

Elastic Properties of Novel Co- and CoNi-Based Superalloys Determined through Bayesian Inference and Resonant Ultrasound Spectroscopy



BRENT R. GOODLET, LEAH MILLS, BEN BALES, MARIE-AGATHE CHARPAGNE, SEAN P. MURRAY, WILLIAM C. LENTHE, LINDA PETZOLD, and TRESA M. POLLOCK

Bayesian inference is employed to precisely evaluate single crystal elastic properties of novel $\gamma - \gamma'$ Co- and CoNi-based superalloys from simple and non-destructive resonant ultrasound spectroscopy (RUS) measurements. Nine alloys from three Co-, CoNi-, and Ni-based alloy classes were evaluated in the fully aged condition, with one alloy per class also evaluated in the solution heat-treated condition. Comparisons are made between the elastic properties of the three alloy classes and among the alloys of a single class, with the following trends observed. A monotonic rise in the c_{44} (shear) elastic constant by a total of 12 pct is observed between the three alloy classes as Co is substituted for Ni. Elastic anisotropy (A) is also increased, with a large majority of the nearly 13 pct increase occurring after Co becomes the dominant constituent. Together the five CoNi alloys, with Co:Ni ratios from 1:1 to 1.5:1, exhibited remarkably similar properties with an average A 1.8 pct greater than the Ni-based alloy CMSX-4. Custom code demonstrating a substantial advance over previously reported methods for RUS inversion is also reported here for the first time. CmdStan-RUS is built upon the open-source probabilistic programming language of Stan and formulates the inverse problem using Bayesian methods. Bayesian posterior distributions are efficiently computed with Hamiltonian Monte Carlo (HMC), while initial parameterization is randomly generated from weakly informative prior distributions. Remarkably robust convergence behavior is demonstrated across multiple independent HMC chains in spite of initial parameterization often very far from actual parameter values. Experimental procedures are substantially simplified by allowing any arbitrary misorientation between the specimen and crystal axes, as elastic properties and misorientation are estimated simultaneously.

<https://doi.org/10.1007/s11661-018-4575-6>

© The Minerals, Metals & Materials Society and ASM International 2018

I. INTRODUCTION

PROPELLED by incredible advances in engineering and materials science—and the promise of greater turbine efficiencies and reduced emissions—the operational limits of Ni-based superalloys have been extended to regimes that include periods at over 90 pct of their melting point,^[1] all while maintaining considerable

mechanical loads in deleterious environmental conditions. Inherent limits exist for future development, as the melting temperature (or more accurately the liquidus and solidus temperatures) is intrinsic to the material and cannot be appreciably increased. This reality has spurred research and development of ordered intermetallic alloys such as NiAl,^[2] Nb-, and Mo-based refractory alloys,^[3–6] and ceramic composites of alumina^[7] and silicon carbide,^[8] all with the goal of supplanting Ni-based superalloys for the most demanding high-temperature applications. However, these alternatives often suffer from poor fracture toughness and processing constraints that make their current use costly and limited,^[3–5,9] especially considering the safety requirements for use in aerospace.

Given the payoff for increased turbine engine operating temperatures and the limitations of current alternatives to Ni-based superalloys, considerable attention has been directed toward potential intermediate

BRENT R. GOODLET, MARIE-AGATHE CHARPAGNE, SEAN P. MURRAY, WILLIAM C. LENTHE and TRESA M. POLLOCK are with the Materials Department, University of California, Santa Barbara, CA 93106, Contact e-mail: bgoodlet@engineering.ucsb.edu LEAH MILLS is with the Department of Materials Science and Engineering, The Ohio State University, Columbus, OH 43210. BEN BALES and LINDA PETZOLD are with the Department of Mechanical Engineering, University of California, Santa Barbara, CA 93106.

Manuscript submitted December 30, 2017.

solutions such as $\gamma - \gamma'$ Co-based superalloys^[10] and hybrid CoNi-based alloys.^[11] The primary motivation for studying these alloys is clear, as the melting point of Co exceeds that of Ni by 40 °C.^[12] They also have the benefit of being cast and processed similar to existing Ni-based superalloys, minimizing the cost of development by leveraging existing infrastructure. Of course there are some outstanding issues that must be addressed before these alloys may become technically relevant, including: a low γ' solvus temperature,^[13] poor oxidation resistance,^[14] and in some cases a less favorable thermal expansion coefficient for adhesion of protective alumina as compared to Ni-based superalloys.^[15] Due to the promise of future operating temperature improvements, this research focuses on measuring the single crystal elastic constants of select Co- and CoNi-based superalloys. Since turbine blades operate under nominally elastic conditions, elastic properties are critically important to design, though have not yet been measured for these new classes of Co-based materials.

In a prior publication,^[16] we demonstrated the first full Bayesian approach to determining elastic constants from simple and non-destructive resonant ultrasound spectroscopy (RUS) measurements. While simultaneously evaluating elastic constants and crystal orientation, this new approach is remarkably robust in comparison to traditional optimization algorithms only estimating elastic constants. The ability to evaluate elastic constants of a parallelepiped specimen cut with arbitrary alignment of the crystallographic axes greatly simplifies specimen preparation requirements and eliminates a major potential source of measurement error.^[16]

The goals of this paper are threefold. First, provide an introduction to RUS Bayesian inference including the basic theory, the distinguishing characteristics from traditional optimization methods, and the nomenclature necessary for discussing and achieving the remaining goals. Second, report the key features of CmdStan-RUS—a custom distribution of the open-source probabilistic programming language Stan,^[17] dedicated to RUS Bayesian inference. Third, demonstrate the capabilities of CmdStan-RUS by determining the elastic properties of nine single crystal γ' containing Co-, Ni-, and CoNi-based alloys, and through parametric studies on the effects of polynomial order, missing modes, and resonance mode selection.

II. EXPERIMENTAL METHODS

A. Materials

The discovery of a stable $\gamma - \gamma'$ two-phase field in the Co-Al-W ternary in 2006 by Sato *et al.*^[12] with morphologically identical microstructure to Ni-based $\gamma - \gamma'$ superalloys^[10] generated considerable interest from the materials science community. Expanding upon the Co-9.2Al-9W at. pct alloy reported by Sato *et al.* the first investigations on the mechanical properties of several $L1_2$ -containing Co-based alloys demonstrated high potential for turbine material application.^[11,15,18,19] The Co-based alloys were based primarily on

quaternary elemental additions to study a variety of effects including: 2 at. pct Ta for precipitate strengthening, 2 at. pct Re for matrix strengthening, 4.5 at. pct Cr for oxidation behavior, 20 at. pct Ni for expanded γ' phase field stability, and 1.5 to 6 at. pct Mo/V for density reduction.^[10] A CoNi-based alloy series explored many of the same alloying effects mentioned previously, but started from a modified base of approximately 30 at. pct Ni to the Co-Al-W ternary alloy, with variable amounts of Ta, Cr, Mo, Al, Ni, and $C + B$. Table I summarizes the chemical composition of each alloy from the previous investigations that were examined as part of this research, with the distinguishing elemental additions in bold.

1. Preparation of experimental materials

Prior to conducting any resonance measurements, the as-received material was solution heat treated and fully aged. Heat treatment schedules devised for the novel Co- and CoNi-based alloys were based on previously measured solidus and solvus temperatures, while the CMSX-4 material (a commercial Ni-based single crystal alloy included for comparison) was heat treated according to industry standards specifying a step-wise solution treatment and a two-step age.^[20] Table II summarizes the heat treatment schedules for each alloy. Solution heat treatments were conducted using a vacuum furnace with a Ti getter for at least 12 hours. Following the solution heat treatment, the material was furnace cooled except for the CMSX-4 material which was furnace quenched with gettered Ar. Co-6Ti, CoNi-A, and CMSX-4, constituting one alloy from each class, were evaluated with RUS in the solution heat-treated condition prior to aging. Then all specimens were sealed under vacuum in quartz tubes and aged in a box furnace according to Table II, followed by an air cool.

Figure 1 demonstrates the consistent microstructure achievable across the three alloy classes by comparing scanning electron microscopy (SEM) micrographs from the Co-6Ti, CoNi-A + , and CMSX-4 alloys in their fully aged conditions. Lower magnification micrographs on the left show aligned dendrites that remain from the directional solidification process, while higher magnification micrographs (center and right) show a significant volume fraction of cuboidal γ' precipitates in a matrix of γ . Microscopy specimens were prepared perpendicular to the nominal [001] growth direction using diamond suspensions down to 1 μm , then vibratory polished using alumina and colloidal silica to achieve a mirror finish. SEM characterization was completed using a FEI SEM with a field emission gun operating at 15 kV using secondary and backscatter electron detectors.

B. Resonant Ultrasound Spectroscopy

1. RUS specimen preparation

Two parallelepiped (PP) specimens for each alloy were machined from cylindrical single crystal castings *via* wire electrical discharge machining (EDM). All PP specimen dimensions were approximately the same at $9 \times 10 \times 11$ mm, though the exact dimensions of each PP varied slightly as the EDM damage layer was

Table I. Alloy Designations and Composition in Atomic Percent

Alloy	Co	Ni	Al	W	Ta	Cr	Ti	Mo	Hf	Re	C	B
CMSX-4	9.9	62.9	12.6	2.1	2.2	7.6	1.3	0.4	< 0.1	1.0	—	—
Co-Ternary	79.9	—	9.4	10.7	—	—	—	—	—	—	—	—
Co-2Ta	79.4	—	8.8	9.8	2.0	—	—	—	—	—	—	—
Co-6Ti	79.0	—	6.7	8.1	—	—	6.2	—	—	—	—	—
CoNi-A	45.9	29.2	9.8	6.3	2.4	6.4	—	—	—	—	—	—
CoNi-A +	45.9	29.2	9.8	6.3	2.4	6.3	—	—	—	—	0.14	0.02
CoNi-B	44.5	28.2	8.8	6.3	2.1	10.1	—	—	—	—	—	—
CoNi-C	38.1	38.0	9.3	6.9	1.4	6.3	—	—	—	—	—	—
CoNi-D	44.8	30.7	9.6	4.8	1.9	6.2	—	2.0	—	—	—	—

Table II. Heat Treatment Schedules for the Various Co-, Ni-, and CoNi-Based Alloys Investigated in this Study

Alloy	Solution	Age
CMSX-4	1277 °C, 4 h → 1287 °C, 2 h → 1296 °C, 3 h → 1304 °C, 2 h → 1313 °C, 2 h → 1316 °C, 2 h → 1318 °C, 2 h	1140 °C, 6 h → 871 °C, 20 h
Co-Ternary	1350 °C, 12 h	950 °C, 100 h
Co-2Ta	1350 °C, 12 h	1000 °C, 100 h
Co-6Ti	1225 °C, 14 h	1000 °C, 100 h
CoNi-based	1215 °C, 12 h	950 °C, 100 h

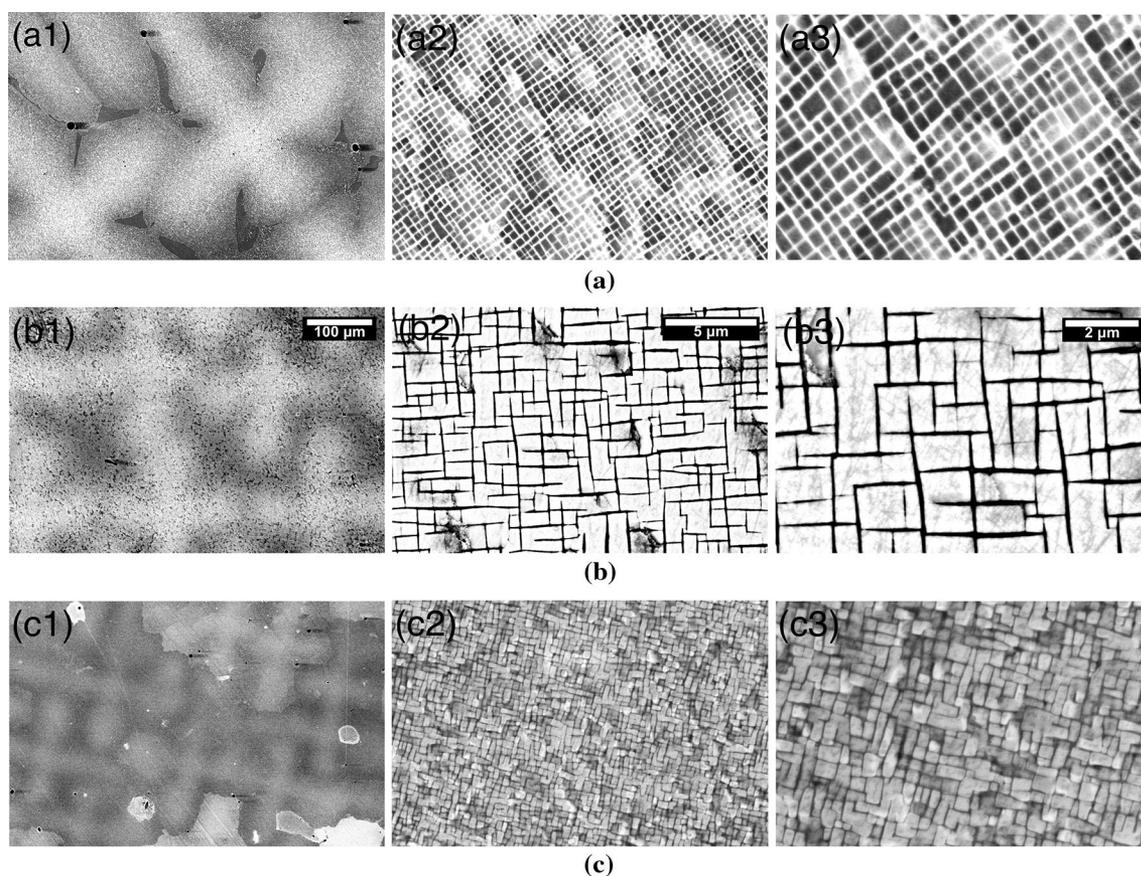


Fig. 1—Backscatter SEM micrographs of the (a) CMSX-4, (b) Co-6Ti, and (c) CoNi-A+ alloys in the solution heat-treated and aged condition. Low magnification micrographs on the left (a1, b1, c1) show the remnant dendritic microstructure, while higher magnification micrographs in the center (a2, b2, c2) and on the right (a3, b3, c3) show the cuboidal $\gamma - \gamma'$ microstructure characteristic of superalloys.

removed from each face by hand grinding with 800 grit sandpaper. Precise measurements of the specimen geometry were made with a Mitutoyo outside micrometer with 0.001 mm precision. The mass of each PP was measured with a digital Acculab scale precise to 0.0001 g, and the density was calculated. Note that the $9 \times 10 \times 11$ mm geometry was devised to both minimize the aspect ratio ensuring bending modes sensitive to non-shear moduli occur along each dimension and at relatively low frequencies, and to limit the likelihood of overlapping resonances that are more likely to occur when PP dimensions are nearly equivalent or half-integer multiples of each other.

Analysis of RUS data is substantially simplified when both the single crystal elastic constants and the arbitrary crystal reference frame rotation are determined simultaneously.^[16] Otherwise, the specimen must be precisely cut from a parent single crystal of known orientation (typically measured with X-ray diffraction techniques) such that the rotation between the crystal and specimen reference frames is nearly (ideally) zero.^[21]

2. RUS equipment

Resonance frequencies were measured with a custom experimental setup provided by the Vibrant Corporation, detailed schematically in Figure 2. This setup included a desktop computer for control and data

visualization, linked to a transceiver with three channels for operating piezoelectric transducers. Coaxial cables connect the transceiver to three identical, custom built, omni-directional piezoelectric transducers (PT) with hemispherical silicon carbide tips to protect the piezoelectric element below. The three PTs were positioned using optical table fixtures affixed to a vibration dampening breadboard, collectively referred to as the transducer cradle and pictured in Figure 3.

To excite a specimen, a swept sinusoidal frequency generated by the transceiver was transmitted to one of the PTs (noted as the drive PT in Figure 2). As the drive frequency nears a resonance frequency of the PP specimen, it begins to resonate with amplified deflections hundreds to thousands of times greater than the displacements produced by the drive transducer. These amplified deflections were then registered by the two additional receive PTs contacting the specimen, with the signal returned to the transceiver and visualized through the computer controller to produce a broadband resonance spectra such as those depicted at the bottom of Figure 2. While the signal amplitudes are generally unreliable given the transducer cradle configuration that allows the specimen to freely lift off of the PTs during a measurement, the frequency values are reliably measured to better than 0.02 pct precision.

The transducer cradle was designed to fully support the PP specimen along three faces at roughly 90-degree angles and near the edge of the PP. Note that the corners of a PP specimen are the points of lowest symmetry^[22–25] and provide the lowest elastic coupling/impedance^[22,26] for minimal biasing of the measured resonance frequencies. For the three transducer cradle configuration described above, contacting the PP specimen on its faces near the corners was found to provide sufficient signal-to-noise ratio with minimal variance in frequency to allow for precise estimates of elastic constants.

3. RUS measurements

All RUS measurements were completed at standard temperature and pressure. To ensure that all of the PP specimen resonance modes were measured, each PP was scanned five times across the range of 70 to 445 kHz, with the specimen removed, rotated, and then replaced on the transducer cradle after each scan. While this procedure occasionally resulted in a single mode being missed from a broadband scan, each of the first 70 resonance modes were effectively measured for all of the PP specimens investigated in this study.

Distilling a list of resonance frequencies from each broadband RUS spectrum was partially automated based on peak amplitude, peak width, separation from neighboring peaks, and the second derivative of the amplitude with respect to frequency. Each of these parameters were set for individual bands (approximately 20 kHz wide) such that the automated process reliably identified most modes. Each scan was then individually verified for any missing or spuriously identified modes before the frequency list was exported from the collection computer for further analysis. Next, an averaged list of frequencies was created for each PP from the five individual frequency lists, inspecting the standard

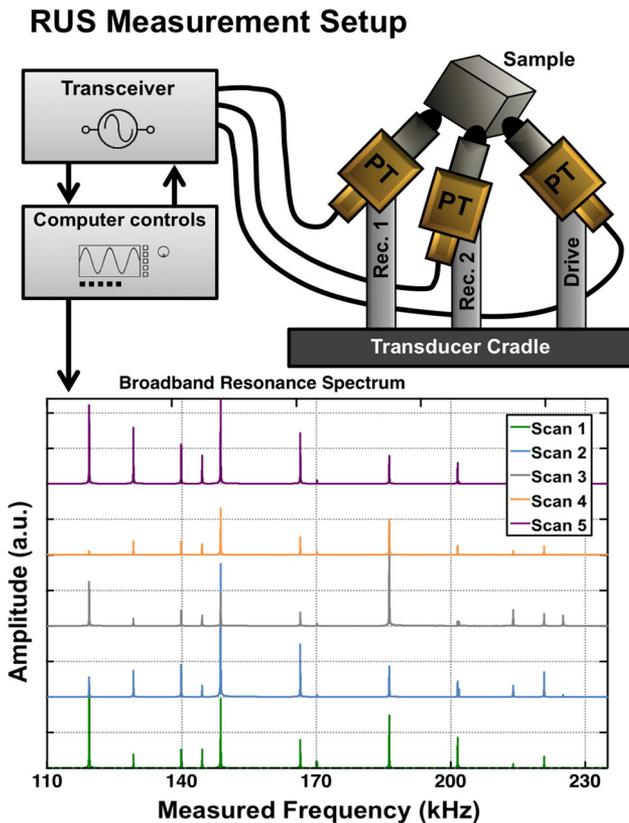


Fig. 2—Diagram of RUS measurement setup and broadband RUS spectrum plots from 5 measurements of the Co-Ternary-B specimen. The spectrum plots are vertically offset to demonstrate the repeatability of the frequency measurement, while the amplitudes are not repeatable and are thus plotted in arbitrary units (a.u.).

deviation across the five scans was within 0.1 kHz to assure that no modes were missed or erroneously recorded. An example mode list for PP specimen CoNi-A+1 (referring to CoNi-A+ PP 1, of 2) is provided in Table III, which clearly demonstrates that measurement precision of ± 0.02 kHz (or less) is reliably achieved. This measurement precision is remarkable given that a single broadband scan of 70 resonance modes is collected in less than 10 minutes. Such reliable precision is what makes RUS an ideal measurement technique for estimating elastic constants.^[21]



Fig. 3—Photo of a parallelepiped specimen positioned on the piezoelectric transducer cradle prior to a broadband RUS measurement.

III. COMPUTATIONAL METHODS

As there is no direct solution to the general inverse problem of determining elastic properties from RUS measurements, elastic properties must be determined indirectly through least-squares optimization^[21,27,28] or statistical inference^[16] methods. Both of these inversion methods ultimately rely on numerous forward problem evaluations,^[16,21] therefore, computational methods begins by introducing the forward problem for calculating resonance frequencies of a specimen in Section III-A. Next, Section III-B details inversion with a short summary on traditional optimization-based methods followed by the distinguishing characteristics of a Bayesian approach. The statistical RUS model is formulated in Section III-D, which details the Bayesian posterior and necessary assumptions. Computation of the posterior with Hamiltonian Monte Carlo (HMC)^[29] is outlined in Section III-E, with example traceplots and initial parameterization discussed in Section III-E-1. Finally, methods of validation are discussed in Section III-E-3, where the concept of posterior predictive distributions is introduced, followed by details regarding posterior predictive plots and how they can be used to diagnose some of the more common problems that may arise during RUS Bayesian inference.

A. The Forward Problem

Visscher's *xyz* algorithm is used to efficiently solve the forward problem by calculating resonance modes of a parallelepiped specimen given its geometry, density, and effective elastic constants.^[30] The effective elastic

Table III. Mean and Standard Deviation (SD) for the First 50 Resonance Modes of CoNi-A + 1, Based on five Broadband RUS Scans

Mode No.	Frequency (kHz) Mean \pm SD	Mode No.	Frequency (kHz) Mean \pm SD
1	135.618 \pm 0.016	26	286.690 \pm 0.013
2	142.851 \pm 0.010	27	290.809 \pm 0.002
3	144.923 \pm 0.010	28	299.620 \pm 0.011
4	154.308 \pm 0.018	29	304.619 \pm 0.013
5	163.844 \pm 0.013	30	308.081 \pm 0.012
6	164.111 \pm 0.008	31	310.982 \pm 0.012
7	168.193 \pm 0.014	32	320.033 \pm 0.015
8	188.061 \pm 0.011	33	321.824 \pm 0.011
9	202.506 \pm 0.010	34	326.095 \pm 0.009
10	204.747 \pm 0.009	35	332.206 \pm 0.013
11	210.693 \pm 0.011	36	334.222 \pm 0.013
12	211.414 \pm 0.012	37	337.300 \pm 0.012
13	211.985 \pm 0.010	38	341.681 \pm 0.011
14	219.423 \pm 0.011	39	344.413 \pm 0.014
15	220.491 \pm 0.014	40	348.664 \pm 0.016
16	237.149 \pm 0.012	41	354.642 \pm 0.012
17	241.558 \pm 0.010	42	355.491 \pm 0.012
18	244.404 \pm 0.016	43	363.661 \pm 0.012
19	246.905 \pm 0.009	44	367.021 \pm 0.012
20	248.223 \pm 0.013	45	367.950 \pm 0.018
21	259.814 \pm 0.014	46	370.402 \pm 0.013
22	266.153 \pm 0.01	47	373.411 \pm 0.012
23	273.195 \pm 0.016	48	379.083 \pm 0.013
24	275.341 \pm 0.013	49	381.878 \pm 0.008
25	285.808 \pm 0.010	50	383.998 \pm 0.011

constants are themselves a product of the single crystal elastic constants and an arbitrary rotation of the crystal reference frame with respect to the specimen reference frame, parameterized in cubochoric coordinates^[31] as addressed in Section III–A–1. The xyz algorithm is a Rayleigh–Ritz method that is named after the simple polynomial basis functions that describe the displacements of the specimen as a product of powers of the Cartesian coordinates with the form: $x^l y^m z^n$.^[30] The powers l , m , and n correspond to non-negative integer values and $l + m + n \leq P$ defines the maximum polynomial order (P). With the size of the forward calculation governed by P , one must be cognizant of the fact that the matrix calculations grow by $S = (P + 1)(P + 2)(P + 3)/6$ such that typical values of $P = 8, 10, 12$ result in matrices of order $S = 165, 286, 455$.

The CmdStan-RUS code is currently limited to a parallelepiped geometry, though Visscher’s polynomial basis functions can be used for a variety of geometries including spheres, hemispheres, spheroids, ellipsoids, cylinders, bells, sandwiches, cones, pyramids, prisms, tetrahedra, octahedra, and others.^[30] However, the parallelepiped geometry is optimal for a variety of reasons including ease of specimen machining and a lack of rotational symmetry that could hinder efforts to estimate crystal orientation. When RUS experimental methods are attentive to minimizing error in the frequency measurements, uncertainty in the geometry of the specimen becomes a dominant source of measurement error that can inhibit precise parameter estimates.

1. Orientations

In Bales *et al.*,^[16] the arbitrary rotation of the crystal reference frame was parameterized as a passive unit quaternion. A unit quaternion is expressed in four dimensions (w, x, y, z) but resides on a 3D manifold characterized by $w^2 + x^2 + y^2 + z^2 = 1$. While parameterization of the rotation as a unit quaternion confounded the default HMC sampler in Stan,^[17] the No U-Turn Sampler,^[32] alternative parameterization with cubochoric coordinates resolved this issue. Cubochoric coordinates are a volume-preserving projection of a regularly gridded cube of side length $\pi^{\frac{2}{3}}$ into a unit ball, followed a homochoric mapping from the ball to a unit quaternion.^[31] Thus, the three components of a cubochoric rotation are valid between the range $-\frac{1}{2}\pi^{\frac{2}{3}}$ to $\frac{1}{2}\pi^{\frac{2}{3}}$.

B. Inversion

Inversion of elastic properties from RUS measured resonance frequencies has historically been accomplished using optimization algorithms that yield maximum likelihood estimates (MLE) of the unknown parameters. Some of the more popular optimization algorithms include the Levenberg–Marquardt (Newton-steepest decent) algorithm popularized by Migliori *et al.*,^[21,22] the derivativeless fixed point iteration (FPI)

method used by Plešek *et al.*,^[33] for finite element (FE) implementations where derivative information is less accessible, and genetic algorithms as used by Remillieux *et al.*^[34] Each of these algorithms perform a least-squares regression of an overdetermined system, with more resonance modes being used than parameters being estimated, by iteratively updating an initial guess set of elastic constants to minimize the sum of squared errors (SSE) between a measured and computed list of resonance frequencies.^[21,22,33,34] Various convergence criteria are then used to terminate the optimization, typically when successive iterations fail to reduce the SSE or after a desired precision is reached.

Instead of pursuing MLEs (point estimates) of unknown parameters by way of optimization, here we formulate the inverse problem as a Bayesian inference which yields probability distributions for the unknown parameters (elastic constants, crystal orientation, *etc.*). When the parameter distributions are normally distributed—a characteristic typically verified a posteriori by fitting the distribution as a Gaussian—the mean and the variance of the distribution provide estimates of the parameter value and its uncertainty. While estimating error in the elastic constants from optimization methods is far less straightforward, Migliori and Sarrao^[21] discuss an empirical approach to determining uncertainty by exploring the curvature of the (SSE) error surface in the vicinity of the MLE. But this approach only estimates the error in linear combinations of parameters that may be difficult to isolate when the curvature of the error surface is shallow. Confronting these same issues, Spoor^[27] employed Monte Carlo methods to estimate parameter uncertainty in a manner similar to this work by assigning normally distributed random perturbations to his measured resonance modes, then characterizing the distribution of the MLEs to infer uncertainty.^[27] Both are meaningful endeavors, but involve complicated additional steps, while uncertainty in the Bayesian estimates are obtained automatically from the posterior distributions.

C. Scaled and Transformed Parameters

For improved numerical stability and HMC sampling, all single crystal elastic constants are scaled by a common factor of $10^{-11}/Pa$ such that a value of 200 GPa ($2.0 \times 10^{11} Pa$), is handled by Stan as a scaled unitless parameter: 2.0. Conveniently, the σ (noise) parameter (in units of kHz) and the three cubochoric orientation parameters did not require scaling. Note that the traceplots and histograms present parameters as they are handled by CmdStan-RUS, while common units of GPa and kHz are used for tables and result summaries.

In addition to scaling parameters, transformed parameters can improve HMC mobility and efficiency. An example of a transformed parameter is provided by first considering the fully anisotropic stiffness tensor in Voigt notation:

$$C_{ij} = \begin{bmatrix} c_{11} & c_{12} & c_{13} & c_{14} & c_{15} & c_{16} \\ c_{12} & c_{22} & c_{23} & c_{24} & c_{25} & c_{26} \\ c_{13} & c_{23} & c_{33} & c_{34} & c_{35} & c_{36} \\ c_{14} & c_{24} & c_{34} & c_{44} & c_{45} & c_{46} \\ c_{15} & c_{25} & c_{35} & c_{45} & c_{55} & c_{56} \\ c_{16} & c_{26} & c_{36} & c_{46} & c_{56} & c_{66} \end{bmatrix}. \quad [1]$$

As the matrix is symmetric, there are 21 independent elastic constants which can be simplified to three independent constants through enforcing cubic elastic symmetry:

$$C_{ij} = \begin{bmatrix} c_{11} & c_{12} & c_{12} & 0 & 0 & 0 \\ c_{12} & c_{11} & c_{12} & 0 & 0 & 0 \\ c_{12} & c_{12} & c_{11} & 0 & 0 & 0 \\ 0 & 0 & 0 & c_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & c_{44} & 0 \\ 0 & 0 & 0 & 0 & 0 & c_{44} \end{bmatrix}. \quad [2]$$

Superalloys cast as single crystals, with a coherent two-phase microstructure possessing a cube-on-cube orientation relationship between the L_{12} precipitates and the FCC matrix,^[1] exhibit cubic elastic symmetry that appears homogeneous on the macroscopic scale. This apparent homogeneity occurs whenever the microstructure is orders of magnitude smaller than the wavelength of the resonance modes. In previous work,^[16] it was observed that the (c_{11}, c_{12}, c_{44}) elastic parameter space led to inferior sampling as the c_{11} and c_{12} parameters were highly correlated. Therefore, the c_{12} parameter was transformed into the Zener anisotropy ratio:^[35]

$$A = \frac{2c_{44}}{c_{11} - c_{12}}, \quad [3]$$

such that sampling now takes place over a transformed (c_{11}, A, c_{44}) elastic parameter space.^[16] The present investigation also used the (c_{11}, A, c_{44}) parameter space, while any alternative elastic symmetry with arbitrary parameter transformations is possible through simple modifications of the Stan input files.

D. Building a Statistical RUS Model

Bayes' rule can be presented in its common form:^[36]

$$\underbrace{P(\theta|M)}_{\text{Posterior}} = \frac{\overbrace{P(M|\theta)}^{\text{Likelihood}} \overbrace{P(\theta)}^{\text{Prior}}}{\underbrace{P(M)}_{\text{Prior/Predictive}}}, \quad [4]$$

where statements of conditional probability $P(\dots|\dots)$, measured data M , and estimated parameters θ , are combined to read " $P(M|\theta)$ " as "the probability of the data given the parameters."^[37]

To perform a Bayesian inference, a statistical model representing the likelihood is first developed and initialized based on prior knowledge and some observed data.

Then new data in the form of previous posterior samples are used to update the prior knowledge in a systematic manner and to inform future posterior evaluations. This iterative learning process minimizes the uncertainty of estimated parameters with regard to the information available and may ultimately allow for precise estimates given enough posterior samples are available and the statistical model accurately describes the problem.

The statistical RUS model employed in this work is described in more detail elsewhere,^[16] so only a cursory description highlighting key elements and assumptions will be made here. The first assumption is that resonance modes $(m_1, m_2, m_3, \dots, m_N)$, in a frequency ordered list of modes M with length N can be modeled as

$$M = m_1, m_2, \dots, m_N = f(L, C, Cu, \rho) + \xi. \quad [5]$$

The forward model $f(\dots)$ is a function of the properties of the modeled specimen (L, C, Cu, ρ) and a normally distributed random noise ξ . Concerning the terms of the forward model, $L = (l_1, l_2, l_3)$ represents the dimensions of the parallelepiped, $C = (c_{11}, c_{12}, c_{44})$ are the single crystal elastic constants (with cubic elastic symmetry assumed for this example), $Cu = (cu_1, cu_2, cu_3)$ are the cubochoric crystal orientation coordinates, and ρ is the mass density of the specimen; collectively expressed as $\theta = L, C, Cu, \rho$ for short. ξ is assumed normally distributed about each mode with a single variance σ^2 (i.e., ξ is the same for every mode) and represents the combined uncertainty in the model and the measurements.^[38] While it is unlikely that the noise is actually the same for every mode, Section III-E-3 discusses a posteriori methods by which this and other assumptions are deemed acceptable.

The formulation of Eq. [5] means that the resonance modes produced by either the forward model or by measurement are themselves random variables.^[16,38] Thus, a statement of the conditional probability of measuring a set of modes M from a specimen with a fixed set of parameters θ may be written in terms of the forward model f as

$$P(M|\theta) \sim \mathcal{N}(f(\theta), \sigma), \quad [6]$$

with the tilde (\sim) indicating that the statement $P(M|\theta)$ "has the probability distribution of" and \mathcal{N} indicating "a normal distribution" of modes with variance σ^2 about $f(\theta)$. In terms of a Bayesian formulation (Eq. [4]), $P(M|\theta)$ is the likelihood. With a rigorous treatment of the mathematics and assumptions, a complete statement of the likelihood can be given as

$$P(M|\theta) = \prod_i \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{-\frac{(f_i(\theta) - M_i)^2}{2\sigma_i^2}}, \quad [7]$$

where i is the index of the computed $(f_i(\theta))$ and measured (M_i) modes.^[16]

Prior probability distributions $(P(\theta))$ in Eq. [4] express prior beliefs about the parameters that can be based on values cited in literature, previous measurements on the same or similar material, or even based on an intuitive understanding of the natural range a parameter may exhibit. Importantly, priors express

knowledge or beliefs about ultimately unknown quantities, they can be estimated from the data itself, and they do not necessarily precede the data in time.^[39] Prior probability distributions can be simple expressions like assuming the c_{11} elastic constant ranges between 0 and 500 GPa with uniform probability:

$$c_{11} \sim \mathcal{U}(0 \text{ GPa}, 500 \text{ GPa}) \quad [8]$$

Or more informed statements like c_{11} being normally distributed about a mean of 250 GPa with a standard deviation of 50 GPa:

$$c_{11} \sim \mathcal{N}(250 \text{ GPa}, 50 \text{ GPa}) \quad [9]$$

Alternative distributions are also available including exponential, Cauchy, Student's-T, and others.^[17,40] With linear combinations of parameters implemented simply as

$$c_{11} - c_{12} \sim \mathcal{N}(100 \text{ GPa}, 50 \text{ GPa}). \quad [10]$$

All of the inferences in this work assumed the same set of weakly informative priors as summarized below:

$$\begin{aligned} \sigma &\sim \mathcal{N}(0 \text{ kHz}, 1.0 \text{ kHz}) \\ c_{11} &\sim \mathcal{N}(250 \text{ GPa}, 75 \text{ GPa}) \\ c_{44} &\sim \mathcal{N}(125 \text{ GPa}, 50 \text{ GPa}) \\ A &\sim \mathcal{N}(3.0, 0.75) \\ c_{11} - c_{44} &\sim \mathcal{N}(125 \text{ GPa}, 50 \text{ GPa}). \end{aligned}$$

Ultimately, we want to find the posterior distributions that express the probability that of our parameters θ take certain values given a measured set of modes M . We now have the likelihood and the priors, but the prior predictive distributions $P(M)$ are unknown. Fortunately, Bayes' rule can be simplified to

$$P(\theta|M) \propto P(M|\theta)P(\theta) \quad [11]$$

when sampling with Markov chain Monte Carlo (MCMC) methods^[41] because posterior samples are generated proportional to the true posterior distribution.^[17] Now with a complete formulation of the inverse problem, the Bayesian posterior is computed with HMC.

E. Computing the Posterior with HMC

Hamiltonian Monte Carlo is a hybrid Markov chain Monte Carlo (MCMC) method^[17,29,42] that avoids the random walk exploration behavior common to other MC methods^[43] by informing the sampler with first-order gradient information.^[32,42] As for the derivatives necessary for Hamiltonian dynamics, they were detailed in our previous work Bales *et al.*,^[16] and are now fully handled within CmdStan-RUS. Like other MCMC methods, HMC involves sequentially building a Markov chain of samples based on memoryless exploration from the current state, with distributions proportional to the true posterior distribution.^[17] With zero consideration

given to previously accepted states, a proposed state at the end of an exploration period is only added to the chain when it is accepted through a Metropolis update based on the relative probability of the proposed state with respect to the current.^[29] The Metropolis update accepts proposed states evaluated with a higher probability than the current state, but may reject a proposed state of lower probability in favor of retaining the current state.

Hamiltonian dynamics provide superior sampling by proposing distant and less correlated samples to the Metropolis update, and are accomplished through introduction of a fictitious momentum term at the beginning of an exploration period.^[29,32] This momentum term can assist in propelling the sampler beyond local minima and bring the sampler towards the steady-state region of the posterior. When a proposed state is rejected, a new momentum term is generated and the exploration is restarted from the retained state.^[29]

1. Initial parameterization and constraints

Initial parameterization of the HMC chains is randomly generated from distributions specified by the priors given in Section III-D. The traceplots in Figure 4 show the first 100 samples for four independent HMC chains of the CmdStan-RUS inversion of CoNi-C1 (the first parallelepiped of alloy CoNi-C). Clearly, the random initial parameterization is only weakly informed by the given priors. For example, the c_{11} parameter initial values range from over 650 GPa for chain 2, to less than 50 GPa for chains 1 and 3. These are by all means "poor initial guess values," from the perspective of RUS inversion literature that consistently cites "good initial guess value" as important for ensuring convergence.^[21,26,27,33,34,44] In fact, the robust convergence behavior of the RUS Bayesian inference framework, despite poor initial values and no efforts to identify measured resonance modes, is one of the key contributions of this work.

As the traceplots in Figure 4 show, the four HMC chains reach the stationary region of the posterior relatively early (in this case, after less than 100 samples) in the 500 sample warm-up period—despite initial parameterization in extremely low probability regions of the posterior. Only the samples generated after the warm-up period are used to characterize the posterior distributions, so the rate at which the stationary region of the posterior is reached is a secondary concern. Each chain is run for a total of 1500 samples, with the first 500 warm-up samples ultimately discarded, a chain that reaches the stationary region before the warm-up period is complete is deemed a success.

The rapidly converging behavior detailed by Figure 4 is not always achieved, though it is by far the most common occurrence. Upwards of 90 pct of the individual chains reached a common stationary region of the posterior within the first half of the warm-up period. Of the approximately 10 pct of chains that fail to reach the stationary region, the vast majority of these cases were

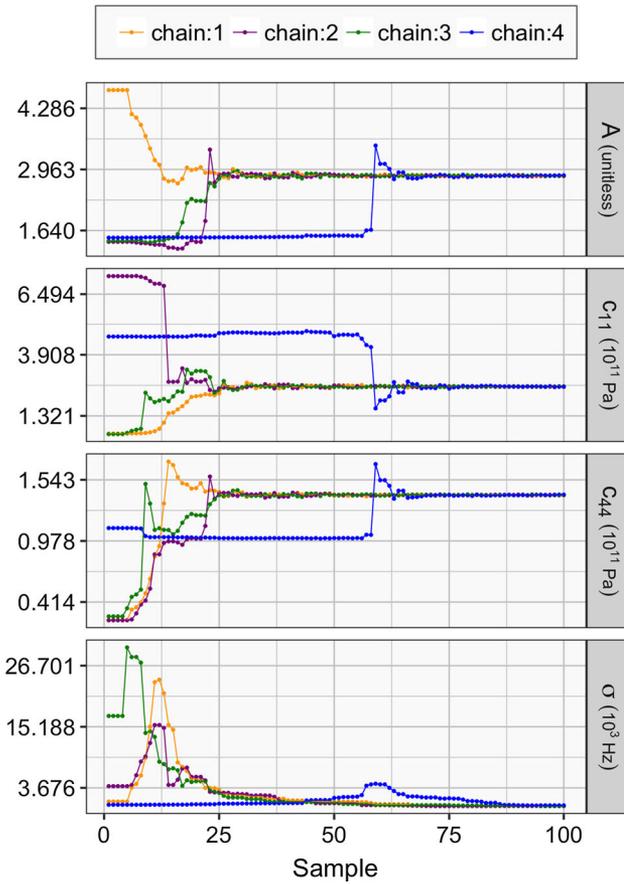


Fig. 4—Traceplots of the first 100 warm-up samples of four independent HMC chains for the CoNi-C1 specimen demonstrating robust convergence to the stationary region of the posterior from remarkably distant, randomly generated, initial conditions.

due to the sampler stalling as it asymptotically approached a constraint specified when parameters are first defined in CmdStan-RUS. Below is a list of constraints applied:

$$\begin{aligned} \sigma, c_{11}, c_{12}, c_{44} &= (0, \infty) \\ A &= (1, \infty) \\ c_{44}/c_{11} &= (0, 3/2) \\ cu_1, cu_2, cu_3 &= (-1.072515, 1.072515). \end{aligned}$$

While all of these constraints are not strictly forbidden, *e.g.*, c_{12} can be negative,^[45] the stiffness matrix (Eq. [2]) must be positive definite which restricts c_{11} , c_{44} , and all constants along the diagonal to be positive. Instead, these constraints were developed through trial and error to minimally restrict the sampler while also mitigating the occurrence of initial parameterizations that would confound the sampler. Fewer constraints afford fewer opportunities for the sampler to find local regions of parameter space where the posterior probability increases as a constraint is approached. For these reasons, it is recommended that implausible yet physically permissible values not be restricted through constraints, but instead discouraged through specifying priors whenever possible as detailed by Section III-D.

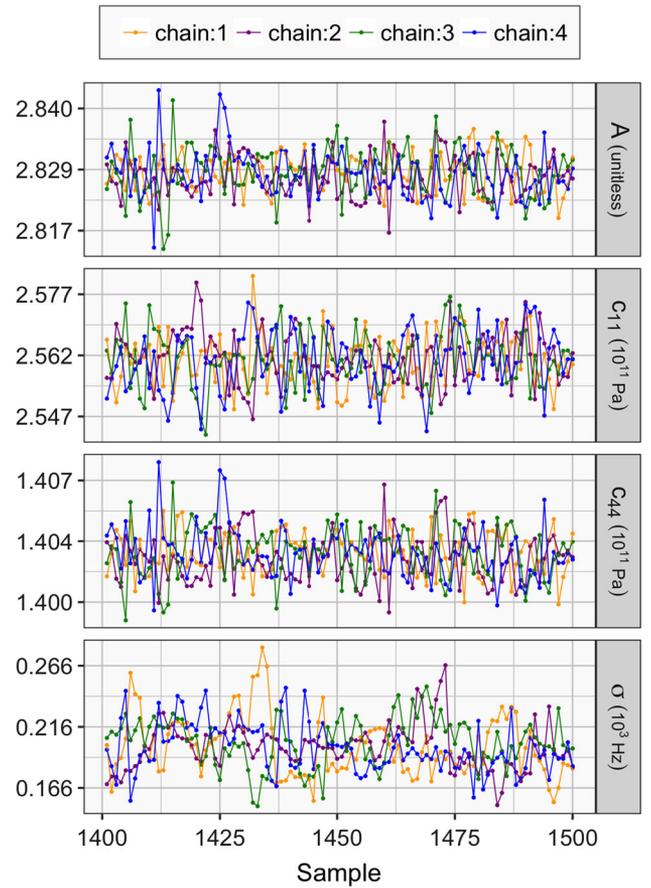


Fig. 5—Traceplots of the last 100 sampling samples of four independent HMC chains for the CoNi-C1 specimen demonstrating desirable uncorrelated sampling behavior of the stationary region of the posterior for the noise and three elastic parameters.

2. Verifying sampling behavior

The full scales on Figure 4 do not allow one to observe the sampling behavior in the stationary region of the posterior. Instead, a detailed view of the last 100 samples of the CoNi-C1 inversion is provided in Figure 5. Here, one can verify quality sampling by observing the following: sufficient mixing as each chain freely explores the posterior, no obvious trends (*i.e.*, stationary), no multimodal tendencies, and mostly uncorrelated behavior between subsequent samples of a given chain. Of the four parameters plotted, the noise parameter σ is clearly being sampled with the least efficiency as it exhibits some correlated sampling behavior. Overall, the traceplots provided in Figures 4 and 5 provide great confidence in our efforts to use HMC to calculate the posterior distributions, with further measures for validation discussed in Section III-E-3.

After verifying quality sampling, all samples generated after the warm-up period across the various HMC chains are combined to generate histograms for visual summary of the results. Figure 6 provides a set of histograms from specimen CoNi-B2 that contain 4000 samples, from four combined chains of 1000 samples each, that are typical of the results collected. Fitting with a Gaussian distribution confirms the normal

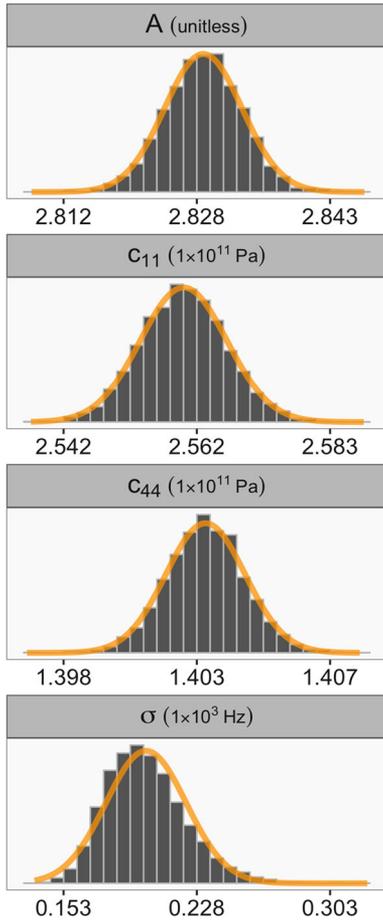


Fig. 6—Histogram populated with 4000 samples (four chains of 1000 samples each) for the CoNi-B2 specimen. Gaussian distributions overlay the histograms to demonstrate the parameters are normally distributed and justify describing the distributions with a mean and standard deviation.

behavior of the parameter distributions around a single mode, and allows for the distributions to be characterized by their mean and standard deviation (SD). With SD being a direct estimate of the uncertainty, it can be reduced through the generation of additional HMC samples according to $1/\sqrt{N_{eff}}$, where N_{eff} is the number of effective samples.^[17]

3. Posterior predictive distributions and validation

A check on the posterior predictive distributions as detailed in this section provides essential information for validating results. Posterior predictive distributions are developed by taking posterior samples of the estimated parameters and generating new resonance modes that reflect the implications of the uncertainty in the parameter estimates. Upon inspection, 95 pct of the measured modes should fall within the 95 pct intervals, otherwise this suggests that there are likely outliers in the measured data or the model is ill-defined.^[16]

As a visual aid, posterior predictive plots like those shown in Figure 7 are generated. Both plots show the difference between the posterior predictive and the measured mode frequencies, designated as the error on the abscissa, for the first 40 modes of specimen

Co-Ternary-B. The length of the bars about each point indicates the 95 pct intervals. With units of kHz, this interval is not the mode invariant noise parameter σ , instead it represents the frequency range about which 95 pct of resonance mode frequencies from the generated posterior predictive distributions exist. With only a 5 pct chance that a measured mode and a posterior predictive mode would differ by more than the interval indicated, a posterior predictive plot with 40 modes might be expected to exhibit 2 modes with an error beyond the 95 pct interval. The difference between the plots in Figure 7 is the polynomial order used for the forward calculation, with $P = 10$ for Figure 7(a) and $P = 12$ for Figure 7(b).

For validation, one should inspect the posterior predictive plot to ensure that there is neither an observable trend in the error, nor any discernible patterns or structure that may indicate an ill-defined model. A positive trend in the error of Figure 7(a) is apparent by the first 12 modes exhibiting a negative error, while six of the last 7 modes are above the zero line. Additionally, the last two modes of Figure 7(a) lie unusually far beyond the 95 pct interval, suggesting either the noise is not independent and identically distributed, or there is a problem with the forward calculation. The fact that the unusual behavior exhibited by modes 39 and 40 in Figure 7(a) is not observed in Figure 7(b) is a clear sign that numerical imprecision of the forward calculation is the root cause, and may lead to biased inference results. For these reasons, the $P = 10$ results are discounted in favor of the $P = 12$ results. Although it is worth noting that the $P = 10$ inference ran for 20.3 hours, while the $P = 12$ inference ran for 60.9 hours. For situations where computational time is of concern, using fewer resonance modes (N) would allow a smaller P without introducing numerical error into the results.

IV. RESULTS AND DISCUSSION

Table IV summarizes the elastic parameter estimates from two parallelepiped specimens (A and B or 1 and 2) of each alloy composition, measured at room temperature, in the solution heat-treated and fully aged condition. Very good agreement is observed between the specimen pairs, with the largest deviation for Co-6Ti. The estimates of the anisotropy ratio (A) for Co-6Ti-A and Co-6Ti-B vary by about 1 pct, and as a result the estimates of c_{12} (which was not estimated directly) vary by just over 2 pct.

A. Influence of Alloy Chemistry on Elastic Constants

First consider a comparison between the five CoNi-based alloys in Table IV. CoNi-C exhibits the lowest anisotropy of the class at 2.83, while CoNi-A+ exhibits the largest A of 2.88. The c_{44} value of CoNi-C is also the smallest of the CoNi-based alloys, about 1 pct smaller than the largest c_{44} again belonging to CoNi-A+. CoNi-C is compositionally the most distinct CoNi alloy with 6.6–11.9 at. pct less Co and 7.3–9.8 at. pct more Ni

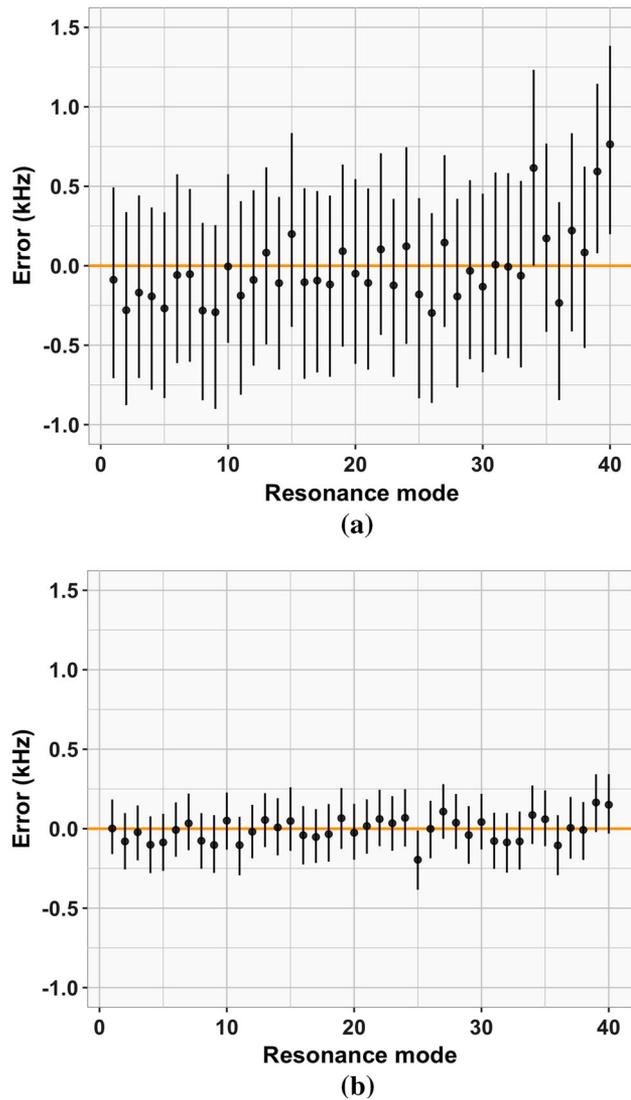


Fig. 7—Co-Ternary-B posterior predictive plots demonstrating (a) undesirable trends in the error when $P = 10$, *i.e.*, polynomial order is insufficiently small, and how (b) using $P = 12$ resolves this issue.

likely contributing to its lower c_{12} , c_{44} , and A values. Interestingly, the c_{11} constant barely varies beyond the approximately 1 GPa uncertainty for the 5 CoNi alloys. Overall, the data in Table IV indicate that all of the CoNi alloys are elastically quite similar, while even moderate alloying additions within the CoNi alloy class have only a small effect on elastic constants.

The Co-based alloy class offers an interesting opportunity to compare three alloys with nearly constant Co concentrations, which appeared to be a significant factor in CoNi-based alloys discussed previously. Overall, the elastic constants of the Co-based alloys differ from each other by a greater extent than the CoNi-based alloys, indicating that Co-based alloy elastic constants are more dependent on small to moderate compositional differences than their CoNi-based counterparts. Interestingly, the Co-6Ti alloy was the least stiff of the three Co-based alloys, exhibiting smaller c_{11} , c_{12} , and c_{44} elastic constants.

The Co-Ternary alloy exhibited the greatest c_{44} and A , while Co-2Ta was marginally stiffer in the c_{11} . With a range varying by 3.3 pct, the c_{44} parameter changed the most as a result of alloying, followed by the anisotropy ratio varying by 2.2 pct among the Co-based alloys.

One of the greatest differences between the moduli of the various alloy classes is the anisotropy ratio exhibited by the Co-based alloys. A increased from a low of about 2.81 for the Ni-based alloy CMSX-4 to around 2.86 for the CoNi-based alloys, or approximately + 1.8 pct. A more dramatic increase was observed between the CoNi-based alloys and the Co-based alloys as A increased by an additional 11 pct, from an average of 2.86 to an average of 3.17. With a nearly negligible increase between CMSX-4 and the CoNi alloy class, the c_{12} elastic constant exhibited a much more significant increase of about 4 to 6 GPa (+ 3 pct) between the CoNi- and the Co-based alloy classes.

Following a similar progression from lowest to highest Co concentration, the c_{44} modulus is observed to consistently increase with increasing Co. Except now, the most dramatic increase in c_{44} is between the Ni-based alloy and the CoNi alloys. A smaller, yet still significant, increase exists between the CoNi- and the Co-based alloy classes. Quantifying this, the c_{44} modulus increased by 10 GPa (+ 8 pct) from CMSX-4 to the CoNi alloy class, before increasing by a more modest 4 to 8 GPa from the CoNi- to the Co-based alloy class.

Initially, c_{11} follows a similar trend with increasing Co, with an increase of about 7 to 8 GPa (or about 3 pct) between CMSX-4 and the average CoNi-based alloy. However, the c_{11} elastic constant fails to increase much beyond the 256 to 257 GPa range exhibited by the CoNi alloy class. With approximately 35 pct more Co than the average CoNi alloy: Co-2Ta exhibits a c_{11} that is only 2 to 3 GPa higher, Co-Ternary is basically equivalent, while Co-6Ti exhibits a c_{11} that is 7 GPa lower than the average CoNi alloy.

B. Effect of Heat Treatments on Elastic Properties

In addition to the results summarized in Table IV, one alloy from each class was evaluated in the solution heat-treated condition to investigate the potential role of changing volume fractions of the FCC and $L1_2$ phases. Specifically, two parallelepipeds of CMSX-4, Co-6Ti, and CoNi-A were solution heat treated, then measured with RUS, aged, and finally measured with RUS again. The results of this heat treatment study are summarized in Table V, with the heat treatment schedules summarized previously in Table II. Clearly, the thermal history does not appreciably affect the elastic properties of the alloys investigated. While a similar result is reported by Parsa *et al.*^[46] who used RUS to study the effect of heat treatments on the elastic properties of a single crystal Ni-based superalloy with a composition similar to that of CMSX-4. These results suggest that the room temperature elastic constants of the ordered and disordered phases of the Co- and CoNi-based alloys are not substantially different from each other, as has been previously reported for Ni-based alloys.^[47,48]

Table IV. Summary of the Best Inferences of Elastic Parameters for Two Parallelepiped Specimens of Each Alloy

Specimen Name	c_{11} (GPa) Mean \pm SD	c_{12} (GPa) Mean \pm SD	c_{44} (GPa) Mean \pm SD	A (unitless) Mean \pm SD
CMSX-4-A	249.0 \pm 2.3	157.0 \pm 2.4	129.2 \pm 0.29	2.811 \pm 0.011
CMSX-4-B	248.3 \pm 2.0	156.3 \pm 2.1	129.6 \pm 0.28	2.817 \pm 0.0097
Co-Ternary-A	257.3 \pm 0.79	164.2 \pm 0.83	149.9 \pm 0.13	3.219 \pm 0.0045
Co-Ternary-B	256.7 \pm 0.58	163.4 \pm 0.60	150.0 \pm 0.10	3.217 \pm 0.0035
Co-2Ta-A	258.8 \pm 1.8	164.5 \pm 1.9	147.7 \pm 0.26	3.134 \pm 0.0085
Co-2Ta-B	259.3 \pm 1.6	165.1 \pm 1.7	147.9 \pm 0.23	3.141 \pm 0.0077
Co-6Ti-A	248.1 \pm 1.4	156.8 \pm 1.4	145.4 \pm 0.24	3.188 \pm 0.0086
Co-6Ti-B	250.7 \pm 1.8	160.2 \pm 1.9	145.2 \pm 0.23	3.211 \pm 0.0086
CoNi-A1	256.7 \pm 0.36	157.6 \pm 0.37	142.1 \pm 0.08	2.868 \pm 0.0029
CoNi-A2	256.4 \pm 0.41	157.6 \pm 0.41	141.9 \pm 0.09	2.873 \pm 0.0032
CoNi-A + 1	256.7 \pm 0.86	157.9 \pm 0.90	142.2 \pm 0.11	2.880 \pm 0.0036
CoNi-A + 2	255.9 \pm 0.66	157.2 \pm 0.70	142.0 \pm 0.09	2.880 \pm 0.0030
CoNi-B1	257.1 \pm 1.3	158.2 \pm 1.4	141.6 \pm 0.21	2.864 \pm 0.0068
CoNi-B2	257.5 \pm 1.0	158.7 \pm 1.1	141.7 \pm 0.17	2.868 \pm 0.0057
CoNi-C1	256.0 \pm 0.67	156.8 \pm 0.70	140.3 \pm 0.13	2.828 \pm 0.0044
CoNi-C2	256.1 \pm 0.75	156.6 \pm 0.79	140.6 \pm 0.14	2.825 \pm 0.0048
CoNi-D1	256.2 \pm 0.60	157.2 \pm 0.62	141.2 \pm 0.09	2.853 \pm 0.0029
CoNi-D2	257.8 \pm 0.78	158.6 \pm 0.82	141.3 \pm 0.12	2.849 \pm 0.0036

Table V. Effect of Solution Heat Treatment and Subsequent Aging Treatment on Elastic Constants

Specimen name	c_{11} (GPa) Mean \pm SD	c_{12} (GPa) Mean \pm SD	c_{44} (GPa) Mean \pm SD	A (unitless) Mean \pm SD
CMSX-4-A-SHT	247.1 \pm 2.1	155.4 \pm 2.2	129.1 \pm 0.28	2.813 \pm 0.010
CMSX-4-A-Aged	249.0 \pm 2.3	157.0 \pm 2.4	129.2 \pm 0.29	2.811 \pm 0.011
CMSX-4-B-SHT	246.1 \pm 1.8	154.3 \pm 1.9	129.4 \pm 0.31	2.819 \pm 0.011
CMSX-4-B-Aged	248.3 \pm 2.0	156.3 \pm 2.1	129.6 \pm 0.28	2.817 \pm 0.0097
Co-6Ti-A-SHT	248.5 \pm 0.77	157.6 \pm 0.80	146.0 \pm 0.14	3.211 \pm 0.0049
Co-6Ti-A-Aged	248.1 \pm 1.4	156.8 \pm 1.4	145.4 \pm 0.24	3.188 \pm 0.0086
Co-6Ti-B-SHT	250.1 \pm 1.1	159.6 \pm 1.1	145.6 \pm 0.13	3.218 \pm 0.0046
Co-6Ti-B-Aged	250.7 \pm 1.8	160.2 \pm 1.9	145.2 \pm 0.23	3.211 \pm 0.0086
CoNi-A1-SHT	255.5 \pm 0.70	156.2 \pm 0.72	142.2 \pm 0.10	2.864 \pm 0.0029
CoNi-A1-Aged	256.7 \pm 0.36	157.6 \pm 0.37	142.1 \pm 0.08	2.868 \pm 0.0029
CoNi-A2-SHT	256.2 \pm 0.43	156.9 \pm 0.44	142.2 \pm 0.10	2.864 \pm 0.0033
CoNi-A2-Aged	256.4 \pm 0.41	157.6 \pm 0.41	141.9 \pm 0.09	2.873 \pm 0.0032

C. Effect of Mode Selection and Polynomial Order

In theory, each specimen has one true set of elastic constants that are probed indirectly by fitting a subset of the lowest-frequency modes, measured *via* RUS, with an appropriately parameterized model. Although even a well-fitting model may yield biased parameter estimates if the measured data contain outliers, or lack necessary information about the complete elastic response. Considering that the shape of a specimen influences the types of modes that occur,^[49] and that the lowest-frequency modes of a specimen are dominated by the smallest (*i.e.*, shear) elastic moduli,^[44,49] it is logical to be concerned that any small sampling of these lowest-frequency modes might provide incomplete information. Although the definition of a small sampling of modes is not obvious and warrants further attention, it is noted in Section II-B-1 that the $9 \times 10 \times 11$ mm parallelepiped geometry was intentionally devised with these concerns in mind.

One obvious way to mitigate the potential of a non-representative sampling of resonance modes being used for a given inversion is to simply use more modes. This assumes that a larger sampling of the lowest-frequency resonance modes would be less influenced by any

single outlier mode, more likely to contain a variety of mode types, and include resonances that are sensitive to the larger moduli that influence higher-frequency modes. Of course, there are a couple of potential problems with this strategy that need to be mentioned. First, the larger the number of modes used (N), the more likely it is that the list contains a missing or spuriously identified mode that could confound the inversion. Fortunately, with the transducer cradle configuration as described, and the fact that each inference is based on an average of five individual broad measurements, missing or spurious modes are quite rare occurrences. Furthermore, Section IV-D demonstrates how a missing mode is readily identifiable using posterior predictive plots, allowing for the measured mode list to be corrected. A second and potentially less obvious issue is that the accuracy of the forward model generally decreases as the order of the resonance modes increases, necessitating a higher polynomial order (P) and substantially greater evaluation time. This is all motivation for a study of inference results as a function of N and P .

Table VI summarizes the parameter estimates from a series of Bayesian inferences of specimen CoNi-A + 2

Table VI. Summary of Parameter Estimates as a Function of the Number of Modes Used for Specimen CoNi-A + 2

Modes Used	Polynomial Order	Noise σ (kHz) Mean \pm SD	c_{11} (GPa) Mean \pm SD	c_{12} (GPa) Mean \pm SD	c_{44} (GPa) Mean \pm SD	A (Unitless) Mean \pm SD
1–10	10	0.24 \pm 0.15	256.7 \pm 4.6	158.0 \pm 4.7	141.7 \pm 0.85	2.869 \pm 0.022
1–15	10	0.134 \pm 0.041	255.3 \pm 1.7	156.5 \pm 1.8	142.0 \pm 0.30	2.875 \pm 0.0080
1–20	10	0.135 \pm 0.029	256.9 \pm 1.3	158.2 \pm 1.3	142.1 \pm 0.14	2.880 \pm 0.0040
1–25	10	0.180 \pm 0.032	255.8 \pm 1.5	157.2 \pm 1.5	142.1 \pm 0.15	2.882 \pm 0.0044
1–30	10	0.216 \pm 0.033	257.4 \pm 1.5	158.9 \pm 1.6	141.9 \pm 0.16	2.880 \pm 0.0051
1–30	12	0.152 \pm 0.025	256.9 \pm 1.1	158.3 \pm 1.1	142.0 \pm 0.12	2.882 \pm 0.0037
1–35	12	0.144 \pm 0.020	256.4 \pm 0.90	157.8 \pm 0.94	142.0 \pm 0.099	2.880 \pm 0.0032
1–40	12	0.138 \pm 0.018	255.9 \pm 0.66	157.2 \pm 0.70	142.0 \pm 0.095	2.880 \pm 0.0030
1–45	12	0.210 \pm 0.024	256.3 \pm 1.0	157.7 \pm 1.1	141.9 \pm 0.13	2.879 \pm 0.0042
1–50	12	0.218 \pm 0.023	256.5 \pm 0.98	158.0 \pm 1.0	141.8 \pm 0.12	2.878 \pm 0.0041
1–60	12	0.242 \pm 0.024	255.6 \pm 1.0	157.0 \pm 1.1	141.7 \pm 0.14	2.875 \pm 0.0046
1–70	12	0.336 \pm 0.031	255.3 \pm 1.2	156.8 \pm 1.3	141.6 \pm 0.18	2.874 \pm 0.0061

Note the results highlighted in bold values are only provided for explanatory purposes as they are either biased by numerical error (*i.e.*, insufficient polynomial order) or insufficient data.

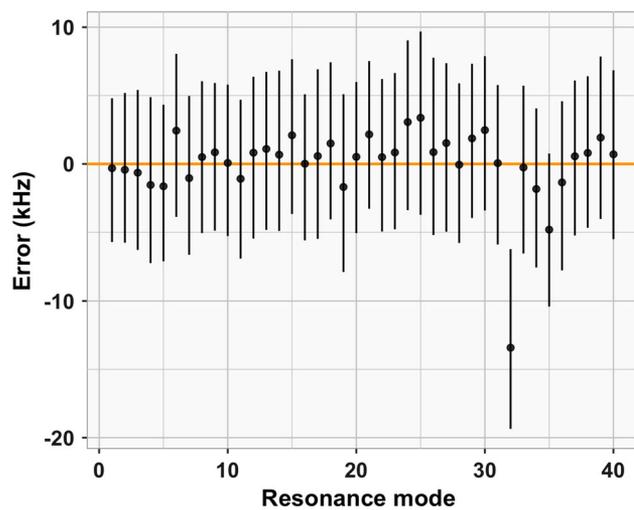


Fig. 8—Posterior predictive plot of CoNi-D1 demonstrating how such a plot can be used to quickly identify which mode was missed during RUS measurement.

using various N and P values. The first takeaway is that consistent results are achieved despite the choice of N , with the best results from the perspective of the lowest uncertainty achieved when $N = 40$. Recalling Figure 7 and the discussion in Section III-E-3, it is apparent that $P = 10$ is insufficient when $N = 40$, thus $P = 12$ is exclusively used for inferences when $N > 30$. The case where $N = 30$ provides an interesting comparison as both $P = 10$ and $P = 12$ inferences were conducted. As this comparison makes clear, $P = 12$ leads to more precise parameter estimates, but the mean values are not appreciably different. A similar observation comes from the $N = 60$ and $N = 70$ inversions where the uncertainty due to an insufficiently small P , as determined through careful inspection of the posterior predictive plots, contributes to slightly worse precision. Fortunately, the accuracy appears minimally affected when numerical error is introduced by an insufficiently small P .

D. Impact of Missing or Spurious Modes

A common issue that confounds RUS inversion efforts is missing or spuriously identified modes from a broadband RUS measurement. This is particularly problematic when the mode is of lower frequency and near the beginning of a frequency ordered list because it displaces every subsequent mode in the list. When prior information about the elastic properties of a specimen is absent and no mode shape measurement capabilities are available, it is almost impossible for one to be certain that a given measured list of modes is complete and without error. Furthermore, conducting an inversion (either an optimization or a Bayesian inference) from a measured frequency list with a missing mode may lead to erroneous elastic property estimates, or unsatisfactorily large uncertainty bounds. Fortunately, by estimating the noise parameter (σ) as part of the Bayesian inference and conducting the validation checks discussed in Section III-E-3, missing or spurious modes are readily identifiable. Then, the spurious mode can easily be excluded to correct the measured frequency list and the inversion repeated for significantly improved results.

For example, mode 32 of specimen CoNi-D1 was not identified prior to running an $N = 40$, $P = 12$ inversion, with results summarized in Table VII. The first indication of a problem with the inversion was the estimate of $\sigma = 2.70 \pm 0.28$ kHz. This noise estimate is an order of magnitude larger than a typical successful inversions, and over 15 times larger than the same inversion after properly accounting for the missing mode. Although a high σ estimate is indicative of an ill-defined model, it does not help with identifying the location of a potentially missing mode. For this, a posterior predictive plot like Figure 8 is instrumental, as it clearly indicates the location of the outlier as mode 32. Given that error is defined here as the modeled mode frequency minus the measured mode frequency, an outlier with a negative error indicates a missing mode while a positive error indicates a spurious mode in the measured frequency list.

Table VII. Summary of Parameter Estimates for Two Inversions of CoNi-D1, With and Without Accounting for Missing Mode 32

Modes Used	Missing Mode 32	Poly. Order	Noise σ (kHz) Mean \pm SD	c_{11} (GPa) Mean \pm SD	c_{12} (GPa) Mean \pm SD	c_{44} (GPa) Mean \pm SD	A (Unitless) Mean \pm SD
1–40	yes	12	2.70 \pm 0.28	283 \pm 25	186 \pm 26	147.0 \pm 3.3	3.04 \pm 0.11
1–40	no	12	0.170 \pm 0.022	256.4 \pm 0.60	157.4 \pm 0.62	141.3 \pm 0.11	2.856 \pm 0.0034

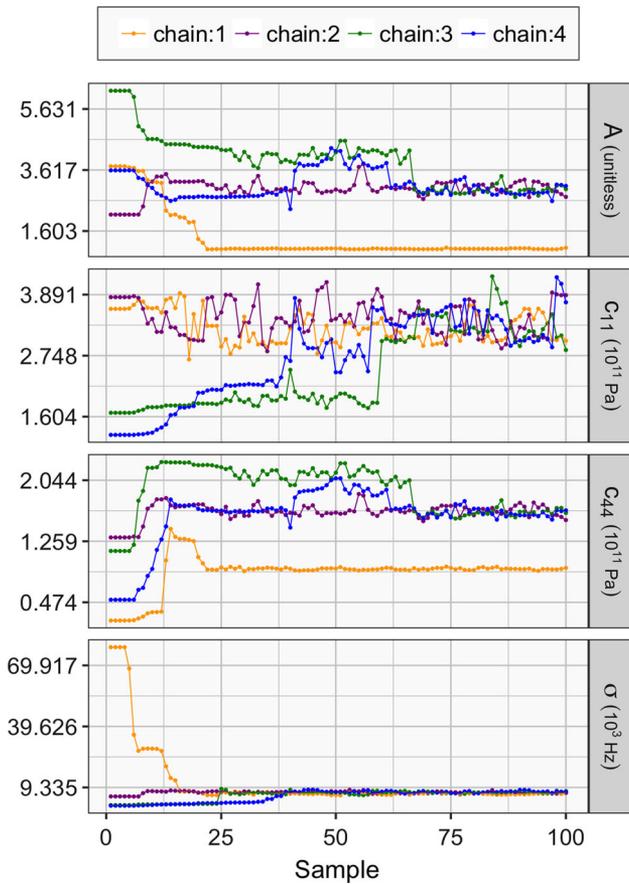


Fig. 9—Traceplots of the first 100 warm-up samples of four independent HMC chains for the CoNi-A1 specimen, without estimating orientation simultaneously.

E. Inversion Without Estimating Orientation

To demonstrate the necessity of estimating crystal orientation and elastic properties simultaneously, a previously successful inversion is repeated with the orientation parameters now fixed. CoNi-A1 was selected for this demonstration because the specimen resulted in one of the best inversions as demonstrated by the precise parameter estimates summarized in Table IV. The effect of fixing the three cubochoctic orientation parameters enforces the (incorrect) assumption that the crystal and specimen axes are perfectly aligned. However, no attempt was made to measure the crystal axes of the parent single crystal before the parallelepiped specimen was fabricated, so it is extremely unlikely that the axes are indeed aligned. This is a perfect example of an ill-defined model, because there should be no combination of c_{11} , A , and c_{44} that will lead to a forward model

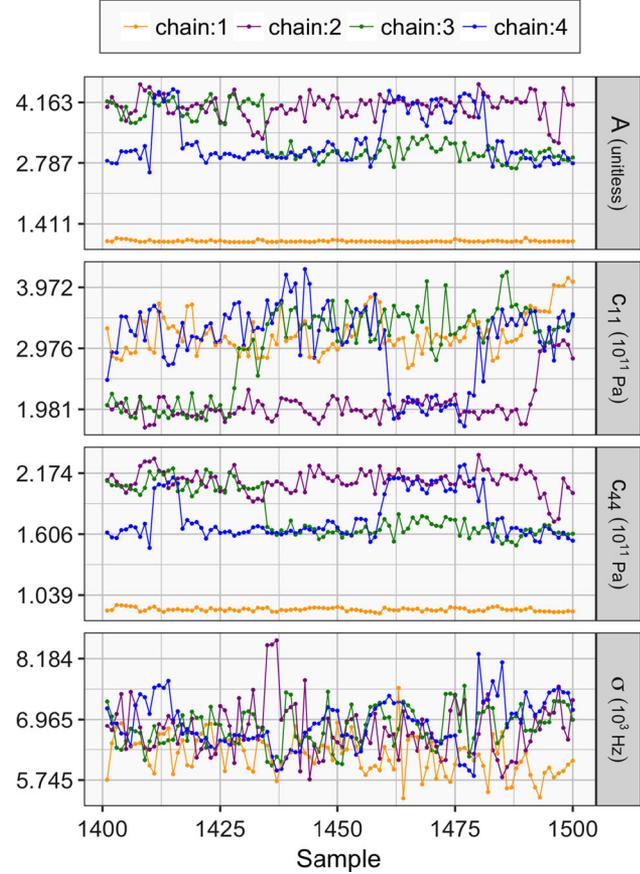


Fig. 10—Traceplots of the last 100 sampling samples of four independent HMC chains for the CoNi-A1 specimen inverted with fixed orientation parameters.

evaluation that remotely agrees with the RUS measured data.

First see the traceplots in Figure 9, which provide the first 100 warm-up samples of the CoNi-A1 inversion with fixed orientation parameters. Given that the initial conditions are randomly generated, there is an initial period of exploration where the parameter estimates are improved. But unlike the rapid convergence of all four chains to a single stationary region of the posterior as Figure 4 shows, the four chains do not agree after the first 100 samples. In fact, the chains never agree on a single set of elastic parameters to explain the measured data, as Figure 10 shows by plotting the last 100 samples from the sampling regime, just as Figure 5 previously plotted the last 100 samples from a successful inversion of CoNi-C1. Multimodal tendencies are demonstrated by the elastic parameter estimates of chains 2–4, while the

Table VIII. Summary of Parameter Estimates for Two Inversions of CoNi-Al, With and Without Estimating Crystal Orientation

Orientation Estimated	Noise σ (kHz) Mean \pm SD	c_{11} (GPa) Mean \pm SD	c_{12} (GPa) Mean \pm SD	c_{44} (GPa) Mean \pm SD	A (Unitless) Mean \pm SD
Yes	0.124 \pm 0.013	256.7 \pm 0.36	157.6 \pm 0.37	142.1 \pm 0.08	2.868 \pm 0.0029
No	6.80 \pm 0.55	299 \pm 61	174 \pm 63	158 \pm 43	2.7 \pm 1.1

estimate of σ around 7 kHz is over an order of magnitude larger than the successful inversions. Table VIII provides a comparison of the two CoNi-Al inversion results as a summary.

V. CONCLUSION

CmdStan-RUS is described and demonstrated here for the first time as a versatile and efficient tool for RUS Bayesian inference of elastic properties. This open-source tool offers significant improvements over traditional RUS inversion schemes, from a simplified implementation to more robust convergence behavior. Notable attributes and advantages are provided below as a summary of the technique:

- No guess values of elastic constants are necessary to initialize an inversion.
- Identifying measured modes and pairing of measured and modeled lists of modes is completely unnecessary.
- Missing or spurious modes are readily identified from posterior predictive plots and easily corrected.
- Random initial parameterization of HMC chains safeguards against biasing results with user-specified values.
- Estimating crystal orientation eliminates the need for costly measurements and simplifies specimen fabrication.
- Statistical modeling tools used by CmdStan-RUS are automatically tuned and work reliably without user input.
- Estimating a noise parameter models a real feature of repeated RUS measurements and reduces the chance of biased parameter estimates.
- Multiple HMC chains from markedly different initial conditions converging to a single stationary region of the posterior provide great confidence in results.
- Consistent results have been demonstrated using a range of user-specified polynomial order and mode count.
- A failed HMC chain is rare, trivial to identify, and may be discarded (or rerun) without biasing results.
- Quantification of parameter uncertainty is intrinsic to a Bayesian approach and comes directly from the posterior distributions.
- While only cubic elastic symmetry was considered here, trivial modifications of input files allow for inversions of lower-symmetry materials as well.

Valuable insights have been gained through the application of CmdStan-RUS to estimate the elastic constants of a variety of single crystal Ni-, Co-, and

CoNi-based γ' containing alloys. Excellent agreement is observed, both among independent chains of a single inference and between the estimated parameters of the two parallelepiped specimens of each alloy. This has led to the following conclusions with regard to the elastic constants of the investigated single crystal materials:

- Co-based alloys exhibited elastic anisotropy (A) 11 pct greater than the CoNi-based alloys and 13 pct greater than the Ni-based alloy CMSX-4.
- The c_{44} elastic constant was observed to increase monotonically as Ni was replaced by Co for a total average increase of 12 pct between CMSX-4 and the Co-based alloys.
- The c_{12} elastic constant was only marginally higher for the CoNi-based alloys as compared to CMSX-4, while a 3 pct increase between the CoNi- and the Co-based alloys contributes to the larger A of the Co-based alloys.
- The CoNi alloys exhibited similar elastic properties while Co-based alloys were more sensitive to alloying.
- Heat treatments designed to modify the $\gamma - \gamma'$ microstructure did not appreciably affect the elastic constants of the alloys investigated.

ACKNOWLEDGMENTS

This research was funded by the U.S. Air Force Research Laboratory (AFRL) through BAA Contract FA8650-15-M-5208, SBIR Contract FA8650-15-M-5074. Vibrant Corporation provided the RUS measurement equipment and many helpful discussions. This paper was cleared for public release by the AFRL, Case Number 88ABW-2017-6331

REFERENCES

1. T.M. Pollock and S. Tin: *J. Propuls. Power*, 2006, vol. 22, pp. 361–74.
2. D. Miracle: *Acta Metall. Mater.*, 1993, vol. 41, pp. 649–84.
3. T.M. Pollock: *Nat. Mater.*, 2016, vol. 15, pp. 809–15.
4. B.P. Bewlay, M.R. Jackson, P.R. Subramanian, and J.C. Zhao: *Metall. Mater. Trans. A*, 2003, vol. 34A, pp. 2043–52.
5. J. Lemberg and R. Ritchie: *Adv. Mat.*, 2012, vol. 24, pp. 3445–80.
6. J. Perepezko: *Annu. Rev. Mater. Res.*, 2015, vol. 45, pp. 519–42.
7. M. Holmquist, R. Lundberg, O. Sudre, A. Razzell, L. Molliex, J. Benoit, and J. Adlerborn: *J. Eur. Ceram. Soc.*, 2000, vol. 20, pp. 599–606.
8. R. Naslain and F. Christin: *MRS Bull.*, 2003, vol. 28, p. 654658.
9. H. Ohnabe, S. Masaki, M. Onozuka, K. Miyahara, and T. Sasa: *Comp. A*, 1999, vol. 30, pp. 489–96.
10. T.M. Pollock, J. Dibbern, M. Tsunekane, J. Zhu, and A. Suzuki: *JOM*, 2010, vol. 62, pp. 58–63.
11. M.S. Titus, A. Suzuki, and T.M. Pollock: *Scripta Mater.*, 2012, vol. 66, pp. 574–77.

12. J. Sato, T. Omori, K. Oikawa, I. Ohnuma, R. Kainuma, and K. Ishida: *Science*, 2006, vol. 312, pp. 90–91.
13. M. S. Titus, A. Suzuki, and T. M. Pollock: *Superalloys 2012*, eds. E. Huron, R. Reed, M. Hardy, M. Mills, R. Montero, P. Portella and J. Telesman, John Wiley & Sons, Inc., Hoboken, NJ, USA, 2012, pp. 823–32.
14. C. Stewart, R. Rhein, A. Suzuki, and T. Pollock: *Superalloys*, Wiley, New York, 2016, pp. 991–99.
15. R. W. Jackson, M. S. Titus, M. R. Begley, and T. M. Pollock: *Surface Coat. Technol.*, 2016, vol. 289, pp. 61–68.
16. B. Bales, B.R. Goodlet, W.C. Lenthe, L. Petzold, and T.M. Pollock: *J. Acoust. Soc. Am.*, 2017, vol. 143, pp. 71–83.
17. Stan Development Team: *Stan Modeling Language User's Guide and Reference Manual, Version 2.15.0*, 2016.
18. M.S. Titus, Y.M. Eggeler, A. Suzuki, and T.M. Pollock: *Acta Mater.*, 2015, vol. 82, pp. 530–39.
19. Y.M. Eggeler, M.S. Titus, A. Suzuki, and T.M. Pollock: *Acta Mater.*, 2014, vol. 77, pp. 352–59.
20. B.C. Wilson, J.A. Hickman, and G.E. Fuchs: *JOM*, 2003, vol. 55, pp. 35–40.
21. A. Migliori and J. Sarrao: *Resonant ultrasound spectroscopy: applications to physics, materials measurements, and nondestructive evaluation*, Wiley-Interscience, New York, 1997.
22. A. Migliori, J. Sarrao, W.M. Visscher, T. Bell, M. Lei, Z. Fisk, and R. Leisure: *Phys. B Condens. Matter.*, 1993, vol. 183, pp. 1–24.
23. H.H. Demarest: *J. Acoust. Soc. Am.*, 1971, vol. 49, p. 768.
24. E. Schreiber, O.L. Anderson, and N. Soga: *Elastic Constants and Their Measurement*, McGraw-Hill, New York, 1973.
25. E. Mochizuki: *J. Phys. Earth*, 1987, vol. 35, pp. 159–70.
26. J. Maynard: *Phys. Today*, 1996, vol. 49, p. 26.
27. P. S. Spoor: Elastic properties of novel materials using pvdf film and resonance ultrasound spectroscopy: Tech. Rep. ADA328037, Pennsylvania State Univ. Dept. of Physics, 1997.
28. R.G. Leisure and F.A. Willis: *J. Phys. Condens. Matter*, 1997, vol. 9, pp. 6001–29.
29. R. Neal: *MCMC Using Hamiltonian Dynamics*, CRC Press, New York, 2011, pp. 113–63.
30. W.M. Visscher, A. Migliori, T.M. Bell, and R.A. Reinert: *J. Acoust. Soc. Am.*, 1991, vol. 90, pp. 2154–62.
31. D. Rosca, A. Morawiec, and M.D. Graef: *Model. Simul. Mater. Sci. Eng.*, 2014, vol. 22, p. 075013.
32. M.D. Hoffman and A. Gelman: *J. Mach. Learn. Res.*, 2014, vol. 15, pp. 1593–23.
33. J. Plešek, R. Kolman, and M. Landa: *J. Acoust. Soc. Am.*, 2004, vol. 116, pp. 282–87.
34. M.C. Remillieux, T.J. Ulrich, C. Payan, J. Rivière, C.R. Lake, and P.-Y. Le Bas: *J. Geophys. Res. Solid Earth*, 2015, vol. 120, pp. 4898–16.
35. C. Zener: *Phys. Rev.*, 1947, vol. 71, pp. 846–51.
36. J. Geweke: *Am. Econ. Rev.*, 2007, vol. 97, pp. 60–64.
37. C.M. Grinstead and J.L. Snell: *Introduction to Probability*, 2nd ed., American Mathematical Society Providence, Rhode Island, 2003.
38. D. Higdon, K. Heitmann, E. Lawrence, and S. Habib: *Using the Bayesian Framework to Combine Simulations and Physical Observations for Statistical Inference*, Wiley, New York, 2010, pp. 87–105.
39. P.K. Kitaniadis: *Bayesian and Geostatistical Approaches to Inverse Problems*, Wiley, New York, 2010, p. 71.
40. Stan Development Team: *CmdStan Interface User's Guide, CmdStan Version 2.16.0*, 2017.
41. W.K. Hastings: *Biometrika*, 1970, vol. 57, pp. 97–109.
42. S. Duane, A.D. Kennedy, B.J. Pendleton, and D. Roweth: *Phys. Lett. B*, 1987, vol. 195, pp. 216–22.
43. N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, and E. Teller: *J. Chem. Phys.*, 1953, vol. 21, pp. 1087–92.
44. P. Sedlák, H. Seiner, J. Zidek, M. Janovská, and M. Landa: *Exp. Mech.*, 2014, vol. 54, pp. 1073–85.
45. J. Schreuer and S. Haussühl: *EMU Notes Mineral.*, 2005, vol. 7, pp. 173–98.
46. A.B. Parsa, P. Wollgramm, H. Buck, C. Somsen, A. Kostka, I. Povstugar, P.-P. Choi, D. Raabe, A. Dlouhy, J. Müller *et al.*: *Adv. Eng. Mater.*, 2015, vol. 17, pp. 216–30.
47. D. Siebörger, H. Knake, and U. Glatzel: *Mater. Sci. Eng. A*, 2001, vol. 298, pp. 26–33.
48. T. Ichitsubo, D. Koumoto, M. Hirao, K. Tanaka, M. Osawa, T. Yokokawa, and H. Harada: *Acta Mater.*, 2003, vol. 51, pp. 4863–69.
49. B.J. Zadler, J.H.L. Le Rousseau, J.A. Scales, and M.L. Smith: *Geophys. J. Int.*, 2004, vol. 156, pp. 154–69.