

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Data Science for Materials Science

Permalink

<https://escholarship.org/uc/item/0qs684d5>

Author

Bales, Benjamin

Publication Date

2019

Supplemental Material

<https://escholarship.org/uc/item/0qs684d5#supplemental>

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Data Science for Materials Science

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Mechanical Engineering

by

Benjamin Bruce Bales

Committee in charge:

Professor Linda Petzold, Chair
Professor Tresa Pollock
Professor Matthew Begley
Professor Frederic Gibou

June 2019

The Dissertation of Benjamin Bruce Bales is approved.

Professor Tresa Pollock

Professor Matthew Begley

Professor Frederic Gibou

Professor Linda Petzold, Committee Chair

June 2019

Data Science for Materials Science

Copyright © 2019

by

Benjamin Bruce Bales

Acknowledgements

The research in Chapter 2 is based upon work supported by the National Science Foundation under Grants No. DMR 1233704 and DMR 1534264. The backscatter electron micrographs of Rene N5 come from data collected by Luke Rettenberg.

The research in Chapter 3 was funded by the U.S. Air Force Research Laboratory (AFRL) through BAA Contract FA8650-15-M-5208, SBIR Contract 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 (DMR-1121053) and NSF CNS-0960316. Vibrant Corporation provided the RUS measurement equipment and many helpful discussions. The original paper was cleared for public release by the AFRL, case number 88ABW-2017-3838.

Curriculum Vitæ

Benjamin Bruce Bales

Education

- 2019 Ph.D. in Mechanical Engineering (Expected), University of California, Santa Barbara.
- 2011 B.S. in Electrical Engineering, University of Tennessee, Knoxville.

Publications

- Bales, B., Pourzanjani, A., Vehtari, A., Petzold, L. (2019). Selecting the Metric in Hamiltonian Monte Carlo. *Submitted to Neural Information Processing Systems*.
- Goodlet, B., Mills, L., Bales, B., Charpagne, M., Murray, S., Lenthe, W. C., Petzold, L. R., Pollock, T. (2018). Elastic Properties of Novel Co and CoNi-Based Superalloys Determined through Bayesian Inference and Resonant Ultrasound Spectroscopy. *Metallurgical and Materials Transactions A* 49(6), 2324-2339.
- Bales, B. B., Goodlet, B., Lenthe, W., Pollock, T., Petzold, L. R. (2018). Bayesian inference of elastic properties with resonant ultrasound spectroscopy. *Journal of the Acoustic Society of America* 143(1), 71-83.
- Bales, B. B., Pollock, T., Petzold, L. R. (2017). Segmentation-Free Image Processing and Analysis of Precipitate Shapes in 2D and 3D. *Modelling Simul. Mater. Sci. Eng.* 25, 045009.
- Bellesia, G., Bales, B. B. (2016). Population dynamics, information transfer, and spatial organization in a chemical reaction network under spatial confinement and crowding conditions. *Physical Review E* 94(4), 042306.
- Bezzola, A., Bales, B. B., Petzold, L. R., & Alkire, R. C. (2014). Numerical Scaling Studies of Kinetically-Limited Electrochemical Nucleation and Growth with Accelerated Stochastic Simulations. *Journal of The Electrochemical Society* 161(8), E3001-E3008.
- Bezzola, A., Bales, B. B., Alkire, R. C., & Petzold, L. R. (2014). An exact and efficient first passage time algorithm for reaction-diffusion processes on a 2D-lattice. *Journal of Computational Physics* 256, 183-197.

Abstract

Data Science for Materials Science

by

Benjamin Bruce Bales

Data analysis in materials science is of increased interest due to the rate at which large datasets can be generated. This thesis covers micrograph analysis, mechanistic modeling, and inference techniques for materials problems.

Segmentation based image analysis techniques are routinely employed for quantitative analysis of complex microstructures containing multiple phases. The downside is that computing reliable segmentations is challenging and, if no special care is taken, segmentation artifacts will make subsequent analysis difficult. Using a two phase nickel-base superalloy microstructure as a model system, we demonstrate a new methodology for analysis of precipitate shapes using a segmentation-free approach based on the histogram of oriented gradients feature descriptor (HOG), a classic tool in image analysis. The benefits of this methodology for analysis of microstructure in two and three dimensions are demonstrated.

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. The algorithms described are demonstrated on RUS data collected from two parallelepiped specimens of structural metal alloys, a specimen of fine-grained polycrystalline Ti-6Al-4V (Ti-64)

with random crystallographic texture and isotropic elastic symmetry and a single crystal Ni-based superalloy CMSX-4 specimen. Our unique contributions are: 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.

Finally, we present a selection criterion for the Euclidean metric adapted during warmup in an HMC sampler. This makes it possible for a sampler to automatically pick the metric based on the model and the availability of warmup draws. Additionally, we present a new adaptation inspired by the selection criterion that requires significantly fewer warmup draws to be effective. The effectiveness of the selection criterion and adaptation are demonstrated on a number of applied problems.

Contents

Curriculum Vitae	v
Abstract	vi
1 Introduction	1
2 Segmentation-Free Image Processing	5
2.1 Introduction	5
2.2 Related Work	8
2.3 Methods	9
2.4 Results	10
2.5 Conclusions	19
3 Resonant Ultrasound Spectroscopy	20
3.1 Introduction	20
3.2 Experimental Methods	25
3.3 Computations	29
3.4 Results	40
3.5 Conclusions	57
4 Selecting the Metric in Hamiltonian Monte Carlo	58
4.1 Introduction	58
4.2 Preliminaries	60
4.3 Implementation Details	64
4.4 Benchmarks	67
4.5 Conclusions	72
5 Conclusions	73

Chapter 1

Introduction

In this thesis we show how data science can be used to advance research in Materials Science and present a method for making it easier to use Hamiltonian Monte Carlo (HMC) as a black-box Bayesian inference method.

To be more specific, the types of materials studied were single crystal superalloys used in the manufacture of gas turbines. The data analyzed were micrographs collected for the purpose of characterizing different superalloy microstructures, and resonance data collected for the purpose of making precise estimates of the elastic constants of novel superalloy compositions.

In terms of contributions, Chapter 2 discusses a technique for extracting information from micrographs without performing an image segmentation, a notoriously difficult piece of most existing analysis pipelines. Chapter 3 describes an approach to estimating elastic constants (with uncertainty estimates) from resonance data, a probabilistic variant of the resonance ultrasound spectroscopy technique (RUS). Chapter 4 introduces an improvement in the statistical techniques used for the inferences in Chapter 3.

The work in Chapter 2 was motivated by the need to automate micrograph analysis, a large driving force in superalloy development. It is difficult to perfect turbine

manufacturing technologies because the components being built and deployed are so expensive. Gas turbines are built to last – a blade in a turbine may be expected to last tens of thousands of hours. How, then, is it possible to understand how choices about the composition of the blade, or treatments of the blade, determine this lifetime?

The answers mostly come from pictures. Scientists take components at various points in their lifetime and cut them open. The question of how might a change in a design variable influence the performance of a component reduces to two constituent questions: how might a change in design variable change a microstructure, and how might that microstructure change influence performance. In this way, micrograph collection provides a window into the life of gas turbines.

Rafting is prototypical of the types of information a micrograph can provide. An example of this phenomenon is given in Figure 1.1. Once this change is quantified numerically, it can provide a proxy for the connection between the design variables and component performance. This can be done in a number of different ways. In Chapter 2 we demonstrate a technique that does not require an image segmentation operation, which can be quite difficult to compute even with high quality micrographs.

A downside of the analysis of Chapter 2 is that it requires cutting up a part. This is not always appropriate. For instance, if a component is to be checked before installation, then a useful check cannot destroy the component. The work in Chapter 3 is motivated largely by this need and falls into the category of non-destructive evaluation (NDE) of materials. The technique in Chapter 3 does not directly answer the question of whether components are ready for use, but it does look at the related question of how to non-destructively evaluate material properties. The technique explored is resonance ultrasound spectroscopy (RUS), the process by which the resonance modes of a linear elastic material are measured and then used to compute the elastic constants of that material.




Figure 1.1: Unrafted microstructure on left and rafted microstructure on right. Micrographs are courtesy Luke Rettenberg.

RUS is a relatively well known technique. What makes the work in Chapter 3 unique is that, along with elastic constants, the method provides, for single crystal materials, an estimate of the misorientation between the crystal and the sample itself. Previously this misorientation would need to be estimated using X-ray diffraction, requiring additional time and special equipment. We use a Bayesian formulation of the problem, which provides uncertainty estimates on all of the model parameters. This is useful because it makes it much easier to evaluate the precision of the experiments. The Bayesian approach to RUS using Markov Chain Monte Carlo (MCMC) is not unique to us [1], however, the use of Hamiltonian Monte Carlo (HMC, an MCMC variant) is. Our follow-up work in [2] ported the code to Stan [3], a well established Bayesian inference package. This increased the flexibility of the probabilistic model and made it possible to more easily switch between different types of RUS models and to explore a larger collection of data more easily. As a comparison, the original work described in Chapter 3 studies two materials, whereas the follow up work [2] studies about fifteen different material/treatment combinations.

Chapter 4 is inspired by the work involved in making the Bayesian approach to

RUS (Chapter 3) possible. The HMC code developed in Chapter 3 is tricky to use because there are a fair number of tuning parameters that must be set for sampling to be efficient. Porting the RUS model to Stan, a software package for Bayesian inference, removed this difficulty and made it possible to scale up the analysis [2] because Stan uses more advanced HMC algorithms and an automated warmup procedure to configure itself automatically.

The work in Chapter 4 contributes directly to the warmup procedure that MCMC packages like Stan depend upon, and in this way impacts more than just Materials Science. MCMC warmup procedures are used to adjust the algorithms automatically to different problems, making it possible to develop black-box inference packages. The particular element of warmup that Chapter 4 targets is the robust estimation of a linear coordinate transform that improves HMC sampling efficiency. A variety of different choices can be made for this coordinate transformation, which presents a selection problem. In some cases, one is better. In another case, another is better. The main observation in this Chapter is that the effectiveness of this coordinate transform can be understood in terms of the stability limit of the leapfrog integrator used internally in HMC. This leads to a technique for picking between the different coordinate transformations, and an alternate method for building the transformation that can work with shorter warmups than the others.

Chapter 2

Segmentation-Free Image Processing¹

2.1 Introduction

The strong driving force for development of rigorous property models for structural materials motivates quantitative analysis of microstructure across a spectrum of alloy systems [4]. Since most engineering materials are multiphase in character, it is usually essential to isolate individual phases for analysis of size, shape and/or distribution in order to input this information into property models. The process for quantifying microstructure typically involves collection of 2D or 3D data on a pixel by pixel basis, followed by a segmentation operation to isolate individual phases within the microstructure. Shape metrics such as volumes, surface areas, or statistical moments [5, 6]) of the resulting precipitates are used to quantify the analysis. These metrics are chosen in part

¹©IOP Publishing. Reproduced with permission. All rights reserved. From: Ben Bales, Tresa Pollock, and Linda Petzold. “Segmentation-free imageprocessing and analysis of precipitate shapes in 2D and 3D”. In: *Modelling and Simulation in Materials Science and Engineering* 25.4 (Apr.2017), p. 045009.

because of their similarity with quantitative microstructure analysis that has been performed manually [7, 8]. The conjecture is that if enough micrographs can be captured and enough precipitates can be characterized, the shape statistics will yield good feature descriptors that can then be used in whatever classification or regression tasks that need to be addressed.

The examples in this chapter are of nickel-based superalloys. For these alloys, it is desirable to develop heat treatment cycles to adjust precipitate shapes for optimization of mechanical properties [7]. A unique feature of this class of alloys is the tendency for the precipitates to undergo directional coarsening during the application of external stresses at elevated temperatures [9], a process known as "rafting". In both cases, measuring the shape of the microstructural precipitates can provide important insights on alloy design and mechanical properties.

The problem with this measurement is that the segmentations are rarely trivial. Especially across data sets, but even within datasets, it can be very difficult to parameterize a segmentation algorithm to produce consistent results. Because the segmentation parameterization can strongly influence the shape statistics and because producing high quality segmentation often requires extensive fine tuning of segmentation parameters, it is difficult to argue that the resultant shape statistics are unbiased (with regards to the segmentation). The artifacts an automated segmentation of a γ - γ' microstructure might produce depend on the imaging modality, but typically include:

1. A large number of single pixel γ or γ' precipitates appear due to detector noise in the scanning electron microscope used to generate the image.
2. Individual precipitates are merged into one large precipitate because the original image does not have high enough resolution for them to be segmented without high level material-specific knowledge.




Figure 2.1: An example backscattered electron micrograph showing γ matrix (light) and γ' precipitates (dark). A , B , and B' highlight regions where the presence of a precipitate is uncertain. A shows four precipitates that have been incorrectly merged by the segmentation algorithm. B and B' highlight areas where maybe there is a precipitate and maybe there isn't. In both cases the segmentation algorithm must make a decision between these two extremes.

Figure 2.1 shows an example image with the first two types of defects. These issues are not unique to superalloy microstructures, and techniques can be developed to address them [10, 11, 12, 13], though it is still very difficult to make them robust, particularly across data sets. The simpler solution, if the information needed and the associated analysis allows, is to employ image analysis approaches that do not rely on segmentation.

The goal of this chapter is to highlight how a tool from computer vision, the Histogram of Oriented Gradients (HOG) feature detector, can be used to solve a wide variety of relevant classification and measurement problems robustly with respect to the difficulties enumerated above. HOG feature detectors have a long, rich history of application in computer vision [14, 15, 16], but to the best of our knowledge have not been used in the study of microstructure.

This chapter is organized as follows. In Section 2.2 we describe related work and in Section 2.3 we outline computation of a HOG feature descriptor. In Section 2.4 we demonstrate the effectiveness of the HOG feature descriptor on a number of relevant

microstructure characterization problems, where microstructural information is available in both 2D and 3D.

2.2 Related Work

There is an extensive existing literature on microstructure analysis including understanding how composition affects microstructure [9] [17], understanding how processing steps affect microstructure [18], and understanding how precipitate shape properties affect strength and other material properties [7], etc. Most of these papers use segmentation based characterization techniques, be they measuring simple areas, aspect ratios, and perimeters [8] or more complicated metrics [5, 6].

There are many similarities between microstructure analysis and the classic problem of shape analysis [19] [20]. Shape analysis in general involves taking outlines of objects and using this information, for example in simple object recognition [20] or pose estimation [19]. A frequent limitation in this process was, much like for the microstructures, obtaining the outlines of the objects. Outlining is simply a segmentation, albeit the images are usually much more complicated than a superalloy microstructure micrograph. This field took a leap forward as techniques were developed to solve the motivating problems directly (pose estimation, object recognition, etc.), without first doing a segmentation. Even though the idea of a segmentation as a blackbox step in shape analysis seemed reasonable, it was limiting. These changes were fueled by the introduction of the Scale-Invariant Feature Transforms (SIFT) [21] for sparse keypoint identification, the repopularization of HOG descriptors by Dalal and Triggs [15], and, more currently, work in neural networks [22]. The inspiration to our current work was the success that shape analysis enjoyed with these segmentation-free feature descriptors. There are, however, reconstruction-based imaging techniques that can directly lead to easier to segment

datasets [23] [24] as compared to the simple micrographs used here.

Several other microstructure characterization techniques that do not involve segmentation have recently been employed, including N-point statistics [25, 26] and SIFT [27]. What, given the established application of N-point statistics and SIFT in microstructure analysis do HOG descriptors provide? Basically, while a quantitative, segmentation-free feature vector is desired, it is also desirable that the feature vectors be easily interpretable by lab scientists. The strength of a segmentation is that the data it produces (the outlines of the precipitates) is easy to directly interpret and understand. Our goal has been to develop a technique that makes both of these scenarios possible: quantitative analysis similar to the N-point statistics and SIFT features, and qualitative analysis similar to that done with segmentations. HOG descriptors, as used here, fill that gap.

2.3 Methods

Computation of the HOG feature descriptor itself is straightforward. This technique is suitable for either standard scanning or transmission electron microscope micrographs. First, an approximate gradient at every point in the image is computed. This is most easily done by applying a light Gaussian blur (just a few pixel radii) to the image and taking finite differences to obtain the gradients. The Gaussian kernel should be large enough to remove the largest detector noise, but not so large that it blurs any important features. Finally, the values of the gradient are summed into a histogram of gradient angles weighted by gradient magnitudes.

For image F with Gaussian kernel G , the gradient at each point, f_{ij} , is given by $\nabla(G * F)_{ij}$. f_{ij} is a vector with magnitude $|f_{ij}|$ and angle $\angle f_{ij}$. To build the histogram over angles, if each bin center is denoted as θ_k with radius δ , then the value of the

histogram W at that bin center is given by

$$W(\theta_k) = \sum_{|\angle f_{ij} - \theta_k| < \delta} |f_{ij}| \quad (2.1)$$

2.4 Results

2.4.1 Comparison of Microstructures

As stated before, a valuable feature of HOGs is the relative ease of their computation as compared to segmentations. Figure 2.2 shows the comparison of two superalloy microstructures from Fährmann[17].

As shown in Figure 2.2, the HOG feature descriptor has peaks pointing in the normal directions of the facets in the top sample of Figure 2.2. This is because the histogram accumulates the magnitudes of gradients, so that where the gradient is large, large values are accumulated. In these superalloy micrographs, the gradients are large at the edges of precipitates. The precipitates in the bottom sample are more spherical, and the HOG feature descriptor reflects this.

The simplest way to use the HOG as a quantitative descriptor instead of qualitative descriptor is to look at the magnitude of the FFT of the HOG feature descriptor and compare the relative amount of energy in different harmonics of the microstructure. The first nine bins of the absolute value of the FFTs of the HOG feature descriptors from Figure 2.2 are shown in Table 2.1.

We can compare how circular the precipitates in the two microstructures are by comparing the amount of energy in bin zero of the magnitude of the FFT to all of the other non-zero bins, and we can compare how square the precipitates in the two microstructures are by comparing the energy in every fourth non-zero frequency bin

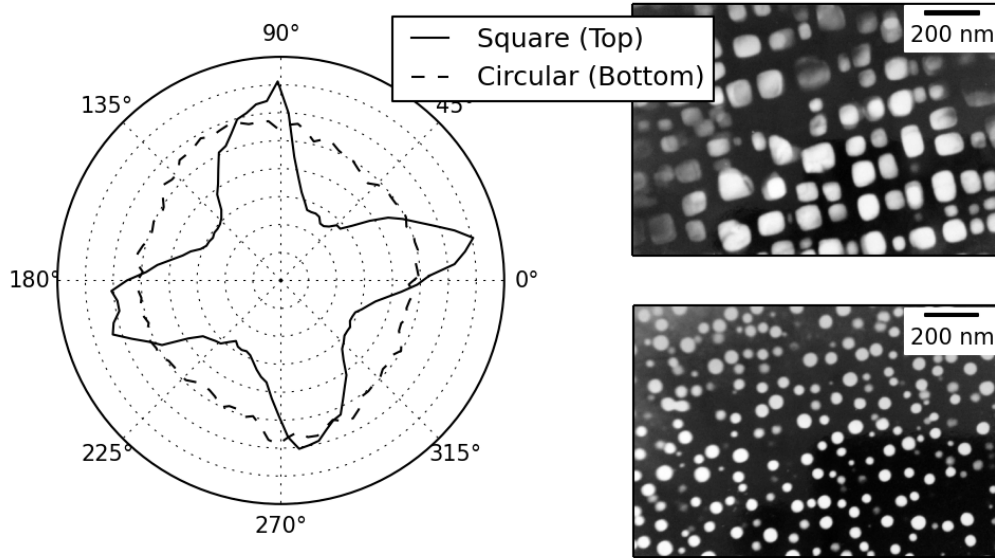


Figure 2.2: Plot of the HOG feature descriptors (on the left) of transmission electron microscope micrographs of two superalloy samples [17]. As can be seen, the precipitates in the top micrograph are more square than the precipitates in the bottom micrograph. A two pixel blur was used in computing the gradients here.

|FFT| of HOG feature descriptors for Figure 2.2

Square (top)	59	0.46	0.70	0.65	13	1.3	0.70	0.19	2.7
Circle (bottom)	88	1.4	2.9	0.29	2.2	0.32	0.29	0.57	0.18
Index	0	1	2	3	4	5	6	7	8

Table 2.1: Magnitudes of the energies in the bins of the normalized HOG feature vector. As can be seen, the circular microstructure has more energy allocated in its zero bin (highlighted in dark grey), and the square microstructure has more energy in the fourth and eighth bin (highlighted in light grey).

HOG Scores for Figure 2.2

	Circle (0Hz signal)	Square (4Hz harmonics)	Layering (2Hz harmonics)
Square (top)	20	0.95	0.97
Circle (bottom)	460	0.29	0.80

Table 2.2: These are the HOG scores for Figure 2.2. As can be seen, the Circle score is much higher for the circular microstructure, and the Square score is much higher for the square microstructure.

(four and eight highlighted in cyan) to the energy in every other non-zero frequency bin. A similar calculation can be performed for 2Hz energies and all the harmonics.

This can detect rafting in microstructures. These three numbers are a quick way to distill microstructure information and compare images in a rotation invariant way. The calculations for the materials from Figure 2.2 are shown in Table 2.2.

It is conceivable that the HOG feature descriptor could be employed to assess elastic anisotropy in a single sample as well, though this is not investigated in this thesis.

There are limits to what information the HOG feature descriptor can extract from a microstructure. For instance, if the cuboidal precipitates in the top sample of Figure 2.2 were not globally aligned with each other, then the HOG feature descriptor for that would appear more uniform like that of the bottom sample. This could happen, for example, if the microstructure is from a polycrystalline sample.

The HOG feature descriptor also does not directly reveal information about scale. For instance, it does not say that precipitates in the top sample are on average larger or smaller than the ones in the bottom sample.

Taking a step back, even though the microstructures in Figure 2.2 appear to be simple squares and circles, the microstructure samples themselves are 3D objects. It is possible, for instance, that if the top sample was cut on an angle the HOG plot would have more or less peaks due to symmetry and the sectioning plane.

2.4.2 Detection of Rafting

The merits of HOG feature descriptors are easily demonstrated in the context of nickel-based superalloys for rafting, a unique tendency in this class of alloys for the precipitates to directionally coarsen during application of stresses at elevated temperatures [9]. Figure 2.3 shows two samples of Rene N5, one unrafted (top) and one rafted (bottom) along with plots of their HOG feature vectors. Table 2.3 shows the FFT-based scores for the rafted microstructure. The biggest change in score from the top (unrafted)

microstructure to the bottom (rafted) one is the amount of energy in 4hz harmonics. The top microstructure has a large fraction of energy there, and the bottom microstructure has basically none.

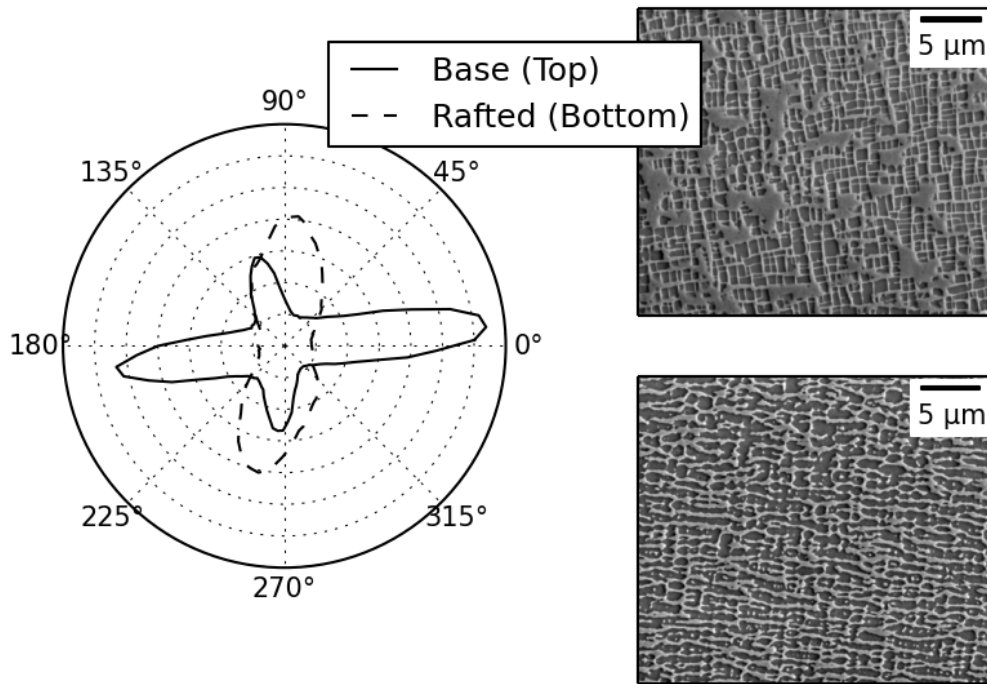


Figure 2.3: Plot of the HOG feature descriptors of BSE micrographs of a microstructure before (top) and after (bottom) rafting. The rafting is very clear in the HOG feature vector plots. A one pixel blur was used in computing the gradients.

Simple HOG Scores for Figure 2.3

	Circle (0Hz signal)	Square (4Hz harmonics)	Layering (2Hz harmonics)
Base (top)	4.4	0.69	0.91
Rafted (bottom)	9.4	0.044	0.98

Table 2.3: These are the simple HOG scores for Figure 2.3. In these two samples, the biggest difference is that the Square score is much higher for the unrafted sample. For the rafted sample, the Square score is lowered but the Layering score remains high.

2.4.3 Analysis of 3D Transformations

HOG descriptors easily transfer to 3D datasets as well. While not as extensively deployed as their 2D counterparts, these feature detectors have found practical use in video datasets for action recognition (two spatial dimensions and one time dimension) [28]. Again, in this application, they enable microstructure analysis without segmentation.

Segmentation in 3D datasets can be quite difficult. All of the same problems with 2D data remain, except now visually verifying segmentations is more involved (requiring either volume rendering or a careful use of contour plots).

A 3D HOG is simply a histogram across two dimensions. For visualization, it is usually desirable to adjust the values in the histogram to account for some bins covering a larger area on the sphere than others so that the values in the histogram are given per-area rather than just as a total sum. This adjustment is used in Figure 2.4 to show a 3D microstructure along with its Histogram of Oriented Gradients plot. As shown in Figure 2.4, there are six distinct peaks corresponding to the 6 faces of the cube-shaped precipitates. Importantly, unlike the 2-D analysis, which would return a different shape based on the sectioning plane, the 3-D analysis would identify a cube shape regardless of sectioning plane. The clarity of the 3D HOG plot in Figure 2.4 demonstrates the robustness of these feature detectors to noise.

In analogy to the FFTs, it is possible to use rotation invariant spherical harmonics [30] as feature vectors for analysis of 3D microstructures. Figure 2.5 along with Table 2.4 show the results on a simulated coarsening experiment done by Wang [31]. The first descriptor in Table 2.4 (“Cube”) comes from looking at every fourth non-zero frequency bin of the rotation invariant spherical harmonics, and second descriptor (“Sphere”) comes from looking at the energy in the zeroth bin compared to everything else. The “Cube“ score remains relatively stable compared to the “Sphere“ score which drops precipitously. This




Figure 2.4: On the left is the 3D HOG descriptor for a 3D BSE dataset (volume rendered on the right) of Rene N4 dataset collected with the Tribeam system [29]. It is the full 3D dataset associated with the image in Figure 2.1, which provides evidence that the HOG descriptor produces easily interpretable results even in the face of large amounts of noise (considerable effort was made to smooth the dataset for the volume rendering).

can be explained by thinking about the precipitate edge curvature remaining constant while the edge lengths increase. The precipitates are cuboids through the whole process, they are just becoming less and less spherical as the microstructure coarsens.

Another simpler way to quantify 3D microstructure is to look at the mass moments of inertia of the HOG feature descriptor itself (computed as if the HOG were a thin-shelled spherical object with mass given by the value at each histogram point).

For a cubic microstructure, there are six peaks in the HOG feature descriptor and three equivalent primary axis of rotation in the spherical HOG object. For a microstructure rafted into a columnar structure, there are only four strong peaks in the HOG feature detector, and likewise two equivalent axes of rotation with large moments of inertia and a third with a smaller moment. For a microstructure rafted into a layer by layer structure, the HOG feature detector has only two strong peaks and there is a single large moment of inertia and two smaller ones for the spherical HOG object. Figure 2.6 shows volume renderings of these two types of rafting that come from simulations done by Wang [31].




Figure 2.5: *A* is the base microstructure, *B* is timestep two in the rafting process, and *C* is timestep eight. The data is from the simulations of Wang et al., [31].

Coarsening experiment

Timestep	Cube	Sphere
1	0.71	21.2
2	0.85	7.4
3	0.87	4.0
4	0.87	2.7
5	0.86	2.0
6	0.84	1.6
7	0.83	1.3
8	0.82	1.1
9	0.82	1.0
10	0.81	0.9

Table 2.4: The Cube and Sphere scores for these microstructures are computed similarly to the Circle and Square scores from Table 2.2 and Table 2.3.

Table 2.5 shows the moment analysis of these experiments which reflects the behavior described above (data also from Wang [31]).

2.4.4 Effects of Sample Drift

Other than shot noise, sample drift is another important source of error in scan based imaging techniques. These distortions are most significant with analysis techniques that depend on feature correlations between images, long microscope exposure times, and high magnifications. For a detailed discussion of distortions in scanning electron




Figure 2.6: *A* is the base microstructure, *B* is the result of rafting to columns, and *C* is the result of rafting to layers. The data is from the simulations of Wang et al., [31].

Moments for HOGs of rafted microstructures

	Columnar rafting (P-Type)			Layered-by-layer rafting (N-Type)		
Time	m_1	m_2	m_3	m_1	m_2	m_3
1	0.743	0.745	0.800	0.723	0.784	0.787
2	0.731	0.733	0.823	0.688	0.802	0.804
3	0.721	0.724	0.841	0.656	0.818	0.820
...						
t	0.669	0.671	0.928	0.325	0.949	0.952
$t + 1$	0.661	0.664	0.939	0.308	0.954	0.956
$t + 2$	0.657	0.659	0.945	0.297	0.957	0.959
$t + 3$	0.654	0.655	0.950	0.288	0.969	0.961
$t + 4$	0.651	0.653	0.953	0.282	0.960	0.962
$t + 5$	0.650	0.651	0.955	0.277	0.962	0.963

Table 2.5: Moments of the 3D HOG feature descriptor treated as a thin-shell object with density given by the value of the HOG (at each step the total mass of this object is scaled to one). In the columnar rafting experiment, the moments slowly transform from all being equal to two smaller moments (m_1 and m_2) and one large one (m_3). In the layered rafting experiment, the moments slowly transform from being similar to one smaller moment (m_1) and two larger ones (m_2 and m_3).

microscopy, see [32] [33]. For the HOG descriptors used in this chapter though, it is hoped that these issues can be largely avoided. First, the descriptors here are not useful for spatially correlating images. Secondly, if image acquisition can be done rapidly relative to the scale of the images acquired, then drift should be minimal. Images collected with faster acquisition times have more shot noise, but the scale of the Gaussian filter in the

derivative approximation of the HOG can be adjusted to accommodate for this noise.

If drift cannot be avoided, naturally the HOGs will be affected. To study the types of error that would appear in the HOG descriptors in the presence of drift, the TEM micrographs in Figure 2.2, the BSE micrographs in Figure 2.3, and the simulated data set in Figure 2.5 were all subjected to artificial drift. In the case of the TEM and BSE micrographs, this corresponds directly with sample drift in the microscope. In the case of the 3D simulated data, which is nothing but a stack of 2D slices, this corresponds more closely with errors in registering a 3D stack of images.

For the micrographs in Figure 2.2 and Figure 2.3, a sequence of shearing distortions were applied that shifted the bottom of the micrograph horizontally some distance with respect to the top. This was done for a range of distortions between zero and five percent of the total width of the images. These experiments were repeated for a vertical shearing distortion of the same length in pixels as the horizontal distortion. For each distortion, the feature descriptors given in Table 2.2 and Table 2.3 were recomputed. The maximum relative error in each of the three feature descriptors with respect to the features computed with the undistorted data was 35% for the Square feature, 9.0% for the Layering feature, and 23% for the Circle feature. We would expect that classifiers based on HOG descriptors should be able to handle these errors, as the values of the feature descriptors that discriminate between images are frequently a factor of five or more different (as is indeed the case in Table 2.2 and Table 2.3).

For the 3D simulated data set in Figure 2.5, a similar set of shear distortions (up to five percent) were tested. For the columnar microstructure, the distortion direction was chosen to be perpendicular to the length of the columns (so the columns leaned to one side) and for the layered microstructure the distortion direction was chosen in the normal direction of the layers (so the layers were no longer level). For each level of distortion, the moments given in Table 2.5 were recomputed. The maximum relative error in each

of the three moments with respect to the moments of the undistorted data was 8.2% for the smallest moment, 2.5% for the middle moment, and 1.6% for the largest moment. Similarly to the 2D case, these errors do not change the conclusions of the analysis of Table 2.5 (the errors in the largest moments are most important, and they are smallest).

2.5 Conclusions

This chapter demonstrates that in many types of basic microstructure analysis it is possible to employ an histogram of oriented gradients feature vector in place of difficult to compute segmentation statistics. While the HOG has limitations, it is easy to compute and is more robust to common noise sources in electron microscopy techniques and can be applied in a number of interesting applications in both 2D and 3D datasets.

Chapter 3

Resonant Ultrasound Spectroscopy ¹

3.1 Introduction

Ultrasonic techniques such as Resonant Ultrasound Spectroscopy (RUS) provide the most accurate characterization of elastic properties [34, 35], as well as superior precision and repeatability compared to static methods [36]. 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 [34, 37]. 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 [34, 37]. RUS methods do not require a plane-wave assumption [34], nor do they require alternative experimental procedures when characterizing elastically anisotropic materials.

Modern experimental procedures for RUS are discussed in great detail by Migliori

¹Reproduced from: Ben Bales, Brent R. Goodlet, William C. Lenthe, Linda Petzold, and Tresa M. Pollock. “Bayesian inference of elastic properties with resonant ultrasound spectroscopy”. In: *The Journal of the Acoustical Society of America* 143.1 (2018), pp. 7183, with the permission of AIP Publishing.

et al. [35, 38, 39], 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 [35]. 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 [34, 35]. 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 computational in nature.

Beyond inverting elastic properties with RUS data, crystallographic orientation can also be determined, as briefly detailed by Sarrao et al. [40]. 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 methods are most often employed to measure the crystallographic orientation of an RUS specimen [35, 40], but this adds considerable complexity and cost to experimental methods as well as an additional source of measurement error. Sarrao et al. [40] 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 towards 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.

3.1.1 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 3D volume of material. Therefore, approximate (numerical) methods must be employed [35]. Meaningful contributions of numerous researchers over many decades including Holland (1968) [41], Demarest (1971) [42], Ohno (1976) [43], and Visscher et al. (1991) [44] 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 [44], 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 [35].

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 [35]. 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, elas-

tic property inversion of RUS data has been accomplished via nonlinear least square optimization algorithms [35, 45, 46]. For this task, the Levenberg-Marquardt (LM) algorithm popularized by Migliori et al. [35, 38] combines a modified Newton method with the steepest descent algorithm [38, 40, 45] 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 [47]. 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 [48, 49]. This problem is not unique to the LM algorithm, as similar concerns exist for genetic algorithms [48, 50, 51].

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. [52] 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 [52]. 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” [52]. 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.

3.1.2 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 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. [1], 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.

3.2 Experimental Methods

3.2.1 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) then carefully ground with 800-grit sandpaper to remove the superficial EDM damage layer. The Ti-64 specimen measured 7.753 x 9.057 x 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 x 13.953 x 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% [38]. The only notable divergence between this work and the typical specimen fabrication procedures outlined by Migliori and Sarrao [35] 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.

3.2.2 Resonant Ultrasound Spectroscopy

RUS Experimental Setup

RUS data were collected using commercially-available RUS equipment developed by the Vibrant Corporation consisting of three primary components: a transceiver, a piezo-

electric transducer (PT) cradle, and a computer control unit. Figure 3.1 shows the configuration of the PT cradle, comprised of three custom-built omnidirectional piezoelectric transducers 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, 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 [34, 35].




Figure 3.1: Piezoelectric transducers configured into a cradle supporting the CMSX-4 parallelepiped specimen.

The “drive PT” in Figure 3.1 is driven with a swept sinusoidal signal from the transceiver to excite the specimen to resonate, while two “receive PTs” convert vibra-

tions 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. With the magnitude of the vibrations amplified by hundreds to thousands of times the driving force amplitude, depending on the ultrasonic attenuation (i.e. damping characteristics) of the material [34, 35]. 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 3.2 shows a broadband resonance spectrum plot collected from the CMSX-4 specimen, with 53 resonance modes across the 200 kHz broadband.

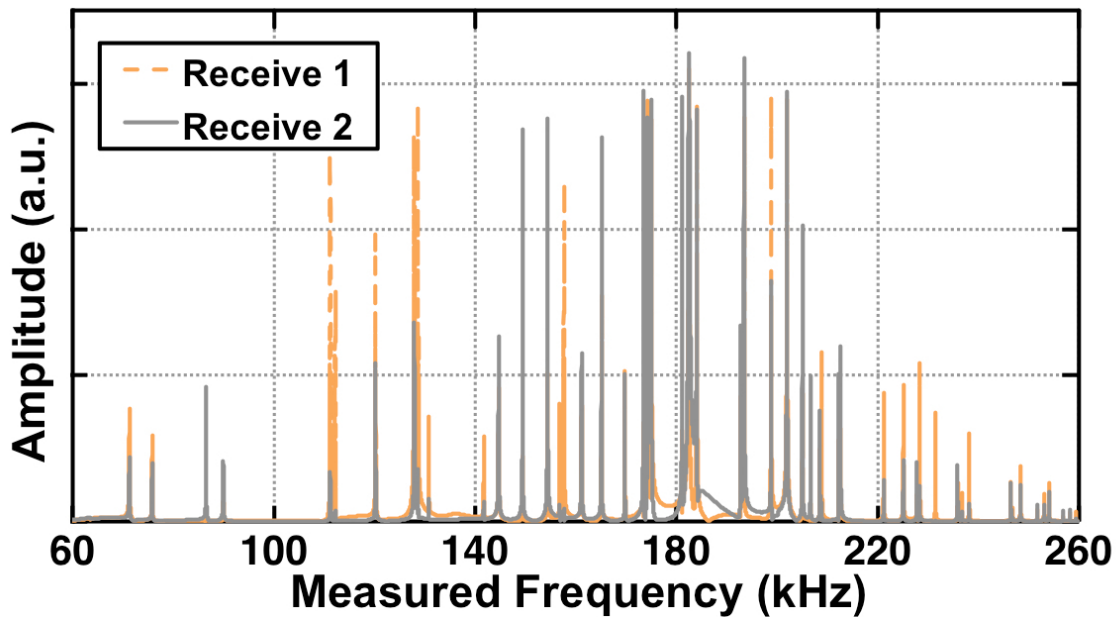


Figure 3.2: RUS broadband spectrum plot collected from the CMSX-4 specimen.

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 [38, 53]. Configuring the piezoelectric transducers into a fixed cradle as detailed by Figure 3.1 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 Figure 3.2 are given with arbitrary units (a.u.). When amplitude information is necessary, to measure acoustic attenuation for instance, parallelepiped specimens are often pinched between two (often planar) PTs [53]. Note that the corners of a parallelepiped are the optimal location for excitation and measurement of resonance frequencies [38]; but for the purpose of determining elastic constants the cradle configuration has proven itself simple and effective.

RUS Data Collection

Broadband RUS spectra were collected at room temperature and at standard atmospheric pressure in accordance with ASTM standard practice 2534-15 [54]. 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 [35], therefore considerable attention was directed towards 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 Figure 3.1. Nevertheless, each broadband measurement was repeated five 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.

3.2.3 Crystallographic Orientation Measurements

To measure the orientation of the CMSX-4 crystal reference frame, a series of x-ray diffraction (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 micron, 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 one degree.

3.3 Computations

The following sections discuss the development of a statistical model for characterizing elastic constants via RUS. Section 3.3.1 reviews the basic mechanics of the system, which

are very similar to previous RUS works [38, 44, 1]. Section 3.3.2 introduces the Bayesian RUS model, and Section 3.3.3 describes how the Bayesian computations are carried out.

3.3.1 Forward Model

Data in this experiment is modeled as:

$$X_1, X_2, \dots, X_N = f(c_{11}, c_{12}, \dots) + \xi \quad (3.1)$$

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 [44] and [1], the specimen in the forward model is approximated as an undamped linear harmonic oscillator with free boundary conditions:

$$Ku = M \frac{\partial^2 u}{\partial t^2} \quad (3.2)$$

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 (FEM) approximation to the problem. Taking the Fourier transform in time (to find the steady state solution) yields:

$$Ku = \omega^2 Mu \quad (3.3)$$

with 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 [44] and [1]). 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 [1] are reproduced in Eqs. 3.4-3.7 (using Einstein notation for tensors):

$$\phi_\lambda(x, y, z) = x^n y^m z^l, \quad (3.4)$$

$$\{\lambda = (n, m, l) \mid n, m, l \in \mathbb{N}^0, n + m + l \leq N\}, \quad (3.5)$$

$$K_{i\lambda, k\lambda'} = \sum_{j,i=1}^3 C_{ijkl} \int_V \epsilon_{ij}(\phi_\lambda) \epsilon_{kl}(\phi_{\lambda'}) dV, \quad (3.6)$$

$$M_{i\lambda, k\lambda'} = \rho \delta_{ik} \int_V \phi_\lambda \phi_{\lambda'} dV. \quad (3.7)$$

In Eqs. 3.4-3.7 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:

$$q = \begin{bmatrix} w^2 + x^2 - y^2 - z^2 & 2(xy - wz) & 2(xz + wy) \\ 2(yx + wz) & w^2 - x^2 + y^2 - z^2 & 2(yz - wx) \\ 2(zx - wy) & 2(zy + wx) & w^2 - x^2 - y^2 + z^2 \end{bmatrix} \quad (3.8)$$

$$C_{ijkl} = q_{ip} q_{jq} C'_{pqrs} q_{kr} q_{ls}. \quad (3.9)$$

For more information, see Section 3.2.11 in Bower [55].

3.3.2 Building a Statistical RUS Model

Since the noise (ξ) in Eq. 3.1 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:

$$P(X_0, X_1, \dots, X_N | c_{11}, c_{12}, \dots) \sim \mathcal{N}(f(c_{11}, c_{12}, \dots), \sigma^2). \quad (3.10)$$

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:

$$\underbrace{P(\theta|X)}_{\text{Posterior}} = \frac{\overbrace{P(X|\theta)}^{\text{Likelihood}} \overbrace{P(\theta)}^{\text{Prior}}}{\underbrace{P(X)}_{\text{Prior Predictive}}}. \quad (3.11)$$

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. 3.1, this chapter assumes that 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 consistently [35] 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 towards estimating the lumped variance parameter.

Thus the complete likelihood can be stated as:

$$P(X|\theta) = \prod_i \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{-\frac{(f_i(\theta)-X_i)^2}{2\sigma_i^2}} \quad (3.12)$$

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 $\sigma^2 = \mathcal{U}(0 \text{ kHz}, \infty \text{ kHz})$.

3.3.3 Computing the Posterior (Hamiltonian Monte Carlo)

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 chapter is Hamiltonian Monte Carlo (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, pro-

ceeds 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 $\int X$ plays in this. X is the measured data. It is fixed, and does not change. Sampling only happens over the parameters ($\int \theta$) in the joint distribution. With this choice of Φ , the Metropolis Monte Carlo 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. Hamiltonian Monte Carlo 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 Hamiltonian Monte Carlo algorithm used in this chapter was taken from [56]. The orientation parameters are handled by the special integrator in [57]. This is necessary because, while quaternions are expressed as vectors with four elements, the lengths of these vectors are constrained to be one. The method outline in [57] allows for efficient sampling of parameters like these. The technique in [56] for using multiple steps sizes was critical for achieving efficient sampling.

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 transformation [58]).

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:

$$\frac{\partial \log P(X|\theta)}{\partial c_{11}} = \sum_i \frac{\partial \log P(X|\theta)}{\partial \omega_i} \frac{\partial \omega_i}{\partial c_{11}} \quad (3.13)$$

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. 3.12):

$$\frac{\partial \log P(X|\theta)}{\partial \omega_i} = \frac{-(\omega_i - X_i)}{\sigma^2}. \quad (3.14)$$

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. 3.3. 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 [59]:

$$\frac{\partial \omega_i^2}{\partial c_{11}} = \nu_i^T \frac{\partial K}{\partial c_{11}} \nu_i \quad (3.15)$$

$$\frac{\partial \omega_i}{\partial c_{11}} = \frac{1}{2\omega_i} \frac{\partial \omega_i^2}{\partial c_{11}} = \frac{1}{2\omega_i} \nu_i^T \frac{\partial K}{\partial c_{11}} \nu_i . \quad (3.16)$$

The partial derivatives of the eigenvalues requires derivatives of the stiffness matrix from Eq. 3.3. These can be obtained from the construction in Eq. 3.6:

$$\frac{\partial K_{i\lambda, k\lambda'}}{\partial c_{11}} = \sum_{j,i=1}^3 \frac{\partial C_{ijkl}}{\partial c_{11}} \int_V \epsilon_{ij}(\phi_\lambda) \epsilon_{kl}(\phi_{\lambda'}) dV , \quad (3.17)$$

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

$$C_{ijkl} = \begin{bmatrix} c_{11} & c_{12} & c_{12} & 0 & 0 & 0 \\ c_{12} & c_{11} & c_{12} & 0 & 0 & 0 \\ c_{12} & c_{12} & c_{11} & 0 & 0 & 0 \\ 0 & 0 & 0 & c_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & c_{44} & 0 \\ 0 & 0 & 0 & 0 & 0 & c_{44} \end{bmatrix} \quad (3.18)$$

and

$$\frac{\partial C_{ijkl}}{\partial c_{11}} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}. \quad (3.19)$$

Combining Eqs. 3.14, 3.16, 3.17, and 3.19 give the necessary expression for $\frac{\log P(X|\theta)}{\partial c_{11}}$. 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:

$$A = \frac{2C_{44}}{C_{11} - C_{12}} \quad (3.20)$$

as defined by Zener [60].

Noise term (σ)

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

$$\frac{\partial \log P(X|\theta)}{\partial \sigma} = -\frac{N}{\sigma} + \sum_i^N \frac{(\omega_i - X_i)^2}{\sigma^3}. \quad (3.21)$$

Lattice-Specimen Orientations

As stated in Section 3.3.1, the lattice-specimen orientation is parameterized as a passive unit quaternion. Though Euler angles are commonly used in electron backscatter diffraction experiments to characterize the lattice-specimen orientation, this parameterization is not trivial to work with in Bayesian inference. In particular, an adjustment must be applied to the posterior density to achieve a uniform prior on the orientation parameter and allowances must be made for the fact that the basic HMC sampler described in [56] operates in an unconstrained coordinate space instead of the constrained space of Euler angles [61]. Quaternions were used instead.

Quaternions, though expressed in four dimensions (w, x, y, z) , live on a three dimensional manifold characterized by:

$$w^2 + x^2 + y^2 + z^2 = 1 . \quad (3.22)$$

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 [57] is used. Without going into detail, Byrne [57] adjusts random momentum generation and timestepping in HMC to keep the quaternion parameters on the manifold (Eq. 3.22). These quaternions can be converted back to Euler angles as needed.

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 $\frac{\partial C_{ijkl}}{\partial w}$ (simply replace c_{11} with w in Eqs. 3.14, 3.16, and 3.17). These require partial derivatives

of Eq. 3.6 and Eq. 3.8:

$$\begin{aligned} \frac{\partial C_{ijkl}}{\partial w} = & \frac{\partial q_{ip}}{\partial w} q_{jq} C'_{pqrs} q_{kr} q_{ls} + \\ & q_{ip} \frac{\partial q_{jq}}{\partial w} C'_{pqrs} q_{kr} q_{ls} + \\ & q_{ip} q_{jq} C'_{pqrs} \frac{\partial q_{kr}}{\partial w} q_{ls} + \\ & q_{ip} q_{jq} C'_{pqrs} q_{kr} \frac{\partial q_{ls}}{\partial w} \end{aligned} \quad (3.23)$$

$$\frac{\partial q}{\partial w} = 2 \begin{bmatrix} w & -z & y \\ z & w & -x \\ -y & x & w \end{bmatrix} . \quad (3.24)$$

3.4 Results

3.4.1 Polycrystalline Ti-64

Polycrystalline Ti-64 elastic constants were inverted using the geometries given in Section 3.2.1 along with the first thirty measured resonance modes. The Ti-64 material was produced in a large-scale billet via an isothermal multi-step "abc" forging process, with the material upset forged along three orthogonal (abc) axes until a cumulative strain of three was achieved. As Zherebtsov et al. [62] report, the result was a homogeneous (texture free) microstructure of globular alpha and alpha-beta subgrains, with an average grain size of approximately 0.4 microns. With such a refined microstructure, smaller than the resonance mode wavelengths by over an order of magnitude, the abc forged Ti-64 behaves as a homogeneous elastic body with isotropic elastic symmetry. This isotropy assumption, however, was not enforced. The standard deviation σ of the noise was constrained to be positive. The initial conditions for the sampler were chosen as $c_{11} = 2.0$,

$A = 1.0$, $c_{44} = 1.0$, and $\sigma = 5.0$ (c_{11} and c_{44} are specified in units of 10^{11} Pa, and σ is specified in units of kHz). The final inverted parameters are given in Table 3.1 along with reference values from Fisher and Renken [63]. Note that Fisher and Renken [63] 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 [64]. Clearly the inverted value of $A = 1.000 \pm 0.002$ demonstrates the Ti-64 material is effectively isotropic.

Table 3.1: Summary of estimated parameters for Ti-64 specimen alongside reference values.

Parameter	Reference[63]	Bayesian estimate ($\mu \pm \sigma$)
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 ¹	$1.000^2 \pm 0.002$

The full estimated posterior distributions for the four parameters (two elastic constants, the anisotropy ratio, and the error term, σ) are shown in Figure 3.3. All the parameters (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 3.1.

In Bayesian modeling, the two basic tools for validating an inversion are traceplots, which are plots of the sequence of states sampled by the MCMC sampler, and posterior predictive distribution plots. The traceplots for the last four thousand posterior samples for the Ti-64 specimen are plotted in Figure 3.4. Extracting information from a traceplot is fairly straightforward. If the distribution of samples in the traceplot is stationary for a long period of time, the chain is assumed (but not guaranteed) to be sampling from the true posterior. Traceplots are usually the easiest place to detect modeling problems. If, for instance, a parameter enters an invalid range or shows a bimodal tendency, it is usually easy to spot in the traceplots. The traceplots in Figure 3.4 show nothing unusual,

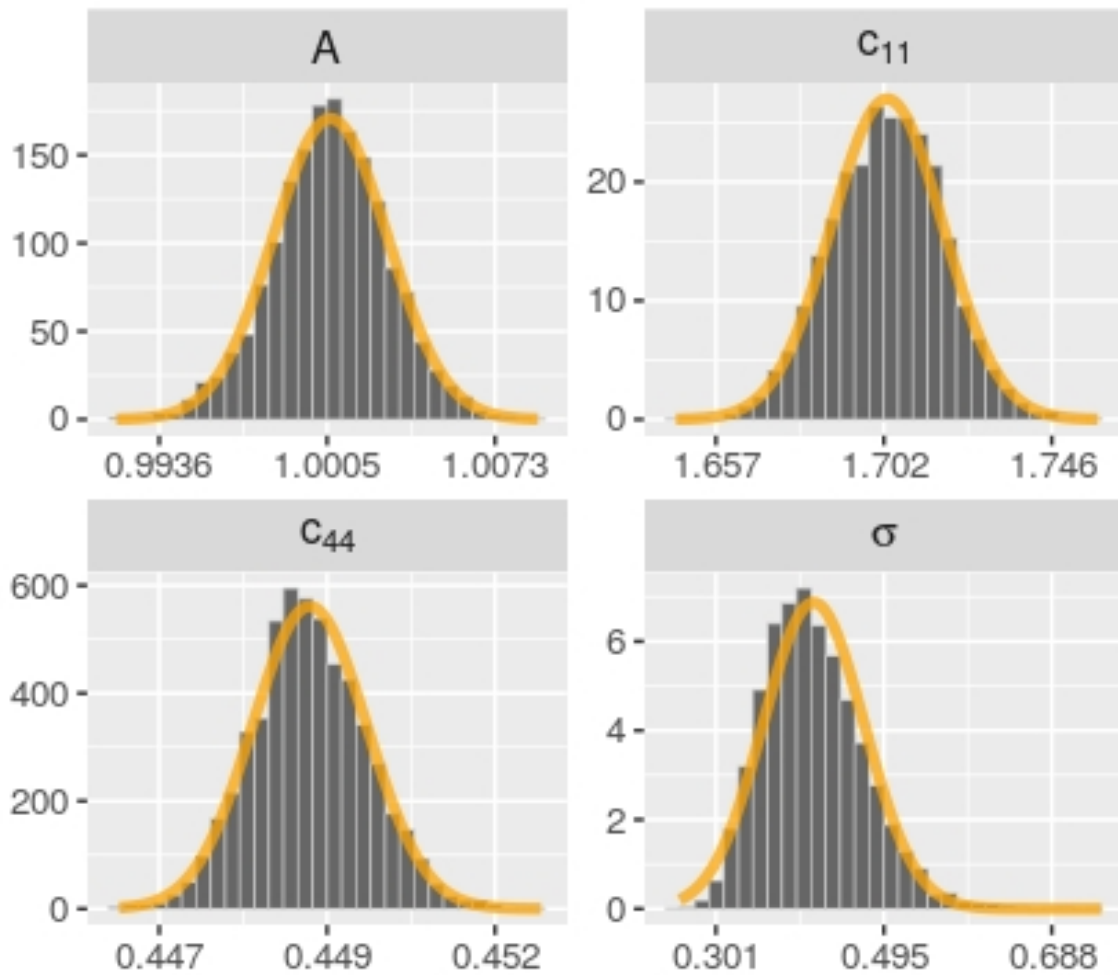


Figure 3.3: Approximate posterior distributions (normalized to a PDF scale) for c_{11} , A , c_{44} , and σ computed from four thousand HMC samples. Normal PDF fits are superimposed to justify the use of mean and standard deviation to characterize the posteriors. c_{11} and c_{44} are given in units of 10^{11} Pa, and σ is given in units of kHz.

suggesting that the samples themselves come from the true posterior of the model.




Figure 3.4: Traceplots (last four thousand samples) for the estimated parameters for the polycrystalline Ti specimen. c_{11} and c_{44} are given in units of 10^{11} Pa, and σ is given in units of kHz.

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

Table 3.2: Ti-64 specimen measured and posterior predictive modes, $(\mu \pm \sigma)$.

Mode	Measured (kHz)	Q Ave. ¹	Posterior ² (kHz)
1	109.076 \pm 0.027	3,832	108.86 \pm 0.43
2	136.503 \pm 0.083	3,209	135.97 \pm 0.43
3	144.899 \pm 0.051	3,157	144.42 \pm 0.41
4	184.926 \pm 0.046	6,482	184.60 \pm 0.44
5	188.476 \pm 0.029	4,781	187.99 \pm 0.43
6	195.562 \pm 0.032	6,828	195.56 \pm 0.44
7	199.246 \pm 0.029	7,016	199.19 \pm 0.44
8	208.460 \pm 0.078	6,162	208.08 \pm 0.43
9	231.220 \pm 0.030	7,241	231.63 \pm 0.44
10	232.630 \pm 0.042	6,759	232.47 \pm 0.44
11	239.057 \pm 0.033	2,874	239.09 \pm 0.45
12	241.684 \pm 0.077	7,148	242.24 \pm 0.44
13	242.159 \pm 0.030	7,367	242.53 \pm 0.44
14	249.891 \pm 0.079	7,436	249.89 \pm 0.42
15	266.285 \pm 0.097	7,771	267.20 \pm 0.46
16	272.672 \pm 0.069	8,376	272.49 \pm 0.44
17	285.217 \pm 0.037	7,984	285.04 \pm 0.45
18	285.670 \pm 0.095	8,008	285.65 \pm 0.43
19	288.796 \pm 0.023	10,400	289.23 \pm 0.45
20	296.976 \pm 0.060	7,651	296.77 \pm 0.42
21	301.101 \pm 0.030	8,687	301.60 \pm 0.46
22	303.024 \pm 0.053	8,854	303.03 \pm 0.43
23	305.115 \pm 0.058	9,296	305.02 \pm 0.44
24	305.827 \pm 0.067	9,509	305.26 \pm 0.47
25	306.939 \pm 0.034	10,706	306.39 \pm 0.45
26	310.428 \pm 0.024	6,946	310.15 \pm 0.44
27	318.000 \pm 0.042	11,199	317.51 \pm 0.45
28	319.457 \pm 0.040	7,600	319.88 \pm 0.47
29	322.249 \pm 0.053	10,011	322.17 \pm 0.42
30	323.464 \pm 0.027	9,986	322.93 \pm 0.44

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. 3.12 corresponding to the missing mode is removed (since that data is not available), and the inference rerun for improved results.

The posterior predictive distributions for the Ti specimen are shown in Table 3.2, 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 traceplots, 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.

3.4.2 Single Crystal CMSX-4

The single crystal CMSX-4 elastic constants and orientation were inverted using the geometries given in Section 3.2.1 and the first thirty measured resonance modes. Instead of running a single, long HMC chain, the inference was broken into warmup and post-warmup stages (as described in [65, Chapter 12]). In the warmup stage, four chains are run from a fixed initial condition $c_{11} = 2.0$, $A = 1.0$, $c_{44} = 1.0$, and $\sigma = 5.0$ to find reasonable parameter estimates (c_{11} and c_{44} are specified in units of 10^{11} Pa, and σ is specified in units of kHz). In the warmup stage the sampler needs to move around very-low probability areas of parameter space, and for stability the HMC timestep 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 timesteps 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 Figure 3.5, while the post-warmup traceplots are shown in Figure 3.6. 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 timestep). Table 3.3 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 Figure 3.6. 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

Table 3.3: Summary of CMSX-4 parameter estimates and 95% posterior interval from each of the four HMC chains of 1000 samples each.

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

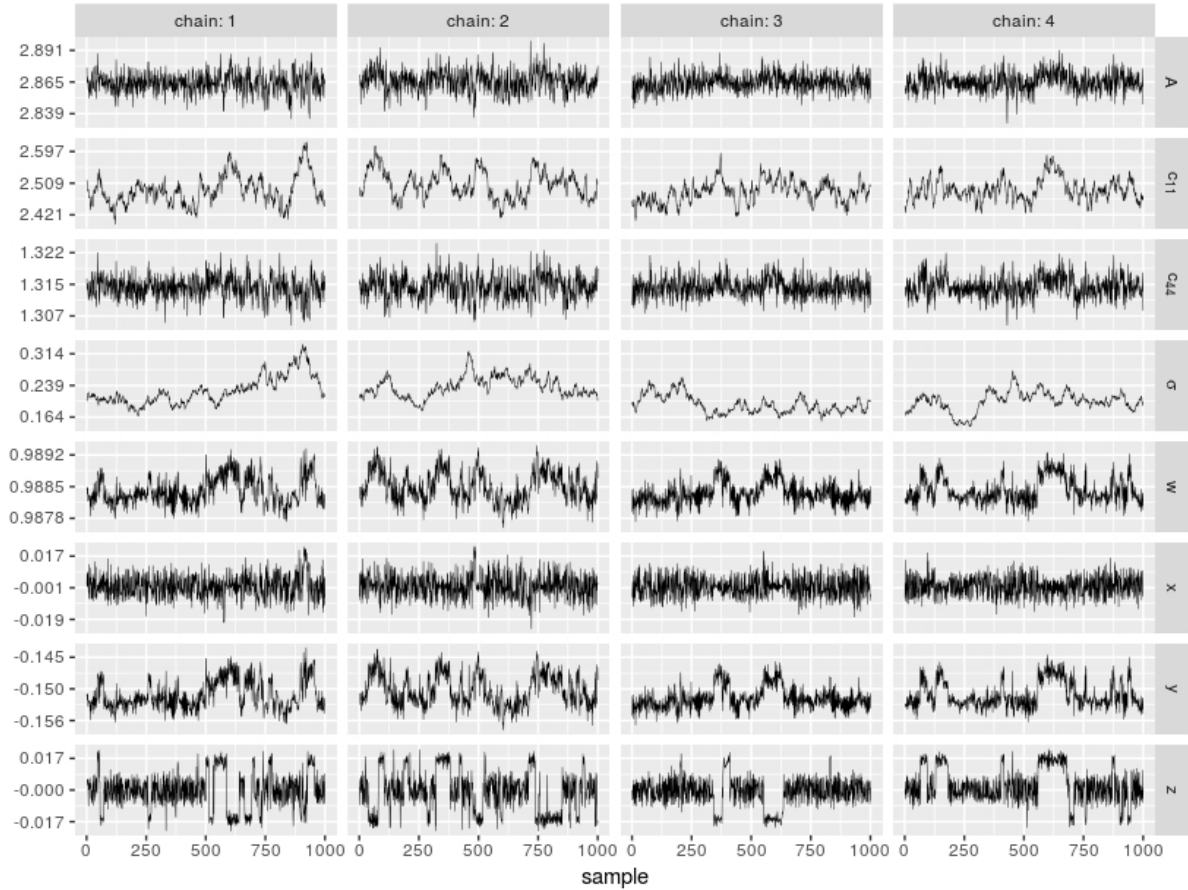


Figure 3.5: Warmup traceplots (four different chains of one thousand samples each) for the estimated parameters in the single crystal CMSX-4 specimen. The rough means of the parameters in these chains were used to initialize four chains with larger time steps to more efficiently explore the posterior. c_{11} and c_{44} are given in units of 10^{11} Pa, and σ is given in units of kHz.

parameter estimate. For example, the standard deviations of the c_{11} and c_{44} parameters are 4.0 GPa and 0.26 GPa respectively. This demonstrates the superior precision of the c_{44} estimate as compared to c_{11} estimate, and is consistent with RUS-based elastic property inversions reported previously by Migliori et al. [38].

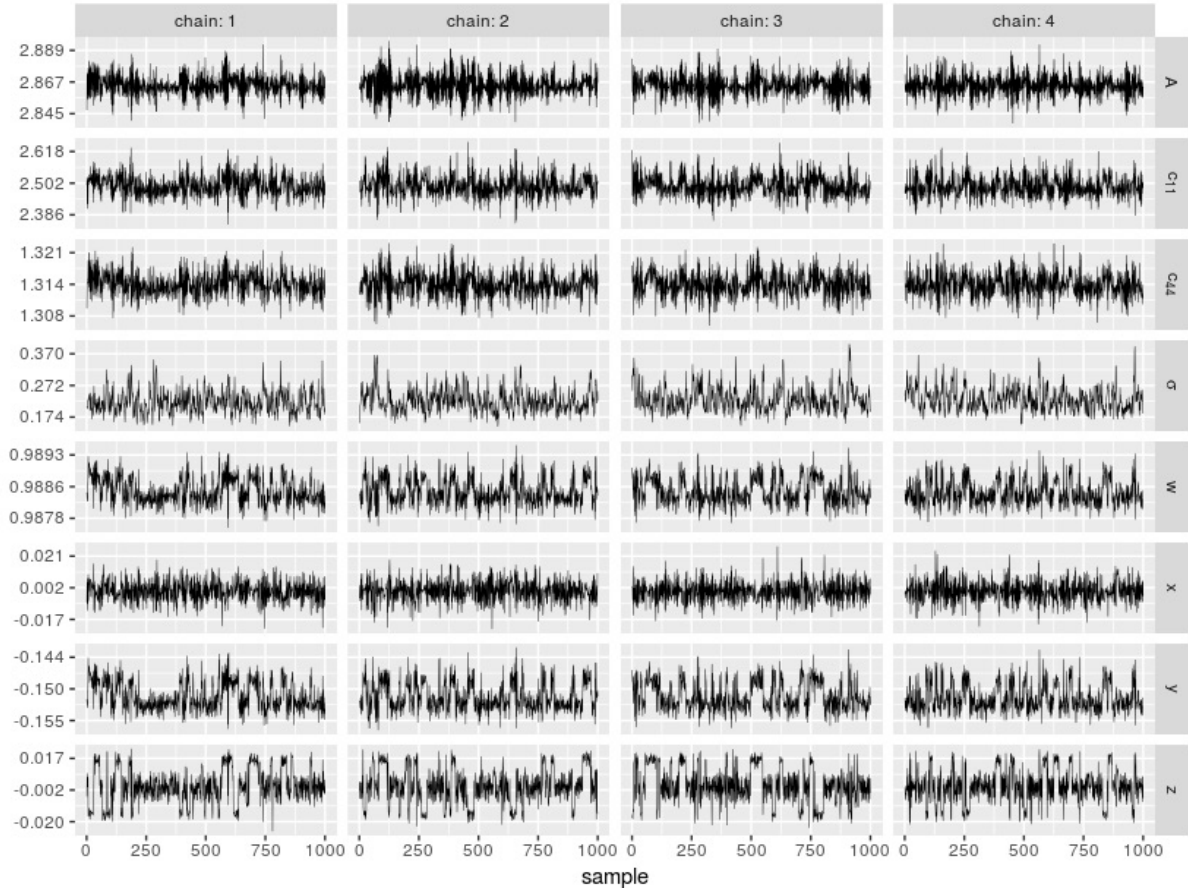


Figure 3.6: Sampling traceplots (four different chains of one thousand samples each) for the estimated parameters of the single crystal CMSX-4 specimen. The elastic constants appear to be unimodal, but the orientation parameters are multimodal. c_{11} and c_{44} are given in units of 10^{11} Pa, and σ is given in units of kHz.

As stated in Section 3.4.1, the posterior predictive distributions elucidate the quality of the fit and are provided for the CMSX-4 specimen in Table 3.4. Alongside these data are the measured frequencies and their standard deviation, as well as the quality

Table 3.4: CMSX-4 measured and posterior predictive modes ($\mu \pm \sigma$).

Mode	Measured (kHz)	Q Ave. ¹	Posterior ² (kHz)
1	71.259 \pm 0.044	2,279	71.22 \pm 0.24
2	75.759 \pm 0.025	2,864	75.69 \pm 0.23
3	86.478 \pm 0.030	3,524	86.46 \pm 0.26
4	89.947 \pm 0.048	3,077	90.00 \pm 0.24
5	111.150 \pm 0.033	4,155	111.06 \pm 0.26
6	112.164 \pm 0.027	4,492	112.01 \pm 0.29
7	120.172 \pm 0.033	5,048	120.32 \pm 0.26
8	127.810 \pm 0.042	4,462	127.98 \pm 0.25
9	128.676 \pm 0.040	4,747	128.64 \pm 0.24
10	130.740 \pm 0.034	5,342	130.75 \pm 0.24
11	141.700 \pm 0.037	5,299	141.79 \pm 0.25
12	144.504 \pm 0.065	5,603	144.36 \pm 0.23
13	149.401 \pm 0.025	5,918	149.52 \pm 0.26
14	154.351 \pm 0.026	5,942	154.42 \pm 0.25
15	156.782 \pm 0.028	6,761	156.97 \pm 0.25
16	157.555 \pm 0.041	6,123	157.57 \pm 0.25
17	161.088 \pm 0.063	6,284	160.97 \pm 0.24
18	165.103 \pm 0.022	6,756	165.21 \pm 0.28
19	169.762 \pm 0.044	6,995	169.77 \pm 0.27
20	173.449 \pm 0.050	7,212	173.28 \pm 0.26
21	174.117 \pm 0.024	6,437	174.13 \pm 0.26
22	174.906 \pm 0.054	6,916	174.68 \pm 0.26
23	181.120 \pm 0.042	6,632	181.54 \pm 0.25
24	182.459 \pm 0.037	7,475	181.87 \pm 0.25
25	183.986 \pm 0.042	7,837	183.81 \pm 0.28
26	192.681 \pm 0.032	7,197	192.83 \pm 0.24
27	193.436 \pm 0.021	7,113	193.71 \pm 0.27
28	198.794 \pm 0.034	7,249	198.95 \pm 0.25
29	201.902 \pm 0.032	8,529	201.89 \pm 0.27
30	205.015 \pm 0.031	8,808	204.85 \pm 0.33

factor for the first 30 resonance modes. As before, the one measured value outside the 95% posterior predictive interval (mode 24) is highlighted in gray, 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.

Table 3.5: Summary of estimated parameters for the CMSX-4 specimen alongside reference values.

Parameter	Reference[66]	Bayesian estimate ($\mu \pm \sigma$)
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

Estimated posterior distributions of the elastic constants (c_{11} , A , and c_{44}) and the noise parameter (σ) are shown in Figure 3.7. 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 3.5. Reference CMSX-4 elastic constants from [66] 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 coordinates [67] before plotting in Figure 3.8. 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

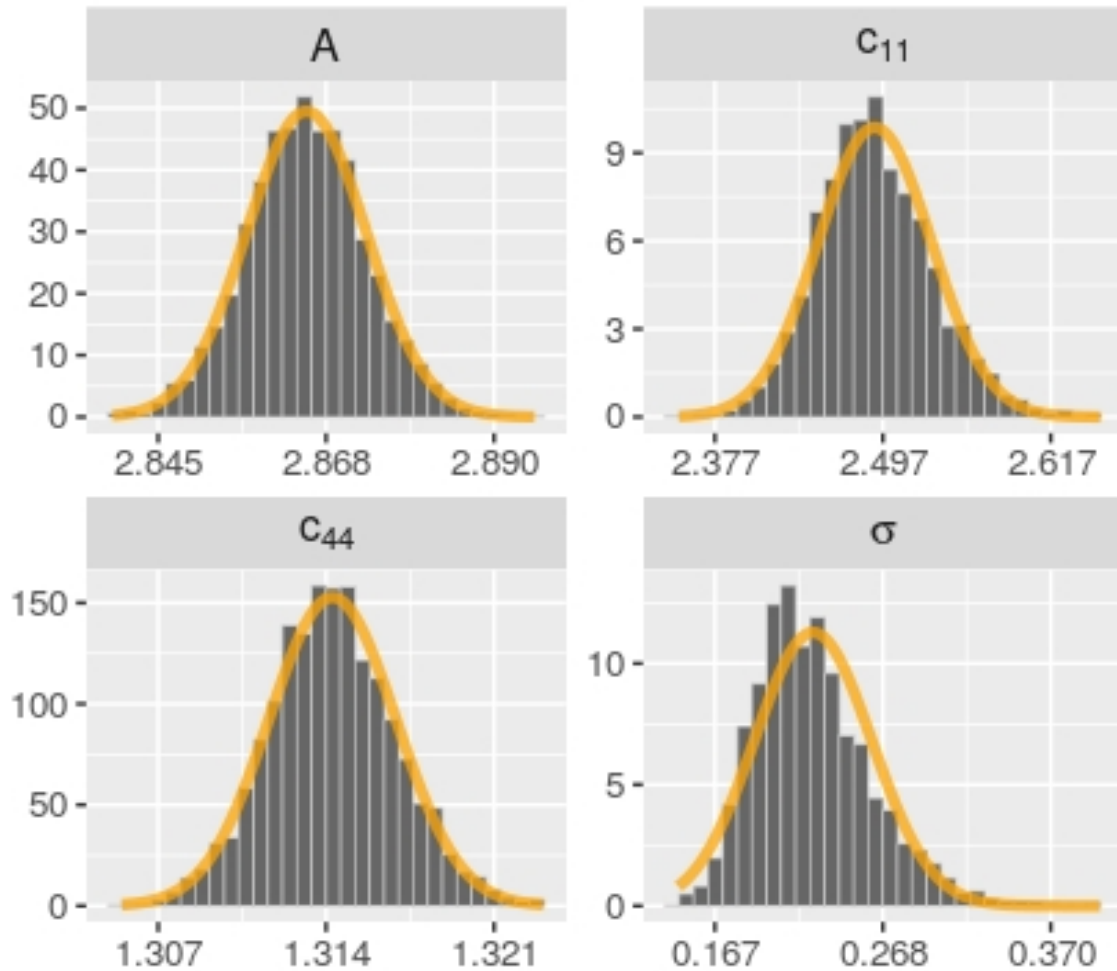


Figure 3.7: Approximate CMSX-4 posterior distributions for c_{11} , A , c_{44} , and noise (σ) parameters, computed from four chains of one thousand HMC samples each. Normal PDF fits are superimposed to justify the use of mean and standard deviation to characterize the posteriors. c_{11} and c_{44} are given in units of 10^{11} Pa, and σ is given in units of kHz.

of approximately one degree as noted in Section 3.2.3 could easily place the measured orientation within a higher-likelihood region of the posterior; given that a rotation of one degree 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 Figure 3.8. Given the overwhelming majority of misorientation angles were calculated as between 1.0 and 2.5 degrees, our confidence in estimating crystallographic orientation from RUS-measured resonance frequencies is high, with results summarized in Table 3.6

Table 3.6: Summary of CMSX-4 specimen crystal orientation from XRD measurement and Bayesian estimate.

Parameter	XRD measurement	Bayesian estimate ($\mu \pm \sigma$)
w	0.987	0.9884 ± 0.0003
x	-0.005 26	0.000 ± 0.0062
y	-0.158	-0.1512 ± 0.0024
z	0.0164	0.001 ± 0.010

To demonstrate that inferring the crystallographic orientation was necessary, an inversion was run without these degrees of freedom. Figure 3.9 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.




Figure 3.8: 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 one degree of rotation.




Figure 3.9: Traceplots for the estimated parameters in the single crystal CMSX-4 specimen without including misorientation estimation. c_{11} and c_{44} are given in units of 10^{11} Pa, and σ is given in units of kHz.

3.4.3 Computation Efficiency

The Ti-64 inversion ran for approximately eight hours on a quad core Intel i7-2600k desktop computer, while the CMSX-4 inversion took a few days. Each HMC sample takes fifty forward evaluations of the RUS model (so one hundred thousand 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 a thousand HMC samples, and two thousand samples from the posterior are computed, the calculation should take a little under two days.

3.4.4 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 [38]. 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 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 [35]. Finally, accounting for missing or spurious modes as described in [1] may facilitate

inversions from RUS spectra that are complicated by a variety of factors including low stiffness or high damping.

3.5 Conclusions

A novel and robust RUS inversion framework for characterizing elastic properties and crystal orientation of parallelepiped specimens using a Bayesian modeling approach and Hamiltonian Monte Carlo 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 degrees 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.

Chapter 4

Selecting the Metric in Hamiltonian Monte Carlo

4.1 Introduction

Hamiltonian Monte Carlo (HMC) methods [68, 56, 69] have proven to be very effective for use in high level Bayesian inference packages like Stan [3] and PyMC3 [70]. The usefulness of HMC is limited by how well it can be adapted to a problem's posterior geometry. Ideally this is done dynamically, but this introduces additional complexity in the form of higher order derivatives and non-linear solves that make it difficult to scale the approaches to large numbers of parameters [71, 72].

For problems where posterior curvature varies little, this dynamic adaptation is also unnecessary. The simpler approach is to build a fixed, linear coordinate transformation during the Monte Carlo warmup phase and use that for sampling. This can be formalized as choosing a Euclidean metric for the sample space [69]. This method would be expected to work on models with unimodal, approximately normal, possibly correlated posteriors.

The basic metric comes from the observation that, when sampling a multivariate

normal posterior, if the metric is set equal to the posterior covariance, sampling becomes much easier. Thus, the metric normally chosen is a variant of the sample covariance [56, 69].

Estimating a full covariance matrix can be just as hard as sampling the desired posterior distribution, thus estimating the full covariance so that the posterior can be sampled is a chicken and egg problem. If the estimate is made badly, it will not improve sampler efficiency. In light of this difficulty, the inverse of a diagonal covariance approximation is often used as the metric because of the ease with which it can be constructed [3]. This is useful for some problems, but highly correlated posteriors, even in low dimensions, will render this ineffective.

The properties of the leapfrog integrator inside HMC can be used as a proxy to understand the effectiveness of HMC on different problems, in terms of the number of effective sample size per second (ESS/s). This line of reasoning leads to the two contributions of this paper: a selection criterion for predicting the effectiveness of a metric on a problem, and a new metric based on the Hessian of the log density that requires fewer warmup draws than the sample covariance to be effective.

This paper is organized as follows. Section 4.2 provides the motivation behind the selection criterion and the adaptation, and Section 4.3 covers the implementation details. Section 4.4 demonstrates the utility of the selection criterion and the robustness of the adaptation on a variety of problems.

4.2 Preliminaries

4.2.1 Hamiltonian Monte Carlo

Hamiltonian Monte Carlo methods begin by defining a Hamiltonian

$$H(q, p) = -\log P(X|q)P(q) + \frac{1}{2}p^T M^{-1}p \quad (4.1)$$

with N position variables, q , and N momentum variables, p , [68, 56, 69]. The position variables correspond to parameters in the distribution that is to be sampled, and transitions in the Markov chain are generated from simulations of trajectories on the Hamiltonian.

In basic HMC, the Hamiltonian is simulated via the leapfrog method [73, 74]. The length of the integrals can either be fixed or adapted dynamically [75, 69]. The Euclidean metric, M , is the degree of freedom that can be adapted in a problem specific way to make sampling efficient.

4.2.2 Linearized Dynamics

The simplest way to analyze the characteristics of a Hamiltonian system is by examining a linearization of the dynamics around a fixed point [76]. In the context of Bayesian inference and HMC, looking at center fixed points correspond to looking at a multivariate normal (Laplace) approximation of the posterior, or a quadratic approximation of the potential energy of the Hamiltonian around the maximum a-posteriori estimate [77].

The dynamics of the Hamiltonian in Eq. 4.1 are [76, 78]

$$\begin{bmatrix} q \\ p \end{bmatrix}' = \begin{bmatrix} M^{-1}p \\ -\nabla_q H \end{bmatrix} = \begin{bmatrix} M^{-1}p \\ \nabla \log P(X|q)P(q) \end{bmatrix}. \quad (4.2)$$

The linearization of this Hamiltonian around a fixed point (q_0, p_0) is

$$\begin{bmatrix} \Delta q \\ \Delta p \end{bmatrix}' = \begin{bmatrix} 0 & M^{-1} \\ -\nabla_{qq}^2 H(q_0) & 0 \end{bmatrix} \begin{bmatrix} \Delta p \\ \Delta q \end{bmatrix}, \quad (4.3)$$

where $\Delta q = q - q_0$, and $\Delta p = p - p_0$. The eigenvalues of the matrix in Eq. 4.3 are given by the zeros of

$$\det(M^{-1}\nabla_{qq}^2 H(q_0) + \omega^2 I) \quad (4.4)$$

and describe the dynamics of the linearized system.

To compute the zeros of this determinant, we can look at the eigenvalues of the simplified system

$$M^{-1}\nabla_{qq}^2 H(q_0)x_i = \lambda_i x_i. \quad (4.5)$$

The eigenvalues of Eq. 4.3 are given by the pairs $\pm\sqrt{\lambda_i}i$. Because the metric must be positive definite, it can be decomposed as $LL^T = M^{-1}$. The substitution $x = Ly$ gives

$$L^T\nabla_{qq}^2 H(q_0)Ly_i = \lambda_i y_i. \quad (4.6)$$

This transformed space has the same eigenvalues as the original (Eq. 4.5). Because $L^T\nabla_{qq}^2 H(q_0)L$ is symmetric, the eigenvalues of Eqs. 4.5 and 4.6 are real.

The dynamics in Eq. 4.3 govern the efficiency of integration in the original Hamiltonian problem and hence the performance of HMC, and can be characterized by the eigenvalue problem in Eq. 4.6.

4.2.3 Leapfrog Stability

Assuming the posterior being sampled is a multivariate normal distribution with fixed covariance, Σ , the Hamiltonian will be

$$H(q, p) = \frac{1}{2}q^T \Sigma^{-1}q + \frac{1}{2}p^T M^{-1}p. \quad (4.7)$$

Given the eigendecomposition $\nabla_{qq}^2 H = \Sigma^{-1} = V\Lambda V^T$, the metric $M = VV^T$ will diagonalize the Hessian in Eq. 4.6. For N parameters (and assuming Σ is strictly positive definite), this reduces the N dimensional dynamical system to N one-dimensional oscillators.

Each of these oscillators can be run separately. From [73], the leapfrog stability limit for each of the n oscillators is

$$\Delta t = \frac{2}{\sqrt{|\lambda|_n}}, \quad (4.8)$$

where λ_n is the square of the angular frequency.

When integrating the complete system, the overall timestep is limited by the smallest single timestep, the size of which is limited by the largest eigenvalue, $|\lambda|_{\max}$. Thus, when integrating in the direction of least curvature, the timestep is too small by a factor inversely proportional to the square root of the ratio of the eigenvalues, compared to an ideal case

$$\frac{\Delta t_{\text{fast}}}{\Delta t_{\text{slow}}} \propto \sqrt{\frac{|\lambda|_{\max}}{|\lambda|_{\min}}}. \quad (4.9)$$

This is the square root of the condition number of Eqs. 4.5 and 4.6. In this way, the problem of selecting the metric can be recast in terms of selecting a good preconditioner for the Hessian of the log density. For this example, the metric $M = \Sigma^{-1}$ would result in a condition number of one.

This makes it possible to ask, for two metrics, M_1 and M_2 , which is better to use?

A practical issue is that the derivation of Eq. 4.9 is based on a linearization around a fixed point. During sampling, there is little chance that the sampler will land exactly on a fixed point. The assumption needed then is that local curvature in the typical set will extrapolate well around the posterior.

4.2.4 Approximate Hessian

We are assuming that the largest eigenvalues and eigenvectors of the Hessian of the log density change very little across the posterior. A low rank approximation to the Hessian can be built by picking out the largest K eigenvalue-eigenvectors pairs of the Hessian

$$\nabla_{qq}^2 H(q_0) \approx \sum_i^K v_i \lambda_i v_i^T. \quad (4.10)$$

Each of these eigenvalues correspond to the curvature of the posterior in the direction given by the eigenvector. Scaling each of the largest directions by the inverse of the corresponding eigenvalues will make it possible to pick larger and larger timesteps without violating the stability condition in Eq. 4.8.

The largest eigenvalue-eigenvector pairs of the Hessian can be computed via the Lanczos algorithm. The Lanczos algorithm converges most quickly for extreme eigenvalues and those with the most separation from their neighboring eigenvalues [79]. This means that the scales and directions in parameter space that are most limiting to the performance of HMC can be identified with relative ease. Conveniently, the Lanczos algorithm only requires Hessian vector products, which can be approximated with gradients from

reverse mode automatic differentiation and finite differences:

$$\nabla_{qq}^2 H(q) \cdot v \approx \frac{\nabla_q H(q + \frac{\Delta x}{2} v) - \nabla_q H(q - \frac{\Delta x}{2} v)}{\Delta x}. \quad (4.11)$$

A diagonal matrix can be used to extend the approximation in Eq. 4.10 to full rank so that it can be used directly as a preconditioner for $\nabla_{qq}^2 H(q)$. Take λ_{K+1} to be the $K + 1^{\text{th}}$ smallest eigenvalue, and call the approximation $A(\nabla_{qq}^2 H(q))$:

$$A(\nabla_{qq}^2 H(q)) = \sum_i^K v_i (\lambda_i - \lambda_{K+1}) v_i^T + \lambda_{K+1} I. \quad (4.12)$$

The inverse of this approximation scales the eigenvector directions by their respective eigenvalues and every other direction by the eigenvalue λ_{K+1} .

4.3 Implementation Details

4.3.1 Selection Criterion

An issue with Eq. 4.9 is that the local curvature of the Hessian, especially around the smaller eigenvalues like (λ_{\min}) , often does not accurately represent the posterior. Figure 4.1 shows a slice from a log density where local curvature is not representative of the whole posterior.

The trick is to replace λ_{\min} with the inverse of the largest eigenvalue of the covariance rescaled under the chosen metric. That is, use the selection criterion

$$\sqrt{|\lambda|_{\max}(L^T \nabla_{qq}^2 H(q) L) \lambda_{\max}(L^{-1} \Sigma L^{-T})} \quad (4.13)$$

instead of Eq. 4.9. For a normal posterior, Eq. 4.9 and Eq. 4.13 are equivalent and the smallest Hessian eigenvalue, λ_{\min} , corresponds to the largest covariance eigenvalue.




Figure 4.1: Slice of a log density plotted along the direction of an eigenvector of the Hessian, centered at a posterior draw (indicated with the dashed line). x is the distance from the posterior draw. The local curvature is positive and not representative of the scale of the log density in this direction.

For non-normal posteriors, the covariance can capture long scale behavior even when the local curvature does not.

The usefulness of Eq. 4.13 is predicated on the assumption that computing the largest eigenvalue of the rescaled covariance is easier than estimating the rescaled covariance ($L^{-1}\Sigma L^{-T}$) itself. The applicability of this type of assumption is explored in [80].

4.3.2 Low Rank Hessian Approximation

Because of differing parameter scales, computing a useful low rank approximation to the Hessian through an eigendecomposition is difficult even for simple posteriors. Instead of computing the approximation as in Eq. 4.12, a diagonal covariance estimate can be used to rescale the problem that makes it much easier to work with

$$D^{-\frac{1}{2}}A(D^{\frac{1}{2}}\nabla_{qq}^2H(q)D^{\frac{1}{2}})D^{-\frac{1}{2}}. \quad (4.14)$$

Figure 4.2 shows the eigenvalues of the Hessian, $\nabla_{qq}^2H(q)$, and the rescaled Hessian, $D^{\frac{1}{2}}\nabla_{qq}^2H(q)D^{\frac{1}{2}}$, to give a sense of what this rescaling does to make the approximation in Eq. 4.12 work better.




Figure 4.2: The eigenvalues of the Hessian of the negative log density evaluated at a point q , $\nabla_{qq}^2 H(q)$, are denoted by circles, and the eigenvalues of the rescaled version of that Hessian, $D^{\frac{1}{2}} \nabla_{qq}^2 H(q) D^{\frac{1}{2}}$, are denoted by plus signs.

4.3.3 Limiting to the Sample Covariance

The inverse of the sample covariance is a good metric as long as enough posterior draws are available. A metric based on local curvature information will require fewer draws to estimate but will be limited by how well that local curvature information extends to the full posterior.

The inverse of the approximation in Eq. 4.14 can be used as an initial estimate of the posterior covariance (Σ_0). Take as a prior

$$P(\Sigma) = \mathcal{W}^{-1}((\nu_0 - d - 1)\Sigma_0, \nu_0), \quad (4.15)$$

where \mathcal{W}^{-1} is an inverse-Wishart distribution and there are d parameters in the problem. Assuming the posterior draws come from a multivariate normal distribution, the posterior on Σ can be computed in closed form. Given n draws Y , we have

$$P(\Sigma|Y) = \mathcal{W}^{-1}((\nu_0 - d - 1)\Sigma_0 + (n - 1)S, \nu_0 + n), \quad (4.16)$$

where S is the sample covariance of Y . As the number of warmup draws increases, Σ

will be more and more weighted toward the sample covariance.

4.4 Benchmarks

A version of Stan [3] was modified to incorporate the selection criterion and adaptation for benchmarking. Unless otherwise noted, warmup is carried out over one thousand draws in the following sequence of steps. First, seventy five draws are used to move the sampler near to the typical set. Next, the metric is iteratively refined over five increasingly sized adaptation windows. The first window lasts for twenty-five draws, the second for fifty draws, the third for one hundred draws, the fourth for two hundred draws, and the last for five hundred draws. The final stage uses fifty draws to finalize the timestep selection.

At the end of each of the five metric adaptation windows, the draws from that window are divided into train (80%) and test (20%) splits. The train split is used to compute the different metrics and the test split is used to evaluate the selection criterion for each. The L variables in Eq. 4.13 come from the training split and W and Σ come from the test split. The selection criterion for each metric is computed at five different random draws from the test set and the maximum of these values is used as the final selection criterion. The chosen metric is recomputed using all of the draws in the given window.

4.4.1 Example Problems

The full the source and data for the example models are available in the Supplementary Material.

The *Kilpisjärvi* model is a three parameter linear regression. Yearly summer temperatures in Kilpisjärvi, Finland are fit as a function of year. The year covariate is not centered, and there is a very high correlation between the slope and intercept parameters

which makes the model difficult to sample [81].

The *Diamonds* model is a twenty-six parameter regression with highly correlated covariates based on the ggplot2 diamonds dataset [82] and the model generated with the software package brms [83]. The correlated covariates lead to posterior correlations. This can be avoided with a reparameterization of the problem [84], or with an appropriate adaptation.

Radon is a three hundred and eighty-nine parameter hierarchical model of radon levels in three hundred eighty-six different counties [81] adapted from [85]. This model is interesting because of the large number of parameters.

The *Accel GP* model is a sixty-six parameter fit to time series data from the mcycle dataset of [86]. The mean and standard deviations of the acceleration were modeled with Gaussian processes using a basis function expansion [87, 88] in brms [83].

The *Accel Splines* model is an eighty-two parameter spline model generated with brms [83] and fit to the same dataset as *Accel GP* with varying mean and standard deviation.

The *Prophet* [89] model used here is a sixty-two parameter time series model of RStan downloads over a few year period. Prophet implements structural time series model with different time resolutions (daily, weekly, etc.).

4.4.2 Results

The results of running thirty-two independent chains of four different adaptations on the example models are given in Table 4.1. The minimum and maximum of the selection criteria (lower is better) computed at the end of the last stage of warmup over all thirty-two chains is given for every model. The minimum and maximum of effective sample size per second (ESS/s) over eight different groups of four chains each was given as a

proxy for the performance. This assumes that the utility of the posterior is limited by the parameter with lowest effective sample size. The maximum four chain \hat{R} computed over all eight groups of chains of the parameter with the lowest effective sample size is reported as a diagnostic. To keep the \hat{R} values low, four thousand post-warmup draws were collected for the *Accel GP* and *Accel Splines* models. In all other cases one thousand post-warmup draws were collected. \hat{R} and the (bulk) effective sample sizes were computed following [90]. All benchmarks were run on an AMD Ryzen 7 2700X desktop.

The switching adaptation, described in the caption of Table 4.1, is competitive in all examples, with two exceptions. First, one of the eight four-chain *Accel GP* inferences done with switching adaptation performed unusually poorly in terms of ESS/s, and secondly another had an alarming number of divergences, one hundred and forty over thirty-two thousand post-warmup draws on the same model. For the *Accel GP* and *Accel Splines* models, all the other non-diagonal adaptation calculations had less than fifty divergences for the same number of post-warmup draws. The diagonal adaptations in comparison had at worst one (for *Accel GP*) and four (for *Accel Splines*) divergences.

Even without introducing the new Hessian-based adaptation scheme, the selection criteria can pick between the established methods (Diagonal and Dense). If enough draws are available such that the full sample covariance is an effective adaptation, the adaptation will pick Dense over Diagonal. Knowing if there were enough draws has been a simple but significant impediment in the deployment of Dense adaptation. All of the adaptations discussed assume posterior curvature is not varying greatly.

The *Radon* model is large enough that the dense matrix-vector products affect sampler performance. Table 4.2 shows the results of enforcing a diagonal sparsity pattern on the metrics, enabling efficient matrix-vector products.

One of the advantages of the low rank adaptation is that it can work with fewer warmup draws than a full sample covariance. To highlight this, the results of running a

MODEL	ADAPTATION	CRIT. [MIN, MAX]	MIN ESS/S [MIN, MAX]	MAX \hat{R}
KILPISJÄRVI	DIAGONAL	[350, 600]	[110, 150]	1.01
	DENSE	[95, 130]	[390, 6100]	1.00
	RANK-1	[1.3, 1.9]	[22000, 33000]	1.00
	SWITCHING	[1.2, 1.7]	[23000, 33000]	1.00
DIAMONDS	DIAGONAL	[490, 670]	[0.61, 0.90]	1.01
	DENSE	[4.7, 5.7]	[160, 190]	1.01
	RANK-1	[2.0, 2.4]	[510, 620]	1.00
	SWITCHING	[1.9, 2.4]	[510, 590]	1.00
ACCEL GP	DIAGONAL	[210, 570]	[2.5, 5.2]	1.00
	DENSE	[41, 240]	[1.1, 6.4]	1.01
	RANK-1	[46, 240]	[3.4, 8.1]	1.00
	SWITCHING	[45, 460]	[3.7, 9.8]	1.00
ACCEL SPLINES	DIAGONAL	[680, 1800]	[0.17, 0.45]	1.02
	DENSE	[90, 420]	[2.3, 5.3]	1.00
	RANK-1	[83, 330]	[1.5, 4.7]	1.00
	SWITCHING	[92, 320]	[0.38, 5.0]	1.02
PROPHET	DIAGONAL	[390, 470]	[1.3, 1.7]	1.00
	DENSE	[13, 15]	[22, 28]	1.00
	RANK-1	[2.8, 3.3]	[24, 29]	1.00
	SWITCHING	[2.8, 3.3]	[15, 28]	1.00
RADON	DIAGONAL	[5.3, 7.2]	[170, 210]	1.00
	DENSE	[67, 93]	[33, 42]	1.00
	RANK-1	[7.4, 9.2]	[120, 140]	1.00
	SWITCHING	[4.3, 6.0]	[180, 210]	1.00

Table 4.1: Four different metrics were benchmarked on six different models. Diagonal and Dense are Stan default metrics. Dense is a full covariance estimate and Diagonal is a diagonal covariance approximation that is much easier to compute. Both are regularized a small amount towards an identity matrix. Rank-1 is an metric from Section 4.3.3 using a rank-1 Hessian approximation, and Switching is a metric that switches between dense, diagonal, rank-1, 2, 4, and 8 Hessian approximation adaptations with and without the modifications described in Section 4.3.3 by choosing the metric with the lowest maximum selection criterion at the end of each warmup window. Only the post-warmup draws are timed. The range of maximum selection criteria (lower is better) from all thirty-two chains are given in the “Crit. [min, max]” column. The ESS/s column characterizes efficiency as effective sample size per second computed over eight groups of four chains each. The maximum four chain \hat{R} computed over eight different groups of four chains is given as a diagnostic.

MODEL	ADAPTATION	CRIT. [MIN, MAX]	MIN N_{EFF}/s	MAX \hat{R}
RADON	DIAGONAL	[5.1, 6.1]	[390, 520]	1.00
	SWITCHING	[4.6, 5.8]	[370, 490]	1.00

Table 4.2: The *Radon* experiments from Table 4.1 are repeated here if the metric is forced to be a diagonal (to take advantage of the efficient matrix-vector multiplies this allows). The Switching adaptation switches between diagonal, and diagonalized rank-1, 2, 4, and 8 Hessian adaptations without the modifications in Section 4.3.3. Only the post-warmup draws are timed.

number of extremely short adaptations on the *Diamonds* and *Prophet* example models are displayed in Figure 4.3. For these experiments, adaptation was limited to the initial seventy-five draws used to get near the typical set, one metric adaptation window, and a final fifty draws to adjust the timestep.

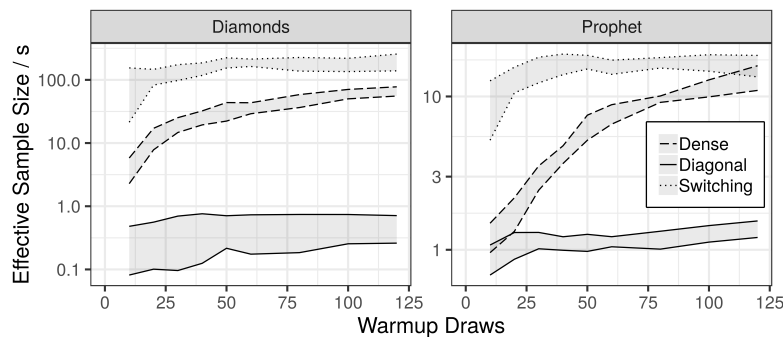


Figure 4.3: Effectiveness of adaptations with short warmup. The Switching adaptation chooses automatically between dense, diagonal, and rank 1, 2, 4, and 8 adaptations from Eq. 4.12 with and without the inverse-Wishart update from Section 4.3.3. Only the post-warmup draws are timed. The maximum and minimum four chain ESS/s across eight groups are plotted as ranges. For more warmup samples, the advantages of switching adaptation go away (as can be seen in Table 4.1).

The short warmup experiments are also useful for understanding the effect of the rank approximation from Eq. 4.12 as well as the inverse-Wishart update in Section 4.3.3. For the most part, these are useful modifications, but neither guarantees a performance increase. Results are plotted in 4.4.

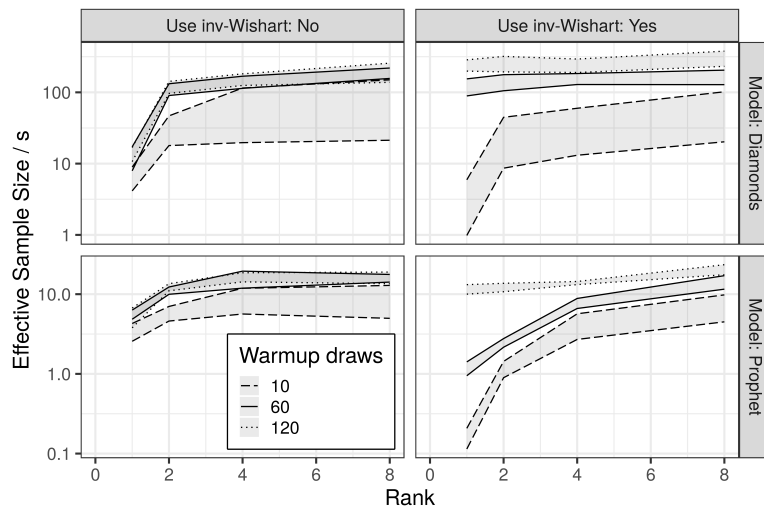


Figure 4.4: Effectiveness of adaptations with short warmup. The rank 1, 2, 4, and 8 adaptations are from Eq. 4.12. The top row includes the inverse-Wishart update from Section 4.3.3 and the bottom does not. Only the post-warmup draws are timed. The maximum and minimum four chain ESS/s across eight groups are plotted as ranges.

4.5 Conclusions

Adapting an effective metric is important for the performance of HMC. This paper outlines a criterion that can be used to automate the selection of an efficient metric from an array of options. In addition, we present a new low-rank adaptation scheme that makes it possible to sample effectively from highly correlated posteriors, even when few warmup draws are available. The selection criterion and the new adaptation are demonstrated to be effective on a number of different models.

All of the necessary eigenvalues and eigenvectors needed to evaluate the selection criterion and build the new adaptation can be computed efficiently with the Lanczos algorithm, making this method suitable for models with large numbers of parameters.

Chapter 5

Conclusions

We have described several novel applications of data science covering micrograph analysis, mechanistic modeling, and Bayesian inference. We described the context in which this research arose, and how the results are applicable in materials science.

First we described the advantages of the histogram of oriented gradients feature descriptor, and how this circumvents the problem in micrograph analysis of producing reliable image segmentations. The effectiveness of this technique was demonstrated on 2D and 3D datasets.

Secondly, we demonstrated the use of Hamiltonian Monte Carlo for characterizing elastic properties and crystal misorientation of superalloy parallelepiped specimens. This algorithm was demonstrated first by estimating the elastic constants for a specimen of fine-grain polycrystalline Ti-6Al-4V (Ti-64) with random crystallographic texture and isotropic elastic symmetry. Secondly, it was tested by estimating the elastic constants and crystallographic orientation for a single crystal Ni-based superalloy CMSX-4 specimen are accurately determined. In both cases, the only data used were measurements of the specimen geometry, mass, and resonance frequencies, substantially simplifying experimental procedures.

Finally, we demonstrated how consideration of the stability limit of the leapfrog integrator used inside HMC leads to the development of a selection criterion for picking the HMC metric, allowing it to run efficiently. Additionally, we presented a new low-rank adaptation scheme inspired by this observation and demonstrated the effectiveness of the selection criterion and adaptation scheme on a number of different applied models. It was shown that all of the necessary calculations can be carried out efficiently through use of the Lanczos algorithm, making this method suitable for models with large numbers of parameters.

Bibliography

- [1] S. Bernard, G. Marrelec, P. Laugier, and Q. Grimal, *Bayesian normal modes identification and estimation of elastic coefficients in resonant ultrasound spectroscopy*, *Inverse Problems* **31** (2015), no. 6 065010.
- [2] B. R. Goodlet, L. Mills, B. Bales, M.-A. Charpagne, S. P. Murray, W. C. Lenthe, L. Petzold, and T. M. Pollock, *Elastic properties of novel co- and con- based superalloys determined through bayesian inference and resonant ultrasound spectroscopy*, *Metallurgical and Materials Transactions A* **49** (June, 2018) 2324–2339.
- [3] B. Carpenter, A. Gelman, M. D. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, and A. Riddell, *Stan: A probabilistic programming language*, *Journal of statistical software* **76** (2017), no. 1.
- [4] M. P. Echlin, W. C. Lenthe, and T. M. Pollock, *Three-dimensional sampling of material structure for property modeling and design*, *Integrating Materials and Manufacturing Innovation* **3** (2014), no. 1 1–14.
- [5] P. G. Callahan, J. P. Simmons, and M. De Graef, *A quantitative description of the morphological aspects of materials structures suitable for quantitative comparisons of 3D microstructures*, *Modelling Simul. Mater. Sci. Eng.* **21** (January, 2013) 015003.
- [6] P. G. Callahan and M. D. Graef, *Precipitate shape fitting and reconstruction by means of 3D Zernike functions*, *Modelling and Simulation in Materials Science and Engineering* **20** (2012), no. 1 015003.
- [7] J. V. Sluytman and T. Pollock, *Optimal precipitate shapes in nickel-base alloys*, *Acta Materialia* **60** (2012), no. 4 1771 – 1783.
- [8] E. E. Underwood, *Microstructural Analysis: Tools and Techniques*, ch. Quantitative Stereology for Microstructural Analysis, pp. 35–66. Springer US, Boston, MA, 1973.
- [9] S. Tin, T. Pollock, and W. King, *Carbon additions and grain defect formation in high refractory nickel-base single crystal superalloys*, *Ann Arbor* **1001** (2000) 48109.

- [10] M. Comer, C. A. Bouman, M. Graef, and J. P. Simmons, *Bayesian methods for image segmentation*, *JOM* **63** (2011), no. 7 55–57.
- [11] J. Simmons, B. Bartha, M. De Graef, and M. Comer, *Development of Bayesian segmentation techniques for automated segmentation of titanium alloy images*, *Microscopy and Microanalysis* **14** (8, 2008) 602–603.
- [12] E. Gulsoy and M. DeGraef, *Advances in the automated segmentation of 3-D microstructures*, *Microscopy and Microanalysis* **15** (7, 2009) 606–607.
- [13] L. Huffman, J. Simmons, M. D. Graef, and I. Pollak, *Shape priors for map segmentation of alloy micrographs using graph cuts*, in *Statistical Signal Processing Workshop (SSP), 2011 IEEE*, pp. 661–664, June, 2011.
- [14] B. Chaudhuri, P. Kundu, and N. Sarkar, *Detection and gradation of oriented texture*, *Pattern Recognition Letters* **14** (1993), no. 2 147 – 153.
- [15] N. Dalal and B. Triggs, *Histograms of oriented gradients for human detection*, in *Computer Vision and Pattern Recognition, 2005*, vol. 1, pp. 886–893, June, 2005.
- [16] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan, *Object detection with discriminatively trained part-based models*, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **32** (Sept, 2010) 1627–1645.
- [17] M. Fährmann, P. Fratzl, O. Paris, E. Fährmann, and W. C. Johnson, *Influence of coherency stress on microstructural evolution in model Ni-Al-Mo alloys*, *Acta Metallurgica et Materialia* **43** (1995), no. 3 1007 – 1022.
- [18] T. Gabb, P. Kantzos, J. Telesman, J. Gayda, C. Sudbrack, and B. Palsa, *Fatigue resistance of the grain size transition zone in a dual microstructure superalloy disk*, *International Journal of Fatigue* **33** (2011), no. 3 414–426.
- [19] M. B. Stegmann and D. D. Gomez, *A brief introduction to statistical shape analysis*, March, 2002.
- [20] S. Belongie, J. Malik, and J. Puzicha, *Shape matching and object recognition using shape contexts*, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **24** (2002), no. 4 509–522.
- [21] D. G. Lowe, *Object recognition from local scale-invariant features*, in *Proceedings of the Seventh IEEE International Conference on Computer Vision*, vol. 2, pp. 1150–1157 vol.2, Sep., 1999.
- [22] A. Krizhevsky, I. Sutskever, and G. E. Hinton, *Imagenet classification with deep convolutional neural networks*, in *Advances in Neural Information Processing Systems*, pp. 1097–1105, 2012.

- [23] E. Faber, D. Martinez-Martinez, C. Mansilla, V. Ocelk, and J. D. Hosson, *Calibration-free quantitative surface topography reconstruction in scanning electron microscopy*, *Ultramicroscopy* **148** (2015) 31 – 41.
- [24] C. Mansilla, M. H. Novais, E. Faber, D. Martínez-Martínez, and J. T. De Hosson, *On the 3d reconstruction of diatom frustules: a novel method, applications, and limitations*, *Journal of Applied Phycology* **28** (2016), no. 2 1097–1110.
- [25] D. T. Fullwood, S. R. Niezgoda, and S. R. Kalidindi, *Microstructure reconstructions from 2-point statistics using phase-recovery algorithms*, *Acta Materialia* **56** (2008), no. 5 942 – 948.
- [26] S. R. Niezgoda, D. M. Turner, D. T. Fullwood, and S. R. Kalidindi, *Optimized structure based representative volume element sets reflecting the ensemble-averaged 2-point statistics*, *Acta Materialia* **58** (2010), no. 13 4432 – 4445.
- [27] B. L. DeCost and E. A. Holm, *A computer vision approach for automated analysis and classification of microstructural image data*, *Computational Materials Science* **110** (2015) 126 – 133.
- [28] A. Klaser, M. Marszalek, and C. Schmid, *A Spatio-Temporal Descriptor Based on 3D-Gradients*, in *BMVC 2008 - 19th British Machine Vision Conference* (M. Everingham, C. Needham, and R. Fraile, eds.), (Leeds, United Kingdom), pp. 275:1–10, British Machine Vision Association, September, 2008.
- [29] M. P. Echlin, A. Mottura, M. Wang, P. J. Mignone, D. P. Riley, G. V. Franks, and T. M. Pollock, *Three-dimensional characterization of the permeability of W-Cu composites using a new “TriBeam” technique*, *Acta Materialia* **64** (2014) 307 – 315.
- [30] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz, *Rotation invariant spherical harmonic representation of 3d shape descriptors*, in *Proceedings of the 2003 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*, SGP '03, (Aire-la-Ville, Switzerland, Switzerland), pp. 156–164, Eurographics Association, 2003.
- [31] N. Zhou, C. Shen, M. Mills, and Y. Wang, *Large-scale three-dimensional phase field simulation of γ' -rafting and creep deformation*, *Philosophical Magazine* **90** (2010), no. 1-4 405–436, [<http://dx.doi.org/10.1080/14786430903081990>].
- [32] A. D. Kammers and S. Daly, *Digital image correlation under scanning electron microscopy: Methodology and validation*, *Experimental Mechanics* **53** (2013), no. 9 1743–1761.
- [33] C. Mansilla, V. Ocelík, and J. T. M. D. Hosson, *A new methodology to analyze instabilities in sem imaging*, *Microscopy and Microanalysis* **20** (12, 2014) 1625–1637.

- [34] R. G. Leisure and F. A. Willis, *Resonant ultrasound spectroscopy*, *J. Phys. Condens. Matter* **9** (July, 1997) 6001–6029.
- [35] A. Migliori and J. Sarrao, *Resonant ultrasound spectroscopy: applications to physics, materials measurements, and nondestructive evaluation*. Wiley-Interscience, New York, 1997.
- [36] M. Radovic, E. Lara-Curzio, and L. Riester, *Comparison of different experimental techniques for determination of elastic properties of solids*, *Materials Science and Engineering: A* **368** (2004), no. 1 56 – 70.
- [37] J. D. Maynard, *The use of piezoelectric film and ultrasound resonance to determine the complete elastic tensor in one measurement*, *Journal of the Acoustical Society of America* **91** (March, 1992) 1754–1762.
- [38] A. 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** (January, 1993) 1–24.
- [39] A. 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** (2005), no. 2005 1–7.
- [40] J. 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** (1994), no. 6 2139–2140.
- [41] R. Holland, *Resonant properties of piezoelectric ceramic rectangular parallelepipeds*, *J. Acoust. Soc. Am.* **43** (1968), no. 5 988.
- [42] H. H. Demarest, *Cube-resonance method to determine the elastic constants of solids*, *J. Acoust. Soc. Am.* **49** (1971), no. 3B 768.
- [43] I. Ohno, *Free vibration of a rectangular parallelepiped crystal and its application to determination of elastic constants of orthorhombic crystals*, *J. Phys. Earth* **24** (1976) 255–379.
- [44] W. 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** (October, 1991) 2154–2162.
- [45] J. 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** (2004), no. 1 282–287.

- [46] B. J. Zadler, J. H. L. Le Rousseau, J. A. Scales, and M. L. Smith, *Resonant ultrasound spectroscopy: theory and application*, *Geophys. J. Int.* **156** (January, 2004) 154–169.
- [47] A. Tarantola, *Inverse Problem Theory and Methods for Model Parameter Estimation*. Society for Industrial and Applied Mathematics, 2005.
- [48] M. 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** (July, 2015) 4898–4916.
- [49] P. Sedlák, 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** (2014), no. 6 1073–1085.
- [50] J. 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** (2007), no. 2 216–227.
- [51] Y. 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** (2006), no. 3 341–357.
- [52] H. Ogi, K. Sato, T. Asada, and M. Hirao, *Complete mode identification for resonance ultrasound spectroscopy*, *Journal of the Acoustical Society of America* **112** (2002), no. 6 2553.
- [53] M. 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* (M. Everett and L. Pedroza, eds.), ch. 3, pp. 97–136. Nova Science Publishers, Hauppauge, NY, 2009.
- [54] *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 (ASTM), West Conshohocken, PA, 2015.
- [55] A. F. Bower, *Applied Mechanics of Solids*. CRC press, Boca Raton, FL, 2009.
- [56] R. M. Neal, *MCMC using Hamiltonian dynamics*, *Handbook of Markov Chain Monte Carlo* **2** (2011), no. 11 2.
- [57] S. Byrne and M. Girolami, *Geodesic Monte Carlo on embedded manifolds*, *Scandinavian Journal of Statistics* **40** (2013), no. 4 825–845.

- [58] Stan Development Team, *Stan Modeling Language Users Guide and Reference Manual, Version 2.14.0*, 2016.
- [59] J. de Leeuw, *Derivatives of generalized eigen systems with applications*, Department of Statistics UCLA (2007).
- [60] C. Zener, *Contributions to the Theory of Beta-Phase Alloys*, *Physical Review* **71** (June, 1947) 846–851.
- [61] A. A. Pourzanjani, R. M. Jiang, B. Mitchell, P. J. Atzberger, and L. R. Petzold, *General Bayesian Inference over the Stiefel Manifold via the Givens Representation*, *arXiv e-prints* (Oct, 2017) arXiv:1710.09443, [arXiv:1710.0944].
- [62] S. V. Zherebtsov, G. A. Salishchev, R. M. Galejev, O. R. Valiakhmetov, S. Yu. Mironov, and S. L. Semiatin, *Production of submicrocrystalline structure in large-scale Ti-6Al-4V billet by warm severe deformation processing*, *Scripta Materialia* **51** (2004), no. 12 1147–1151.
- [63] E. S. Fisher and C. J. Renken, *Single-Crystal Elastic Moduli and the hcp bcc Transformation in Ti, Zr, and Hf*, *Physical Review* **135** (July, 1964) A482–A494.
- [64] R. Hill, *The Elastic Behaviour of a Crystalline Aggregate*, *Proc. Phys. Soc. Sect. A* **65** (1952), no. 5 349–354.
- [65] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian Data Analysis, 3rd Edition*. CRC Press, Boca Raton, FL, 2014.
- [66] D. Siebörger, H. Knake, and U. Glatzel, *Temperature dependence of the elastic moduli of the nickel-base superalloy CMSX-4 and its isolated phases*, *Material Science and Engineering: A* **298** (January, 2001) 26–33.
- [67] D. 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*, *Modeling and Simulation in Materials Science and Engineering* **22** (2014), no. 7 075013.
- [68] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth, *Hybrid Monte Carlo*, *Physics letters B* **195** (1987), no. 2 216–222.
- [69] M. Betancourt, *A conceptual introduction to Hamiltonian Monte Carlo*, *arXiv preprint arXiv:1701.02434* (2017).
- [70] J. Salvatier, T. V. Wiecki, and C. Fonnesbeck, *Probabilistic programming in Python using PyMC3*, *PeerJ Computer Science* **2** (2016) e55.

- [71] M. Girolami and B. Calderhead, *Riemann manifold Langevin and Hamiltonian Monte Carlo methods*, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73** (2011), no. 2 123–214, [<https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-9868.2010.00765.x>].
- [72] M. Betancourt, *A general metric for Riemannian manifold Hamiltonian Monte Carlo*, in *Geometric Science of Information*, pp. 327–334. Springer, 2013.
- [73] B. Leimkuhler and S. Reich, *Simulating Hamiltonian Dynamics*, vol. 14. Cambridge University Press, 2004.
- [74] E. Hairer, C. Lubich, and G. Wanner, *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*, vol. 31. Springer Science & Business Media, 2006.
- [75] M. D. Hoffman and A. Gelman, *The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo*, *Journal of Machine Learning Research* **15** (2014), no. 1 1593–1623.
- [76] L. Perko, *Differential equations and dynamical systems*, vol. 7. Springer Science & Business Media, 2013.
- [77] A. Gelman, H. S. Stern, J. B. Carlin, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian Data Analysis*. Chapman and Hall/CRC, 2013.
- [78] J. José and E. Saletan, *Classical Dynamics: A Contemporary Approach*. AAPT, 2000.
- [79] C. D. Meyer, ed., *Matrix Analysis and Applied Linear Algebra*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.
- [80] A. Loukas, *How close are the eigenvectors of the sample and actual covariance matrices?*, in *Proceedings of the 34th International Conference on Machine Learning*, vol. 70, pp. 2228–2237, August, 2017.
- [81] A. Gelman and J. Hill, *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Analytical Methods for Social Research. Cambridge University Press, 2006.
- [82] H. Wickham, *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York, 2016.
- [83] P.-C. Brkner, *brms: An R package for Bayesian multilevel models using Stan*, *Journal of Statistical Software* **80** (2017), no. 1 1–28.
- [84] Stan Development Team, *Stan Users Guide Version 2.18*. 2018.

- [85] C. Fonnesbeck, “A Primer on Bayesian Multilevel Modeling using PyStan.” <https://mc-stan.org/users/documentation/case-studies/radon.html>, 2016. Accessed: 2019-03-30.
- [86] W. Venables and B. Ripley, *Modern Applied Statistics with S-PLUS*. Springer, 3 ed., 1999.
- [87] G. Riutort Mayol, M. R. Andersen, P. Bürkner, and A. Vehtari, *Hilbert space methods to approximate Gaussian processes using Stan, In preparation*. (2019).
- [88] A. Solin and S. Särkkä, *Hilbert space methods for reduced-rank Gaussian process regression, arXiv preprint arXiv:1401.5508* (2014).
- [89] S. J. Taylor and B. Letham, *Forecasting at scale, The American Statistician* **72** (2018), no. 1 37–45.
- [90] A. Vehtari, A. Gelman, D. Simpson, B. Carpenter, and P.-C. Bürkner, *Rank-normalization, folding, and localization: An improved \hat{R} for assessing convergence of MCMC, arXiv e-prints* (March, 2019) arXiv:1903.08008, [arXiv:1903.0800].