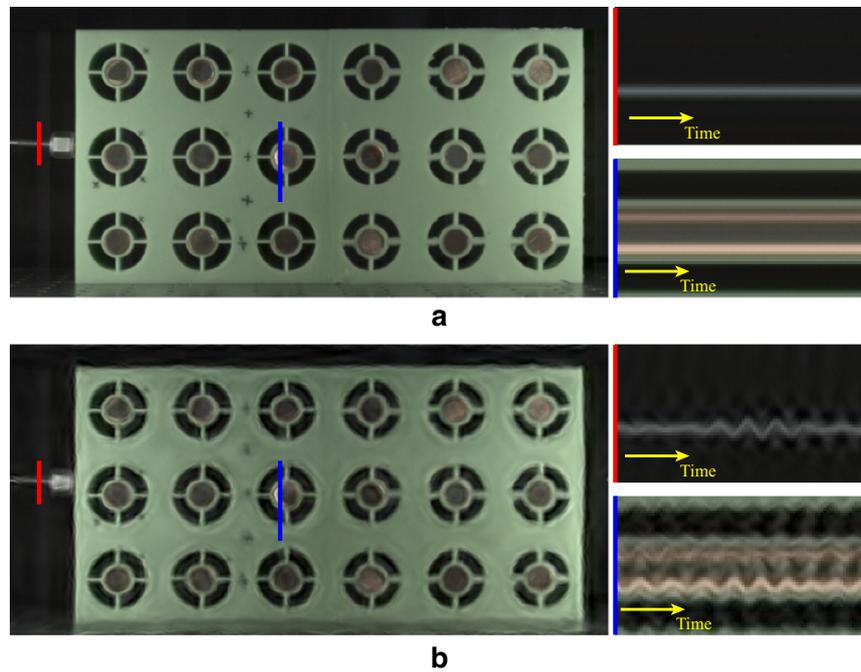


Fig. 56. Magnification of a spatial smooth and temporally filtered noise can look like a real signal. (A) Frames and time slices from a synthetic 300-frame video created by replicating a single frame 300 times and adding a different realistic noise pattern to each frame. (B) Corresponding frame and time slices from the synthetic video motion magnified $600\times$ in a temporal band of 40 Hz to 60 Hz. Time slices from the same parts of each video are shown on the right for comparison.

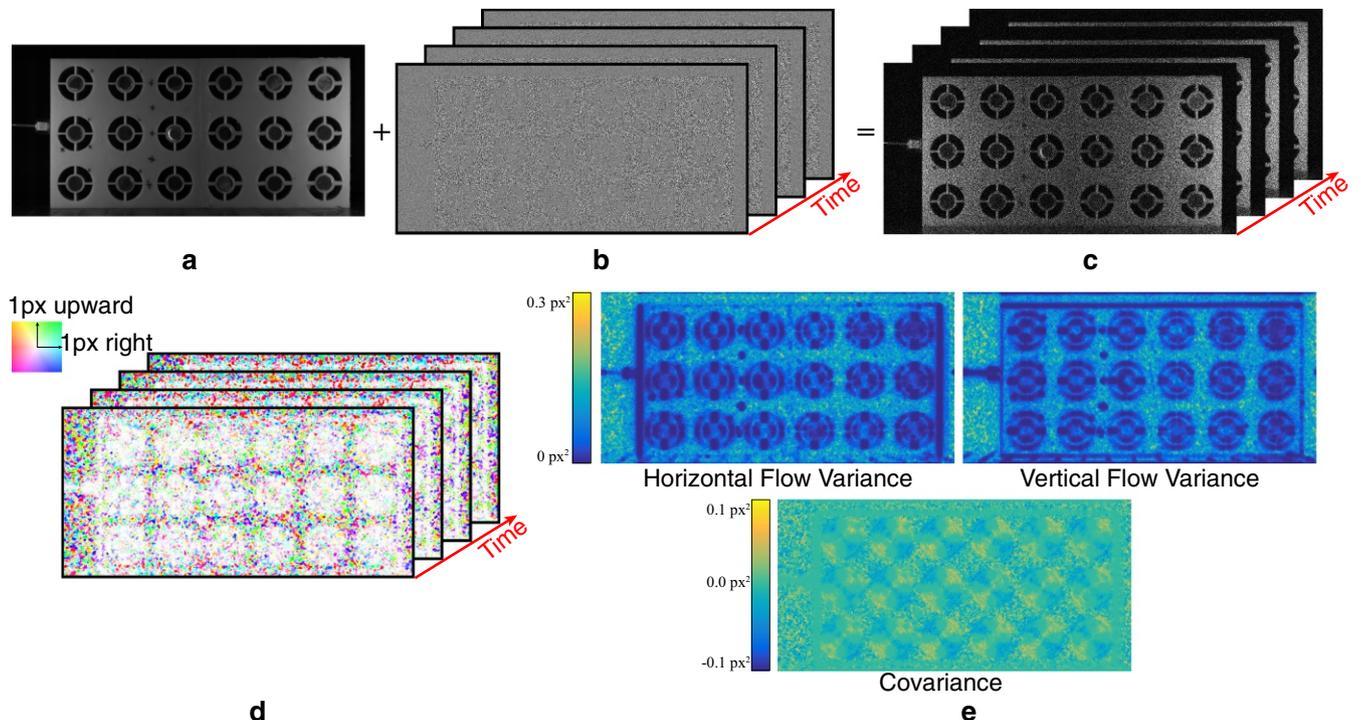


Fig. 57. Using a probabilistic simulation to compute the noise covariance of the motion estimate. (A) A single frame from an input video, in this case, of an elastic metamaterial. (B) Simulated, but realistic noise (contrast enhanced $80\times$). (C) Synthetic video with no motions, consisting of the input frame replicated plus simulated noise (noise contrast-enhanced $80\times$). (D) Estimated motions of this video. (E) Sample variances and sample covariance of the vertical and horizontal components of the motion are computed to give an estimate of how much noise is in the motion estimate.

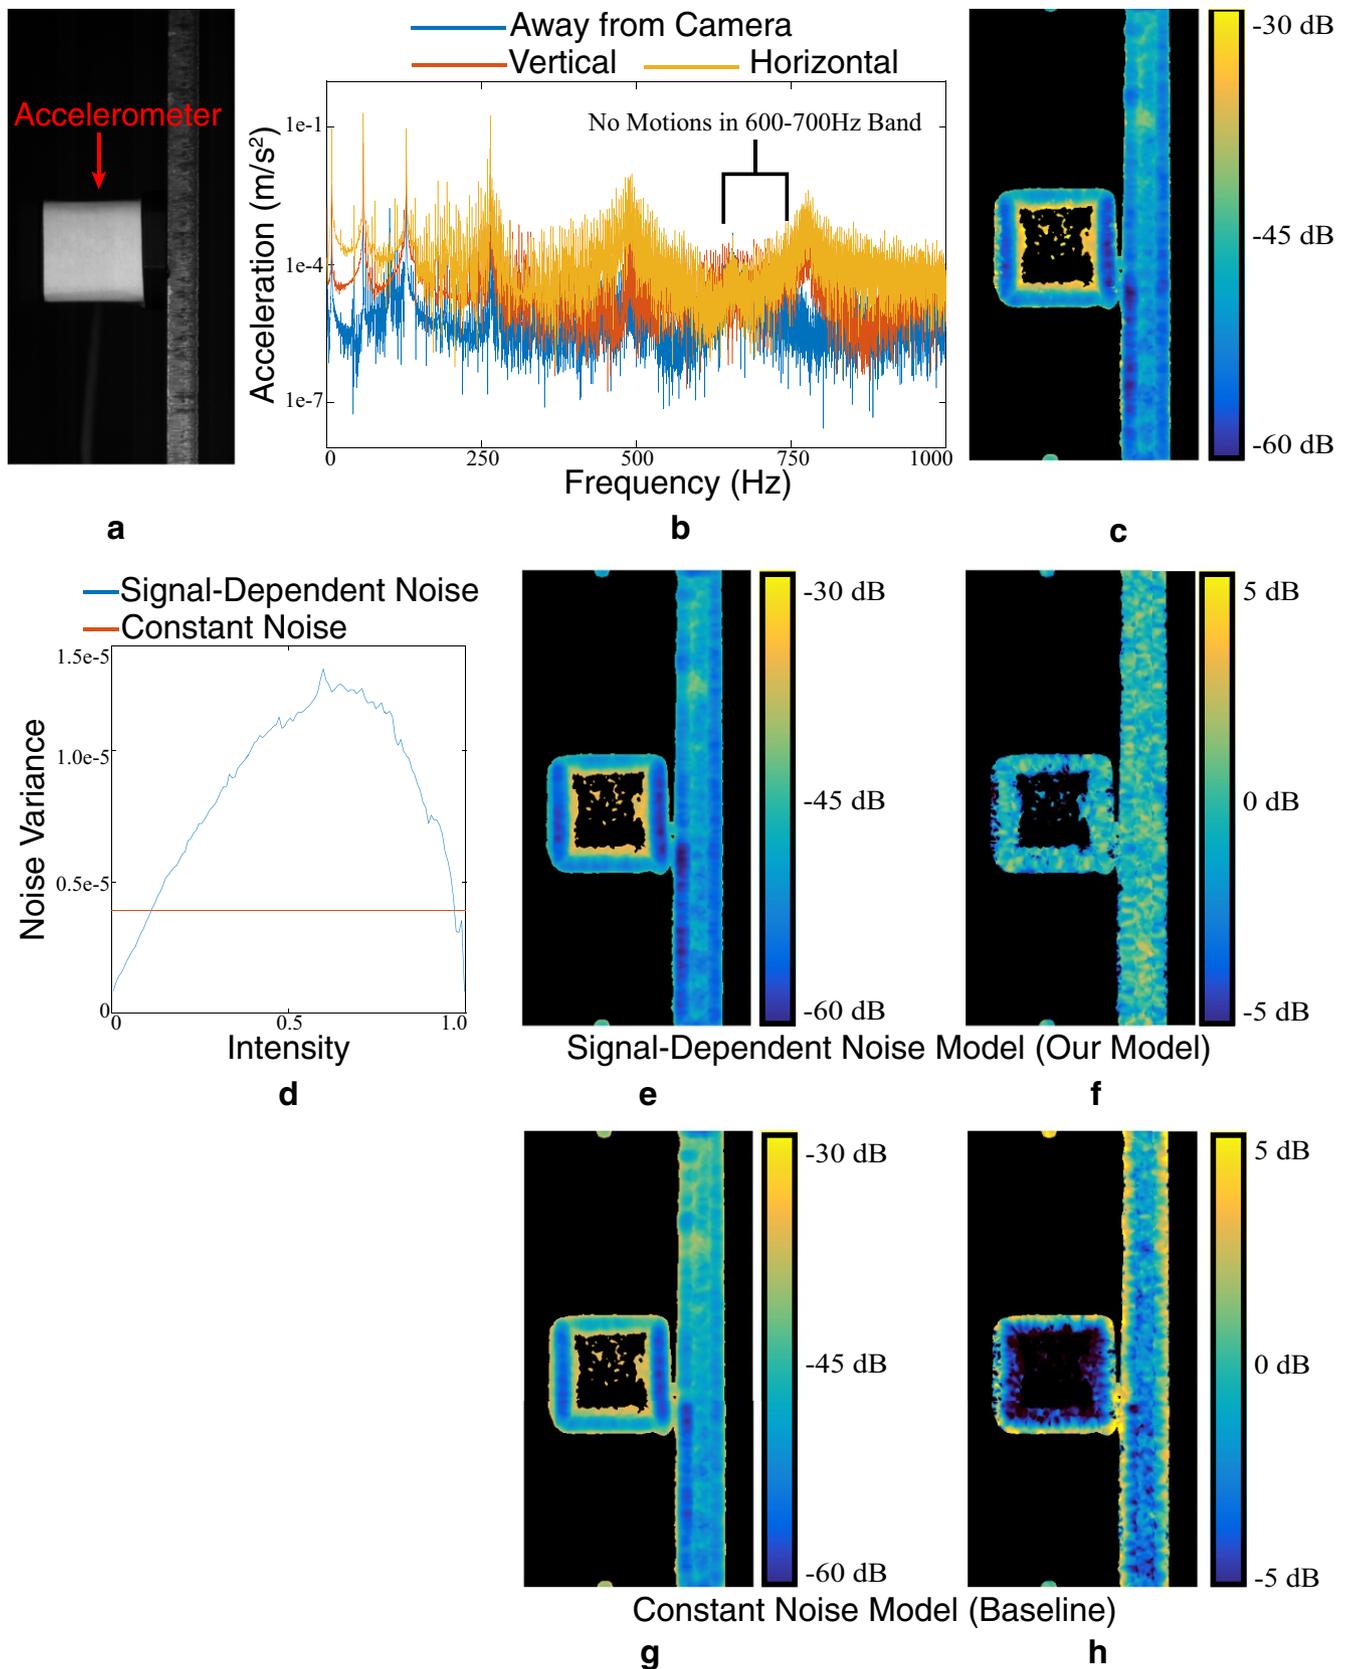
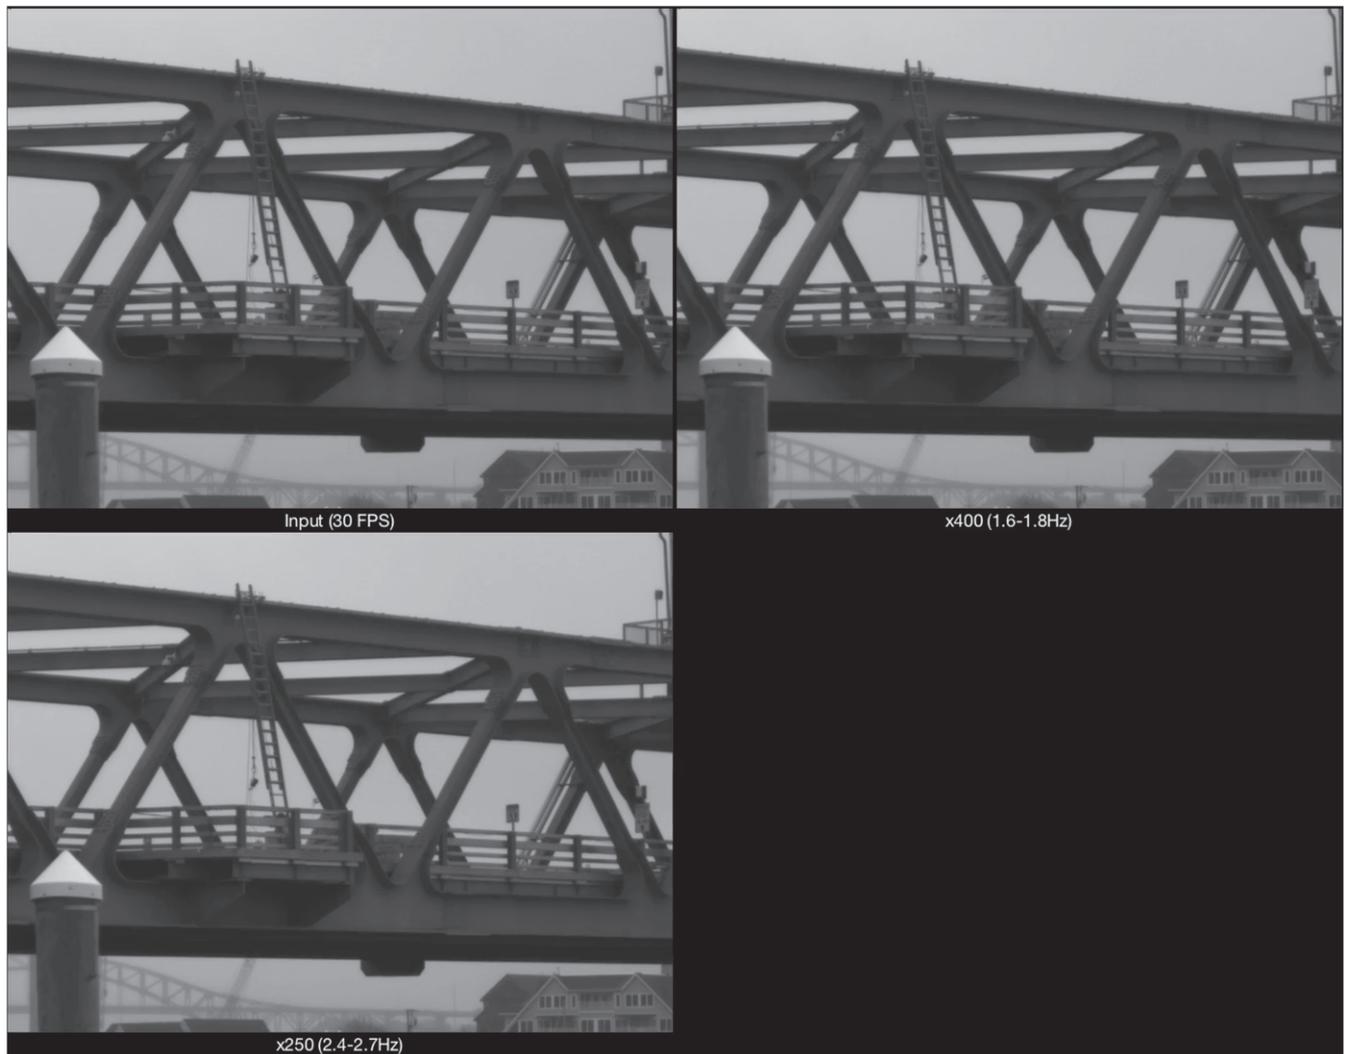


Fig. S9. Validation of our noise estimation on real data. (A) A frame from a real video of an accelerometer attached to a beam. (B) The accelerometer shows there are no motions in the frequency band 600 Hz to 700 Hz. (C) The variance of our motion estimate in the 600- to 700-Hz band serves as a ground-truth measure of noise, as there are no motions. (D) The estimated noise level vs. intensity for a signal-dependent noise model and a constant noise model. (E) The noise estimate produced by our Monte Carlo simulation with a signal-dependent model. All variances are of the motions projected onto the direction of least variance. Textureless regions, where the motion estimation is not meaningful, have been masked out in black. (F) Difference in decibels between ground truth and E. (G) Noise estimate produced by the Monte Carlo simulation with a constant noise model. (H) Difference in decibels between ground truth and G.



Movie S2. The input bridge video is concatenated with two motion-magnified videos revealing different modal shapes of the bridge. Motions within a 1.6- to 1.8-Hz frequency band are amplified 400 times to produce the video on the right, in which the first bending mode is revealed. Motions within a 2.4- to 2.7-Hz frequency band are amplified 250 times to produce the video on the bottom, in which the first torsional mode is revealed. The impact of the central span (not shown in the video) occurs ~ 5 s after the video's start.

[Movie S2](#)

