

Bayesian inference of elastic properties with resonant ultrasound spectroscopy

Ben Bales, Linda Petzold, Brent R. Goodlet, William C. Lenthe, and Tresa M. Pollock

Citation: *The Journal of the Acoustical Society of America* **143**, 71 (2018);

View online: <https://doi.org/10.1121/1.5017840>

View Table of Contents: <http://asa.scitation.org/toc/jas/143/1>

Published by the *Acoustical Society of America*

Articles you may be interested in

[Room acoustic modal analysis using Bayesian inference](#)

The Journal of the Acoustical Society of America **141**, 4480 (2017); 10.1121/1.4983301

[A model for acoustic vaporization dynamics of a bubble/droplet system encapsulated within a hyperelastic shell](#)

The Journal of the Acoustical Society of America **143**, 23 (2018); 10.1121/1.5019467

[Effects of background noises on nonlinear dynamics of a modelled thermoacoustic combustor](#)

The Journal of the Acoustical Society of America **143**, 60 (2018); 10.1121/1.5020059

[Time reversal focusing of elastic waves in plates for an educational demonstration](#)

The Journal of the Acoustical Society of America **141**, 1084 (2017); 10.1121/1.4976070

[No correlation between headphone frequency response and retail price](#)

The Journal of the Acoustical Society of America **141**, EL526 (2017); 10.1121/1.4984044

[Paradigmatic variation of vowels in expressive speech: Acoustic description and dimensional analysis](#)

The Journal of the Acoustical Society of America **143**, 109 (2018); 10.1121/1.5018433

Bayesian inference of elastic properties with resonant ultrasound spectroscopy

Ben Bales and Linda Petzold

Department of Mechanical Engineering, University of California, Santa Barbara, California 93106-5070, USA

Brent R. Goodlet,^{a)} William C. Lenthe, and Tresa M. Pollock

Materials Department, University of California, Santa Barbara, California 93106-5050, USA

(Received 23 August 2017; revised 17 November 2017; accepted 27 November 2017; published online 8 January 2018)

Bayesian modeling and Hamiltonian Monte Carlo (HMC) are utilized to formulate a robust algorithm capable of simultaneously estimating anisotropic elastic properties and crystallographic orientation of a specimen from a list of measured resonance frequencies collected via Resonance Ultrasound Spectroscopy (RUS). Unlike typical optimization procedures which yield point estimates of the unknown parameters, computing a Bayesian posterior yields probability distributions for the unknown parameters, and HMC is an efficient way to compute this posterior. The algorithms described are demonstrated on RUS data collected from two parallelepiped specimens of structural metal alloys. First, the elastic constants for a specimen of fine-grain polycrystalline Ti-6Al-4V with random crystallographic texture and isotropic elastic symmetry are estimated. Second, the elastic constants and crystallographic orientation for a single crystal Ni-based superalloy CMSX-4 specimen are accurately determined, using only measurements of the specimen geometry, mass, and resonance frequencies. The unique contributions of this paper are as follows: the application of HMC for sampling the Bayesian posterior of a probabilistic RUS model, and the procedure for simultaneous estimation of elastic constants and lattice-specimen misorientation. Compared to previous approaches these algorithms demonstrate superior convergence behavior, particularly when the initial parameterization is unknown, and enable substantially simplified experimental procedures. © 2018 Acoustical Society of America. <https://doi.org/10.1121/1.5017840>

[VMK]

Pages: 71–83

I. INTRODUCTION

Ultrasonic techniques such as Resonant Ultrasound Spectroscopy (RUS) provide the most accurate characterization of elastic properties,^{1,2} as well as superior precision and repeatability compared to static methods.³ Pulse-echo ultrasonic methods assess characteristic elastic wave propagation velocities of a material via time-of-flight measurements, and with a plane-wave assumption provide simply-defined relationships for elastically isotropic media.^{1,4} However, when the material is elastically anisotropic, pulse-echo experimental methods are complicated and often require multiple specimens, with parallel faces coincident to planes of elastic symmetry, and numerous independent velocity measurements along particular crystallographic directions.^{1,4} RUS methods do not require a plane-wave assumption,¹ nor do they require alternative experimental procedures when characterizing elastically anisotropic materials.

Modern experimental procedures for RUS are discussed in great detail by Migliori *et al.*,^{2,5,6} but generally involve excitation of a specimen with vibrations from a piezoelectric element that is in physical contact with the specimen. When the drive frequency of the piezoelectric element approaches a natural vibrational mode frequency of the specimen, a

resonance condition develops from constructive interference of opposite-traveling elastic waves to generate a standing wave throughout the specimen.² This standing wave leads to deflections of the specimen surface that are magnified in amplitude, potentially thousands of times greater than the drive amplitude, and are easily recorded by a second contacting piezoelectric element.^{1,2} When collecting a broadband RUS measurement the lowest-frequency mode is first identified, then the specimen is excited through a continuous range of greater frequencies until a desired quantity of modes are collected. These characteristic resonance frequencies are then provided to an inversion algorithm for estimating elastic properties. Today, the greatest impediments to broad application of resonance methods for elastic property evaluation are not empirical, but computational in nature.

Beyond inverting elastic properties with RUS data, crystallographic orientation can also be determined, as briefly detailed by Sarrao *et al.*⁷ The techniques developed here incorporate the ability to simultaneously estimate elastic properties and crystal orientations when the crystal reference frame is misaligned with the specimen reference frame. Simply machining a single crystal specimen along the crystal growth direction is insufficient, as only a few degrees of misalignment can lead to unacceptably large uncertainty in modulus estimates. X-ray diffraction (XRD) methods are most often employed to measure the crystallographic orientation of a RUS specimen,^{2,7} but this adds considerable complexity

^{a)}Electronic mail: bgoodlet@engineering.ucsb.edu

and cost to experimental methods as well as an additional source of measurement error. Sarrao *et al.*⁷ were first to report inverting crystal orientation and elastic moduli simultaneously, and rightly note the added value provided by such a capability. However, little guidance is offered toward reproducing their results beyond instructing the reader to proceed with “proper caution” by performing the inversion with “slower convergence steps.” The present research details an alternative and robust approach for simultaneously estimating crystal orientation and elastic properties.

A. Computational considerations for inversion

While straightforward theoretically, there is considerable difficulty in the practical implementation of a RUS inversion framework. First, no general analytical solution exists for the computation of resonance frequencies for a three-dimensional (3D) volume of material. Therefore, approximate (numerical) methods must be employed.² Meaningful contributions of numerous researchers over many decades including Holland,⁸ Demarest,⁹ Ohno,¹⁰ and Visscher *et al.*¹¹ have culminated in a generalized numerical approach for solving the forward problem based on variational methods. The *xyz* method of Visscher *et al.* is applicable to most simple specimen geometries and requires minimal computational resources,¹¹ given that all of the requisite information about the specimen geometry and material properties are provided. Ultimately, the viability of any inverse method for evaluating elastic properties from RUS measurements depends on an efficient and accurate forward calculation method.²

As the resonance frequencies of a specimen are dependent on its shape, elastic constants, density, external forces, and the orientation of the elastic body, the deconvolution of unknown parameters from resonance frequencies is far from a trivial task. In fact, no closed form solution to this problem exists whereby unknown attributes of the specimen geometry or material properties are computed directly from a measured list of resonance frequencies.² Therefore, inverse methods are employed to find values for the unknown parameters that bring a forward calculation of resonance frequencies into sufficient agreement with the resonance frequencies measured via RUS. Historically, elastic property inversion of RUS data has been accomplished via nonlinear least square optimization algorithms.^{2,12,13} For this task, the Levenberg-Marquardt (LM) algorithm popularized by Migliori *et al.*^{2,5} combines a modified Newton method with the steepest descent algorithm^{5,7,12} to perform the inversion.

The least squares approach for overdetermined systems is popular due to its computational simplicity, but often suffers from a lack of robustness to outliers in the data set.¹⁴ Another unfortunate characteristic of optimization-based inversion methods is that different initial parameterizations (i.e., “best guess” elastic moduli and optimizer-specific parameters) can lead the optimization algorithm to alternative solutions, while “poor” initial guess values may preclude convergence entirely.^{15,16} This problem is not unique to the LM algorithm, as similar concerns exist for genetic algorithms.^{15,17,18}

An ideal inversion framework would be robust to uncertainty in the initial parameterization, noise in the measured data, misidentified modes, as well as missing or spurious modes, and would consistently converge to the correct solution. Ogi *et al.*¹⁹ demonstrate an optimizer-based inversion framework capable of reliable convergence without the benefit of quality initial guess moduli. But their framework first requires proper mode identification via laser Doppler interferometry mapping of resonance mode shapes.¹⁹ Indeed, collecting additional data may simplify the inversion procedure, as would preparing a specimen with its crystal axes aligned with the specimen axes. But these methods only complicate experimental procedures while adding considerable cost. Ogi *et al.* go on to conclude that “correct mode identification is essential for successful, optimum determination of material coefficients.”¹⁹ However, as this research will demonstrate, mode identification, quality initial guess moduli, and x ray measurements of crystal orientation are by no means essential elements of a robust RUS inversion framework.

B. Implications of a Bayesian approach to inversion

A considerable downside to Bayesian techniques is the increased computation time, as the number of forward calculations necessary for Hamiltonian Monte Carlo (HMC) sampling of the posterior distributions is significantly greater than optimizer-based methods. However, thanks to advances in computing, Bayesian estimates of elastic constants and orientation parameters from measured resonance frequencies is now possible.

Most classic RUS computations provide only point estimates of elastic constants; that is, single number estimates for each parameter in the RUS model regardless of the amount or precision of data. Point estimates are unsatisfactory in many inverse problems because they do not give information about how well a fit worked or how well a parameter is known. Bayesian techniques can be used to avoid these problems by systematically estimating uncertainty. It is very reasonable that a point estimate produces answers that are “good enough” in controlled experiments, but it is difficult to develop confidence in the methods for experiments where prior knowledge is sparse.

In a manner similar to that of Bernard *et al.*,²⁰ the approach developed here improves upon classical RUS inversion by reformulating the problem as a Bayesian inference and characterizing the unknown parameters through sampling the resultant posterior distribution. The capability to simultaneously estimate the orientation and elastic properties of elastically anisotropic bodies offers further improvement to classical RUS inversion techniques, simplifying specimen preparation procedures and eliminating a potential source of measurement error. Key advantages and disadvantages of a Bayesian formulation will be highlighted, along with experimental and computational considerations helpful for replication. Ultimately this work intends to demonstrate robust convergence behavior irrespective of initial parameterization and easy-to-interpret uncertainties for all parameter estimates.

II. EXPERIMENTAL METHODS

A. Specimen preparation

Regular parallelepiped specimens of fine-grained polycrystalline Ti-6Al-4V (Ti-64) and single crystal Ni-based superalloy CMSX-4 were machined via wire EDM (electrical discharge machining) and then carefully ground with 800-grit sandpaper to remove the superficial EDM damage layer. The Ti-64 specimen measured $7.753 \times 9.057 \times 13.199$ mm, with a mass of 4.0795 g and a calculated density of 4402 kg/m^3 . The CMSX-4 parallelepiped dimensions were $11.959 \times 13.953 \times 19.976$ mm with a mass of 29.0041 g and a calculated density of 8701.4 kg/m^3 . With randomly oriented grains, the Ti-64 specimen exhibits isotropic elastic symmetry while the single crystal CMSX-4 material possesses cubic symmetry. Minimizing geometric defects during specimen fabrication and precise measurement of the geometry and mass are important for minimizing the uncertainty in the HMC parameter estimates; with a “good” parallelepiped geometry, according to Migliori *et al.*, exhibiting dimensional errors less than 0.1%.⁵ The only notable divergence between this work and the typical specimen fabrication procedures outlined by Migliori and Sarrao² is that no attempt was made to align the axes of the parallelepiped to the crystallographic axes of the CMSX-4 material. As orientation and elastic constants will be determined simultaneously through inversion, any arbitrary misorientation between the crystal and specimen axes is allowed.

B. RUS

1. RUS experimental setup

RUS data were collected using commercially-available RUS equipment developed by the Vibrant Corporation (Albuquerque, NM) consisting of three primary components: a transceiver, a piezoelectric transducer (PT) cradle, and a computer control unit. Figure 1 shows the configuration of the PT cradle, comprised of three custom-built omnidirectional PTs held in a tripod configuration with adjustable optical table fixtures affixed to a vibration-damped breadboard. The PTs were custom built and consist of a cylindrical brass housing encasing a piezoelectric element, electrical leads,

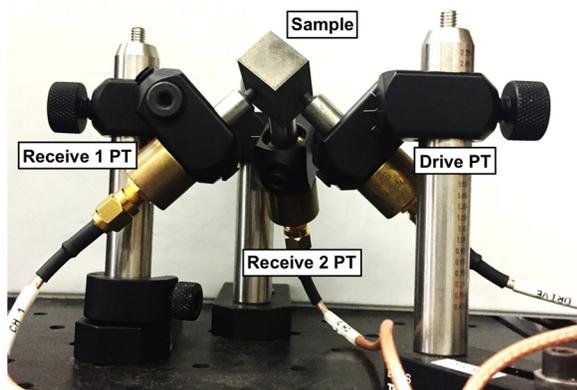


FIG. 1. (Color online) PTs configured into a cradle supporting the CMSX-4 parallelepiped specimen.

and a wear-resistant hemispherical silicon carbide tip. The parallelepiped specimens freely rest upon the silicon carbide-tipped transducers, with no couplant necessary to facilitate the transmission of vibrations between the transducers and the specimen.^{1,2}

The “drive PT” in Fig. 1 is driven with a swept sinusoidal signal from the transceiver to excite the specimen to resonate, while two “receive PTs” convert vibrations from the specimen back to an electrical signal that is returned to the transceiver and computer control unit for analysis. As a principle of mechanical resonance, a standing elastic wave develops throughout the specimen when the drive frequency approaches a resonance frequency of the specimen. This results in sample deflections that are amplified by hundreds to thousands of times that of the drive force amplitude, depending on the ultrasonic attenuation (i.e., damping characteristics) of the material.^{1,2} Plotting the signal registered by the two receive PTs as a function of the drive frequency yields a broadband RUS spectrum plot, with each peak indicating the frequency of a unique resonance mode. Figure 2 shows a broadband resonance spectrum plot collected from the CMSX-4 specimen, with 53 resonance modes across the 200 kHz broadband.

2. RUS measurement considerations

It is important to minimize the magnitude and the variability of any external forces on the specimen while collecting RUS data because these external forces can interfere with the free vibrations of the specimen and affect the frequencies that are measured.^{5,21} Configuring the PTs into a fixed cradle as detailed by Fig. 1 serves to minimize contact force variability between the specimen and the transducers across multiple measurements and specimen sizes. However, as the specimen is free to deflect away from the PT cradle during resonance the signal amplitude information is generally unreliable. For this reason the broadband RUS data plotted in Fig. 2 are given with arbitrary units. When amplitude information is necessary to measure acoustic attenuation, for instance, parallelepiped specimens are often pinched between two (often planar) PTs.²¹ Note that the corners of a parallelepiped are the optimal location for excitation and measurement of resonance frequencies,⁵ but for the purpose of determining elastic constants the cradle configuration has proven itself simple and effective.

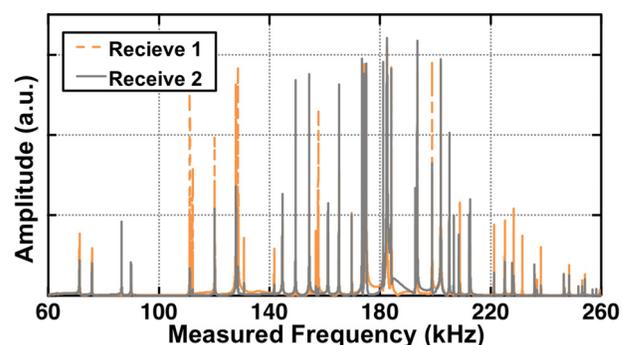


FIG. 2. (Color online) RUS broadband spectrum plot collected from the CMSX-4 specimen.

3. RUS data collection

Broadband RUS spectra were collected at room temperature and at standard atmospheric pressure in accordance with ASTM standard practice 2534-15.²² The broadband scans ranged from 60–260 kHz for the CMSX-4 specimen and from 100–375 kHz for the Ti-64 specimen. Distillation of the broadband resonance spectrum into a list of resonance frequencies was automated for consistency and verified through inspection as the data were collected. Missing or spuriously identifying a mode can greatly confound the process of determining elastic constants,² therefore considerable attention was directed toward ensuring that measured lists of resonance frequencies accurately reflect the specimen from which they were collected. Occasionally a resonance peak will be obscured by a higher-amplitude peak in close proximity, but rarely are modes completely missed by both receive PTs of the cradle configured as shown in Fig. 1. Nevertheless, each broadband measurement was repeated 5 times, with the specimen removed and replaced on the PT cradle after each measurement to ensure that any unintentional specimen-transducer interactions related to specimen placement would not be repeated. From the multiple broadband measurements a single average list of resonance frequencies was created, from which all computations and property estimates were based.

C. Crystallographic orientation measurements

To measure the orientation of the CMSX-4 crystal reference frame, a series of XRD measurements were collected using a Rigaku Smartlab High-Resolution Diffractometer with motorized RX-RY stage. First, the broad face of the specimen was prepared in accordance with standard metallographic techniques: wet grinding with 1200-grit paper, mechanical polishing via diamond suspension to 0.25 μm , and a final electrochemical etch. Then the specimen was affixed to the RX-RY stage of the diffractometer, which tilted in two orthogonal directions while Bragg peaks were recorded for $\{100\}$, $\{110\}$, and $\{111\}$ crystallographic planes. A least squares fit of the stage position, corresponding to the crystal plane normal vectors, determined the crystal orientation: (0.987, -0.00526 , -0.158 , 0.0164) as a passive unit quaternion (ordered as $wxyz$) and an uncertainty of approximately 1° .

III. COMPUTATIONS

Sections III A–III C discuss the development of a statistical model for characterizing elastic constants via RUS. Section III A reviews the basic mechanics of the system, which are very similar to previous RUS works.^{5,11,20} Section III B introduces the Bayesian RUS model, and Sec. III C describes how the Bayesian computations are carried out.

A. Forward model

Data in this experiment are modeled as

$$X_1, X_2, \dots, X_N = f(c_{11}, c_{12}, \dots) + \xi, \quad (1)$$

where X_1, X_2, \dots, X_N are the measured resonance modes, f is the forward model which computes the resonance modes of

the specimen given the necessary elastic constants, and ξ is a noise term that represents the combined uncertainty in fabrication and measurement of the specimen. Following Refs. 11 and 20, the specimen in the forward model is approximated as an undamped linear harmonic oscillator with free boundary conditions. This results in the generalized eigenvalue problem

$$Ku = \omega^2 Mu, \quad (2)$$

where K is the stiffness matrix, M is the mass matrix, and u is a vector of displacements in three dimensions. K and M are computed from either a Rayleigh-Ritz or finite element method (FEM) approximation to the problem. The square roots of the eigenvalues (ω) of this generalized eigenvalue problem being the measured resonance modes.

Herein, a Rayleigh-Ritz solver with polynomial basis is used to compute these eigenvalues (again following the derivations in Refs. 11 and 20). An FEM solver was tested as well, though it was significantly slower than the Rayleigh-Ritz method for the simple geometry (parallelepiped) used here. The equations for the basis polynomials, stiffness matrix (K) and mass matrix (M) from Ref. 20 are reproduced in Eqs. (3)–(6) (using Einstein notation for tensors)

$$\phi_\lambda(x, y, z) = x^n y^m z^l, \quad (3)$$

$$\{\lambda = (n, m, l) \mid n, m, l \in \mathbb{N}^0, n + m + l \leq N\}, \quad (4)$$

$$K_{i\lambda, k\lambda'} = \sum_{j,i=1}^3 C_{ijkl} \int_V \epsilon_{ij}(\phi_\lambda) \epsilon_{kl}(\phi_{\lambda'}) dV, \quad (5)$$

$$M_{i\lambda, k\lambda'} = \rho \delta_{ik} \int_V \phi_\lambda \phi_{\lambda'} dV. \quad (6)$$

In Eqs. (3)–(6) above, N is the maximum order of polynomials used in the resonance approximation (usually 10–14), where C_{ijkl} is the stiffness tensor.

If the crystal lattice is not aligned with the specimen axes, then C_{ijkl} must be adjusted. If the rotation from the specimen axes to the crystal axes is represented as a passive unit quaternion with elements $wxyz$, and C'_{pqrs} are the elastic constants of an aligned specimen, then the effective elastic constants of the rotated specimen, C_{ijkl} can be computed as follows:

$$q = \begin{bmatrix} w^2 + x^2 - y^2 - z^2 & 2(xy - wz) & 2(xz + wy) \\ 2(yx + wz) & w^2 - x^2 + y^2 - z^2 & 2(yz - wx) \\ 2(zx - wy) & 2(zy + wx) & w^2 - x^2 - y^2 + z^2 \end{bmatrix}, \quad (7)$$

$$C_{ijkl} = q_{ip} q_{jq} C'_{pqrs} q_{kr} q_{ls}. \quad (8)$$

For more information, see Sec. 3.2.11 in Bower.²³

B. Building a statistical RUS model

Since the noise (ξ) in Eq. (1) was modeled as a random variable, the outputs (X_1, X_2, \dots, X_N) are also random variables. Assuming the noise of each mode is normally distributed with mean zero and variance σ^2 , the probability of

measuring a set of resonance modes (X_0, X_1, \dots, X_N) given the elastic constants (c_{11}, c_{12} , etc.) can be written using the forward model, f , as

$$P(X_0, X_1, \dots, X_N | c_{11}, c_{12}, \dots) \sim \mathcal{N}(f(c_{11}, c_{12}, \dots), \sigma^2). \quad (9)$$

This equation, usually written in shorthand as $P(X|\theta)$, is known as the likelihood and is the probability of measuring a set of data given some fixed parameters. For an inverse problem, it is the opposite relation, $P(\theta|X)$, or the probability that parameters take certain values given the measured data. $P(\theta|X)$ is known as the posterior distribution, and in a Bayesian interpretation represents the uncertainty in a set of parameters given the data. The posterior can be computed from the likelihood and any prior knowledge about the parameters by using Bayes' rule

$$\underbrace{P(\theta|X)}_{\text{Posterior}} = \frac{\overbrace{P(X|\theta)}^{\text{Likelihood}} \overbrace{P(\theta)}^{\text{Prior}}}{\underbrace{P(X)}_{\text{Prior Predictive}}}. \quad (10)$$

The $P(\theta)$ term is called the prior because it is specified to contain the prior beliefs about the probabilities of certain parameters (which could be as simple as requiring a parameter to be positive, or something much more complicated). $P(X)$ is the prior predictive distribution. It can be computed from the likelihood and the prior ($P(X) = \int P(X|\theta)P(\theta)d\theta$), but for the Monte Carlo computations here can be regarded as a normalization constant and ignored.

If the model and data match well, the posterior distributions on the parameters will be tight, and it will be easy to extract estimates for the parameters. However, from the outset, it is unclear how informative the posterior will actually be. It is possible, for instance, to have higher confidence in one parameter than another, or to have multiple values of a parameter that give equally likely explanations for the data. Because of this uncertainty, it is important to compute the full posterior $P(\theta|X)$ and work with confidence intervals rather than just returning a single estimate.

Returning to Eq. (1), this paper assumes the noise (ξ) is distributed normally about every resonance mode with a single variance (i.e., the scale of the noise does not change for each mode). The primary justification for picking this model comes *a posteriori* by checking that the model explains the data well with few outliers. Of course such a check is not always so simple, as the standards for “explaining the data well” and the definition of an outlier are very application specific. In this work the RUS measurement noise (presumably from inconsistent placement of the specimen on the transducer cradle or specimen transducer interactions) was much smaller than the noise inherent to the specimen itself (presumably from specimen fabrication). Unfortunately, specimen fabrication is not easily repeatable in a manner that would allow for multiple independent samples, making it difficult to ever fully justify these assumptions. A seemingly unavoidable feature of RUS measurements is that a

few of the lowest-frequency resonance modes are more difficult to measure consistently² which could also cause problems with the mode-invariant noise assumption.

Following the assumptions stated above, given that one set of resonance modes is available, every measurement goes toward estimating the lumped variance parameter. Thus the complete likelihood can be stated as

$$P(X|\theta) = \prod_i \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{-(f_i(\theta) - X_i)^2 / 2\sigma_i^2}, \quad (11)$$

where $f_i(\theta)$ is the i th computed resonance mode. By collecting many independent resonance mode measurements, the estimates for θ can be tightened to suitable levels.

The prior term [$P(\theta)$] can be used to specify prior information about a parameter. For instance, it makes sense to assume that the c_{11} elastic constant is positive, somewhere between zero and a few hundred gigapascals. This can be expressed by a uniform prior distribution $P(c_{11}) = \mathcal{U}(0 \text{ GPa}, 500 \text{ GPa})$. Likewise, perhaps a parameter is known to some precision, in which case a normal prior like $P(c_{11}) = \mathcal{N}(200 \text{ GPa}, 10 \text{ GPa})$ is reasonable. For this work, the prior on the variance parameter was set to $\sigma^2 = \mathcal{U}(0 \text{ kHz}, \infty \text{ kHz})$.

C. Computing the posterior (HMC)

Given a likelihood and prior, it is trivial to use Bayes' theorem to write out an expression for the posterior. However, evaluating this expression is difficult because the dimension of θ can be large and the cost of evaluating the likelihood high. It is possible, though, to approximate the posterior by drawing samples from it using Monte Carlo techniques.

The Monte Carlo technique used in this paper is HMC. To understand the results it will be useful to quickly review the characteristics of HMC and Markov Chain Monte Carlo (MCMC) methods in general. Perhaps the most common MCMC method is Metropolis Monte Carlo (Metropolis MC). In physics terms, Metropolis MC generates a sequence of samples s_0, s_1, \dots, s_N that represent states drawn from a thermodynamic equilibrium. The Metropolis algorithm, like all MCMC methods, proceeds sequentially. That is, state s_{i-1} is used to generate s_i . The jump from s_{i-1} to s_i is chosen randomly, and the decision to keep or reject the new state s_i is made based on the difference in an energy function $\Delta\Phi = \Phi(s_i) - \Phi(s_{i-1})$, representing the transition energy from state s_{i-1} to state s_i . If the Metropolis algorithm accepts and rejects are handled properly in accordance with the energy function Φ , then the sequence of states generated by the process will have physical meaning with regards to the thermal equilibrium of the simulated system.

In statistical applications, the sequence of states, s_0, s_1, \dots, s_N , is replaced with a sequence of parameterizations, $\theta_0, \theta_1, \dots, \theta_N$, and Φ is set equal to $\log P(X|\theta)P(\theta)$, the log of the joint distribution. A common pitfall for newcomers is wondering what role X plays in this. X is the measured data. It is fixed and does not change. Sampling only happens over the parameters (θ) in the joint distribution. With this choice of Φ , the Metropolis MC method will generate a sequence of

parameterizations where each parameterization, θ , is drawn from a distribution proportional to the true posterior. These samples can then be used to approximate the true posterior.

Application of Metropolis MC is mostly limited for computational reasons. In practice, Metropolis MC does not efficiently explore parameter space due to how it randomly selects new parameterizations with very little regard to the problem at hand. HMC addresses this issue where possible by using the gradient of $\log P(X|\theta)P(\theta)$ to select new states more intelligently. Compared to Metropolis MC (and many other MCMC methods), HMC generates posterior samples much more efficiently. The need for the derivatives of $\log P(X|\theta)P(\theta)$ limits HMC's applicability in general, but the necessary derivatives are available in the forward model used here. A key parameter for HMC is the time step, which will determine how efficiently the chain can move around parameter space. The time step cannot be too large though, or the HMC chain will go unstable and always reject new states.

Without going into too much detail, the HMC algorithm used in this paper was taken from Ref. 24 with adaptations for orientations by Ref. 25. The technique in Ref. 24 for using multiple steps sizes was critical for achieving efficient sampling.

1. Necessary derivatives

The gradient of $\log P(X|\theta)P(\theta)$ is derived here. All derivatives are computed with the chain rule. $P(\theta)$ is assumed equal to one to simplify the math (the non-negativity constraint on σ^2 is controlled with a parameter transformation²⁶).

2. Elastic constants

With $P(\theta)$ equal to one, the $\log P(X|\theta)P(\theta)$ term is simplified to $\log P(X|\theta)$. Using the chain rule to write the partial derivative of the log-likelihood with respect to the parameter c_{11} gives the sum

$$\frac{\partial \log P(X|\theta)}{\partial c_{11}} = \sum_i \frac{\partial \log P(X|\theta)}{\partial \omega_i} \frac{\partial \omega_i}{\partial c_{11}} \quad (12)$$

where ω_i is the i th resonance mode, or the i th index of $f(\theta)$. The partial derivative with respect to ω_i is easy to compute [given the likelihood in Eq. (11)]

$$\frac{\partial \log P(X|\theta)}{\partial \omega_i} = \frac{-(\omega_i - X_i)}{\sigma^2}. \quad (13)$$

The partial derivative of ω_i with respect to c_{11} (or any elastic constant) requires derivatives of the eigenvalues (the ω_i^2 's) of Eq. (2). Given a number of distinct eigenvalues (ω_i^2) and orthonormal eigenvectors (ν_i), the derivative of the i th eigenvalue can be computed as in Ref. 27,

$$\frac{\partial \omega_i^2}{\partial c_{11}} = \nu_i^T \frac{\partial K}{\partial c_{11}} \nu_i, \quad (14)$$

$$\frac{\partial \omega_i}{\partial c_{11}} = \frac{1}{2\omega_i} \frac{\partial \omega_i^2}{\partial c_{11}} = \frac{1}{2\omega_i} \nu_i^T \frac{\partial K}{\partial c_{11}} \nu_i. \quad (15)$$

The partial derivatives of the eigenvalues requires derivatives of the stiffness matrix from Eq. (2). These can be obtained from the construction in Eq. (5),

$$\frac{\partial K_{i\lambda, k\lambda'}}{\partial c_{11}} = \sum_{j,i=1}^3 \frac{\partial C_{ijkl}}{\partial c_{11}} \int_V \epsilon_{ij}(\phi_\lambda) \epsilon_{kl}(\phi_{\lambda'}) dV, \quad (16)$$

while $\partial C_{ijkl}/\partial c_{11}$ arises from the specific symmetry of the system. For a cubic crystal specimen, C_{ijkl} and $\partial C_{ijkl}/\partial c_{11}$ are given by

$$C_{ijkl} = \begin{bmatrix} c_{11} & c_{12} & c_{12} & 0 & 0 & 0 \\ c_{12} & c_{11} & c_{12} & 0 & 0 & 0 \\ c_{12} & c_{12} & c_{11} & 0 & 0 & 0 \\ 0 & 0 & 0 & c_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & c_{44} & 0 \\ 0 & 0 & 0 & 0 & 0 & c_{44} \end{bmatrix} \quad (17)$$

and

$$\frac{\partial C_{ijkl}}{\partial c_{11}} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}. \quad (18)$$

Combining Eqs. (13), (15), (16), and (18) give the necessary expression for $\log P(X|\theta)/\partial c_{11}$. This can be repeated for the other elastic constants as well.

Finally, instead of estimating the three cubic stiffness parameters, c_{11} , c_{12} , and c_{44} directly, a simple parameter transformation was employed for improved mobility and HMC sampling. The transformed parameter space, c_{11} , A , and c_{44} was used herein, with A being the cubic anisotropy ratio

$$A = \frac{2C_{44}}{C_{11} - C_{12}}, \quad (19)$$

as defined by Zener.²⁸

3. Noise term (σ)

The partial derivative of the log-likelihood with respect to the noise term σ is given by

$$\frac{\partial \log P(X|\theta)}{\partial \sigma} = -\frac{N}{\sigma} + \sum_i \frac{(\omega_i - X_i)^2}{\sigma^3}. \quad (20)$$

4. Lattice-specimen orientations

As stated in Sec. III A, the lattice-specimen orientation is parameterized as a passive unit quaternion. Though this complicates the calculations slightly as compared to Euler angles, parameterizing the problem with Euler angles produced unsatisfying results. With Euler angles, due to the degrees of freedom collapsing around the poles, it was difficult to determine whether the orientations were converging.

Quaternions, though expressed in four dimensions (w, x, y, z), live on a 3D manifold characterized by

$$w^2 + x^2 + y^2 + z^2 = 1. \quad (21)$$

In other words, not all combinations of four real numbers make a valid quaternion. In order to sample correctly on this manifold, the Geodesic HMC algorithm in Ref. 25 is used. Without going into detail, Byrne and Girolami²⁵ adjusts random momentum generation and time stepping in HMC to keep the quaternion parameters on the manifold [Eq. (21)].

Computing the partial derivatives of the log-likelihood with respect to an orientation parameter, for instance w , is the same as for an elastic constant up to the term $\partial C_{ijkl}/\partial w$ [simply replace c_{11} with w in Eqs. (13), (15), and (16)]. These require partial derivatives of Eqs. (5) and (7),

$$\begin{aligned} \frac{\partial C_{ijkl}}{\partial w} &= \frac{\partial q_{ip}}{\partial w} q_{jq} C'_{pqrs} q_{kr} q_{ls} + q_{ip} \frac{\partial q_{jq}}{\partial w} C'_{pqrs} q_{kr} q_{ls} \\ &+ q_{ip} q_{jq} C'_{pqrs} \frac{\partial q_{kr}}{\partial w} q_{ls} + q_{ip} q_{jq} C'_{pqrs} q_{kr} \frac{\partial q_{ls}}{\partial w}, \quad (22) \end{aligned}$$

$$\frac{\partial q}{\partial w} = 2 \begin{bmatrix} w & -z & y \\ z & w & -x \\ -y & x & w \end{bmatrix}. \quad (23)$$

IV. RESULTS

A. Polycrystalline Ti-64

Polycrystalline Ti-64 elastic constants were inverted using the geometries given in Sec. II A along with the first 30 measured resonance modes. While the fine grain microstructure and random crystallographic texture of the Ti-64 material would lead one to guess the elastic symmetry to be isotropic, an isotropy assumption was not enforced. The standard deviation of the noise was constrained to be positive. The initial conditions for the sampler were chosen as ($c_{11} = 2.0$, $A = 1.0$, $c_{44} = 1.0$, $\sigma = 5.0$). The final inverted parameters are given in Table I along with reference values from Fisher and Renken.²⁹ Note that Fisher and Renken²⁹ provide single crystal stiffness values for pure titanium, which are commonly accepted as sufficient for Ti-64. These elastic constants were then converted to isotropic moduli using a Voigt-Reuss-Hill average scheme.³⁰ Clearly the inverted value of $A = 1.000 \pm 0.002$ demonstrates the Ti-64 material is effectively isotropic.

The full estimated posterior distributions for the four parameters (two elastic constants, the anisotropy ratio, and the error term, σ) are shown in Fig. 3. All the parameters

TABLE I. Summary of estimated parameters (mean \pm standard deviation) for Ti-64 specimen alongside reference values.

Parameter	Ref. 29	Bayesian estimate
c_{11}	1.651×10^{11} Pa	$(1.703 \pm 0.015) \times 10^{11}$ Pa
c_{44}	4.330×10^{10} Pa	$(4.492 \pm 0.001) \times 10^{10}$ Pa
σ	—	$(0.414 \pm 0.058) \times 10^3$ Hz
A	1.000 ^a	1.000 ^b \pm 0.002

^aThe reference anisotropy ratio is equal to 1 as a result of the Voigt-Reuss-Hill (Ref. 30) polycrystalline average.

^bThe estimated anisotropy ratio was allowed to vary.

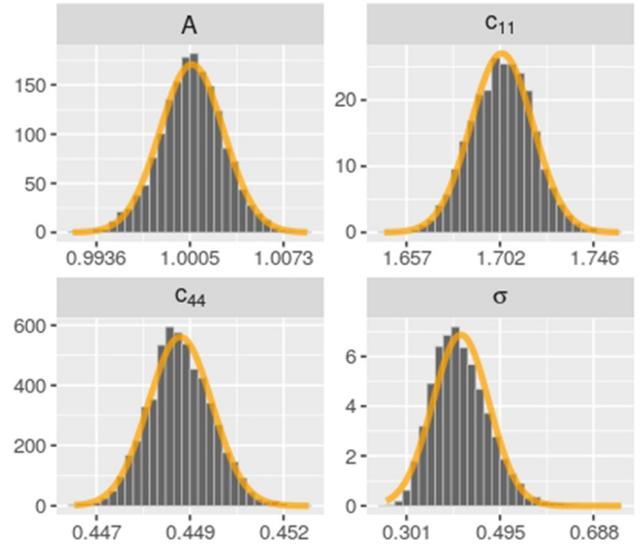


FIG. 3. (Color online) Approximate posterior distributions (normalized to a PDF scale) for c_{11} , A , c_{44} , and σ computed from 4000 HMC samples. Normal PDF fits are superimposed to justify the use of mean and std to characterize the posteriors.

(c_{11} , A , c_{44} , and σ) are well approximated by the superimposed normal distributions. This (visually estimated) quality of the fits justifies the use of mean and standard deviation, summarized in Table I.

In Bayesian modeling, the two basic tools for validating an inversion are trace plots, which are plots of the sequence of states sampled by the MCMC sampler, and posterior predictive distribution plots. The trace plots for the last 4000 posterior samples for the Ti-64 specimen are plotted in Fig. 4. Extracting information from a trace plot is fairly straightforward. If the distribution of samples in the trace plot is stationary for a long period of time, the chain is assumed (but not guaranteed) to be sampling from the true posterior. Trace plots are usually the easiest place to detect modeling problems. If, for instance, a parameter enters an invalid range or

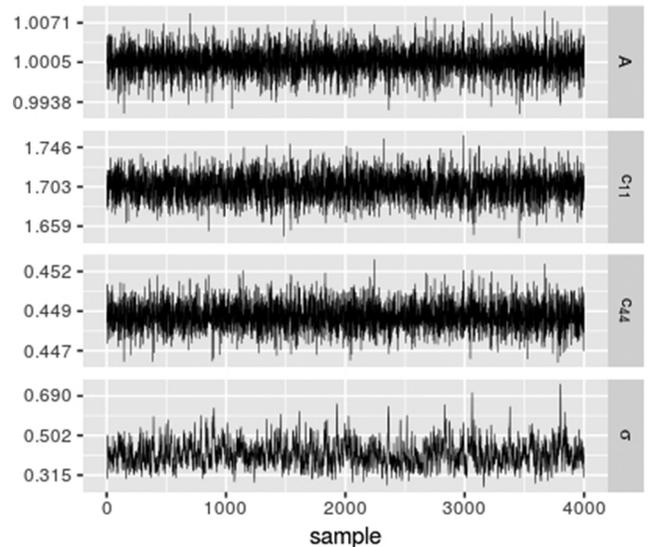


FIG. 4. Traceplots (last 4000 samples) for the estimated parameters for the polycrystalline Ti specimen.

shows a bimodal tendency, it is usually easy to spot in the trace plots. The trace plots in Fig. 4 show nothing unusual, suggesting that the samples themselves come from the true posterior of the model.

The second tool for validating a Bayesian modeling process are the posterior predictive distributions. The posterior predictive distributions are the distributions generated if samples from the approximated posterior distribution are used to generate new resonance modes. The quality of fit can be evaluated by comparing the posterior predictive distributions to the measured resonance modes. In the context of RUS measurements, evaluating the “quality of fit” means making sure there are not many outliers in the data (for instance, by making sure 95% of the data is within the 95% posterior predictive intervals) and identifying any systematic biases in the posterior predictive means. One example of this would be if the first ten posterior predictive resonance modes were all estimated with mean frequencies less than their respective measured modes. Errors should appear random. Any structure suggests a problem in the model or the data. For RUS, this could include missing or spurious resonance modes as part of the measured frequency list that will be readily identifiable when compared to the posterior predictive distributions. If a missing mode is identified, the term in the likelihood in Eq. (11) corresponding to the missing mode is removed (since that data are not available), and the inference rerun for improved results.

The posterior predictive distributions for the Ti specimen are shown in Table II, along with the first 30 measured resonance modes for comparison. As can be seen, all but one measured mode is in the 95% posterior interval and the estimated noise level is on par with what could be expected by measuring other specimens produced with the same specifications. The average quality factor (Q Ave.) and the standard deviation of the measured modes is provided for context about the repeatability and quality of the RUS measurements. As the specimen was removed, rotated, and then replaced on the transducer cradle between each broadband measurement, the standard deviation of the measured frequencies varied more significantly than they would have if the specimen was not removed between measurements. But having successfully identified each of the first 30 modes with an average standard deviation of 0.05 kHz, the procedure is deemed acceptable. While the measurement uncertainty is approximately one-quarter of the noise estimate, itself representing a combination of (measured and modeled) sources.

Given the combined evidence from the trace plots, the posterior predictive distributions, and the posterior distributions themselves, it is reasonable to conclude that the model describes the data well and that statements based on the computed posterior distributions should represent the physical system.

B. Single crystal CMSX-4

The single crystal CMSX-4 elastic constants and orientation were inverted using the geometries given in Sec. II A and the first 30 measured resonance modes. Instead of running a single, long HMC chain, the inference was broken

TABLE II. Ti-64 specimen measured and posterior predictive modes (mean \pm standard deviation).

Mode	Measured (kHz)	Q Ave. ^a	Posterior ^b (kHz)
1	109.076 \pm 0.027	3832	108.86 \pm 0.43
2	136.503 \pm 0.083	3209	135.97 \pm 0.43
3	144.899 \pm 0.051	3157	144.42 \pm 0.41
4	184.926 \pm 0.046	6482	184.60 \pm 0.44
5	188.476 \pm 0.029	4781	187.99 \pm 0.43
6	195.562 \pm 0.032	6828	195.56 \pm 0.44
7	199.246 \pm 0.029	7016	199.19 \pm 0.44
8	208.460 \pm 0.078	6162	208.08 \pm 0.43
9	231.220 \pm 0.030	7241	231.63 \pm 0.44
10	232.630 \pm 0.042	6759	232.47 \pm 0.44
11	239.057 \pm 0.033	2874	239.09 \pm 0.45
12	241.684 \pm 0.077	7148	242.24 \pm 0.44
13	242.159 \pm 0.030	7367	242.53 \pm 0.44
14	249.891 \pm 0.079	7436	249.89 \pm 0.42
15	266.285 \pm 0.097	7771	267.20 \pm 0.46
16	272.672 \pm 0.069	8376	272.49 \pm 0.44
17	285.217 \pm 0.037	7984	285.04 \pm 0.45
18	285.670 \pm 0.095	8008	285.65 \pm 0.43
19	288.796 \pm 0.023	10 400	289.23 \pm 0.45
20	296.976 \pm 0.060	7651	296.77 \pm 0.42
21	301.101 \pm 0.030	8687	301.60 \pm 0.46
22	303.024 \pm 0.053	8854	303.03 \pm 0.43
23	305.115 \pm 0.058	9296	305.02 \pm 0.44
24	305.827 \pm 0.067	9509	305.26 \pm 0.47
25	306.939 \pm 0.034	10 706	306.39 \pm 0.45
26	310.428 \pm 0.024	6946	310.15 \pm 0.44
27	318.000 \pm 0.042	11 199	317.51 \pm 0.45
28	319.457 \pm 0.040	7600	319.88 \pm 0.47
29	322.249 \pm 0.053	10 011	322.17 \pm 0.42
30	323.464 \pm 0.027	9986	322.93 \pm 0.44

^aThe average quality factor (Q Ave.) is provided along with the standard deviation of the measured frequency as an indicator of the quality and repeatability of the RUS measurements.

^bEstimated means and standard deviation of the posterior predictive distribution were generated using the last 200 samples from the HMC trajectory (Fig. 4). The only mode outside of the 95% posterior intervals is mode 15.

into warmup and post-warmup stages (as described in Ref. 31, Chap. 12). In the warmup stage, four chains are run from a fixed initial condition ($c_{11} = 2.0$, $A = 1.0$, $c_{44} = 1.0$, $\sigma = 5.0$) to find reasonable parameter estimates. In the warmup stage the sampler needs to move around very-low probability areas of parameter space, and for stability the HMC time step must be kept relatively small. After running the chains long enough to reach the high probability region of parameter space, new chains were initialized with larger HMC time steps to more efficiently explore the posterior. In both stages, multiple chains are run to verify that the HMC is converging to the same solution, while only the samples from the post-warmup stage are used for the inferences.

The warmup traceplots are shown in Fig. 5, while the post-warmup traceplots are shown in Fig. 6. As can be seen in these figures, all the chains in both warmup and post-warmup stages are sampling the same region of parameter space, while a key difference between the plots is that the post-warmup chains more aggressively explore the c_{11} and σ parameter spaces (due to the larger HMC time step). Table

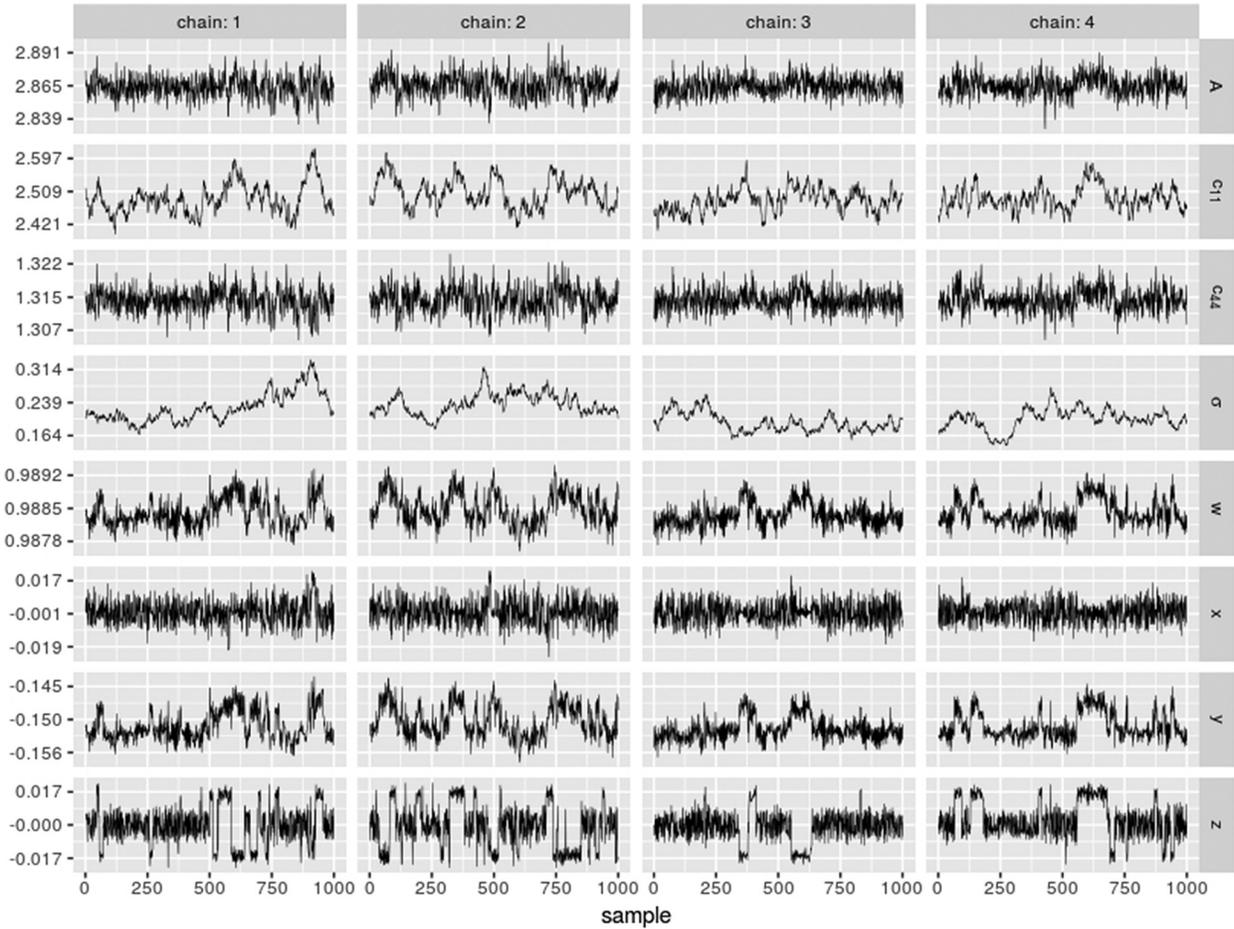


FIG. 5. Warmup traceplots (four different chains of 1000 samples each) for the estimated parameters in the single crystal CMSX-4 specimen. The rough means of the parameters in these chains were used to initialize four chains with larger time steps to more efficiently explore the posterior.

III summarizes the material parameter estimates from each of the four 1000-sample HMC chains, as well as the sum total of the four chains presented in Fig. 6. The consistency between the mean values of the four chains is a testament to the reliability of the inference, while the standard deviation associated with each mean demonstrates the precision of the parameter estimate. For example, the standard deviation of the c_{11} and c_{44} parameters are 4.0 and 0.26 GPa, respectively. This demonstrates the superior precision of the c_{44} estimate as compared to the c_{11} estimate, and is consistent with RUS-based elastic property inversions reported previously by Migliori *et al.*⁵

As stated in Sec. IV A, the posterior predictive distributions elucidate the quality of the fit and are provided for the CMSX-4 specimen in Table IV. Alongside these data are the measured frequencies and their standard deviation, as well as the quality factor for the first 30 resonance modes. As before, the one measured value outside the 95% posterior predictive interval (mode 24) is indicated in bold, while the average measured standard deviation of 0.03 kHz is approximately one-eighth the value of the noise estimate. Given the relatively small standard deviation associated with each mode, having one of 30 modes outside the 95% posterior predictive interval is not an issue.

Estimated posterior distributions of the elastic constants (c_{11} , A , and c_{44}) and the noise parameter (σ) are shown in Fig. 7. As each parameter exhibits a normal distribution, it makes sense to summarize the parameter estimates by their mean and standard deviation as provided in Table V. Reference CMSX-4 elastic constants from Ref. 32 are provided for context, and agree well with the estimates of this work.

Due to the symmetry inherent to crystals, in particular cubic crystals, there are numerous symmetrical representations for a given crystal orientation. Therefore it is possible, indeed common, for two orientation quaternions with substantially different component values to represent crystal orientations that are physically very close to each other. In order to visualize the posterior of the crystallographic orientation and produce two-dimensional plots, orientation quaternion were transformed into cubochoric coordinates³³ before plotting in Fig. 8. These plots detail the location of the measured orientation in orange with respect to each of the 4000 HMC sample orientations plotted in black. Although it is not particularly good that the measured crystallographic orientation is away from the bulk of the posterior, uncertainty in the XRD measurement of approximately 1° as noted in Sec. II C could easily place the measured

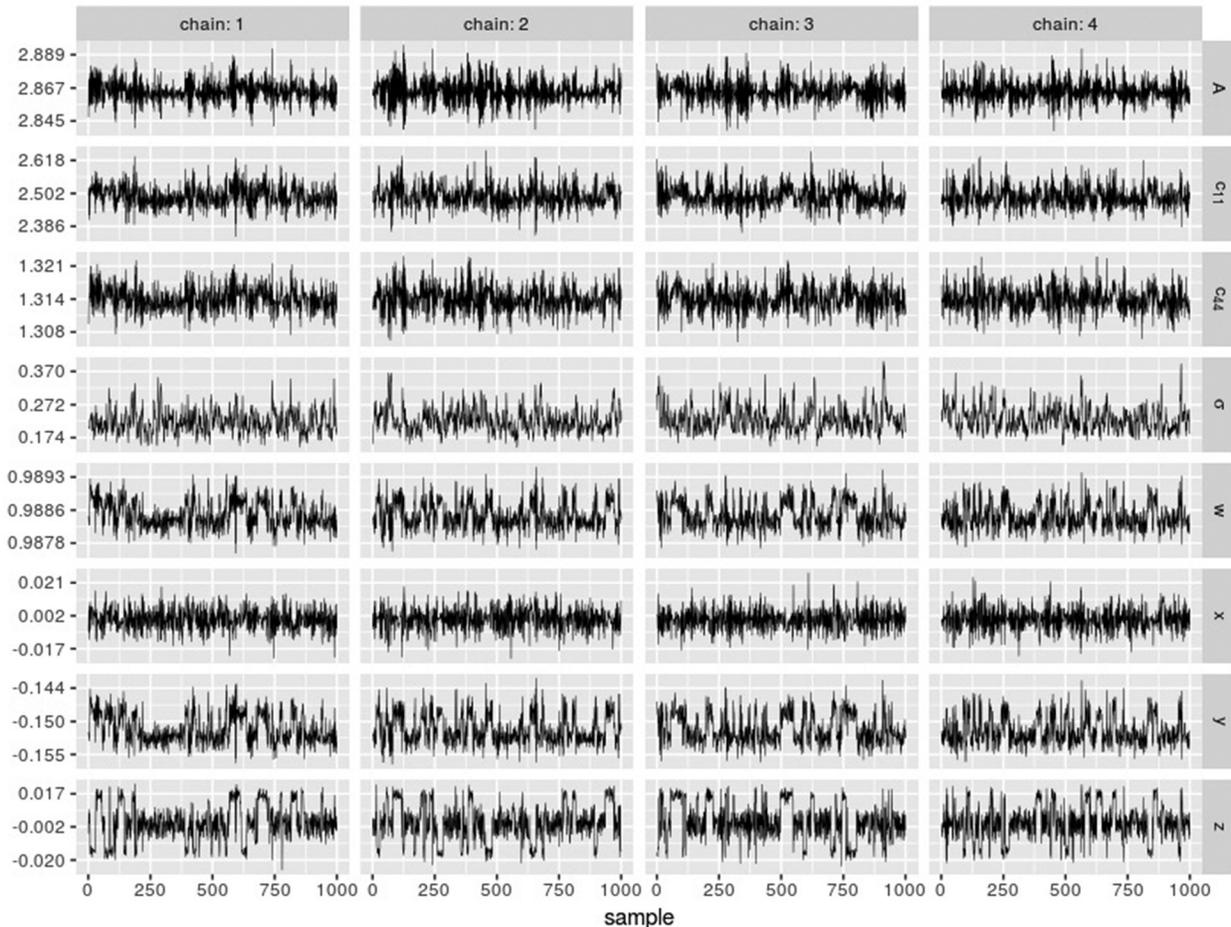


FIG. 6. Sampling traceplots (four different chains of 1000 samples each) for the estimated parameters of the single crystal CMSX-4 specimen. The elastic constants appear to be unimodal, but the orientation parameters are multimodal.

orientation within a higher-likelihood region of the posterior; given that a rotation of 1° corresponds roughly to a distance of 0.01 in cubochoric coordinates. Furthermore, the total misorientation angle (in degrees) between the measured orientation and each of the inverted orientations is detailed by the histogram in the bottom right of Fig. 8. Given the overwhelming majority of misorientation angles were calculated as between 1.0° and 2.5° , our confidence in estimating crystallographic orientation from RUS-measured resonance frequencies is high, with results summarized in Table VI.

To demonstrate that inferring the crystallographic orientation was necessary, an inversion was run without these degrees of freedom. Figure 9 shows the results. As can be seen, c_{11} is not converging to a steady state distribution and the inversion fails. The scale of σ indicates the fit is not good compared to the inversion with the orientation parameters included.

Given that all four warmup and post-warmup chains converge to similar physically realistic solutions and that the posterior predictive distributions align well with the measured data, it is reasonable to conclude that this model describes the data well and that statements based on the computed posterior distributions should represent the physical system. This is confirmed with the reference elastic constants and measured crystallographic orientations.

C. Computation efficiency

The Ti-64 inversion ran for approximately 8 h on a quad core Intel i7-2600k desktop computer, while the CMSX-4 inversion took a few days. Each HMC sample takes 50 forward evaluations of the RUS model (so 100 000 total for the examples presented here). Since the eigensolve in the forward model evaluation constitutes a large majority of the total evaluation time, it is simple to estimate how long an inversion might take if the forward model is available. If the forward model takes a second to evaluate, the sampler converges to a steady posterior distribution within 1000 HMC samples, and 2000 samples from the posterior are computed, the calculation should take a little under 2 days.

D. Avenues for future work

It should be possible to parameterize certain aspects of the specimen geometry in the same way as the elastic constants and estimate them on the fly, though there are known identifiability problems with the eigenvalue problem.⁵ It also might make sense to replace the forward model with an approximate model to decrease computation time, particularly as the number of unknown parameters is increased. Various techniques could be employed for this, such as Gaussian processes (a frequently used technique in Statistics

TABLE III. Summary of CMSX-4 parameter estimates and 95% posterior interval from each of the four HMC chains of 1000 samples each.

Parameter (units)	Chain	Estimate mean \pm st. dev.	Posterior int.	
			2.5%	97.5%
c_{11} (10^{11} Pa)	1	2.492 \pm 0.040	2.418	2.574
	2	2.492 \pm 0.041	2.415	2.580
	3	2.493 \pm 0.041	2.418	2.578
	4	2.489 \pm 0.040	2.417	2.569
	1-4	2.492 \pm 0.040	2.417	2.576
c_{44} (10^{11} Pa)	1	1.3145 \pm 0.0025	1.3099	1.3196
	2	1.3144 \pm 0.0027	1.3089	1.3203
	3	1.3145 \pm 0.0026	1.3092	1.3198
	4	1.3143 \pm 0.0026	1.3094	1.3194
	1-4	1.3144 \pm 0.0026	1.3093	1.3197
A	1	2.8652 \pm 0.0075	2.8509	2.8801
	2	2.8651 \pm 0.0089	2.8483	2.8831
	3	2.8652 \pm 0.0079	2.8493	2.8804
	4	2.8650 \pm 0.0078	2.8492	2.8805
	1-4	2.8651 \pm 0.0081	2.8493	2.8811
w	1	0.9885 \pm 0.0003	0.9879	0.9891
	2	0.9884 \pm 0.0003	0.9879	0.9891
	3	0.9884 \pm 0.0003	0.9879	0.9891
	4	0.9884 \pm 0.0003	0.9879	0.9891
	1-4	0.9884 \pm 0.0003	0.9878	0.9891
x	1	0.0000 \pm 0.0062	-0.0117	0.0119
	2	0.0002 \pm 0.0063	-0.0121	0.0130
	3	-0.0002 \pm 0.0061	-0.0114	0.0116
	4	0.0000 \pm 0.0063	-0.0124	0.0128
	1-4	0.0000 \pm 0.0062	-0.0118	0.0124
y	1	-0.1510 \pm 0.0025	-0.1549	-0.1461
	2	-0.1512 \pm 0.0025	-0.1552	-0.1460
	3	-0.1511 \pm 0.0024	-0.1550	-0.1464
	4	-0.1514 \pm 0.0024	-0.1549	-0.1462
	1-4	-0.1512 \pm 0.0024	-0.1550	-0.1461
z	1	0.001 \pm 0.010	-0.018	0.019
	2	0.000 \pm 0.010	-0.018	0.019
	3	0.001 \pm 0.010	-0.018	0.019
	4	0.001 \pm 0.010	-0.018	0.019
	1-4	0.001 \pm 0.010	-0.018	0.019

for interpolation), or generalized polynomial chaos (a common method for uncertainty quantification in engineering). While there is the issue that these methods require fairly good priors to be usable, inference on approximate models is much faster and could be used to quickly evaluate specimen characteristics without running a full Bayesian inversion. The noise models could be improved to account for outliers or greater variability in certain modes—as is often reported for the first few lowest-frequency modes.² Finally, accounting for missing or spurious modes a la Ref. 20 may facilitate inversions from RUS spectra that are complicated by a variety of factors including low stiffness or high damping.

V. CONCLUSION

A novel and robust RUS inversion framework for characterizing elastic properties and crystal orientation of paralleliped specimens using a Bayesian modeling approach

TABLE IV. CMSX-4 specimen measured and posterior predictive modes (mean \pm standard deviation).

Mode	Measured (kHz)	Q Ave. ^a	Posterior ^b (kHz)
1	71.259 \pm 0.044	2279	71.22 \pm 0.24
2	75.759 \pm 0.025	2864	75.69 \pm 0.23
3	86.478 \pm 0.030	3524	86.46 \pm 0.26
4	89.947 \pm 0.048	3077	90.00 \pm 0.24
5	111.150 \pm 0.033	4155	111.06 \pm 0.26
6	112.164 \pm 0.027	4492	112.01 \pm 0.29
7	120.172 \pm 0.033	5048	120.32 \pm 0.26
8	127.810 \pm 0.042	4462	127.98 \pm 0.25
9	128.676 \pm 0.040	4747	128.64 \pm 0.24
10	130.740 \pm 0.034	5342	130.75 \pm 0.24
11	141.700 \pm 0.037	5299	141.79 \pm 0.25
12	144.504 \pm 0.065	5603	144.36 \pm 0.23
13	149.401 \pm 0.025	5918	149.52 \pm 0.26
14	154.351 \pm 0.026	5942	154.42 \pm 0.25
15	156.782 \pm 0.028	6761	156.97 \pm 0.25
16	157.555 \pm 0.041	6123	157.57 \pm 0.25
17	161.088 \pm 0.063	6284	160.97 \pm 0.24
18	165.103 \pm 0.022	6756	165.21 \pm 0.28
19	169.762 \pm 0.044	6995	169.77 \pm 0.27
20	173.449 \pm 0.050	712	173.28 \pm 0.26
21	174.117 \pm 0.024	6437	174.13 \pm 0.26
22	174.906 \pm 0.054	6916	174.68 \pm 0.26
23	181.120 \pm 0.042	6632	181.54 \pm 0.25
24	182.459 \pm 0.037	7475	181.87 \pm 0.25
25	183.986 \pm 0.042	7837	183.81 \pm 0.28
26	192.681 \pm 0.032	7197	192.83 \pm 0.24
27	193.436 \pm 0.021	7113	193.71 \pm 0.27
28	198.794 \pm 0.034	7249	198.95 \pm 0.25
29	201.902 \pm 0.032	8529	201.89 \pm 0.27
30	205.015 \pm 0.031	8808	204.85 \pm 0.33

^aThe average quality factor (Q Ave.) is provided along with the standard deviation of the measured frequency as an indicator of the quality and repeatability of the RUS measurements.

^bEstimated means and standard deviation of the posterior predictive distribution were generated using the last 200 samples from the HMC trajectory (Fig. 6). Only mode 24 is outside the 95% posterior intervals.

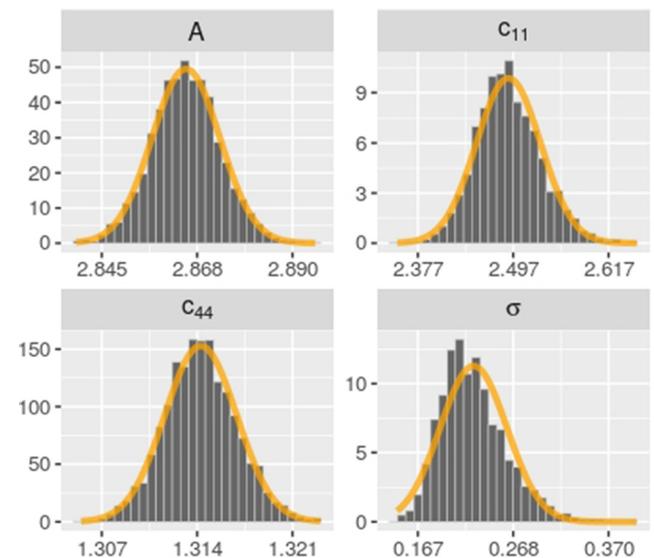


FIG. 7. (Color online) Approximate CMSX-4 posterior distributions for c_{11} , A, c_{44} , and noise (σ) parameters, computed from four chains of 1000 HMC samples each. Normal PDF fits are superimposed to justify the use of mean and standard deviation to characterize the posteriors.

TABLE V. Summary of estimated parameters (mean \pm standard deviation) for the CMSX-4 specimen alongside reference values.

Parameter	Ref. 32	Bayesian estimate
c_{11}	2.52×10^{11} Pa	$(2.492 \pm 0.040) \times 10^{11}$ Pa
c_{44}	1.31×10^{11} Pa	$(1.314 \pm 0.003) \times 10^{11}$ Pa
σ	—	$(0.229 \pm 0.037) \times 10^3$ Hz
A	2.88	2.865 ± 0.008

TABLE VI. Summary of CMSX-4 specimen crystal orientation from XRD measurement and Bayesian estimate (mean \pm standard deviation).

Parameter	XRD measurement	Bayesian estimate
w	0.987	0.9884 ± 0.0003
x	-0.00526	0.000 ± 0.0062
y	-0.158	-0.1512 ± 0.0024
z	0.0164	0.001 ± 0.010

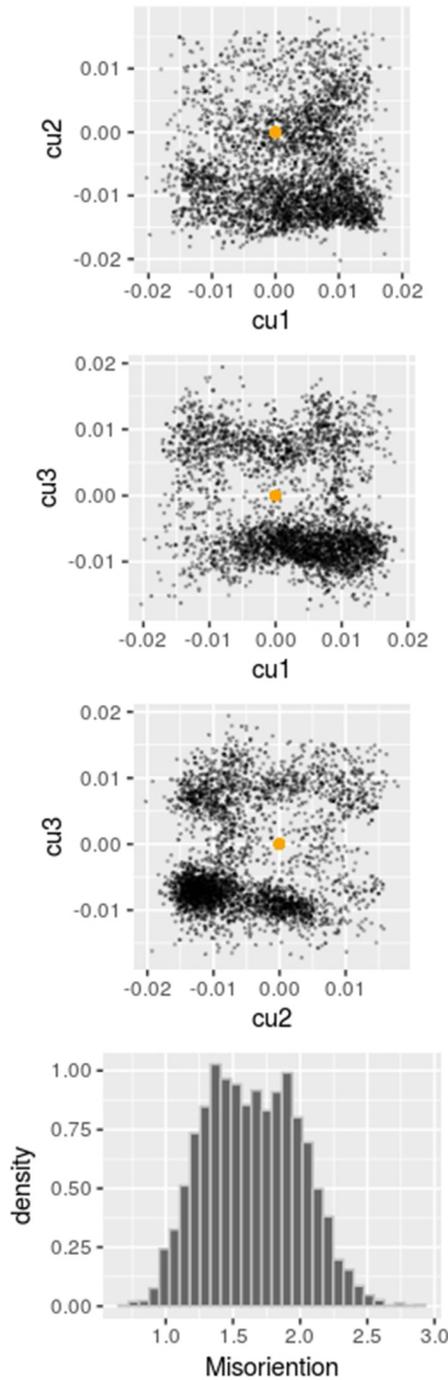


FIG. 8. (Color online) Crystallographic orientation posterior plotted in cubochoric coordinates. The distribution of minimum misorientation angle (in degrees) between the measured and estimated orientations is at the bottom. Note that a difference of one-hundredth on any cubochoric axis corresponds roughly to 1° of rotation.

and HMC sampling has been developed. The inversion framework was tested with two experimental datasets for validation: a fine-grained Ti-64 specimen and a single crystal specimen of Ni-based superalloy CMSX-4 with misaligned specimen-crystal reference frames. The Ti-64 specimen exhibited elastic isotropy with c_{11} and c_{44} stiffness constants in agreement with literature values for a Voigt-Reuss-Hill average of randomly-oriented grains. Inversion of the CMSX-4 data yielded accurate estimates of the three independent elastic moduli in strong agreement with literature values, while the crystallographic orientation was determined to within approximately 2° of XRD measured values. While the current version of the code requires substantially greater computation time as compared to conventional RUS inversion schemes, it also provides the following advantages:

- (1) Built-in uncertainty estimates on all parameters.
- (2) Simultaneous estimation of elastic constants and crystal orientation.
- (3) Simplified requirements for specimen preparation.
- (4) Robustness to common RUS inversion problems like misidentified resonance modes and initial parameterization uncertainty.

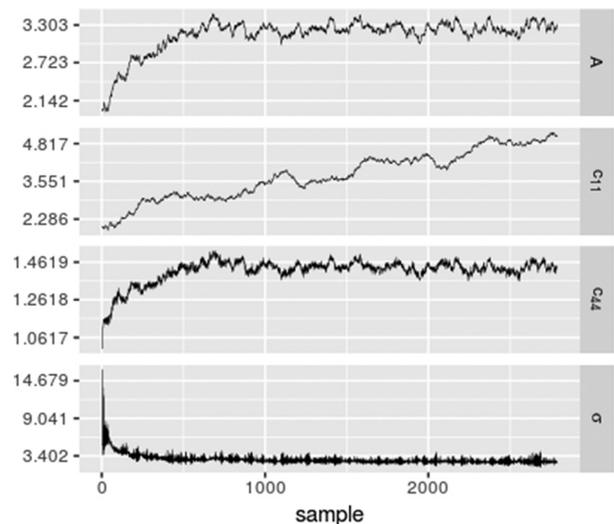


FIG. 9. Traceplots for the estimated parameters in the single crystal CMSX-4 specimen without including misorientation estimation.

ACKNOWLEDGMENTS

This research was funded by the U.S. Air Force Research Laboratory (AFRL) through BAA Contract No. FA8650-15-M-5208, SBIR Contract No. FA8650-15-M-5074, and U.S. DOE Grant No. DE-SC0008975. California Nanosystems Institute (CNSI) shared facilities: the Center for Scientific Computing and the MRL X-Ray Facility were utilized in conducting this research, with CNSI funding from NSF MRSEC (Grant No. DMR-1121053) and NSF CNS-0960316. Vibrant Corporation provided the RUS measurement equipment and many helpful discussions. This paper was cleared for public release by the AFRL, Case No. 88ABW-2017-3838.

- ¹R. G. Leisure and F. A. Willis, "Resonant ultrasound spectroscopy," *J. Phys. Condens. Matter* **9**(28), 6001–6029 (1997).
- ²A. Migliori and J. Sarrao, *Resonant Ultrasound Spectroscopy: Applications to Physics, Materials Measurements, and Nondestructive Evaluation* (Wiley-Interscience, New York, 1997), Chaps. 3–7.
- ³M. Radovic, E. Lara-Curzio, and L. Riester, "Comparison of different experimental techniques for determination of elastic properties of solids," *Mater. Sci. Eng. A* **368**(1–2), 56–70 (2004).
- ⁴J. D. Maynard, "The use of piezoelectric film and ultrasound resonance to determine the complete elastic tensor in one measurement," *J. Acoust. Soc. Am.* **91**(3), 1754–1762 (1992).
- ⁵A. Migliori, J. Sarrao, W. M. Visscher, T. Bell, M. Lei, Z. Fisk, and R. Leisure, "Resonant ultrasound spectroscopic techniques for measurement of the elastic moduli of solids," *Phys. B Condens. Matter* **183**(1–2), 1–24 (1993).
- ⁶A. Migliori and J. D. Maynard, "Implementation of a modern resonant ultrasound spectroscopy system for the measurement of the elastic moduli of small solid specimens," *Rev. Sci. Instrum.* **76**(12), 121301 (2005).
- ⁷J. L. Sarrao, S. R. Chen, W. M. Visscher, M. Lei, U. F. Kocks, and A. Migliori, "Determination of the crystallographic orientation of a single crystal using resonant ultrasound spectroscopy," *Rev. Sci. Instrum.* **65**(6), 2139–2140 (1994).
- ⁸R. Holland, "Resonant properties of piezoelectric ceramic rectangular parallelepipeds," *J. Acoust. Soc. Am.* **43**(5), 988–997 (1968).
- ⁹H. H. Demarest, "Cube-resonance method to determine the elastic constants of solids," *J. Acoust. Soc. Am.* **49**(3B), 768–775 (1971).
- ¹⁰I. Ohno, "Free vibration of a rectangular parallelepiped crystal and its application to determination of elastic constants of orthorhombic crystals," *J. Phys. Earth* **24**, 355–379 (1976).
- ¹¹W. M. Visscher, A. Migliori, T. M. Bell, and R. A. Reinert, "On the normal modes of free vibration of inhomogeneous and anisotropic elastic objects," *J. Acoust. Soc. Am.* **90**(4), 2154–2162 (1991).
- ¹²J. Plešek, R. Kolman, and M. Landa, "Using finite element method for the determination of elastic moduli by resonant ultrasound spectroscopy," *J. Acoust. Soc. Am.* **116**(1), 282–287 (2004).
- ¹³B. J. Zadler, J. H. L. Le Rousseau, J. A. Scales, and M. L. Smith, "Resonant ultrasound spectroscopy: Theory and application," *Geophys. J. Int.* **156**(1), 154–169 (2004).
- ¹⁴A. Tarantola, *Inverse Problem Theory and Methods for Model Parameter Estimation* (SIAM, Philadelphia, PA, 2005), 57 pp.
- ¹⁵M. C. Remillieux, T. J. Ulrich, C. Payan, J. Rivière, C. R. Lake, and P.-Y. Le Bas, "Resonant ultrasound spectroscopy for materials with high damping and samples of arbitrary geometry," *J. Geophys. Res. Solid Earth* **120**(7), 4898–4916 (2015).
- ¹⁶P. Sedláč, H. Seiner, J. Zidek, M. Janovská, and M. Landa, "Determination of all 21 independent elastic coefficients of generally anisotropic solids by resonant ultrasound spectroscopy: Benchmark examples," *Exp. Mech.* **54**(6), 1073–1085 (2014).
- ¹⁷J. Vishnuvardhan, C. V. Krishnamurthy, and K. Balasubramaniam, "Genetic algorithm reconstruction of orthotropic composite plate elastic constants from a single non-symmetric plane ultrasonic velocity data," *Compos. Part B Eng.* **38**(2), 216–227 (2007).
- ¹⁸Y. Zhang, L. O. Hall, D. B. Goldgof, and S. Sarkar, "A constrained genetic approach for computing material property of elastic objects," *IEEE Trans. Evol. Comput.* **10**(3), 341–357 (2006).
- ¹⁹H. Ogi, K. Sato, T. Asada, and M. Hirao, "Complete mode identification for resonance ultrasound spectroscopy," *J. Acoust. Soc. Am.* **112**(6), 2553–2557 (2002).
- ²⁰S. Bernard, G. Marrelec, P. Laugier, and Q. Grimal, "Bayesian normal modes identification and estimation of elastic coefficients in resonant ultrasound spectroscopy," *Inverse Probl.* **31**(6), 065010 (2015).
- ²¹M. Landa, H. Seiner, P. Sedlak, L. Bicanova, J. Zidek, and L. Heller, "Resonant ultrasound spectroscopy close to its applicability limits," in *Horizons World Physics*, Vol. 268, edited by M. Everett and L. Pedroza (Nova Science Publishers, Hauppauge, NY, 2009), Chap. 3, pp. 97–136.
- ²²ASTM E2534-15, "Standard practice for process compensated resonance testing via swept sine input for metallic and non-metallic parts," Vol. 03.03 (American Society for Testing and Materials, West Conshohocken, PA, 2015), pp. 1–8.
- ²³A. F. Bower, *Applied Mechanics of Solids* (CRC Press, Boca Raton, FL, 2009), pp. 79–81.
- ²⁴R. M. Neal, "Mcmc using Hamiltonian dynamics," in *Handbook of Markov Chain Monte Carlo*, edited by S. Brooks, A. Gelman, G. Jones, and X.-L. Meng (Chapman & Hall/CRC Press, London, 2011), Vol. 2(11), Chap. 5, <https://arxiv.org/abs/1206.1901> (Last viewed July 1, 2017).
- ²⁵S. Byrne and M. Girolami, "Geodesic Monte Carlo on embedded manifolds," *Scandinavian J. Stat.* **40**(4), 825–845 (2013).
- ²⁶Stan Development Team, *Stan Modeling Language Users Guide and Reference Manual*, Version 2.14.0 (2016), <http://mc-stan.org/>, pp. 399–415 (Last viewed August 10, 2017).
- ²⁷J. de Leeuw, "Derivatives of generalized eigen systems with applications," Department of Statistics UCLA (2007).
- ²⁸C. Zener, "Contributions to the theory of beta-phase alloys," *Phys. Rev.* **71**(12), 846–851 (1947).
- ²⁹E. S. Fisher and C. J. Renken, "Single-crystal elastic moduli and the hcp → bcc transformation in Ti, Zr, and Hf," *Phys. Rev.* **135**(2A), A482–A494 (1964).
- ³⁰R. Hill, "The elastic behaviour of a crystalline aggregate," *Proc. Phys. Soc. Sect. A* **65**(5), 349–354 (1952).
- ³¹A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian Data Analysis*, 3rd ed. (CRC Press, Boca Raton, FL, 2014), p. 297.
- ³²D. Siebörger, H. Knake, and U. Glatzel, "Temperature dependence of the elastic moduli of the nickel-base superalloy CMSX-4 and its isolated phases," *Mater. Sci. Eng. A* **298**(1–2), 26–33 (2001)6.
- ³³D. Rosca, A. Morawiec, and M. De Graef, "A new method of constructing a grid in the space of 3d rotations and its applications to texture analysis," *Model. Simul. Mater. Sci. Eng.* **22**(7), 075013 (2014).