

Robust Bayesian Calibration of a RANS Model for Jet-in-Crossflow Simulations

Jaideep Ray * and Sophia Lefantzi †

Sandia National Laboratories, Livermore, CA 94550-0969, USA

Srinivasan Arunajatesan ‡ and Lawrence Dechant‡

Sandia National Laboratories, Albuquerque, NM 87185-5800, USA

Compressible jet-in-crossflow interactions are poorly simulated using Reynolds-Averaged Navier Stokes (RANS) equations. This is due to model-form errors (physical approximations) in RANS as well as the use of parameter values simply picked from literature (henceforth, the nominal values of the parameters). Previous work on the Bayesian calibration of RANS models has yielded joint probability densities of $\mathbf{C} = (C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$, the most influential parameters of the RANS equations. The calibrated values were far more predictive than the nominal parameter values and the advantage held across a range of freestream Mach numbers and jet strengths. In this work we perform Bayesian calibration across a range of Mach numbers and jet strengths and compare the joint densities, with a view of determining whether compressible jet-in-crossflow could be simulated with either a single joint probability density or a point estimate for \mathbf{C} . We find that probability densities for $C_{\epsilon 2}$ agree and also indicate that the range typically used in aerodynamic simulations should be extended. The densities for $C_{\epsilon 1}$ agree, approximately, with the nominal value. The densities for C_μ do not show any clear trend, indicating that they are not strongly constrained by the calibration observables, and in turn, do not affect them much. We also compare the calibrated results to a recently developed analytical model of a jet-in-crossflow interaction. We find that the values of \mathbf{C} estimated by the analytical model delivers prediction accuracies comparable to the calibrated joint densities of the parameters across a range of Mach numbers and jet strengths.

Nomenclature

$(\mathbf{u}_e, \mathbf{v}_e)$ Experimental counterpart of $(\mathbf{u}_m, \mathbf{v}_m)$

$(\mathbf{u}_m, \mathbf{v}_m)$ Modeled streamwise velocity deficit and normalized vertical velocities at probes

$(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ Parameters in the $k - \epsilon$ RANS model requiring calibration

\mathbf{C}_a Analytical estimates of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$

\mathbf{C}_{nom} Nominal value of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$

\mathcal{R} Physically realistic part of the parameter space

CVP Counter-rotating vortex pair

JIC Jet-in-crossflow

I. Introduction

Jet-in-crossflow interactions occur in a myriad of natural and engineering settings.¹ In aerodynamics, spin jets are often used to maneuver launch vehicles, setting up a compressible jet-in-crossflow interaction. Experimental investigations² have shown that the exhaust from the spin rockets (the “jet”) rolls into a complex configuration of counter-rotating vortex pairs (CVPs) that significantly modify the pressure distribution over the stabilizing fins of the flight vehicle. Detailed experimental investigations of compressible

*Technical Staff, Quantitative Modeling and Analysis, MS 9159

†Technical Staff, Quantitative Modeling and Analysis, MS 9152, Senior Member

‡Technical Staff, Aerosciences Department, MS 0825, Senior Member

JIC interactions, for a range of Mach numbers and jet strengths, have been performed experimentally,^{3–5} referred to here as the “Beresh” experiments. They have been compared to simulations performed using Reynolds-Averaged Navier Stokes (RANS) models.⁶ These comparisons show that RANS simulations have poor predictive skill - the simulated jet rolls into a CVP that is far too strong and penetrates into the freestream far more than what is observed in experimental data. These predictive errors arise from two sources. The first concerns model-form errors i.e., due to approximations of turbulence physics in RANS equations. The second source of errors is the use of parameters in RANS equations simply picked from literature, which are often derived by calibrating to canonical incompressible flows. It is unclear which is the larger source of predictive uncertainty.

One way of improving the predictive skill of RANS in JIC simulations is to obtain calibrated parameters. Due to the model-form errors and limited experimental data, it is not always possible to estimate these parameters with much certainty. Consequently, Bayesian calibration techniques are used, as they model the parameters to be estimated as random variables and infer their probability density functions (PDFs) from observational data. The PDFs capture the uncertainty in the estimation. Bayesian estimation of $k - \epsilon$ RANS models has been performed for flow over flat plates under favorable and adverse pressure gradients⁷ and for turbulent flow in urban canyons.⁸ Alternatively, one may assume that the poor predictive skill of RANS equations may be ameliorated by augmenting some of the model equations with a spatially variable source term which is learned by calibrating (i.e., full-field inversions for the source term) to high-fidelity data e.g., Direct Numerical Simulation solutions of flows similar to the flow of interest. The dependence of the source term on local flow properties is then learned using the calibrated RANS solutions as the training data. This approach has been used for 2D flows over bumps in channels and airfoils.^{9–12}

A completely different way of improving the predictive skill of RANS JIC simulations is to develop expressions/estimates of critical parameters via an *analytical* approach. In a companion paper,¹³ we derive a self-similar solution of the jet (after a JIC interaction) that is applicable in the farfield and obtain estimates of 3 $k - \epsilon$ parameters $\mathbf{C} = (C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ analytically, without using Beresh measurements and without performing any type of fitting. Thus these measurements are free of any conflation of model-form and parametric inadequacies. The parameter estimates $\mathbf{C}_a = \{0.1, 2.0, 1.34\}$ are somewhat different from the nominal values of the parameters $\mathbf{C}_{nom} = \{0.09, 1.92, 1.44\}$ typically used in $k - \epsilon$ RANS simulations, but produced predictions of the flowfield that were significantly better than those obtained using \mathbf{C}_{nom} . This determination was performed using the Beresh experimental data, but was limited to the interaction of a Mach 0.8 crossflow interacting with a jet with a jet-to-crossflow momentum ratio of 10.2.

In a previous paper¹⁴ we addressed the Bayesian calibration of compressible JIC interactions using data from one of the Beresh experiments. The experiment consists of introducing a Mach number $M = 3.73$ jet of 9.53 mm diameter into a $M = 0.8$ freestream in a wind tunnel test section. The jet-to-freestream momentum ratio J is 10.2. The jet bends into the flow and rolls into a CVP as it flows downstream. PIV (velocity) measurements are available on the mid-plane, the longitudinal plane of symmetry and a transverse/cross plane that slices through the CVP. Experimentally measured vorticity on the crossplane was used, via Bayesian calibration, to obtain a 3D joint PDF (JPDF) for $\mathbf{C} = (C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$, the parameters to which the flow was found to be most sensitive. An optimal parameter set $\mathbf{C}_{opt} = \{0.1025, 2.099, 1.416\}$ was also obtained, which is quite different from \mathbf{C}_{nom} . Samples of \mathbf{C} drawn from the calibrated JPDF provided predictions of the flowfield that were far more accurate than those obtained using \mathbf{C}_{nom} ; this was verified by comparing with (Beresh) experimental velocity measurements on the mid-plane and crossplane. The improvement in predictive skill was retained, when examined against experimental data, at other M and J values, though the agreement was not as good as the $M = 0.8, J = 10.2$ case.

Clearly, then, \mathbf{C}_a , \mathbf{C}_{opt} and the calibrated JPDF from the ($M = 0.8, J = 10.2$) calibration study are all more predictive than \mathbf{C}_{nom} . However the predictive skill of \mathbf{C}_{opt} and the calibrated JPDF from the ($M = 0.8, J = 10.2$) calibration study is seen to drop as we apply them to (M, J) combinations other than (0.8, 10.2). The ultimate aim of this study is to find a replacement for \mathbf{C}_{nom} in $k - \epsilon$ simulations of compressible JIC interactions that is robust across a $M - J$ range. In this study we perform calibration for \mathbf{C} for four (M, J) combinations for which we have experimental measurements in Refs. 3–5; they are (M, J) =, (0.6, 10.2), (0.7, 10.2), (0.8, 10.2) and (0.8, 16.7). These calibrations will result in four JPDFs which are specific to the experiments they have been calibrated to. If the JPDFs are similar, it may be possible to smooth them into a single JPDF that would hold across the (M, J) covered by the Beresh experiments. If they are very dissimilar, it may be practical to find a single point estimate that would be predictive across the same (M, J), and we will explore whether \mathbf{C}_a can serve in this role.

II. Problem formulation

II.A. Experimental and computational setup

The experimental and computational setup used in this calibration study have been described fully elsewhere¹⁴ and we provide a summary below. The data used here is obtained from a set of wind tunnel experiments conducted by Beresh et al.^{3,4}. A schematic of the test section is shown in Fig. 1. The experiments involve injecting a $M = 3.73$ jet, of diameter 9.53 mm, into a freestream from the bottom of the test section. The freestream Mach number is denoted M and the ratio of jet-to-freestream momenta is J . Experiments are conducted for $M = 0.6, 0.7$ and 0.8 while $J = 10.2$. For $M = 0.8$, experiments are conducted at $J = 16.7$ and 5.6 , in addition to 10.2 . The jet bends into the freestream and rolls into a CVP. PIV measurements are made on two planes, as shown in Fig. 1. The plane of symmetry is called the mid-plane, whereas the transverse plane slicing through the CVP is called the crossplane. Velocity measurements are available on the mid-plane for all the five experimental datasets being considered in this study. Velocity measurements (and therefore streamwise vorticity) are available only for the three $M = 0.8$ test cases (i.e., for $J = 5.6, 10.2, 16.7$) and was used as the calibration observable in our previous work.¹⁴ Since we will perform calibration for four (M, J) combinations (not all of which have crossplane measurements), we will use measurements of streamwise velocity u and vertical velocity v on the mid-plane as our calibration observable. Measurements on the mid-plane are made at 5 stations, with the first station being 200 mm downstream of the jet and the rest at intervals of 50 mm. At each location, measurements are made at 63 vertically distributed points called “probes”. The crossplane is 321.8 mm downstream of the jet.




Figure 1. Schematic of the test section showing the orifice where the jet is introduced and the mid- and crossplane where experimental measurements are made.

The calibration is performed using the compressible form of the $k - \epsilon$ RANS equation¹⁵ and is fully described in Ref. 14. The most important parameters that affect the flowfield are $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$, which we will refer to as **C**. Our flow solver, SIGMA CFD (Sandia Implicit Generalized Mult-Block Analysis Code for Fluid Dynamics), uses a Roe-TVD flux scheme with a minmod limiter for spatial discretization. Time integration is carried out using a first order point-implicit scheme. The calculations are initialized using a first-order spatial scheme, time-marched for 5000 timesteps, and relaxed to convergence using the second order scheme for 25000 timesteps. Time marching to steady state was carried out using local time stepping with a gradual CFL ramp to accelerate convergence. A multi-block mesh with approximately 10 million grid cells is used; mesh convergence studies are in Ref. 6. Details of initial and boundary conditions and mesh refinement near the walls are in Ref. 14.

II.B. The inverse problem

Let \mathbf{y}_e be a vector (of length N_p) of experimental observations, measured at a set of N_p locations (“probes”). Let $\mathbf{y}_m(\mathbf{C})$ be model predictions of the same, produced by a parameter setting **C**. They are related by $\mathbf{y}_e = \mathbf{y}_m(\mathbf{C}) + \boldsymbol{\epsilon}$ where $\boldsymbol{\epsilon}$ is a combination of measurement and model-form error. We make a modeling assumption that the errors at the probes are uncorrelated, independently and identically distributed as a

zero-mean Gaussian i.e. $\epsilon = \{\epsilon_i\}, \epsilon_i \sim \mathcal{N}(0, \sigma^2)$. σ^2 thus provides a crude measure of the model - data misfit after calibration. If the measurement errors are low, σ^2 provides an estimate of the model-form error.

Let $P(\mathbf{C}, \sigma^2 | \mathbf{y}_e)$ be the joint probability density function of the parameters and the model - data misfit, conditional on the observed data \mathbf{y}_e . Let $\Pi_1(\mathbf{C})$ and $\Pi_2(\sigma^2)$ be our prior belief regarding the distribution of \mathbf{C} and σ^2 . The likelihood of observing \mathbf{y}_e , given a parameter setting \mathbf{C} , $\mathcal{L}(\mathbf{y}_e | \mathbf{C})$, is given by

$$\mathcal{L}(\mathbf{y}_e | \mathbf{C}, \sigma^2) \propto \frac{1}{\sigma^{N_p}} \exp\left(-\frac{\|\mathbf{y}_e - \mathbf{y}_m(\mathbf{C})\|_2^2}{2\sigma^2}\right).$$

By Bayes' theorem, the calibrated distribution (or posterior distribution) of (\mathbf{C}, σ^2) can be given as

$$P(\mathbf{C}, \sigma^2 | \mathbf{y}_e) \propto \mathcal{L}(\mathbf{y}_e | \mathbf{C}, \sigma^2) \Pi_1(\mathbf{C}) \Pi_2(\sigma^2) \propto \frac{1}{\sigma^{N_p}} \exp\left(-\frac{\|\mathbf{y}_e - \mathbf{y}_m(\mathbf{C})\|_2^2}{2\sigma^2}\right) \Pi_1(\mathbf{C}) \Pi_2(\sigma^2) \quad (1)$$

The actual model outputs i.e., components of \mathbf{y}_m , are the streamwise velocity deficit u_{def} and normalized velocity v_{norm} defined as

$$u_{def} = \frac{U_{max} - u}{U_\infty(x)} \quad v_{norm} = \frac{v}{U_\infty},$$

where U_{max} is the maximum streamwise velocity at an x-location and U_∞ is the freestream velocity. Let \mathbf{u}_m be the predicted *streamwise velocity deficit* u_{def} and \mathbf{v}_m be the *normalized vertical velocity* v_{norm} for all the probes. Let \mathbf{u}_e and \mathbf{v}_e be their experimental counterparts over all the probes. The specific form of Eq. 1 used in this study is:

$$P(\mathbf{C}, \sigma^2 | \mathbf{u}_e, \mathbf{v}_e) \propto \frac{1}{\sigma^{N_p}} \exp\left(-\frac{\|(\mathbf{u}_e - \mathbf{u}_m(\mathbf{C}))/K_u\|_2^2}{2\sigma^2}\right) \exp\left(-\frac{\|(\mathbf{v}_e - \mathbf{v}_m(\mathbf{C}))/K_v\|_2^2}{2\sigma^2}\right) \Pi_1(\mathbf{C}) \Pi_2(\sigma^2), \quad (2)$$

where $K_u = \max(\mathbf{u}_e)$ and $K_v = \max(\mathbf{v}_e)$. Normalization of the model-data mismatch using K_u and K_v ensures that the contributions from u_{def} and v_{norm} are equally weighted in the expression for the posterior distribution.

As in Ref. 14, the inverse problem is solved via sampling. We use a Markov chain Monte Carlo (MCMC) method called Delayed Rejection Adaptive Metropolis (DRAM; Ref. 16), as implemented in the R¹⁷ package FME¹⁸ to draw samples of $\{\mathbf{C}, \sigma^2\}$. We reconstruct $P(\mathbf{C}, \sigma^2 | \mathbf{y}_e)$ empirically by kernel density estimation.¹⁹ The prior density for σ^2 is defined in terms of its reciprocal i.e., $\Pi_2(\sigma^{-2})$ and is modeled using a Gamma prior i.e., $\sigma^{-2} \sim \Gamma(k, \theta)$, where $k = 1, \theta = 1$. The inverse Gamma prior for σ^2 is a conjugate prior which simplifies sampling of σ^{-2} via a Gibbs sampler. Also, the prior is virtually non-informative for $\sigma^2 > 5$. Raftery-Lewis method²⁰ implemented in the R package mcgibbsit.²¹ The MCMC method requires $O(10^4)$ samples to construct $P(\mathbf{C}, \sigma^2 | \mathbf{y}_e)$, each of which requires a 3D RANS model evaluation to provide $\mathbf{y}_m(\mathbf{C})$. Since this is impractical, we will develop a surrogate model, a polynomial that maps the dependence of our calibration variables $(\mathbf{u}_m, \mathbf{v}_m)$ on \mathbf{C} . The surrogate model will serve as a computationally inexpensive proxy for SIGMA CFD. We consider the following bounds on \mathbf{C} taken from Ref. 15:

$$0.06 \leq C_\mu \leq 0.12, \quad 1.7 \leq C_{\epsilon 2} \leq 2.1, \quad \text{and} \quad 1.2 \leq C_{\epsilon 1} \leq 1.7. \quad (3)$$

II.C. Priors and surrogates

An informative prior: While it is tempting to combine the bounds in Eq. 3 into a uniform distribution in the cuboid \mathcal{C} in $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ space, we shall refrain from doing. Arbitrary combinations of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ from \mathcal{C} may not be physical, and consequently, we will use the discrepancy $\mathbf{y}_e - \mathbf{y}_m(\mathbf{C})$ to choose a physically realistic region $\mathcal{R} \subset \mathcal{C}$ to serve as $\Pi_1(\mathbf{C})$. Details on how $\Pi_1(\mathbf{C})$ is constructed are in Ref. 14 and we provide a summary here.

Constructing the surrogates requires us to generate a training dataset that captures the variation of \mathbf{u}_m and \mathbf{v}_m over \mathcal{C} . We draw 2744 ($= 14^3$) samples of \mathbf{C} from \mathcal{C} using a quasi Monte Carlo space-filling (Halton sequence) method and conduct 3D RANS simulations with them using SIGMA CFD. Many of the samples are non-physical and the simulations do not converge to a steady state. Others (2628 samples) do, but yield flowfields that are quite unlike transonic, high Reynolds number flows. In order to isolate values of \mathbf{C} that yield realistic flowfields we compute the RMS (root mean square) error between $(\mathbf{u}_e, \mathbf{v}_e)$ and $(\mathbf{u}_m, \mathbf{v}_m)$ for all the successful simulations and preserve the top 25%, resulting in $0.25 \times 2628 = 657$ samples. A random subset

of these \mathbf{C} samples are plotted in Fig. 2 (left) and occupy a certain section of \mathcal{C} i.e., they are not uniformly scattered in \mathcal{C} . This is the physically realistic part of \mathcal{C} , called \mathcal{R} , and forms the basis of constructing the prior density $\Pi_1(\mathbf{C})$. As in Ref. 14, we define

$$\Pi_1(\mathbf{C}) = \begin{cases} 1 & \text{if } \mathbf{C} \in \mathcal{R} \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

The 2628 successful runs (of which 657 mark out \mathcal{R}) are used to train a binary support vector machine classifier. The procedure and software to do so are in Ref. 14. A misclassification rate less than 10% is a requirement for building a successful classifier. The classifier is used as the implementation of the prior $\Pi_1(\mathbf{C})$ as described in Eq. 4 and is used within the MCMC procedure to infer the posterior density $P(\mathbf{C}, \sigma^2 | \mathbf{u}_e, \mathbf{v}_e)$.




Figure 2. Left: A set of \mathbf{C} points in \mathcal{C} that produce physically realistic flowfields i.e., they approximately mark out \mathcal{R} . Right: A plot of the probes on the mid-plane where experimental data are available. The filled circles are the probes for which we could build accurate surrogate models. They track the spatial evolution of the CVP, and do not cover the boundary layer at the test section walls..

Surrogate models: A surrogate model is defined as a polynomial curve-fit between a model prediction at a probe i.e., u_{def} and v_{norm} , that captures their dependence on \mathbf{C} . We use the samples of \mathbf{C} that define \mathcal{R} , and the flowfields they yield, to learn the polynomial model. The procedure is described in detail Ref. 14. We start with a cubic polynomial in (C_μ, C_{e2}, C_{e1}) space to map \mathbf{C} to u_{def} (or v_{norm}). The fitting is performed via least-squares regression and simplified via incrementally dropping terms and computing the Akaike Information Criterion. Many cubic terms are removed and, for a few probes, the surrogate reduces to a quadratic one. However, simply being able to construct a surrogate does not imply that it is accurate. As in Ref. 14, we compute the predictive error of the surrogate via repeated random sub-sampling validation (a form of cross validation) and retain only those probes whose surrogate models have less than 15% error (for both u_{def} and v_{norm}). In Fig. 2 (right) we plot the probes on the mid-plane that were retained for the ($M = 0.8, J = 10.2$) case. As the figure shows, of the $63 \times 5 = 315$ probes for which we have experimental measurements, we could make surrogate models (of acceptable accuracy) for only 107. These “model-able” probes follow the trajectory of the CVP. Note that we do not keep any probes in the boundary layer - the calibration is simply aimed at simulating the jet and its roll-up accurately.

II.D. Bayesian calibration

Having implemented the prior $\Pi_1(\mathbf{C})$ via a classifier and constructed polynomial surrogates for \mathbf{u}_m and \mathbf{v}_m , we solve the inverse problem (Eq. 2) for the posterior density $P(\mathbf{C}, \sigma^2 | \mathbf{u}_e, \mathbf{v}_e)$. We will use $P_{(\mathbf{u}_e, \mathbf{v}_e)}$ as an abbreviation for $P(\mathbf{C}, \sigma^2 | \mathbf{u}_e, \mathbf{v}_e)$ and P_{ω_e} for $P(\mathbf{C}, \sigma^2 | \omega_e)$ which we estimated in our previous work.¹⁴ Here, ω_e is the experimentally observed vorticity on the crossplane due to the CVP. We plot the marginalized




Figure 3. Marginalized posterior PDFs of C_μ , $C_{\epsilon 2}$ and $C_{\epsilon 1}$ as inferred using $(\mathbf{u}_e, \mathbf{v}_e)$ (marked “UV”) and using streamwise vorticity (marked “Omega”). The corresponding priors are plotted using symbols. The vertical line is the nominal value.

posterior and prior PDFs in Fig. 3. For comparison, we also plot the marginalized PDFs obtained using measurements of vorticity on the crossplane¹⁴ for the same interaction. We see that the two posterior PDFs are close but not identical. Further, the priors that we use here (plotted using \circ) are different from those used in our previous study,¹⁴ plotted with \triangle . This is especially true for $C_{\epsilon 1}$. The comparison ensures that (1) $(\mathbf{u}_e, \mathbf{v}_e)$ are as informative about $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ as vorticity on the crossplane and (2) the classifier and surrogates were constructed with a degree of accuracy that is comparable to our previous work.¹⁴ Therefore, the calibration methodology using $(\mathbf{u}_e, \mathbf{v}_e)$, as described above, can be applied to other JIC interactions i.e., those with different (M, J) combinations.

Next we check the effect of the difference in PDFs, plotted in Fig. 3, on the flowfield, primarily to check its sensitivity to $\{(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})\}$. We sample 100 realizations of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ from $P_{(\mathbf{u}_e, \mathbf{v}_e)}$ and P_{ω_e} separately, perform RANS simulations and compare their results. In Fig. 4 (top) we plot the streamwise velocity deficit u_{def} and normalized vertical velocity v_{norm} at streamwise locations $x/D_j = 21, 31.5$ and 42 , where $D_j = 9.53$ mm is the diameter of the jet. Results due to \mathbf{C}_{nom} , the nominal value of the parameters, and \mathbf{C}_a , parameters estimated analytically (more on this in Sec. III) in Ref. 13 are also plotted. We find that the ensemble mean predictions using $P_{(\mathbf{u}_e, \mathbf{v}_e)}$ and P_{ω_e} are practically indistinguishable, which is remarkable given that the two PDFs were calibrated to different observations (mid-plane velocity and crossplane vorticity respectively). We also plot the profiles due to the analytically predicted values \mathbf{C}_a from Ref. 13, which also agree with experiments. In Fig. 5, we plot the corresponding vorticity fields due the right vortex of the CVP on the crossplane. The experimental vorticity field is plotted with white contours. Again, in Fig. 5 (top left) we see that \mathbf{C}_{nom} predicts the experimental data badly, whereas as \mathbf{C}_a (Fig. 5, top right) results in a marked improvement. In Fig. 5 (bottom left and right) we plot the predictions using $\mathbf{C}_{opt,1}$ and $\mathbf{C}_{opt,2}$, where $\mathbf{C}_{opt,1} = \{0.1025, 2.099, 1.416\}$ is the $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ set from the 100 drawn from




Figure 4. Profiles of u_{def} (top) and v_{norm} (bottom) at three stations $x/D_j = 21, 31.5$ and 42 for a ($M = 0.8, J = 10.2$) interaction. The symbols \circ are the experimental measurements, the solid blue and red lines are the ensemble means computed using 100 C samples from P_{ω_e} and $P_{(u_e, v_e)}$ respectively. The dashed red line is the prediction using C_{nom} and the dashed black line is computed using C_a from Ref. 13.




Figure 5. Comparison of the crossplane vorticity field (due to the right vortex of the CVP) as reproduced using various parameters for a ($M = 0.8, J = 10.2$) interaction. The contour are the vorticity field from experiments.⁴ Top left: The flood plot is the vorticity field predicted by C_{nom} . Top right: Vorticity field produced by C_a .¹³ Bottom left: Vorticity field computed using $C_{opt,1} = \{0.1025, 2.099, 1.416\}$, obtained from P_{ω_e} .¹⁴ Bottom right: Vorticity field from $C_{opt,2} = \{0.11, 1.86, 1.66\}$, obtained from $P_{(u_e, v_e)}$.

P_{ω_e} that achieves the best reproduction of crossplane experimental vorticity (see description in Ref. 14) and $\mathbf{C}_{opt,2} = \{0.11, 1.86, 1.66\}$ is its counterpart drawn from $P_{(\mathbf{u}_e, \mathbf{v}_e)}$ that reproduces the mid-plane streamwise velocity deficit and normalized vertical velocity best. We see that the predictions using $\mathbf{C}_{opt,1}$ and $\mathbf{C}_{opt,2}$ are indistinguishable from each other and agree quite well with the experimental vorticity. They are also much improved compared to the \mathbf{C}_{nom} simulations.

Thus it is clear that predictions using $P_{(\mathbf{u}_e, \mathbf{v}_e)}$ improves the agreement of RANSI simulations with experiments. While that is not surprising when considering u_{def} and v_{norm} (since they were used in the calibration), the crossplane vorticity predictions are remarkably good. Further, the mid-plane and crossplane predictions also match predictions using P_{ω_e} , indicating the calibration using mid-plane velocity measurements is similar in quality to those obtained using vorticity measurements in Ref. 14. Thus we can proceed with calibration using mid-plane quantities, which are available for the four cases being considered in this study.

II.E. Combining JPDFs

The $J = 4$ separate JPDFs obtained using Bayesian calibration will yield distributions $g_j(\mathbf{C}, \sigma^2), j = 1 \dots J$. If they are similar, the could be combined into a “smoothed” distribution $G(\mathbf{C}, \sigma^2)$ that would be predictive across the (M, J) space spanned by the Beresh experiments. This could be a practical way of summarizing the results of Bayesian calibration.. We will use Bayesian model averaging (BMA) to do so, closely following the arguments in Ref. 22.

We sample $N = 100$ realizations of \mathbf{C} from $g_j(\mathbf{C}, \sigma^2)$, leading to $\mathbf{C}_l, l = 1 \dots L, L = NJ$. We use the surrogate models used in Bayesian calibration to predict u_{def} and v_{norm} at M probes for which we have surrogate models for all $K = 4$ experiments. Let the predictions generated using \mathbf{C}_l be \mathbf{z}_l , a vector $2MK$ long; the factor of 2 indicates the two predictions u_{def} and v_{norm} that we obtain at each probe.

We view the (collection of) surrogate models, each seeded with a different \mathbf{C}_l , as an ensemble of L models. Let \mathbf{z}_{obs} be the observations of u_{def} and v_{norm} from the K experiments at the same probes. We establish that

$$\mathbf{z}_{obs} \sim \mathcal{N}(\mathbf{z}_l, \Sigma^2)$$

where Σ is unknown and $\mathcal{N}(\cdot)$ is a Gaussian distribution. This equation can be used to calculate the “importance” of model l (equivalently, \mathbf{C}_l), given Σ ; if \mathbf{z}_{obs} and \mathbf{z}_l are very different, the likelihood of sampling \mathbf{z}_{obs} will be low. We formulate the smoothed version as

$$\begin{aligned} G(\mathbf{C}, \sigma^2) &= \sum_l w_l q_l(\mathbf{z}_l), \\ \mathbb{E}(\mathbf{z}_{obs} | \mathbf{z}_l) &= \sum_l w_l \mathbf{z}_l, \quad \sum_l w_l = 1 \end{aligned} \quad (5)$$

Here q_l serves as an abbreviation for $\mathcal{N}(\mathbf{z}_l, \Sigma^2)$; w_l are the weights of each model/ \mathbf{C}_l . Eq. 5 can be solved using Expectation-Maximization for (w_l, Σ) as done in Ref. 22. The weighted collection of \mathbf{C}_l can be used to construct a JPDF for \mathbf{C} that retains the region of the \mathbf{C} -space that is predictive across a (M, J) range. However, this assumes that highly weighted \mathbf{C}_l will cluster. If they do not, BMA will results in a complex, multimodal distribution that would not be any simpler than J separate JPDFs. In Sec. III we will investigate if $G(\mathbf{C}, \sigma^2)$ can be computed or if \mathbf{C}_a is the only way of summarizing the JPDFs arising from Bayesian calibration.

III. Results

III.A. Evaluating \mathbf{C}_a

We perform Bayesian calibration using mid-plane velocity measurements (strictly, streamwise flow deficit and normalized vertical velocity, as captured in Eq. 2) from the Beresh experiments.³⁻⁵ These are performed for $(M, J) = (0.6, 10.2), (0.7, 10.2), (0.8, 10.2)$ and $(0.8, 16.7)$. In Fig. 6 we plot the marginalized PDFs (i.e., marginalized versions of $P(\mathbf{C}, \sigma^2 | \mathbf{u}_e, \mathbf{v}_e)$) obtained from the four Bayesian calibrations. The nominal values \mathbf{C}_{nom} are plotted using dashed vertical lines and the analytical estimate \mathbf{C}_a using a solid vertical line. It is clear that $C_{\epsilon 2}$ is much higher than the nominal value and one could argue that the upper bound for $C_{\epsilon 2}$ should be increased to allow a better calibration. The nominal estimate for $C_{\epsilon 1}$ agrees with the peak for many of the calibration cases except for the $(0.6, 10.2)$ case. The PDFs for C_{μ} do not show any consistent




Figure 6. Marginalized posterior PDFs inferred for JIC interactions corresponding to five different (M, J) combinations. The dashed vertical line is C_{nom} whereas the solid line is the analytical estimate C_a .

trend, though we do see 3 peaks in the region $C_\mu > 0.11$. Post-calibration, the residual model-data mismatch (quantified by σ) is also plotted in Fig. 6; it is clear that the $(M = 0.6, J = 10.2)$ has the worst fit whereas the $(M = 0.8, J = 10.2)$ and $(M = 0.7, J = 10.2)$ have the best fits (lower σ).

Fig. 6 shows that the analytical estimate for $C_{\epsilon 2}$ and C_μ are closer to the peaks of the PDFs than the nominal one, whereas the opposite holds true for $C_{\epsilon 1}$. The three parameters exert different degrees of influence on the flowfield and it is unclear whether the predictive skill of \mathbf{C}_a may be similar to the calibrated JPDF. Consequently we discuss the flowfields that the calibrated $P_{(\mathbf{u}_e, \mathbf{v}_e)}$ produce vis-à-vis experiments, \mathbf{C}_a and \mathbf{C}_{nom} . Note that we have already shown in Ref. 13 that \mathbf{C}_a is more predictive than \mathbf{C}_{nom} .

Comparisons for the $(M = 0.8, J = 10.2)$ case are plotted in Fig. 4 and Fig. 5. The accuracy of the $P_{(\mathbf{u}_e, \mathbf{v}_e)}$ distribution is assessed using the ensemble mean of predictions obtained using 100 \mathbf{C} samples drawn from it, as described in Sec. II.D. In Fig. 4 we see that the velocity profiles using ensemble means from $P_{(\mathbf{u}_e, \mathbf{v}_e)}$ match well with experimental data; they are also indistinguishable from the predictions using \mathbf{C}_a . This holds true for both u_{def} and v_{norm} . Further, they are an immense improvement over \mathbf{C}_{nom} predictions. We next check the ability of $P_{(\mathbf{u}_e, \mathbf{v}_e)}$ and \mathbf{C}_a to reproduce the crossplane vorticity field. We obtain $\mathbf{C}_{opt,2} = \{0.11, 1.86, 1.67\}$ from $P_{(\mathbf{u}_e, \mathbf{v}_e)}$, and in Fig. 5 (bottom right), we see plot its vorticity field vis-à-vis experimental data (plotted as white contours). The agreement with the experiment is very good, and certainly much better than the one obtained using \mathbf{C}_{nom} (Fig. 5, top left). In Fig. 5 (top right) we also plot the vorticity field produced by \mathbf{C}_a . Comparing with experimental data, we see that the agreement is good. Further, the predictions using \mathbf{C}_a and $\mathbf{C}_{opt,2}$ are very similar. Clearly, then, \mathbf{C}_a can serve as a proxy for $P_{(\mathbf{u}_e, \mathbf{v}_e)}$ for the $(M = 0.8, J = 10.2)$ interaction.

In Fig. 7 and 8, we perform the same comparison for a $(M = 0.7, J = 10.2)$ interaction. $\mathbf{C}_{opt,2} = \{0.12, 2.1, 1.45\}$ is obtained in a manner identical to the one described in Sec. II.D. In Fig. 7, we plot the

streamwise flow deficit u_{def} (on top) and the normalized vertical velocity v_{norm} below. The predictions using \mathbf{C}_a are indistinguishable from the ensemble mean computed using 100 samples drawn from $P_{(\mathbf{u}_e, \mathbf{v}_e)}$. They are both closer to the experimental data than the predictions using \mathbf{C}_{nom} . In Fig. 8 (left), we plot the vorticity field using $\mathbf{C}_{opt,2}$ with predictions using \mathbf{C}_a overlaid using white contours. The agreement is very close. In Fig. 8 (right), we compute the circulation, centroid and radius of gyration of the vorticity field predicted by the samples drawn from $P_{(\mathbf{u}_e, \mathbf{v}_e)}$, normalized by their counterparts computed using \mathbf{C}_a . The ratios show less than 5% deviation, showing that (1) the joint PDF of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ plotted in Fig. 6 do not result in much variation in the vorticity field and (2) \mathbf{C}_a can reproduce an approximation to the ensemble predictions quite accurately. This general statement also holds true for the investigation into the ($M = 0.6, J = 10.2$) interaction. This does not automatically imply that the predicted vorticity field is accurate; unfortunately, experimental measurements of crossplane velocities do not exist for this jet-in-crossflow interaction to allow verification of $P_{(\mathbf{u}_e, \mathbf{v}_e)}$.

While not shown in this paper, the ($M = 0.6, J = 10.2$) and ($M = 0.8, J = 16.7$) interactions show the same agreement between mid-plane experimental measurements, \mathbf{C}_a predictions and ensemble means generated using \mathbf{C} samples from $P_{(\mathbf{u}_e, \mathbf{v}_e)}$.




Figure 7. Profiles of u_{def} (top) and v_{norm} (bottom) at three stations $x/D_j = 21, 31.5$ and 42 for a ($M = 0.7, J = 10.2$) interaction. The symbols \circ are the experimental measurements, the solid red lines are the ensemble means computed using 100 \mathbf{C} samples from $P_{(\mathbf{u}_e, \mathbf{v}_e)}$. The dashed red line is the prediction using \mathbf{C}_{nom} and the dashed black line is computed using \mathbf{C}_a from Ref. 13.

III.B. Constructing $G(\mathbf{C}, \sigma^2)$

We implemented the BMA method for combining \mathbf{C}_l samples using u_{def} and v_{norm} observations from $K = 4$ experiments. $M = 82$ probes for which we could make surrogate models; there are 315 probes in all. Therefore we seek to estimate $NJ = 100 \times 4 = 400$ w_l using $2 \times M \times K = 656$ observations. In Fig. 9 (left) we plot the distribution of weights; most are very small. In fact, twenty-two models/ \mathbf{C}_l combinations account for 99.9% of the weights. In Fig. 9 (right) we plot the 400 \mathbf{C}_l samples; they are clearly disjoint, and unlikely to lead to a simple distribution. The twenty-two \mathbf{C}_l samples that account for most of the weight are also plotted; they do not cluster. Thus it is unlikely that there exists a distribution in \mathbf{C} -space that could




Figure 8. Left: Comparison of the crossplane vorticity field (right vortex of the CVP) as reproduced using $\mathbf{C}_{opt,2} = \{0.12, 2.1, 1.45\}$ computed using $P_{(u_e, v_e)}$ from a $(M = 0.7, J = 10.2)$ interaction against predictions using \mathbf{C}_a , which is overlaid using white contours. Right: Circulation of the vortex, the horizontal and vertical positions of the centroid of the vorticity distribution and radius of gyration, computed using 100 samples from $P_{(u_e, v_e)}$, normalized by the counterparts computed using \mathbf{C}_a .

summarize the collection of JPDFs obtained (see Fig. 6) and \mathbf{C}_a may be the only useful summary.

III.C. Discussion

The results presented here are best viewed within the context provided by our previous publications,^{14,23} and our companion paper Ref. 13. These studies seek to uncover the cause behind RANS’s inability to simulate compressible JIC interactions with much precision. There can be two possible causes; (1) the use of inappropriate values of RANS constants ($C_\mu, C_{\epsilon 2}, C_{\epsilon 1}$) and (2) approximations (“missing physics”) in RANS which can be removed only by augmenting the RANS equations. In our studies, we consider both the possibilities.

In Ref. 14, we developed a Bayesian calibration method to estimate $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ from experimental measurements on the crossplane of a JIC interaction. The calibration improved predictions immensely (vis-à-vis predictions using \mathbf{C}_{nom} , the nominal/literature values of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$), though it was limited to the mean flow; turbulent stresses were not substantially improved. Further, the calibration was performed for a Mach 0.8 crossflow interacting with a jet with a jet-to-crossflow momentum ratio of 10.2. The optimal parameters, referred to in this paper as $\mathbf{C}_{opt,1}$, were quite different from \mathbf{C}_{nom} . Some results from this calibration are in Sec. II. However, Bayesian calibration merely fits a model to data, and there is no guarantee that the JPDF (or optimal parameter values) so inferred have any physical basis. In addition, Bayesian calibration could easily conflate model-form errors (RANS approximations) with parametric uncertainty i.e., the newly inferred parameters could simply be compensating for RANS’ “missing physics”.

We address this question in two ways. In our companion paper,¹³ we developed an analytical model for JIC interactions, which have yielded \mathbf{C}_a , the analytical counterpart of $\mathbf{C}_{opt,1}$. Since \mathbf{C}_a does not involve fitting to data, there is no question of conflating model-form and parametric uncertainties. \mathbf{C}_a is close to $\mathbf{C}_{opt,1}$, as estimated for the $(M = 0.8, J = 10.2)$ interaction, indicating that Bayesian calibration yields physically defensible results and is not merely a “curve fit”.

The second approach used to examine the question is developed in this paper. We calibrated RANS to four JIC experiments and showed that predictions of the mean flow using the calibrated JPDF and \mathbf{C}_a agree; furthermore, both are close to experimental measurements. In addition, both result in a tremendous improvement in predictions over \mathbf{C}_{nom} . This holds for all for experiments that span a range of Mach numbers ($M = \{0.6, 0.7, 0.8\}$) and $J = \{10.2, 16.7\}$). These imply that Bayesian calibration yield physically based




Figure 9. Left: Histogram of $\log_{10}(w_i)$, showing that most of the weights are small. Right: Distribution of C_1 used in BMA. The red C points are chosen from the JPDF computed from $(M = 0.6, J = 10.2)$, the blue ones from $(M = 0.7, J = 10.2)$, the green ones from $(M = 0.8, J = 10.2)$ and magenta symbols from the JPDF from $(M = 0.8, J = 10.2)$ interaction. The large black circles are the C_1 combinations that were selected by BMA. There is no clustering.

JPDFs and the lack of RANS' predictive skill for JIC was likely due to the use of inappropriate values of $(C_\mu, C_{\epsilon 2}, C_1)$.

As a check whether the model-form errors could be a *leading* cause of RANS prediction errors, we explored augmenting the linear eddy viscosity model (EVM) used in SIGMA CFD with quadratic and cubic terms. That investigation, that involved both the enrichment of the EVM and Bayesian calibration,²³ and tested for a $(M = 0.8, J = 10.2)$ interaction, showed that the predictive skill was comparable to our results in Ref. 14, which was purely a product of Bayesian calibration. This also somewhat corroborates our finding in this paper that better RANS constant could improve its JIC predictions. Note that after calibration, any lack of predictive skill is purely due to model-form errors, and they are manifest most clearly in the turbulent stresses.¹⁴ More research is necessary to quantify them and uncover the degree to, and the mechanism by which they impact the mean flow quantities investigated here.

The four JPDFs inferred from the experimental datasets considered here are quite different and cannot be smoothed into a useful summary JPDF via a weighted average i.e., a mixture model. Instead, \mathbf{C}_a serves as an adequate summary. The reason is as follows. $C_{\epsilon 2}$ is the most sensitive parameter, and all four calibrations indicate a value of 2.0 or higher. This is captured by \mathbf{C}_a which has it at 2.0. The PDF for C_1 peaks at a value close to \mathbf{C}_{nom} and \mathbf{C}_a is not very different either. C_μ exhibits very different values in the four calibrations, but the mean flow is not very sensitive to it. There is some indication that C_μ should be higher than its value in \mathbf{C}_{nom} and \mathbf{C}_a accommodates that slight evidence.

IV. Conclusions

In this paper, we explore the usefulness of \mathbf{C}_a , an analytical value of RANS constants $(C_\mu, C_{\epsilon 2}, C_1)$ with respect to predictive jet-in-crossflow simulations. We have done so for crossflow Mach numbers between 0.6 and 0.8, and for jet-to-crossflow momentum ratios of 10.2, and 16.7. We performed Bayesian calibration using data from four separate experiments, so that we could obtain very good RANS predictions, and found that \mathbf{C}_a could reproduce them very closely. Further, both RANS and \mathbf{C}_a predictions matched experiments and were far better than predictions using the \mathbf{C}_{nom} , the nominal values of $(C_\mu, C_{\epsilon 2}, C_1)$. This held true on the mid-plane, as well as the crossplane, for the two experiments for which we have crossplane (experimental) data. \mathbf{C}_a is thus a useful summary of the four joint probability density functions inferred from Bayesian calibration. Model-averaging the four densities into a smoothed summary distribution failed, mostly due to

the fact that the Bayesian calibrations had only one thing in common - the correction of the most important constant $C_{\epsilon 2}$. This correction was also adequately captured by \mathbf{C}_a .

This work, along with our companion paper¹³ and our previous papers (Ref. 14,23) lead us to believe that the main cause of RANS prediction errors for jet-in-crossflow may be due to inappropriate RANS constants. Calibration or the use of \mathbf{C}_a improve mean flow predictions immensely (certainly vis-à-vis \mathbf{C}_{nom} , the nominal values of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$), but does little to improve turbulent stresses.¹⁴ Having removed prediction errors due to inappropriate RANS constant, any residual errors are therefore due to model-form inadequacies i.e., missing physics in RANS. A better treatment of the turbulence models i.e., the equations for k , ϵ , the eddy viscosity model and the closures in the momentum and energy equations, are required to remedy this shortcoming.

Acknowledgment

This work was supported by Sandia National Laboratories' Advanced Scientific Computing (ASC) Verification and Validation program. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U. S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

References

- ¹Mahesh, K., "The interaction of jets with crossflows," *Annual Reviews of Fluid Mechanics*, Vol. 45, 2013, pp. 379–407.
- ²Beresh, S. J., Heineck, J. T., Walker, S. M., Schairer, E. T., and Yaste, D. M., "Planar velocimetry of jet/fin interaction on a full-scale flight vehicle configuration," *AIAA Journal*, Vol. 45, No. 8, 2007, pp. 1827–1840.
- ³Beresh, S. J., Henfling, J. F., Erven, R. J., and Spillers, R. W., "Penetration of a transverse supersonic jet into a subsonic compressible crossflow," *AIAA Journal*, Vol. 43, No. 2, 2005, pp. 379–389.
- ⁴Beresh, S. J., Henfling, J. F., Erven, R. J., and Spillers, R. W., "Crossplane velocimetry of a transverse supersonic jet in a transonic crossflow," *AIAA Journal*, Vol. 44, No. 12, 2006, pp. 3051–3061.
- ⁵Beresh, S. J., Henfling, J. F., Erven, R. J., and Spillers, R. W., "Turbulent characteristics of a transverse supersonic jet in a subsonic compressible crossflow," *AIAA Journal*, Vol. 43, No. 11, 2005, pp. 2385–2394.
- ⁶Arunajatesan, S., "Evaluation of two-equation RANS models for simulation of jet-in-crossflow problems," *50th AIAA Aerospace Sciences Meeting*, 2012.
- ⁷Edeling, W. N., Cinnella, P., Dwight, R. P., and Bijl, H., "Bayesian estimates of parameter variability in $k - \epsilon$ turbulence model," *Journal of Computational Physics*, Vol. 258, 2014, pp. 73–94.
- ⁸Guillas, S., Glover, N., and Malki-Epshtain, L., "Bayesian calibration of the constants of the $k - \epsilon$ turbulence model for a CFD model of street canyon flow," *Computer Methods in Applied Mechanics and Engineering*, Vol. 279, 2014, pp. 536–553.
- ⁹Duraisamy, K., Zhang, Z. J., and Singh, A. P., "New approaches in turbulence and transition modeling using data-driven techniques," *Proceedings of the 53rd AIAA Aerospace Sciences Meeting, Kissimmee, FL*, 2015, Paper No. AIAA-2015-1284.
- ¹⁰Tracey, B., Duraisamy, K., and Alonso, J. J., "A machine learning strategy to assist turbulence model development," *Proceedings of the 53rd AIAA Aerospace Sciences Meeting, Kissimmee, FL*, 2015, Paper No. AIAA-2015-1287.
- ¹¹Zhang, Z. J. and Duraisamy, K., "Machine learning methods for data-driven turbulence modeling," *Proceedings of the 22nd AIAA Computational Fluid Dynamics Conference*, 2015, Paper No. AIAA-2015-2460.
- ¹²Parish, E. and Duraisamy, K., "Quantification of turbulence modeling uncertainties using full field inversions," *Proceedings of the 22nd AIAA Computational Fluid Dynamics Conference*, 2015, Paper No. AIAA-2015-2459.
- ¹³DeChant, L., Ray, J., Lefantzi, S., Ling, J., and Arunajatesan, S., " $k - \epsilon$ Turbulence Model Parameter Estimates Using an Approximate Self-similar Jet-in-Crossflow Solution," 2016, Submitted, AIAA Aviation 2017.
- ¹⁴Ray, J., Lefantzi, S., Arunajatesan, S., and Dechant, L., "Bayesian Parameter Estimation of a $k - \epsilon$ Model for Accurate Jet-in-Crossflow Simulations," *AIAA Journal*, Vol. 54, No. 8, 2016, pp. 2432–2448.
- ¹⁵Brinckman, K. W., Calhoon, W. H., and Dash, S. M., "Scalar fluctuation modeling for high-speed aeropropulsive flows," *AIAA Journal*, Vol. 45, No. 5, 2007, pp. 1036–1046.
- ¹⁶Haario, H., Laine, M., Mira, A., and Saksman, E., "DRAM-Efficient adaptive MCMC," *Statistics and Computing*, Vol. 16, No. 4, 2006, pp. 339–354.
- ¹⁷R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2012, ISBN 3-900051-07-0.
- ¹⁸Soetaert, K. and Petzoldt, T., "Inverse Modelling, Sensitivity and Monte Carlo Analysis in R Using Package FME," *Journal of Statistical Software*, Vol. 33, No. 3, 2010, pp. 1–28.
- ¹⁹Silverman, B. W., *Density Estimation for Statistics and Data Analysis*, Chapman and Hall, 1986.
- ²⁰Raftery, A. and Lewis, S. M., "Implementing MCMC," *Markov Chain Monte Carlo in Practice*, edited by W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, Chapman and Hall, 1996, pp. 115–130.
- ²¹Warnes, G. R. and with contributions by Robert Burrows, *mcgibbsit: Warnes and Raftery's MCGibbsit MCMC diagnostic*, 2011, R package version 1.0.8.
- ²²Raftery, A. E., Gneiting, T., Balabdaoui, F., and Polakowski, M., "Using Bayesian Model Averaging to Calibrate Forecast Ensembles," *Monthly Weather Review*, Vol. 133, No. 5, 2005, pp. 1155–1174.

²³Ray, J., Lefantzi, S., Arunajatesan, S., and Dechant, L., "Learning an eddy viscosity model using shrinkage and Bayesian calibration: A jet-in-crossflow case study," *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part B: Mechanical Engineering*, Accepted, March 2017.