Posterior Covariance Matrix Approximations

Schmid, Abigail Coyle
Andrews, Stephen Arthur

To be published in: Journal of Verification, Validation and Uncertainty Quantification

DOI to publisher's version: 10.1115/1.4065378

POSTERIOR COVARIANCE MATRIX APPROXIMATIONS

Abigail C. Schmid[1,2], Stephen A. Andrews[1,*]

[1]Los Alamos National Laboratory, Los Alamos, NM
[2]University of Colorado, Boulder, CO

ABSTRACT

The Davis Equation of State (EOS) is commonly used to model thermodynamic relationships for High Explosive reactants. Typically, the parameters in the EOS are calibrated, with uncertainty, using a Bayesian framework and Markov Chain Monte Carlo (MCMC) methods. However, MCMC methods are computationally expensive, especially for complex models with many parameters. This paper provides a comparison between MCMC and less computationally expensive Variational methods (Variational Bayesian and Hessian Variational Bayesian) for computing the posterior distribution and approximating the posterior covariance matrix based on heterogeneous experimental data. All three methods recover similar posterior distributions and posterior covariance matrices. This study demonstrates that for this EOS parameter calibration application, the assumptions made in the two Variational methods significantly reduce the computational cost but do not substantially change the results compared to MCMC.

Keywords: Bayesian, Uncertainty Quantification, MCMC, Variational Inference, Equation of State

1. INTRODUCTION

Equations of State (EOS) are physics models which describe the relationship between thermodynamic quantities for a material. The parameters in an EOS are usually calibrated based on experimental data for the material of interest. Often, the Davis EOS [1–3] is used to model High Explosive (HE) reactants and the parameters are calibrated using experimental data from the HE and its components. Having an accurate description of the EOS is essential for running large scale simulations.

Typically, calibrations of EOS parameters, with uncertainty quantification, are performed in a Bayesian framework, using Markov Chain Monte Carlo (MCMC) methods to compute the posterior (e.g. [4],[5]). However MCMC approaches are computationally expensive and it can be difficult to assess whether the solution has converged to the posterior distribution, especially for high dimensional problems. This can pose significant challenges in settings with expensive simulations and complex models with many parameters to calibrate. An alternative method to characterize the posterior is using Variational Bayesian methods. These approaches are usually more computationally efficient than MCMC methods [6]. However Variational Bayesian methods provide an approximation of the posterior distribution, whereas MCMC methods can, in the limit, fully characterize the posterior.

Depending on the setting, the approximations made in the Variational methods may or may not be appropriate. Direct comparisons between Variational Bayesian and MCMC methods for model calibration have been performed with experimental data in the fields of astrophysics [7], cosmology [8], and psychology [9], but this comparison has not been carried out for the calibration of EOSs.

This paper calibrates the Davis Reactants Equation of State parameters, with uncertainty, for the high explosive PBX 9501. The calibration is posed as a Bayesian inverse problem which uses a heterogeneous set of experimental data. This paper also provides a direct comparison of MCMC and Variational Bayesian methods for computing the posterior and estimating the posterior covariance matrix, for the Davis Reactants equation of state, an example of a parametric physics model.

Section 2 describes the Equation of State and parameters to calibrate, the data that comprises the heterogeneous dataset, and the Bayesian inverse problem. Section 3 explains the three methods used to compute the posterior distributions and approximate the posterior covariance matrix. Section 4 shows the results of the parameter calibrations, estimates of the posterior covariance matrix, and uncertainties. Section 5 provides conclusions from this work.

2. PROBLEM DESCRIPTION

This study aims to compare three methods of computing the posterior covariance matrix, as a method of uncertainty quantification, in the context of calibrating the parameters in the Davis Reactants Equation of State (EOS) for the high explosive (HE) PBX 9501 using a heterogeneous dataset.

*Corresponding author: saandrews@lanl.gov

2.1 Equation of State

The Davis Reactants Equation of State (EOS) describes the relationship between thermodynamic quantities like specific internal energy, pressure, and density. The model was originally developed by Davis [1] and subsequently improved by Wescott *et al.* [2] and Aslam [3]. The model is a Mie-Grüneisen form and is based on the reference isentrope (denoted with a subscript r) through pressure $P = 0$ and density $\rho = \rho_0$, where ρ_0 is the reference ambient density. The equations here represent the functional form used by Aslam [3].

The EOS provides functions for the pressure P and temperature T as a function of density ρ and specific internal energy e,

$$P(\rho, e) = P_r(\rho) + \rho\Gamma(\rho)(e - e_r(\rho)) \tag{1}$$

$$T(\rho, e) = T_r(\rho)\left\{\frac{1+\alpha_{st}}{C_v T_r(\rho)}\left[e - e_r(\rho)\right]\right\}^{\frac{1}{1+\alpha_{st}}}. \tag{2}$$

The EOS defines the reference isentrope pressure by

$$P_r(\rho) = \frac{\rho_0 A^2}{4B}\begin{cases}\exp(4By) - 1 & y < 0\\ \sum_{j=1}^{3}\frac{(4By)^j}{j!} + C\frac{(4By)^4}{4!} + \frac{y^2}{(1-y)^4} & y \geq 0\end{cases} \tag{3}$$

where $y = 1 - \rho_0/\rho$ is the compression factor, A is the sound speed in the material at ambient conditions, B is the slope of the Hugoniot curve, and C is a high particle speed correction for the Hugoniot curve. The energy on the reference curve is,

$$e_r(\rho_i) = \int_0^{\rho_i}\frac{P_r(\rho)}{\rho^2}d\rho + E_0, \tag{4}$$

where E_0 is the stored chemical potential energy in the material. The temperature along the reference isentrope is,

$$T_r(\rho) = \begin{cases} T_0\left(\frac{\rho}{\rho_0}\right)^{\Gamma_r^0} & \rho \leq \rho_0\\ T_0\exp(-Zy)\left(\frac{\rho}{\rho_0}\right)^{(\Gamma_r^0+Z)} & \rho > \rho_0\end{cases} \tag{5}$$

where T_0 is the reference temperature and Z describes how the Grüneisen gamma changes with density. Additionally, the specific heat capacity is given by,

$$C_{vr}(\rho_0, T) = C_v\left(\frac{T}{T_0}\right)^{\alpha_{st}}, \tag{6}$$

where α_{st} relates how the specific heat changes with temperature. The Grüneisen parameter is defined as,

$$\Gamma_r(\rho) = \begin{cases}\Gamma_r^0 & \rho < \rho_0\\ \Gamma_r^0 + Zy & \rho \geq \rho_0\end{cases}. \tag{7}$$

The Grüneisen gamma is strongly related to the volumetric thermal expansion coefficient β,

$$\Gamma_r^0 = \sqrt{\frac{1}{4\beta^2 T_0^2} + \frac{A^2}{C_v T_0}} - \frac{1}{2\beta T_0}. \tag{8}$$

Thus, to fully define the reactants EOS, the parameters A, B, C, ρ_0, Γ_r^0, C_v, α_{st}, T_0, E_0, and Z need to be specified. Here the values $\rho_0 = 1.836\,\mathrm{g\,cm^{-3}}$, $T_0 = 297\,\mathrm{K}$, $E_0 = 5.602\,08\,\mathrm{kJ\,g^{-1}}$, and $Z = 0$ were chosen to match the values used in [10], based on the experimental conditions for the data included in these analyses. The remaining parameters were calibrated with the experimental data using the methods described below.

2.2 Data

The calibration used a heterogeneous experimental dataset with three different types of experiments. The experimental data include tests for the thermal expansion, specific heat capacity, and Hugoniot response of the HE and its components. Four Hugoniot [11–14], one thermal expansion [15], and one specific heat capacity [16] test were used. These data were collected over the span of several decades by different researchers across different facilities, so they represent a broad set of knowledge and testing capabilities. Additionally, the experimental data were provided without errors. In the present study, we assumed a fixed level of uncertainty for each data type to account for this lack of information. This assumption was made to focus on the goal of estimating the posterior covariance matrix using the three methodologies described in Section 3. An interesting direction of future work could be to estimate the unknown experimental errors from the data in the calibration procedure.

2.2.1 Hugoniot. The Hugoniot data were experimental measurements of shocked states of the high explosive material. These experiments measured the response of the reactants to shocks which did not immediately initiate a chemical reaction, and are meant to be representative of the solid reactants rather than the gaseous detonation products. The data were reported as sets of shock strengths and the associated particle speed imparted on the material by the shock. These experiments were simulated by the thermo-chemical code magpie [17]. The simulations resulted in shock strength-particle speed curves corresponding to each experimental dataset. These data and simulations were most informative for calibrating the parameters A, B, and C. The data were gathered by four different researchers who each used different approaches to compute the shock speed and particle speed and whose experiments were performed at different facilities over the course of several decades. There were no errors reported on any of these measurements, so a constant uncertainty of 10 % of the experimental observation was assumed. This level of uncertainty was chosen to reflect the natural spread in the data, as seen in Figure 2 and to have good overlap between the model evaluated with the parameters corresponding to the mode of the prior and most of the experimental data points.

2.2.2 Thermal Expansion. The thermal expansion data were obtained from experiments and were reported as functions of the isotropic coefficient of linear thermal expansion as a function of temperature. The coefficient of volumetric thermal expansion was three times this value and was used to create the change in density along a 1 atmosphere isobar. Again, magpie was used to simulate these experiments as density-temperature curves. These data and simulations were most informative for calibrating the A, B, C, and Γ_r^0 parameters. These data had no reported errors and

a 1 % uncertainty in the experimental observation was assumed. This level of uncertainty was chosen to have good agreement between the model evaluated with the parameters corresponding to the mode of the prior and the experimental data points as shown in Figure 3.

2.2.3 Specific heat. The specific heat data was obtained from chemical kinetic simulations of HMX molecules [18] whose results were converted to specific heat values in [16]. Magpie was used to simulate the specific heat capacity-temperature curves for these experiments. These data and simulations were most informative for calibrating the parameters C_v and α_{st}. These data did not have associated errors and a constant value of $\pm 400\,\mathrm{J\,kg^{-1}\,K^{-1}}$ was assumed. A constant uncertainty, rather than fraction value was used as the data contained values near zero, leading to spuriously low uncertainties in this region. This level of uncertainty was chosen to provide a match to the model evaluated with the parameters corresponding to the mode of the prior and the experimental data points as shown in Figure 3.

2.3 Bayesian Inverse Problem

The calibration is posed as a Bayesian inverse problem, seeking the most likely set of parameters θ given the experimental data Y. These values maximized the probability of the posterior distribution $\mathcal{P}(\theta|Y)$, defined by Bayes' Theorem,

$$\mathcal{P}(\theta|Y) \propto \mathcal{P}(Y|\theta)\mathcal{P}(\theta). \tag{9}$$

$\mathcal{P}(\theta)$ is the prior distribution over the model parameters and encodes knowledge about their values before considering data. $\mathcal{P}(Y|\theta)$ is the likelihood of the data given a set of parameters.

To calibrate the EOS, the prior and likelihood were chosen to be multivariate normal distributions based on the assumptions of the Variational Bayesian methods described in Sections 3.1 and 3.2. The prior mode values (see Table 1) were set based on a previous calibration of the EOS using similar datasets for the same HE [10]. To evaluate the likelihood, the experimental data and magpie simulation results were used. Additionally, it was assumed that the experimental uncertainty was uncorrelated. Including a model for the correlation between the uncertainties in different experiments could be a valuable direction of future work. Altogether, the posterior probability computed for all three methods is given as,

$$\log\left(\mathcal{P}\left(\theta|\mathcal{Y}\right)\right) = \sum_{k=1}^{n_k} \log\left(\mathcal{N}(\mu_k(\theta)|y_k, \Sigma_k)\right) + \log\left(\mathcal{N}(\theta|\hat{\theta}, \Sigma_\theta)\right), \tag{10}$$

where the first term comes from the likelihood, Σ_k is the covariance matrix for the k^{th} experiment y_k, and $\mu_k(\theta)$ is the mapping function between the model terms and the experimental data. The second term is from the prior where $\hat{\theta}$ is the mode of the prior and Σ_θ is the prior covariance matrix.

The model parameters were bounded to be within ±75% of the prior values. This was chosen to ensure the parameter values were feasible and positive. Additionally, since the prior parameter values were selected from a previous calibration [10], the posterior values were not expected to change significantly.

The posterior distribution mode and covariance matrix were estimated using three methods, a Variational Bayesian approach [19], computing the Hessian directly, and sampling via Markov Chain Monte Carlo. Sections 3.1, 3.2, and 3.3 provide additional details.

3. METHODS

3.1 Variational Bayesian

The first method used to compute the posterior and estimate the posterior covariance was the Variational Bayesian (VBayes) approach detailed in [19]. Briefly, the method assumes that the prior and likelihood are normal distributions and formulates the calibration as an optimization problem. This method makes a major assumption that there is an affine mapping between the model parameters θ and the experimental data y_k for each of the k experiments,

$$\theta \xrightarrow{\mu_k} y_k, \tag{11}$$

and that this mapping μ_k can be represented as a first order Taylor series approximation of some general non-linear function f

$$\mu_k(\theta + \eta) = f(\theta) + \eta \underbrace{\frac{\partial f}{\partial \theta}}_{D_k} + \mathcal{O}(\eta^2), \tag{12}$$

where η is a small perturbation in the model parameters θ and D_k is the sensitivity of the k^{th} mapping function to θ. This assumes that the underlying function is linear, while the true function which used the code magpie to evaluate the thermodynamic loci described in Section 2.2 was non-linear. Subsequent tests will examine the suitability of this assumption for these data.

With these assumptions, the posterior probability of (9, 10) is reduced to a quadratic problem which can be solved with very few evaluations of the mapping function μ.

The posterior covariance is then reduced to,

$$\Sigma_p^{-1} = \Sigma_\theta^{-1} + \sum_{k=1}^{n_k} D_k^T \Sigma_k^{-1} D_k, \tag{13}$$

where Σ_θ, Σ_k, and Σ_p are the covariance matrices for the prior, k^{th} experiment, and posterior, respectively. This form assumes the expected value of the error $\mu_k(\theta) - y_k$ is zero.

This approach is advantageous when the function is computationally expensive as the posterior mode and covariance can be computed without having to compute a Hessian directly. However, the assumption in (12) ignores any non-linearities in the mapping function and this may affect estimates of the posterior covariance. Previous work [19] has compared this implementation of the VBayes method to MCMC for an EOS model with synthetic data, but a direct comparison has not been made for EOS models using experimental data.

3.2 Hessian Variational Bayesian (SQP)

If the assumption in (12) is not made, then a traditional optimization method can be used to find the point which maximizes the log-posterior probability (10). Though any constrained optimization method could be used, for this work the SLSQP [20] method as implemented by Scipy [21] was used.

The Laplace approximation can again be used to estimate the posterior covariance, but here the Hessian must be evaluated directly according to

$$\Sigma_p^{-1} = \frac{\partial^2 P\left(\theta|Y\right)}{\partial\theta^2}. \tag{14}$$

Without the approximations discussed in Section 3.1, the Hessian Variational Bayesian method requires more function evaluations than the VBayes approach. Additionally, after the optimization is complete the Hessian has to be evaluated in (14) to calculate the posterior covariance, which is a $O\left(m^2\right)$ operation.

Note, this approach does not assume an affine mapping between the model parameters θ and the data y_k and the non-linearities of this mapping function are represented in the optimal solution and its posterior distribution.

3.3 Markov Chain Monte Carlo

An alternative approach to calculate the posterior distribution is using Markov Chain Monte Carlo (MCMC) methods. Here a Metropolis MCMC algorithm [22] which only uses feasible draws from the proposal distribution was used to construct the posterior. The proposal distribution was a multivariate, zero mean normal with a scaled version of the prior covariance matrix. The prior covariance matrix was scaled by a factor of 0.4 to have a sampling acceptance rate of 25.66 %, which is close to the best practice acceptance rate given in [23]. Five chains were run for 30,000 samples each, resulting in 150,000 samples overall. The chains were started at a small random perturbation of the prior values in Table 1. A conservative burn-in period of 15,000 samples was selected. Based on the procedure from [23] and the implementation in the numpyro python package [24, 25], the effective number of independent samples per chain was 1,500, or a step size of 10 between samples to ensure they are uncorrelated. All subsequent analyses used the post burn-in, independent samples.

Figure 7 shows the MCMC chains with the post-processed samples for each parameter, and demonstrates that for each parameter, there is good mixing within and between the chains. Figure 1 shows the marginal PDFs of the MCMC samples as histograms on the diagonal and the correlation between samples in the off-diagonal. Each of the marginal histograms indicates the samples are roughly normally distributed.

4. RESULTS

4.1 Optimal Models

The mode of the prior and posterior distribution for each model parameter is reported in Table 1. There was very little difference between the optimal values obtained from the three different calibration methods, and all three calibrations were in close agreement with the previous calibration [10], which was used as the prior. Figures 2 and 3 show that the models evaluated using the model parameter sets corresponding to the mode of the posterior distributions are similar to the models evaluated with the prior mode parameters. This was expected since the prior mode values are from a previous calibration. For the Hugoniot data, all four models are consistent at low particle speeds. At the higher particle speeds, the prior model captures the higher values of associated shock speeds, the MCMC posterior model captures the lower shock speeds, and the variational posterior models are in between. For the specific heat data, all four models evaluated with the distribution modes are consistent at the lower temperatures and the MCMC posterior model predicts slightly lower specific heat capacity values than the other three models at the higher temperatures. With the thermal expansion data, the prior and two variational posterior models are consistent across the whole range of temperatures considered, whereas the MCMC posterior model captures slightly lower densities at the lower temperatures and slightly higher densities at the higher temperatures.

TABLE 1: MODEL PARAMETER PRIOR AND POSTERIOR MODES.

	Prior	VBayes	SQP	MCMC
A / km s^{-1}	2.300	2.329	2.314	2.284
B	3.300	2.931	2.944	2.756
C	0.200	0.199	0.196	0.181
Γ	0.838	0.828	0.824	0.649
C_v / kJ g^{-1} K^{-1}	0.00107	0.00106	0.00104	0.00104
α_{st}	0.366	0.388	0.393	0.358

4.2 Posterior Covariance

For the Variational Bayesian method, the Laplace Approximation was used to estimate the posterior covariance matrix. For the Hessian Variational Bayesian method, the Laplace approximation with the Hessian computed directly was used to estimate the posterior covariance. Finally, for the MCMC approach, the posterior covariance matrix was estimated from the samples of the posterior. The covariance matrices Σ can be decomposed into a correlation matrix R and diagonal matrix with the posterior variance Λ [26].

$$\Sigma = \Lambda R \iff R = \Lambda^{-\frac{1}{2}}\Sigma\Lambda^{-\frac{1}{2}} \tag{15}$$

The correlation matrices found with each approach for approximating the posterior covariance matrix are shown in Figure 4. All three have a similar structure, with some differences in the magnitudes of the weaker correlations between parameters. The parameters A and B are strongly negatively correlated, as are C_v and α_{st}. The other parameters are weakly or not correlated. These results agree with the scatter plots in Figure 1.

To assess the marginal uncertainty in each of the model parameters, the standard deviation of the posterior was computed from the diagonal matrices with the variance. These are shown in Figure 5. The parameters C, Γ_r^0, α_{st}, and B have the largest uncertainties, while A and C_v have the lowest. Even the largest uncertainty is only about 10%.

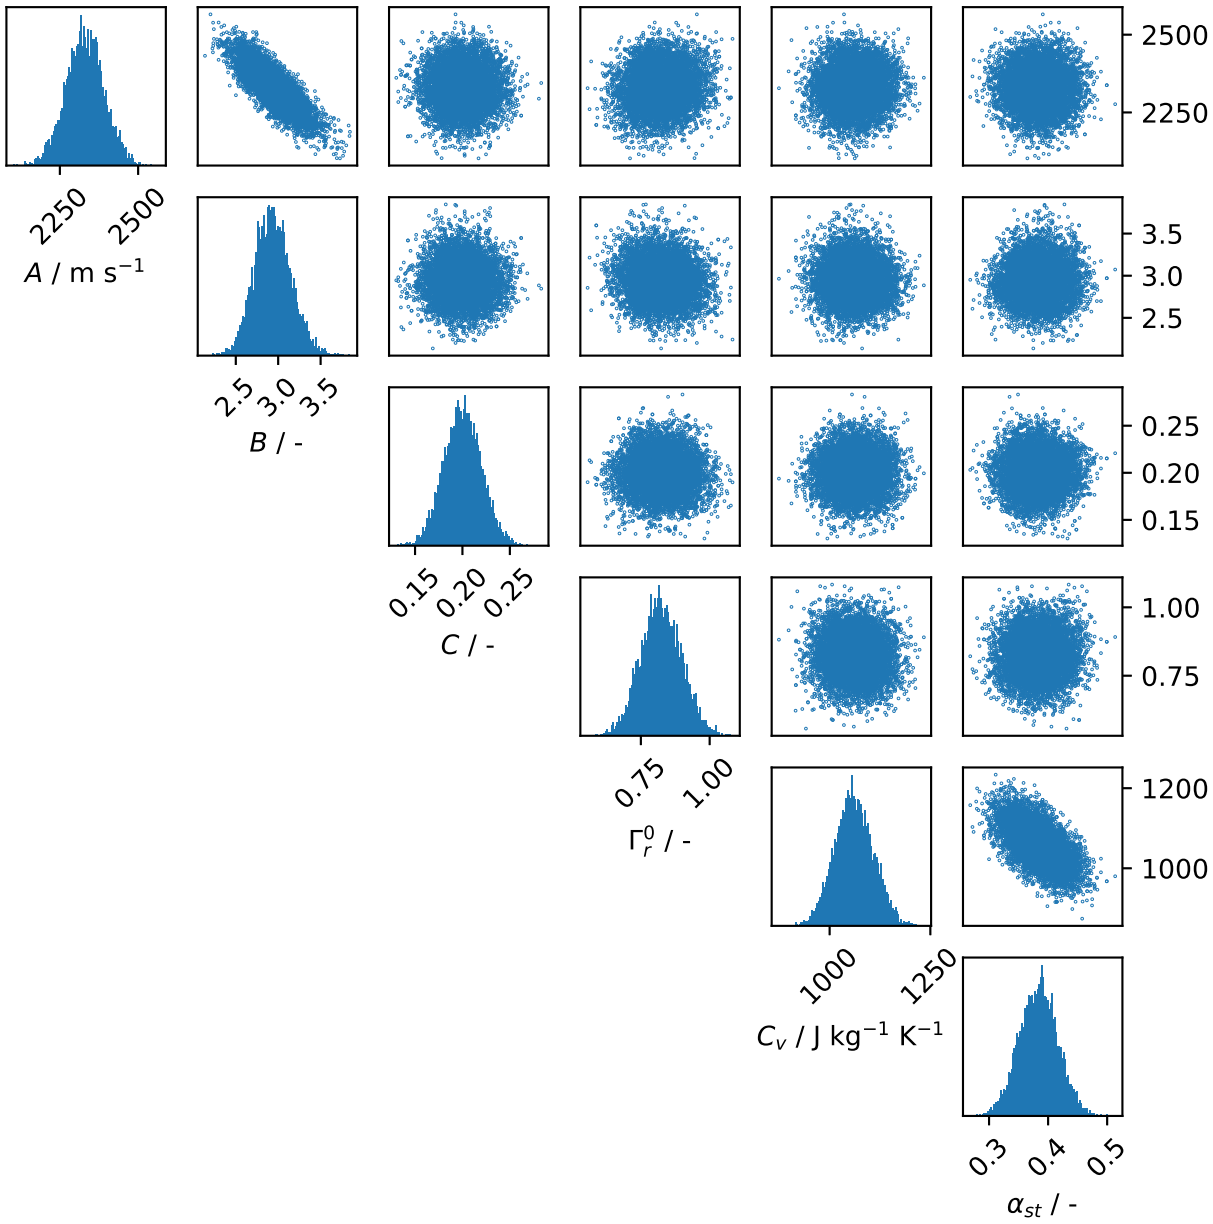

FIGURE 1: DIAGONAL: MARGINAL HISTOGRAMS OF MCMC SAMPLES AFTER POST PROCESSING. OFF DIAGONAL: CORRELATION BETWEEN SAMPLES. THE MARGINAL HISTOGRAMS SHOW ALL THE PARAMETERS ARE ROUGHLY NORMALLY DISTRIBUTED. THE SCATTER PLOTS SHOW THAT MOST OF THE PARAMETERS ARE UNCORRELATED, EXCEPT FOR NEGATIVE CORRELATIONS BETWEEN *A* AND *B*, AND C_v AND α_{st}.

A spectral decomposition analysis was also performed on the posterior covariance matrices to understand their structure according to the eigenvalues and eigenvectors. Figure 6 indicates that the eigenvalues of the three covariance matrices are almost identical since the five eigenvalue ranks are approximately the same value for each of the three methods.

Figure 8 shows the first four eigenvectors of the posterior covariance matrices broken down by vector component and matrix approximation method. Each component in the eigenvectors corresponds to one of the calibrated EOS parameters. The eigenvectors are similar across the three methods in that the dominant component of each eigenvector is the same regardless of the matrix approximation method. Moreover, the dominant components of each eigenvector are aligned with the parameters which have the largest uncertainties shown in Figure 5. That is, the largest component of the first eigenvector is related to the parameter C, in the second and third eigenvectors the largest components are Γ_r^0 and α_{st}, and the fourth eigenvector has its largest component associated with the B parameter. This structure is reflected in all three methods of estimating the posterior covariance matrix, with only small differences in the magnitudes.

All three methods estimated similar values for the mode of the

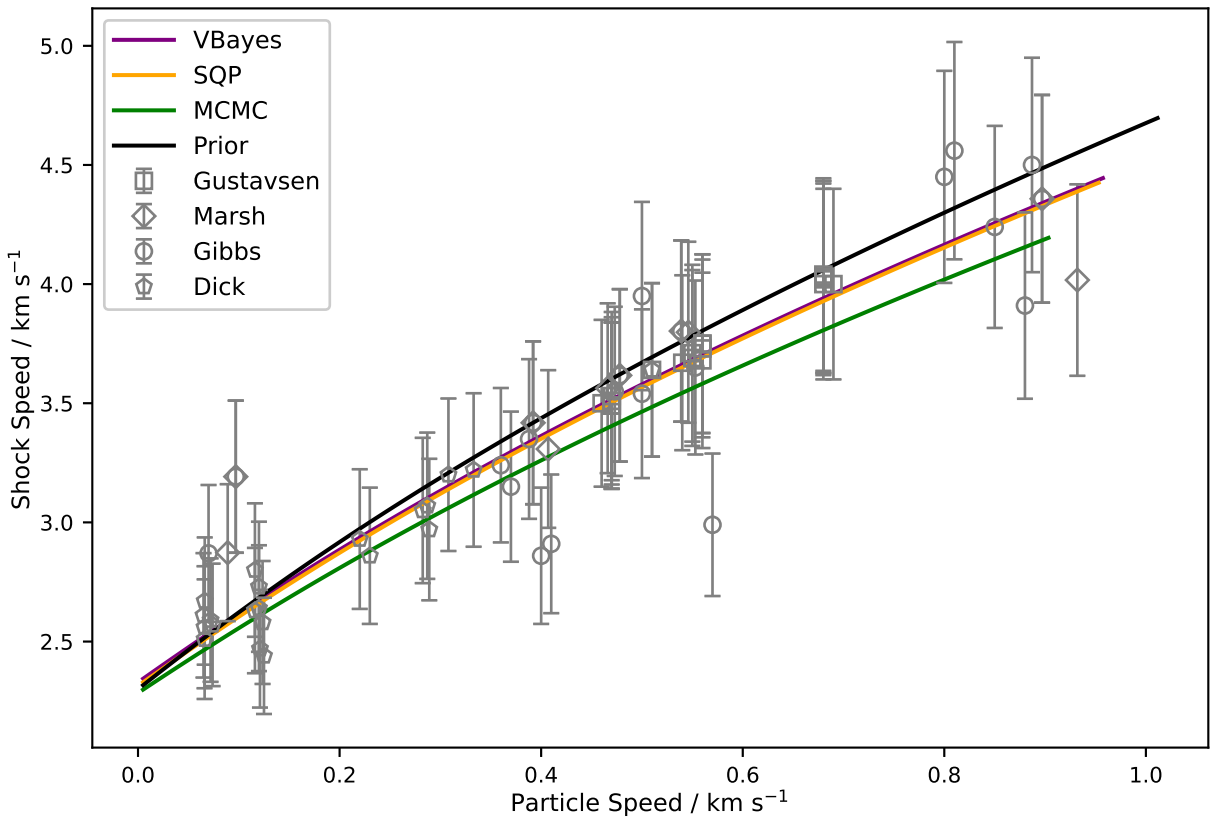

FIGURE 2: HUGONIOT DATA WITH EOS MODEL EVALUATED USING MODEL PARAMETER SETS CORRESPONDING TO THE MODE OF THE PRIOR AND POSTERIOR DISTRIBUTIONS. DATA POINTS FROM THE FOUR HUGONIOT EXPERIMENTS (GRAY SHAPE MARKERS) ARE SHOWN WITH THE ASSUMED 10% EXPERIMENTAL UNCERTAINTY (GRAY ERROR BARS). THE BLACK LINE IS THE MODEL EVALUATED WITH THE PRIOR PARAMETER VALUES. THE MODEL EVALUATED WITH THE VBAYES, SQP, MCMC POSTERIOR PARAMETER VALUES ARE THE SOLID PURPLE, ORANGE, AND GREEN LINES, RESPECTIVELY.

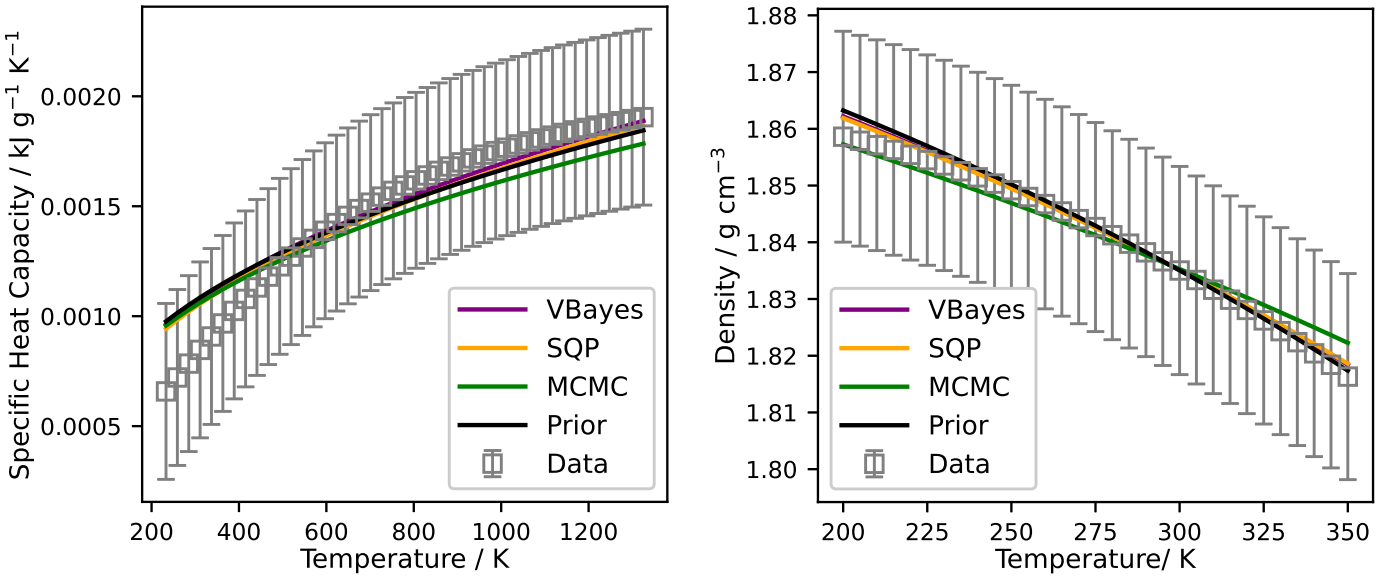

(a) Specific Heat Capacity [16]

(b) Thermal Expansion [15]

FIGURE 3: MODEL FIT TO THE SPECIFIC HEAT CAPACITY AND THERMAL EXPANSION DATASETS.

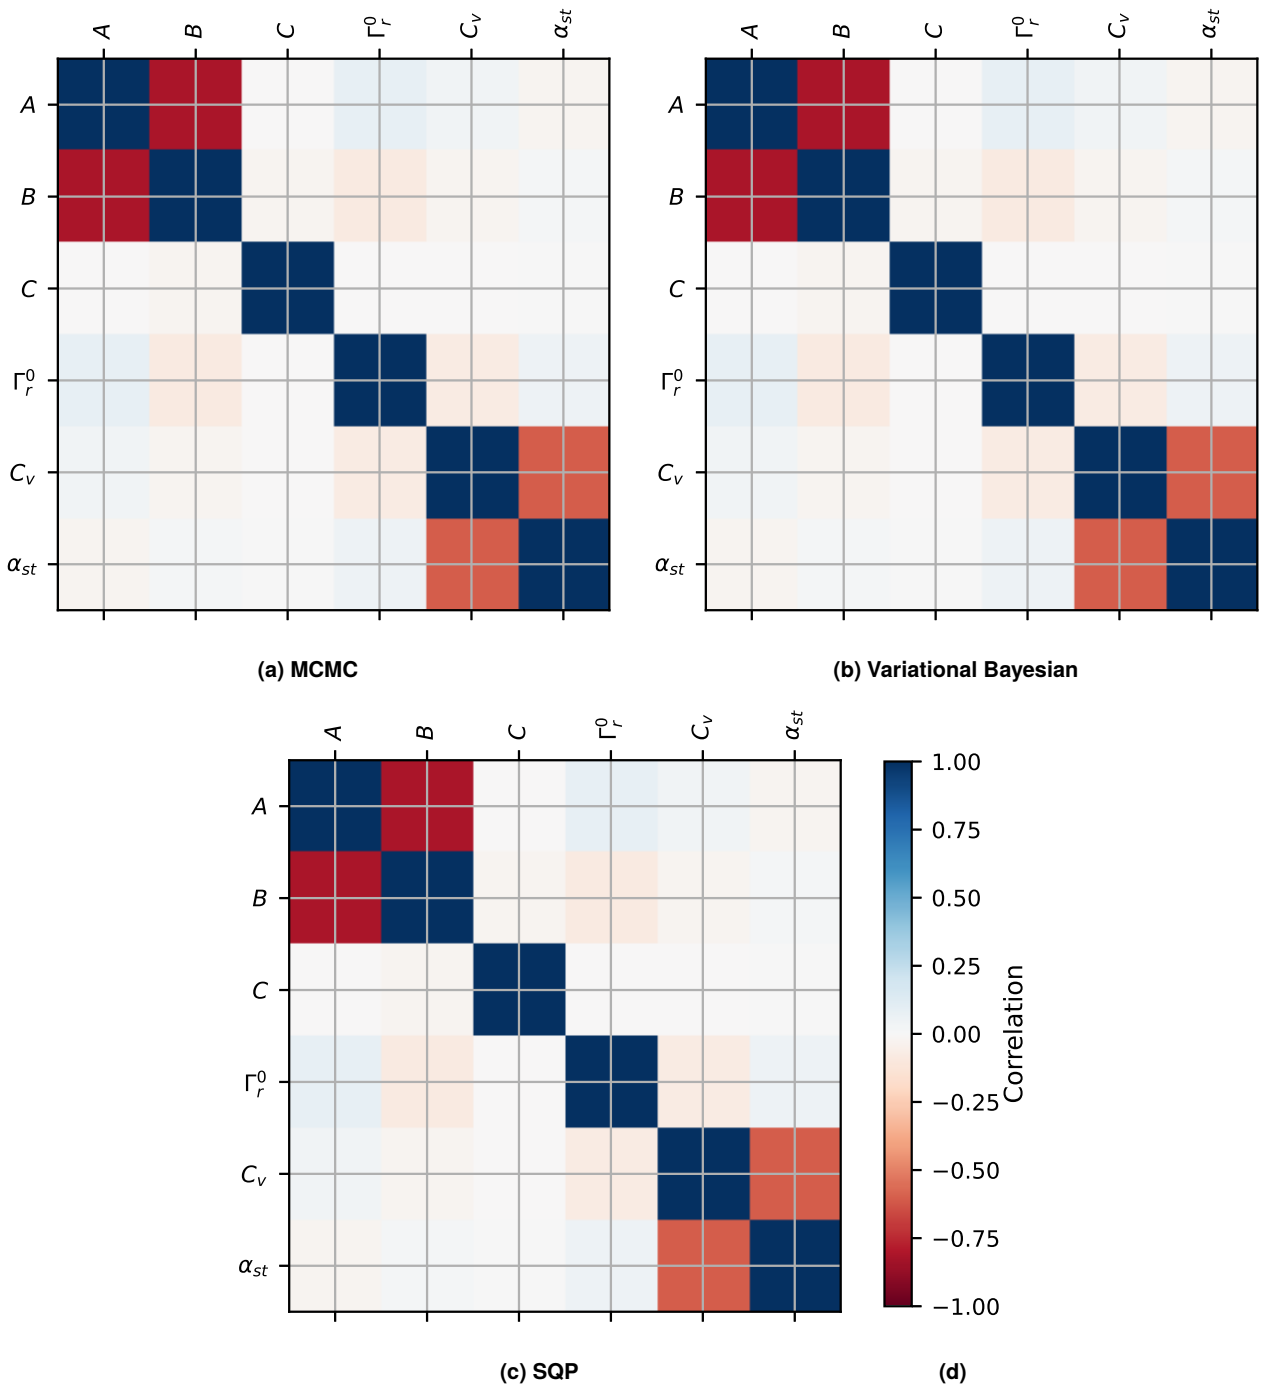

FIGURE 4: CORRELATION MATRICES. DARK BLUE INDICATES STRONG POSITIVE CORRELATION AND DARK RED IS STRONG NEGATIVE CORRELATION.

posterior distributions, which when used to evaluate the EOS model captured the behavior of the experimental data well. Additionally, the spectral decomposition revealed that the posterior covariance matrices had similar structures, regardless of the method used to estimate the matrix.

In terms of computational cost, the VBayes method used 34 function evaluations, the SQP optimization used 80 function evaluations, and the MCMC algorithm used 1,500 independent function evaluations per chain for a total of 7,500 independent function evaluations. This analysis used a simple Metropolis implementation of the MCMC algorithm, which is not the state of the art for these methods in terms of efficiency. In order for a state of the art MCMC algorithm to be more efficient than the Variational approaches, the independent function evaluations would need to be reduced by two orders of magnitude compared with the Metropolis approach used here.

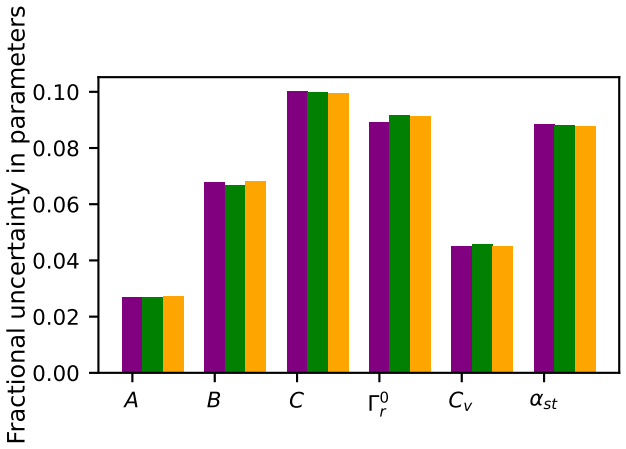

FIGURE 5: MARGINAL UNCERTAINTY ON EACH PARAMETER - POSTERIOR STANDARD DEVIATION NORMALIZED BY THE PRIOR PARAMETER VALUE. THE MCMC UNCERTAINTIES ARE SHOWN IN PURPLE, THE VARIATIONAL BAYESIAN UNCERTAINTIES ARE SHOWN IN GREEN, AND THE SQP UNCERTAINTIES IN ORANGE.

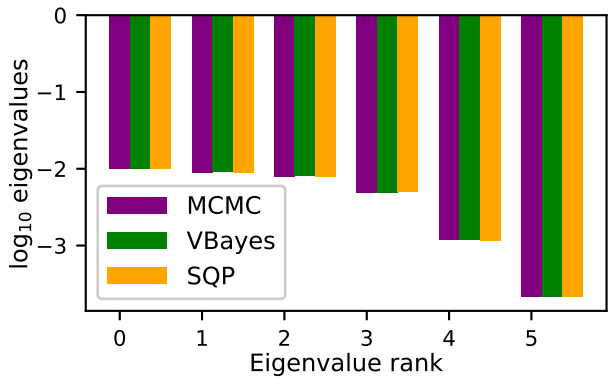

FIGURE 6: COMPARISON OF EIGENVALUES FROM SPECTRAL DECOMPOSITION. THE MCMC EIGENVALUES ARE SHOWN IN PURPLE, THE VARIATIONAL BAYESIAN EIGENVALUES ARE SHOWN IN GREEN, AND THE SQP EIGENVALUES IN ORANGE.

Both the vector of mode values and the structure of the covariance matrices were very similar for the three methods considered. These results are encouraging as they show that the assumptions made by the methods described in 3.1 do not significantly impact the assessment of the magnitude or shape of the uncertainty in the model parameters. However, there were some features of the EOS calibration problem examined in this paper which may have improved the agreement between the methods and which may not be generally applicable to all problems. The choice of using a normal distribution for the prior and likelihood in the MCMC approach was made in order to match the assumptions of the Variational methods and keep the problem the same across methods for the sake of comparison. MCMC, generally, can represent much more complicated posterior distributions if different choices are made for the probability distribution of the prior and likelihood. However, in this case, there was no feature of the data or the model which suggested the use of anything other than a normal distribution, with one exception, the parameters Γ_r^0 and α_{st}. These parameters are dimensionless, of order one, and positive, which may have been better represented by an inverse gamma distribution, for example. However, considering the marginal distributions in Figure 1, none of the samples for these variables approach zero, so there is almost zero probability contained in the infeasible region. In this case, the normal distribution remains a good choice for these variables and there is no need to consider a distribution which excludes negative values. The other feature of the data which was advantageous was that these data were all obtained from high quality calibration experiments. It was reasonable to expect that a single model could fit all the experiments from the heterogeneous data set and the results shown in Figures 2 and 3 show that all three methods found models which were a good fit to all the available data with no significant outliers. This outcome was expected from this dataset but is not generally true for all datasets and the presence of a significant bias between sets of experimental data would violate the assumption made in Section 3.1 that the expected value of the error between the mapping function and the data was zero.

5. CONCLUSIONS

This paper provides a comparison between three methods of calibrating EOS model parameters and approximating the posterior covariance matrix. The calibration was posed as an Bayesian inverse problem leveraging a heterogeneous dataset to estimate the EOS parameters with uncertainty and the posterior covariance matrix.

This analysis found that the Variational Bayesian, Hessian Variational Bayesian, and MCMC methods had good agreement in the mode posterior values for each parameter. Evaluating the EOS at the mode posterior values for the parameters resulted in as good of a fit to the experimental data as the model evaluated with the high quality prior mode values. Additionally, all three methods estimated posterior covariance matrices with similar structures. The uncertainties in the parameters and spectral analysis results were also consistent across the three methods. Estimating the posterior distribution and posterior covariance matrix with the two Variational methods did not lead to a loss or change in information about the structure of the posterior covariance matrix compared to using MCMC.

In this setting where the parameters, likelihood, and data can be modeled as multivariate normal distributions, and the experimental data are consistent, the calibration and uncertainty quantification of parameters in parametric physics models via Variational methods and the Laplace approximation is reasonable. These more computationally efficient methods may be advantageous when the posterior probability function is too computationally expensive to evaluate using MCMC methods.

NOMENCLATURE

$\mathcal{N}$	Normal distribution	
α_{st}	Describes how the specific heat changes with temperature	
β	Volumetric thermal expansion coefficient	K^{-1}
Σ_θ	Covariance matrix for prior	
η	Perturbation to the θ vector	
Σ_k	Covariance matrix for k^{th} experiment	
Σ_p	Covariance matrix for posterior	
Γ_r^0	Grüneisen gamma	
θ	Set of model parameters	
ρ_0	Reference ambient density	$\mathrm{kg\,m}^{-3}$
$\hat{\theta}$	Mode of the prior distribution	
μ_k	Mapping function for k^{th} experiment	
A	Sound speed in the material,	$\mathrm{m\,s}^{-1}$
B	Slope of the Hugoniot curve	
C	High particle speed slope correction	
C_v	Material specific heat capacity at the reference temperature for constant volume	$\mathrm{J\,kg}^{-1}\,\mathrm{K}^{-1}$
D_k	Sensitivity of the k^{th} mapping function to θ	
E_0	Stored chemical potential energy	$\mathrm{kJ\,g}^{-1}$
f_k	Non-linear function for simulating the k^{th} experiment	
P	Pressure	Pa
$\mathcal{P}$	Probability	
R	Correlation matrix	
T	Temperature	K
T_0	Reference temperature	K
Y	Set of all experimental data	
y	Compression factor	
y_k	Data for k^{th} experiment	
Z	Describes how the Grüneisen gamma changes with density	

REFERENCES

[1] Davis, William C. "Complete Equation of State for Unreacted Solid Explosive." *Combustion and Flame* Vol. 120 No. 3 (2000): pp. 399–403. DOI 10.1016/S0010-2180(99)00112-1.

[2] Wescott, B. L., Stewart, D. Scott and Davis, W. C. "Equation of state and reaction rate for condensed-phase explosives." *Journal of Applied Physics* Vol. 98 No. 5 (2005): p. 053514. DOI 10.1063/1.2035310.

[3] Aslam, Tariq D. "Shock temperature dependent rate law for plastic bonded explosives." *Journal of Applied Physics* Vol. 123 No. 14 (2018): p. 145901. DOI 10.1063/1.5020172.

[4] Lindquist, Beth A. and Jadrich, Ryan B. "Uncertainty quantification for a multi-phase carbon equation of state model." *Journal of Applied Physics* Vol. 131 No. 15 (2022): p. 155104. DOI 10.1063/5.0087210.

[5] Lindquist, Beth A., Jadrich, Ryan B., Heras Rivera, Juampablo E. and Rondini, Lucia I. "Uncertainty quantification for high explosive reactant and product equations of state." *Journal of Applied Physics* Vol. 134 No. 7 (2023): p. 075102. DOI 10.1063/5.0157842.

[6] Blei, David M., Kucukelbir, Alp and McAuliffe, Jon D. "Variational Inference: A Review for Statisticians." *Journal of the American Statistical Association* Vol. 112 No. 518 (2017): pp. 859–877. DOI 10.1080/01621459.2017.1285773.

[7] Gunapati, Geetakrishnasai, Jain, Anirudh, Srijith, P. K. and Desai, Shantanu. "Variational inference as an alternative to MCMC for parameter estimation and model selection." *Publications of the Astronomical Society of Australia* Vol. 39 (2022): p. e001. DOI 10.1017/pasa.2021.64.

[8] Rizzato, Matteo and Sellentin, Elena. "Extremely expensive likelihoods: a variational-Bayes solution for precision cosmology." *Monthly Notices of the Royal Astronomical Society* Vol. 521 No. 1 (2023): pp. 1152–1161. DOI 10.1093/mnras/stad638.

[9] Natesan, Prathiba, Nandakumar, Ratna, Minka, Tom and Rubright, Jonathan D. "Bayesian Prior Choice in IRT Estimation Using MCMC and Variational Bayes." *Frontiers in Psychology* Vol. 7 (2016). DOI 10.3389/fpsyg.2016.01422.

[10] Aslam, T. D., Price, Matthew A., Ticknor, Christopher, Coe, Joshua D., Leiding, Jeffery A. and Zocher, Marvin A. "AWSD callibration for the HMX based explosive PBX 9501." *AIP Conference Proceedings* Vol. 2272 No. 030001 (2020). DOI 10.1063/12.0000891.

[11] Marsh, P., Stanley. *LASL Shock Hugoniot Data*. University of California Press (1980).

[12] Gibbs, R., Terry and Popolato, Alphonse. *LASL Explosive Property Data*. University of California Press (1980).

[13] Gustavsen, Richard L., Sheffield, Stephen A. and Alcon, Robert R. "Shock Initiation of New and Aged Pbx 9501 Measured with Embedded Electro- Magnetic Particle Velocity Gauges." LA- 13634-MS. Los Alamos National Lab. 1999. DOI 10.2172/10722.

[14] Dick, J. J., Martinez, A. R. and Hixson, R. S. "Plane impact response of PBX 9501 and its components below 2 GPa." Technical Report No. LA-13426-MS. Los Alamos National Laboratory, Los Alamos, NM. 1998. DOI 10.2172/663187.

[15] Thompson, Darla G., Brown, Geoffery W., DeLuca, Racci, Giambra, Anna M. and Sandstrom, Mary M. "Thermal Expansion of PBX 9501 and PBX 9502 Plastic-Bonded Explosives." *Proceedings of the North American Thermal Analysis Society*. 2009. Lubbock TX.

[16] Menikoff, R. "Detonation waves in PBX 9501." *Combustion Theory and Modelling* Vol. 10 No. 6 (2006): pp. 1003–1021. DOI 10.1080/13647830600851754.

[17] Ticknor, C., Andrews, S. A. and Leiding, J. A. "Magpie: A new thermochemical code." *AIP Conference Proceedings* Vol. 2272 No. 1 (2020): p. 030033. DOI 10.1063/12.0000785.

[18] Lyman, John L, Liau, Yeong-Cherng and Brand, Holmann V. "Thermochemical functions for gas-phase, 1,3,5,7-tetranitro-1,3,5,7-tetraazacyclooctane (HMX), its condensed phases, and its larger reaction products." *Combustion and Flame* Vol. 130 No. 3 (2002): pp. 185–203. DOI 10.1016/S0010-2180(02)00364-4.

[19] Andrews, Stephen A. and Fraser, Andrew M. "Estimating physics models and quantifying their uncertainty using optimization with a bayesian objective functional." *ASME Journal of Verification Validation and Uncertainty Quantification* Vol. 4 No. 1 (2019): pp. 011002–1 – 01102–15. DOI 10.1115/1.4043807.

[20] Kraft, Dieter. "Algorithm 733: TOMP–Fortran Modules for Optimal Control Calculations." *ACM Trans. Math. Softw.* Vol. 20 No. 3 (1994): p. 262–281. DOI 10.1145/192115.192124.

[21] Virtanen, Pauli, Gommers, Ralf, Oliphant, Travis E., Haberland, Matt, Reddy, Tyler, Cournapeau, David, Burovski, Evgeni, Peterson, Pearu, Weckesser, Warren, Bright, Jonathan, van der Walt, Stéfan J., Brett, Matthew, Wilson, Joshua, Millman, K. Jarrod, Mayorov, Nikolay, Nelson, Andrew R. J., Jones, Eric, Kern, Robert, Larson, Eric, Carey, C J, Polat, İlhan, Feng, Yu, Moore, Eric W., VanderPlas, Jake, Laxalde, Denis, Perktold, Josef, Cimrman, Robert, Henriksen, Ian, Quintero, E. A., Harris, Charles R., Archibald, Anne M., Ribeiro, Antônio H., Pedregosa, Fabian, van Mulbregt, Paul and SciPy 1.0 Contributors. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python." *Nature Methods* Vol. 17 (2020): pp. 261–272. DOI 10.1038/s41592-019-0686-2.

[22] Metropolis, Nicholas, Rosenbluth, Arianna W., Rosenbluth, Marshall N., Teller, Augusta H. and Teller, Edward. "Equation of State Calculations by Fast Computing Machines." *The Journal of Chemical Physics* Vol. 21 No. 6 (1953): pp. 1087–1092. DOI 10.1063/1.1699114.

[23] Gelman, Andrew, Carlin, John B., Stern, Hal S., Dunson, David B., Vehtari, Aki and Rubin, Donald B. *Bayesian Data Analysis*, 3rd ed. CRC Press, Boca Raton, FL (2013).

[24] Phan, Du, Pradhan, Neeraj and Jankowiak, Martin. "Composable Effects for Flexible and Accelerated Probabilistic Programming in NumPyro." (2019). URL arXiv:1912.11554.

[25] Bingham, Eli, Chen, Jonathan P., Jankowiak, Martin, Obermeyer, Fritz, Pradhan, Neeraj, Karaletsos, Theofanis, Singh, Rohit, Szerlip, Paul A., Horsfall, Paul and Goodman, Noah D. "Pyro: Deep Universal Probabilistic Programming." *J. Mach. Learn. Res.* Vol. 20 (2019): pp. 28:1–28:6. URL http://jmlr.org/papers/v20/18-403.html.

[26] Andrews, Stephen A. and Wilson, Brandon M. "Variational Bayesian Calibration of a PTW Material Strength Model of OFHC Copper." *Proceedings of ASME 2023 Verification, Validation and Uncertainty Quantification Symposium*. VVS2023-108829. 2023. DOI 10.1115/VVUQ2023-108829.

[27] Hunter, John D. "Matplotlib: a 2D graphics environment." *Computing in Science & Engineering* Vol. 9 (2007): pp. 21–29. DOI 10.1109/MCSE.2007.55.

[28] Harris, Charles R., Millman, K. Jarrod, van der Walt, Stéfan J, Gommers, Ralf, Virtanen, Pauli, Cournapeau, David, Wieser, Eric, Taylor, Julian, Berg, Sebastian, Smith, Nathaniel J., Kern, Robert, Picus, Matti, Hoyer, Stephan, van Kerkwijk, Marten H., Brett, Matthew, Haldane, Allan, Fernández del Río, Jaime, Wiebe, Mark, Peterson, Pearu, Gérard-Marchant, Pierre, Sheppard, Kevin, Reddy, Tyler, Weckesser, Warren, Abbasi, Hameer, Gohlke, Christoph and Oliphant, Travis E. "Array programming with NumPy." *Nature* Vol. 585 (2020): p. 357–362. DOI 10.1038/s41586-020-2649-2.

ACKNOWLEDGEMENTS

The authors acknowledge support by the Department of Energy, National Nuclear Security Administration, Predictive Science Academic Alliance Program (PSAAP) under Award Number DE-NA0003962 and by the Department of Energy's Advanced Simulation and Computing program. This work was performed at the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001)

The authors would like to thank the developers of the following open source software packages which were used in the preparation of this manuscript: matplotlib [27], numpy [28], and scipy [21]. Approved for public release LA-UR-23-31558 and DOPSIR 24-T-0067.

APPENDIX A. MCMC POST PROCESSING

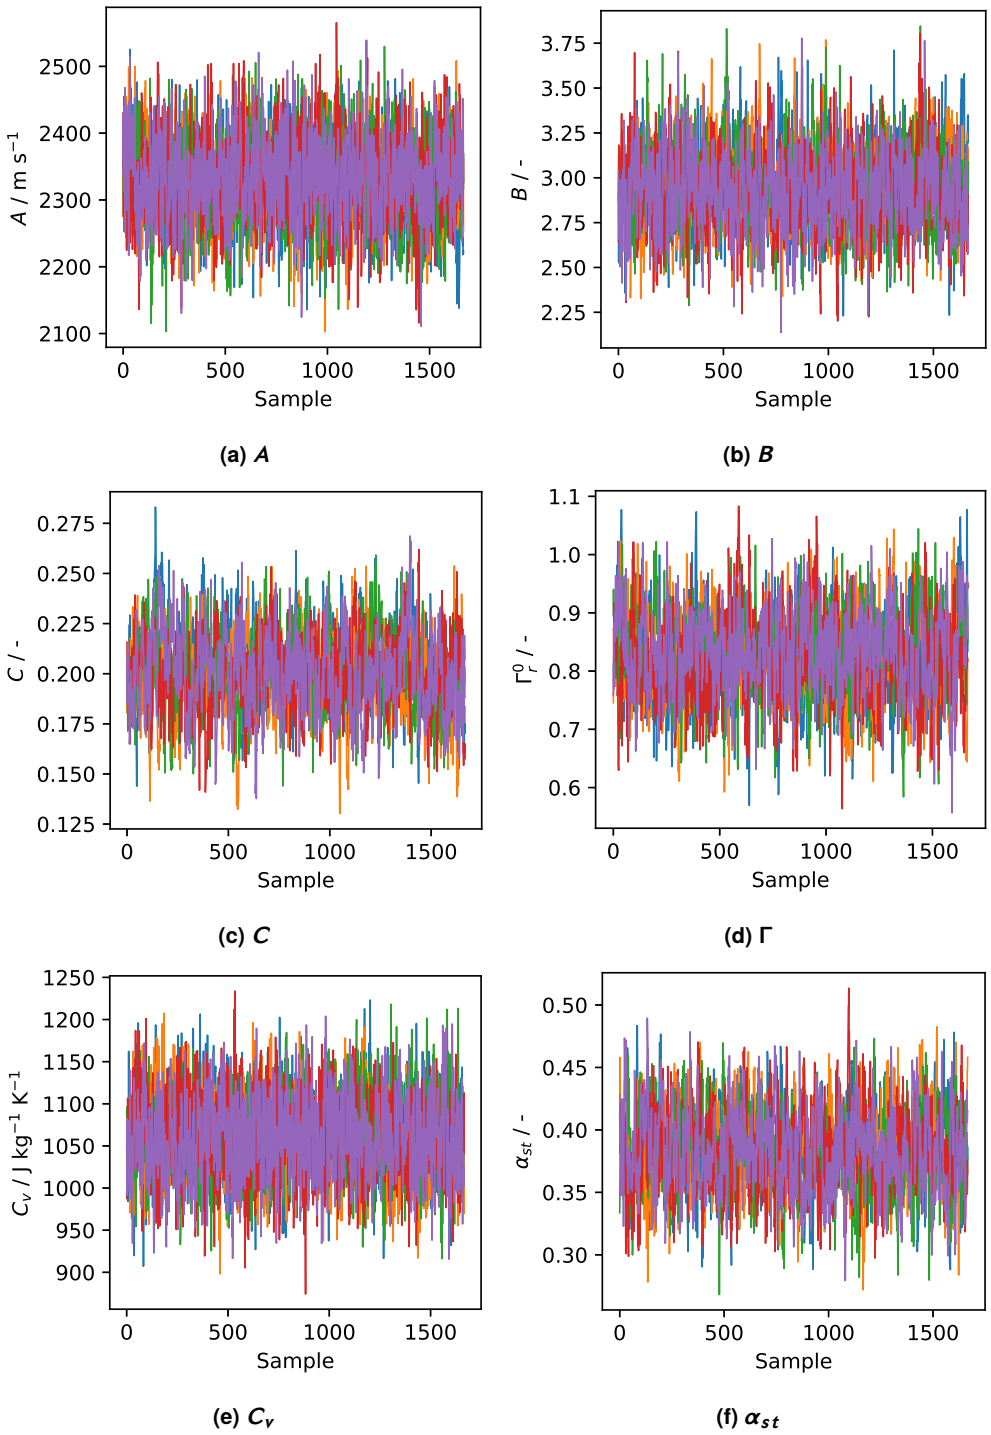

FIGURE 7: INDEPENDENT SAMPLES FROM THE MCMC CHAINS AFTER THE BURN-IN PERIOD SHOWING GOOD MIXING WITHIN AND BETWEEN CHAINS. EACH COLOR REPRESENTS A DIFFERENT CHAIN.

APPENDIX B. SPECTRAL DECOMPOSITION

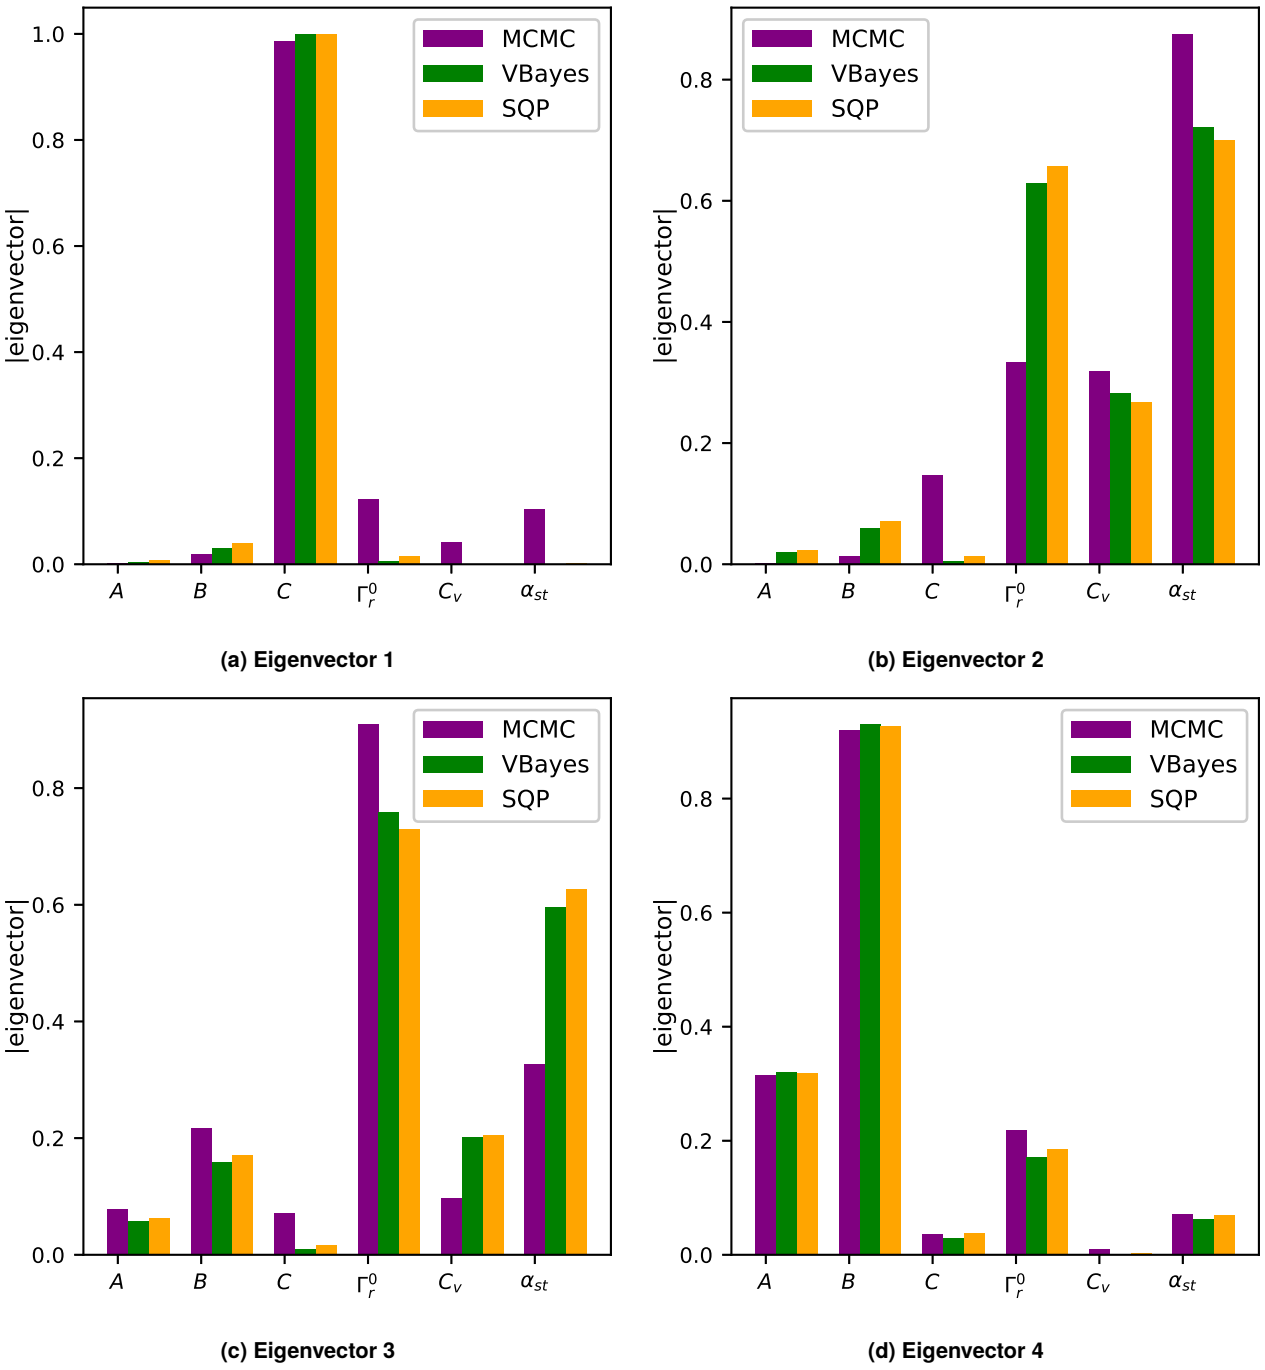

(a) Eigenvector 1

(b) Eigenvector 2

(c) Eigenvector 3

(d) Eigenvector 4

FIGURE 8: ABSOLUTE VALUE OF THE FIRST FOUR EIGENVECTORS FROM SPECTRAL DECOMPOSITION. THE MCMC EIGENVECTORS ARE SHOWN IN PURPLE, THE VARIATIONAL BAYESIAN EIGENVECTORS ARE SHOWN IN GREEN, AND THE SQP EIGENVECTORS IN ORANGE.