



# Motion microscopy for visualizing and quantifying small motions

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

<sup>a</sup>Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139; <sup>b</sup>Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139; <sup>c</sup>Harvard-MIT Program in Health Sciences and Technology, Cambridge, MA 02139; <sup>d</sup>Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139; <sup>e</sup>Google Research, Google Inc. Cambridge, MA 02139; <sup>f</sup>Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139; <sup>g</sup>School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138; <sup>h</sup>Department of Mechanical Engineering, Johns Hopkins University, Baltimore, MD 21218; and <sup>i</sup>Hopkins 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

**M**otion 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  $\theta$ . 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.

<sup>1</sup>Present address: Google Research, Google Inc. Mountain View, CA 94043.

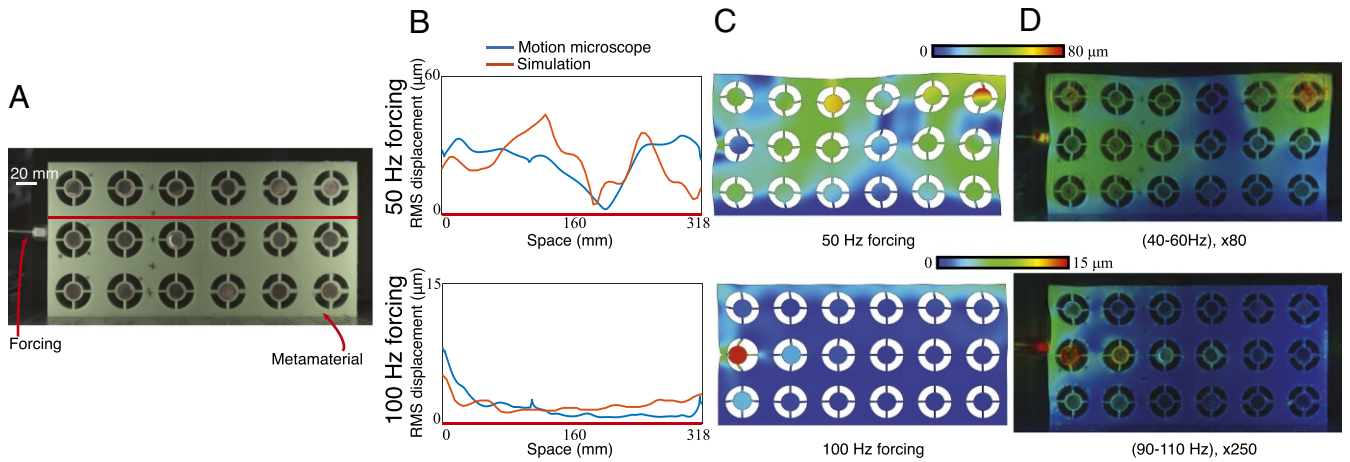
<sup>2</sup>To whom correspondence should be addressed. Email: billf@mit.edu.

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









**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  $A_{r_i, \theta_i}^2$ .

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  $19 \times 19$  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  $\sigma^2$ , 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  $\mu$ s, 25  $\mu$ s, 50  $\mu$ s, 100  $\mu$ s, 200  $\mu$ s, 300  $\mu$ s, 400  $\mu$ s, 490  $\mu$ 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  $\mu$ 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:

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

where  $\mathbf{V}_S$  is the desired interpolated field,  $\mathbf{V}$  is the estimated motion field,  $\Sigma_V$  is its covariance,  $\mathcal{N}(\mathbf{x})$  is the four-pixel neighborhood of  $\mathbf{x}$ , and  $\lambda_S$  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  $7.39 \times 10^5$  kPa, and density 1,050 kg·m<sup>3</sup> (Abaqus parameters C10 = 221.7 kPa, D1 =  $2.71 \times 10^{-9}$  Pa<sup>-1</sup>). We modeled the copper core with shear modulus  $4.78 \times 10^7$  kPa, bulk modulus  $1.33 \times 10^8$  kPa, and density 8,960 kg·m<sup>3</sup> (Abaqus parameters C10 =  $2.39 \times 10^7$  kPa, D1 =  $1.5 \times 10^{-11}$  Pa<sup>-1</sup>). 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.

- 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.
- 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.

