

ENHANCING PERFORMANCE PREDICTION ACCURACY OF HIGH STRENGTH
ALLOYS VIA UNCERTAINTY QUANTIFICATION

A Dissertation

by

JAYLEN R. JAMES

Submitted to the Graduate and Professional School of
Texas A&M University
in partial fulfillment of the requirements for the degree of
DOCTOR OF PHILOSOPHY

Chair of Committee,	Raymundo Arróyave
Co-Chair of Committee,	Douglas Allaire
Committee Members	Ankit Srivastava
	Ibrahim Karaman
Head of Department,	Ibrahim Karaman

May 2022

Major Subject: Materials Science and Engineering

Copyright 2022 Jaylen R. James

ABSTRACT

Uncertainty Quantification (UQ) and its subsequent propagation are powerful tools for estimating material property and performance distributions. As the paradigm of materials discovery within an Integrated Computational Materials Engineering framework continues to mature, so does the need for UQ to validate behavior predictions of new and existing materials. In this work, UQ is combined with physics informed models, high fidelity simulations, and statistical optimization techniques to characterize the high-strain-rate response of AF9628 and the dwell fatigue life of Ti-6Al-4V. Three specific instances involving the characterization of these materials using uncertainty quantification, and propagation, are presented.

First, a technique of information fusion called Reification, is used to combine constitutive models and experimental data to obtain a high-strain-rate performance boundary of AF9628. A range of model parameter values is also determined and is then propagated through a high fidelity simulation software to further estimate material property boundaries.

In the second study, the distribution of material density obtained via the high fidelity simulation is re-weighted, using a newly developed technique, Probability Law Optimized Weights (PLOW). The purpose of PLOW is to adjust a proposal distribution generated from a sub-set of inputs such that it more accurately reflects the target distribution generated from the full set of inputs. The method is particularly useful when computational limits prevent evaluation of the entire input set.

In the third study, 2-dimensional measurements of Microtextured Regions (MTRs) in Ti-6Al-4V are used to infer the size of the 3-dimensional MTRs from which they are sectioned. This distribution of estimated sizes is then propagated through a dwell fatigue model to study the influence of measurement type on the expected fatigue life for the near alpha titanium alloy.

DEDICATION

To my mom, Belinda, dad, Timothy, two brothers, Timothy II and Thaddeus, and my grandparents on both sides. Also to my extended family and closest friend Jesus.

ACKNOWLEDGMENTS

First I would like to thank Dr. Raymundo Arróyave and Dr. Douglas Allaire. They have invested so many hours into my development as a research scientist and as a person. I am very grateful for this.

Special thanks to Dr. Eric Payton, Dr. Manny Gonzales, Dr. Reji John, Austin Gerlt, and so many others at the Air Force Research Lab. They have helped me in so many ways that I will never forget.

Other specific acknowledgments include Dr. Ankit Srivastava who has taught me so much about fundamental and advanced materials concepts, Dr. Ibrahim Karaman who has helped provide funding for my research, and the entire Materials Science and Engineering Department at Texas A&M University.

CONTRIBUTORS AND FUNDING SOURCES

Contributors

This work was supported by a dissertation committee consisting of advisor Dr. Raymundo Arróyave, co-advisor Dr. Douglas Allaire, Dr. Ankit Srivastava of the Materials Science & Engineering Department, and Dr. Ibrahim Karaman, head of the Materials Science & Engineering Department.

The data analyzed for Chapter 1 was provided by Dr. Manny Gonzales. The impact simulations depicted in Chapters 1 and 2 were conducted by Austin Gerlt. Both Dr. Manny Gonzales and Austin Gerlt are from the Air Force Research Laboratory in Dayton, Ohio. The reweighting calculations for Chapter 2 were done in part by Meet Sanghvi of the Mechanical Engineering Department at Texas A&M University.

All other work conducted for the dissertation was completed by the student independently.

Funding Sources

Graduate study was supported by the Graduate Diversity and Pathways fellowships from Texas A&M University and the Minority Leadership Program at the Air Force Research Laboratory.

TABLE OF CONTENTS

	Page
ABSTRACT	ii
DEDICATION	iii
ACKNOWLEDGMENTS	iv
CONTRIBUTORS AND FUNDING SOURCES	v
TABLE OF CONTENTS	vi
LIST OF FIGURES	viii
LIST OF TABLES.....	xi
1. INTRODUCTION	1
2. PARAMETER RANGE ESTIMATION AND PROPERTY UNCERTAINTY ANALYSIS FOR HIGH STRAIN RATE MATERIAL RESPONSE VIA MODEL FUSION	5
2.1 Abstract.....	5
2.2 Introduction.....	5
2.3 Methods.....	8
2.3.1 Experimental Procedure.....	9
2.3.2 Determining Johnson-Cook Parameters	10
2.3.3 Determining Zerilli-Armstrong Parameters	12
2.3.4 Reification Technique	13
2.3.5 Determining the Reified Model's Parameters	15
2.3.6 Propagation to Taylor Anvil Simulations.....	16
2.3.6.1 Performing MCMC and Medoids	16
2.3.6.2 Performing Taylor Anvil Simulations	17
2.4 Results	18
2.4.1 Experimental Data and Model Fits	18
2.4.2 Model Fusion Parameter Results	19
2.4.3 Taylor Anvil Simulation Results	23
2.5 Discussion and Future Work.....	30
2.6 Conclusions.....	31
3. OPTIMIZED UNCERTAINTY PROPAGATION ACROSS HIGH FIDELITY TAYLOR ANVIL SIMULATION	33

3.1	Abstract	33
3.2	Introduction.....	33
3.2.1	The cost of sampling high fidelity simulations.....	35
3.3	Methods.....	37
3.3.1	Experimental setup	37
3.3.2	Obtaining Target Data Set Values.....	38
3.3.3	Hydrocode simulations to obtain target distribution of material density	38
3.3.4	Probability Law Optimized Importance Weights	41
3.3.4.1	Background	41
3.3.4.2	Approach	42
3.3.5	Description of the 'Toy Problem'	45
3.4	Results	49
3.4.1	Target material density distribution generated from Taylor-Anvil simulations	49
3.4.2	Toy Problem Results	52
3.5	Discussion	54
3.6	Conclusions.....	57
4.	INCREASED FATIGUE LIFE PREDICTION ACCURACY OF Ti-6Al-4V USING UN-FOLDING	59
4.1	Abstract	59
4.2	Introduction.....	59
4.3	Estimation Types	63
4.3.1	Unfolding	63
4.3.2	Expectation Maximization	66
4.4	Methods.....	67
4.4.1	Obtaining MTR Measurements	67
4.4.2	Coding and Use of Unfolding.....	68
4.4.3	Coding and Use of EM	69
4.4.4	Sampling from the Distributions.....	71
4.4.5	Creating A vs N plots - Crack Growth Model	74
4.5	Results and Discussion.....	75
4.5.1	Bi-Variate Histograms of MTR Data	75
4.5.2	CDFs of Lengths	76
4.5.3	Crack Length vs Number of Cycles distributions	79
4.6	Conclusion.....	81
4.7	Acknowledgements	84
5.	SUMMARY AND CONCLUSIONS	85
5.1	Further Study	85
	REFERENCES	86

LIST OF FIGURES

FIGURE	Page
2.1 Framework used to determine parameter value ranges and simulation output distributions.	7
2.2 Depiction of the 2-dimensional Taylor anvil simulation. The specimen (Rod) dimensions are noted as well as the location at which the response data is recorded for the present study.	18
2.3 Monotonic tension response for AF9628 tested at room temperature and strain rate of $10^{-2} s^{-1}$	19
2.4 Dynamic stress-strain tension response of AF9628. The experimental data points from each of the five experiments are shown in green. The filled purple circles are the average stresses of the 5 experimental data points.	20
2.5 (a) Reification mean and 95% confidence interval (CI) plotted with the average experimental data, and High-Strain-Rate Models fit to the data. (b) Reification mean with the upper and lower JC models representing the response of the reification mean. The parameters of the upper and lower models are the mode of the MCMC chain values excluding the burn-in region. (c) Same as (b) except the upper and lower ZA models are displayed. The parameters of the upper and lower models are the mode and mean MCMC chain values with the burn-in region removed, respectively. For (b) and (c) the chain statistic (mean or mode) used was that which minimized the root mean squared error between the reified mean and the model. The exact parameter values for each model are recorded in Table 2.3.	21
2.6 Last frame of Taylor anvil impact simulations. Each image shows the material's response from the perspective of the (a) von Mises stress (b) triaxiality (c) pressure and (d) temperature. All screenshots are shown at the same impact time of 20 microseconds.	23
2.7 Cumulative distributions of outputs from Taylor anvil simulations. In each plot the distributions produced using the Johnson-Cook Lower (JC-L), Johnson-Cook Upper (JC-U), Johnson-Cook alone (JC), and Zerilli-Armstrong alone (ZA) models are shown.	27

2.8	Cumulative distributions of re-weighted outputs from Taylor anvil simulations with the cumulative distributions of the original outputs. In each plot the distributions produced using the Johnson-Cook Lower (JC-L), Johnson-Cook Upper (JC-U), and Johnson-Cook alone (JC) are shown. Also displayed are the re-weighted CDFs (JC-L Re, JC-U Re, JC Re) which are based on the previously mentioned models. . . .	29
3.1	Process used to analyze the effect of the PLOW method on the high fidelity Taylor-Anvil simulation and at scale using the Zerilli-Armstrong model.	36
3.2	Illustration of the process to construct a measure matrix, like A in Equation 3.2. On the left, a bivariate distribution space is displayed corresponding to a model with two parameters for simplicity. In the case of inputs, the dots on the surface represent the co-variance of points the investigator is able to run. The borders show each parameters individual distribution. On the right, circles S_j (hyper-spheres in higher dimensions) are drawn over the space to populate the measure matrix.	43
3.3	Illustration of the Kolmogorov–Smirnov statistic. The red and blue lines are the two CDFs being compared, while the black arrow is a measurement of the largest distance between the two distributions. It is commonly known as the K-S statistic, but is also referred to as the D-value (distance value) in the present work.	48
3.4	Last frame of Taylor-Anvil simulation movie created using hydro-code software. The pressure throughout the specimen and anvil are represented by color. The simulation was informed by one ZA parameter set of the available 499.	49
3.5	Cumulative distribution of specimen density values predicted by the Taylor Anvil simulation. Each line, except one, represents a proposal distribution determined by an input sub-sample size as indicated in the legend. One line represents the target distribution. Naturally, as the size of sub-sampled input points used to construct the proposal distribution approaches 500, the closer the proposal distribution will be to the target distribution.	50
3.6	D-value versus sub-sample size trend for TA simulations. The sub-sample size represents the number of input points used to construct the proposal distribution. At each X-value, 100 2-SampleK-S statistic values (D-values) were calculated between the target distribution and a proposal distribution constricted using a different drawing of points from the pool of 499 total input points. The average and standard deviation of the 100 D-values is plotted.	52
3.7	D-value versus sub-sample size trend for toy problem. A comparison between a random sub-sample selection with PLOW applied after the selection (Weighted) and one without PLOW applied (Unweighted) is shown. Like other D-value versus proposal sample size plots in this work, the average and standard deviation of 100 D-values is plotted at each sub-sample size value.	53

3.8	D-value versus sub-sample size trends for the toy problem. A comparison of the convergence to the target distribution between random sub-sample selection with PLOW applied (Weighted-Random), random sub-sample selection without PLOW (Unweighted-Random), k-medoids sub-sample selection with PLOW applied (Weighted - Kmedoids), and k-medoids sub-sample selection without PLOW (Unweighted - Kmedoids) are shown.....	54
4.1	Method to compare the approaches.....	60
4.2	Size of MTRs with relation to region in plot.....	62
4.3	Histograms of the 2D, Unfolded, and EM-Chan probabilities of the synthetic data. The grid size of the histograms are 10 by 10.....	72
4.4	Illustration showing how crack growth model works in 2 Dimensions.	73
4.5	a) Bi-variate histogram of the experimentally measured MTR major principal semi-axis. b) Unfolded histogram of experimentally measured MTR data with Expectation Maximization applied. c) Same as b) but Expectation Maximization is not applied. The histograms are of size 20 by 20.	77
4.6	Values of 3D and Chan-EM are from sample of 400K points. Chan-EM probabilities are generated from 400K experimental data points uniformly sampled without replacement. Displayed in each sub-figure are (a) Major Principal Semi-Axes or Mapsas, (b) Minor Principal Semi-Axes or Mipsas, and (c) Eccentricities.....	78
4.7	Cumulative distribution of cycle at fracture predictions for 1 million iterations of curves per data source.	81

LIST OF TABLES

TABLE	Page
2.1 Tension test conditions.	9
2.2 Monotonic stress test results	19
2.3 Parameter values of the models used in the study.....	22
2.4 Statistical significance comparison of Taylor anvil simulation outputs between model pairs. Values in the cells are the D -values calculated between the two output CDFs. Shaded cells represent statistical similarity between the two models for the specified output (i.e. A D -value so small that H_0 , the null hypothesis that two distributions are from the same sample, can not be rejected in a Komolgorov-Smirnov 2-Sample test). The CDFs of each of the outputs are shown in Figure 2.7.....	24
2.5 Statistical significance comparison of re-weighted Taylor anvil simulation outputs between model pairs. Values in the cells are the D -values calculated between the two output CDFs. Shaded cells represent statistical similarity between the two models for the specified output (i.e.A D -value so small that H_0 , the null hypothesis that two distributions are from the same sample, can not be rejected in a Komolgorov-Smirnov 2-Sample test). The CDFs of select outputs are shown in Figure 2.8.	25
3.1 Zerilli-Armstrong constants obtained using gradient based minimization against experimental data.....	39
4.1 Important number of cycle statistics.....	80

1. INTRODUCTION

The act of putting a number to how confident one is in a measurement is a broad definition of quantifying uncertainty. A simple example of the theory is demonstrated when a number of people are asked to guess how many of an object, say marbles, are in a jar. One person would estimate one measurement while another would estimate a value slightly different. The more people who take a guess, the more likely that the average of their guesses is the correct amount. In this case, the rules of the game prevent one from being precise by taking out each object and counting them one by one. In a scientific setting, a rule like this is analogous to randomness in a system due to factors that may be known, but are very difficult to account for in the experiment. This is aleatoric uncertainty [1]. The value of a well-designed experiments is its ability to isolate some aspect so that one can study the effects of that specific thing. Removing all sources of aleatoric uncertainty is possible, but very difficult. Additionally, there are some aspects of an experiment that effect the result, but one is not aware of it to account for it. This is epistemic uncertainty [1].

The concept of uncertainty quantification (UQ) is closely related to the concept of a distribution of a random variable in statistics. Utilizing UQ to account for unknowns gained more relevance in the early 1990s when investigators were interested in incorporating physical knowledge about a system of interest with their statistical knowledge [1]. The advent of computers becoming more powerful and cheaper also enabled widespread adoption, as scientists could easily run thousands to millions of iterations of a simulation with slightly different settings each time. This approach allows the investigator to mitigate many sources of uncertainty related to parameter values, the selected model, or errors in observational measurements [2]. The outcome is a confidence interval in the output space that can help guide design decisions. This information can also be propagated to further models in a linked framework, a process known as Uncertainty Propagation (UP), to study its influence on quantities of interest (QOIs) further along the chain. Advancements in computing, including parallel computing frameworks, ever increasing processor speeds, and the development of processing hardware have allowed UQ and UP to be implemented across more complex simu-

lations and at larger scales [3].

Generally, many scientific studies address UQ unknowingly through experimental design. Because the goal of an experiment is to study the response of a system by a predetermined input, a good design isolates the effect of an independent variable(s) on some dependent one(s). Often in materials science, a model is developed to represent the experimental data. However, there is likely still some source of unknown influence on the system no matter how meticulous one is in ensuring conditions between experiments are identical. Therefore, many authors studying materials systems will include an uncertainty analysis in their report. It shows that the investigator(s) took as much in to account about the work as possible, within the limits of their ability to measure possible changes between experiments.

Similar to statistics, the concept of UQ is very general and thus useful in a wide range of industries including weather prediction [4], healthcare [5], and additive manufacturing [6]. It is also widely used in failure analysis and mitigation of structural components [7]. Its application is limited only by the imagination. From a computational standpoint, UQ is often used in conjunction with Integrated Computational Materials Engineering (ICME), an approach to designing materials using a combination of databases, computational models and simulations at multiple length scales [8, 9, 10]. Given the dependence of high-level simulations on each other, it is important that uncertainties associated with each type of simulation are appropriately accounted for and propagated, as to not compound errors between them. As more software is developed to make increasingly detailed connections between model chains, UP is critical for proper validation and verification of models and simulations [11].

The overall purpose of the work is to examine the impact of uncertainty in evaluating material performance. I present three cases where uncertainty is quantified at various stages in the evaluation process and propagated to reveal where the performance lies. In the first work, the aim is to quantify the performance of AF9628 at high strain rates using a combination of information sources including experimental data and models. An envelope of expected performance values for the newly developed material is presented in the stress-strain space and its influence on a down-

stream simulation is analyzed. Calibrating widely used models to this material's behavior extends the material's utility to broader applications. The goal of the second work is to assess the impact of an optimized weighting algorithm on the output distribution of a QOI. The algorithm is demonstrated across a high fidelity simulation, allowing one to obtain an accurate distribution of the QOI with a subset of all the available inputs. This allows scientists to be more confident in the distribution of their results if they are limited to a certain amount of evaluations due to computing constraints. The purpose of the third work is to identify if the predicted life of a Ti-6Al-4V specimen is influenced by the measurement techniques used. The approach considers the uncertainty associated with the slight differences in microstructure present in multiple samples of the same bulk material.

Each topic included is in support of advancing, not only materials knowledge, but also the mission of the Air Force (AF). The discovery of the high strength steel AF9628 already positions the AF ahead of the curve. The additional analysis of the material's behavior at high strain rates provides engineers with the knowledge required to make informed material selection decisions in applications of defense or offense. Applying the optimized weights approach across high fidelity simulations has the potential to reduce costs associated with running such simulations. If there are bottlenecks on supercomputer clusters, adequate UQ can still be accomplished via optimized weights. Finally, studying the influence of the current measurement techniques on Ti-6Al-4V component lifetimes is also important. How a microstructural feature is measured can affect the maintenance schedule of a fleet of equipment. Being more knowledgeable about the true life distribution of a part can lead to longer part life or more proactive maintenance procedures.

The benefits of UQ and UP in the context of materials science span the entire process structure property framework, operating at multiple structure hierarchy levels. At the thermodynamic level, Honarmandi et al. demonstrated that sources of uncertainty, like that in the models used to predict phase stability, can be accounted for and propagated into the phase diagram of a material. Additionally, various groups have shown that uncertainty in processing conditions can directly influence the expected material properties [12, 13]. Such is the case for additively manufactured materials

as well [14], which are gaining widespread adoption in various industries. Once manufactured and in service, quantifying uncertainty can also help assess the risk of failure associated with a given material subject to a specified loading condition [8]. The process of performing UQ is also not complex to implement, often times consisting of sampling parameter values from a Gaussian distribution. When implemented computationally, an investigator sidesteps the need to conduct large amounts of repeated physical experiments. Conducting such experiments physically can quickly become impractical in time and cost compared to the scale that can be achieved with computational methods. Indeed, the chance of unknown influences impacting experimental results increases with repeated physical experiments, as components of the apparatus wear down for example. Nevertheless, the experiments contain the most physically accurate response of the system, increasing their importance in the validation process.

In this set of work, UQ and UP are used to evaluate the performance of AF9628 and Ti-6Al-4V in a high strain rate and dwell fatigue applications, respectively. In the first work, an information fusion technique, Reification, is utilized to generate uncertainty bounds for the high strain rate response of the material [2]. It combines the available information from experiments and models calibrated to the material. Markov Chain Monte Carlo provides an efficient way to fit the models and generate distributions of the model parameters. The second work implements the Probability Law Optimized Weights (PLOW) approach based on work by Sanghvi et al. [15] to estimate a property of AF9628 across a high fidelity Taylor-Anvil simulation. Lastly, a geometric algorithm known as Unfolding [16] is employed to estimate the three-dimensional size of detrimental inclusions in Ti-6-4 from two-dimensional measurements obtained in-part via electron back-scatter diffraction. A statistical technique known as Expectation Maximization is also applied to the data set to compare any differences in part life predicted between the two methods.

2. PARAMETER RANGE ESTIMATION AND PROPERTY UNCERTAINTY ANALYSIS FOR HIGH STRAIN RATE MATERIAL RESPONSE VIA MODEL FUSION

2.1 Abstract

The dynamic stress-strain response of a material can be described by a number of different models of varying fidelity. However, the inherent fidelity of the model itself can be hindered by experimental variability. A reification approach is presented for fitting a model to experimental measurements with inherent scatter, and is demonstrated for determining optimum parameters of a Johnson-Cook (J-C) and Zerilli-Armstrong (Z-A) models. A fused model is generated by combining the portions of the data that the JC and ZA model fit the best, and generating a reified dataset which represents a “best-guess” sampling of the possible stress-strain response of a high strength steel. Under the reification framework, a best guess range of high strain rate parameter values are presented for each model, given the unavoidable variability of SHPB results and limitations of models to fit the data. The reified J-C models, representing the upper and lower bounds of reification, are then compared with the J-C and Z-A models fit to the mean response of the stress-strain experimental curves via evaluation of material properties determined using CTH.

2.2 Introduction

Understanding the response of structural components to high-strain-rate deformation is essential to a wide range of applications. Predictive finite element analysis and design of structures require the use of realistic constitutive models of the high-strain-rate response of the material. Unfortunately, high-strain-rate experiments are subject to considerable sources of uncertainty that are, in general, difficult to deconvolute and/or isolate. This uncertainty poses a significant challenge for the calibration of constitutive models. Consequently, these calibration limitations make reliably predicting the mechanical response of a structure subjected to high-strain-rate deformation ($10^3 - 10^6 \text{ s}^{-1}$) challenging [17].

Many of the most commonly used constitutive models that describe a materials' high-strain-

rate behavior do not account for deformation mechanisms equally. Furthermore, many of these models are designed for intermediate strain rate regimes ($10^2 - 10^4 \text{ s}^{-1}$) [17], and cannot be easily extrapolated to higher strain-rates due to different deformation mechanisms activated and complex triaxial states of stress caused by shock wave propagation. Beyond this strain rate, one-dimensional strain becomes increasingly difficult to maintain in bulk specimens, making model predictions of the stress-strain plot beyond the very high-strain-rate regime less accurate [18].

The many models currently available vary in their fidelity and their degree of agreement with the underlying experiments. The different model assumptions imply that they may agree with experimentally-determined materials response at different regimes or stages of the deformation process. In essence, one could expect that each individual model contains some useful information about the system being investigated—in the present case, a high strength steel under high-strain-rate deformation—while not being able to describe the response of the material in its entirety. Since all models may predict different regimes of material response with varying fidelity, a logical corollary is to combine or ‘fuse’ them to maximize their ability to describe natural phenomena, even in the presence of considerable experimental uncertainty. A potential approach to carry out this model fusion is *Reification* [2], which exploits the statistical correlation between different models as well as between the models themselves and the experimental data in order to build an effective fused model. The fused information can then be further parameterized by refitting constitutive model forms for implementation in modeling and simulation codes. The constitutive models (with uncertain model parameters) can then be used to examine the response of structures subject to high-strain-rate deformation, effectively propagating the uncertainty arising from fitting constitutive models to uncertain experiments, providing more robust predictions with uncertainty bounds.

Two ubiquitous constitutive model forms used for predicting dynamic material response include the Johnson-Cook (JC) and Zerilli-Armstrong (ZA) models [19, 20]. Typically, experiments performed using the Split Hopkinson Pressure Bar (SHPB) method are employed to determine the high-strain-rate parameters of the models [21]. The Johnson-Cook model is the most widely-

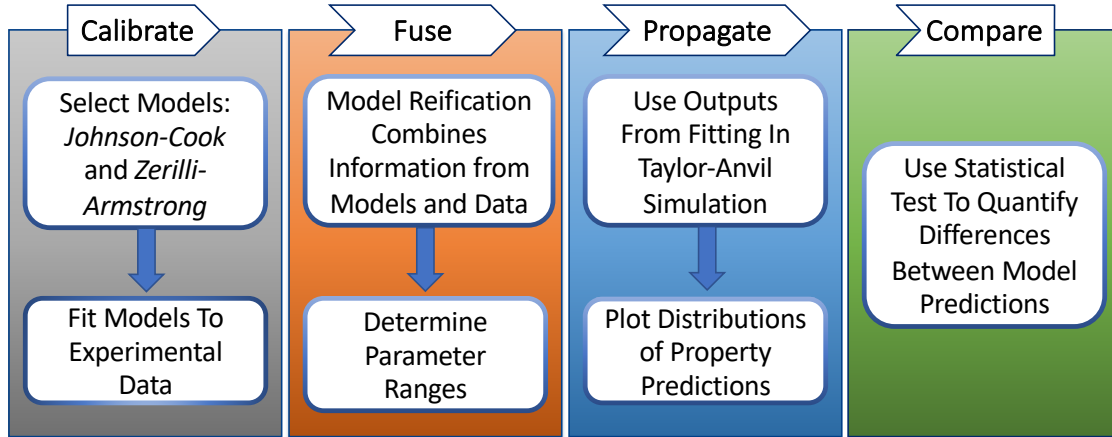


Figure 2.1: Framework used to determine parameter value ranges and simulation output distributions.

used due to its relative simplicity and ease of calibration given limited experimental data. The Zerilli-Armstrong model, on the other hand, can take into account microstructural effects on the plasticity of the material, if such information is available. Both models have been used extensively in scientific literature.

For example, Banerjee et al. used the Johnson-Cook model to predict the behavior of an armored steel in the finite element code ABAQUS [22, 23]. They also propose a modified version of the Johnson-Cook damage model to predict behavior for a wider variety of loading conditions than with the original formulation. Ultimately, the model was shown to have good agreement with experiments. However experiments served as the only information source to calibrate the models, although information from existing calibrations of a similar material to the JC model was available [24].

The Zerilli-Armstrong model, a more microstructurally sensitive constitutive model than JC, has also been used to predict the stress-strain behavior of materials under high-strain-rates, e.g. in tantalum [25]. There is good agreement with experimental data, but the equation loses accuracy near $4.9 \times 10^3 \text{ s}^{-1}$. Mirzaie et al. recently made a slight modification to the ZA equation to account for dynamic recovery and recrystallization that occurs in carbon steel during tensile tests at high temperatures (roughly 900 — 1100 °C) [26]. It has been shown that the fitting allows the predicted

response to “bend” closer to the experimental data, however this process isn’t applied to strain rates higher than 0.1 s^{-1} . Here again, external work to modeling material behavior using the ZA equation has been cited by the authors, but only experimental data was used in calibrating the model.

Each of these models are constructed under different assumptions, ignoring the possibility that they may provide complementary information in different regimes of the materials response. Moreover, the high-strain rate behavior of materials is challenging to properly parameterize due to the considerable experimental variability in Split Hopkinson Pressure Bar tests as nominally identical experimental conditions can result in significantly different measured responses. Given the uncertainty in experiments and the incompleteness of the most commonly used constitutive models to high-strain-rate deformation, it is worth exploring approaches to maximize the useful information derived from these models in combination with available experimental data.

In the present work, the reification method is utilized to combine the Johnson-Cook and Zerilli-Armstrong model forms, and experimental data to determine a range of appropriate parameter values capable of modeling the stress-strain response of a new high strength steel, AF9628 [27], at high-strain-rates. Experimental data is used as ground truth in determining the reified response of the system. First the JC and ZA equations are calibrated with temperature effects taken into account. The models are then fused, which generates a fused dataset, herewith called the “reified model”. The reified model is then partitioned into a confidence bound, capturing the material’s experimental variability in stress. The JC and ZA model forms are then re-calibrated to the partitions of the fused dataset using a Markov Chain Monte Carlo (MCMC) calibration approach. Lastly, the parameterized JC and ZA model forms are employed in a Taylor anvil (TA) simulation to compare the response distributions predicted by each of the models.

2.3 Methods

Figure 2.1 shows an overview of the steps used in the current work. In order to create the reified model, it was first necessary to determine the constants of both the JC and ZA equations by calibrating them to experimental data; the calibration step. Multiple tests at nominally identical

Table 2.1: Tension test conditions.

Experiment ID	Gauge Diameter [in.]	Sample Gauge Length [in.]	Strain Rate [s^{-1}]	Temperature [$^{\circ}C$]
Experiment 1	0.125	0.350	0.01	20
Experiment 2	0.125	0.350	0.01	20
Experiment 3	0.110	1.142	1440	300
Experiment 4	0.110	1.140	1400	300
Experiment 5	0.110	1.139	1430	300
Experiment 6	0.122	1.139	1350	300
Experiment 7	0.110	1.140	1440	300

input conditions were conducted to calibrate the JC and ZA models to capture the experimental and material variability. Next, model refication was utilized to fuse information from the models and experiments; the fusion step. The JC and ZA models were then re-calibrated to the resulting reified data points via MCMC. The resulting parameter values were propagated into Taylor anvil simulations; the propagation step. Finally, the outputs from the Taylor anvil simulations were compared via hypothesis testing; the comparison step. This general framework is applicable to other cases, even beyond materials science, and is fully extensible.

2.3.1 Experimental Procedure

Experimental data for this work was obtained using Split Hopkinson Pressure Bar (SHPB) tests [18] performed at the University of Dayton Research Institute (UDRI) in Dayton, OH. The SHPB experiment probes the high-strain-rate response of materials by accelerating a striker bar into an input bar, which can place the specimen into either compression or tension depending on the transmitting bar configuration. The impact transmits a pressure wave through the material and the stress-strain response under the dynamic loading condition is measured by strain gauges mounted on the bars, assuming dynamic equilibrium. In the present work, the test configuration is indirect tension where a compression wave propagates through the input bar and specimen, reflects at a momentum trap, then returns as a tension wave which fractures the specimen.

For the tests, five subscale E8 cylindrical dogbone specimens of AF9628 were prepared using

careful machining practice and low stress grinding to ensure a machine finish $\leq 32 \mu\text{m}$ roughness for use in the SHPB apparatus, fitted with maraging steel bars. Tight dimensional tolerances, surface finishes, and stereoscopic measurements were made to ensure specimen pedigree and experimental quality. The specimen dimensions and test conditions are shown in Table 2.1. Additionally, the input conditions controlling the striker bar were kept constant. Thus, any variations in output conditions for the five high-strain-rate tests stemmed from material and experimental variability. Variations in strain rate arise from variation in impact velocity and losses in the system.

2.3.2 Determining Johnson-Cook Parameters

The JC model was selected due to its ubiquity and ease of use. Its constants can be determined using a few uni-axial tests at various strain rates. Overall, the model pre-supposes that the total contributions of strain hardening, strain rate effect, and strain rate sensitivity can all be decoupled in a multiplicative fashion with each parameter being independent. The model has the form shown in Equation 2.1:

$$\sigma_p = \left(A + B\epsilon^n \right) \left(1 + C \ln \left(\frac{\dot{\epsilon}}{\dot{\epsilon}_o} \right) \right) \left(1 - \left(T^* \right)^m \right). \quad (2.1)$$

where A is the yield strength of the material, B , n , C , and m are model constants, ϵ is the plastic strain, $\dot{\epsilon}$ is the strain rate, and $\dot{\epsilon}_o$ is the reference strain rate [22]. T^* , the homologous temperature [28], has the form

$$T^* = \frac{T - T_o}{T_m - T_o}, \quad (2.2)$$

where T is the temperature the material is subject to or being tested at, T_o is the room temperature, and T_m is the melting temperature of the material [22].

Monotonic Constants: A , B , and n

Johnson and Cooks' model has three terms that account for different mechanisms which affect the

strength of the material. The first term

$$(A + B\epsilon^n)$$

accounts for the slow ($\approx 1 \times 10^{-2} \text{ s}^{-1}$) strain rate response of the system and is recognized as the Ludwik-Hollomon equation for the plastic term of the stress-strain response [29]. To find the constants A , B , and n , two specimens were tested in tension at a slow, monotonic strain rate of 10^{-2} s^{-1} . Each test was performed in accordance with the ASTM E8 standard [30]. Dimensions of the specimen are displayed in Table 2.1 under Experiment 1 and 2. The constant A , in Equation 2.1, was set equal to the average 0.2% yield stress from the tests conducted at 10^{-2} s^{-1} .

The maximum engineering stress value of each specimen marked the end of the plastic region. The set of points, from the yield stress to the end of the plastic region, was used for fitting the B and n constants via a least squares optimization routine. With these three constants, the material response of the system could be modeled. The other multiplicative terms in the JC equation simply act as modifiers to the first term in order to account for the effect of high-strain-rates and temperature changes on the strength of the material.

Strain Rate and Temperature Constants: C and m

A Markov Chain Monte Carlo approach was used to calculate the strain rate constant C and temperature constant m , of the JC model

$$\left(1 + C \ln \left(\frac{\dot{\epsilon}}{\dot{\epsilon}_o}\right)\right) \left(1 - \left(T^*\right)^m\right),$$

while holding the other parameters, A , B , and n , at their previously determined values. The model was calibrated against the 5 high-strain-rate tests performed at 300 °C. This data was used in the likelihood distribution via the sum of squares error function. The average value from the five SHPB experiments of $\approx 1412 \text{ s}^{-1}$ was used for the strain rate.

2.3.3 Determining Zerilli-Armstrong Parameters

Unlike the Johnson-Cook model, the Zerilli-Armstrong model takes into account some microstructural effects on the stress of the material. There are in fact, two separate ZA equations; one for alloys with a BCC structure and one for alloys with an FCC structure. The equation for the BCC case, which is applicable to the low-alloy tempered martensitic steel AF9628, is

$$\sigma = C_0 + C_1 e^{-C_3 T + C_4 T \ln \dot{\epsilon}} + C_5 \epsilon^n. \quad (2.3)$$

Here, $C_0 = \Delta\sigma'_G + kl^{-0.5}$, which represents the microstructural contributions to the yield strength of the material. Within C_0 , k and $l^{-0.5}$ represent the microstructural stress intensity factor and the inverse square root of the average grain size diameter of the material, respectively. T , $\dot{\epsilon}$, and ϵ represent the absolute temperature at which one would like to predict stress, the strain rate, and the strain, in that order. The terms C_1 , C_3 , C_4 , C_5 , and n are material constants.

The Zerilli-Armstrong equation is unique in that the thermal component of the equation can be derived without considering the strain for BCC metals [20]. Interestingly, this thermally activated stress component, is associated with the movement of a single dislocation [31]. This is contrary to FCC metals where thermal activation behavior is dependent on strain. Lastly, the term $C_5 \epsilon^n$ accounts for the strain hardening component of stress in BCC metals and is assumed to have a power law dependence on strain.

The first term in the ZA model was added to the microstructural component, and set to the same value as A in the JC model, as this value represents the same quantity, the yield stress of the material. Determining this value with the obtained experimental data proved useful in the current work since the k and l values for the microstructure had not been measured for the working material. Each of the other terms were determined using MCMC against the experimental data, while holding C_0 constant. Because the equation is additive, the model could be fit to each of the remaining parameters simultaneously. The absolute temperature at which the experiments were done was set as the value for T . Like the JC model, the average strain rate of the five SHPB

experiments was used for ϵ during calibration.

2.3.4 Reification Technique

Model reification can be applied to an arbitrary number of information sources where the statistical correlation between the predictions of the information sources is known. The case of two models will be used for the following explanation and for the presented work. Once the correlation coefficient is calculated, it is then used to calculate the mean and variance of the fused estimate given by Equation 2.4 and 2.5, respectively. This mean and variance are based on a combination of the different information sources. The derivation of the expressions

$$\mathbb{E}[y] = \frac{(\sigma_2^2 - \rho\sigma_1\sigma_2)\bar{f}_1(\mathbf{x}^*) + (\sigma_1^2 - \rho\sigma_1\sigma_2)\bar{f}_2(\mathbf{x}^*)}{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2}, \quad (2.4)$$

$$Var(y) = \frac{(1 - \rho^2)\sigma_1^2\sigma_2^2}{(\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2)}, \quad (2.5)$$

can be found in Winkler [32]. Here, σ_1^2 and σ_2^2 represent the variances of the input models, while ρ represents the correlation coefficient.

To find the correlation coefficient, it is first assumed that the outputs of the models to be fused, $f_1(x)$ and $f_2(x)$, are normally distributed with a mean \bar{f}_i and variance σ_i , where $i = \{1, 2\}$. The outputs of the models can be written as $f_1(\mathbf{x}^*) \sim \mathcal{N}(\bar{f}_1(\mathbf{x}^*), \sigma_1^2)$ and $f_2(\mathbf{x}^*) \sim \mathcal{N}(\bar{f}_2(\mathbf{x}^*), \sigma_2^2)$ where \bar{f}_1 and \bar{f}_2 are the output values of the respective functions given certain input values, \mathbf{x}^* . The distribution in the output is a result of the sources of uncertainty that are added to the function output given the input \mathbf{x}^* and can be written as $f_i(\mathbf{x}^*) = \bar{f}_i(\mathbf{x}^*) + \delta_i(\mathbf{x}^*)$ where $\delta_i(\mathbf{x}^*)$ represents uncertainty associated with the quantity of interest modeled by the function. Sources of uncertainty include random errors in the test data to which a model is to be fit as well as model inadequacy. The model parameters themselves can also be a source of uncertainty when allowed to vary as a joint distribution between each other. Further details on the types of uncertainty that can affect model prediction have been discussed by Thomison and Allaire [2].

Next, one of the models is selected to represent the true sampling of the population of possible

responses of the system, and the covariance between this model and the remaining model is calculated. In the case of model 1 being considered real, the covariance between the two models can be represented as

$$\mathbb{E}[\tilde{f}_1(\mathbf{x}^*)\tilde{f}_2(\mathbf{x}^*)] = \sigma_1^2. \quad (2.6)$$

This is used to calculate the correlation coefficient shown in Equation 2.7:

$$\rho_1(\mathbf{x}^*) = \frac{(\sigma_1^2)}{(\sigma_1\sigma_2)} = \frac{\sigma_1}{\sqrt{\left(\bar{f}_1(\mathbf{x}^*) - \bar{f}_2(\mathbf{x}^*)\right)^2 + \sigma_1^2}}. \quad (2.7)$$

Next, the second model is considered real and Equation 2.7 is used again to find $\rho_2(\mathbf{x}^*)$. These values are then averaged using

$$\bar{\rho}(\mathbf{x}^*) = \frac{(\sigma_2^2)}{(\sigma_1^2 + \sigma_2^2)}\rho_1(\mathbf{x}^*) + \frac{(\sigma_1^2)}{(\sigma_1^2 + \sigma_2^2)}\rho_2(\mathbf{x}^*). \quad (2.8)$$

The mean correlation coefficient, $\bar{\rho}(x^*)$, is then used in Equations 2.4 and 2.5 to determine the mean and variance of the fused model. In its current version, the output of the reification algorithm is a set of points within the response space of interest. The points include the mean and a confidence region for the prediction of data above and below the mean.

As shown in Equation 2.8, the characteristics of the constituent models influence the mean correlation coefficient, which in turn influences the behavior of the fused model. A visual example depicting the relationship between the fused model and the characteristics of the constituent models has been shown by Thomison and Allaire [2]. Generally, the higher the mean correlation coefficient is, the more the fused model would predict away from both models. Intuitively, this makes sense in that if two constituent models are highly correlated, they both would predict the same response, be it too low or too high compared to the experimental value.

2.3.5 Determining the Reified Model's Parameters

Because experimental data was available, it was incorporated into the fusion process to increase the accuracy of the fused model. The average of the 5 high-strain-rate, high temperature experimental data points served as the ground truth in the algorithm. We assume that the stress response is normally distributed and that the average stress of the 5 replicated experiments at each strain would best represent the true response of the material.

To obtain the average stress at each strain value, splines were first fit to each set of experimental data points separately to ensure co-registration of the data at identical strain values. Once the splines were fit to each experimental data set, a stress value was taken from each of the 5 specimen at periodically spaced strain values. The five stresses at each strain were then averaged. The average stresses are used to represent the real-world response in the reification process [2].

Once the information sources were fused, a nonphysical artifact was discovered in the resulting reification output. The artifact, discussed further in Section 2.4.2, is a result of the reification algorithm switching its “trust” from the JC model at lower strains to the ZA model at higher strains. Indeed, at lower the strains, the error between the JC model and experimental data is less than that of the ZA model and experimental data. The opposite is true at higher strains. Removing this artifact by partitioning the reified mean resulted in two possible stress-strain responses, which are possible samples from the population distribution of possible stress-strain outcomes. In the present work, they represent a lower and upper bound estimate of the true response of the material.

The JC and ZA models were each fit to both the upper and lower bounds of the reified data for use in the multimaterial Eulerian hydrocode CTH (Version 11.1, Sandia National Laboratory). The details of the fitting are described in Section 2.3.6. Fitting JC and ZA models to the upper and lower bounds of the reification output would provide the benefit of making a direct comparison between the parameter values of the models fit to the experimental data versus the same models fit to the reified mean data.

2.3.6 Propagation to Taylor Anvil Simulations

Once the models were re-calibrated to the fused, or reified, data points, the parameter values could be propagated into the TA simulation. This section describes the details of the process.

2.3.6.1 *Performing MCMC and Medoids*

To determine if a statistical difference between the predictions of the models existed, the outputs from the MCMC chain were used as the model parameter values within the Hydrocode [33]. MCMC provided a large set of possible parameter values to be used for each model type. Additionally, a wide range of parameter values could be accessed given that the MCMC algorithm is not as vulnerable to getting trapped in local minima in contrast to a gradient based minimization procedure.

The MCMC analysis was done with a uniform prior distribution on all of the parameters for each model. Reasonable upper and lower parameter bounds were set based on values found in literature for steel alloys similar in strength to AF9628. From here, at least 1 million parameter sets were calculated for each of the model types: JC, ZA, JC-Upper, JC-Lower, ZA-Upper, and ZA-Lower. Model types with a dash are named with regard to the data that equation was calibrated against. For instance, the JC-Upper model refers to the JC equation being fit to the upper bound of the reified data points. For fitting to the upper and lower reification data, all parameters were allowed to vary to capture the behavior of the reified mean ($\mathbb{E}[y]$) as accurately as possible. The JC and ZA models on the other hand, were fit to the high-strain-rate experimental data. In the JC and ZA fits, only the high-strain-rate parameters C and m (for JC) and C_1, C_3, C_4, C_5 , and n (for ZA) were allowed to vary.

The data each model was fit to was used in the likelihood function of MCMC. After each run of MCMC, 200,000 values were removed from the beginning of the parameter chain sets in order to exclude the burn-in region. The result was six tables of 800,000 parameter value sets corresponding to the six models used. In the present work, the term “point” is used to describe one parameter value set. For example, one point from an MCMC parameter calibration to the JC model would be

a vector of 2 values.

Due to the computational expense of running the remaining 800,000 points, k-medoids was used to down-select points to a manageable set[34]. The k-medoids approach was favorable in that it could divide the parameter value sets into k clusters desired by the user. Once partitioned, the algorithm selects a representative value from each cluster that already exists within the original data set (Contrast this with the k-means algorithm which calculates an average value to represent each cluster). K-medoids removes the possibility of parameter values being selected that may not have been determined via MCMC.

Of the 800,000 parameter sets generated from MCMC, 500 were selected using from 500 clusters partitioned by k-medoids. Simultaneously, the data set can still produce a well defined distribution of values once run through the hydrocode. Another benefit of using 500 clusters is that good representative values can be selected to represent each cluster. This is due to the fact that there are more points within each cluster for the algorithm to choose from. The situation is analogous to determining the mean of a random variable by the central limit theorem.

2.3.6.2 *Performing Taylor Anvil Simulations*

Once the parameter values were selected, Taylor anvil simulations were conducted to compare the material models. The Taylor anvil simulation consists of accelerating a cylindrical specimen into an anvil. 2D simulations were conducted for this study. Ultimately, the goal of utilizing the Taylor anvil simulation was to determine the material's response under uni-axial stress conditions. A two-dimensional visualization of the Taylor anvil simulation is shown in Figure 2.2.

As shown in Figure 2.2, each simulation was set with a line of tracer particles at the front of the rod to record time-varying data. The node locations are depicted in Figure 2.2 as a vertical, black line on the rod. These nodes were set to record temperature, pressure, particle speed and position in the x and y directions, and the stress deviator invariant, J_2 . Additional values recorded with each simulation were the impact rod's final length and the final diameter of the impact end.

CTH was used to perform the Taylor anvil simulations. Mie-Grüneisen equations of state were determined for the steel alloy, while the JC and ZA constitutive model constants were modified

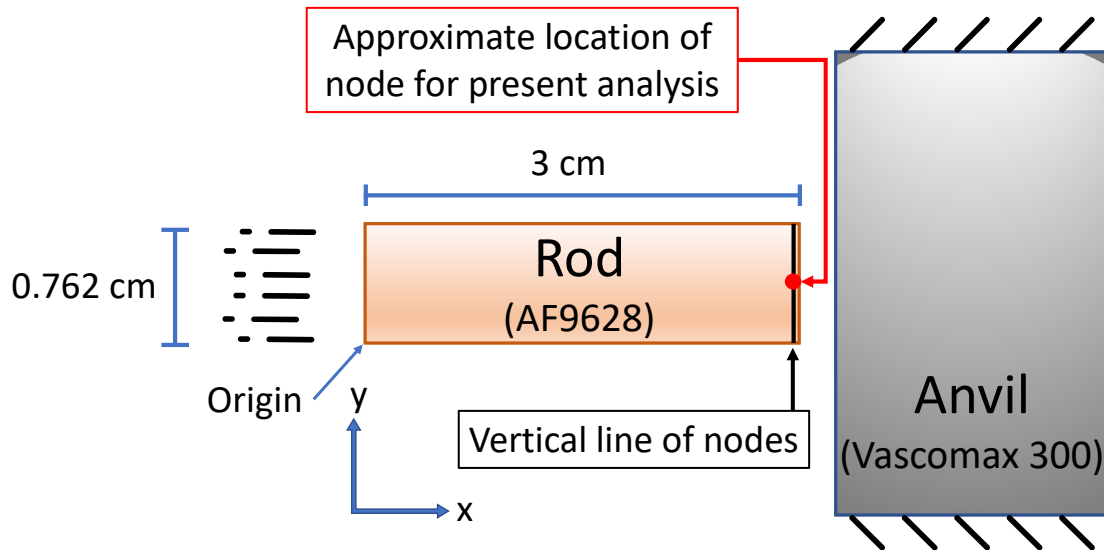


Figure 2.2: Depiction of the 2-dimensional Taylor anvil simulation. The specimen (Rod) dimensions are noted as well as the location at which the response data is recorded for the present study.

based on the fitting procedure results. For the anvil, the material constants of the alloy Vascomax 300 were used. The simulation would be run for the JC, ZA, JC-Upper, and JC-Lower models.

2.4 Results

2.4.1 Experimental Data and Model Fits

The alloy, AF9628, was tested at different strain rates as shown in Table 2.1. Figure 2.3 shows negligible variability at $0.01s^{-1}$ for the monotonic tension tests performed at room temperature. The ultimate tensile stress, 0.2% offset yield stress, and modulus for both monotonic tension tests are displayed in Table 2.2.

In Figure 2.4, the stress-strain plots of the 5 SHPB tension experiments which were run at a nominal temperature of $300^{\circ}C$ is shown. The striker bar impact velocity induced strain rates that ranged from $1350 s^{-1}$ to $1440 s^{-1}$. The output data was used for calibrating the JC and ZA equations. The constants for each model are tabulated in Table 2.3.

As shown in Figure 2.4, the experimental points in the plot also display a strain hardening effect with an increase in strain, which is commonly observed in tempered martensitic steel. However,




Figure 2.3: Monotonic tension response for AF9628 tested at room temperature and strain rate of 10^{-2} s^{-1} .

Table 2.2: Monotonic stress test results

Experiment	UTS [MPa]	0.2% Offset YS [MPa]	E [GPa]
1	1722	1330	216
2	1702	1335	208

there is significant variability between nominally similar experiments. This variation in response hints to some variable that may not be accounted for between experiments such as small differences in experimental setup, slight microstructural differences between specimen, some other unknown factor, or a combination of factors.

2.4.2 Model Fusion Parameter Results

Figure 2.5a shows the JC and ZA models plotted together, along with the average of the 5 specimen tension data at 300°C. The average is taken because it represents the best estimate of the

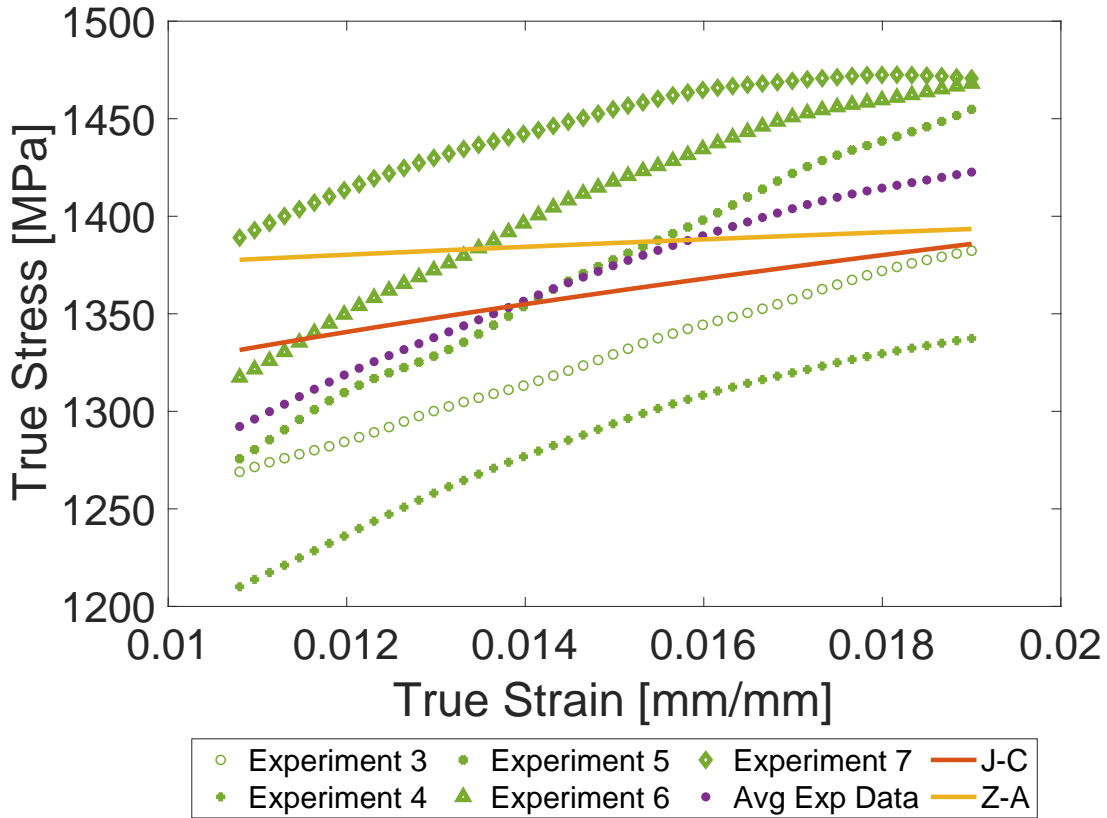
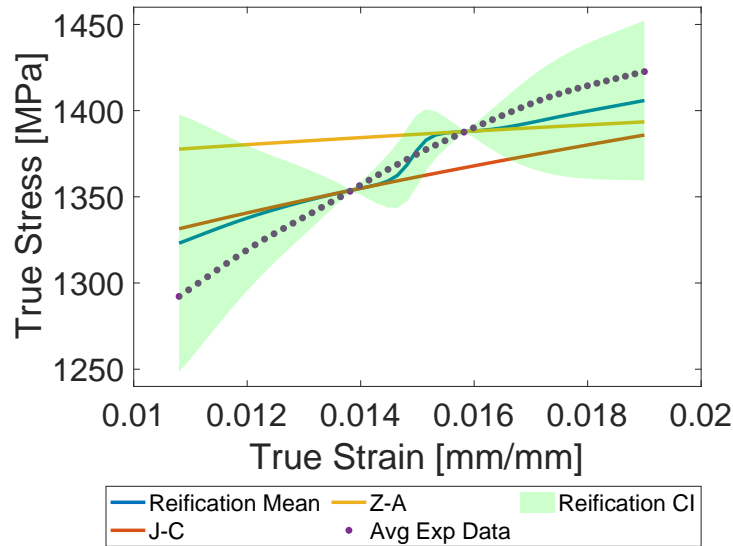



Figure 2.4: Dynamic stress-strain tension response of AF9628. The experimental data points from each of the five experiments are shown in green. The filled purple circles are the average stresses of the 5 experimental data points.

material response given the limited, and highly variable, data. With this information, the reified mean could be calculated and plotted. The 95% confidence interval for the reified model is also plotted as a shaded region, which represents where the reified model expects experimental data to exist given the input conditions.

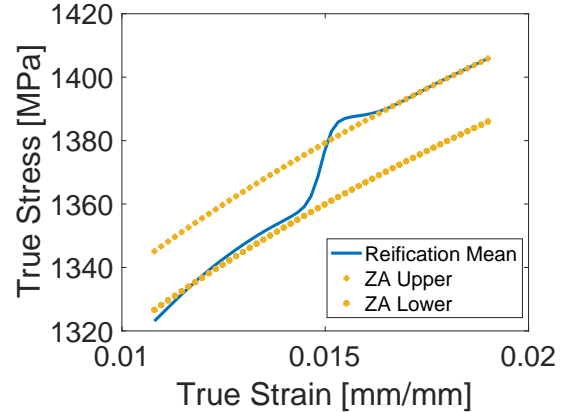
As shown in Figure 2.5, the curved profile of the reification mean encompasses both the JC and ZA models. More explicitly, at lower strains, the reification mean is lower than the prediction of both models and at higher strains, the reification mean is higher than the prediction of both models. The reification mean profile also abruptly rises in the middle strains (roughly 0.014 to 0.016). This could suggest that the material is experiencing a period of hardening illustrated by a sudden increase in stress at roughly 0.015 [mm/mm]. However, such hardening can only physically occur



(a)



(b)



(c)

Figure 2.5: (a) Reification mean and 95% confidence interval (CI) plotted with the average experimental data, and High-Strain-Rate Models fit to the data. (b) Reification mean with the upper and lower JC models representing the response of the reification mean. The parameters of the upper and lower models are the mode of the MCMC chain values excluding the burn-in region. (c) Same as (b) except the upper and lower ZA models are displayed. The parameters of the upper and lower models are the mode and mean MCMC chain values with the burn-in region removed, respectively. For (b) and (c) the chain statistic (mean or mode) used was that which minimized the root mean squared error between the reified mean and the model. The exact parameter values for each model are recorded in Table 2.3.

Table 2.3: Parameter values of the models used in the study.

System	A [MPa]	B [MPa]	n	-	C	m
Johnson-Cook Upper	1332.9	1565.8	0.3263	-	0.0367	0.5249
Johnson-Cook Lower	348.1	1632.5	0.1257	-	0.0159	1.1336
Johnson-Cook	1332.5	1565.8	0.4000	-	6.19×10^{-5}	1.2856
System	C_0 [MPa]	C_1 [MPa]	n	C_3 [K ⁻¹]	C_4 [K ⁻¹]	C_5 [MPa]
Zerilli-Armstrong Upper	984.3	7323.4	0.2755	0.6522	0.0036	1256.2
Zerilli-Armstrong Lower	992.7	3902.4	0.2896	0.5386	0.0365	1239.2
Zerilli-Armstrong	1332.5	7128	0.5309	0.6696	0.0858	500

if the material experiences a change in phase, temperature, or strain rate during the test, which was not the case for these experiments. As a result, it is more likely that the abrupt rise in stress is a nonphysical artifact of the reification procedure.

However, the experimental data and its average show strain hardening behavior that is not constant. The model forms may not sufficiently capture the hardening response in the parabolic form and the material constants may show some dependency on strain level. Considering the abrupt rise in stress predicted by the reification mean and restrictive form of the models, it became clear that fitting the dynamic models to partitions of the reification mean would best capture the experimental variability. Furthermore, the parameter values of the models calibrated to the partitions could serve similar to uncertainty bounds for AF9628 under the present loading conditions. The fits of these models to the upper and lower reification mean using MCMC are shown in Figures 2.5b and 2.5c. The chain values from the MCMC run were propagated into the Taylor anvil simulation. It is important to note that the upper and lower reification mean does not represent the scatter in the raw experimental data.

The parameter values for the two models using the upper and lower bounds of the reification mean in the MCMC process are shown in Table 2.3. They can also be used as ranges for acceptable parameter values, given that one runs a SHPB tension experiment at 300 °C and roughly $1000s^{-1}$ on the alloy, or a similar one to, AF9628. Given the variability of the data and limits of the models to represent reality, these parameters are the best possible guesses for the material using the




Figure 2.6: Last frame of Taylor anvil impact simulations. Each image shows the material’s response from the perspective of the (a) von Mises stress (b) triaxiality (c) pressure and (d) temperature. All screenshots are shown at the same impact time of 20 microseconds.

Johnson-Cook or Zerilli-Armstrong models in a Reification framework. Additionally, the values for the Johnson-Cook parameters are similar in magnitude to those determined by Hopson et al., who calibrated a Johnson-Cook model to Eglin steel, a steel alloy similar both in composition and typical thermal processing route to AF9628 [35]. Consequently, it can be assumed that these values are valid for the Zerilli-Armstrong model since they are both fit to the same upper and lower data sets.

2.4.3 Taylor Anvil Simulation Results

Figure 2.6 shows a screenshot for one of the Johnson-Cook informed simulations. Each screenshot shows the response of the specimen and anvil plotting different field variables, such as the von Mises Stress, triaxiality, pressure, and temperature. Each image shows the final time step of the simulation. The shape is a result of the interactions between the impact induced shock-waves and the specimen’s boundaries. Temperature plots reveal shear banding at instability regions near the unconfined lateral deformation area.

Table 2.4: Statistical significance comparison of Taylor anvil simulation outputs between model pairs. Values in the cells are the D -values calculated between the two output CDFs. Shaded cells represent statistical similarity between the two models for the specified output (i.e. A D -value so small that H_0 , the null hypothesis that two distributions are from the same sample, can not be rejected in a Komolgorov-Smirnov 2-Sample test). The CDFs of each of the outputs are shown in Figure 2.7.

System	Length	Width	Pressure	Temperature	X-Velocity	Y-Velocity	J2P	Density
JCU vs JCL	0.3480	0.5860	0.5320	0.6000	0.1820	0.1540	0.6580	0.5440
JC vs JCL	0.1345	0.1908	0.0938	0.2547	0.2597	0.0796	0.2077	0.0768
JC vs JCU	0.3828	0.4770	0.4991	0.4410	0.2170	0.1486	0.5031	0.4871
ZA vs JCU	0.7355	1.0000	0.6079	1.0000	0.8558	0.1265	0.6957	0.7400
ZA vs JCL	0.6855	0.8660	0.2714	0.5300	0.8198	0.1495	0.4737	0.3000
JC vs ZA	0.6293	0.7635	0.2926	0.5952	0.9078	0.1804	0.3086	0.3587

The distributions of outputs from the Taylor anvil simulations are shown in Figure 2.7. The cumulative distribution functions (CDFs) represent the effect of parameter selection on simulation output prediction. They also show the probability of certain values being taken on by the material, based on the model type used. The outputs predicted during the simulation include the final length and width of the specimen on the impact end. Additionally, the pressure, temperature, particle velocity in the x-direction, particle velocity in the y-direction, J_2 stress deviator invariant (converted to von Mises stress, σ_v), and density were all recorded in the last frame of the simulation at the location indicated in Figure 2.2. In the present work, only the distributions of the final length, width at impact end, von Mises stress, and temperature are shown. The length, width, and temperature were selected for display because they represent quantities about the specimen that could be recorded during a real world TA experiment. Von Mises stress, on the other hand, gives insight into whether the material will yield or not.

Interestingly, for most of the simulation outputs, the upper and lower bounds of reification, represented by JC-Upper and JC-Lower, respectively, bound the response of the JC model. This reflects the behavior of said models in the the stress-strain space, previous to propagating through the TA simulation. Additionally, the distributions from the TA simulations show how slight changes

Table 2.5: Statistical significance comparison of re-weighted Taylor anvil simulation outputs between model pairs. Values in the cells are the D -values calculated between the two output CDFs. Shaded cells represent statistical similarity between the two models for the specified output (i.e. A D -value so small that H_0 , the null hypothesis that two distributions are from the same sample, can not be rejected in a Komolgorov-Smirnov 2-Sample test). The CDFs of select outputs are shown in Figure 2.8.

System	Length	Width	Pressure	Temperature	X-Velocity	Y-Velocity	J2P	Density
JCU vs JCL	0.3461	0.5915	0.5954	0.6193	0.1455	0.0150	0.6895	0.6046
JC vs JCL	0.1137	0.1913	0.0435	0.2499	0.2236	0.0167	0.1996	0.1167
JC vs JCU	0.3807	0.4656	0.3634	0.4086	0.1694	0.1342	0.4468	0.5085

in parameter values can cause differences in simulation results, emphasizing the importance of proper parameter selection for JC and ZA models for a given set of experimental conditions. For example, in Figure 2.7C the ZA outputs for Final von Mises stress are spread from roughly 1300 MPa to 2200 MPa. Each of these outputs, and the predicted lengths in between, originate from a different set of parameter values selected using k-medoids. Moreover, when comparing the values for ZA-Upper and ZA-Lower in Table 2.3, the differences between their parameters, such as n and C_5 for instance, appear to be minor. However, these seemingly small differences in magnitude for each parameter can cause a wide range of predictions when propagated into a simulation program.

Furthermore, the spread of the points in the plots of Figure 2.7 show how predictable each model is compared to the others. In essentially every case, the JC-Upper model predicts outputs most consistently. For example, in Figure 2.7A, all of the predicted Final Length values for the JC-Upper model lie between 2.75 and 2.8 cm. Each of the other models predict values from 2.75 cm to beyond 2.8 cm. The JC-Upper model also reports the lowest median value for a given output in most cases. The ZA model, on the other hand, predicts with the most variance in all but the von Mises case, Figure 2.7C. In Figure 2.7D, the ZA model predicts Final Temperature values from ≈ 330 K to ≈ 375 K, while the others predict values from only ≈ 320 K to ≈ 355 K.

As previously mentioned, a large amount of potential inputs to the TA simulation were not run in order to reduce computational costs. However, statistical techniques can be used to supplement this issue by estimating the distribution that would have resulted if all inputs were able to be

run. The probability measure optimized importance weights (PLOW) is one such method [15]. The PLOW method is able to estimate the full, or target, distribution using the values that were evaluated in the model, or simulation, of interest. Cumulative distributions generated by applying the PLOW method to the data produced by the JC-informed simulations are shown in Figure 2.8.

The distributions determined using PLOW match closely with those determined using only the known outputs. This applies to each variable and model source. The ZA model is not shown here as the JC, JC-U, and JC-L provided enough model sources to analyze the effect of applying PLOW to the data set. In these cases, JC-U (original and re-weighted) again predict more precisely than their counterparts in that their distributions are more narrow than any of the other sets of points. Additionally, the medians of each distribution are not exactly equal for each of the outputs. Usually, JC-U (original and re-weighted) predict lower values, followed by JC (original and re-weighted), and then JC-L (original and re-weighted). This corresponds closely to the prediction order in the stress-strain space among these models as illustrated in Figure 2.5.

Additionally, the points in Figures 2.7 and 2.8 are plotted on a normal probability scale on the ordinate axis. Because the data points form an approximately straight line, it appears that the JC-U data follows a normal distribution for the Final Length, Width, and Temperature outputs. For the JC model, the data appears to follow a normal distribution for the final length, width, and von Mises stress. Despite these few alignments, there was no complete consistency with regard to normality among any of the models and the outputs plotted. Each of the data sets were plotted on log-normal, Rayleigh, and Weibull probability scales in addition to the normal scale. These did not show any major consistency between models and or predicted outputs and thus are not shown here.

The distributions from each simulation output were compared using the Komolgorov-Smirnov (KS) 2-Sample test. Tables 2.4 and 2.5 show the D -statistic for each KS test. This test can be used to determine if the data from two distributions are from the same sample. It is a non-parametric test which removes the requirement that the distributions being compared have to be normal or have equal variance such as in the case of the student's t-test. Another benefit of the test is that the two

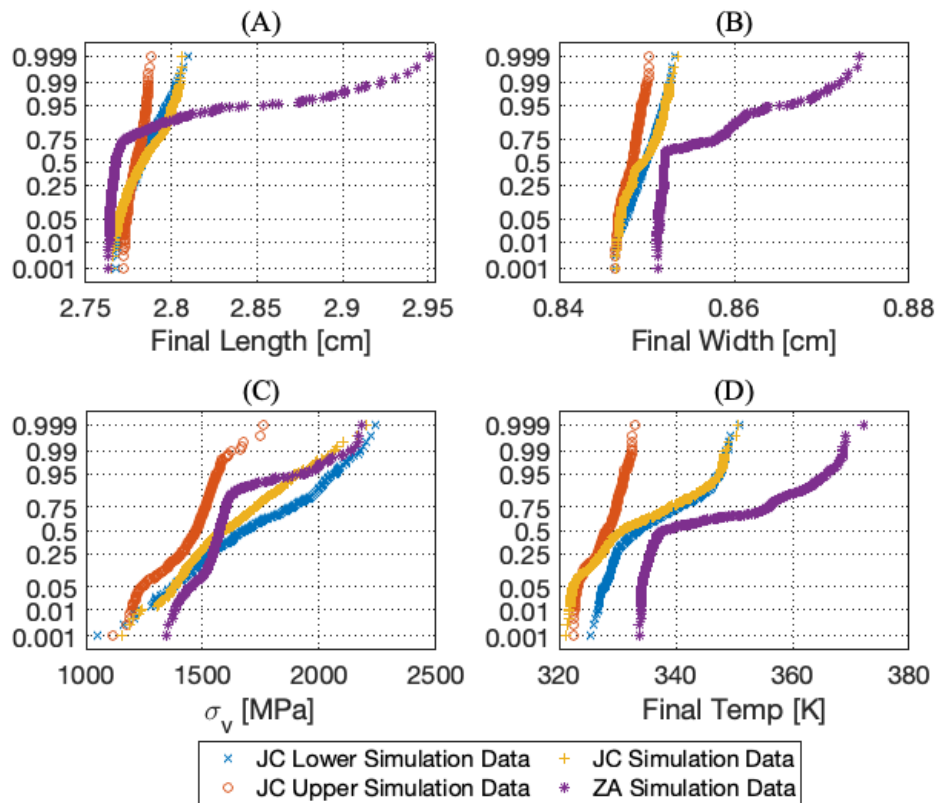


Figure 2.7: Cumulative distributions of outputs from Taylor anvil simulations. In each plot the distributions produced using the Johnson-Cook Lower (JC-L), Johnson-Cook Upper (JC-U), Johnson-Cook alone (JC), and Zerilli-Armstrong alone (ZA) models are shown.

samples do not have to be equal in size. The statistic for the KS 2-Sample test is a distance value D , which is expressed as

$$D_{m,n} = \max(|ECDF_1 - ECDF_2|).$$

The value represents the maximum distance between the two empirical CDF's, $ECDF_1$ and $ECDF_2$.

When the samples being compared are large (roughly greater than 30) and the expression

$$D_{m,n} > c(\alpha) \sqrt{\frac{m+n}{m \times n}}, \quad (2.9)$$

is true, then the null hypothesis can be rejected [36]. Here, m and n refer to the sizes of the first and second samples. The null hypothesis, H_0 , is that the two distributions are from the same sample. In the cases presented here, the significance level $\alpha = 0.5$. This corresponds to a $c(\alpha)$ value of 1.358 according to the KS test statistical significance calculation [36]. For each of the tests, the size of the two samples compared were either 499 or 500 measurements making the expression on the right of Inequality 2.9, equal to 0.085.

After PLOW was applied to all of the output distributions, they again were compared using the KS 2-Sample test. These comparisons consisted of the Reified Upper Bound versus the Lower Bound, the Johnson-Cook model vs the Reified Upper Bound, and the Johnson-Cook vs the Reified Lower Bound. Figure 2.8 shows how the shape of the re-weighted CDFs diverge slightly from the original CDFs of the JC, JC-Upper, and JC-Lower cases. Only JC was done since the JC, JC-Upper, and JC-Lower cases proved the point that if all the data were able to be run, there would still be a statistical difference between the models. Ultimately, most of the 24 tests (8 tests per comparison), were determined to have been sampled from populations with different distributions. This means that the different models, unique only through slight differences in parameter values, caused statistically different results to occur.




Figure 2.8: Cumulative distributions of re-weighted outputs from Taylor anvil simulations with the cumulative distributions of the original outputs. In each plot the distributions produced using the Johnson-Cook Lower (JC-L), Johnson-Cook Upper (JC-U), and Johnson-Cook alone (JC) are shown. Also displayed are the re-weighted CDFs (JC-L Re, JC-U Re, JC Re) which are based on the previously mentioned models.

2.5 Discussion and Future Work

Characterizing the behavior of newly discovered materials using traditional experimental techniques is necessary for bench-marking against other materials. Here, the dynamic response of AF9628 is assessed and two widely used models are calibrated against it. Doing so is certainly not novel, but it is significant, as it allows for prediction of the material's behavior in various engineering systems.

Naturally, there is uncertainty associated with a material's experimental response. Accounting for all sources of uncertainty in an experiment is next to impossible, especially for dynamic experiments on materials where slight microstructural differences influence behavior. Still, the models developed to predict material behavior are still valuable. Employing a framework like the one in the present work to account for uncertainty in models can help reveal a more confident envelope of parameter values for such models. Reification enables this by allowing the investigator to predict a quantity of interest using multiple sources of information about the system simultaneously.

Reification has many benefits, but statistician, George Box's famous quote that "All models are wrong, but some are useful" applies here. The reification mean is designed to predict closer to values where experiments and models agree the most. With this understanding, the implications of the reification mean should be interpreted on a case by case basis. In the present work it serves to bound the response of the material. Researchers interested in using the framework should be mindful of the numerical origin of reification and that what it predicts is not always representative of the mechanisms involved in the physical experiment.

Future work includes combining additional models in the reification process. The addition of constitutive models would have a large influence on the profile of the reification mean. In such a case, the reification mean is not guaranteed to encompass the response of the individual models as it did in the present work. In some instances, the mean profile could be used to represent the stress-strain response in full. In this case, the profile can then be used to define a novel constitutive model that combines the prediction capability of its constituent models.

It is also important to note that the effect of mesh sizing was not investigated in the present

work. The sizing of the mesh in the Taylor anvil simulation could alter the distributions produced by the simulations shown in Figures 2.7 and 2.8. The difference, if any, could be minor on the final length and width predictions, as these were calculated based on the positions of particles that form the boundaries of the specimen. However, other predicted outputs could change on account of increased simulation accuracy and thus, its effect on outputs could be of focus in future studies.

Additionally, the distribution of the parameter values used in the MCMC process could also alter the results. In the present work, a uniform distribution was used for each parameter value in the MCMC process in each case of fitting. As more experimental data on the behavior of AF9628 at high-strain-rates becomes available, it can be used to inform the distribution of model parameter values. This data could enable simulations, that use the models as inputs, to make more accurate predictions.

2.6 Conclusions

- A set of parameter ranges for the Johnson-Cook and Zerilli-Armstrong models were obtained using information fusion for the material AF9628 (Table 2.3).
- A distribution of predicted outputs for AF9628 has been obtained via uncertainty propagation. The distributions were also found to be mostly statistically different between the original, upper, and lower models.
- The re-weighted simulation output distributions showed that if all of the inputs were able to be run, the predicted value distributions would still be mostly statistically different between the JC, JU-U, and JC-L models.

The parameter values presented here represent the best possible ranges, using reification, given the variability of the data and limits of the models to represent reality. When compared to values from literature for Eglin steel [35], the values calculated using the methods stated above were similar in magnitude. This helps validate the results of the calculations since AF9628 is closely related to the Eglin steel alloy.

Understanding the appropriate range of parameter values enables engineers to have more confidence in their designs. Like the Taylor Anvil simulations described here, the information about AF9628 can feed into a number of codes that describe systems outside of high strain rate applications.

Because most simulation outputs generated by the upper and lower parameter values for a given model are statistically different, they can approximately represent the upper and lower bounds of values one would expect the TA simulation to predict for AF9628.

Acknowledgements

The funding for this work was provided by the Minority Leaders - Research Collaboration Program at the Air Force Research Laboratory's Wright-Patterson Air Force Base in Ohio. Split Hopkinson Pressure Bar Experiments were performed by the University of Dayton Research Institute (UDRI). A special thank you to Pejman Honarmandi for help in verifying Markov Chain calculations and the propagation of uncertainty.

Data Sharing

The data that support the findings of this study are available from the Air Force Research Laboratory (AFRL). Restrictions apply to the availability of these data, which were used under license for this study. Data are available from the authors upon reasonable request and with the permission of AFRL.

3. OPTIMIZED UNCERTAINTY PROPAGATION ACROSS HIGH FIDELITY TAYLOR ANVIL SIMULATION

3.1 Abstract

In computational materials research, and more specifically uncertainty analysis, there are often situations where the cost of running all available combinations of inputs becomes impractical. At this point, investigators must sacrifice output distribution accuracy by down-selecting from the inputs for computational expediency. Recently, a method was developed which can change the output probability distribution by applying weights and using information directly from the proposal and target probability measures. In this work, the method is applied to a dynamic simulation that serves as a link between material structure and properties. It is demonstrated in two contexts: randomly down-selecting inputs and “intelligently” down-selecting inputs via k-medoids. A simplified constitutive model example is also provided. It is shown that applying optimized weights on a sub-sample is able to more accurately reveal the target distribution in both the random selection and k-medoid selection cases.

3.2 Introduction

The central paradigm in materials science rests on the existence of process-structure-property-performance (PSPP) relationships, which constitute the foundation to materials design[37]. When experiments are too costly to run, simulation-based approaches can be used, although in general linkages along PSPP chains are developed via experiments, or through a combination of experiments and simulations, mediated through data analytics techniques [38, 39].

Even in simulation-focused approaches, a limited number of experiments are necessary in order to properly parameterize the (hopefully predictive) models. In the case of materials behavior under quasi-static loading conditions, uncertainty tends to arise from the stochastic nature of materials microstructure. This uncertainty is somewhat controllable and bounded by the stochastic nature of the microstructure of the material being investigated. In the case of polycrystalline metals, the

spread in the mechanical response when comparing separate experimental runs is not significant, provided the materials are nominally the same—i.e. they have very similar distribution in chemical and microstructural attributes.

Under high-strain rate conditions, however, extreme sensitivity of experimental conditions can lead to dramatic variance in the observed materials response. Such uncertainties can be quite significant, even when comparing the response of nominally identical materials under virtually indistinguishable experimental runs making the prediction of properties and performance of metals subject to high strain rate deformation difficult. Thus, constitutive equations that are designed to predict a material's stress, given a strain, are less reliable for prediction [28], [20], [40]. To account for uncertainties in experiments, techniques such as Markov Chain Monte Carlo (MCMC) can be used to sample the parameter space in the constitutive equations based on a small set of experimental data [41].

In the context of materials design, uncertainty quantification (UQ) and uncertainty propagation (UP) play important roles [42, 43, 44]. Under nearly every experimental condition, uncertainty in the measured value can stem from slight variations in material (micro)structure (aleatory) and/or the physics associated with a ground truth realization that are not accounted for in a model (epistemic). Such uncertainties are difficult or impossible to account for directly [45, 46, 47, 48], and are the subject of UQ. We note that these sources of uncertainty exist not only in physical experimentation, but also in computational experiments. In either case, it is important to account for these variations in measured results so that they can be propagated accurately and proper calibration and validation and verification (V&V) of the models is carried out.

Once this information is gathered, it can be used to define distributions of probable values for the parameters of a model that predict a quantity of interest (QOI). Subsequently, the model can be run multiple times with different parameter values taken from these distributions. This creates a distribution on the output which can serve as an estimate of material performance. Such a process is referred to as uncertainty propagation [49, 50, 51]. Ultimately, data on performance ranges for a material's behavior can determine an intended application and also provide a factor of safety for

a minimum or maximum performance expectation. Furthermore, when high fidelity simulations are available, this information can be determined before a part is ever manufactured.

3.2.1 The cost of sampling high fidelity simulations

More often than not, UP is carried out using Monte Carlo integration or the Markov Chain Monte Carlo approach [41], which can generate parameter values for models based on any experimental realizations available to the investigator. Other approaches include local expansion-, functional expansion-, and numerical integration-based methods [15]. Depending on the computational resources and time one has available, the result of this technique can be hundreds of thousands or even more than a million different parameter sets for the equation under investigation. Returning to the general problem of simulating the response of a material under high strain rates, each parameter set can be used to inform a different hydrocode simulation to predict how a material will perform in certain loading conditions and what properties it is expected to have. Such loading configurations include the Taylor-Anvil and spall plate impact configurations, and these can be modeled with a suitable hydrocode. The present work investigates the Taylor-Anvil configuration (termed CTH code in this work) as a validation of the constitutive models fit to intermediate strain rate ($\sim 10^3 \text{ s}^{-1}$) tensile tests.

High fidelity hydrocode simulations accounting for material response may be computationally expensive, and therefore we can only probe a limited number of models from an available set. This makes the use of MCMC-based approaches that typically need to sample the parameter space hundreds of thousands or millions of times impractical. Rather than carrying out an extensive parameter sampling, a solution would be instead to extract sub-samples out of the total number of observations that can be drawn from the parameter distribution space. While this sub-sampled distribution does reduce computation time, it also can misrepresent the output distribution, depending on whether the sub-sample is weighted and method used. If the sub-sampled data does not reflect the ground truth parameter distribution, there cannot be proper propagation of uncertainty. Different methods of sampling exist, such as K-medoids [52], Latin-Hypercube [53], or simply random selection—we note that such misrepresentation will occur with any method of sampling that does

not properly account for the importance of a given sample relative to the full, ground-truth distribution.

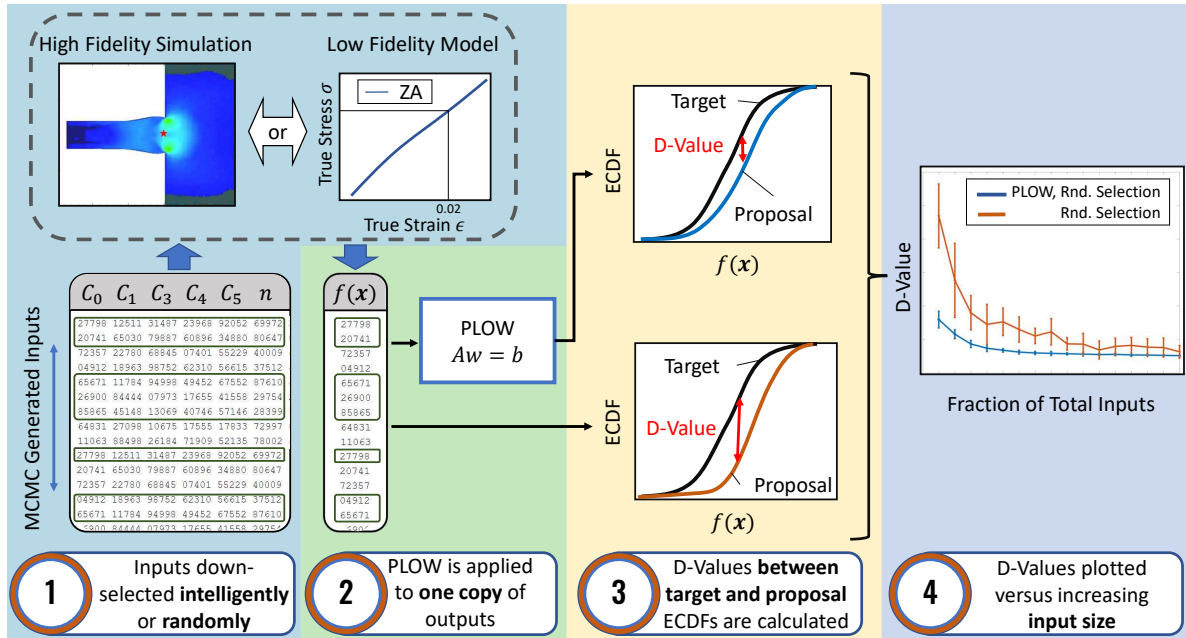


Figure 3.1: Process used to analyze the effect of the PLOW method on the high fidelity Taylor-Anvil simulation and at scale using the Zerilli-Armstrong model.

The sheer volume of possible values to try can pose a problem if the subsequent simulations that take the values as inputs are expensive to run in time or computational costs. In such situations, an investigator must down-select values from the UP generated pool to run in the simulation and obtain the output distribution. When this is done, the distribution that would have materialized if all of the inputs were run is not observed. The distribution obtained from the down-selected inputs will be referred to as the proposal distribution and the distribution that comes from running all of the inputs will be referred to as the target distribution. Depending on the number of down-selected inputs, the proposal distribution can deviate substantially from the target distribution.

Noting the challenge in appropriately down-selecting UP-generated input values, the present effort investigates the effect of applying the probability law optimized importance weights (PLOW)

method across a simulation chain to improve the fidelity of the down-selected sub-samples.¹ The procedure has been shown to work well in models of up to 10 dimensions [15]. However it has not yet been applied across a high-fidelity simulation that is often used for material property estimation. The shape of the distributions of the sub-samples and the full data set will be compared using the Kolmogorov-Smirnov (KS) statistic, or “d-value” [55] Specifically, the predicted density of a material after a simulated Taylor-Anvil test is studied. Additionally, a toy problem will be used to compare the effects of the PLOW re-weighting method on a larger data set than can be captured in the hydrocode given the resource allotment for this work. It is expected that the KS statistic, used to compare the shapes of the sub-sample and target distributions, will be lower for the PLOW sub-samples than for non-PLOW sub-samples, thus providing a higher fidelity sampling of the UP-generated input values for the simulations. Figure 3.1 shows an overview of the steps taken in the present work to perform the analysis.

3.3 Methods

3.3.1 Experimental setup

The present work is focused on predicting the high-strain-rate response on a low alloy, high performance steel AF9628 [56, 57, 58], which shows a similar mechanical response to a related armor steel – Eglin Steel [59]. To obtain the intermediate rate stress-strain response for AF9628, five Split-Hopkinson Pressure Bar (SHPB) [21] experiments were performed using an indirect tension configuration [60] at a temperature of 300°C, performed with identical input conditions targeting a strain rate of 1500 s⁻¹. The AF9628 material was heat treated from the as-received condition using the patented heat treatment schedule [61, 62, 63]. Five subscale E8 tension specimens of AF9628 were carefully machined with tight tolerances, and a low stress grinding (LSG) process was used to reduce the sample geometry and surface finish variability across all specimens. The five tests were controlled, as much as possible, for all input variables including material composition, microstructure, geometry, and impact bar velocity, among others. This was to ensure that any

¹The name is partially inspired by the soil-themed Earth Mover’s Distance metric analogy, which is a measure of distance between two distributions [54].

variability in the stress-strain response came from aleatoric sources in the experimental apparatus and not from neglecting known effects.

The resulting measured strain rates for the 5 experiments were 1440, 1400, 1430, 1350, and 1440 s⁻¹. The stress-strain data was fit to a Zerilli-Armstrong (ZA) constitutive model using a minimization routine that combined sequential quadratic programming and MCMC, assuming an average strain rate of 1412 s⁻¹ for the parameterization (MCMC was implemented via the MCMC toolbox for Matlab package [64]).

3.3.2 Obtaining Target Data Set Values

This paper focuses on the Body Centered Cubic (BCC) ZA model shown in Equation 3.1.

$$\sigma = C_0 + C_1 e^{-C_3 T + C_4 T \ln \dot{\epsilon}} + C_5 \dot{\epsilon}^n \quad (3.1)$$

Here, C_1 , C_3 , C_4 , C_5 , and n are material constants determined through fitting to experimental data. The C_0 constant accounts for the stress contribution due to grain size. If grain size information is unavailable, this term can be set to the yield strength of the material. This is the case for the current work so it is used in the starting set of parameter values for the MCMC process. T , ϵ_p , and $\dot{\epsilon}_p$ are the experiment temperature, equivalent plastic strain, and plastic strain rate respectively.

Table 3.1 shows the ZA parameter values obtained from minimization on the experimental data. These values are used as the starting point for MCMC sampling. Once the parameters were determined, MCMC was performed in order to obtain a distribution of possible parameter values for the model. This method allowed the 6-dimensional parameter space to be explored more completely, giving the model a chance to converge to values that better fit the data set. This also allowed for the uncertainty of the parameters to be captured.

3.3.3 Hydrocode simulations to obtain target distribution of material density

High fidelity impact simulations are needed to exercise the ZA constitutive model and obtain the target output distributions of relevant performance variables. The present work investigates the Taylor Anvil impact configuration, wherein a cylindrical rod composed of the material under study

Parameter	C_0 [MPa]	C_1 [MPa]	C_3 [K ⁻¹]	C_4 [MPa]	C_5	n
Value	1332.5	7.128E3	0.6626	0.0882	499.99	0.5309

Table 3.1: Zerilli-Armstrong constants obtained using gradient based minimization against experimental data.

impacts a rigid anvil. The impact conditions create a one-dimensional uniaxial stress loading condition, where stress waves propagate through the rod and their interactions with the rod boundaries determine the final deformed state of the rod. The unconfined nature of the deformation leads to triaxial states of stress at the extremes of the deformed rod, which makes the Taylor Anvil test a useful validation tool to test and exercise constitutive models [65]. The Taylor Anvil test was simulated in the multi-material Eulerian hydrocode CTH (Version 12.0, Sandia National Laboratory) [66], employing the ZA constants obtained from MCMC sampling and holding all other input conditions constant.

The Taylor Anvil impact simulation models an AF9628 rod impacting a Vascomax C300 anvil, as shown in Figure 3.4 where the pressure was output at the final frame of the simulation. Relevant time-histories of thermodynamic and mechanical field variables (e.g. Pressure, stress, temperature, density, etc.) can be obtained from each CTH simulation over the length of the AF9628 Taylor rod impacting a Vascomax C300 anvil. The center point at the impact face of the rod was selected for monitoring all field variable data. The AF9628 Taylor rod was modeled with a Mie-Grüneisen equation of state (EOS) with parameters obtained from Neel et al. [56], and the anvil was modeled with a Mie-Grüneisen EOS with parameters obtained from the Los Alamos Shock Compendium [67]. The simulations were conducted at an impact velocity of 200 m/s and ran for 20 μ s. An automatic mesh refinement (AMR) algorithm was employed with a final mesh size of 100 μ m at the smallest element. Fracture was not modeled in the present work.

Simulations were run on a 48 core high performance computing system with Intel Xeon Platinum 8168 cores. Exact wall time was not recorded per simulation, however the estimated run time for a single simulation was 15 minutes. Even with run time scaling linearly with total number of

simulations, the run time can quickly become impractical for larger numbers of runs. For instance, at just 500 simulations, the total wall time can exceed 125 hours.

It is often necessary to propagate the uncertainty when performing simulations of an experiment to account for different sources of error. MCMC was performed for the ZA parameter fitting for one million iterations, enough to assume adequate convergence. The result of the MCMC fitting is a matrix of values with one million rows and six columns. Each row is an estimated value for the six parameters in the ZA model shown in Equation 3.1. In the current work, the word “points” is used to describe a single six column row in the matrix. Each one of these points is a potential input to the TA simulation via the ZA model. Afterwards, an inferred burn-in region of 200,000 points was removed from this data set, leaving 800,000 possible parameters for each constant. Once the parameter values were obtained, they were used to inform a CTH simulation of a Taylor Anvil test to obtain the output deformation distributions for each ZA iteration. Due to limits on computational time, only a subset of the 800,000 points generated could be run, since a separate simulation would need to be run for each parameter set and each simulation is computationally expensive. 500 points was chosen as a reasonable number of simulations that could fit within the high performance computing time allotment constraints for this work.

K-medoids, a method for sub-sampling data, was used to reduce the data set to 500 parameter sets. Essentially, the K-medoids algorithm first partitions the data set into k different groups. A single value from each partition is then selected to represent the group of points in that partition. Furthermore, the algorithm ensures that the selected value is within the original data set. K-medoids works even if each data point is multi-dimensional. In the case of this work, the total 800,000 value data set was partitioned into 500 groups. A set of ZA parameter values was then selected from each group, and this set of 500 data points was then used to inform 500 different Taylor Anvil simulations. Notably, one simulation did not finish, leaving 499 simulations to analyze. Although it is unclear what precisely caused the run not to finish, the most likely reason could be that the simulation simply didn’t finish within the HPC hours allotted to this project.

Properties that can be extracted from this simulation include density, particle (material) veloc-

ity, final specimen length, and final specimen width, among many others. With the 499 tests, a distribution of the predictions for a given property could be plotted. For this work, the density of the material was chosen as the single output due to its relatively low sensitivity to changes in constitutive parameters. This will permit a good assessment of the down-select methodology employed in the present work. Although 499 points were plotted, this sub-sample could be showing an untrue representation of the form of the entire distribution of values. This is especially true due to the fact that 499 points is a very small subset compared to the original 800,000 point dataset. It will be shown below that PLOW can be used to uncover a truer form of the distribution from the 800,000 point dataset. This is shown to perform even better than using a method like K-medoids to intelligently select the data.

3.3.4 Probability Law Optimized Importance Weights

3.3.4.1 Background

The Probability Law Optimized importance Weighting (PLOW) approach, first proposed by Sanghvi et al. [15], employs a least squares formulation to determine a set of empirical importance weights to achieve a change of probability measure. The notion of using importance weights to simulate the moving of a probability distribution has its roots in optimal transport and the concept of the Earth mover’s or Wasserstein metric [68]. The objective of the PLOW approach is to estimate statistics from a target distribution of interest, by using random samples generated from a different proposal distribution for which samples are more easily attained. The approach taken here works directly with the probability law of the proposal and target random variables (which may be vector-valued), from which only samples from each are needed. For completeness, we note that the probability law of a random variable x is a mapping, $\mathcal{P} : \mathbb{B} \rightarrow [0, 1]$, where \mathbb{B} is the Borel σ -algebra over the domain of the random variable $x \in \mathbb{R}$. The probability law allows us to consider the probabilities of different sets over the domain of a random variable (or vector) rather than on cumulative probability up to a point as done in previous work (see, e.g., Ref. [69]). The result is an approach more capable of achieving high dimensional probability measure change

than current state-of-the-art methods. Such a method can enable efficient and accurate propagation of uncertainty through model chains of unknown input and output regularity, such as those often encountered in PSPP linkages in materials science.

3.3.4.2 Approach

The PLOW approach described here enables re-weighting to be applied across model chains [51, 47, 70, 71]. Model chains refer to frameworks that employ multiple models sequentially where the outputs of one model become the inputs of the next model. PLOW utilizes decomposition-based approaches to calculate weights, which allows it to be used across model chains without losing accuracy in the target distribution prediction. Because of this, investigators can use UP more reliably. The approach is based on the change of probability measure. The general idea is that for the propagation of uncertainty, there is a desired input distribution we wish to propagate through a model. We refer to this distribution as the *target distribution* [41]. In many cases however, the target distribution of an upstream model is not known, or was sampled without knowing the underlying distribution of the model’s parameters, which can negatively affect the following model. PLOW can be used to estimate the target distribution of such upstream models, given an already existing set of function evaluations composing a *proposal distribution*.

To obtain the target distribution, it is common to use traditional importance re-weighting methods such as density ratio estimation (e.g., importance sampling), typically determined using kernel density estimation [72] or in approximate form using approaches based on calculating an L_2 -norm [69]. With PLOW, the importance weights are calculated using only the samples of the proposal distribution or (if available) target distribution information. This approach enables the user to determine the regions in the input space over which to calculate importance weights. This leads to excellent scalability of the PLOW method in comparison to current re-weighting methods that are subject to numerical ill-conditioning at high dimensions.

Importance weights using PLOW are computed as the solution to a linear system of equations that can be written as,

$$Aw = \mathbf{b}, \tag{3.2}$$

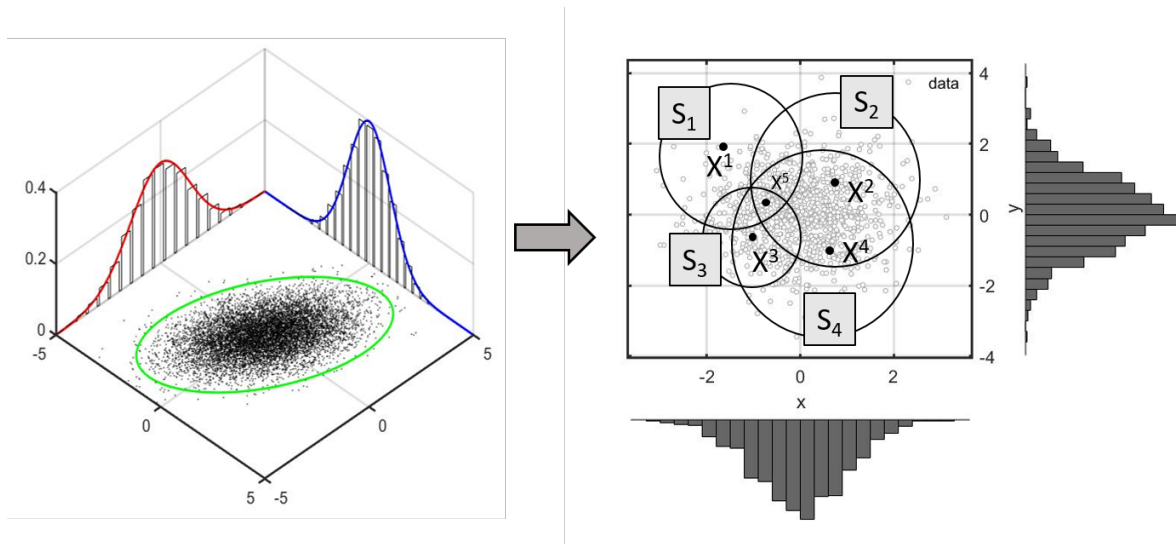


Figure 3.2: Illustration of the process to construct a measure matrix, like A in Equation 3.2. On the left, a bivariate distribution space is displayed corresponding to a model with two parameters for simplicity. In the case of inputs, the dots on the surface represent the co-variance of points the investigator is able to run. The borders show each parameters individual distribution. On the right, circles S_j (hyper-spheres in higher dimensions) are drawn over the space to populate the measure matrix.

The construction of the A , w , and b is described in the following paragraphs. Once constructed, the solution of Equation 3.2 can be calculated via least squares using the normal equations noted below.

The matrix A and the vector b of Equation 3.2 are determined using specified regions within the domain of the random variables whose measure we seek to change. Here, A represents the presence, or absence, of each measured input sample point with respect to each hypersphere. The vector b represents the joint probability of points contained within a specific hypersphere. A graphical depiction of this region selection procedure is shown in Figure 3.2. A distribution space that represents the input values (black dots) available for each parameter of the model in question is shown in the left portion of the figure. In this case, the model would have two input parameters, represented by the two axes making up the base of the graph. In this example, the points are normally distributed for both parameters, although this is generally not known or necessary prior to applying PLOW. This illustration can also represent the distribution of points that make up the

proposal distribution, or model output space. The difference being that there would be less points on the proposal distribution space due to the experimenter's inability to run all input conditions through their model. For simplicity of describing the generation of the matrices only four points and four regions are considered.

The regions, S_1, S_2, S_3, S_4 shown on the right side of Figure 3.2 are deterministically-generated to capture every point in the space. In the present case, the regions are in the shape of circles. However, in higher dimensions, they would be hyperspheres. Their size and location are determined based on the samples in the space alone. The process of generating these regions involves growing a hypersphere around each proposal point so that all points are within at least one region. This ensures that the process is scalable with respect to problem dimensionality, which is essential for higher dimensional problems where hyperspheres do a poor job of space filling. We note that because each point has a hypersphere associated with it, we ensure each point is represented in the matrix A . In very high dimensional problems (e.g., $\mathcal{O}(100)$ dimensions), the use of hyperspheres will lead to large regions of the input space not being represented in A . It is a topic of future work to explore the use of region shapes better suited to the concentrations of measure that occur in these high dimensional cases. In our approach, we grow the regions until the union of the regions covers the entire space. In practice, this typically yields a matrix A with dependent columns, leading to a non-unique solution to the least squares problem. To recover dependent columns, we identify non-unique rows in the matrix A , which identify those points that are currently indistinguishable from one another. For these rows, we reduce the radius of the respective hyperspheres until the rows are unique. We then add a region that covers the entire space to ensure the weights are associated with a probability measures (that is, the probability of the entire input space is still unity). This approach leads to independent columns and the matrix A has more rows than columns. Thus, the normal equations of Equation 3.2 shown below, have a unique solution. It is a topic of future work to explore that non-unique solution case that can occur when sample points are nearby one-another. It is expected that this could lead to a capability to delete input samples and potentially lead to enhanced computational efficiency.

In the example of Figure 3.2, for clarity, there are five proposal points shown as X^i , where $i \in \{1, 2, \dots, 5\}$. The hypersphere, S_1 , contains only X^1 among the proposal samples. This region, S_1 , and its proposal sample contents are captured as a row vector within a matrix A_{ij} by populating a 1 in entry $[j, i]$ if proposal sample i is in hypersphere j (in this case, $j \in \{1, 2, \dots, 4\}$). Therefore, we have entry $A_{11} = 1$, since sample 1 is in region 1. The remaining random sample points, X^2 , X^3 , X^4 , and X^5 are equal to 0 for S_1 . This gives $A_1 = [1, 0, 0, 0]$ for the first row of the A matrix.

Filling in data according to Figure 3.2 for the remaining points and hyperspheres, the following matrix would be generated:

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 1 \end{bmatrix}$$

The \mathbf{b} vector is created to contain the probability of a random sample landing in region S_j under the target distribution as the j th element of the vector. This can be computed via target samples or via a known target distribution with traditional integration. For the example, the \mathbf{b} vector is $\mathbf{b} = [\mathbb{P}(S_1), \mathbb{P}(S_2), \mathbb{P}(S_3), \mathbb{P}(S_4)]^T$, where $\mathbb{P}(S_j)$ is the probability of a target sample being in the region S_j . Once A and \mathbf{b} are created, the weights, \mathbf{w} , can be determined from the normal equations of Eq. 3.2, which can be used to find the weights that produce the least squares solution. Specifically, we compute the weights as $\mathbf{w} \leftarrow (A^T A)^{-1} A^T \mathbf{b}$, which can be solved via elimination. Further details on this process can be found in Ref. [15].

3.3.5 Description of the 'Toy Problem'

To apply the function to the entire data set of 800,000 points, a less computationally expensive model than a full Taylor-Anvil simulation in CTH was needed. The toy model was selected to be the Zerilli-Armstrong stress-strain function itself $\sigma = \sigma(\epsilon_p, \dot{\epsilon}_p, T)$. The stress at an equivalent plastic strain of 0.02 was chosen as the output. The 1 million ZA model evaluations, from which

the 800,000 points originate, were performed using the MCMC toolbox for MATLAB package, allowing the output of each evaluation to be saved. On a quad-core workstation using Intel i7-7700 cores, the run time for the full 1 million evaluations was roughly 35 minutes. The simplicity of the model makes it significantly less time consuming to run than the TA simulations. The process for generating the distribution is as follows:

1. The 6 parameter values in one set, of the possible 800,000 parameterizations, were substituted into ZA model equation 3.1
2. A value of 0.02 was input as the strain value
3. The corresponding stress was obtained and saved
4. Steps 1-3 are repeated 800,000 times for each parameter set
5. The resulting values are binned and plotted as a distribution.

The distribution of 800,000 points would serve as the known, true distribution because it includes all of the data points. Once this distribution was obtained, its shape would be compared to that of varying subset sizes of this data set. An issue arose however, in using this data set in that a large number of values were identical. To obtain a target data set of unique values, the replicate data was removed leaving roughly 200,000 parameter sets. Furthermore, after applying the PLOW method on an increasing number of sub-sample distributions, there came a point where increasing the number of values in the sub-sample resulted in a negligible change in shape compared to the target distribution. At this point, the sub-sample and the target distribution shapes were essentially the same. The *threshold* sub-sample size of 15,000 points was therefore selected for the target distribution size. Consequently, in order to show the convergence of the d-value and to compare with other methods, the proposal sub-sample sizes needed to be altered.

Additionally, issues arose with using the k-medoids algorithm to select roughly 500 points from a data set. Therefore, for the d-value comparison of the Toy Problem, a random selection method of the sub samples was used for comparison to the 15,000 target sample size. For this

case, the sub-sample sizes began at 500 points, and incremented to 1000 by 500. At this point the sub-sample increment was changed to 1000 and continued to increase to the target sample size of 15,000 points.

In order to compare the performance of PLOW on K-medoids for the toy problem, a target data set of 500 was selected from the full 15,000 points. The sub-sample sizes here started at 50 and incremented by 50 up to 450. From here, the increment was changed to 10. This continued until the subsample size reached 499.

The 2-sample Kolmogorov-Smirnov (KS) test was used to compare the shape of the target data sets to the smaller subset distributions. This test was used because it does not require the distributions to have a certain shape or standard deviation, i.e., it is a non-parametric test. For this test, the distributions to be compared are first converted to empirical Cumulative Distribution Functions (eCDFs). The vertical distance between the two eCDFs are measured across each point on the horizontal axis and the maximum value among these is selected. An example of this value is shown in Figure 3.3. Normally, this value becomes the KS-Statistic which is tested against the null hypothesis that the two models are sampled from the same distribution.

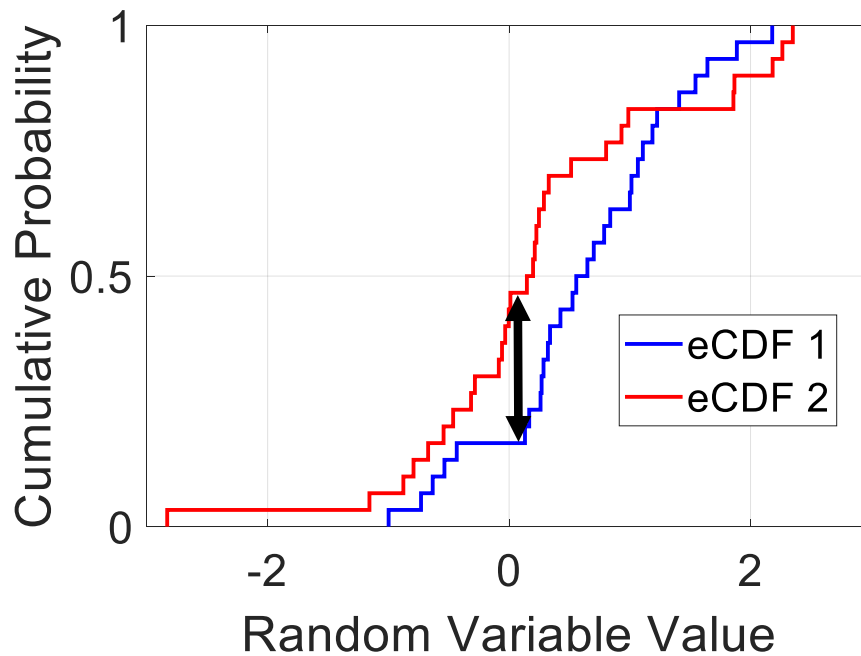


Figure 3.3: Illustration of the Kolmogorov–Smirnov statistic. The red and blue lines are the two CDFs being compared, while the black arrow is a measurement of the largest distance between the two distributions. It is commonly known as the K-S statistic, but is also referred to as the D-value (distance value) in the present work.

For this work, however, the KS-Statistic [73] is used as a measure of convergence of the weighted proposal distribution to the target distribution. That is, the smaller the KS-Statistic, the closer the distribution in question (our weighted proposal distribution) is to the distribution being tested against (the target distribution). This follows from the Glivenko-Cantelli theorem, which demonstrates that the KS-Statistic converges to zero with probability one when the distributions are the same and the number of samples tends to infinity. Thus, as the KS-statistic reduces, the two distributions in question are closer (in distribution), hence the use of the KS-Statistic as a convergence measure.

3.4 Results

3.4.1 Target material density distribution generated from Taylor-Anvil simulations

The CTH simulations of the Taylor Anvil impact test provided field variable histories throughout the AF9628 rod. The equilibrium density was calculated and noted to converge to a stable value by the end of the simulation. Only the input parameters for the ZA constitutive model were varied for each run and all other parameters noted in section 3.3.3. The predicted field variables are sensitive to the constitutive model input parameters to different degrees. Thus, each unique set of parameter values obtained using MCMC for the ZA equation is expected to produce different outcomes in pressure, temperature, and every other thermodynamic and mechanical property. The proposed technique can be applied to each property to propagate uncertainty for future model chains.

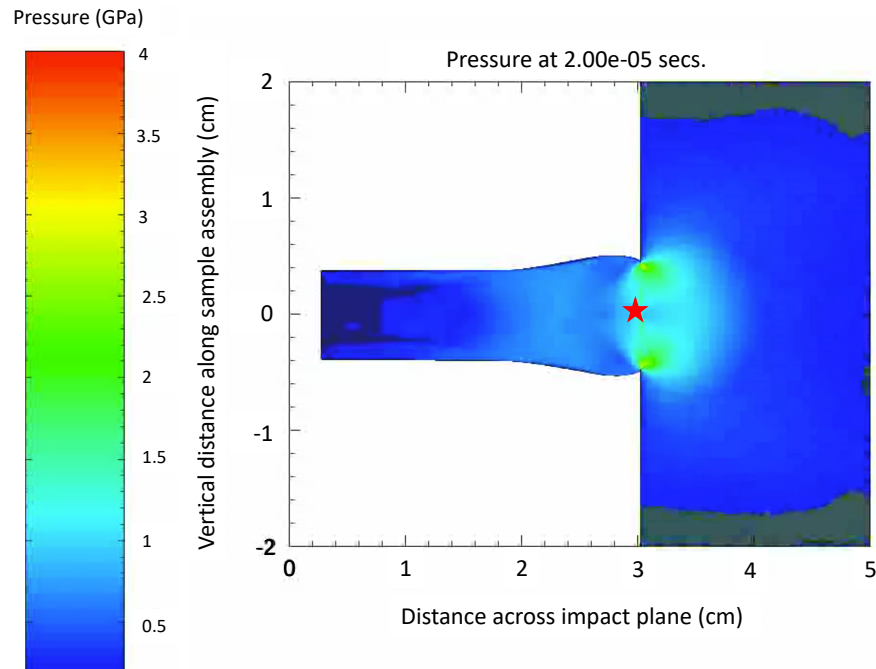


Figure 3.4: Last frame of Taylor-Anvil simulation movie created using hydro-code software. The pressure throughout the specimen and anvil are represented by color. The simulation was informed by one ZA parameter set of the available 499.

For the current work, the density of the specimen at a tracer located at the center of the surface of impact was used as the property of interest. Once the 499 Taylor Anvil simulations were run, a distribution of the density calculated from each simulation could be plotted. Figure 3.5 displays the distribution of all 499 of the density values in the form of a solid, dark blue line. Also each of the 499 simulations were informed by a unique set of ZA parameter values.




Figure 3.5: Cumulative distribution of specimen density values predicted by the Taylor Anvil simulation. Each line, except one, represents a proposal distribution determined by an input sub-sample size as indicated in the legend. One line represents the target distribution. Naturally, as the size of sub-sampled input points used to construct the proposal distribution approaches 500, the closer the proposal distribution will be to the target distribution.

From this dataset of density values, we can select subsets of values and apply the PLOW method to the subset to compare to other distribution subsets. Figure 3.5 shows how the subset distributions of density values begin to match the full set as the number of values in the subset approaches the total dataset size of 499 values. Although the convergence with increasing

sub-sample size is an intended outcome, it is not easily seen when viewed as CDF distributions. Therefore, comparing the distributions using a different method was necessary.

Additionally, the distributions are not symmetric across the mean or mode. Also, there are also many values predicted below the mode that cause the graph to be left-skewed. Comparing such a distribution to another of a similar shape would require a non-parametric testing method, as mentioned previously. To do this comparison, the testing parameter of the Komolgorov Smirnov test (KS-Test) was used [74].

As shown in Figure 3.5, sub-samples of the original 499 sample distribution were taken in sizes of 100, 150, 200, 250, 300, 350, 400, 450, 460, 470, 480, and 490. At each sub-sample size, the K-medoid distribution either had PLOW applied or was left unmodified. The two were then compared to the known 499 sample distribution. This was repeated 100 times for each sub-sample size. This allowed a mean and standard deviation to be calculated for each. Figure 3.6 shows the results of the KS-Statistic versus sample size calculations. The sub-samples that did not have the PLOW method applied to them were labeled *Unweighted*, while the samples that did were labeled *Weighted*.




Figure 3.6: D-value versus sub-sample size trend for TA simulations. The sub-sample size represents the number of input points used to construct the proposal distribution. At each X-value, 100 2-SampleK-S statistic values (D-values) were calculated between the target distribution and a proposal distribution constructed using a different drawing of points from the pool of 499 total input points. The average and standard deviation of the 100 D-values is plotted.

The convergence of the KS-statistic with increasing sample size shows how the shapes of the distributions are becoming more and more similar until they are almost exact. Of course, this convergence occurs even when using the k-medoids (or unweighted) method. However, when PLOW is applied to the proposal distribution, the difference between the proposal and target distributions is lower for all sub-sample sizes. That is, the distribution tends to converge faster towards the ground truth distribution.

3.4.2 Toy Problem Results

After applying the PLOW method proved effective in more accurately representing the relatively small target distribution of 499 values, it was of high interest to apply the method to a much larger target distribution. Doing so could test for any biases that may arise due to the size of the

target data set.




Figure 3.7: D-value versus sub-sample size trend for toy problem. A comparison between a random sub-sample selection with PLOW applied after the selection (Weighted) and one without PLOW applied (Unweighted) is shown. Like other D-value versus proposal sample size plots in this work, the average and standard deviation of 100 D-values is plotted at each sub-sample size value.

As previously mentioned, the ZA constitutive equation was directly evaluated at a plastic strain value of 0.02 varying sample input parameters and serves as the toy problem to compare with the CTH simulations. Unlike with the CTH model, more parameter sets from MCMC could be used and tested because querying the ZA model is many times cheaper than running a CTH simulation. Additionally, the same methods used to obtain parameter set values for the simulation tests were used for the toy problem. In this case however, many of the 800,000 parameters obtained using MCMC were duplicate value sets.

Again, once the duplicate values were removed from the total dataset, we were left with 15,000 unique values. The d-value for the subset revealed a virtually identical distribution when com-

paring a subset distribution size of 200,000 with a 15,000 subset size. Figure 3.7 shows how the d-value evolves with sub-sample size for the toy problem. Figure 3.8 shows the same plot along with k-medoid sampled data. The weighted k-medoid sampled data shows a better convergence than the random sampled inputs.




Figure 3.8: D-value versus sub-sample size trends for the toy problem. A comparison of the convergence to the target distribution between random sub-sample selection with PLOW applied (Weighted-Random), random sub-sample selection without PLOW (Unweighted-Random), k-medoids sub-sample selection with PLOW applied (Weighted - Kmedoids), and k-medoids sub-sample selection without PLOW (Unweighted - Kmedoids) are shown.

3.5 Discussion

Overall, the method shows promise in fine-tuning distributions generated using a sub-selection of inputs and outputs, even if the model to determine the outputs is complex with no explicit function to map inputs to outputs.

In every case of the d-value vs sub-sample size plots, applying the PLOW method re-weights the outputs such that the distribution is closer in shape to the target distribution than without PLOW

applied, on average. Additionally, in both the CTH simulation and toy problem cases, the biggest difference between the two methods occurs in the smallest sub-sample size. This is where applying the PLOW method has the biggest impact on altering the distribution shape. The results also show that PLOW works across the CTH simulation, which is a more complex model than prior applications of the PLOW method.

The plots of the random and k-medoid selection methods with and without PLOW applied (Figure 3.6) show that increased accuracy can be achieved even after intelligently down-selecting inputs that best represent the population of available inputs. Furthermore, for nearly every sub-sample size, re-weighting the distribution changes its shape towards that of the target distribution enough that the standard deviations associated with re-weighting and not re-weighting do not overlap. For example, at a sub-sample size fraction of 0.5, the vertical bars representing standard deviation for the Weighted - Medoid and Weighted - Random d-values are between roughly 0.01 and 0.015. On the other hand, the vertical bars for the Unweighted - Medoid and Unweighted - Random are between 0.02 and 0.04.

Another takeaway is that when k-medoids is used alone, it has the ability to select input points that bring the output distribution shape closer to the known, target distribution reasonably well compared to a purely random selection. This can be observed by comparing the mean d-value of the blue line (Unweighted-Medoid) to that of the yellow line (Unweighted-Random) for sub-sample fraction sizes of 0.3 through 0.5, and 0.7 to 0.8.

Such gains in accuracy are very small between the random vs medoid selection processes when PLOW has been applied. The average d-values nearly overlap for all sub-sample fraction sizes between the two lines in Figure 3.6 (purple vs orange lines). The result signifies that PLOW has the ability to find a similar "optimum" shape regardless of the sub-sampling process used.

In the case of the toy problem with only the random sub-sample selection shown (Figure 3.7), the PLOW method creates proposal distributions similar in shape to the known, target distribution, even in small sub-sample sizes. For example, even when just 3% of the total data set is used as a sub-sample, the average d-value is roughly 0.017 for the weighted data, which is much lower than

the average d-value of about 0.046 of the unweighted data set. Keep in mind that the lower the d-value, the closer the sub-sampled distribution is to the known output. The standard deviations (vertical lines at each sub-sample size fraction) are also very narrow compared to the unweighted case, highlighting the precision of the PLOW method.

For both the toy problem and CTH cases, the mean d-value decreases almost exponentially with constant increases in sub-sample size. Such a trend can be used to define an optimum amount of sub-samples to run when computational resources are limited. For example, if a decrease in the average d-value is less than a pre-determined range with an incremental increase in sub-sample size, the final sub-sample size would be optimal. It is unclear whether this same trend remains at smaller sample sizes. For instance, if a trend turns out to be linear when the target distribution is less than 499 samples, the optimal fraction of sub-samples to run would be the most that the investigator could afford. Figure 3.7 shows that the exponentially decreasing behavior holds for target distributions of up to 150,000 points, given the ZA model. Generally, however, the smaller the proposal distribution size, the greater the difference PLOW makes on the distribution.

Additionally, as shown in Figures 3.6, 3.7, and 3.8, as the sub-sample size increases, the difference in mean d-value between the weighted and unweighted distributions approach zero. This convergence proves that the target distribution is the same for both the unweighted and weighted data sets. Overall, the weighted data set consistently performs better than the unweighted one in essentially every ratio value. The only exceptions are the few ratios where the standard deviation of the unweighted plot reaches below the mean of the weighted plot.

The results point toward the use of the PLOW method to obtain the most accurate distribution of a random variable when the number of possible inputs exceed the computational resources to sample the entire distribution. Such situations are prevalent in material design exploration where the possible input values for parameters like composition and annealing temperature in heat treatment studies are continuous (i.e. can take an infinite number of values), which can lead to very different bulk properties [75, 76, 77].

Although the PLOW method creates a more identical shape to the target distribution than with-

out it, this difference is rather small. To strengthen its credibility, a test of the effectiveness of PLOW on sub-sample data selected from a more disparate target distribution would be required. For the examples shown in the current work, the ZA parameter sets contained unique values, however these parameter sets were very close in value as a result of performing MCMC/random selection within a pre-determined range. One may not always have a data set to evaluate in which the input conditions are very close among candidate options.

3.6 Conclusions

In this work, a method for predicting the true output distribution of a quantity of interest from only a limited number of evaluations is presented. The PLOW method has been applied to a simple constitutive function and a more sophisticated hydrodynamic simulation to demonstrate its flexibility in different applications. Such a method has the potential to reduce the number of input samples needed to best represent the output distribution of interest, especially considering the cost of high fidelity simulations. Major findings include (1) the ability of PLOW to better predict a QoI distribution across a high fidelity simulation and (2) The PLOW method's ability to predict the distribution of a QoI with improved accuracy, even when a selection algorithm is used to obtain a representative subset of inputs. Additionally, the gains in prediction accuracy were most prevalent in cases where computational resources limit calculations to only a small set of inputs. Further, the method is able to perform this calculation using less information (only the inputs and their corresponding outputs) compared to other re-weighting methods. Here, the effectiveness of the method has been demonstrated on a problem where material performance is the quantity of interest. However, the approach is general enough that it can be applied to problems across essentially any discipline. Ultimately, it has the potential to expand the capabilities of scientists without access to expensive computational resources.

Acknowledgements

The authors acknowledge the Air Force Research Laboratory core funding in supporting data collection for this study. R. Arroyave and D. Allaire also acknowledge the National Science Foun-

dation through Grant No. 2001333. M. Gonzales gratefully acknowledges support from the Air Force Office of Scientific Research lab task no. 18RXCOR042 (Dr. Martin Schmidt, Program Officer). This work was supported in part by high-performance computer time and resources from the DoD High Performance Computing Modernization Program.

4. INCREASED FATIGUE LIFE PREDICTION ACCURACY OF Ti-6Al-4V USING UNFOLDING

4.1 Abstract

The fatigue life of near-alpha titanium alloys in service is a function of dwell time and microstructure. Clusters of alpha phase with similar c-axis orientations, known as micro-textured regions (MTRs), provide a local environment which facilitates initiation and growth of sub-surface cracks. Models exist for prediction of fatigue crack growth rate taking into account both dwell time and microtextured region parameters such as size; however, to date only 2-dimensional (2D) measurements have been used as inputs to this inherently 3-dimensional (3D) problem. In the present work, the MTR regions are assumed to be in the shape of prolate spheroids and the major and minor axis lengths are measured in cross-section. Then, the 3D size and shape distributions are estimated using the expectation-maximization and the Cruz-Orive spheroid unfolding algorithm. The magnitude of uncertainty in the unfolding results and implications for 2D fatigue life models for Ti-6Al-4V will be discussed.

4.2 Introduction

Titanium-aluminium alloys are often used in turbine engines such as those for commercial aircraft. Near alpha-titanium alloys, like Ti-6Al-4V in particular, get their value from being both lightweight and strong. Despite, their strength, these alloys are known to vary largely in their cyclic loading properties and time dependent properties [78]. They are generally resistant to deformation in aspects such as fatigue and creep deformation. However, studies have shown their fatigue life can significantly decrease depending on how the part is loaded [79, 80, 81]. Therefore, monitoring the health of an engine as a function of time spent in detrimental loading conditions is important to the safe transport of goods and, more importantly, human lives [78]. Effective monitoring, and even prevention, starts at the microstructural level.

Researchers have looked closely at loading conditions and material makeup to understand their

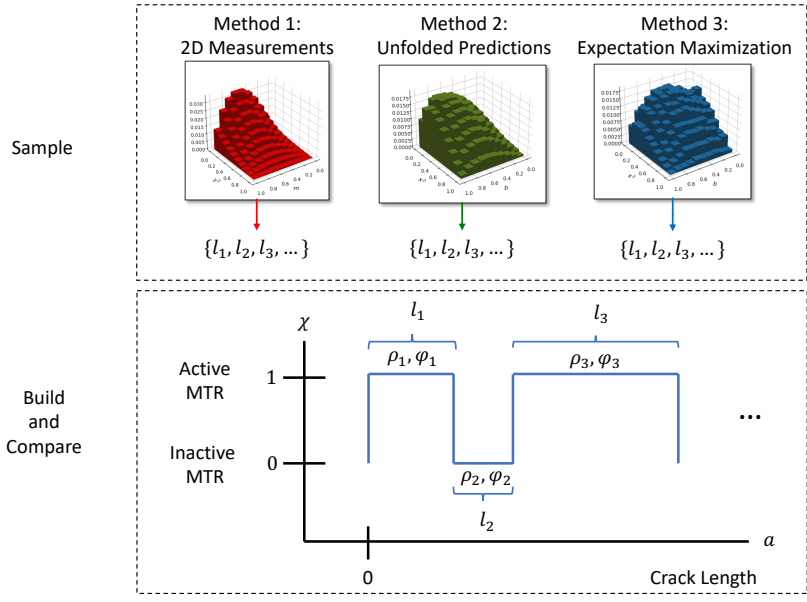


Figure 4.1: Method to compare the approaches.

effects on crack growth rate and have found that in dwell fatigue, growth rates of small cracks were faster in regions of high texture versus low texture [81], resulting in a significant debit in part life [82]. In near alpha-titanium alloys, regions can exist within the microstructure where there are many grains that are oriented in the same direction. This texture is the result of α phase grains which are usually present at low aluminum concentrations. Such regions of texture are commonly referred to as Micro-Textured Regions (MTRs). They have been measured to range in size from 100 microns to 10,000 microns [83] and have been linked to faceted fracture surfaces [84].

Further experiments have shown that these Hexagonal-Close-Packed (HCP) α grain regions' influence on crack growth rate can vary depending on how closely the c-axis direction is oriented along the axis of applied stress [81]. Although there is some consensus on what characteristics of MTRs affect growth rate, the mechanism by which they do it is not fully agreed upon. Evans and Banche's modified Stroh model of a pileup of dislocations between nearly orthogonal grains is frequently referenced in related literature [85, 86], while others believe grain orientation data of dwell fatigue fracture surfaces tell a different story [84]. Whatever the mechanism, a key component to accurate life prediction is capturing the most influential attributes and incorporating them

appropriately in models.

Jha et al. developed a model to shed more light on the likelihood of dwell fatigue failure based on the type of loading [87]. An enhancement to this model predicts part life as a function of MTR size, grouping density, and c-axis inclination [88]. The model, shown in Equation 4.1,

$$\frac{da}{dN} = \left(\frac{da}{dN} \right)_{cyc} + \chi_I * t_h * \left(\frac{da}{dt} \right)_{dwell} \quad (4.1)$$

will be used in the present study (see Equation 4.3 for the fully expanded version). Here, χ_I is the indicator function, which is used to determine if the characteristics of a given MTR meet the user defined criteria necessary for it to contribute to crack growth during dwell fatigue. The dwell time in hours is represented by t_h . The first rate term $\left(\frac{da}{dN} \right)_{cyc}$ is the cyclic small crack contribution. Expanded out, it has the following form

$$\left(\frac{da}{dN} \right)_{cyc} = e^c \Delta K^n,$$

where c and n are crack growth constants and ΔK is the change in stress intensity during the cycle. It is modeled to represent a through crack, so it has the form

$$\Delta K = (\sigma_{max} - \sigma_{min}) * \sqrt{\pi * (a + a_0)}$$

where σ_{max} and σ_{min} are the maximum and minimum stresses imparted on the material during one cycle, respectfully. In [7], Jha et al. gives more details about the small crack growth model (CGM) and information about the distributions of the parameters C and n for Ti-6Al-2Sn-4Zr-6Mo (a similar alloy to Ti-6Al-4V of the present study). The second rate term, $\left(\frac{da}{dt} \right)$ is the small crack growth contribution due to dwell fatigue

$$\left(\frac{da}{dt} \right)_{dwell} = e^{C_d} K_{max}^{n_d}$$

where C_d and n_d are dwell fatigue crack growth parameters and K_{max} is the maximum stress

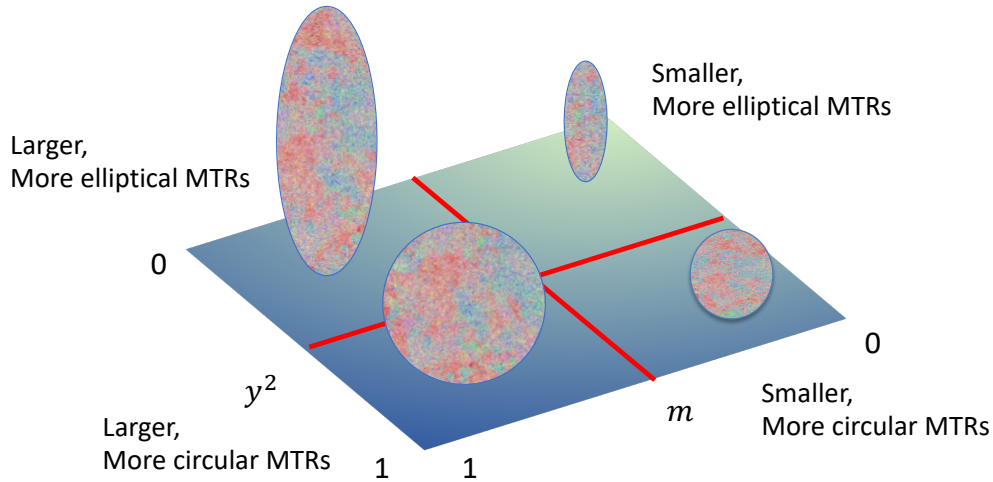


Figure 4.2: Size of MTRs with relation to region in plot.

intensity factor experienced near the crack tip during the dwell time. It is also modeled based on a through crack so it has the following form

$$K_{max} = \sigma_{max} * \sqrt{\pi * a + a_0}$$

The χ_I parameter is crucial to incorporating MTR features into the model. With inclination angle and grouping density information, it is able to account for anisotropy of HCP alpha phase where the hardest grains are those with a c-axis parallel to the loading direction and the softest grains are those with a c-axis perpendicular to the loading direction [86]. The CGM does not incorporate, however, the path of the crack because correlations to growth rate and crack path have not been found in previous studies [81]. Ultimately, it is concluded that the size, grouping density, and c-axis orientation with respect to the loading direction are the biggest contributing factors to accelerated crack growth.

The focus of the current work is to understand how much influence 3-dimensional MTR size, or length, has on the predicted dwell fatigue life of Ti-6Al-4V. In the form of a question: If one were able to measure the major principal semi-axes (mapsas) of prolate-spheroid shaped MTRs in 3 dimensions (3D), would that create a different result in part life than using the mapsas of ellipse

shaped MTRs in 2 dimensions (2D)? Currently, the most direct way to measure MTR size is in 2 dimensions via electron back-scatter diffraction (EBSD). The outcome of this study will reveal if the current standards of collecting and using 2D MTR measurements is sufficient for accurate life prediction. It is important to note that the study is purely numerical, since only a comparison between the measurement techniques is needed.

Stereology, “the science of geometric sampling with applications to the statistical analysis of microstructures in biology and materials science” [89], provides a great starting point for systematically addressing the question. Two techniques, Unfolding and Maximum Likelihood via Expectation Maximization (EM), can be used in the context of stereology to make estimates of mapsa lengths in 3 dimensions, from 2-dimensional data [90, 89, 91, 92]. These techniques are used in the present study to answer the question presented.

Because MTRs are essentially spread isotropically throughout the material, the ellipse measurements obtained from a single section of the bulk material consist of random sections of the 3D parent MTRs. As a result, many of the 2D ellipses within the data set will originate from intersections of spheroids where the section is not coincident with the major principal semi-axis of the spheroid. Such an ellipse would have a smaller major principal semi-axis length in 2D than its parent spheroid. The hypothesis is that the estimation methods will predict longer mapsas and thus shorter fatigue lives compared to the 2-dimensional distribution.

What follows is an explanation of the two approaches employed to estimate 3D size distributions. Next, the methods used to create size estimations and crack growth curves are described. Finally, the results of the study are shown and discussed.

4.3 Estimation Types

4.3.1 Unfolding

Unfolding is a stereological technique used to estimate the true, or 3-dimensional, size distribution of particles within a body based on a size distribution of the particles from a 2-dimensional or 1-dimensional section of the body [93]. Since its creation, the process of unfolding has been

improved upon and modified for different particle geometries. One method, for example, has been developed that performs unfolding on one section of an object, assuming the parent shape is a sphere [94, 95].

One of the most direct approaches to obtaining a distribution of the 3-dimensional size of inclusions is by taking multiple sections of the inclusion-containing body, measuring some geometric feature of the inclusions in each section, then stitching those sections together to reconstruct the underlying inclusion. Here each of these sections must be parallel to each other. This approach is referred to as the dissector method, as at least two sections are needed to perform the analysis [96]. While this is a direct way of measuring particle size, the method can be biased if sections are taken in a region of the material where inclusions are oriented in the same direction, although they may be isotropically distributed within the body.

Another approach to unfolding is to take three perpendicular sections of a specimen and, using geometric relationships, infer the 3D shape of the inclusions. Such an approach is referred to as the tri-sector method [97]. This method is less biased than the dissector method. However, both the dissector and tri-sector approaches are time and resource intensive, as many sections must be prepared and characterized. Furthermore, they are destructive to the body being studied. Prior to such techniques, it was clear to mathematicians that the ability to make size estimations of 3-dimensional particles from 2-dimensional measurements, or to unfold data, from a wider variety of shapes using one section would be very valuable.

S. D. Wicksell was the first to develop an approach to estimating the size-shape distributions of spheroidal particles (3D spheroids) from their ellipsoidal measurements (2D ellipses) using one section. He first did so with spheroidal particles spread throughout a matrix, later expanding the algorithm to include prolate and oblate spheroids [94, 98]. Following Wicksell, others began to improve upon his approach and apply it to specific cases including Cruz-Orive in his work [16] published roughly 50 years later.

While there are ways to estimate 3D inclusions based on multiple sections, we are limited to data from one section of the Ti-6Al-4V material in the present study. Furthermore, the 2D

MTR data is assumed to be ellipsoidal [99] and therefore prolate spheroidal in shape in 3D. These attributes about our data pointed to the work of Cruz-Orive (CO) who developed a set of equations that can perform unfolding for oblate or prolate spheroids, given a frequency distribution of major and minor principal semi-axis lengths determined by a section of the material.

The output of CO's algorithm, shown in Equation 4.2, is a multi-variate histogram of the minor principal semi-axis lengths and eccentricities of the parent spheroids [90].

$$g_{ij} = (\bar{H}/\Delta) \sum_{\alpha=i}^s \sum_{\beta=j}^k p^{i,\alpha} * f_{\alpha,\beta} * q^{\beta,j}, (i, j = \overline{1,10}) \quad (4.2)$$

Here, $p^{i,\alpha}$ and $q^{\beta,j}$ represent the elements of the size and shape adjustment factors respectfully. $f_{\alpha,\beta}$ represents the bi-variate histogram (m, y^2) where m is the minor principal semi axis of an ellipse and $y^2 = 1 - (m/M)^2$ is the shape component of that ellipse. Here M is the largest measured major principal semi-axis length of all of the ellipses. The bar over 1,10 denotes the selected bin ranges of i and j which, in this case, are from 1 to 10. The random variable ranges are $\{0 < m \leq B\}$ and $\{0, <= y^2 < 1\}$ where B is the largest minor principal semi-axis value measured in the data set. The variables, s and k , represent the total number of bins of the input histogram for the size and shape components, respectfully. The value of each variable is set by the user. \bar{H} represents the overall mean caliper diameter and $\Delta = B/s$. To simplify the analysis, the major principal semi-axis values are normalized, making B equal to 1. For more details on the algorithm, the reader is referred to the source paper [90].

Once all of the elements of the g_{ij} matrix are calculated, a distribution such as that shown in Figure 4.3a and 4.3d is produced. This figure is a reproduction of Cruz-Orive's synthetic input distribution (Figure 4.3a and 4.3d) and the unfolded output (Figure 4.3b and 4.3e) using Python. For the unfolded histogram, the random variables are labeled x^2 and b . These correspond to the shape component of a 3-dimensional MTR and its minor principal semi-axis length, respectively. Furthermore, the shape component, $x^2 = 1 - (\frac{b}{a})^2$, where a is the major principal semi-axis of the spheroid. Given the type of data available and the inferred shape of MTRs, Cruz-Orive's method

of unfolding has been applied. It is from the unfolded histogram that an MTR's major principal semi-axis can be back calculated with the expectation that the estimated value more accurately represents the MTR in 3 dimensions.

4.3.2 Expectation Maximization

At its core, Expectation Maximization (EM) employs an iterative optimization technique to determine the parameters of the likelihood function, given a set of observed random variable measurements. In literature, the likelihood function is often referred to as the unobservable aspect of an observable set of data. Determining the function is useful for predicting the chance that a random variable will take a particular value in the range of possible values. When one maximizes the likelihood, he or she is searching the function parameter space for values that most accurately describe the measured data. Do and Batzoglou provide a great explanation of the EM process and walk through a practical example of how it can be applied [100].

One of the earliest implementations of estimating the likelihood function from incomplete data was by Hartley in 1958. He developed an improved methodology of calculating the underlying parameters of a data set (i.e. calculating the maximum likelihood parameters) with missing data [101]. Later, Dempster et al. followed up Hartley's work by providing the derivation for applying maximum likelihood to cases when data was missing from a data set *via* expectation maximization [102]. Even techniques for smoothing out the estimation from EM were developed to account for some issues of regular EM which can give "noisy" or "spiky" estimates of the unknown space [103]. Outside of statistics, Vardi et al. applied EM to the medical field, helping radiologists estimate the rate of occurrence of positron emissions in the body, and in turn enhance images of organ activity [104].

Others saw an opportunity to apply these techniques to materials science problems. For example, Baaske et al. studied the estimation of the parameters of a function that described the distribution of spheroids using maximum likelihood after unfolding [91]. The distribution however is based on a bivariate histogram which is not the most realistic representation of the size-shape distribution of prolate spheroidal shaped inclusions within a matrix. Studies have shown that the

size distribution of MTRs in a near-alpha titanium alloy is not normally distributed. Rather, the distribution is skewed with smaller major axis and minor axis lengths occurring more frequently than larger ones [83].

Unlike Baaske et al., Chan and Qin developed an algorithm for applying EM in a non-parametric way. They first considered the case of estimating the likelihood of a spheroidal distribution using a mixture of 1-, 2-, and 3-dimensional data about the spheroidal inclusions. The use of EM helped overcome the problem of sampling bias and indirect measurements that can come from obtaining lower dimensional measurements from 3D spheroids. Later, in the same work, the authors expand the EM approach to the case of ellipsoidal measurements. This expanded approach is derived from the distribution function describing the size-shape frequency of prolate ellipsoids in a material developed originally by CO [92]. In the current paper, Chan and Qin's EM approach is used for all figures with a legend element labeled "EM". References to the EM approach specifically developed by Chan and Qin will be referred to as "Chan-EM" in the text.

4.4 Methods

4.4.1 Obtaining MTR Measurements

The elliptical dimensions of the MTRs were obtained from a near alpha-titanium alloy, Ti-6Al-4V. The alloy was sectioned in order to analyze the microstructure using electron back-scatter diffraction (EBSD). A combination of the EBSD data and the software DREAM.3D [105] enabled the quantification of 678,963 MTR minor and major principal axis lengths. More information on the criteria chosen to identify MTRs are described in [83].

Because the specific shape of these regions in 3D can be irregular and distorted due to processing history, it is assumed that the MTRs have the general shape of a prolate spheroid. A single section of this spheroid would produce an ellipse (excluding the instance where it is sliced coincident with the minor semi axis which would produce a circle). Simplifying the geometry is also helpful for modeling purposes.

4.4.2 Coding and Use of Unfolding

From the data set of 2D MTR size measurements, it is of interest to infer the 3D size and shapes of the parent prolate spheroids. Cruz-Orive's model for estimating the bi-variate size-shape histogram of prolate spheroids from the size-shape histogram of ellipses is programmed into Python in order to quickly process the large experimental data set.

The input of the algorithm is a bi-variate histogram, therefore the experimental data was binned in such a way that it created a bi-variate histogram. The Freedman-Diconis rule was first employed to determine the number of bins for each variable [106]. Using the rule, the number of recommended bins was 2,113 and 147 for m and y^2 , respectively. Together, this would create a bi-variate histogram of (2,113 by 147) 310,611 bins. The approximate number of points per cell would be 678,995 points divided by 310,611 bins resulting in 2.18 points per bin. This value is too small considering that the purpose of binning is to obtain a general shape of the distribution of data. The rule of thumb for selecting the number of bins for a data set is that the average number of data per interval should be greater than the number of intervals.

For the current study, a constant grid size of 20 by 20 was selected. Even so, the codes developed by the authors was written to allow the user to select the number of bins according to their data set size. For this study, the effect of bin sizing was not accounted for and was thus held constant for the calculations.

Validation of Unfolding Code

The input histogram data, provided in Cruz-Orive's paper, was replicated and passed into the unfolding function in order to validate the implementation of the code. Figure 4.3 a and b show the replicated input and corresponding output of Cruz-Orive's algorithm, respectively. Visually, the inputs and outputs from the coded function, and that shown in Figure 1 in CO's paper are in very good agreement [90]. Before applying the method to the experimental data however, it was important to understand how the algorithm performed when the input data did not exactly resemble the synthetic input data supplied by Cruz-Orive. This was done by sub-sampling the data from a distribution shaped like CO's input distribution.

A sample of values, equal to a predetermined sub-sample size, was taken, binned, and passed into the unfolding function. This process was repeated for various amounts of sub-sample sizes. The result of the test was a continuously decreasing Residual Sum as the number of input samples was increased. This represents how the unfolded histogram predictions get closer to the known solution as the input sample size approaches that of the known input.

Upon plotting the unfolded histograms from many of the sub-sampled data sets, it became clear that some bins predicted negative values. This issue stemmed from the sub-sampled data leaving some bins in the multi-variate space empty or very near 0, which would cause numerical issues with unfolding (i.e. low valued negative bins that do not sum to positive values). For clarity, all histograms shown in the paper and created with CO's unfolding function are plotted with any negative bin values set to zero. Such bins have a non-physical meaning. Additionally, after testing the code with larger input data sets, it became clear that a major drawback of CO's method is that it is ill-conditioned, meaning that very small changes in the input values can result in large changes in the outputs. (CO confirmed this directly via email.)

The sensitivity of the CO approach prompted the need for an alternate, more robust 3D size estimation approach. Ultimately, the problem of 3D size estimation is one of estimating the parameters of a "hidden" distribution that described the observable 2-dimensional ellipse data. Therefore, maximum likelihood via expectation maximization was used as an alternative method to compare with the unfolding method.

4.4.3 Coding and Use of EM

Maximum likelihood, a proven statistical technique, provides a way to estimate the unknown or "missing" [102] function parameters most likely to produce a set of experimental data points. The method can be applied to cases where it is expected that multiple distribution functions are responsible for the observed data. If taking the integral of a likelihood function is intractable, expectation maximization can be used to estimate the parameters of the function. At its core, EM is an optimization method with the objective to maximize the expectation, or mean, of the likelihood function with the observed data. When employing EM, it is helpful to select an appropriate

likelihood distribution function if there is prior knowledge about the shape of the distribution. A more general approach, however, would be to use an EM process that is distribution free, meaning it does not assume a functional form of the likelihood during the optimization process.

Such is the case with Chan and Qin's estimation algorithm [92] that considers the case of estimating the 3D size-shape distribution of prolate spheroids from a 2D size-shape ellipse distribution. Unlike CO's unfolding approach, the inputs of the Chan-EM algorithm are the individual occurrence probabilities associated with each ellipse measurement. With each iteration of the Chan-EM process, the probability of each ellipse is "adjusted" to reflect that of the 3D size-shape distribution of its parent spheroid. Once programmed into Python, the Chan-EM algorithm was used to calculate the estimated 3D size-shape distribution of the prolate spheroids.

One drawback of Chan's approach is that it can become extremely memory or time intensive (depending on implementation) as the amount of inputs to process increases. Much effort and time was put into optimizing the code to strike a balance between available time and random access memory (RAM) available. To reduce time, the algorithm was coded using matrix and vector multiplication to reduce for-loops, resulting in a $400\times$ speed increase. Additionally, packages such as Numba were integrated to help speed up the code even further. To account for RAM, a number of strategies were implemented including deleting high memory usage variables when available and using the highest memory capabilities offered by TAMU Supercomputing resources.

Despite these efforts, the highest amount of data points that could be run from the total of 678,963 was 400,000; roughly 60% of the data. The data was sub-sampled without replacement and included the highest and lowest experimental value from the total data set.

Validation of EM Code

Prior to the current work, the EM algorithm proposed by Chan and Qin had not been programmed, nor tested. (Chan admitted this via email.) Once coded here, the algorithm was tested using the synthetic data set proposed by CO in [90]. The idea was that the Chan-EM algorithm should produce the same, or a similar, output distribution as CO's method, if implemented correctly. The result of the test is shown in Figure 4.3c and 4.3f. Compare those histograms to b and

e of Figure 4.3 which are the known unfolded histograms of the synthetic outputs.

To apply Chan-EM to a set of data, the random variable measurements that make up the data are required as inputs. Because these values are not provided by CO which make up his synthetic input histogram, a random sample was required that would represent the known input histogram as closely as possible. These inputs would be used in the Chan-EM algorithm to produce a histogram of the 3D MTR size estimates. The residual sum (RS) of bin heights between the two histograms served as a metric to gauge how close the Chan-EM and the known unfolded histograms were to each other. With increasing sample sizes, the predicted Chan-EM histogram approached the shape of the known unfolded histogram. Although the outputs were similar at sample sizes of 200,000 inputs, slight differences still existed between the two distributions which prevented the RS from reaching 0. This difference in histogram values are likely the result of non-identical inputs for the two approaches.

Because the input random variable measurements for CO's method do not exist, it is very difficult to recreate a data set that, when binned, makes the same histogram shape as CO's synthetic input histogram. Also, CO's synthetic input histogram was made "perfectly" for his approach and does not reflect a real data set. Since the method is ill-conditioned, CO had to craft the inputs in such a way that they would create a "good" or "reasonable" looking output. Lastly, the Chan-EM method was never validated against an experimental data set, so there are still some questions about it's accuracy.

4.4.4 Sampling from the Distributions

Two steps were required prior to incorporating data into the crack growth models. First, grouping density values were sampled from a log normal distribution with a mean and standard deviation of -1.90975 and 0.30897, respectfully. Second, c-axis inclination values were generated by sampling from a mixture of two log normal distributions truncated to range from 0 to 90 °s. The mean and standard deviation for one of the distributions was 4.22733 and 0.334652, respectfully. The other distribution had a mean of 3.55245 and standard deviation of 0.726819. The values generated here will be used in the construction of the χ_I variable.

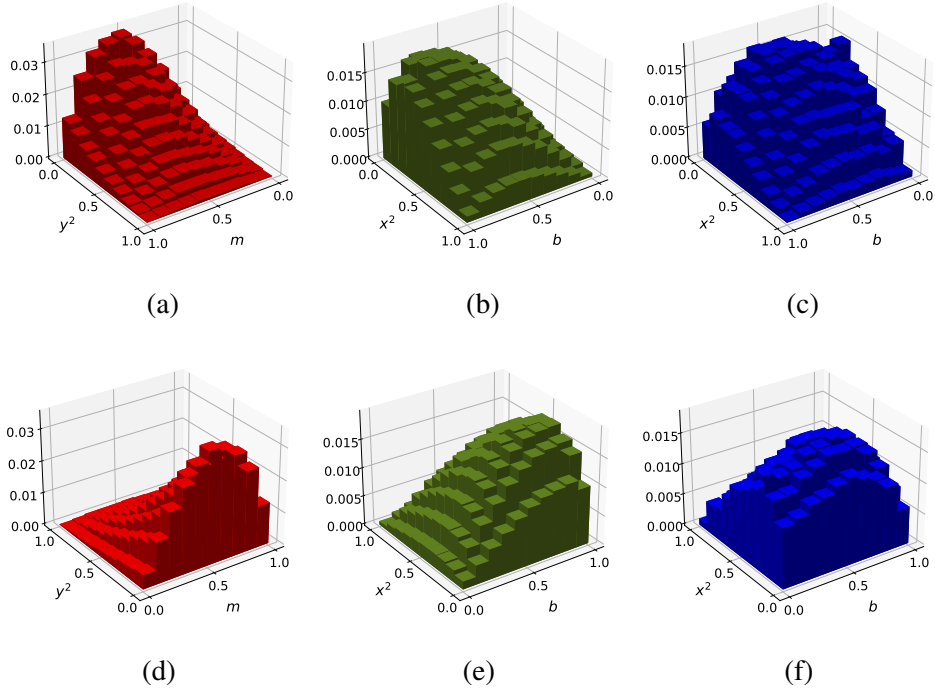


Figure 4.3: Histograms of the 2D, Unfolded, and EM-Chan probabilities of the synthetic data. The grid size of the histograms are 10 by 10.

Next, the experimental data is binned into a 20 by 20 celled bi-variate histogram. One variable, the minor principal semi-axis value (mipsa) (m) is taken directly from the measurements. The eccentricity (y^2) is calculated from the the mipsa and largest measured major principal semi-axis (mapsa) (M) value of the data set using

$$y^2 = 1 - \left(\frac{m}{M}\right)^2$$

Next, the bi-variate histogram of 2-dimensional data is passed into the unfolding function, resulting in a histogram showing the estimated frequencies of 3D prolate spheroid size-shape classes in the material. The same bivariate histogram of 2-dimensional data is passed into the Chan-EM function to produce another estimate of the 3D prolate spheroid distribution. The three histograms can be plotted: one representing the size and shape distributions of the two-dimensional MTR sec-




Figure 4.4: Illustration showing how crack growth model works in 2 Dimensions.

tion measurements, one of the unfolded estimates, and the third of the Chan-EM estimates. Figure 4.5 displays each of the histograms for the experimental data.

From here, the Numpy random sampling function is employed to sample m and y^2 values from the unfolded and Chan-EM histograms. It is also used to sample b and x^2 values from the 2D ellipse histogram. The sampled mipsa lengths and shape parameters from each distribution type are used to back-calculate the mapsa lengths.

Finally, the chi function, χ_I , is constructed using a combination of mapsas, grouping densities, and c-axis inclinations. The bottom half of Figure 4.1 shows an example of how this data is used to determine the size and activity of the MTRs. Mapsa lengths are represented as horizontal lines. As for the MTR sizes, a minor principal semi-axis length and an eccentricity are simultaneously sampled randomly from the appropriate distributions. The two values are then used to back-calculate a major principal semi-axis length which is used as the length of the MTR. Each length is paired with a grouping density and c-axis inclination. During the chi function construction process, a value is randomly selected from the distributions of the c-axis orientation and grouping density variables. A two state plot is then generated which represents a crack growing through one MTR (state 1), leaving the MTR, and then entering another (state 2). The left hand side of Figure 4.4 shows this information graphically. For the model, it is assumed that a growing crack is propagating through

MTRs at all times. However, not every MTR is active so growth through inactive MTRs represent the crack growing through a normal region of the matrix.

If the c-axis inclination is below 20° with the loading axis and the grouping density above 0.2, then the MTR is considered Active, the indicator function value, χ_I becomes 1, and the dwell time effect on crack growth is considered. [88]. Experiments have shown that facets of HCP grains can range from 10-15 degrees in misalignment with the principal loading direction [107]. With χ_I constructed, it can be substituted into Equation 4.1, the crack growth model. It is important to note that Figure 4.4 simplifies the path of the crack into a straight line. In reality, a propagating crack would not follow such an ideal path. The model compensates for this however, through the semi-random alignment of active and inactive MTRs.

4.4.5 Creating A vs N plots - Crack Growth Model

To represent the progression of crack growth in the present study, crack length versus number of cycles plots were generated. They were formed using the model described by Jha et al. [88] method shown in Equation 4.3 via a crack length increment approach.

$$da = e^C * \left((\sigma_{max} - \sigma_{min}) * \sqrt{\pi * (a + a_0)} \right)^n * dN + \chi_I * t_h * \left(e^{C_d} * \sigma_{max} * \sqrt{\pi * a + a_0} \right)^{n_d} * dt \quad (4.3)$$

In the equation, dt and dN are set as 1. An illustration of the way the model incorporates data is shown in Figure 4.4. Overall, the algorithm is a 4 step process:

1. The value da is calculated using the crack growth model shown in Equation 4.3. Here C and n are held constant at -27.19 and 3.75, respectfully. The C_d and n_d values are randomly selected from two normal distributions with mean and standard deviation -33.7765 and 0.0193 for C_d , and 4.093 and 0.0059 for n_d . Sampling in this way allows the influence of the dwell fatigue contribution to be isolated and observed. In the model, dN and dt are set to 1. Each iteration is identified by 1 cycle and when the dwell time contribution is activated, a differ-

ence of one time unit passes. The initial length of the crack, a_0 , was set to 10.8E-6 meters (or 10.8 microns) which is based on the average size of alpha particles documented in Pilchak et al. [81].

2. The length calculated in step 1 is added to the current length of the crack. This becomes the new total length of the crack, a .
3. The total length of the crack is plotted on the a vs N plot at the corresponding number of total cycles where fracture is expected to occur. An illustration of this step is shown on the right side of Figure 4.4. The line connecting the points are for illustration purposes only.
4. The code checks to see if the total length of the crack is still within the same MTR as in the previous cycle. If so, the value of χ_I stays the same. If not, it is changed to the opposite value. With each iteration, the code also checks if the stress concentration at the crack tip is greater than a critical value set as $60 \text{ MPa} \sqrt{m}$. If it is not, the process repeats. If it is, the process ends.

The crack growth model algorithm is repeated for a given C and n value in Equation 4.3. It is also repeated for each data set, i.e. a set of a vs N curves is generated from each histogram (2D, Unfolding based 3D estimation, and Chan-EM based 3D estimation). The final result is a distribution of fatigue life which is informed by the MTR parameters.

4.5 Results and Discussion

4.5.1 Bi-Variate Histograms of MTR Data

The 2D, 3D (or Unfolded), and EM (or Chan-EM) distributions for the experimental data are shown in Figure 4.5. Figure 4.5a shows the bivariate distribution of the more than 600,000 experimental 2D ellipse measurements. As shown, the shape of the histogram is very different than that of Cruz-Orive's synthetic input data set. This makes sense given that his data set is not modeled after a material system of a nucleating phase.

The histograms share a similar shape when observed on the minor principal semi-axis direction or margin. In each, the largest number of MTRs are small in length fraction (i.e. contained within the m or b bin ranges of 0.0 to 0.05). Bin heights taper off significantly in the 0.05 to 0.1 range and continue to decrease rapidly as bin ranges approach 1. The shape indicates that a very large fraction of data that consist of smaller MTRs. The shape of the histogram does reflect the known behavior of a nucleating phase in a material. Generally, many areas begin to nucleate into the next phase with a small number growing large enough to combine with other nucleation sites to form a larger nucleating phase. In addition to the smaller MTRs outweighing the larger ones in number, they also outweigh them in influence of crack growth. Any MTRs that would accelerate the growth of the crack would only do so for a few microns. This is the case with each histogram and as a result, their predicted number of cycles at fracture are very similar.

A major difference between the histograms however, is the gradient in eccentricities (y^2 and x^2). The histogram displaying experimental data (Figure 4.5a,d) has a significant amount of eccentricities ranging from 0 to 1 in the m fraction size range of 0.0 to 0.05. The Chan-EM histogram preserves some of this distribution, but does not predict significant bin heights at smaller eccentricities.

According to Z-axis heights, it is clear that the unfolding predictions (Figure 4.5b,e) have the least spread over the $b-x^2$ space. Prior to plotting the histograms, the bin heights are normalized. Because less bins occupy cells outside of the front right corner of the space (0.0 - 0.1 b , 0.7-1 x^2), the fraction of sizes contained within it are higher than the 2D and EM histograms. The zeroing out of bins could be contributing to the added weight (or increased bin height) of non-zero bins. The ill-conditioning of the CO method makes it less robust than the EM approach for making size predictions.

4.5.2 CDFs of Lengths

The cumulative distributions of MTR lengths sampled from each histogram are plotted in Figure 4.6. The mipsa and eccentricities are simultaneously sampled from the appropriate bi-variate size-shape distributions in Figure 4.5. Mapsas are back-calculated using the mipsa and eccentricity




Figure 4.5: a) Bi-variate histogram of the experimentally measured MTR major principal semi-axis. b) Unfolded histogram of experimentally measured MTR data with Expectation Maximization applied. c) Same as b) but Expectation Maximization is not applied. The histograms are of size 20 by 20.

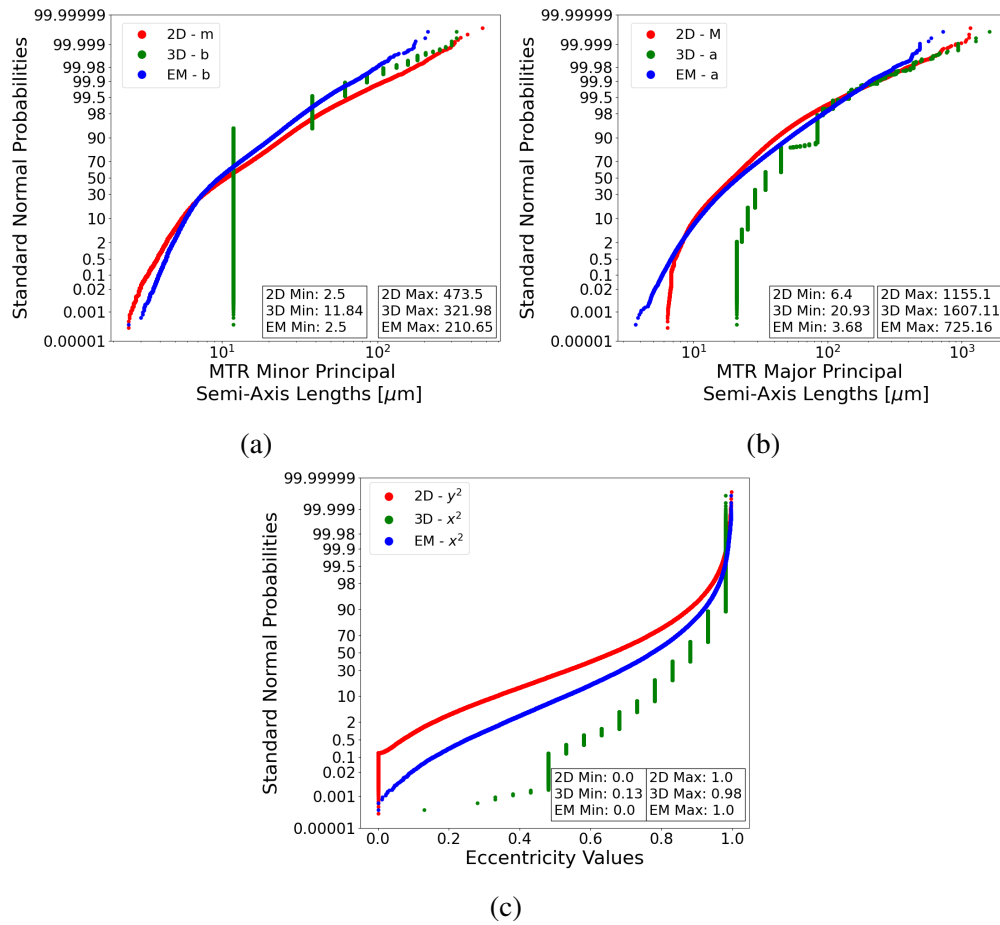


Figure 4.6: Values of 3D and Chan-EM are from sample of 400K points. Chan-EM probabilities are generated from 400K experimental data points uniformly sampled without replacement. Displayed in each sub-figure are (a) Major Principal Semi-Axes or Mapsas, (b) Minor Principal Semi-Axes or Mipsas, and (c) Eccentricities.

values and plotted in Figure 4.6c. For each data source displayed in the figures, the spread of points in size are a direct result of the bin heights in the corresponding histograms.

Of the three data sources, the 3D (or unfolded) points are the most sparse in terms of size or shape variation. The lack of variation is a direct result of having less bins to select from during the sampling process. Likewise, less bins are a result of negative bins being zeroed out after unfolding is applied. The sensitivity of the CO algorithm is not able to handle the non-ideal, experimental data well. Outputs of the model fluctuate significantly if they stray even slightly from the ideal, synthetic input data set.

Observe Figures 4.6b and 4.6a which represent the principal semi-axis dimensions of the MTRs. Here, the right tails of the 3D and EM estimation types are less populated than the experimental data. The result reveals that the estimation algorithms predict smaller MTR dimensions in 3D, than those measured from a 2D section. Put another way, it shows that a random section of a spheroid is more likely to result in an ellipse with smaller dimensions than the spheroid itself. This is in alignment with Wicksell's statement that "Apparent diameters are as a rule smaller than the true diameters" [94]. The impact of this can be seen in Figure 4.6b. For the cases of EM and 3D, it is estimated that there are much less MTRs with mapsas greater than roughly 500 microns compared to the 2D data. It shows how the experimental data reports that there are a lot of large MTRs when actually, those are more likely to be smaller as the estimation methods predict.

4.5.3 Crack Length vs Number of Cycles distributions

Figure 4.7 summarizes the predictions of the number of cycles at failure for each of the three methods. It plots the values in the form of cumulative distributions. The figure is constructed by taking the number of cycles at which the length of the crack is expected to cause fracture in the bulk material. It is assumed that fracture occurs when the stress concentration, K at the tip of the crack exceeds the yield stress of the material. Yield stress is used because the expression for fracture in the crack growth model is representing a through crack. In this condition, plane strain is not present.

The number of cycles is not very different between the types of length estimations. Such is

the case even for less populated numbers of cycles (i.e. less than 46,000 and more than 50,000) predicting the shortest, most conservative values. The closeness of the points is more apparent in the inset box of Figure 4.7, where the vertical line marks the B.001 life, or book life, a statistic used by the U.S. Air Force to report that probability of a component failing [108]. In the case of 1 million iterations, it represents the 1 thousandth number of cycles value from the minimum predicted life of the entire set ordered from least to greatest. (The B.1 and B.01 book lives represent the tenth lowest value of 10,000 values and the one hundredth lowest value of 100,000 values, respectively.) One line is shown for each B.001 life of the data sources, because the points are so close that the line is thick enough to capture them all. It is not until the very edges of the data set that the points begin to spread out from each other.

Points found near the left edge of Figure 4.7 are important to analyze because they represent the shortest life spans of the part. They are the result of a crack growing through a greater number of active MTRs with larger mapsas, causing the crack to grow faster than normal in a given amount of time. The data shows us however, that each data source has a similar size distribution of both larger and smaller MTRs. It also shows that sizes of the larger, more detrimental MTRs are captured well with just 2-dimensional ellipse data. It is likely that the over 600,000 measurements in 2D contribute to accurately capturing the 3D size distribution as well. From a statistical standpoint, it is well established that the larger the sample size of a random variable, the more accurate are the estimators (i.e. mean, variance, etc.) that describe the distribution. With so many MTRs intersecting the section where measurements are taken, there is likely a large population of MTRs, within the 680,000+, that were cut very near the major principal semi-axis of the 3D spheroid.

Data Source	B.1 Cycles	10K Median	B.01 Cycles	100K Median	B.001 Cycles	1 Million Median
2D	45609 - M	47651 - L	45638 - L	47660 - H	45643 - L	47659 - H
3D (Unfolded)	45613 - H	47670 - H	45645 - M	47659 - M	45645 - H	47657 - L
EM (Chan-EM)	45565 - L	47666 - M	45656 - H	47658 - L	45643 - L	47657 - L

Table 4.1: Important number of cycle statistics.

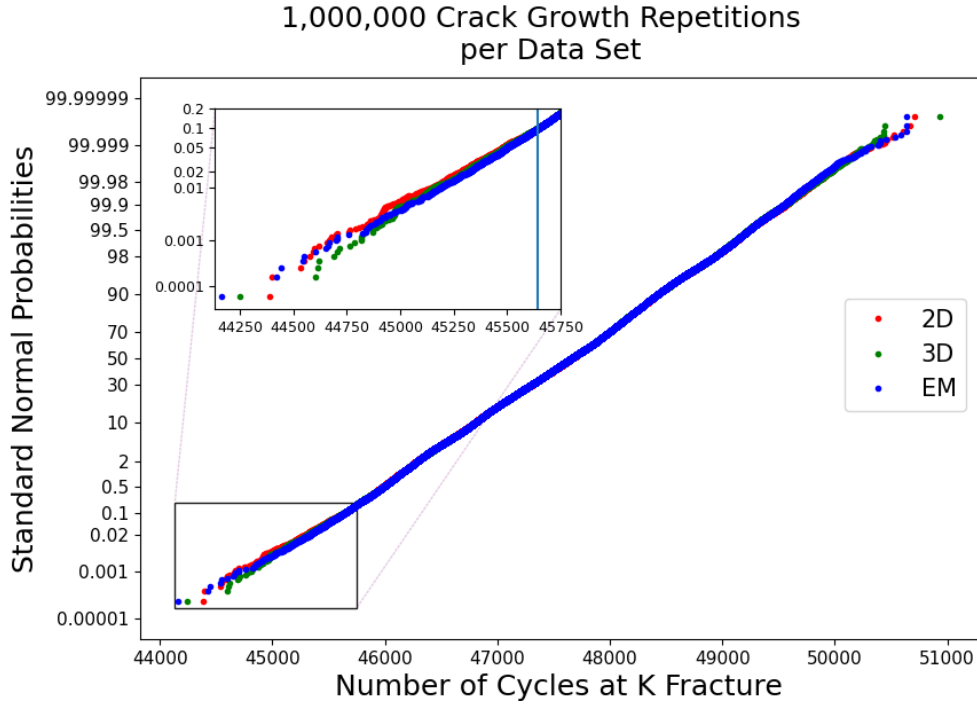


Figure 4.7: Cumulative distribution of cycle at fracture predictions for 1 million iterations of curves per data source.

In addition to the 1 million repetitions case, two additional cases of crack growth distributions were plotted; a 10,000 and a 100,000 repetition case. Their plots are not shown in the paper, but their important statistical values for the number of cycles at fracture are tabulated in Table 4.1. As shown in the table, as the iteration size increases (10,000 to 100,000 to 1 million), the B.1, B.01, and B.001 values get closer to each other. At 1 million iterations, there is little difference between the B.1, B.01, or B.001.

4.6 Conclusion

To properly characterize the size of 3D particles in a matrix, one does not necessarily need an estimate of the size and shape of the inclusions in 3-dimensions. The extra cost associated with obtaining the 3D distribution may not justify the increase in precision, considering that the change resulting from the 3D and EM measurements are negligible. The difference in length of MTRs between 2D and 3D is simply too small to support the increase in accuracy of measurements. For

those interested in characterizing the size-shape of inclusions for other tests, the combination of EBSD and DREAM3D is a high throughput option for capturing a large amount of data quickly. In fact, the abundance of measurements in 2D can perhaps make up for an absence of 3D measurements. In practice, this translates to increased confidence in mechanics making decisions on replacing a part, even when referencing crack growth curve data generated based on a single section of an engine component.

Another benefit of using 2D data is that there is a degree of a safety factor built into its cycle life prediction. This is because its largest mapsa lengths are more populated than those of the 3D and EM estimators. Any crack growth curve estimates that are based on a single section would theoretically be more likely to predict shorter lives than those constructed with 3D measured MTR data.

Although the histograms predict fatigue life similarly, if one needed to estimate 3D MTR size the Chan-EM method should be used. CO's full unfolding algorithm is fast by comparison, but its sensitivity to inputs makes it impractical for experimental use. Moreover, its sensitivity to inputs can cause it to predict negative frequencies that are nonphysical. Alternatively, the Chan-EM approach incorporates the functional form of the spheroidal distribution determined by CO, but applies it in a framework that doesn't predict negative values and is distribution free on the outputs. Currently, the only observed drawback for Chan-EM is that for large amounts of data, the method can be slow (or memory intensive depending on how it is coded).

Limitations in the current study were encountered, but addressed. One issue was that some MTR lengths captured by the DREAM3D software were smaller than the known resolution of the microscope used for EBSD. Therefore, measurements with a minor principal axis ($2 * m$) smaller than 5 microns were removed before conducting the study. Doing so would allow the results to resemble reality more closely. This size filter reduced the data from 678,995 to 678,963 measurements.

In terms of the Chan-EM algorithm, there was no application of the method for the prolate-spheroid case prior to this work. As a result, the algorithm was not able to be verified when first

coded nor after edits were made to increase the speed. Without the availability of a known data set, CO's input and output histograms were the most fitting for verification purposes. It has been shown in the current work that the algorithm performs well in predicting the synthetic output. Furthermore, now that the algorithm has been coded and optimized, it can be used by others in future studies.

Optimizing the code for increased speed created another barrier to calculation in that a large amount of RAM was required to calculate over the entire data set. To overcome this challenge, the experimental data was down-sampled to 400,000 points before applying the Chan-EM algorithm. To capture a representative sample, a uniform sampling scheme without replacement was applied. Fortunately, the extremes of the lengths (min and max) were captured in the sub sample and over 58% of the experimental data was used.

Due to limitations in computing power, the experimental histograms were restricted to a 20 by 20 grid space. Given the large set of data, one could argue that the grid spacing could be too coarse to accurately represent the distributions. However, the shape of the distributions, specifically the 2D data, still reasonably reflects the known distribution shape for a phase nucleating in a material. Using a 20 by 20 grid also created a baseline for which code optimizations could be compared and allowed for a large total of 400 cells across the space. Lastly, the code features an option for the user to easily set the amount of total cells they would like to use. Future researchers can adjust the amount of bins based on their amount of data and their computing resources.

For future work, the crack growth curve predictions can be plotted against experimental curves for Ti-6Al-4V. Such a test would allow a model to be classified as better or worse at prediction than others. Because the current work is a comparison study between the models, experimental data was not required to answer the research question.

Another aspect of future work could be in further verification of the 3D size estimation methods. Each method can be tested against a synthetically generated set of prolate spheroid-shaped particles in a volume that has been sectioned. A data set like this can replace the arbitrary input histogram described by CO and used in the current work. Any alignment of the estimated distribu-

tions with the synthetically generated one would ultimately reveal a more accurate representation of the crack-growth-dominated lifetime of the titanium alloy. Consequently, an engineer would know more precisely, what crack length would be unacceptable for an expected number of cycles in a given design.

4.7 Acknowledgements

Portions of this research were conducted with the advanced computing resources provided by Texas A&M High Performance Research Computing.

5. SUMMARY AND CONCLUSIONS

The implications of uncertainty quantification and propagation are demonstrated in the context of materials analysis. As shown, the frameworks used are general and therefore applicable beyond only materials science. AF9628 is a promising material, achieving high strengths from relatively cheap compositional constituents. Engineers looking to design military grade equipment can use the information documented here to verify applicability in their designs. The PLOW framework adds a proven tool for efficient uncertainty propagation to the Air Force. The method could one day be used within a super-computing scheduler to give users the option to run a sub-set of their inputs if resources are occupied by competing calculations. It has also been shown that 2-dimensional measurements of microtextured regions within the near alpha titanium alloy Ti-6Al-4V are sufficient for predicting fatigue life. The results give specialists increased confidence in crack growth predictions built using data measured in two-dimensions.

5.1 Further Study

Future studies involve comparing the numerical results calculated in the present work to experimental data. Such a comparison is critical to determining the accuracy of the models proposed. For predicting the high strain rate behavior of AF9628, additional constitutive models could be included in the fusion process beyond just the Johnson-Cook and Zerilli-Armstrong models. Additionally, the method of reification can be applied to the problem of predicting fatigue lifetimes. The problem has multiple information sources that help predict the same QoI, the length of the crack at fracture. Perhaps a more narrow distribution could be achieved by incorporating prior information about the system along with available models.

REFERENCES

- [1] T. J. Sullivan, *Introduction to Uncertainty Quantification*, vol. 63. Springer, 2015.
- [2] W. D. Thomison and D. L. Allaire, “A Model Reification Approach to Fusing Information from Multifidelity Information Sources,” *American Institute of Aeronautics and Astronautics*, Jan. 2017.
- [3] R. C. Smith, *Uncertainty quantification: theory, implementation, and applications*. Computational science and engineering series, Philadelphia: Society for Industrial and Applied Mathematics, 2013.
- [4] D. A. Stainforth, T. Aina, C. Christensen, M. Collins, N. Faull, D. J. Frame, J. A. Kettleborough, S. Knight, A. Martin, J. M. Murphy, C. Piani, D. Sexton, L. A. Smith, R. A. Spicer, A. J. Thorpe, and M. R. Allen, “Uncertainty in predictions of the climate response to rising levels of greenhouse gases,” *Nature*, vol. 433, pp. 403–406, Jan. 2005. Number: 7024 Publisher: Nature Publishing Group.
- [5] S. Sankaran and A. L. Marsden, “A Stochastic Collocation Method for Uncertainty Quantification and Propagation in Cardiovascular Simulations,” *Journal of Biomechanical Engineering*, vol. 133, Feb. 2011.
- [6] M. Mahmoudi, G. Tapia, K. Karayagiz, B. Franco, J. Ma, R. Arroyave, I. Karaman, and A. Elwany, “Multivariate Calibration and Experimental Validation of a 3D Finite Element Thermal Model for Laser Powder Bed Fusion Metal Additive Manufacturing,” *Integrating Materials and Manufacturing Innovation*, vol. 7, pp. 116–135, Sept. 2018.
- [7] S. K. Jha, R. John, and J. M. Larsen, “Incorporating small fatigue crack growth in probabilistic life prediction: Effect of stress ratio in Ti–6Al–2Sn–4Zr–6Mo,” *International Journal of Fatigue*, vol. 51, pp. 83–95, June 2013.

- [8] J. H. Panchal, S. R. Kalidindi, and D. L. McDowell, “Key computational modeling issues in Integrated Computational Materials Engineering,” *Computer-Aided Design*, vol. 45, pp. 4–25, Jan. 2013.
- [9] N. R. Council, D. o. E. a. P. Sciences, N. M. A. Board, and C. o. I. C. M. Engineering, *Integrated Computational Materials Engineering: A Transformational Discipline for Improved Competitiveness and National Security*. National Academies Press, Sept. 2008. Google-Books-ID: JlfSnH8jI4UC.
- [10] M. F. Horstemeyer, *Integrated Computational Materials Engineering (ICME) for Metals: Using Multiscale Modeling to Invigorate Engineering Design with Science*. John Wiley & Sons, July 2012.
- [11] P. Honarmandi and R. Arróyave, “Uncertainty Quantification and Propagation in Computational Materials Science and Simulation-Assisted Materials Design,” *Integrating Materials and Manufacturing Innovation*, vol. 9, pp. 103–143, Mar. 2020.
- [12] B. Kouchmeshky and N. Zabarás, “The effect of multiple sources of uncertainty on the convex hull of material properties of polycrystals,” *Computational Materials Science*, vol. 47, pp. 342–352, Dec. 2009.
- [13] A. Molkeri, D. Khatamsaz, R. Couperthwaite, J. James, R. Arróyave, D. Allaire, and A. Srivastava, “On the importance of microstructure information in materials design: PSP vs PP,” *Acta Materialia*, vol. 223, p. 117471, Jan. 2022.
- [14] F. Rizzi, R. E. Jones, J. A. Templeton, J. T. Ostien, and B. L. Boyce, “Plasticity models of material variability based on uncertainty quantification techniques,” *arXiv:1802.01487 [physics]*, Nov. 2017. arXiv: 1802.01487.
- [15] M. Sanghvi, P. Honarmandi, V. Attari, T. Duong, R. Arroyave, and D. L. Allaire, “Uncertainty propagation via probability measure optimized importance weights with application to parametric materials models,” in *AIAA Scitech 2019 Forum*, 2019.

- [16] L.-M. C. Cruz-Orive, "Particle size-shape distributions: The general spheroid problem: I. Mathematical model," *Journal of Microscopy*, vol. 107, pp. 235–253, Aug. 1976.
- [17] K. Ramesh, "High Strain Rate and Impact Experiments," in *Springer Handbook of Experimental Solid Mechanics*, pp. 929–960, 01 2008.
- [18] J. Field, S. Walley, W. Proud, H. Goldrein, and C. Siviour, "Review of experimental techniques for high rate deformation and shock studies," *International Journal of Impact Engineering*, vol. 30, no. 7, pp. 725–775, 2004.
- [19] G. R. Johnson and W. H. Cook, "A constitutive model and data for materials subjected to large strains, high strain rates, and high temperatures," *Proc. 7th Int. Sympo. Ballistics*, pp. 541–547, Apr. 1983.
- [20] F. J. Zerilli and R. W. Armstrong, "Dislocation-mechanics-based constitutive relations for material dynamics calculations," *Journal of Applied Physics*, vol. 61, no. 5, pp. 1816–1825, 1987.
- [21] H. Kolsky, "An investigation of the mechanical properties of materials at very high rates of loading," *Proceedings of the Physical Society. Section B*, vol. 62, no. 11, pp. 676–700, 1949.
- [22] A. Banerjee, S. Dhar, S. Acharyya, D. Datta, and N. Nayak, "Determination of Johnson Cook Material and Failure Model Constants and Numerical Modelling of Charpy Impact Test of Armour Steel," *Materials Science and Engineering: A*, vol. 640, pp. 200 – 209, 2015.
- [23] M. Smith, *ABAQUS/Standard User's Manual, Version 6.9*. United States: Dassault Systèmes Simulia Corp, 2009.
- [24] Z. Xu and F. Huang, "Plastic Behavior and Constitutive Modeling of Armor Steel over Wide Temperature and Strain Rate Ranges," *Acta Mechanica Solida Sinica*, vol. 25, pp. 598–608, Dec. 2012.

- [25] F. J. Zerilli and R. W. Armstrong, "Description of tantalum deformation behavior by dislocation mechanics based constitutive relations," *Journal of Applied Physics*, vol. 68, no. 4, pp. 1580–1591, 1990.
- [26] T. Mirzaie, H. Mirzadeh, and J.-M. Cabrera, "A simple zerilli–armstrong constitutive equation for modeling and prediction of hot deformation flow stress of steels," *Mechanics of Materials*, vol. 94, pp. 38–45, 2016.
- [27] R. A. Abrahams, "Low Alloy High Performance Steel," 2016. US Patent no. US20160369362A1.
- [28] G. R. Johnson and W. H. Cook, "Fracture characteristics of three metals subjected to various strains, strain rates, temperatures and pressures," *Engineering Fracture Mechanics*, vol. 21, no. 1, pp. 31–48, 1985.
- [29] R. Kwesi Nutor, "Using the Hollomon Model to Predict Strain-Hardening in Metals," *American Journal of Materials Synthesis and Processing*, vol. 2, no. 1, p. 1, 2017.
- [30] E28 Committee, "Test methods for tension testing of metallic materials," vol. 03.01, 2016.
- [31] H. Conrad, "Yielding and Flow of the Body Centered Cubic Metals at Low Temperatures," p. 73, 1963.
- [32] R. L. Winkler, "Combining probability distributions from dependent information sources," *Management Science*, vol. 27, no. 4, pp. 479–488, 1981.
- [33] B. Hayes, "First links in the markov chain," *American Scientist*, p. 7, 2013.
- [34] L. Kaufman and P. J. Rousseeuw, *Finding groups in data: an introduction to cluster analysis*, vol. 344. John Wiley & Sons, 2009.
- [35] M. V. Hopson, C. M. Scott, and D. Lambert, "Applicability of statistical flaw distributions of eglin steel for fracture calculations," *Procedia Engineering*, vol. 103, pp. 213–220, 2015.
- [36] D. E. Knuth, *The Art of Computer Programming, Volume 2 (3rd Ed.): Seminumerical Algorithms*. USA: Addison-Wesley Longman Publishing Co., Inc., 1997.

- [37] C. S. Smith, *A search for structure. Selected essays on science, art and history*. The MIT Press, 1983.
- [38] X. Shi, W. Zeng, Y. Sun, Y. Han, Y. Zhao, and P. Guo, “Microstructure-tensile properties correlation for the ti-6al-4v titanium alloy,” *Journal of Materials Engineering and Performance*, vol. 24, no. 4, pp. 1754–1762, 2015.
- [39] B. L. DeCost, M. D. Hecht, T. Francis, B. A. Webler, Y. N. Picard, and E. A. Holm, “UHCSDB: UltraHigh carbon steel micrograph DataBase: Tools for exploring large heterogeneous microstructure datasets,” *Integrating Materials and Manufacturing Innovation*, vol. 6, no. 2, pp. 197–205, 2017.
- [40] S. R. Bodner and Y. Partom, “Constitutive equations for elastic-viscoplastic strain-hardening materials,” *Journal of Applied Mechanics*, vol. 42, no. 2, p. 385, 1975.
- [41] C. Robert and G. Casella, *Monte Carlo statistical methods*. Springer Science & Business Media, 2013.
- [42] J. H. Panchal, S. R. Kalidindi, and D. L. McDowell, “Key computational modeling issues in integrated computational materials engineering,” *Computer-Aided Design*, vol. 45, no. 1, pp. 4–25, 2013.
- [43] D. L. McDowell and S. R. Kalidindi, “The materials innovation ecosystem: a key enabler for the materials genome initiative,” *Mrs Bulletin*, vol. 41, no. 4, pp. 326–337, 2016.
- [44] R. Arroyave, S. Shields, C.-N. Chang, D. Fowler, R. Malak, and D. Allaire, “Interdisciplinary research on designing engineering material systems: Results from a national science foundation workshop,” *Journal of Mechanical Design*, vol. 140, no. 11, p. 110801, 2018.
- [45] W. Chen, L. Baghdasaryan, T. Buranathiti, and J. Cao, “Model validation via uncertainty propagation and data transformations,” *AIAA journal*, vol. 42, no. 7, pp. 1406–1415, 2004.
- [46] A. Chernatynskiy, S. R. Phillpot, and R. LeSar, “Uncertainty quantification in multiscale simulation of materials: A prospective,” *Annual Review of Materials Research*, vol. 43, pp. 157–182, 2013.

- [47] S. Ghoreishi and D. Allaire, “Adaptive uncertainty propagation for coupled multidisciplinary systems,” *AIAA Journal*, pp. 3940–3950, 2017.
- [48] D. Allaire and K. Willcox, “A mathematical and computational framework for multifidelity design and analysis with computer models,” *International Journal for Uncertainty Quantification*, vol. 4, no. 1, 2014.
- [49] S. H. Lee and W. Chen, “A comparative study of uncertainty propagation methods for black-box-type problems,” *Structural and Multidisciplinary Optimization*, vol. 37, no. 3, p. 239, 2009.
- [50] K. Li and D. Allaire, “A compressed sensing approach to uncertainty propagation for approximately additive functions,” in *ASME 2016 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, pp. V01AT02A027–V01AT02A027, American Society of Mechanical Engineers, 2016.
- [51] S. Amaral, D. Allaire, and K. Willcox, “A decomposition-based approach to uncertainty analysis of feed-forward multicomponent systems,” *International Journal for Numerical Methods in Engineering*, vol. 100, no. 13, pp. 982–1005, 2014.
- [52] L. Kaufman and P. Rousseeuw, “Clustering by means of medoids,” 1987.
- [53] M. D. McKay, R. J. Beckman, and W. J. Conover, “A comparison of three methods for selecting values of input variables in the analysis of output from a computer code,” *Technometrics*, vol. 42, no. 1, pp. 55–61, 2000.
- [54] Y. Rubner, C. Tomasi, and L. J. Guibas, “The Earth Mover’s Distance as a Metric for Image Retrieval,” *International Journal of Computer Vision*, p. 23, 2000.
- [55] F. J. M. Jr., “The kolmogorov-smirnov test for goodness of fit,” *Journal of the American Statistical Association*, vol. 46, no. 253, pp. 68–78, 1951.
- [56] C. Neel, S. Gibbons, R. Abrahams, and J. House, “Shock and spall in the low-alloy steel af9628,” *Journal of Dynamic Behavior of Materials*, vol. 6, no. 1, pp. 64–77, 2020.

- [57] S. Gibbons, R. Abrahams, M. Vaughan, R. Barber, R. Harris, R. Arroyave, and I. Karaman, “Microstructural refinement in an ultra-high strength martensitic steel via equal channel angular pressing,” *Materials Science and Engineering: A*, vol. 725, pp. 57–64, 2018.
- [58] R. Abrahams, “The science of cost-effective materials design-a study in the development of a high strength, impact resistant steel,” in *APS Shock Compression of Condensed Matter Meeting Abstracts*, pp. O7–005, 2017.
- [59] M. Dilmore, “Eglin steel: researchers developed a new high-strength, high-performance, low-cost steel.,” *AFRL Technology Horizons*, vol. 4, no. 4, pp. 48–49, 2003.
- [60] N. Brar, V. Joshi, and B. Harris, “Constitutive model constants for al7075-t651 and al7075-t6,” in *Shock Compression of Condensed Matter - 2009* (M. Elert, W. Buttler, M. Furnish, W. Anderson, and W. Proud, eds.), pp. 945–948, Melville, NY: American Institute of Physics, 2009.
- [61] R. Abrahams, “Low alloy high performance steel,” 2017.
- [62] V. Sinha, E. J. Payton, M. Gonzales, R. A. Abrahams, and B. S. Song, “Delineation of prior austenite grain boundaries in a low-alloy high-performance steel,” *Metallography, Microstructure, and Analysis*, vol. 6, no. 6, pp. 610–618, 2017.
- [63] V. Sinha, M. Gonzales, R. Abrahams, B. Song, and E. Payton, “Correlative microscopy for quantification of prior austenite grain size in AF9628 steel,” *Mater. Charact.*, vol. 159, p. 109835, 2020.
- [64] M. Laine, “MCMC Toolbox for Matlab,” Jan. 2022.
- [65] K. T. Ramesh, “High rates and impact experiments,” in *Springer Handbook of Experimental Solid Mechanics*, pp. 929–960, Springer US, 2008.
- [66] J. McGlaun, S. Thompson, and M. Elrick, “Cth: A three-dimensional shock wave physics code,” *International Journal of Impact Engineering*, vol. 10, no. 1, pp. 351 – 360, 1990.
- [67] S. P. Marsh, *LASL Shock Hugoniot Data*. University of California Press, 1980.

- [68] C. Villani, *Optimal transport: old and new*, vol. 338. Springer, 2009.
- [69] S. Amaral, D. Allaire, and K. Willcox, “Optimal l2-norm empirical importance weights for the change of probability measure,” *Statistics and Computing*, vol. 27, no. 3, pp. 625–643, 2017.
- [70] S. Amaral, D. Allaire, and K. Willcox, “A decomposition approach to uncertainty analysis of multidisciplinary systems,” in *12th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference and 14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, p. 5563, 2012.
- [71] S. Amaral, D. Allaire, E. D. L. R. Blanco, and K. E. Willcox, “A decomposition-based uncertainty quantification approach for environmental impacts of aviation technology and operation,” *AI EDAM*, vol. 31, no. 3, pp. 251–264, 2017.
- [72] P. Billingsley, *Probability and measure*. John Wiley & Sons, 2008.
- [73] Wikipedia contributors, “Kolmogorov–smirnov test — Wikipedia, the free encyclopedia,” 2020. [Online; accessed 15-May-2020].
- [74] F. J. Massey, “The kolmogorov-smirnov test for goodness of fit,” *Journal of the American Statistical Association*, vol. 46, no. 253, pp. 68–78, 1951.
- [75] D. Khatamsaz, A. Molkeri, R. Couperthwaite, J. James, R. Arróyave, D. Allaire, and A. Srivastava, “Efficiently exploiting process-structure-property relationships in material design by multi-information source fusion,” *Acta Materialia*, vol. 206, p. 116619, Mar. 2021.
- [76] M. Ranaiefar, P. Honarmandi, L. Xue, C. Zhang, A. Elwany, I. Karaman, E. J. Schwalbach, and R. Arroyave, “A differential evaporation model to predict chemistry change of additively manufactured metals,” *Materials & Design*, vol. 213, p. 110328, Jan. 2022.
- [77] L. Johnson, M. Mahmoudi, B. Zhang, R. Seede, X. Huang, J. T. Maier, H. J. Maier, I. Karaman, A. Elwany, and R. Arróyave, “Assessing printability maps in additive manufacturing of metal alloys,” *Acta Materialia*, vol. 176, pp. 199–210, Sept. 2019.

- [78] J. Larsen, S. Jha, C. Szczepanski, M. Caton, R. John, A. Rosenberger, D. Buchanan, P. Golden, and J. Jira, “Reducing uncertainty in fatigue life limits of turbine engine alloys,” *International Journal of Fatigue*, vol. 57, pp. 103–112, 2013.
- [79] J. M. Larsen, B. D. Worth, C. G. Annis, and F. K. Haake, “An assessment of the role of near-threshold crack growth in high-cycle-fatigue life prediction of aerospace titanium alloys under turbine engine spectra,” *International Journal of Fracture*, vol. 80, pp. 237–255, Apr. 1996.
- [80] C. Szczepanski, S. Jha, J. Larsen, and J. Jones, “Microstructural Influences on Very-High-Cycle Fatigue-Crack Initiation in Ti-6246,” *Metallurgical and Materials Transactions A*, vol. 39, pp. 2841–2851, Dec. 2008.
- [81] A. L. Pilchak, A. Hutson, W. J. Porter, D. Buchanan, and R. John, “On the cyclic fatigue and dwell fatigue crack growth response of ti-6al-4v,” in *Proceedings of the 13th World Conference on Titanium* (V. Venkatesh, A. L. Pilchak, J. E. Allison, S. Ankem, R. Boyer, J. Christodoulou, H. L. Fraser, M. A. Imam, Y. Kosaka, H. J. Rack, A. Chatterjee, and A. Woodfield, eds.), pp. 993–998, John Wiley & Sons, Inc., 2016.
- [82] D. Norfleet, J. Williams, S. Ghosh, M. Mills, and S. Rokhlin, “2007_vol_2_pres_114.pdf.”
- [83] A. L. Pilchak, J. Shank, J. C. Tucker, S. Srivatsa, P. N. Fagin, and S. L. Semiatin, “A dataset for the development, verification, and validation of microstructure-sensitive process models for near-alpha titanium alloys,” *Integrating Materials and Manufacturing Innovation*, vol. 5, no. 1, pp. 259–276, 2016.
- [84] A. L. Pilchak and J. C. Williams, “Observations of Facet Formation in Near- α Titanium and Comments on the Role of Hydrogen:,” tech. rep., Defense Technical Information Center, Fort Belvoir, VA, Oct. 2010.
- [85] W. Evans and M. Bache, “Dwell-sensitive fatigue under biaxial loads in the near-alpha titanium alloy IMI685,” *International Journal of Fatigue*, vol. 16, no. 7, pp. 443–452, 1994.

- [86] M. Cuddihy, A. Stapleton, S. Williams, and F. Dunne, "On cold dwell facet fatigue in titanium alloy aero-engine components," *International Journal of Fatigue*, vol. 97, pp. 177–189, Apr. 2017.
- [87] S. Jha, J. Larsen, and A. Rosenberger, "Towards a physics-based description of fatigue variability behavior in probabilistic life-prediction," *Engineering Fracture Mechanics*, vol. 76, no. 5, pp. 681–694, 2009.
- [88] S. K. Jha, J. C. Tucker, J. M. Larsen, R. John, and A. L. Pilchak, "Probabilistic dwell fatigue life prediction in microtextured ti-6al-4v." The Minerals, Metals, & Materials Society Annual Meeting.
- [89] L. M. Cruz-Orive, "Stereology: A Historical Survey," *Image Analysis & Stereology*, vol. 36, p. 153, Dec. 2017.
- [90] L. Cruz Orive, "Particle size-shape distributions: The general spheroid problem," *Journal of Microscopy*, vol. 112, pp. 153–167, 1978.
- [91] M. Baaske, F. Ballani, and A. Illgen, "On the estimation of parameters of a spheroid distribution from planar sections," *Spatial Statistics*, vol. 26, pp. 83–100, 2018.
- [92] K. C. G. Chan and J. Qin, "Nonparametric maximum likelihood estimation for the multi-sample wicksell corpuscle problem," *Biometrika*, vol. 103, no. 2, pp. 273–286, 2016.
- [93] W. L. Nicholson and K. R. Merckx, "Unfolding Particle Size Distributions," *Technometrics*, vol. 11, pp. 707–723, Nov. 1969.
- [94] Wicksell, "The Corpuscle Problem: A Mathematical Study of a Biometric Problem," *Biometrika*, p. 17, 1925.
- [95] E. J. Payton, "Revisiting sphere unfolding relationships for the stereological analysis of segmented digital microstructure images," *Journal of Minerals and Materials Characterization and Engineering*, vol. 11, no. 3, pp. 221–242, 2012.
- [96] P. R. Mouton, *Unbiased Stereology: A Concise Guide*. JHU Press, 2011.

- [97] L. M. Karlsson and A. M. Gokhale, “Estimating Mean Linear Intercept Length Using the Trisector,” *Acta Stereol*, pp. 53–58, 1996.
- [98] Wicksell, “The Corpuscle Problem. Second Memoir. Case of Ellipsoidal Corpuscles,” 1926.
- [99] A. L. Pilchak, “A simple model to account for the role of microtexture on fatigue and dwell fatigue lifetimes of titanium alloys,” *Scripta Materialia*, vol. 74, pp. 68–71, 2014.
- [100] C. B. Do and S. Batzoglou, “What is the expectation maximization algorithm?,” *Nature Biotechnology*, vol. 26, pp. 897–899, Aug. 2008.
- [101] H. Hartley, “Maximum Likelihood Estimation from Incomplete Data,” *Biometrics*, p. 22, 1958.
- [102] A. P. Dempster, N. M. Laird, and D. B. Rubin, “Maximum likelihood from incomplete data via the *EM* algorithm,” *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 39, no. 1, pp. 1–22, 1977.
- [103] B. W. Silverman, M. C. Jones, J. D. Wilson, and D. W. Nychka, “A Smoothed Em Approach to Indirect Estimation Problems, with Particular Reference to Stereology and Emission Tomography,” *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 52, pp. 271–303, Jan. 1990.
- [104] Y. Vardi, L. A. Shepp, and L. Kaufman, “A Statistical Model for Positron Emission Tomography,” *Journal of the American Statistical Association*, vol. 80, pp. 8–20, Mar. 1985.
- [105] M. A. Groeber and M. A. Jackson, “DREAM.3d: A digital representation environment for the analysis of microstructure in 3d,” *Integrating Materials and Manufacturing Innovation*, vol. 3, no. 1, pp. 56–72, 2014.
- [106] D. Freedman and P. Diaconis, “On the histogram as a density estimator:L 2 theory,” *Zeitschrift fr Wahrscheinlichkeitstheorie und Verwandte Gebiete*, vol. 57, pp. 453–476, Dec. 1981.

- [107] V. Sinha, M. J. Mills, J. C. Williams, and J. E. Spowart, "Observations on the faceted initiation site in the dwell-fatigue tested ti-6242 alloy: Crystallographic orientation and size effects," *Metallurgical and Materials Transactions A*, vol. 37, pp. 1507–1518, May 2006.
- [108] L. Christodoulou and J. M. Larsen, "Using materials prognosis to maximize the utilization potential of complex mechanical systems," *JOM*, vol. 56, pp. 15–19, Mar. 2004.