Bayesian calibration of strength parameters using hydrocode simulations of symmetric impact shock experiments of Al-5083

David J. Walters (iD), Ayan Biswas, Earl C. Lawrence, Devin C. Francom, Darby J. Luscher (iD), D. Anthony Fredenburg, Kelly R. Moran, Christine M. Sweeney (iD), Richard L. Sandberg, James P. Ahrens, and C. A. Bolme (iD)

View Online   Export Citation   CrossMark

ARTICLES YOU MAY BE INTERESTED IN

AIP Publishing

Bayesian calibration of strength parameters using hydrocode simulations of symmetric impact shock experiments of Al-5083

David J. Walters,[1,a)] Ayan Biswas,[2] Earl C. Lawrence,[2] Devin C. Francom,[2]
Darby J. Luscher,[1] D. Anthony Fredenburg,[3] Kelly R. Moran,[4] Christine M. Sweeney,[2]
Richard L. Sandberg,[5] James P. Ahrens,[2] and C. A. Bolme[6]

[1]*Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*
[2]*Computer, Computational, and Statistical Sciences Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*
[3]*Theoretical Design Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*
[4]*Department of Statistical Sciences, Duke University, Durham, North Carolina 27708, USA*
[5]*Center for Integrated Nanotechnologies, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*
[6]*Shock and Detonation Physics, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*

Predictive modeling of materials requires accurately parameterized constitutive models. Parameterizing models that describe dynamic strength and plasticity require experimentally probing materials in a variety of strain rate regimes. Some experimental protocols (e.g., plate impact) probe the constitutive response of a material using indirect measures such as free surface velocimetry. Manual efforts to parameterize constitutive models using indirect experimental measures often lead to non-unique optimizations without quantification of parameter uncertainty. This study uses a Bayesian statistical approach to find model parameters and to quantify the uncertainty of the resulting parameters. The technique is demonstrated by parameterizing the Johnson-Cook strength model for aluminum alloy 5083 by coupling hydrocode simulations and velocimetry measurements of a series of plate impact experiments. Simulation inputs and outputs are used to calibrate an emulator that mimics the outputs of the computationally intensive simulations. Varying the amount of experimental data available for emulator calibration showed clear differences in the degree of uncertainty and uniqueness of the resulting optimized Johnson-Cook parameters for Al-5083. The results of the optimization provided a numerical evaluation of the degree of confidence in model parameters and model performance. Given an understanding of the physical effects of certain model parameters, individual parameter uncertainty can be leveraged to quickly identify gaps in the physical domains covered by completed experiments. *Published by AIP Publishing.* https://doi.org/10.1063/1.5051442

I. INTRODUCTION

Many industries (e.g., aerospace, automotive, defense) benefit from the ability to accurately predict the high-rate loading behavior of materials. Understanding this response of materials has been the focus of continuous research for several decades. Physical experimentation with materials was one of the first forms of exploration of materials' properties that continues today and will carry on into the future. The ability to predict the constitutive response of a material in a simple experiment, combined with numerical modeling and simulation, enables estimating the response of more complex systems in which physical experimentation is limited or impossible.

Inelastic deformation modes often dominate the deviatoric response during shock and impact loading of materials. Strength and plasticity models which are used to represent much of this behavior commonly employ a numerical approximation of the flow stress of the material. Many models[1–3] for the flow stress, Y_f, used for approximating dynamic plastic deformation are expressed as a function of the accumulated

effective plastic strain, ε_p; plastic strain rate, $\dot{\varepsilon}_p$; temperature, T; and model specific parameters, $\mathcal{P}$, i.e.,

$$Y_f = \hat{Y}_f(\varepsilon_p, \dot{\varepsilon}_p, T; \mathcal{P}).\qquad(1)$$

When such parameters are not known directly from physical principles, they may be determined by comparing the predicted response of the material under controlled experimental conditions to the corresponding measurement from the experiment. For example, a fit to measured stress versus strain data can be obtained by comparing predictions of the stress state given a set of physical conditions, $\sigma_{calc}(\varepsilon_p, \dot{\varepsilon}_p, T; \mathcal{P})$ to those observed experimentally, $\sigma_{exp}(\varepsilon_p, \dot{\varepsilon}_p, T)$. Convenient error norms employed in the materials science community include the mean relative difference, i.e.,[4]

$$\delta = \frac{1}{N}\sum_{i=1}^{N}\frac{\left|\sigma_{calc}\left(\{\varepsilon_p, \dot{\varepsilon}_p, T\}_i; \mathcal{P}\right) - \sigma_{exp}\left(\{\varepsilon_p, \dot{\varepsilon}_p, T\}_i\right)\right|}{\sigma_{exp}\left(\{\varepsilon_p, \dot{\varepsilon}_p, T\}_i\right)},\quad(2)$$

or the square root of the sum of square errors, i.e.,[5,6]

$$\delta = \left\{\sum_{i=1}^{N}\left[\sigma_{calc}(\varepsilon_i, \ldots) - \sigma_{exp}(\varepsilon_i, \ldots)\right]^2\right\}^{1/2}\quad(3)$$

---
a)Author to whom correspondence should be addressed: djwalters@lanl.gov

such that the optimal parameter set, $\mathcal{P}^*$, minimizes the estimated error norm over the set of experiments, i.e., $\frac{\partial \delta}{\partial \mathcal{P}} = 0$.

Empirically-derived models such as the Johnson-Cook model[1] or the Zerilli-Armstrong model[2] receive widespread application due to their relatively few parameters and broad range of applicability. However, these models have difficulty covering a large range of conditions (ε_p, $\dot{\varepsilon}_p$, T) with a single set of unique parameters for any given material.[4] More sophisticated models with the aim of covering a larger physical range (e.g., MTS,[7–10] BCJ,[11–13] PTW,[3,14] and others[15–17]) have the consequence of employing more parameters than those previously cited. Parameterizing these more complex models would only become marginally more difficult if direct comparisons between stress and strain could be made with a set of experiments.

While the material state variables, stress, strain, and temperature, as well as strain rate, can be measured and/or controlled during quasistatic [$\dot{\varepsilon} = O(10^{-8} - 10^0)\text{s}^{-1}$] or split Hopkinson pressure bar (SHPB) dynamic experiments [$\dot{\varepsilon} = O(10^2 - 10^4)\text{s}^{-1}$], such measurements cannot be directly obtained from experiments that characterize higher strain-rate deformation, such as plate impact experiments [$\dot{\varepsilon} \, O(10^6)\text{s}^{-1}$]. Instead, simulations are used to estimate observable quantities such as free surface velocity measured using VISAR[18] or lattice strain using X-ray diffraction[19] to infer the internal state of a material. The computational cost of running simulations generally increases with the implementation of more complex material models and with application to more complex experimental configurations. For example, Winey and Gupta[20,21] used velocimetry measurements to calibrate models for the anisotropic shock response of single crystal LiF and PETN, Addessio *et al.*[22] calibrated a model for α to γ phase transformation in RDX, and Versino and Bronkhorst[23] used velocimetry to estimate parameters for a micromechanics model of ductile spall in tantalum. Prime *et al.*[24] developed an approach using hydrodynamic calculations of Richtmyer-Meshkov instability to assess the flow stress at high strain rates in polycrystalline copper. In these and the broader literature, very little description is provided on the process of estimating of parameters. Presumably, the parameters are often estimated by manually adjusting them until they produce an acceptable visual match. Brown and Hund[25] is a notable exception from the statistics literature. They follow a procedure similar to the one that we use, although they handle the functional nature of the output in a different manner. Ali *et al.*[26] also used a more rigorous approach to estimate the uncertainty in equation of state parameters by employing iterative forward analysis. They used both Monte Carlo and perturbation methods to propagate the experimental uncertainties into the equation of state parameters.[26]

An increasing number of experiments and/or simulations are required to explore the full domain of parameters as the number of model parameters increases. Coupled with increasingly challenging simulations, finding optimized parameters can be a slow and tedious endeavor. The likelihood of finding non-unique optimizations also increases when more parameters are involved and the experimental data are limited. Additionally, optimization routines that directly minimize

the error norms highlighted above do not provide a probability distribution of the parameter space characterizing the uncertainty of the model inputs or outputs.

Bayesian model calibration[27,28] (BMC) provides a solution. This rigorous statistical approach provides a framework to compare a computationally intensive physics model with experimental data in order to estimate the input parameters that make the model best match the data. In addition to estimates of the best parameters, the procedure also provides uncertainty about these estimates in terms of a probability distribution. Further, the procedure can be made to account for estimates, with uncertainty, of systematic discrepancies between the computational physics model and experimental measurement. The approach has been successful in application areas including cosmology[29] and nuclear density functional theory.[30] The key ingredient of this approach is an *emulator*, a statistical model that attempts to mimic the computationally intensive physics model, but much faster. Based on a training set, the emulator gives predictions, with uncertainty, for the physics model at new inputs in a fraction of a second, whereas the physics model may take minutes, hours, or even weeks. The fast predictor can then be used as part of a Bayesian estimation scheme to give the desired results.

Figure 1 provides a generic illustration of the approach with a toy example having a single data point and a single unknown parameter. Experimental data with a value of 0.8 is observed and shown as a horizontal line with error bars shown around it in light gray. The data are compared to our physics model, which is shown as the sigmoid curve to determine which value of the input θ (*x*-axis) calibrates the model to match the data. However, the physics model is computationally expensive, so it can only be evaluated at four values of θ. An interpolating regression model is fit through those points that can be used to predict the physics model at all
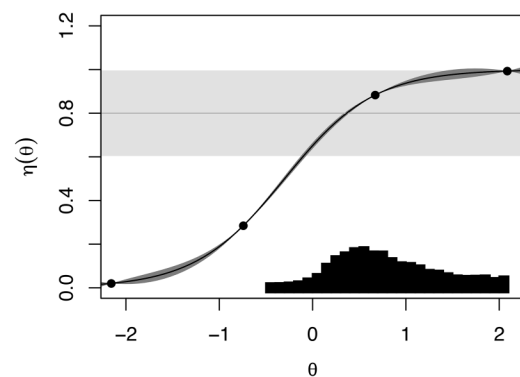

FIG. 1. Illustration of model calibration with a toy example. Experimental data with a value of 0.8 is shown as a horizontal line with error bars displayed around it in light gray. The physics model shown as the sigmoid curve is used to determine which value of the input θ (*x*-axis) calibrates the model to match the data. The computationally expensive physics model is evaluated at four values of θ. An interpolating regression model is fit through those points to predict the physics model at all other values of θ. This regression model, called an *emulator* can be evaluated quickly enough to explore the range of the input. It also returns an estimate of its own uncertainty about the predictions for the physics model, shown by the bulges around the sigmoid between the physics model evaluations. The emulator compares with the experimental data to estimate a distribution for θ that produces probabilistically reasonable matches, shown as the histogram on the *x*-axis.

TABLE I. Impact experiment parameters and results as reported by Boteler and Dandekar.[31]

Experimental data			Elastic compression			Final compression				
Shot no.	Flyer/target (thickness, mm)	V_{impact} (km/s)	σ_{HEL} (GPa)	u_e (km/s)	ρ_e (g/cm^3)	U_s (km/s)	σ_{final} (GPa)	u_p (km/s)	ρ (g/cm^3)	$\frac{\rho_0}{\rho}$
0104S	3.950/5.965	0.2040	0.6254	0.0360	2.683	5.452	1.584	0.1020	2.7159	0.9823
0105S	3.957/9.940	0.3549	0.5844	0.0337	2.682	5.510	2.697	0.1775	2.7542	0.9687
0106S	3.969/5.961	0.4838	0.6429	0.0370	2.683	5.718	3.766	0.2419	2.7837	0.9584

other values of θ. This regression model, which is called an *emulator* can be evaluated quickly enough to explore the range of the input. It also returns estimates of its own uncertainty about the predictions for the physics model. These are shown by the bubbles that bulge out around the sigmoid between the physics model evaluations. The emulator can now be compared with the experimental data to estimate a distribution for θ that produces probabilistically reasonable matches. This distribution is shown as the histogram on the *x*-axis.

This approach has several advantages over the mostly hand-tuned velocimetry-matching work cited earlier. First, the procedure is grounded in statistical principles, so the resulting uncertainties have meaningful statistical interpretations. Second, the procedure is fully automated. Third, it is easy to incorporate multiple observations. Finally, parameter estimates are given with uncertainty and correlation, rather than single valued best fits.

In this paper, we demonstrate the utility of Bayesian model calibration by applying it to a case study using aluminum symmetric plate impact data from Boteler and Dandekar[31] and the Johnson-Cook plasticity model.[1] Each experiment is simulated many times over a range of Johnson-Cook model parameters. These simulations are used to train the emulator, which is compared with data to produce distributions of plausible parameter values and predictions.

The Johnson-Cook model was selected for this purpose due to its relative simplicity and wide use in the materials science community.[5,32–35] Other more physically based models such as PTW[3] and those put forth by McDowell[15] and Luscher *et al.*[17] would be well suited for this analysis but are unnecessarily complicated for the demonstration presented here.

Section II presents the relevant details of the experiments and the Johnson-Cook calculations. Section III gives an overview of our statistical approach (further detail is provided in Appendixes A and B). The results of this procedure are shown in Sec. IV. Finally, we conclude with a discussion of the performance of our approach and its application in other settings in Sec. V.

II. EXPERIMENT AND SIMULATION DETAILS

A. Boteler and Dandekar experimental data

A series of transmission impact experiments were conducted by Boteler and Dandekar[31] to investigate the Hugoniot Equation of State (EOS) and the Hugoniot Elastic Limit (HEL) of Al-5083. Using a 102 mm bore single-stage

light gas-gun, impact pressures between 1.5 and 8.0 GPa were achieved in symmetric and asymmetric normal impacts (within 1 mrad of arc).[31] Of the eight transmission impact tests performed by Boteler and Dandekar,[31] we selected three shots to calibrate the Johnson-Cook parameters for Al-5083. The three shots chosen were all symmetric impact tests; thus involving only one material for characterization rather than introducing the uncertainty of a second material.

Shots 0104S, 0105S, and 0106S from Boteler and Dandekar[31] exhibited a range of specimen thicknesses and impact velocities. All samples were 38 mm in diameter.[31] A summary of the shot characteristics of the chosen specimens are reproduced in Table I to include values such as: flyer and target thicknesses; flyer velocity at impact (V_{impact}); stress, particle velocity, and density at the HEL (σ_{HEL}, u_e, and ρ_e, respectively); and the shock velocity, stress, particle velocity, and density at the fully compressed state (U_s, σ_{final}, u_p, and ρ, respectively). The velocimetry traces were recorded using VISAR interferometry[18] (with reported 1% precision over the entire range of compression) and are shown in Fig. 2. Charged pins were used to obtain V_{impact} with an uncertainty of less than 2%.[31] The VISAR traces were recorded with 0.5 ns resolution providing high temporal resolution for the Johnson-Cook parameter optimization.

B. Hydrocode simulations

Data for training the emulator to compute optimized parameters and their uncertainty is provided by a series of hydrocode simulations. The experimental shots taken by Boteler and Dandekar[31] are simulated using a multiphysics continuum hydrodynamics code, FLAG. FLAG is a research code developed by Los Alamos National Laboratory to
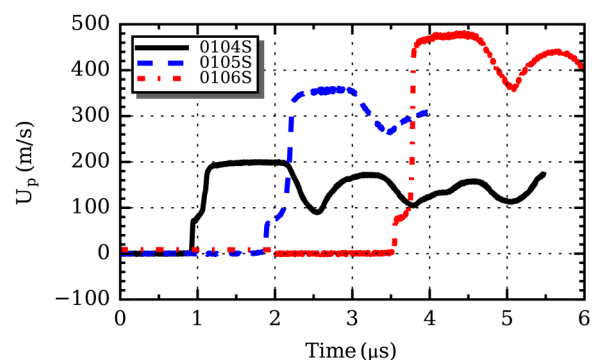
FIG. 2. Experimental velocimetry profiles of symmetric impact shots 0104S, 0105S, and 0106S from Boteler and Dandekar.[31]

simulate high strain-rate and large deformation response of materials and components.[36–38] The following will discuss details regarding material constitutive models, simulation and solver details, and selection of parameterized variables.

1. Constitutive models

All experiments conducted by Boteler and Dandekar[31] that are used here involve aluminum alloy 5083 with strain hardening treatment designated H131. The volumetric response of the material in the simulation is provided by a Mie-Grüneisen equation of state, while the deviatoric response is modeled by the aforementioned Johnson-Cook strength model. Each model is described in detail below.

a. Equation of state. The equation of state for compression is described using a Mie-Grüneisen form in which the reference cold curve is defined using a cubic function of compression ratio. This results in a polynomial equation which describes the pressure, P, in a material as a function of density, ρ, and specific internal energy, E_m, according to

$$P = c_1\mu + c_2\mu^2 + c_3\mu^3 + \left(\gamma_0 + \gamma_1\mu + \gamma_2\mu^2\right)\rho_0 E_m. \quad (4)$$

Here, $\mu = \rho/\rho_0 - 1$ relates the current density to the reference density, ρ_0, c_i are the cold curve pressure coefficients, and γ_i are coefficients defining a quadratic variation of the Grüneisen parameter with compression; over the range of pressures expected, these linear and quadratic coefficients are assumed zero, $\gamma_1 = \gamma_2 = 0$. The remaining terms are summarized in Table II.

b. The Johnson-Cook plasticity model. The Johnson-Cook model for describing dynamic strength and plasticity is attractive for this demonstration due to the relatively small number of parameters needed (5) to fit the model. The flow stress is defined as[1]

$$Y_f = \left(A + B\varepsilon_p^n\right)\left[1 + C\ln\frac{\dot{\varepsilon}_p}{\dot{\varepsilon}_0}\right]\left[1 - (T^*)^m\right]. \quad (5)$$

In this equation, A is the initial yield stress and B and n control strain hardening effects related to the equivalent plastic strain, ε_p. The dependence of flow stress on plastic strain rate, $\dot{\varepsilon}_p$, is affected by parameter C and related to a nominal strain rate set here as, $\dot{\varepsilon}_0 = 1.0\,\mathrm{s}^{-1}$. Thermal softening behavior is

described in terms of the homologous temperature T^* defined (by Johnson and Cook[1]) as

$$T^* = \frac{T - T_{\mathrm{ref}}}{T_{\mathrm{melt}} - T_{\mathrm{ref}}} \quad (6)$$

and the exponent m, where T, T_{ref}, and T_{melt} are the current, reference, and melt temperatures of the material.

Parameterizing the Johnson-Cook model can be accomplished by tuning the 5 parameters described above to bring simulation results into agreement with a set of dynamic experiments (e.g., gas-gun impact, Split Hopkinson bar pressure tests, etc.) conducted on the desired material. Previous parameterization efforts of this model on Al-5083 were conducted by Gray *et al.*[4] for a series of low [$\dot{\varepsilon} = O(10^{-8} - 10^0)\,\mathrm{s}^{-1}$] and intermediate [$\dot{\varepsilon} = O(10^2 - 10^4)\,\mathrm{s}^{-1}$] strain rates. Gray *et al.*[4] published four different sets of optimized parameters as shown in Table III which resulted from different restrictions on the experimental data set as indicated in the table. These sets of parameters were discussed qualitatively, but as with many parameterizations currently in the literature, numerical estimates of the uncertainty were not provided. The results from the previous parameterization are compared to the parameters generated in this work by Bayesian calibration to gas-gun plate impact experiments at much higher strain rates conducted by Boteler and Dandekar.[31] The identified model parameters could vary depending on the range of conditions examined as demonstrated by the variability of model parameters identified by Gray *et al.*[4] for Al-5083 listed in Table III.

2. Simulation setup

The simulations use a one dimensional Lagrangian mesh with a contact slideline interface between the aluminum flyer plate and the aluminum target. Since each experimental shot has different flyer and target thicknesses (see Table I), a consistent element size of approximately $4\,\mu\mathrm{m}$ per cell is chosen to achieve a balance between adequate resolution of the shock wave and computation time. The simulations are run for a total of $4.0\,\mu\mathrm{s}$ (enough time to reach the saturation velocity) with an automatic time stepping routine and a specified initial time

TABLE II. Fixed parameters for the equation of state described in Eq. (4).

Parameters	Nominal value
c_1 (MBar)[a]	0.7797
c_2 (MBar)[a]	0.2239
c_3 (MBar)[a]	1.668
γ_0 (-)[b]	2.13
ρ_0 (g cm^{-3})	2.668
T_{ref} (K)	298
c_v (J kg^{-1} K^{-1})	900

[a]Reference 39.
[b]Reference 40.

TABLE III. Parameterization of the Johnson-Cook model by Gray *et al.*[4] for Al-5083 subjected to strain rates between 0.001 and $7000\,\mathrm{s}^{-1}$ and temperatures between 77 and 473 K.

Parameters	All $\dot{\varepsilon}$, all T[a]	High $\dot{\varepsilon}$, all T[b]	All $\dot{\varepsilon}$, $T > $ RT[c]	High $\dot{\varepsilon}$, $T > $ RT[d]
A (MPa)	210	200	270	170
B (MPa)	620	600	470	425
C	0.0125	0.0200	0.0105	0.0335
n	0.375	0.380	0.600	0.420
m	1.525	1.500	1.200	1.225

[a]Full range of strain rates and temperatures, T^* uses $T_{\mathrm{ref}} = 0$ K.
[b]Strain rates $> 2000\,\mathrm{s}^{-1}$ and the full range of temperatures, T^* uses $T_{\mathrm{ref}} = 0$ K.
[c]Full range of strain rates and temperatures above reference temperature, T^* uses $T_{\mathrm{ref}} = 298$ K.
[d]Strain rates $> 2000\,\mathrm{s}^{-1}$ and temperatures above reference temperature, T^* uses $T_{\mathrm{ref}} = 298$ K.

step size of $\Delta t_{\text{init}} = 1 \times 10^{-7} \mu$s and a maximum step size of $\Delta t_{\text{max}} = 2 \times 10^{-3} \mu$s. The flyer velocity observed in the experimental data is applied as an initial condition to the region associated with the aluminum flyer and positioned with no initial gap between the flyer and the target. Due to the $\pm 2\%$ uncertainty reported by Boteler and Dandekar,[31] the initial flyer velocity is included as one of the variable parameters for Bayesian calibration.

Artificial viscosity is also an important factor enabling numerical simulation of shocked materials. Artificial viscosity acts to smear the shock front resulting from an impact across multiple cells within a model. Without it, only a single cell at a time would be impacted by the shockwave inducing a discontinuous change in internal state (e.g., pressure, density, etc.). This often causes an individual cell to collapse completely in one timestep due to the discontinuous nature of the change induced to that cell—therefore, causing the simulation to crash.

To alleviate this problem, the two term Barton artificial viscosity[41] is employed with constant parameters for the linear term, $q_1 = 0.10$, and the quadratic term, $q_2 = 1.0$. These parameters were optimized manually to produce the steepest (most realistic) shockwave profiles in the velocimetry traces while damping out overshoot and ringing behavior of the particle velocity at the edge of the plastic plateau (cf. Fig. 4). These parameters remained fixed for all simulations conducted in this study.

3. Selection of variables to include in parameterization

Uncertainty permeates several aspects of both the model and the experimental conditions associated with measurements. Indeed, there is uncertainty in the appropriate equation of state and strength model to use, and the corresponding parameters of each. In order to maintain a simple and narrow focus for this effort, we do not consider uncertainty (or parameter variations) in all possible aspects of the model. For example, although there is (certainly) uncertainty in the parameters for the equation of state,[25] they were held fixed in this work. We primarily restrict our attention to the variation of model parameters for the Johnson-Cook strength model, although there are a few additionally included variables noted here. The statistical approach we use (described in Sec. III) requires a simulation suite spread over plausible ranges of these parameters. Here, we describe those parameters and their ranges.

The shear modulus, G, provides a coupling between the volumetric and deviatoric behavior. Since it physically depends upon the density and temperature of the material,[42] the shear modulus increases and decreases with increases in density and temperature, respectively.

The implementation of the Johnson-Cook model in FLAG is associated with a fixed, constant shear modulus, i.e., $G = G_0$ is independent of density and temperature. Because the velocity of a longitudinal wave, e.g., the elastic precursor, depends upon the shear modulus, neglecting its density (and consequently pressure) dependence will result in incorrect wave arrival times for simulations that span multiple impact velocities since each results in a different compressed density or shock pressure. Therefore, we employ a constant shear modulus within the simulations of a particular shot, but allow shots at higher shock pressures to have a larger constant value as described below.

Additionally, Boteler and Dandekar[31] describe a $\pm 2\%$ uncertainty of the flyer velocity at impact. The impact velocity significantly affects the final velocity of the target free surface. Therefore, the flyer velocity for each shot was varied around the corresponding nominal values to include this source of experimental uncertainty.

Table IV shows the selected ranges of parameters varied for the Bayesian calibration. Parameters are sampled between these ranges 1000 times according to a Latin hypercube design. Each of the three shots (S104S, S105S, and S106S) utilized this set of 1000 parameter realizations while using the fixed equation of state and setup parameters described previously. In Table IV, simulations for shot S104S utilize the variables A, B, C, n, m, V_1, and G_1, S105S utilize the variables A, B, C, n, m, V_2, G_1, and Δ_2, and S106S utilize the variables A, B, C, n, m, V_3, G_1, Δ_2, and Δ_3. In order to keep the basis between the sets of simulations the same, the average shear modulus for shot S105S is computed from $G_2 = G_1 + \Delta_2$ and for shot S106S $G_3 = G_2 + \Delta_3 = G_1 + \Delta_2 + \Delta_3$.

The size of the simulation suite, 1000 runs, is chosen to provide sufficient coverage of the parameter space. This number greatly exceeds the rule-of-thumb number of ten points per input dimension.[43] This rule-of-thumb makes assumptions about smoothness and effect sparsity that are difficult to verify, so we attempted to make a conservative choice by choosing a much larger number. The cross-validation results shown in Fig. 6 suggest that our choice was more than sufficient.

Free surface velocity time history results of the 1000 simulations along with the corresponding experiment are shown in Figs. 11–13 (sims in blue, experiment in red). The experimental results have each been shifted in time to approximately line up with the simulations. The shift is necessary due to unknown timing fidelity provided by the experimentalists in Boteler and Dandekar.[31] The equation of state may also be a small source of uncertainty. However, as will be described later, absolute timing accuracy for the feature finding routines is unnecessary; only relative timing between features is important.

TABLE IV. Variable parameter ranges utilized for Bayesian calibration. The first set of variables refer the Johnson-Cook strength model, the middle set refer to the experimental uncertainty of impact velocity, and the last set refer to changes in the average shear modulus across each shot.

Parameters	Min. value	Max. value
A (MBar)	0.0005	0.01
B (MBar)	0.0005	0.01
C (-)	0.0	0.03
n (-)	0	1.5
m (-)	0	3
V_1 (cm μs^{-1})	0.0192	0.0216
V_2 (cm μs^{-1})	0.0334	0.0376
V_3 (cm μs^{-1})	0.0455	0.0513
G_1 (MBar)	0.2	0.5
Δ_2 (MBar)	0	0.1
Δ_3 (MBar)	0	0.2

III. BAYESIAN MODEL CALIBRATION

A. Bayesian statistics

Bayesian statistics is an approach to statistics in which all uncertainty is captured with a probability distribution. In effect, this means that anything with uncertainty is treated as a random variable whose distribution we attempt to estimate. Under the Bayesian framework, a statistical model is built from two components. One component is called the likelihood, denoted $f(y \mid \theta)$ and it describes the distribution of the data, y, conditional on the parameters, θ. It is important to note that "data" in this context is an output observable resulting from either a measurement (e.g., experimental data) or a model (e.g., simulation result). The second component, called the prior distribution, denoted $\pi(\theta)$, describes initial beliefs about the marginal distribution of the unknown parameter. Prior distributions can be used to represent information based on past experiments and scientific intuition, or be relatively uninformative, containing little more than an acceptable range of parameters and necessary physical constraints. Following the rules of probability, the two components are multiplied together to give a joint distribution for unknowns and observables, which can then be conditioned on the observations to give a distribution for the unknowns given data,

$$p(\theta \mid y) \propto f(y \mid \theta)\pi(\theta). \tag{7}$$

This distribution, called the posterior distribution, can be sampled using Markov chain Monte Carlo[44,45] (MCMC). MCMC is a procedure that produces a correlated series of samples from a distribution. It is a powerful technique that can be used for any distribution, although it may not always be efficient. It is also inherently serial and may require many evaluations of the posterior distribution. The resulting sample can be used to estimate quantities of interest like the expected value and quantiles.

B. Model calibration

Here, we try to provide an introduction and a high-level view of the approach. More detail is provided in Appendixes A and B and a thorough description of the BMC approach is given by Higdon *et al.*[28]

Let $M(\theta)$ represent our implementation of the Johnson-Cook model with input parameters $\theta = (A, B, C, m, n, V_1, V_2, V_3, G_1, \Delta_2, \Delta_3)$. Let $y_i, \forall \ i \in \{104, 105, 106\}$ be the velocimetry profile, or a vector of extracted features, from each experiment. We will assume that the experimental data for each shot follows a multivariate Gaussian distribution with a mean vector given by the prediction from FLAG at some best value of θ,

$$y_i \sim N[M(\theta^*), \Sigma], \tag{8}$$

where Σ is a possibly unknown covariance matrix. In our case, we will assume that Σ is a diagonal matrix where the standard deviations are fixed at 1% of the observations, a value based on typical VISAR behavior.[18] In principle, prior distributions could be assigned to the unknown parameters and the posterior sampled with MCMC evaluating (or calling) the hydrocode simulations directly. However, at the resolution

used for these simulations, each computation requires on the order of several minutes, while an acceptable MCMC sample may require as many as 1×10^6 calls to the hydrocode (standard single site Metropolis updates would require a call for every parameter in every MCMC iteration, thus 100 000 MCMC samples for a model with 10 parameters would require 1×10^6 hydrocode evaluations). Thus, the direct evaluation of the hydrocode simulation for each MCMC sample is computationally intractable.

Instead, we expand the scope of the statistical inference to include direct FLAG hydrocode simulations. In addition, to estimating the unknown parameters, we also estimate the output of FLAG over a large region of its input space. To do this, we supplement the set of observables with a carefully selected set of FLAG input-output pairs. Our statistical model for this part of the inference is based on a Gaussian process[46] (GP). A GP can be viewed as a distribution on functions. When points along a function are observed, we can estimate the predictive distribution for other points along the function.

Figure 3 illustrates the basic idea of a Gaussian process and its use in this context. The top panel shows random draws from a Gaussian process with one input dimension. The smoothness of these curves is controlled by the statistical parameters of the distribution. The middle panel shows three points that we take as observations of the unknown function represented by the curve. The bottom panel shows draws of a Gaussian process that are conditioned to pass through those observations.

Let z be a function of some d-dimensional vector of inputs x. A Gaussian process says that any collection of $\vec{z}$ follows a mean-zero multivariate Gaussian distribution with a highly structured covariance matrix that depends on the values of x. There are many possible choices, but we use $\Sigma = \sigma^2 R(X; \vec{\beta})$, where

$$R_{i,j} = \prod_{k=1}^{d} \rho_k^{4(\theta_{i,k} - \theta_{j,k})^2}. \tag{9}$$

Thus, the correlation between any z_i and z_j is a function of the weighted distance between inputs θ_i and θ_j. As the
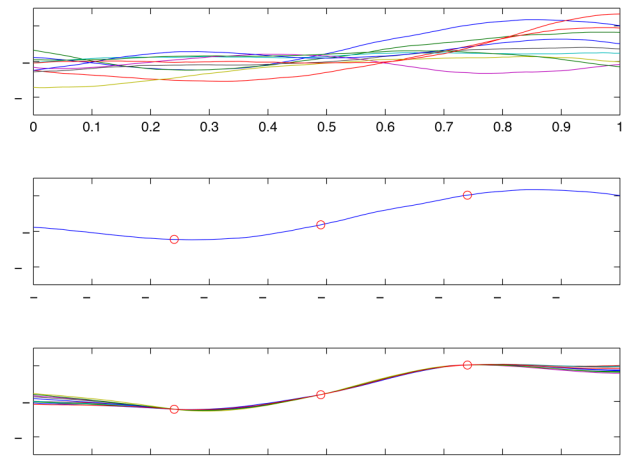

FIG. 3. Top: Unconditional draws from a Gaussian process. Center: Observed points along a potential function. Bottom: Gaussian process draws conditioned on the observed points.
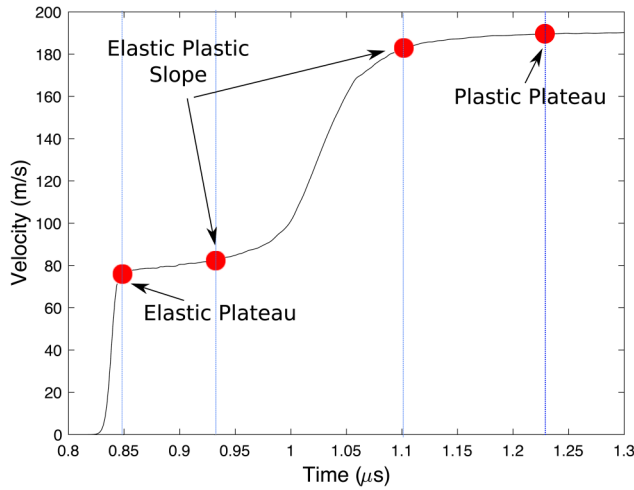
FIG. 4. The four features extracted from each curve for the purposes of analysis. They are the height of the elastic and plastic plateaus and two points along the curve between these plateaus that define the slope of this part of the curve. These are computed from the simulations and experiments for all three shots.

distance goes to zero, the correlation goes to one and z_i and z_j become identical. As the distance goes to infinity, z_i and z_j become completely uncorrelated.

Of course, the output of a computationally intensive physics model is often multivariate as it is in our case. For such cases, we first use the training runs to compute an orthogonal basis using a singular value decomposition (the result is basically the same as principal component directions). The training runs are projected onto this basis and the weights are modeled as functions of the inputs with an independent GP model for each set of weights.

This additional layer inference for the emulator is combined with the model for the data in Eq. (8). MCMC can then be applied to simultaneously estimate both the unknown physics parameters and the hyperparameters of the Gaussian process (sets of σ^2 and ρ for various dimensions and basis weights).

C. Feature extraction

The conventional approach to analyzing shock experiments is based on matching features in velocimetry curves. For these experiments, we focus on four physically meaningful

points along each velocimetry curve, shown in Fig. 4. The height of the elastic plateau, often called the Hugoniot elastic limit, marks the onset of plastic flow. The height of the plastic plateau is related to the final Hugoniot particle velocity. The approximate slope of the curve between the end of the elastic plateau and the beginning of the plastic plateau is related to the strain rate of the material as it is compressed. Scientists can usually identify these locations by eye. Here, we present an algorithm to extract them automatically based on modeling the velocimetry curves as a linear combination of Gaussian cumulative distribution functions.

A curve for which we would like to find features is denoted as $y(\tau)$, where $\tau \in \mathcal{T}$ and τ_0 is the infimum of $\mathcal{T}$. Here, τ denotes time and $\mathcal{T}$ is the time interval for which the simulation/experiment was run. Now, after standardizing $y(\tau)$ to have minimum zero and maximum one, we model $y(\tau)$ as the cumulative density function (CDF) of a mixture of Gaussian densities. Below, we describe how the features we are interested in correspond to the tails of certain mixture components. First, consider the Gaussian mixture CDF,

$$y(\tau) = \int_{\tau_0}^{\tau} \sum_{i=1}^{k} \phi_i N(x \mid \mu_i, \sigma_i^2) dx, \qquad (10)$$

where $\sum_{i=1}^{k} \phi_i = 1$ and $N(x \mid \mu, \sigma^2) = (2\pi\sigma^2)^{-1/2}$ $\exp\{-0.5(x-\mu)^2/\sigma^2\}$. Here, the values of $\{\phi_i, \mu_i, \sigma_i^2\}_{i=1}^{k}$ are unknown, but fitted with statistical techniques described in Appendix B. We group the mixture components into "early" and "late" components depending on their means. The jump-off and elastic plateau features are then the 0.01 and 0.99 quantiles of the early components. Similarly, the slope features are the 0.1 and 0.9 quantiles of the late components. An example of early and late grouping of components, as well as the corresponding feature locations, is shown in Fig. 5. The plastic plateau is the 0.995 percentile of the original CDF $y(\tau)$.

IV. RESULTS

A. Emulator performance

As discussed above, part of the calibration process produces a statistical approximation to FLAG for the experiments that we simulated. Figure 6 summarizes the results of the emulation process using cross-validation. Cross-validation
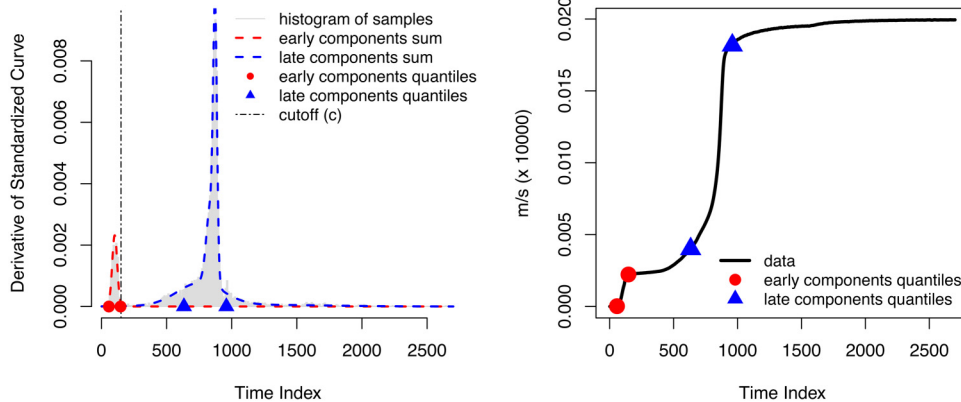

FIG. 5. Left, description of how the five component mixture is separated into two groups—early and late components, based on whether the bulk of each component mass is below a cutoff c. The tails of the sum of early components identify the first two features, while the sum of the late components identify the third and fourth features. Right, identified features on top of the original data.
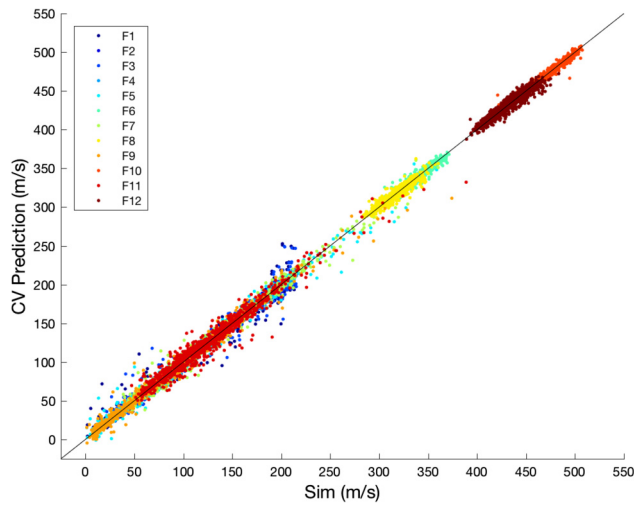
FIG. 6. Comparison of cross-validation predictions of individual feature values versus the corresponding feature value from simulation for all 12 features (four points along each of the three shot curves) in each simulation. RMSE for all 12 features over all (1000) simulation cases is 5.5 m/s. The RMSE computed for individual features over all simulations cases range from 1.8 m/s to 8.6 m/s.

proceeds by rebuilding the emulator with 999 of the simulations and using this emulator to predict the remaining simulation. This process is repeated 1000 times holding out a different run each time. The predictions for each held out simulation are in Fig. 6 with the true simulation values on

the x-axis and the prediction on the y-axis. Thus, points that fall close to the $y = x$ line are predicted well. Each color shows results for a different feature. Although the emulator prediction is not perfect, it generally predicts the physics-based FLAG simulation very well. The root mean squared error (RMSE) across all of the simulations and features is 5.5 m/s. Individual feature RMSEs vary around the all-feature average RMSE fairly closely, with the worst feature-level RMSE being 8.6 m/s.

B. Single shot results

The first three rows of Fig. 7 show the parameter estimation results for the Johnson-Cook parameters based on separate calibrations to each experimental shot. Recall that the MCMC procedure produces a sample of points from the posterior distribution. We use an MCMC sample of size 250 000 for each single-shot calibration. The distributions in each row can be interpreted as the distribution of parameters that give a good match to the experimental data from the corresponding shot. Note that the full distribution is seven dimensional, but we are not showing the results for velocity and shear modulus of each shot. A few insights on the parameterization are evident from these plots. First, some of these parameters show considerable uncertainty based only on individual shots. For example, Shot 104 essentially returns the prior distribution for B (the posterior spans the entire range with no strong peaks), which indicates that the experimental data
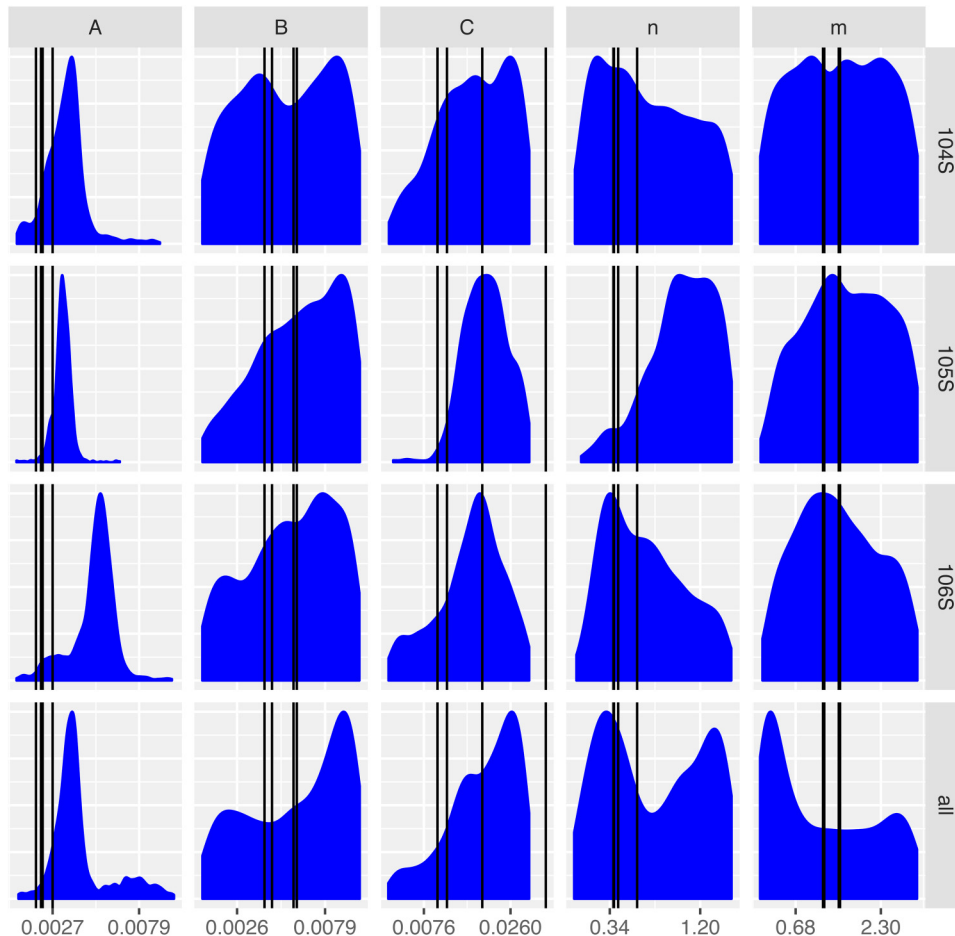

FIG. 7. Summary of Johnson-Cook parameter posterior distributions for calibrations to each of the three shots individually and to all three shots simultaneously. Parameters A and B have units of MBar, all other parameters are unitless. The single-shot calibration results are based on 250 000 MCMC samples and the joint calibration results are based on 175 000 MCMC samples. The black lines show the corresponding four optimal parameters fit by Gray *et al.*[4] Note that the highest value of C from Gray is actually outside the range of our prior distribution.

provide almost no constraint on the parameter B. The other two shots moderately favor higher values of *B*, but leave considerable posterior probability over the whole range. An approach that simply returns the best fitting value would not reveal that many other values are also plausible. Second, the experiments appear to provide conflicting information on the exponential *n* parameter. It may be expected that in isolation each shot, individually probing only a narrow range of state space, provides ambiguous or conflicting information on the strain hardening and temperature sensitivity terms in the model. One advantage of the Bayesian approach is the ease with which these experiments can be combined.

C. Three shot results

The last row of Fig. 7 and all of Fig. 8 summarize the parameter information based on matching all three experimental shots simultaneously. These results are based on an MCMC sample of size 175 000.

The last row of Fig. 7 demonstrates how combining data from across shots can handle conflicting information and reduce uncertainty. The exponent parameters *n* and, to a lesser extent, *m* have bimodal marginal distributions.

Figure 8 shows more information on the joint distribution. Plots along the diagonal of these figures show the density estimates of the posterior for each parameter individually. The off-diagonal plots show contours of the estimates of the marginal bivariate posteriors for each pair of parameters. Collectively, these give a sense of correlations and trade-offs between the parameters. For example, G_1 demonstrates some correlation with the velocity estimates. There is some evidence that one of the modes of *n* is associated with a wider range for (i.e., larger uncertainty in) *C* than the other mode.

Some of the parameters, for example, Δ_3, have posterior distributions with a mode located close to one of the prior parameter bounds. This can indicate some potential problems that we discuss later. In this case, this result is not overly concerning because the calibrated model matches the data reasonably well.
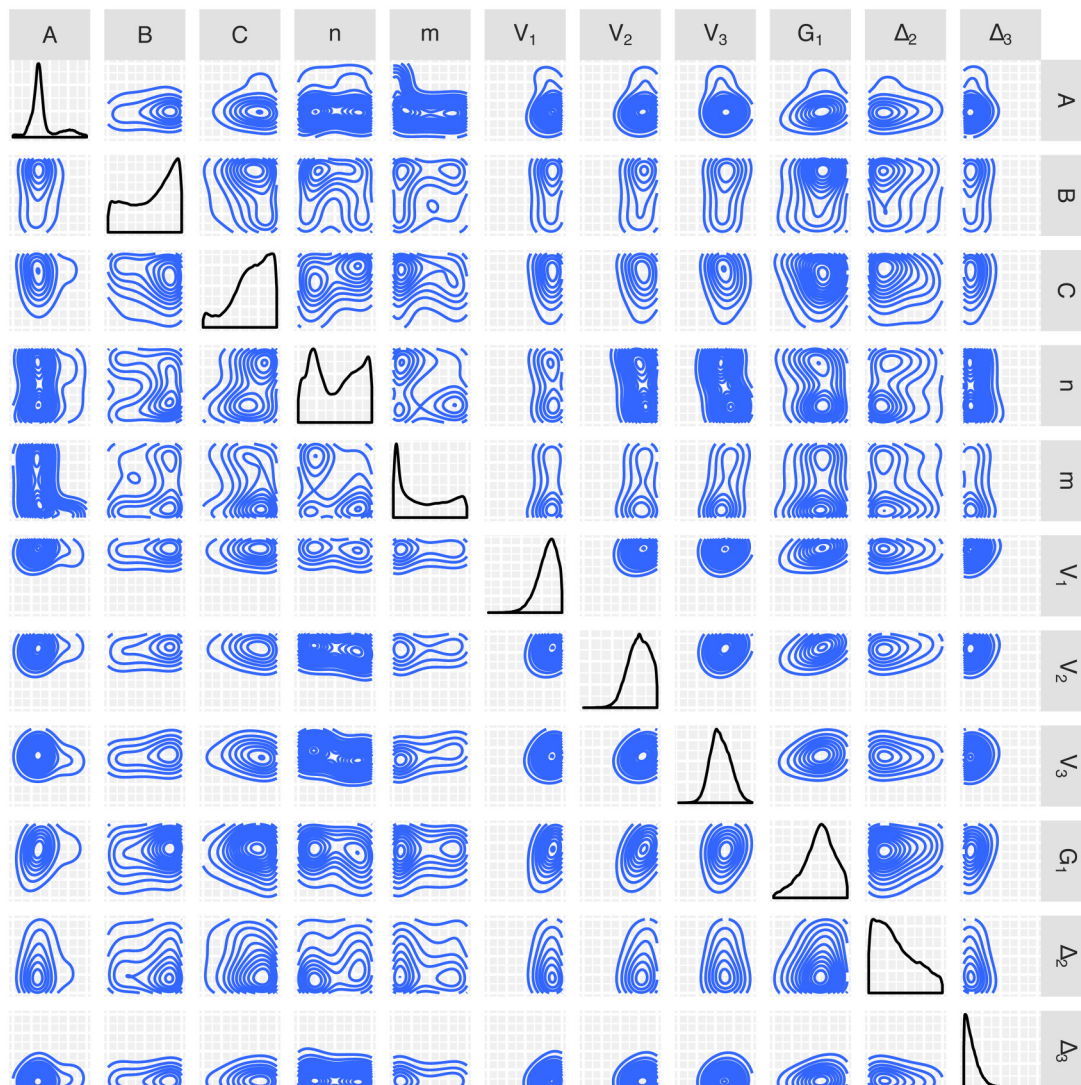

FIG. 8. Summary of 175 000 MCMC samples from the posterior distributions of input parameters following parameter estimation attempting to match all three experimental shots simultaneously. The range in each parameter corresponds to the parameter input bounds from Table IV. (A full page version of this figure is provided in the supplementary material.)

Figure 9 shows several pieces of information in the feature space of the three shots (four points along each of the three shot curves). The diagonal shows results for each feature individually and the off-diagonal plots show the relationship between pairs of features. The red points indicate experimental results. The blue points show the results from the 1000 training simulations. The green points show predictions based on propagating the sample shown in Fig. 8 through the emulator. These predictions do a good job of recovering the experiment with adequate uncertainty. Figure 10 shows the same thing, but using the actual FLAG calculation. Again, the predictions do a good job of recovering the experimental data, but the uncertainty here is somewhat larger. In part, this reflects the emulator error from Fig. 6. This error may be larger on average because of the parameter estimates near the edge where emulator uncertainty can be larger than average. Nevertheless, the

reconstruction is good and represents a rigorous accounting of all sources of error.

Figures 11–13 compare the velocity time histories for each of the three shots, respectively. As before, the blue traces correspond to results from the training simulations, the green traces are from the FLAG simulations using 100 posterior samples of parameter space from Figure 8, and the red trace represents the experimental measurement. Note that the experimental time history (red trace) has been shifted by a constant time because of a timing offset that we are unable to resolve. Aside from absolute timing, the posterior traces appear to represent a reasonable set of predictions, accounting for uncertainty. The heights of the elastic and plastic plateaus, as well as the approximate shape of the transitional rise between elastic precursor and following shock, all seem to be well matched. The variability (green) includes uncertainty from the measurement uncertainty, parameter uncertainty, and emulation uncertainty.
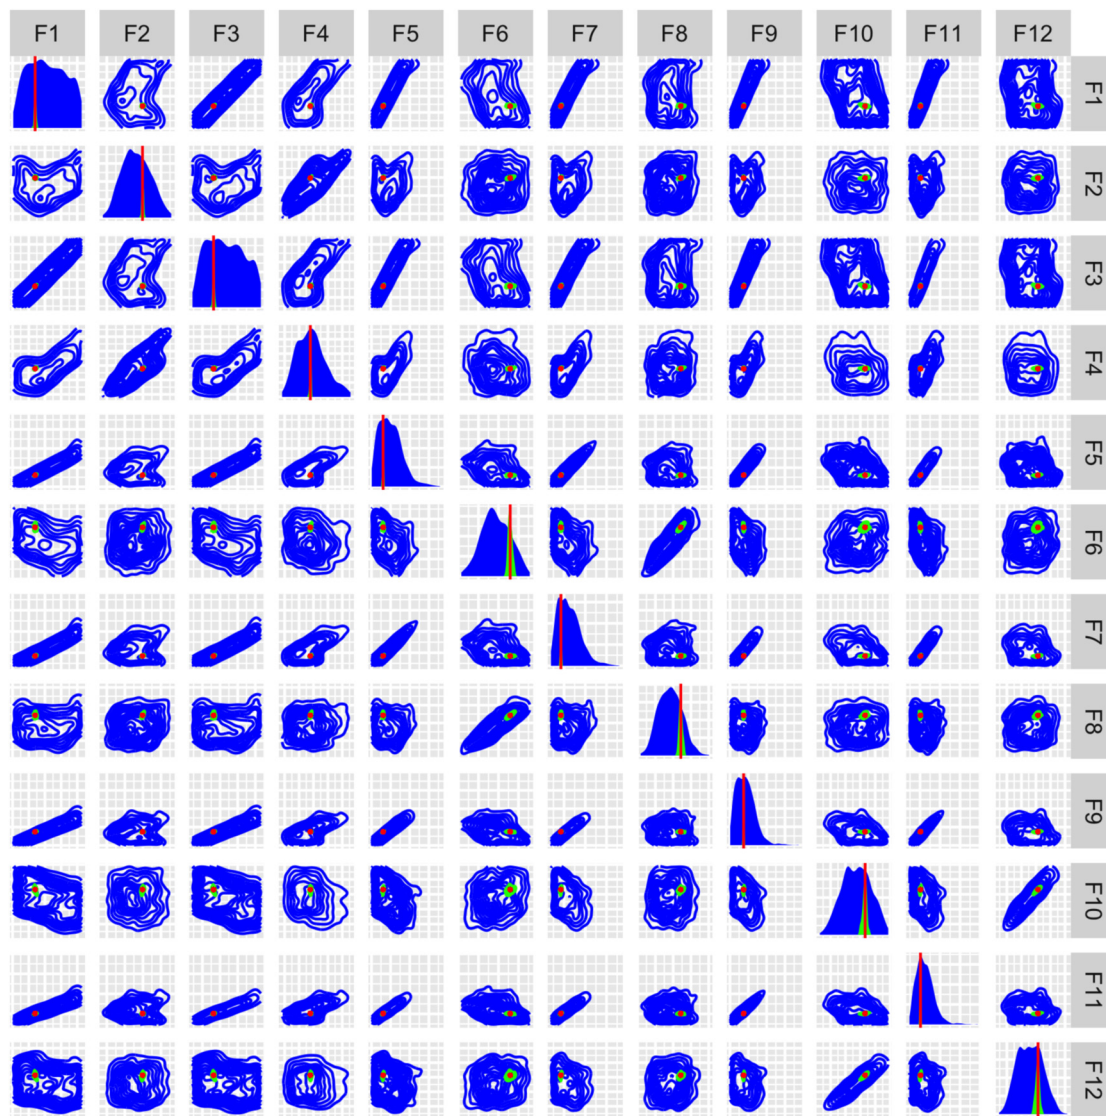

FIG. 9. Information in the feature space of the three shots (four points along each of the three shot curves). The diagonal shows results for each feature individually and the off-diagonal plots show the relationship between pairs of features. The red points and lines indicate experimental results. The blue densities show the results from the 1000 training simulations. The green densities show predictions based on propagating the sample shown in Fig. 8 through the emulator. These predictions do a good job of recovering the experiment with adequate uncertainty. All units are in m/s. See Fig. 6 for approximate ranges. (A full page version of this figure is provided in the supplementary material.)
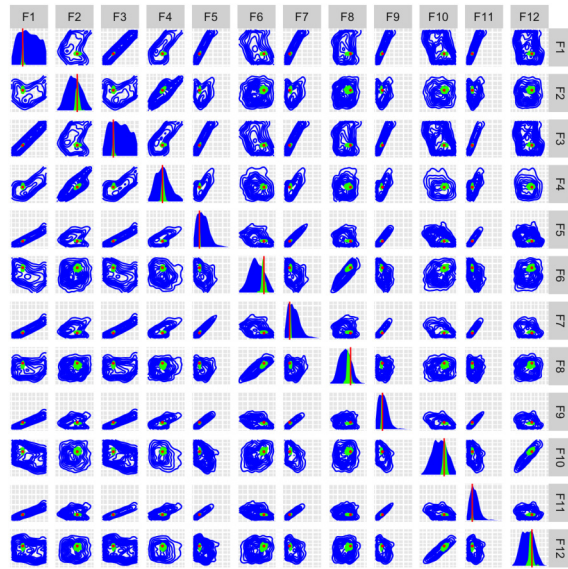
FIG. 10. Information in the feature space of the three shots (four points along each of the three shot curves). The diagonal shows results for each feature individually and the off-diagonal plots show the relationship between pairs of features. The red points and lines indicate experimental results. The blue densities show the results from the 1000 training simulations. The green densities show predictions based on propagating the sample shown in Fig. 8 through FLAG. These predictions show more uncertainty than is shown in Fig. 9 reflecting the emulator error. All units are in m/s. See Fig. 6 for approximate ranges. (A full page version of this figure is provided in the supplementary material.)

V. DISCUSSION

A. Comparison with Gray *et al.*

Optimizing a model's parameters accurately has been done in numerous places for several different materials and experiments using a variety of techniques. Like these other methods, the Bayesian emulator can provide accurate parameter predictions. The Bayesian methods presented here also provide a robust measurement of uncertainty. Its accuracy in this particular case can be validated with a comparison to the
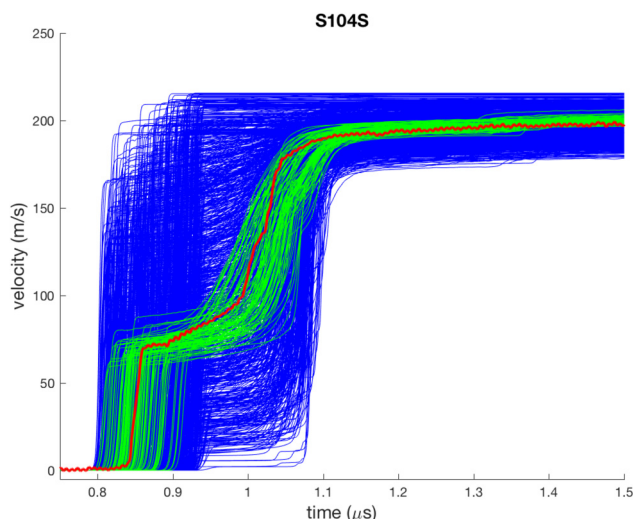

FIG. 12. Comparison of free surface velocity time histories for Shot 105S. Training simulations in blue. Experimental data in red. Simulations using 100 posterior samples from Fig. 8 in green. Note that the timing of the experiment has been adjusted by a constant ($0.45\,\mu s$ subtracted from the times) to align with the simulations.

optimized parameters reported by Gray *et al.*[4] Figure 7 shows that within the numerical uncertainty reported by the emulator, the optimized Johnson-Cook parameters overlap with those listed in Table III.[4] These results affirm that the material models used for the simulations to train and calibrate the emulator appropriately predict the material response given reasonable parameters. When taken with the RMSE cross validation in Fig. 6, these results also reflect the accuracy of the emulator at representing the numerical output of the simulations conducted here.

B. Using uncertainty to guide experiment

The uncertainty associated with the optimized parameters can be a useful tool for guiding further experimentation or
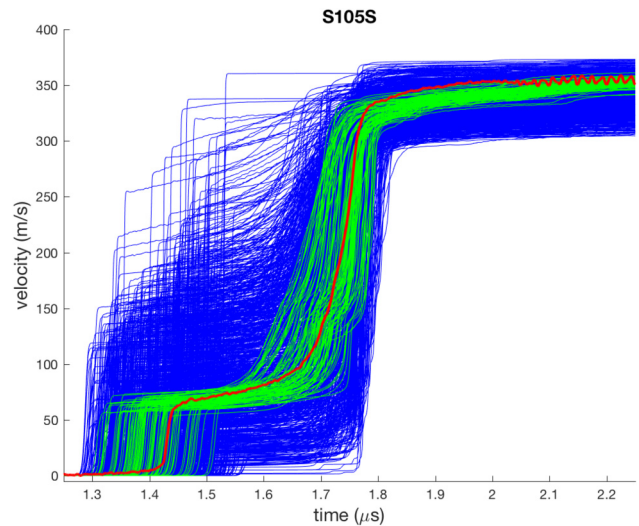

FIG. 11. Comparison of free surface velocity time histories for Shot 104S. Training simulations in blue. Experimental data in red. Simulations using 100 posterior samples from Fig. 8 in green. Note that the timing of the experiment has been adjusted by a constant ($0.08\,\mu s$ subtracted from the times) to align with the simulation results.
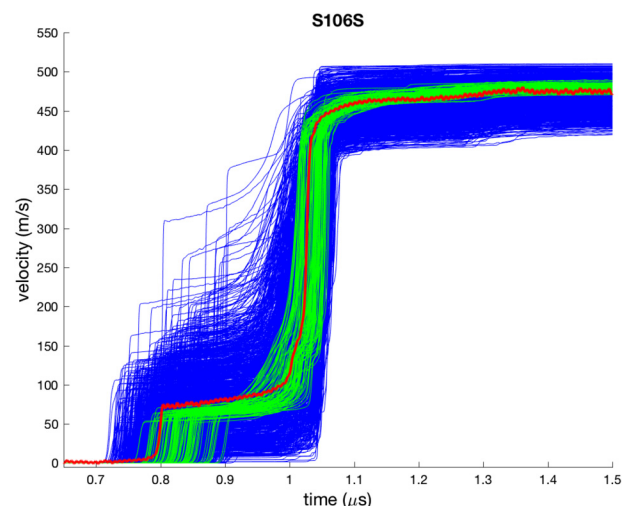

FIG. 13. Comparison of free surface velocity time histories for Shot 106S. Training simulations in blue. Experimental data in red. Simulations using 100 posterior samples from Fig. 8 in green. Note that the timing of the experiment has been adjusted by a constant ($2.75\,\mu s$ subtracted from the times) to align with the simulations.

simulation. The Johnson-Cook flow stress model [see Eq. (5)] empirically associates its parameters with physical aspects of the material response (described previously in Sec. II B 1 b). It is apparent from the posterior distributions that the only parameter well constrained when optimizing to individual shots was A, the initial yield stress. The remaining Johnson-Cook parameters contained a wider band of uncertainty— likely due to probing a limited space of strain, strain rate, and temperature with just a single shot. Without constraint from experimental data sufficiently spanning this physical space of strain, strain rate, and temperature, several equally "good" sets of parameters can non-uniquely reproduce the experiments; thus, parameterization to individual shots led to a large degree of uncertainty in the model parameters. When all three experimental shots were considered, the uncertainty in several parameters was reduced by the additional constraint of experimental data that more broadly spanned the physical space.

Some uncertainty remained after including three experimental shots. For example, focusing on the marginal parameter distributions in Figs. 7 and 8, both exponential terms n and m have a bi-modal distribution of calibrated parameters. The exponent n refers to the characteristic power-law strain hardening behavior and m refers to the power-law temperature softening of the material. The bi-modal distribution could suggest similar modeled outcomes resulting from n picked from the top of its range and m picked from the bottom of its range or vice versa.

Further examination of the physical data set up to replicate the three experimental shots showed the ranges of relevant physical states achieved by the material. For example, the different impact velocities resulted in simulated maximum strain rates of approximately $0.7 \, \mu s^{-1}$, $1.75 \, \mu s^{-1}$, and $3.4 \, \mu s^{-1}$ for shots 0104S, 0105S, and 0106S, respectively—a well explored range of strain rates. A more limited range of temperatures were probed with all shots experiencing the same initial temperature of 298 K and reaching maximums of 309 K, 322 K, and 365 K, respectively. The variability in maximum strain achieved was quite small with simulated maximums of approximately 0.015, 0.020, and 0.028 for each respective shot. The limited range of physical states of temperature and strain likely contributed to the wide uncertainty exhibited by parameters most related to temperature and strain.

An envisioned application of this framework is to provide relatively rapid feedback and evolution of an experimental design during an ongoing sequence of experiments, for example, suggesting in this case that additional shots be performed at different initial temperatures to probe a wider temperature space to reduce uncertainty in the distributions of m. Alternatively, redesigning the specifications of the flyer and/or target could be implemented to achieve significantly different ranges of strain.

C. Making sense of the results

Undesirable levels of uncertainty or unexpected optimized parameter distributions could also indicate the need to adjust aspects of the emulator training data. For certain optimized patterns, such as when distributions of parameters

favor the boundary of the initial parameter range, one may need to provide a larger range of possible values. In addition to providing a robust technique for parameter optimization, the tabulated uncertainty and prediction of possible outputs could provide an indication of the validity of the training data when compared to reality. Perhaps, the constitutive models used to describe the material behavior inadequately capture the physics of the problem. In this case, we can add a discrepancy term to the statistical model as described in Higdon *et al.*[28] This term can estimate systematic biases between the physics model and the experiments. These biases may reflect this missing physics and correct for it empirically.

As noted earlier, posterior parameter distributions that are pushed up against prior boundaries can indicate a problem. Such occurrences are consistent with several possible explanations. First, it may simply be the case that this is the best distribution to characterize the particular parameter given the experimental data. More problematically, it may indicate that our prior bounds were badly chosen, or it may indicate that the physics model is not capable of predicting all aspects of the experimental data. In this latter case, the calibration process may favor parameter values near their bounds where the Gaussian process emulator has a large uncertainty about what the physics simulation should be. In other words, there may be very little uncertainty about the fact that the physics model does not match the data over parameters safely within the bounds, therefore, the most plausible region is that with the most uncertainty about what the physics model might predict. These two potential problems are related because our prior bounds may have excluded parts of the parameter space that would have enabled the model to match the experimental data. A further test of these conditions is to verify whether or not the posterior parameter distribution produces predictions that match the experimental data reasonably well. Figure 14 demonstrates this by recreating Fig. 1, but with the model unable to match the data (see Fig. 1 for an explanation). The resulting posterior for the parameter is pushed against the boundary to compensate as much as possible, which is why this behavior can be an indicator of model inadequacy.
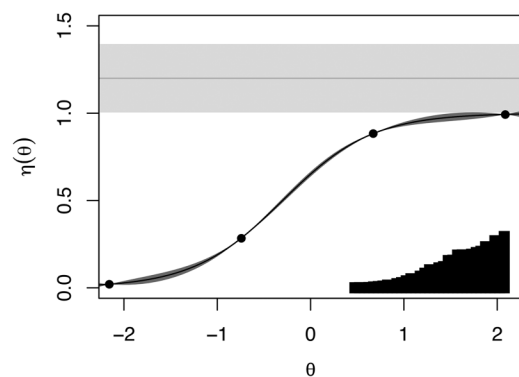

FIG. 14. Illustration of model calibration when the physics model is not capable of matching the experimental data over the prior ranges. Unlike the demonstration in Fig. 1, the model is unable to match the experimental data. See Fig. 1 for an explanation of the details of this figure.

In our current study, the output predictions of both the calibrated emulator and simulator match the features of the experimental velocimetry curves well (Figs. 9 and 10) indicating a good performance of the underlying constitutive model. This gives us a strong indication that we are not suffering from the type of model inadequacy just discussed, despite some of the parameter posterior distributions being close to the edges of prior boundaries.

Along with the uncertainty of the optimized parameters, Figs. 11–13 show how the uncertainty can propagate to the predicted output. Another utility of the emulator is the ability to calibrate across multiple data sets. This was demonstrated here by optimizing parameters for individual experimental shots, then combining the data from all three shots to improve the optimization by narrowing the uncertainty in several parameters. Theoretically, other types of experiments (such as those conducted by Gray *et al.*[4]) could also be incorporated to augment the set of data available for training the emulator.

VI. SUMMARY

This effort is particularly directed toward developing parameter estimation approaches that can be used in the context of dynamic materials experiments (e.g., gas-gun plate impact) to rapidly inform experimentalists and modelers on the implications to model parameters from experimental data including velocimetry. Consequently, we focused this effort on Al-5083 because of the availability of relevant velocimetry data,[31] as well as previous model parameterizations from dynamic experiments in other strain-rate regimes. The Johnson-Cook strength model and its associated parameters for the plastic deformation of Al-5083 were chosen for this study due its simplicity, the relatively wide use of the model, and the availability of other parameterizations and experimental data. The EOS was selected to maintain consistency with the previous work by Boteler and Dandekar.[31]

The power of Bayesian emulation was demonstrated here by reasonably optimizing material parameters, predicting the resulting output related to experiments, and quantifying the variability of the optimized parameters and predicted output. This tool can be applied across multiple data sets (experiments) to quickly evaluate optimal parameters and model adequacy and suggest areas of further exploration through the computation of uncertainty. These tools might be especially well suited to dynamic compression sciences at next generation light sources such as X-ray free electron lasers, where dynamic experiment rates are rapidly accelerating.[47–49]

SUPPLEMENTARY MATERIAL

Full page versions of Figs. 8–10 are provided as supplementary material to better illustrate small details contained within the plots.

ACKNOWLEDGMENTS

APPENDIX A: DETAILS OF THE GAUSSIAN PROCESS CALIBRATION APPROACH FOR MULTIVARIATE OUTPUT

Let $\eta(t)$ represent the FLAG output for d-dimensional input vector t. We will assume that there is some true input vector θ that predicts our actual experimental value y, within measurement error ϵ

$$y = \eta(\theta) + \epsilon, \quad \epsilon \sim N(0, \Sigma_y). \quad (A1)$$

We assume a prior distribution for θ, denoted $\pi(\theta)$ that is uniform over some hyper-rectangle $\mathcal{C}$ based on upper and lower bounds for each input parameter. This gives a posterior for the unknown parameters of

$$p(\theta \mid y) \propto L[y \mid \eta(\theta)]\pi(\theta), \quad (A2)$$

where

$$L[y \mid \eta(\theta)] = \exp\left\{\frac{1}{2}[y - \eta(\theta)]'\Sigma_y^{-1}[y - \eta(\theta)]\right\} \quad (A3)$$

is the likelihood and comes the Gaussian model for the data. Because $\eta(t)$ is slow and unknown over most of the input space, we include it in our inference. Let $\eta^* = [\eta(t_1), \ldots, \eta(t_m)]$ be the simulation output from a collection of inputs $t_1, \ldots, t_m$. This will add a component to both the likelihood and the prior

$$\pi[\theta, \eta(\cdot) \mid y, \eta^*] \propto L[y \mid \eta(\theta)] \cdot L[\eta^*$$
$$\mid \eta(\cdot)] \cdot \pi[\eta(\cdot)] \cdot \pi(\theta). \quad (A4)$$

For convenience, we assume that we are working with standardized output. In practice, we standardize the training simulations and the experimental data by the mean vector and a scaling factor computed from the simulation training set.

Our estimate for $\eta(\cdot)$ is called an emulator. For simplicity, scale the inputs such that $t \in [0, 1]^d$. It starts with a basis decomposition

$$\eta(t) = \sum_{i=1}^{q} \phi_i w_i(t) + \upsilon, \quad t \in [0, 1]^d, \quad (A5)$$

where $\{\phi_1, \ldots, \phi_q\}$ are orthogonal basis vectors, the $w_i(t)$ are weights depending on the input, and υ is an error term, assumed to be small, that accounts for deficiencies in the basis representation. This formulation assumes that the structure of the simulation output, as represented by the basis set, is constant over the input and that the weights vary with the input. In practice, this seems to work well.

Each $w_i(\cdot)$ is modeled with a Gaussian process

$$w_i(t) \sim N\left[0, \lambda_{wi}^{-1}R(t; \rho_i)\right], \quad (A6)$$

where λ_{wi} is the marginal precision (inverse variance) and $R(t; \rho_i)$ is a correlation matrix with entries dependent given

by the correlation function

$$\text{Corr}[w_i(t), w_i(t')] = \prod_{k=1}^{p} \rho_{ik}^{4(t_k - t'_k)^2}. \tag{A7}$$

Let $w_i = [w_i(t_1), \ldots, w_i(t_m)]'$ for $i = 1, \ldots, q$. Further, let $R(t; \rho_i)$ be the $m \times m$ correlation matrix resulting from applying Eq. (A7) to each pair of input settings, where the p-vector ρ_i gives the correlation distances for each of the input dimensions. Then, we can write

$$\begin{pmatrix} w_1 \\ \vdots \\ w_q \end{pmatrix} \sim N \left[\begin{pmatrix} 0 \\ \vdots \\ 0 \end{pmatrix}, \begin{pmatrix} \lambda_{w1}^{-1} R(t; \rho_1) & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \lambda_{wq}^{-1} R(t; \rho_q) \end{pmatrix} \right]. \tag{A8}$$

More succinctly, we can write $w \sim N(0, \Sigma_w)$.

The unknown GP parameters are given gamma priors for the precisions and beta priors for the correlations

$$\pi(\lambda_{wi}) \propto \lambda_{wi}^{a_w - 1} e^{-b_w \lambda_{wi}}, \quad i = 1, \ldots, q,$$

$$\pi(\rho_{ik}) \propto \rho_{ik}^{a_\rho - 1}(1 - \rho_{ik})^{b_\rho - 1}, \quad i = 1, \ldots, q, \quad k = 1, \ldots, p. \tag{A9}$$

The basis decomposition should result in the variance of the $w_i(t)$ being near one. Thus, we choose $a_w = b_w = 5$. Additionally, we typically expect that only a subset of the inputs will have a large effect on the output. This implies $\rho_i k$ is often near one. For this reason, we choose $a_\rho = 1$ and $b_\rho = 0.1$, which gives $\text{Pr}(\rho_{ik} < 0.98) \approx \frac{1}{3}$ *a priori*.

Returning to the error term υ in the basis decomposition, we will treat this as independent Gaussian. Let η be the vector obtained by stacking all of the $n(t_i)$ and let $\Phi = [I_m \otimes \phi_1; \cdots; I_m \otimes \phi_q]$. Now, we can write

$$\eta \mid w, \lambda_\eta \sim N\left(\Phi w, \frac{1}{\lambda_\eta} I\right). \tag{A10}$$

we give λ_η a Gamma prior with parameters (a_η, b_η).

The experimental data are now modeled using the basis representation as well

$$y = \Phi_y w(\theta) + \epsilon, \tag{A11}$$

where $w(\theta)$ is the q-vector $[w_1(\theta), \ldots, w_q(\theta)]'$, the basis weights evaluated at the unknown best input, and Φ_y are the basis vectors interpolated onto the index for the physical data, if different than the simulations. We reparameterize the variance of the data in terms of a precision, Σ_y^{-1} as $\lambda_y W_y$, where λ_y can be estimated as scaling parameter. This gives

$$y \mid w(\theta), \lambda_y \sim N[\Phi_y w(\theta), (\lambda_y W_y)^{-1}], \tag{A12}$$

$$\lambda_y \sim Ga(a_y, b_y). \tag{A13}$$

We constrain the scaling parameter near one by setting $a_y = b_y = 5$.

We can now write the entire posterior distribution. First, define

$$a_\eta^* = a_\eta + \frac{m(n_\eta - q)}{2},$$

$$b_\eta^* = b_\eta + \frac{1}{2}\eta'[I - \Phi(\Phi'\Phi)^{-1}\Phi']\eta,$$

$$\hat{w} = (\Phi'\Phi)^{-1}\Phi'\eta,$$

$$\hat{w}_y = (\Phi'_y W_y \Phi_y)^{-1}\Phi'_y W_y y,$$

$$a_y^* = a_y + \frac{1}{2}(n - q),$$

$$b_y^* = b_y + \frac{1}{2}(y - \Phi_y \hat{w}_y)' W_y (y - \Phi_y \hat{w}_y),$$

$$\Lambda_y = \lambda_y \Phi'_y W_y \Phi_y,$$

$$\Lambda_\eta = \lambda_\eta \Phi'\Phi,$$

$$I_q = q \times q \text{ identity matrix},$$

$$\Sigma_{w_y w} = \begin{pmatrix} \lambda_{w1}^{-1} R(\theta, \theta^*; \rho_1) & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \lambda_{wq}^{-1} R(\theta, \theta^*; \rho_q) \end{pmatrix},$$

$$\hat{z} = \begin{pmatrix} \hat{w}_y \\ \hat{w} \end{pmatrix},$$

$$\Sigma_{\hat{z}} = \begin{pmatrix} \Lambda_y^{-1} & 0 \\ 0 & \Lambda_\eta^{-1} \end{pmatrix} + \begin{pmatrix} I_q & \Sigma_{w_y w} \\ \Sigma'_{w_y w} & \Sigma_w \end{pmatrix}.$$

The posterior can now be written as

$$\pi(\lambda_\eta, \lambda_w, \rho, \lambda_y, \theta | \hat{z})$$

$$\propto |\Sigma_{\hat{z}}|^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}\hat{z}'\Sigma_{\hat{z}}^{-1}\hat{z}\right\} \times \lambda_\eta^{a_\eta^* - 1} e^{-b_\eta^* \lambda_\eta}$$

$$\times \prod_{i=1}^{q} \lambda_{wi}^{a_w - 1} e^{-b_w \lambda_{wi}} \times \prod_{i=1}^{q} \prod_{k=1}^{p} \rho_{ik}^{a_\rho - 1} (1 - \rho_{ik})^{b_\rho - 1}$$

$$\times \lambda_y^{a_y^* - 1} e^{-b_y^* \lambda_y} \times I[\theta \in C], \tag{A14}$$

where C denotes the p-dimensional rectangle defined by the prior parameter ranges.

The best fit inputs and all of the statistical parameters can be sampled using Markov chain Monte Carlo.

APPENDIX B: DETAILS OF THE VELOCIMETRY FEATURE EXTRACTION

Typical methods for finding $\{\phi_i, \mu_i, \sigma_i^2\}_{i=1}^{k}$ assume that we have independent and identically distributed samples from $g(x)$, say $x_1, \ldots, x_n \sim g$, where $g(x) = \sum_{i=1}^{k} \phi_i N(x \mid \mu_i, \sigma_i^2)$. While we do not have samples, we can easily generate them using the probability integral transform. To demonstrate this, consider the random variable $x \sim g$. Then, a new random variable $z = y(x)$ has Uniform(0,1) distribution, since y is the CDF of x. We can generate a sample $x_1, \ldots, x_n \sim g$ by first generating $z_1, \ldots, z_n$ and using the inverse transformation to get $x_i = y^{-1}(z_i)$ for $i = 1, \ldots, n$. While we do not have an
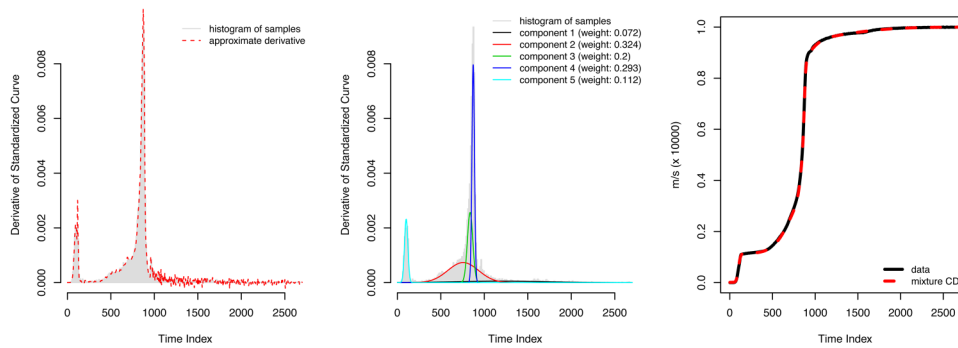
FIG. 15. Using the probability integral transform, samples from a density function look like the derivative of the curve of interest (left plot). The density function is approximated with a mixture of Gaussians (middle plot), which can be integrated to get an approximation of the curve of interest (right plot).

analytical version of $y(\tau)$ to invert, our observed $y(\tau)$ is on a dense enough grid of t values that we can get a good approximation of the inverse function y^{-1} by linearly interpolating the $[y(\tau), \tau]$ pairs.

We use a sample of 10 000 values from $g(x)$ to find $\{\phi_i, \mu_i, \sigma_i^2\}_{i=1}^k$ using an expectation maximization algorithm. Where it may seem like two mixture components ($k = 2$) would be sufficient for our feature finding purposes, we find that using more components allows us to (1) capture tail behavior that is heavier than Gaussian and (2) accurately model the CDF when there are clearly more than two jumps in $y(\tau)$. Specifically, we use five mixture components ($k = 5$). The mixture components for one curve are shown in Fig. 15.

Upon obtaining our five components defined by $\{\phi_i, \mu_i, \sigma_i^2\}_{i=1}^5$, we split them into two groups, which we call the early and late groups. If $\mu_i + 2\sigma_i < c$, component i is assigned to the early group, indicating that most of the mass of the component occurs before c. The cutoff c is visually identifiable when all the data are plotted together (a time point between the elastic plateau feature and the first elastic-plastic slope feature).

Inherent in the algorithm is the assumption of monotonicity. While not physically justified in this case, this is a practical assumption given the shapes of the velocity curves that we observe. The curves have some high frequency signal that can have negative slope. However, that signal is very small relative to the low frequency signal and can be ignored for our purposes. This can present problems when we build the inverse CDF function y^{-1}, because the inverse is not well-defined. The resulting linear interpolation is slightly unstable, but that instability is small in comparison to the general shape of the function. To use this approach for non-monotonic curves, we could allow for negative weights, essentially treating the Gaussian densities as basis functions.

[1] G. R. Johnson and W. H. Cook, in *Seventh International Symposium on Ballistics, The Hague, The Netherlands* (1983), pp. 541–547; available at http://www.worldcat.org/title/proceedings-seventh-international-symposium-on-ballistics-the-hague-the-netherlands-19-21april-1983/oclc/9951901.

[2] F. J. Zerilli and R. W. Armstrong, J. Appl. Phys. **61**, 1816 (1987).

[3] D. L. Preston, D. L. Tonks, and D. C. Wallace, J. Appl. Phys. **93**, 211 (2003).

[4] G. T. Gray, S. R. Chen, W. Wright, and M. F. Lopez, "Constitutive equations for annealed metals under compression at high strain rates and high temperatures," Technical Report No. LA-12669-MS, 1994.

[5] T. Børvik, O. S. Hopperstad, T. Berstad, and M. Langseth, Eur. J. Mech. A/Solids **20**, 685 (2001).

[6] S. M. Goh, M. N. Charalambides, and J. G. Williams, Mech. Time-Dependent Mater. **8**, 255 (2004).

[7] H. Mecking and U. Kocks, Acta Metallurg. **29**, 1865 (1981).

[8] Y. Estrin and H. Mecking, Acta Metallurg. **32**, 57 (1984).

[9] Y. Estrin and H. Mecking, Scr. Metallurg. **19**, 451 (1985).

[10] P. Follansbee and U. Kocks, Acta Metallurg. **36**, 81 (1988).

[11] D. J. Bammann, Appl. Mech. Rev. **43**, S312 (1990).

[12] D. Bammann, M. Chiesa, and G. Johnson, "A strain rate dependent flow surface model of plasticity," Sandia National Laboratory Technical Report No. SAND90-8227, 1990.

[13] D. Bammann, M. Chiesa, A. McDonald, W. Kawahara, and J. Dike, in *Failure Criteria and Analysis in Dynamic Response* (ASME, Metals Park, OH, 1990), p. 7.

[14] D. L. Preston, "PTW Mod 1," Los Alamos National Laboratory Technical Report No. LA-UR-10-06436, 2010.

[15] R. A. Austin and D. L. McDowell, Int. J. Plast. **27**, 1 (2011).

[16] R. A. Austin and D. L. McDowell, Int. J. Plast. **32–33**, 134 (2012).

[17] D. Luscher, F. Addessio, M. Cawkwell, and K. Ramos, J. Mech. Phys. Solids **98**, 63 (2017).

[18] L. M. Barker and R. E. Hollenbach, J. Appl. Phys. **43**, 4669 (1972).

[19] K. J. Ramos, B. J. Jensen, A. J. Iverson, J. D. Yeager, C. A. Carlson, D. S. Montgomery, D. G. Thompson, K. Fezzaa, and D. E. Hooks, J. Phys. Conf. Ser. **500**, 142028 (2014).

[20] J. Winey and Y. Gupta, J. Appl. Phys. **99**, 023510 (2006).

[21] J. Winey and Y. Gupta, J. Appl. Phys. **107**, 103505 (2010).

[22] F. Addessio, D. Luscher, M. Cawkwell, and K. Ramos, J. Appl. Phys. **121**, 185902 (2017).

[23] D. Versino and C. A. Bronkhorst, Comput. Methods Appl. Mech. Eng. **333**, 395 (2018).

[24] M. B. Prime, W. T. Buttler, M. A. Buechler, N. A. Denissen, M. A. Kenamond, F. G. Mariam, J. I. Martinez, D. M. Oró, D. W. Schmidt, J. B. Stone *et al.*, J. Dyn. Behav. Mater. **3**, 189 (2017).

[25] J. L. Brown and L. B. Hund, J. R. Stat. Soc. Ser. C (Applied Stat.) **67**, 1023–1045 (2018).

[26] S. J. Ali, R. G. Kraus, D. E. Fratanduono, D. C. Swift, and J. H. Eggert, J. Appl. Phys. **121**, 195901 (2017).

[27] M. C. Kennedy and A. O'Hagan, J. R. Stat. Soc. B (Stat. Methodol.) **63**, 425 (2001).

[28] D. Higdon, J. Gattiker, B. Williams, and M. Rightley, J. Am. Stat. Assoc. **103**, 570 (2008).

[29] D. Higdon, E. Lawrence, K. Heitmann, and S. Habib, in *Statistical Challenges in Modern Astronomy V* (Springer, 2012), pp. 41–57.

[30] J. McDonnell, N. Schunck, D. Higdon, J. Sarich, S. Wild, and W. Nazarewicz, Phys. Rev. Lett. **114**, 122501 (2015).

[31] J. M. Boteler and D. P. Dandekar, J. Appl. Phys. **100**, 054902 (2006).

[32] T. Børvik, M. Langseth, O. S. Hopperstad, and K. A. Malo, Int. J. Impact Eng. **22**, 855 (1999).

[33] R. Winzer and A. Glinicka, Eng. Trans. **59**, 85 (2011).

[34] J. K. Holmen, O. S. Hopperstad, and T. Børvik, Int. J. Impact Eng. **108**, 136 (2017).

[35] Y. K. Xiao, H. Wu, Q. Fang, W. Zhang, and X. Z. Kong, Mater. Des. **133**, 237 (2017).

[36] D. E. Burton, in *Nuclear Explosives Code Developers Conference* (Lawrence Livermore National Laboratory, Livermore, CA, 1992).

[37] D. E. Burton, in *Symposium on Advanced Numerical Methods for Lagrangian Hydrodynamics* (Los Alamos National Laboratory, Los Alamos, NM, 2007), Vol. LA-UR-07-7547.

[38] E. J. Caramana, D. E. Burton, M. J. Shashkov, and P. P. Whalen, J. Comput. Phys. **146**, 227 (1998).

[39] G. Hauver and A. Melani, "The Hugoniot of 5083 aluminum," Ballistic Research Laboratories Technical Report No. 2345, 1973.

[40]G. R. Fowles, J. Appl. Phys. **32**, 1475 (1961).

[41]L. G. Margolin, "A centered artificial viscosity for cells with large aspect ratios," Lawrence Livermore National Laboratory Technical Report No. UCRL-53882, 1988.

[42]L. Burakovsky, C. W. Greeff, and D. L. Preston, Phys. Rev. B **67**, 094107 (2003).

[43]J. L. Loeppky, J. Sacks, and W. J. Welch, Technometrics **51**, 366 (2009).

[44]G. Casella and E. I. George, Am. Stat. **46**, 167 (1992).

[45]S. Chib and E. Greenberg, Am. Stat. **49**, 327 (1995).

[46]C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning* (MIT Press, Cambridge, MA, 2006), Vol. 1.

[47]E. Weckert, Int. Union Crystallogr. J. **2**, 230 (2015).

[48]R. Schoenlein, P. Abbamonte, F. Abild-Pedersen, P. Adams, M. Ahmed, F. Albert, R. A. Mori, A. Anfinrud, A. Aquila, M. Armstrong *et al.*, SLAC Report No. SLAC-R-1053, 2015.

[49]U. Zastrau, M. McMahon, K. Appel, C. Baehtz, E. Brambrink, R. Briggs, T. Butcher, B. Cauble, B. Chen, H. Damker, C. Deiter, J. Eggert, K. Falk, L. Fletcher, S. H. Glenzer, S. Göde, M. Harmand, A. Higginbotham, Z. Konôpková, D. Kraus, H.-P. Liermann, M. Nakatsutsumi, P. A., G. Priebe, R. Redmer, A. Schropp, R. Smith, P. Sperling, I. Thorpe, and S. Toleikis, "Conceptual design report: Dynamic laser compression experiments at the HED instrument of European XFEL," European X-Ray Free-Electron Laser Facility GmbH Technical Report No. REPORT-2017-004, XFEL.EU TR-2017-001, 2017.