Learning an eddy viscosity model using shrinkage and Bayesian calibration: A jet-in-crossflow case study

Jaideep Ray[*]
MS 9159, PO Box 969, Sandia National Laboratories, Livermore, California 94550.
Email: jairay@sandia.gov

Sophia Lefantzi
MS 9152, PO Box 969, Sandia National Laboratories, Livermore, California 94550.
Email: slefant@sandia.gov

Srinivasan Arunajatesan
MS 0825, PO Box 5800, Sandia National Laboratories, Albuquerque, New Mexico 87185.
Email: sarunaj@sandia.gov

Lawrence Dechant
MS 0825, PO Box 5800, Sandia National Laboratories, Albuquerque, New Mexico 87185.
Email: ljdecha@sandia.gov

We demonstrate a statistical procedure for learning a high-order eddy viscosity model from experimental data and using it to improve the predictive skill of a Reynolds-Averaged Navier Stokes simulator. The method is tested in a 3D, transonic jet-in-crossflow configuration. The process starts with a cubic eddy viscosity model developed for incompressible flows. It is fitted to limited experimental jet-in-crossflow data using shrinkage regression. The shrinkage process removes all terms from the model, except an intercept, a linear term and a quadratic one involving the square of the vorticity. The shrunk eddy viscosity model is implemented in a Reynolds Averaged Navier-Stokes simulator and calibrated, using vorticity measurements, to infer three parameters. The calibration is Bayesian and is solved using a Markov chain Monte Carlo method. A three-dimensional probability density distribution for the inferred parameters is constructed, thus quantifying the uncertainty in the estimate. The phenomenal cost of using a 3D flow simulator inside a Markov chain Monte Carlo loop is mitigated by using surrogate models ("curve-fits"). A support vector machine classifier is used to impose our prior belief regarding parameter values, specifically to exclude non-physical parameter combinations. The calibrated model is compared, in terms of its predictive skill, to simulations using uncalibrated linear and cubic eddy viscosity models. We find that the calibrated model, with one quadratic term, is more accurate than the uncalibrated simulator. The model is also checked at a flow condition at which the model was not calibrated.

Nomenclature

C   Parameters to be calibrated
$C_{\varepsilon 2}, C_{\varepsilon 1}$ Parameters in the $k - \varepsilon$ model, to be calibrated
$\mathcal{N}(\mu, \sigma^2)$ Normal distribution with mean μ and standard deviation σ
J Jet-to-crossflow momentum ratio
$\mathcal{R}$ The physically relevant part of the **C** parameter space
S Strain-rate tensor
$\boldsymbol{\Omega}$ Vorticity tensor
AIC   Akaike Information Criterion
DRAM   Delayed Rejection Adaptive Metropolis
CEVM   Cubic eddy Viscosity Model
CVP   Counter-rotating Vortex Pair
EVM   Eddy Viscosity Model
JPDF   Joint Probability Density Function

---

[*]Corresponding author

LASSO   Least Absolute Shrinkage and Selection Operator
LEVM   Linear Eddy Viscosity Model
LS / TS   Learning Set / Testing Set
MAP   Maximum a posteriori
MCMC   Markov chain Monte Carlo
MSE   Mean Square Error
NLEVM   Nonlinear Eddy Viscosity Model
QEVM   Quadratic Eddy Viscosity Model
SI   Supporting Information
SVMC   Support Vector Machine Classifier

1   Introduction

Reynolds-Averaged Navier Stokes (RANS) models are the workhorse of aerodynamic design calculations due to their computational speed vis-à-vis other simulation methods e.g., Large Eddy Simulations (LES). RANS obtains its computational celerity by approximating many of the turbulent processes in the flow. The approximations include equations governing the evolution of the turbulent kinetic energy k, its dissipation rate ε and a host of empirical closure models. The empirical closures also contain parameters whose values are obtained by calibrating to simple canonical flows [1,2]; we will refer to these values as the "nominal" parameter values. Despite their widespread use, RANS models are not particularly predictive for complex turbulent interactions such as jet-in-crossflow, especially at high Reynolds and Mach numbers. Predictive inaccuracy arises from two causes. Firstly, sub-optimal values of model parameters, simply picked from literature, are unlikely to be satisfactory for complex flows. This can be rectified by tuning parameters to relevant experimental/high-fidelity data [3,4]. The second cause, model-form error, arises from "missing physics" such as assumptions of isotropy, neglect of history effects, simplification of various terms in the model equations and assumptions of linear relationship between the turbulent stresses and velocity field strain rates. Here we focus on the last approximation, namely, the linear relationship between the turbulent stresses and the velocity strain rate. In flows with strong curvature and anisotropy, this simple relationship (called the linear eddy viscosity model, LEVM) is unlikely to suffice. The problem can be alleviated by enriching the LEVM with higher-order terms, e.g., quadratic and cubic terms [5], but they in turn introduce more parameters that require calibration. With limited experimental data, properly calibrating *all* parameters becomes impossible because of the danger of overfitting.

In this work, we devise a principled procedure for enriching the LEVM, and calibrating (tuning) parameters to obtain a predictive RANS model for a high Reynolds number interaction of a compressible (Mach number M = 3.73) jet with a transonic (M = 0.8) crossflow. The experiment has been described in [6–8], and it is poorly simulated by RANS (with LEVM) [9]. We hypothesize that simultaneously enriching the LEVM with higher-order terms and calibrating the parameters to experimental JIC (jet-in-crossflow) measurements could yield a high-order RANS simulator with predictive skill superior to RANS with LEVM or the cubic eddy viscosity model (CEVM) described in [5]. The techical difficulty is twofold. First, it is unlikely that the experimental measurements are sufficiently informative to allow the estimation of all coefficients in a high-order RANS model (9 parameters for a cubic model). Consequently, we will devise a rigorous method for selecting the higher-order terms with which to enrich the LEVM. Secondly, the approximations in the RANS models, and the limited experimental measurements may not allow the estimation of parameters with much certainty. To address this, we will adopt a Bayesian calibration approach, as developed by us for calibrating RANS models [4], and use it to estimate RANS and eddy viscosity model (EVM) parameters as a joint probability density function (JPDF). The JPDF will capture the uncertainty in the estimate. To the best of our knowledge, such a data-driven procedure to "discover" and configure a high-order RANS model, commensurate with the information content of an experimental dataset, currently does not exist.

2   Background

Jet-in-crossflow (JIC): Jets perpendicular to the long axis of slender aerodynamic bodies are used to spin the vehicle for various applications. These jets convect downstream and generate a complex vortex system that can interact with the fins of the vehicle. This interaction generates a pressure distribution on the fins, which can result in the production of torque that can counter the effects of the jets, thereby reducing their effectiveness. Thus it is critical to accurately predict this interaction for effective aerodynamic design of such vehicles. In an engineering design setting, these predictions are obtained using RANS models with LEVMs (for computational tractability). Their predictive accuracy for JIC simulations leave much to be desired. A previous investigation [9] with $k - \omega$ RANS models showed that the simulations overpredicted turbulent intensities as well as the penetration of the jet into the crossflow. Further, the predicted vortex system is "fatter", indicating that the turbulent diffusion is too large. LES simulations of the same JIC configurations obtained much better agreement with experiments [10,11], but they are too computationally expensive for routine design investigations.

RANS with LEVM: We have previously investigated if the predictive skill of RANS simulations of JIC could be improved via calibration [4]. There, as well as in this paper, we use a compressibility-corrected version of the $k - \varepsilon$ RANS model; descriptions with an embedded LEVM are in [12,13]. As part of this study, we will enrich this particular form of the RANS equations with the CEVM in [5]. Further, we will calibrate two parameters, $C_{\varepsilon 2}, C_{\varepsilon 1}$ in the ε evolution equa-

tion; see Eq. 2 in [4] for details. The nominal values of the parameters in the $k - \varepsilon$ RANS equations are obtained from simple flows and there are inconsistencies between their values in various incarnations of the $k - \varepsilon$ model. Further, these parameters do not agree with values derived from simple incompressible-flow experiments either; see [4, 14] for a discussion. This uncertainty arises from the tendency to treat these parameters as "universal" constants that could be estimated from *any* flow, whereas in reality, they are tunable constants with flow-dependent optimal values.

There have been previous efforts to tune RANS parameters using a Bayesian approach (i.e., in the form of a PDF) as well as to estimate the model-form error. These were primarily performed for simple flows (flat-plate boundary layers, channel flows etc.) [14, 15] due to the enormous cost of Markov chain Monte Carlo (MCMC; [16]) solutions of the Bayesian inverse problems set up to estimate the parameters as PDFs. Investigations into more complicated flows have been performed by assuming that the inaccuracies in RANS predictions were entirely due to model-form errors. The studies in [17–21] augment their turbulence transport model (e.g., the evolution equation for turbulent viscosity in the Spalart-Allmaras model, or the equation for k in $k - \varepsilon$ ones) with a spatially-dependent source term that is estimated from LES data via a gradient descent algorithm. The source terms are thereafter modeled using the local flow variables using machine-learning methods. The same philosophy is employed in [22] where turbulent stress anisotropy was modeled using spatial basis functions and calibrated using iterated ensemble Kalman filters. Ref. [23] develops models of turbulent stress anisotropy in terms of local flow variables using random forest models, whereas [24] develops distance metrics to quantify the difference between a set of training flows and a flow where the trained turbulent models would be used. This distance is used to judge the suitability of the trained models in the particular use case. Ref. [25] develops a genetic algorithm to generate functional forms that could be used to relate local flow properties to turbulent stress anisotropy.

The use of Bayesian estimation of parameters in a complex flow i.e., one that requires the RANS simulator to be replaced with a surrogate model to overcome the phenomenal cost of MCMC solutions, was performed in [26]. Four $k - \varepsilon$ parameters, and a model-form error, were estimated from experimental measurements of turbulent kinetic energy k in a wind-tunnel simulation of flow in a urban canyon. Gaussian Process surrogates were used as proxies for the 3D RANS model. They had to excise a part of the parameter space since it was non-physical. A similar approach was used in our previous work [4] on the estimation of three RANS model parameters and a crude estimate of the model-form error for the JIC dataset used in this study. The paper describes in detail how we isolated the realistic part of the parameter space (called $\mathcal{R}$) and modeled it using a support vector

machine classifier (SVMC). It served as a prior distribution for the parameters being estimated. The 3D RANS model was replaced with polynomial surrogates, with stringent accuracy requirements that could not always be satisfied; consequently, all locations ("probes") with measurements could not be used in the parameter estimation effort. The surrogates and SVMC were combined within an adaptive MCMC algorithm (called DRAM, [27]) to yield JPDFs. They were tested for predictive accuracy in 3 different JIC interactions. The same calibration framework and software will be reused in this study. Note that all the studies above used the LEVM in their RANS simulators and none attempted to enrich it in a data-driven manner; however, the studies in [17–21, 28] performed something similar - they enriched the turbulent transport model instead.

Nonlinear eddy viscosity models (NLEVM): While previous studies have tried to improve parameter values, they have not addressed some fundamental shortcomings of the LEVM itself. LEVMs are incapable of accommodating anisotropy in turbulent stresses, and are also insensitive to curvature in the mean-flow streamlines e.g., bending of the jet in JIC and swirl e.g., CVP. Quadratic terms (leading to quadratic eddy viscosity models, QEVM) allow anisotropy in the normal turbulent stresses while cubic terms sensitize them to swirl and streamline curvatures. Reviews of QEVM and CEVM (collectively, nonlinear eddy viscosity models, NLEVM) as well as inter-comparisons within $k - \varepsilon$ models are available for bluff bodies [29], shock-boundary layer interactions [30] and turbulent jets [31]. NLEVMs model the turbulent stress τ_{ij} as (Eq. 2 in [5])

$$-\tau_{ij} = \overline{u_i' u_j'} = \underbrace{\frac{2}{3}\delta_{ij}k - \nu_T S_{ij}}_{\text{Intercept \& Linear}} + \underbrace{\nu_T \frac{k}{\tilde{\varepsilon}} \sum_{l=1}^{3} c_l f_l(\boldsymbol{S}, \boldsymbol{\Omega})}_{\text{Quadratic}}$$

$$+ \underbrace{\nu_T \left(\frac{k}{\tilde{\varepsilon}}\right)^2 \sum_{l=4}^{7} c_l f_l(\boldsymbol{S}, \boldsymbol{\Omega})}_{\text{Cubic}}, \quad (1)$$

where $\nu_T = C_\mu k^2 / \tilde{\varepsilon}$,

$$S_{ij} = \left(\frac{\partial}{\partial x_j}U_i + \frac{\partial}{\partial x_i}U_j\right) \quad , \quad \Omega_{ij} = \left(\frac{\partial}{\partial x_j}U_i - \frac{\partial}{\partial x_i}U_j\right)$$

$$\tilde{S} = \frac{k}{\tilde{\varepsilon}}\sqrt{\frac{S_{ij}S_{ij}}{2}} \quad , \quad \tilde{\Omega} = \frac{k}{\tilde{\varepsilon}}\sqrt{\frac{\Omega_{ij}\Omega_{ij}}{2}}$$

$$\tilde{\varepsilon} = \varepsilon - 2\nu \left(\frac{\partial k^{1/2}}{\partial x_j}\right)^2 \text{ and } C_\mu = C_\mu(\tilde{S}, \tilde{\Omega}). \quad (2)$$

Various QEVMs differ in terms of the coefficients $c_1, \ldots, c_3$; also, C_μ is a constant in most of them. CEVMs too vary in

the values used for c_l, as well as the functional form for C_μ in terms of $\boldsymbol{S}$ and $\boldsymbol{\Omega}$. The coefficients and functional forms are summarized in [29, 30]. In bluff-body flows, NLEVMs produce results closer to experimental measurements, primarily because they simulate the stagnation points (irrotational regions of large strain rates) accurately [29]. Simple turbulent jets are better simulated using LEVM (rather than NLEVM) except for impinging jets [31]. For shock-boundary layer simulations, [30] showed that $C_\mu = C_\mu(\boldsymbol{S})$ provided better results and cubic terms were not very useful for those flows. Note that these jets did not have significant curved streamlines or swirl. Studies in [32] found that QEVM, calibrated to different flows, resulted in a huge variability in c_1, c_2, c_3 and could not be used as a "universal" model. They formulated Eq. 1 instead. While it captures some of the anisotropy in normal turbulent stresses, it is not sufficient for flows with significant heat transfer through boundaries, and require an extra evolution equation for the anisotropy. Note that there are some NLEVMs that do not adhere to the series structure in Eq. 1, e.g. [33].

Experimental data: The wind-tunnel experiment that provides us with calibration data is described in [6–8]. A schematic of the test section is in Fig. 1 of [4].

The test section is 0.5m long and has a square cross-section of side 304.8mm. The flow is from left to right, with variable Mach numbers from 0.5 to 0.8. The supersonic jet at Mach 3.73 is introduced into the flow from a nozzle mounted on the floor of the tunnel test section. Particle Image Velocimetry (PIV) is used to measure velocities on the spanwise mid-plane (plane of symmetry) extending over a region 200 mm to 400 mm downstream of the jet. For some cases, measurements are also made on the cross-plane, perpendicular to the flow direction, where the vorticity field is best captured. Measurements inside a window $\mathcal{W}$, $0 \leq z \leq 0.04\text{m}, 0.03\text{m} \leq y \leq 0.11\text{m}$, which captures one of CVP (see Fig. 1 in [4]) are used for calibration; it contains a 8×28 grid of observations. Here y is the vertical coordinate axis. We will call these locations "probes". A more comprehensive description of the experimental setup is in [4].

3 Eddy Viscosity Model

The CEVM described in Eq. 1 has 7 parameters and it is unlikely the information content of our experimental dataset will allow their estimation *in toto*. Consequently, we simplify the EVM. The measurements on the midplane at made at 5 downstream locations; each location contains 63 measurement points ("probes"). Thus, the midplane measurements provide us with $5 \times 63 = 315$ "probe" locations where τ_{11}, τ_{22} and τ_{12} measurements exist. Further, PIV measurements of the mean flow on the midplane provide us with a well-resolved velocity measurement, with experimental values of S_{ij} and Ω_{ij} computed at the 315 midplane probes. This allows us to set up a linear problem for CEVM-predicted tur-
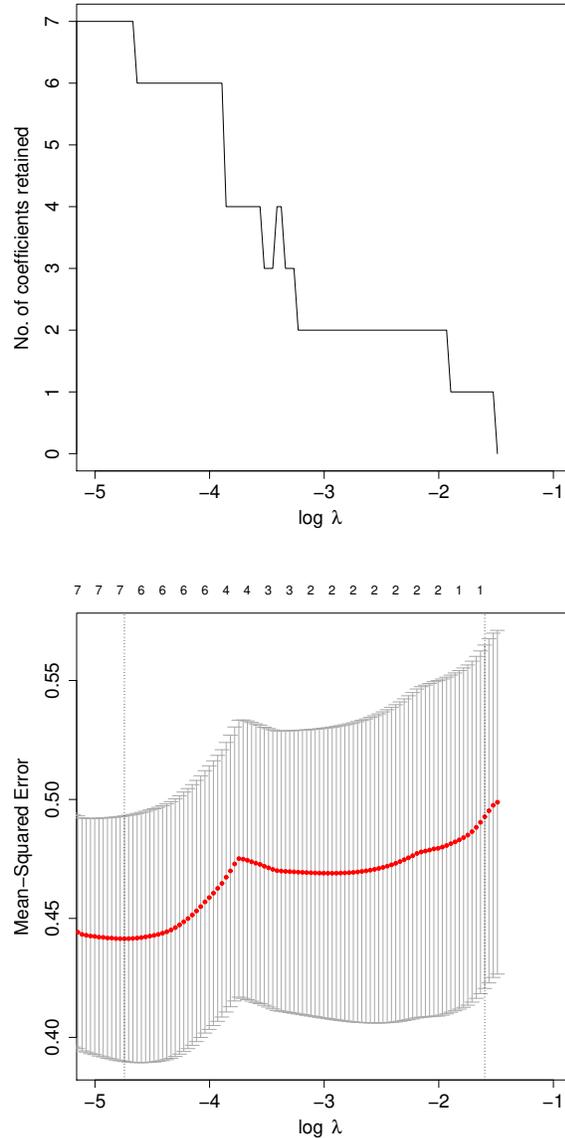
Fig. 1. Top: Removal of CEVM coefficients c_i as λ is increased. Bottom: Mean prediction error and ± 2 standard deviation bounds of prediction error, as a function of $\log(\lambda)$. λ_{min} (left vertical line) and λ_{1se} (right vertical line) are also shown.

bulent stresses $\mathbf{V} = \mathbf{A}(\boldsymbol{S}, \boldsymbol{\Omega}, \tilde{\varepsilon})\mathbf{c}$, where the matrix $\mathbf{A}(\boldsymbol{S}, \boldsymbol{\Omega}, \tilde{\varepsilon})$ contains the strain-rate and vorticity terms in Eq. 1 evaluated at the probes and $\mathbf{c} = \{c_1, \ldots, c_7\}$. $\mathbf{A}(\boldsymbol{S}, \boldsymbol{\Omega}, \tilde{\varepsilon})$ has as many rows as the number of observations (nominally, thrice the number of probes). Unfortunately, $\tilde{\varepsilon}$ cannot be measured (and is typically available only inside a RANS simulation), and consequently we approximate it.

We assume that the rate of production of turbulent kinetic energy is balanced by the dissipation, allowing us to

set $\epsilon = \tau_{ij}\frac{\partial}{\partial x_j}U_i$; thereafter $\tilde{\epsilon}$ can be computed using Eq. 2. This assumption holds true only very approximately and at best provides an order-of-magnitude estimate of $\tilde{\epsilon}$. For many probes, $\tilde{\epsilon}$ assumes (non-physical) negative values and the probes have to be eliminated. Further, probes in the non-turbulent freestream too have to be excised, as both $\boldsymbol{S}$ and $\boldsymbol{\Omega}$ are nearly zero there. In addition, we remove the probes that lie near the wall, in an attempt to obtain a CEVM that is tuned for the CVP. Removing these measurements leaves us with about 170 measurements that have large $\boldsymbol{S}, \boldsymbol{\Omega}$ and τ_{ij}. Consolidating the measured, non-zero τ_{ij} into $\mathbf{V}^{obs}$, we pose the estimation of $\mathbf{c}$ as one of shrinkage:

$$\underset{\mathbf{c}}{\text{minimize}} \quad \|\mathbf{V}^{obs} - \mathbf{A}(\boldsymbol{S},\boldsymbol{\Omega},\tilde{\epsilon})\mathbf{c}\|_2^2 + \lambda\|\mathbf{c}\|_1 \qquad (3)$$

The first term in the objective function in Eq. 3 drives $\mathbf{c}$ to values that minimize the discrepancy between measurements $\mathbf{V}^{obs}$ and the CEVM prediction $\mathbf{A}(\boldsymbol{S},\boldsymbol{\Omega},\tilde{\epsilon})\mathbf{c}$. The second term drives the optimizer to set elements of $\mathbf{c}$ to zero as long as the first term does not degrade. This allows us to "shrink" the CEVM, retaining only the terms that can be supported by the observational data, and prevents overfitting the CEVM model to limited $\mathbf{V}^{obs}$ data. The penalty λ plays a crucial role and its optimal value is obtained via 11-fold cross-validation (CV). LASSO (Least Absolute Shrinkage and Selection Operator, [34]), as implemented in R [35] via the `glmnet` package [36] is used to solve the linear inverse problem (Eq. 3). The measured turbulent stresses $\mathbf{V}^{obs}$ and the CEVM-predicted ones $\mathbf{V}$ contain non-dimensionalized normal and shear stresses $\widetilde{\tau_{ij}} = \tau_{ij}/T_{ij}$, where $T_{ij} = \overline{|\tau_{ij}^{obs}|}$ is the mean of absolute values of the measured turbulent stresses. Thus a different normalization constant is computed for τ_{11}, τ_{22} and τ_{12}.

In Fig. 1 (top) we plot the number of coefficients retained in the CEVM model as λ is increased; it is clear that as the penalty increases, $\mathbf{c}$ becomes increasing sparse, retaining the most important coefficients c_i required to minimize the $\| : \|_2$ term. For extreme λ, the CEVM reduces to the intercept and linear terms. In Fig. 1 (bottom), we plot the results of the 11-fold CV performed for increasing values of λ. For each value of λ we obtain 11 mean-square errors (MSE); their average and ± 2 standard deviation bounds are plotted. We identify λ_{min}, the value for which we obtain the minimum average MSE (left vertical line); the digits on top of the plot show that this corresponds to a CEVM with all coefficients retained. However, we also obtain λ_{1se}, a value of λ for which the mean MSE is one standard deviation away from the one obtained via λ_{min}. Further, this model requires the retention of just one coefficient, c_3, in the CEVM. This corresponds to the term with Ω_{ij}^2; in a strongly vortical flow such as JIC, this is not entirely surprising.

Thus the NLEVM that we will use consists of 2 terms - the linear term in Eq. 1 and *one* quadratic term involving Ω_{ij}^2. Note that we do not show the *value* of c_3 anywhere; we merely use LASSO to perform a data-driven simplification of the CEVM. This is because the gross approximation used to obtain $\tilde{\epsilon}$, for use in Eq. 3, makes the value of c_3 so obtained untrustworthy. We will calibrate it, along with $C_{\epsilon2}$ and $C_{\epsilon1}$, later in Sec. 5.

Note that it is possible that the quadratic term was retained because of its ability to correct the errors in the coarse approximation of $\tilde{\epsilon}$, and consequently has little ability to favorably impact flow predictions. This will be investigated in Sec. 5 where we will calibrate c_3, the factor that multiplies the quadratic term, using flow observations. If it turns out that c_3's prior and posterior distributions do not differ much, then the flow contains no information on the quadratic term in the NLEVM, and its retention was driven by the coarse approximation of $\tilde{\epsilon}$. If the converse is true, i.e., c_3 is constrained by flow observations, then the quadratic term merits being included in an EVM.

4 Posing the Calibration Problem

We seek to infer $\mathbf{C} = (c_3, C_{\epsilon2}, C_{\epsilon1})$ from measurements of vorticity on the crossplane, obtained from the experiments reviewed in Sec. 2. The method we adopt for doing so is described in [4], though there it was used to calibrate a LEVM-based RANS. Reference [4] also contains a verification test of the Bayesian parameter estimation method and software framework used in this study. We provide a summary below.

4.1 Bayesian Inverse Problem for Parameters

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 $\mathbf{C}$. 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. $\boldsymbol{\epsilon} = \{\epsilon_i\}, \epsilon_i \sim \mathcal{N}(0,\sigma^2)$. σ^2 thus provides a crude measure of the model - data misfit after calibration.

Let $P(\mathbf{C},\sigma^2|\mathbf{y}_e)$ be the JPDF of the parameters and the model - data misfit that we seek; it is also called the "posterior" density. 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

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

and by Bayes's theorem,

$$P(\mathbf{C}, \sigma^2 | \mathbf{y}_e) \propto \mathcal{L}(\mathbf{y}_e | \mathbf{C}, \sigma^2) \, \Pi_1(\mathbf{C}) \, \Pi_2(\sigma^2)$$

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

$$\Pi_1(\mathbf{C}) \, \Pi_2(\sigma^2). \qquad (4)$$

This inverse problem is solved via MCMC sampling. Identical to [4], we use DRAM as implemented in the R [35] package FME [37]. The sufficiency of samples is checked using the Raftery-Lewis method [38] implemented in the R package mcgibbsit [39]. A conjugate inverse Gamma prior is used for $\Pi_2(\sigma^2)$, and is the same as in [4]. The samples of $\{\mathbf{C}, \sigma^2\}$ drawn by the sampler are used to develop histograms or PDFs using kernel density estimation. The MCMC method may require $O(10^4)$ samples to construct $P(\mathbf{C}, \sigma^2 | \mathbf{y}_e)$, each of which requires a 3D RANS simulation of JIC to provide $\mathbf{y}_m(\mathbf{C})$. This is impractical and we replace the RANS simulator with a surrogate (also called a statistical emulator) to serve as a computationally inexpensive proxy.

4.2 Surrogate Models

A surrogate model, for the purposes of this study, is the functional dependence between vorticity ω_x (or Ω_{23}) in the streamwise direction as predicted by the RANS model, and the parameters $\mathbf{C}$ that cause it. As in [4], this dependence is posed as a cubic polynomial in $(c_3, C_{\varepsilon 2}, C_{\varepsilon 1})$, and learned by fitting (linear regression) to a training corpus of RANS runs.

We define a cuboidal parameter space $\mathcal{C}_1$ given by the following bounds: $c_3 \in [0.1, 3.5]$, $C_{\varepsilon 2} \in [1.7, 2.5]$ and $C_{\varepsilon 1} \in [1.2, 1.7]$. The last two bounds were obtained from [26]. We draw 2744 samples using a quasi-Monte Carlo sampler (Halton sequence) and seed our RANS simulator SIGMA (see [4] for a description of the numerical scheme and test of mesh convergence; our 3D finite volume meshes have about 10 million cells). Most of the parameter combinations are non-physical and only 222 / 2744 runs survived i.e. converged. These 222 are deemed to define a more realistic parameter space $\mathcal{C}_2$ and we draw more samples from it to provide a training corpus of sufficient size. We train a SVMC with radial kernels (see [4] for a description) to segregate $\mathcal{C}_2$ in $\mathcal{C}_1$ and identify 1500 $\mathbf{C}$ samples that lie in $\mathcal{C}_2$. We run these samples again; 374 / 1500 samples survived. The 374 in turn define $\mathcal{C}_3$ and we generate 1500 points from it using a SVMC. These are run and 1275 / 1500 runs survive. These 1275 samples are deemed to be a parameter space $\mathcal{C}_4$ where a RANS simulator may be expected to run to completion. Note that we use the CEVM in [5], with all CEVM parameters except c_3 set to zero.

The runs using samples drawn from $\mathcal{C}_3$ do not necessarily provide flowfields that resemble physical ones; however, they do provide a vorticity field, however unrealistic, on the
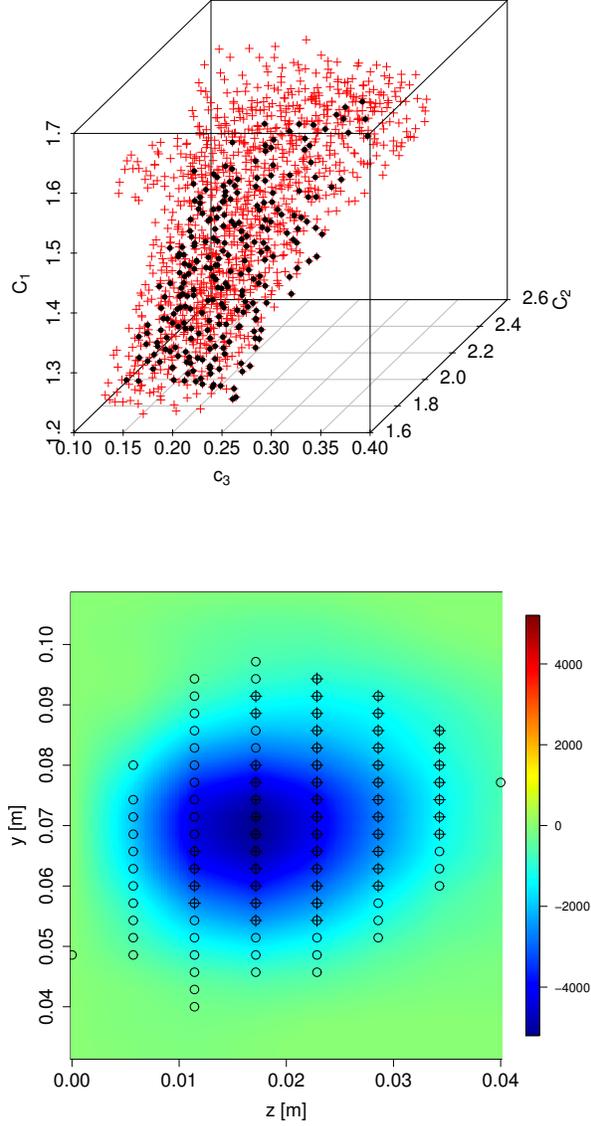
Fig. 2. Top: Plot of the physically realistic part of the parameter space, $\mathcal{R}$ (filled circles) along with the region where the RANS simulator runs without crashing $\mathcal{C}_3$ (crosses). Bottom: The experimental vorticity field as a color plot, with locations with large vorticity (○) and the subset of probes with accurate surrogate models (+).

crossplane. We compare these vorticity fields with the experimental one, compute the root-mean-square-error, and retain the best 20% of the runs. The parameter space occupied by these samples is the physically realistic part $\mathcal{R}$. In Fig. 2 (top) crosses plot out $\mathcal{C}_3$ whereas the filled circles denote $\mathcal{R}$. The supporting information (SI) contains projections on the relevant 2D planes.

The $0.2 \times 1275 = 255$ RANS runs constituting $\mathcal{R}$ are

used to construct surrogate models. We first identify the subset of $8 \times 28 = 224$ probes in $\mathcal{W}$ where vorticity is large (see [4] for details). At these probes, we impute a cubic polynomial in $(c_3, C_{\varepsilon 2}, C_{\varepsilon 1})$ to model vorticity and fit it via linear regression. We use Akaike Information Criterion (AIC) to simplify the model; for most probes, the models reduce to quadratic polynomials in $(c_3, C_{\varepsilon 2}, C_{\varepsilon 1})$. The surrogate models so obtained are not very robust (small changes in the training corpus could lead to very different surrogates). Consequently, as in [4], we perform 50 rounds of repeated random subsampling (a form of crossvalidation) to compute robust surrogates. Surrogates whose testing set errors are less than 10% and where the ratio of learning and testing set errors deviate from 1 by less than 15% are deemed fit for calibration purposes. Thus accurate surrogate models could not be constructed for all probes with large vorticity. We plot the testing set errors and the ratios in Fig. 2 of the SI. In Fig. 2 we plot the 90 probes inside $\mathcal{W}$ with large vorticity. A subset of 50 / 90 probes, plotted with "plus" signs, are the ones where we could construct accurate surrogate models. Thus $N_p = 50$.

Altogether, about 6000 3D RANS JIC simulations were conducted. Each simulation takes approximately 12,000 core-hours on a PowerPC A2 processor.

4.3 Constructing an Informative Prior $\Pi_1(\mathbf{C})$

The surrogate models constructed in Sec. 4.2 are only valid inside $\mathcal{R} \subset \mathcal{C}_4$ and consequently the posterior density (Eq. 4) must be solved within that constraint. We employ the prior belief:

$$\Pi_1(\mathbf{C}) = \begin{cases} 1 \text{ for } \mathbf{C} \in \mathcal{R} \\ 0 \text{ otherwise.} \end{cases} \quad (5)$$

We enforce this prior belief using a SVMC.

Fundamentally, we seek a binary function $\zeta(\mathbf{C})$ in $(c_3, C_{\varepsilon 2}, C_{\varepsilon 1})$ space that could be used to decide whether a parameter combination $\mathbf{C} \in \mathcal{R}$. We set $\zeta(\mathbf{C}) = 1$ at the 255 / 1275 that define $\mathcal{R}$ and $\zeta(\mathbf{C}) = -1$ elsewhere. Then $\zeta(\mathbf{C}) = 0$ denotes $\partial \mathcal{R}$, the demarcation of the physically realistic part of the parameter space. This function is approximated by the SVMC with radial kernels, as described in [4]. We test the SVMC's accuracy via repeated random subsampling cross-validation. We randomly split the training data into Learning Set / Testing Set (LS/TS) pairs, with 85% of the training data lying in the LS. The SVMC is trained on the LS and is tested on the TS to obtain the mis-classification rate. This is performed 50 times, with different LS/TS pairs, and an average mis-classification rate is computed. We find it to be 6.5%.

The SVMC is used to compute $\zeta_m(\mathbf{C})$, the approximation of $\zeta(\mathbf{C})$, at arbitrary $\mathbf{C}$, and identifies whether $\mathbf{C} \in \mathcal{R}$.

Thereafter, the computation of $\Pi_1(\mathbf{C})$ is straightforward via Eq. 5.

5 Results

The Bayesian inverse problem in Eq. 4 requires 80,000 MCMC steps to converge. These samples are thinned by a factor of 5 (i.e., 1 in 5 samples is retained) and used to describe the posterior density. In Fig. 3, we plot the marginalized PDFs for all the parameters as well as σ^2. The nominal values of the parameters are also plotted. We see that the nominal value for $C_{\varepsilon 1}$ approximately agrees with the maximum a posteriori value (MAP, peak of the PDF), but for the rest of the parameters the MAP values and the nominal ones are quite different. In fact, the nominal values are in the tails of the prior distribution, indicating their suboptimality. The PDFs of all three parameters are bimodal, another sign of the complexity of the posterior density. Further, the marginalized PDF for c_3 is quite different from its prior, indicating that the observations contain information on c_3. In addition, we plot the model-data misfit, that has units of the calibration variable/observable, the vorticity on the crossplane. The variation of the magnitude of vorticity predicted by the best set of parameters ($\mathbf{C}_{opt}$; see below) i.e., its $5^{\text{th}}, 25^{\text{th}}, 50^{\text{th}}, 75^{\text{th}}$, and 95^{th} percentiles are $\{0.007, 0.039, 0.76, 2.4, 4.1\} \times 10^3 \text{ sec}^{-1}$. The experimental counterparts are $\{0.0093, 0.069, 0.184, 0.89, 3.8\} \times 10^3 \text{ sec}^{-1}$. The sign of the vorticity is negative. Comparing σ with the 95^{th} percentile of the vorticity magnitude, we see that σ is quite large (about 20%). Note that a slight offset in two sharply defined vorticity fields (such as the CVP) can result in a very large σ. A better comparison is a plot of the two vorticity fields; this is discussed below in Fig. 6. In Fig. 3 in the SI, we plot 2D marginal densities for the parameters. We see that they are complex and strong correlations exist between all, and especially for $C_{\varepsilon 2}$ and $C_{\varepsilon 1}$. The two modes in the density are also clearly evident.

First we check the accuracy of the calibration via a posterior predictive test. We randomly select 100 $(c_3, C_{\varepsilon 2}, C_{\varepsilon 1}, \sigma^2)$ samples from those picked by the MCMC sampler and predict the vorticity at the probes with accurate surrogate models. Note that these predictions also contain a realization of the model-data mismatch ε, corresponding to $\mathcal{N}(0, \sigma^2)$. These predictions are then normalized by their experimental counterparts. We use this ensemble to compute the $5^{th}, 50^{th}$ and 95^{th} percentiles, of the prediction at the probes. In Fig. 4, we plot the normalized vorticity at the probes where they are available. The filled $\circ$ is the median prediction and the error bars span the 5^{th}-95^{th} percentile range. The horizontal line at 1 denotes the experimental data. It is clear that the experimental values are consistently bracketed by the error bars. Note that much of the variability (the width of the error bars) is due to ε; Fig. 4 in SI contains the same plot, but without ε (we call this the "pushed-forward
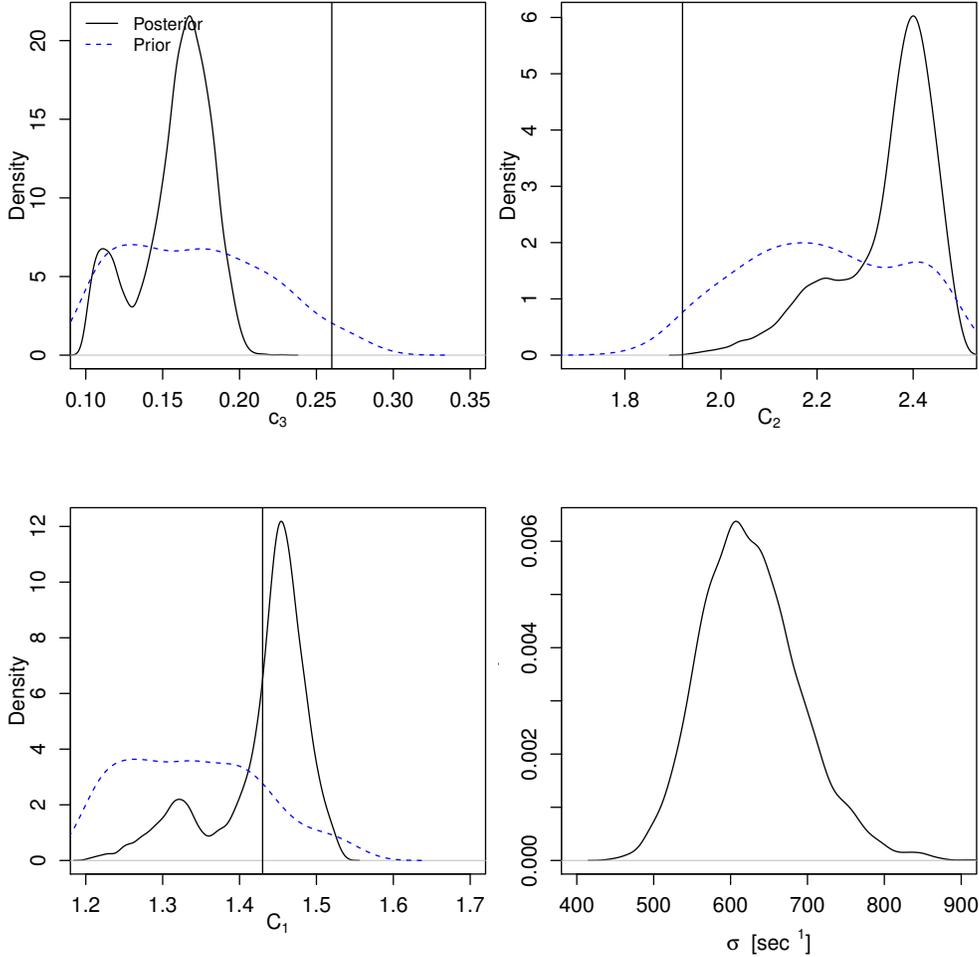
Fig. 3. Marginalized PDFs of $c_3, C_{\varepsilon 2}, C_{\varepsilon 1}$ and σ^2. The dashed line is the prior density (due to $\mathcal{R}$ and the solid line denotes the posterior density. The vertical lines are the nominal values of the parameters. The figure at the bottom right shows magnitude of the data - model misfit. As a comparison to σ, the 95^{th} percentile of the (experimental) vorticity magnitude is $3.8 \times 10^3 \text{sec}^{-1}$.

posterior"). As is clear, the shorter error bars in the "pushed-forward posterior" tests sometimes do not bracket the experimental measurements. Most of ε is due to model-form error (as opposed to measurement uncertainty) and shows the impact of approximate modeling of turbulent processes in the $k - \varepsilon$ model. Next we investigate the response of the RANS-CEVM simulator (RANS with an embedded CEVM with 7 parameters) to the 100 $(c_3, C_{\varepsilon 2}, C_{\varepsilon 1})$ samples from the posterior density. All CEVM parameters except c_3 are set to zero. We perform the simulations and summarize the stream-wise vorticity field inside $\mathcal{W}$ by the circulation, the vorticity field's centroid in the crossplane and the radius of gyration. These provide a point-vortex approximation of the vorticity field, and we refer to them as the "point-vortex metrics" (PVM). These metrics are normalized by their experimental counterparts and plotted in Fig. 5. No ε is added to them. The outline of the boxes are the first and third quartiles of the

predictions, whereas the horizontal line is the median. The whiskers denote the limits beyond which we have outliers. The small variability of the PVM are a reflection of the short error bars in Fig. 4 in the SI. The horizontal line at $y = 1$ indicates the experimental value. The open $\circ$ are predictions using RANS-CEVM, seeded with $c_1, \ldots, c_7$ from [5], while the filled $\diamond$ are predictions with RANS-LEVM, seeded with the nominal values $\{C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1}\} = \{0.09, 1.92, 1.44\}$. We see that while using nominal values of the parameters, LEVM is *better* than CEVM. This is somewhat unexpected as the severe bending of the jet and the CVP should have played to the strengths of CEVM. However, selecting one quadratic term in the CEVM and calibrating it results in a RANS simulator that is clearly superior to either of the two uncalibrated models. This seems to indicate that the inferiority of CEVM may be due to the values of $c_1, \ldots, c_7$ in [5] being unsuitable for transonic JIC interactions. Post-calibration, the PVM
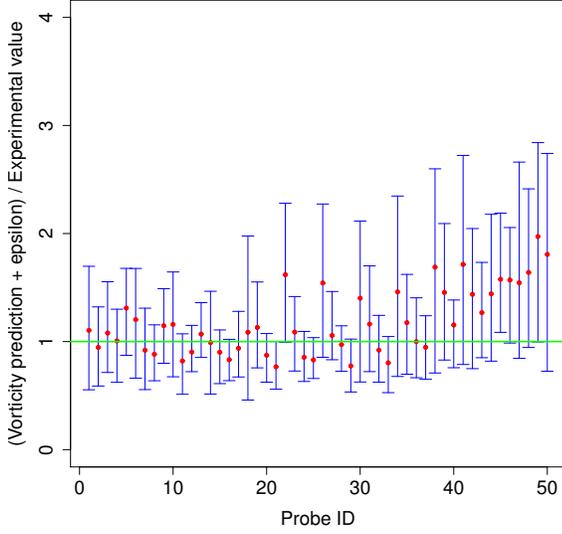
Fig. 4. Results of the posterior predictive test using 100 samples. We plot the predicted vorticity normalized by the measured value at the probes with accurate surrogate models. The horizontal line indicates the experimental measurement. The error bars span the 5^{th} - 95^{th} percentile range and the filled ∘ are the median prediction.
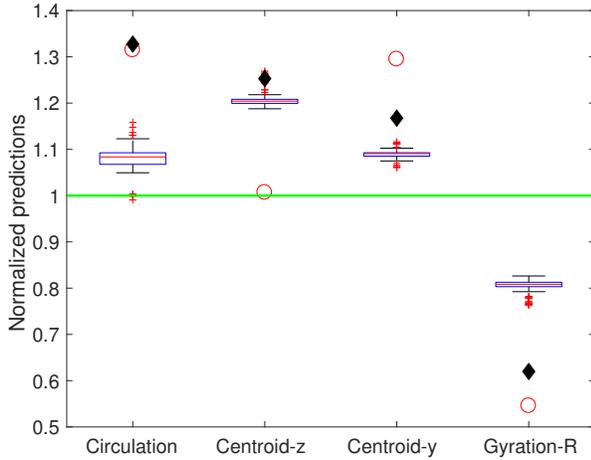
Fig. 5. Box-and-whisker plots of the posterior samples' runs with RANS-CEVM, for the PVM. The horizontal line denotes experimental results. The open ∘ are prediction using nominal CEVM parameters [5] whereas the filled ◇ are predictions using an LEVM with nominal parameters. Model predictions are normalized by their experimental counterparts.

have errors less than 20%, whereas predictions with nominal parameter values lead to 30-40% error. The ensemble of 100 runs also allowed us to obtain an optimal value of $\mathbf{C}$, $\mathbf{C}_{opt} = (c_3, C_{\varepsilon 2}, C_{\varepsilon 1})_{opt} = \{0.169, 2.464, 1.496\}$, that yielded PVM that were closest to 1. We will refer to this as our "Best case".

In Fig. 6, we investigate the vorticity field on the cross-plane further. We plot the streamwise vorticity field predicted by RANS-LEVM (left) and RANS-CEVM (right) using their nominal values. Overlaid on them is the experimental vorticity field, plotted as contours. As is clear, the agreement is not very good; further, the vorticity field produced by the CEVM is substantially worse than LEVM. On the right is the vorticity field predicted by $\mathbf{C}_{opt}$. The improvement is substantial, further reinforcing the unsuitability of RANS-CEVM (with nominal parameter values) for JIC.

The improvement in vorticity field predictions after the estimation of $(c_3, C_{\varepsilon 2}, C_{\varepsilon 1})$ is somewhat expected since that was the calibration variable. We now investigate whether calibration improves the entire flowfield. In order to do so, we compare the streamwise velocity deficit $((U_{max} - u)/U_\infty)$ and normalized vertical velocity (v/U_∞) profiles as a function of the normalized height y/D_j with experimental measurements. Here U_∞ is the freestream velocity (286 m/s) and D_j is the jet diameter (9.83 mm). U_{max} is the maximum velocity in test section and is slightly higher than U_∞ due to addition of extra mass by the jet. We use the midplane flowfields simulated using the 100 $(c_3, C_{\varepsilon 2}, C_{\varepsilon 1})$ samples described above to compute the streamwise velocity deficit and the vertical velocity. We summarize the ensemble's predictions using the mean. These plots are shown in Fig. 7 at three locations 200, 300 and 400 mm downstream of the jet. The top row contains the velocity deficit and the lower row the vertical velocity. Predictions with the uncalibrated parameters (LEVM or CEVM) i.e., the dotted and dashed lines, lead to a jet that sits too far above the true (experimental) jet; this is clear in both the deficit and vertical velocity plots. Further, the vertical velocity is over-predicted (bottom row), indicating a CVP that is too strong. This is in agreement with the circulation being over-predicted (before calibration) as seen in Fig. 5. It is also clear that the LEVM results agree with the experimental data better than the CEVM ones before calibration. However, the improvement in agreement, after calibrating $(c_3, C_{\varepsilon 2}, C_{\varepsilon 1})$, is quite substantial. The mean streamwise velocity deficit and vertical velocity are quite close to the experimental data, with the agreement in vertical velocity being superior, compared to the streamwise velocity deficit. In addition, the agreement in vertical velocity improves as we proceed downstream. Further, the prediction using $\mathbf{C}_{opt}$ (+ symbols in Fig. 7) is almost identical to the ensemble mean.

Finally we check if the calibration performed using a $M = 0.8$ crossflow is applicable at other flow conditions. We use the 100 $(c_3, C_{\varepsilon 2}, C_{\varepsilon 1})$ samples to simulate a JIC interaction with $M = 0.7$ crossflow. The experimental dataset [6] contains measurements of velocity on the midplane for a $M = 0.7$ interaction, but we do not have any crossplane measurements. Consequently we plot the equivalent of Fig. 7 for the $M = 0.7$ interaction in Fig. 8. The findings of Fig. 7 hold here too - LEVM (nominal parameter values) provides
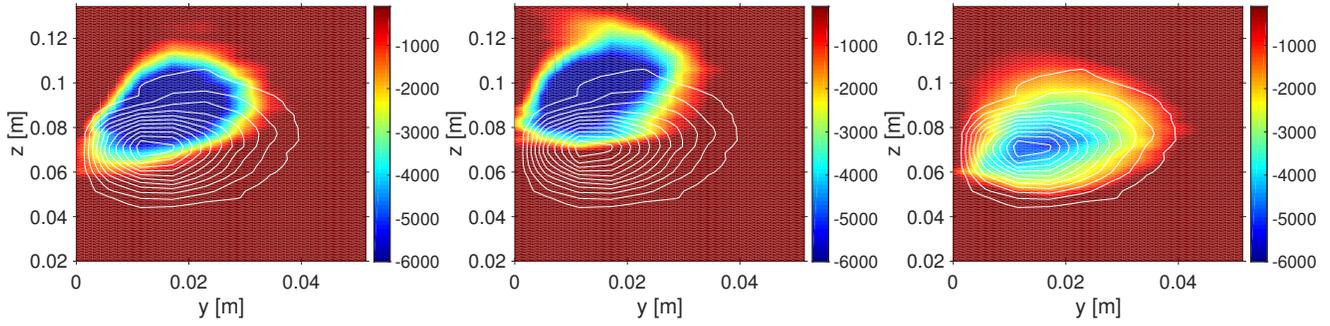
Fig. 6. Plots of simulated streamwise vorticity field (as a flood plot) with contours of experimental vorticity overlaid. Left: Simulations using RANS-LEVM driven by nominal parameters. Middle: RANS-CEVM, driven by the nominal parameters in [5]. Right: Predictions using $\mathbf{C}_{opt}$. The improvement is stark. Note that the scales of the vertical and horizontal axes are different.

better predictions than CEVM (nominal parameter values) and the calibration produced with $M = 0.8$ crossflow leads to a remarkably good prediction (via the ensemble mean) of the $M = 0.7$ experimental results, especially for the vertical velocity (bottom row). Thus there is some degree of robustness in the calibration of $(c_3, C_{\varepsilon2}, C_{\varepsilon1})$.

5.1 Discussion

Predictive errors in RANS simulations are due to, among other causes, the shortcomings of the EVM and the approximate closure relationships in the evolution equations for k and ε. In the results above, we see that simply sensitizing the turbulent stresses to strain-rate, vorticity, streamline curvature etc. serves little purpose unless the parameters accompanying the model are also properly calibrated. The lack of free parameters to calibrate thus renders the simple LEVM more predictive for JIC than CEVM when the nominal parameter values are used. Our EVM which has only one quadratic term in CEVM "turned on" (the one multiplying Ω_{ij}^2) provides superior results, but only after calibration. Further since we calibrate $C_{\varepsilon2}$ and $C_{\varepsilon1}$ (which appear in the evolution equation of ε) in addition to c_3, it is not clear whether our calibrated EVM or the ε-equation is responsible for this improvement.

In our previous work [4], we had addressed the improvement of a RANS equation (with LEVM), calibrated to the same experimental data. There we estimated $(C_\mu, C_{\varepsilon2}, C_{\varepsilon1})$ and obtained an improvement in the predictive skill similar to the one seen in this paper (or perhaps even slightly better). Our previous paper allowed us to directly control the dependence of turbulent stresses on strain-rate (via the LEVM) during calibration whereas in the present study, we control its dependence on vorticity (via our NLEVM). However, both the studies provide a similar degree of improvement in predictive skill, inviting the speculation that it may be $(C_{\varepsilon2}, C_{\varepsilon1})$, the parameters in the ε-equation, that play the dominant role in drawing the flowfield closer to experimental measurements. Further, the PDFs for $C_{\varepsilon2}$ and $C_{\varepsilon1}$, drawn

from these two studies (see Fig. 3 and Fig. 5 in [4]), show that their MAP values are somewhat similar - for $C_{\varepsilon2}$ they are far greater than the nominal value of 1.92 whereas $C_{\varepsilon1}$ is quite close to the nominal one of 1.44. The insensitivity of these crucial parameters to the form of the EVM tends to argue that the impact of c_3 (or C_μ in LEVM) may be of secondary importance, and the equations for k and ε may be more important for improving the predictive skill of RANS for JIC interactions.

The impact of adding the Ω_{ij}^2 term to the EVM is clearly seen in Fig. 7 and Fig. 8. As the CVP evolves downstream, it widens i.e., a larger fraction of the flowfield contains large vorticity. The vorticity field modulates τ_{ij}. As is seen in the two figures, agreement of the vertical velocity clearly improves downstream. This dependence was not seen in our previous study [4]. The results above also show that the calibration of $(c_3, C_{\varepsilon2}, C_{\varepsilon1})$ improves the predictive skill of the entire flowfield and not just streamwise vorticity. Further, it is robust across changes in the crossflow Mach number. This is in line with what was observed in our previous paper which did not address enrichment of the EVM [4].

6 Conclusions

In this study we investigate whether enriching a standard linear eddy viscosity model with higher order terms could reduce the model-form errors in RANS simulations of jet-in-crossflow interactions. In particular, we choose a cubic eddy viscosity model [5]. We find that simulations using the linear eddy viscosity model are actually *superior* to those using the cubic model, as long as both are seeded with parameters obtained from literature i.e., they use nominal parameter values. This highlights the crucial role played by the eddy viscosity model's parameters, and the importance of calibrating them to relevant experimental data. Thus the type of model enrichment that may be possible is inextricably linked to the available measurements. We devise a statistical procedure, based on shrinkage, to discover a high-order eddy viscosity model
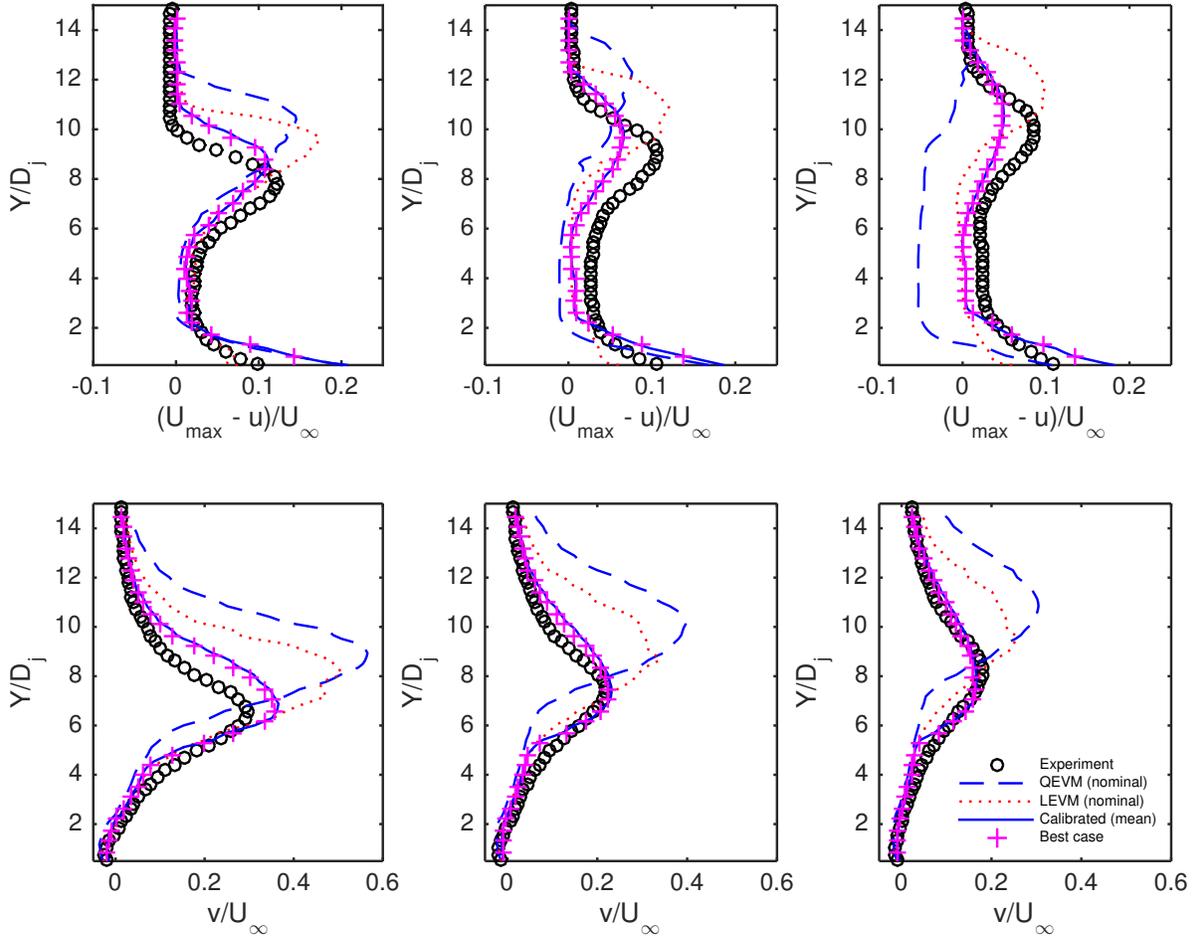
Fig. 7. Plots of streamwise velocity deficit (top row) and vertical velocity (bottom row) at three locations 200, 300 and 400 mm downstream of the jet. Experimental data is plotted using symbols, LEVM (nominal) using the dotted line, the CEVM (nominal) using the dashed line and the ensemble mean of 100 samples from the posterior using the solid line. The $+$ symbols are the predictions using $\mathbf{C}_{opt}$. The crossflow is $M = 0.8$.

in a purely data-driven manner. We find the our experimental dataset can support an additional term in the eddy viscosity model that is quadratic in vorticity. Given the strongly vortical nature of the jet-in-crossflow interaction, this is not surprising.

We embed the new eddy viscosity model in a RANS simulator and calibrate its parameter (c_3) along with two other parameters, $C_{\varepsilon 2}$ and $C_{\varepsilon 1}$, that appear in the evolution equation for ε, the dissipation rate of the turbulent kinetic energy. We use vorticity measurements on the crossplane of a jet-in-crossflow interaction as the calibration observable. The Mach number M of the crossflow is 0.8. The calibration is Bayesian and we compute a joint PDF for the parameters being estimated. We use surrogate models to circumvent the enormous computational expense of Bayesian solutions of inverse problems. We also use informative priors to excise non-physical parts of the parameter space. We find that calibration immensely improves the predictive skill of the sim-

ulator. This improvement is seen in the entire flowfield (and not just vorticity on the crossplane) and holds even when the calibrated parameters are use to predict a jet-in-crossflow interaction with a $M = 0.7$ crossflow. Thus the calibration has some degree of robustness.

The joint calibration of $(c_3, C_{\varepsilon 2}, C_{\varepsilon 1})$ in this study, taken in isolation, does not allow us to ascribe the improvement in predictive skill to the addition of the Ω_{ij}^2 term in the eddy viscosity model. However, in a previous study [4], we had calibrated a RANS simulator with a linear eddy viscosity model to the same experimental dataset. The parameters calibrated were $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$. The improvements in predictive skill and the robustness of the calibration are very similar to the ones obtained in this study. Further, the PDFs for $C_{\varepsilon 2}$ and $C_{\varepsilon 1}$ also show some agreement. Thus the addition of a Ω_{ij}^2 term in the eddy viscosity model had a rather muted effect. This was unexpected given the strongly vortical nature of the jet-in-crossflow interaction, where the primary coherent structure is
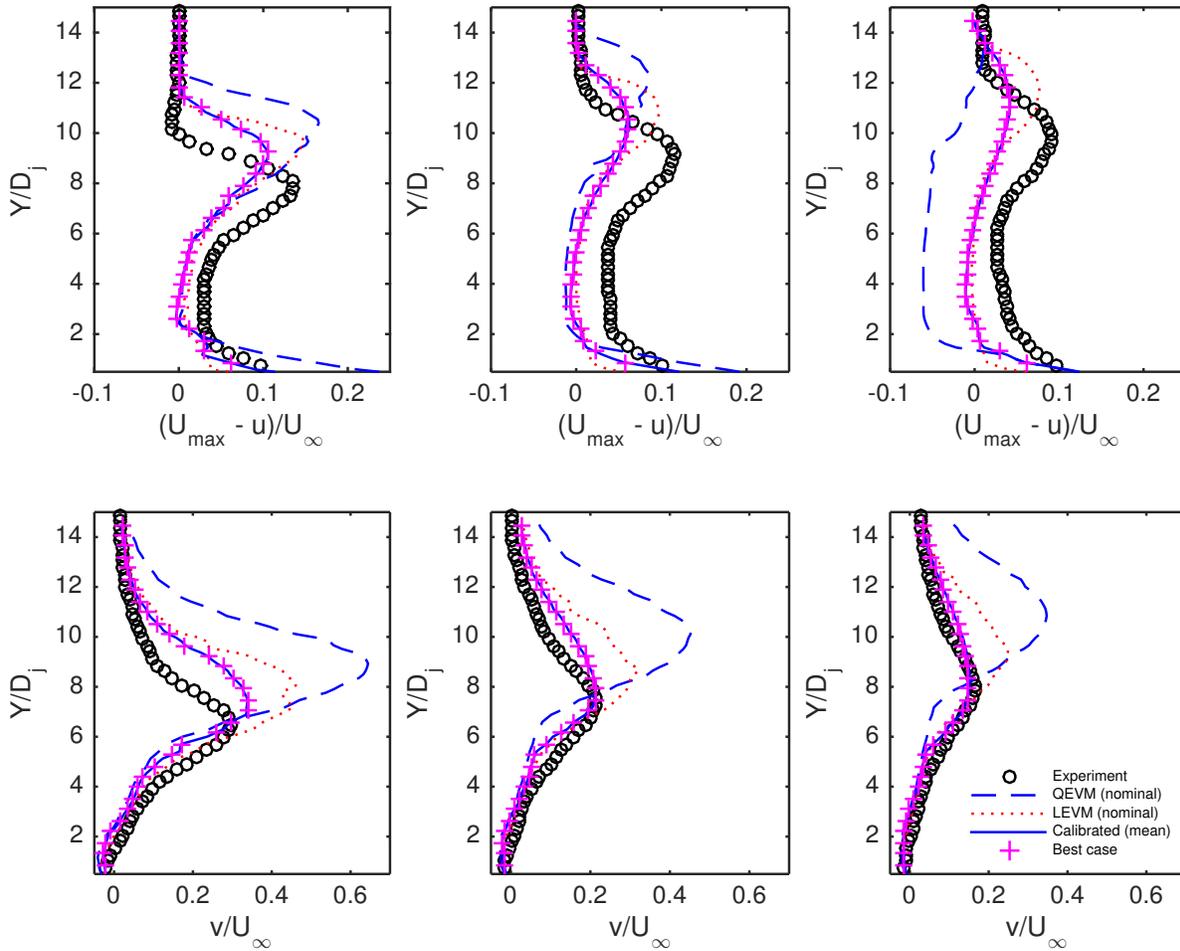
Fig. 8. Plots of streamwise velocity deficit (top row) and vertical velocity (bottom row) at three locations 200, 300 and 400 mm downstream of the jet. Experimental data is plotted using symbols, LEVM (nominal) using the dotted line, the CEVM (nominal) using the dashed line and the ensemble mean of 100 samples from the posterior using the solid line. The + symbols are the predictions using $\mathbf{C}_{opt}$. The crossflow is $M = 0.7$.

a counter-rotating vortex pair. It supports the notion that the modification of the closure models (in the ε-equation) by $C_{\varepsilon 2}$ and $C_{\varepsilon 1}$ plays a significant, if not dominant, role in improving the predictive skill of RANS models. It also implies that the closure models may be equally important contributors to the model-form errors in RANS jet-in-crossflow simulations, along with the eddy viscosity model.

While the enrichment of the eddy viscosity model with a quadratic vorticity term may not have caused large global changes in flowfield, it may have intense local effects e.g., inside the counter-rotating vortex pair. This could be evidenced in a RANS simulation of a jet-in-crossflow interaction where vorticity measurements on the crossplane are available. This particular measurement is not available in our $M = 0.7$ dataset, and the investigation is left for future work.

References

[1] Poroseva, S., and Iaccarino, G., 2001. Simulating separated flows using the $k - \varepsilon$ model. Tech. rep., Center for Turbulence Research, Annual Research

Briefs. Available at `https://ctr.stanford.edu/annual-research-briefs-2001/`.

[2] Sjögren, T., and Johansson, A. V., 2000. "Development and calibration of algebraic nonlinear models for terms in the Reynolds stress transport equation". *Physics of Fluids,* **12**(6), pp. 1554–1572.

[3] Durbin, P. A., 1995. "Separated flow computations using the $k - \varepsilon - v^2$ model". *AIAA Journal,* **33**, pp. 659–664.

[4] Ray, J., Lefantzi, S., Arunajatesan, S., and Dechant, L., 2016. "Bayesian parameter estimation of a k-epsilon model for accurate jet-in-crossflow simulations". *American Institute of Aeronautics and Astronautics Journal,* **54**(8), pp. 2432–2448. `http://dx.doi.org/10.2514/1.J054758`.

[5] Craft, T. J., Launder, B. E., and Suga, K., 1996. "Development and application of a cubic eddy-viscosity model of turbulence". *International Journal of Heat and Fluid Flow,* **17**(2), pp. 108–115.

[6] Beresh, S. J., Henfling, J. F., Erven, R. J., and Spillers, R. W., 2005. "Penetration of a transverse supersonic jet into a subsonic compressible crossflow". *AIAA Journal,* **43**(2), pp. 379–389.

[7] Beresh, S. J., Henfling, J. F., Erven, R. J., and Spillers, R. W., 2005. "Turbulent characteristics of a transverse supersonic jet in a subsonic compressible crossflow". *AIAA Journal,* **43**(11), pp. 2385–2394.

[8] Beresh, S. J., Henfling, J. F., Erven, R. J., and Spillers, R. W., 2006. "Crossplane velocimetry of a transverse supersonic jet in a transonic crossflow". *American Institute of Aeronautics and Astronautics Journal,* **44**(12), pp. 3051–3061.

[9] Arunajatesan, S., 2012. "Evaluation of two-equation RANS models for simulation of jet-in-crossflow problems". In 50th AIAA Aerospace Sciences Meeting. AIAA Paper 2012-1199.

[10] Chai, X., and Mahesh, K., 2011. "Simulations of high speed turbulent jets in crossflow". In 49th AIAA Aerospace Sciences Meeting, Orlando Florida. AIAA Paper 2010-4603.

[11] Genin, F., and Menon, S., 2010. "Dynamics of sonic jet injection into supersonic crossflow". *Journal of Turbulence,* **11**, pp. 1–13.

[12] Brinckman, K. W., Calhoon, W. H., and Dash, S. M., 2007. "Scalar fluctuation modeling for high-speed aeropropulsive flows". *AIAA Journal,* **45**(5), pp. 1036–1046.

[13] Papp, J. L., and Dash, S. M., 2001. "Turbulence model unification and assessment for high-speed aeropropulsive flows". In Proceedings of the 39th AIAA Aerospace Sciences Meeting and Exhibit. Paper No. AIAA-2001-0880.

[14] Edeling, W. N., Cinnella, P., Dwight, R. P., and Bijl, H., 2014. "Bayesian estimates of parameter variability in $k - \varepsilon$ turbulence model". *Journal of Computational Physics,* **258**, pp. 73–94.

[15] Cheung, S. H., Oliver, T. A., E. E. Prudencio, S. P., and Moser, R. D., 2011. "Bayesian uncertainty analysis with applications to turbulence modeling". *Reliability Engineering and System Safety,* **96**, pp. 1137–1149.

[16] Gilks, W. R., Richardson, S., and Spiegelhalter, D. J., eds., 1996. *Markov Chain Monte Carlo in Practice.* Chapman & Hall.

[17] Duraisamy, K., Zhang, Z. J., and Singh, A. P., 2015. "New approaches in turbulence and transition modeling using data-driven techniques". In Proceedings of the 53rd AIAA Aerospace Sciences Meeting, Kissimmee, FL. Paper No. 2015-1284.

[18] Zhang, Z. J., and Duraisamy, K., 2015. "Machine learning methods for data-driven turbulence modeling". In Proceedings of the 22nd AIAA Computational Fluid Dynamics Conference. Paper No. 2015-2460.

[19] Zhang, Z., Duraisamy, K., and Gumerov, N. A., 2015. "Efficient Multiscale Gaussian Process Regression using Hierarchical Clustering". *ArXiv e-prints*.

[20] Pratap Singh, A., Medida, S., and Duraisamy, K., 2016. "Machine Learning-augmented Predictive Modeling of Turbulent Separated Flows over Airfoils". *ArXiv e-prints*.

[21] Singh, A. P., and Duraisamy, K., 2016. "Using field inversion to quantify functional errors in turbulence closures". *Physics of Fluids,* **28**(4), p. 045110.

[22] Xiao, H., Wu, J.-L., Wang, J.-X., Sun, R., and Roy, C., 2016. "Quantifying and reducing model-form uncertainties in reynolds-averaged navier–stokes simulations: A data-driven, physics-informed bayesian approach". *Journal of Computational Physics,* **324**, pp. 115 – 136.

[23] Wang, J.-X., Wu, J.-L., and Xiao, H., 2016. "Physics-Informed Machine Learning for Predictive Turbulence Modeling: Using Data to Improve RANS Modeled Reynolds Stresses". *ArXiv e-prints*.

[24] Wu, J.-L., J.-X., J.-X. W., Xiao, H., and Ling, J., 2016. "Physics-Informed Machine Learning for Predictive Turbulence Modeling: A Priori Assessment of Prediction Confidence". *ArXiv e-prints*, July.

[25] Weatheritt, J., and Sandberg, R., 2016. "A novel evolutionary algorithm applied to algebraic modifications of the RANS stress–strain relationship". *Journal of Computational Physics,* **325**, pp. 22 – 37.

[26] Guillas, S., Glover, N., and Malki-Epshtein, L., 2014. "Bayesian calibration of the constants of the $k - \varepsilon$ turbulence model for a CFD model of street canyon flow". *Computer Methods in Applied Mechanics and Engineering,* **279**, pp. 536–553.

[27] Haario, H., Laine, M., Mira, A., and Saksman, E., 2006. "DRAM-Efficient adaptive MCMC". *Statistics and Computing,* **16**(4), pp. 339–354.

[28] Parish, E., and Duraisamy, K., 2015. "Quantification of turbulence modeling uncertainties using full field inversions". In Proceedings of the 22nd AIAA Computational Fluid Dynamics Conference. Paper No. 2015-2459.

[29] Bazdidi-Tehrani, F., Mohammadi-Ahmar, A., Kianamsouri, M., and Jadidi, M., 2015. "Investigation of various non-linear eddy viscosity turbulence models for simulating flow and pollutant dispersion on and around a cubical model building". *Building Simulation,* **8**(2), pp. 149–166.

[30] Loyau, H., Batten, P., and Leschziner, M. A., 1998. "Modelling shock/boundary-layer interaction with nonlinear eddy-viscosity closures". *Flow, Turbulence and Combustion,* **60**(3), pp. 257–282.

[31] Balabel, A., and El-Askary, W. A., 2011. "On the performance of linear and nonlinear k-ε turbulence models in various jet flow applications". *European Journal of Mechanics B / Fluids,* **30**(3), pp. 325–340.

[32] Craft, T. J., Launder, B. E., and Suga, K., 1997. "Prediction of turbulent transitional phenomena with a nonlinear eddy-viscosity model". *International Journal of Heat and Fluid Flow,* **18**(1), pp. 15–28.

[33] Gatski, T. B., and Speziale, C. G., 1992. On explicit algebraic stress models for complex turbulent flows. Tech. Rep. 92-58, Institute for Computer Application in Science and Engineering.

[34] Hastie, T., Tibshirani, R., and Friedman, J., 2008. *The elements of statistical learning*, 2 ed. Springer Series in Statistics. Springer.

[35] R Core Team, 2012. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.

[36] Friedman, J., Hastie, T., and Tibshirani, R., 2010. "Regularization paths for generalized linear models via coordinate descent". *Journal of Statistical Software,* **33**(1), pp. 1–22.

[37] Soetaert, K., and Petzoldt, T., 2010. "Inverse modelling, sensitivity and monte carlo analysis in R using package FME". *Journal of Statistical Software,* **33**(3), pp. 1–28.

[38] Raftery, A., and Lewis, S. M., 1996. "Implementing MCMC". In *Markov Chain Monte Carlo in Practice*, W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, eds. Chapman and Hall, pp. 115–130.

[39] Warnes, G. R., and with contributions by Robert Burrows, 2011. *mcgibbsit: Warnes and Raftery's MCGibbsit MCMC diagnostic*. R package version 1.0.8.

List of figure captions

Figure 1. Top: Removal of CEVM coefficients c_i as λ is increased. Bottom: Mean prediction error and ± 2 standard deviation bounds of prediction error, as a function of $\log(\lambda)$. λ_{min} (left vertical line) and λ_{1se} (right vertical line) are also shown

Figure 2. Top: Plot of the physically realistic part of the parameter space, $\mathcal{R}$ (filled circles) along with the region where the RANS simulator runs without crashing $\mathcal{C}_3$ (crosses). Bottom: The experimental vorticity field as a color plot, with locations with large vorticity (○) and the subset of probes with accurate surrogate models (+).

Figure 3. Marginalized PDFs of $c_3, C_{\varepsilon 2}, C_{\varepsilon 1}$ and σ^2. The dashed line is the prior density (due to $\mathcal{R}$ and the solid line denotes the posterior density. The vertical lines are the nominal values of the parameters. The figure at the bottom right shows magnitude of the data - model misfit. As a comparison to σ, the 95[th] percentile of the (experimental) vorticity magnitude is $3.8 \times 10^3 \text{sec}^{-1}$.

Figure 4. Results of the posterior predictive test using 100 samples. We plot the predicted vorticity normalized by the measured value at the probes with accurate surrogate models. The horizontal line indicates the experimental measurement. The error bars span the 5[th] - 95[th] percentile range and the filled ○ are the median prediction.

Figure 5. Box-and-whisker plots of the posterior samples' runs with RANS-CEVM, for the PVM. The horizontal line denotes experimental results. The open ○ are prediction using nominal CEVM parameters [5] whereas the filled ◇ are predictions using an LEVM with nominal parameters. Model predictions are normalized by their experimental counterparts.

Figure 6. Plots of simulated streamwise vorticity field (as a flood plot) with contours of experimental vorticity overlaid. Left: Simulations using RANS-LEVM driven by nominal parameters. Middle: RANS-CEVM, driven by the nominal parameters in [5]. Right: Predictions using $\mathbf{C}_{opt}$. The improvement is stark. Note that the scales of the vertical and horizontal axes are different.

Figure 7. Plots of streamwise velocity deficit (top row) and vertical velocity (bottom row) at three locations 200, 300 and 400 mm downstream of the jet. Experimental data is plotted using symbols, LEVM (nominal) using the dotted line, the CEVM (nominal) using the dashed line and the ensemble mean of 100 samples from the posterior using the solid line. The $+$ symbols are the predictions using $\mathbf{C}_{opt}$. The crossflow is $M = 0.8$.

Figure 8. Plots of streamwise velocity deficit (top row) and vertical velocity (bottom row) at three locations 200, 300 and 400 mm downstream of the jet. Experimental data is plotted using symbols, LEVM (nominal) using the dotted line, the CEVM (nominal) using the dashed line and the ensemble mean of 100 samples from the posterior using the solid line. The $+$ symbols are the predictions using $\mathbf{C}_{opt}$. The crossflow is $M = 0.7$.

List of table captions

There are no tables.

Supporting Information

Learning an eddy viscosity model using shrinkage and Bayesian calibration:

A jet-in-crossflow case study

J. Ray, S. Lefantzi, S. Arunajatesan & L. Dechant, Sandia National Labs
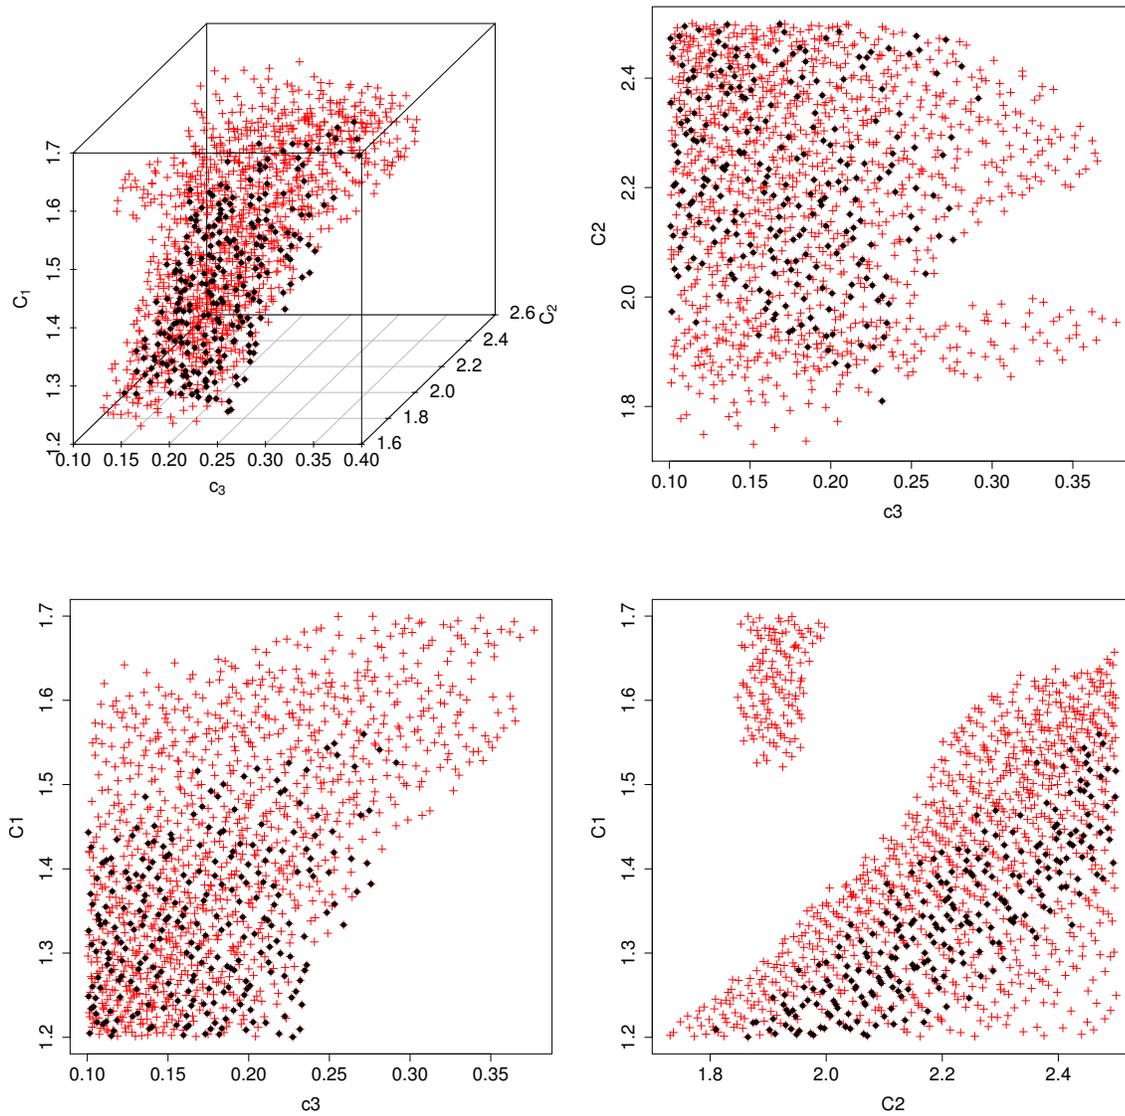
Figure 1: Top left: The set of samples constituting $\mathcal{R}$ (filled ○) and $\mathcal{C}_3$ (+ signs). The rest of the figures plot the samples as projected on 2D planes
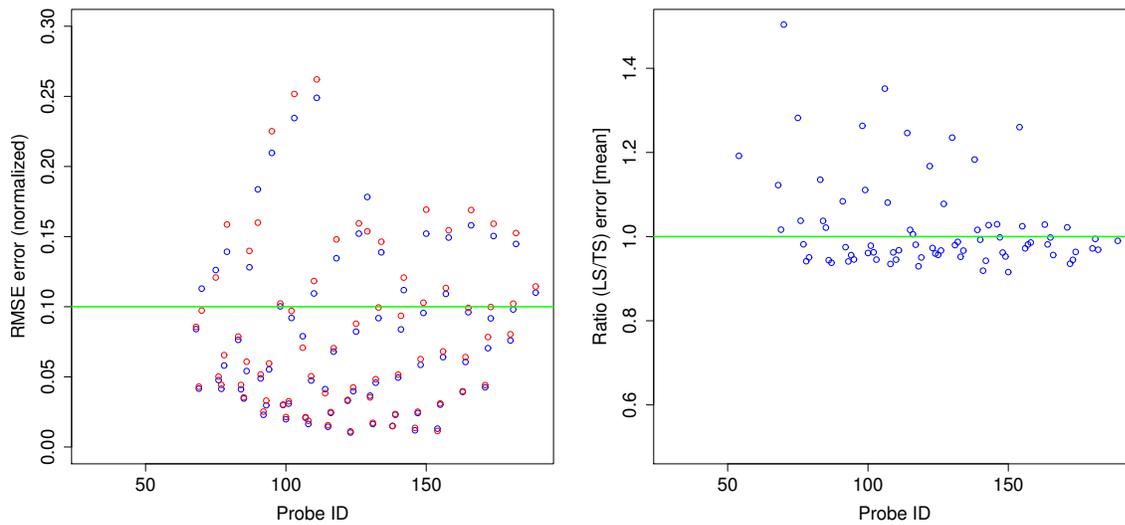
Figure 2: Left: Results from CV when making surrogate models for 90 probes. Root mean squared relative error, averaged over 50 rounds of CV, are plotted. Blue circles denote results from the Learning Set and red circles are represent errors from the Testing Set. We see many probes where relative errors (surrogate modeling errors normalized by RANS results) are less than 10% (horizontal line). Further, both types of errors are about the same in magnitude for most probes. Right: Ratio of RMSE error, TS/LS. Ratios around 1 (±0.15) , as denoted by the horizontal line, denote robust surrogate models that do not substantially overfit the data.
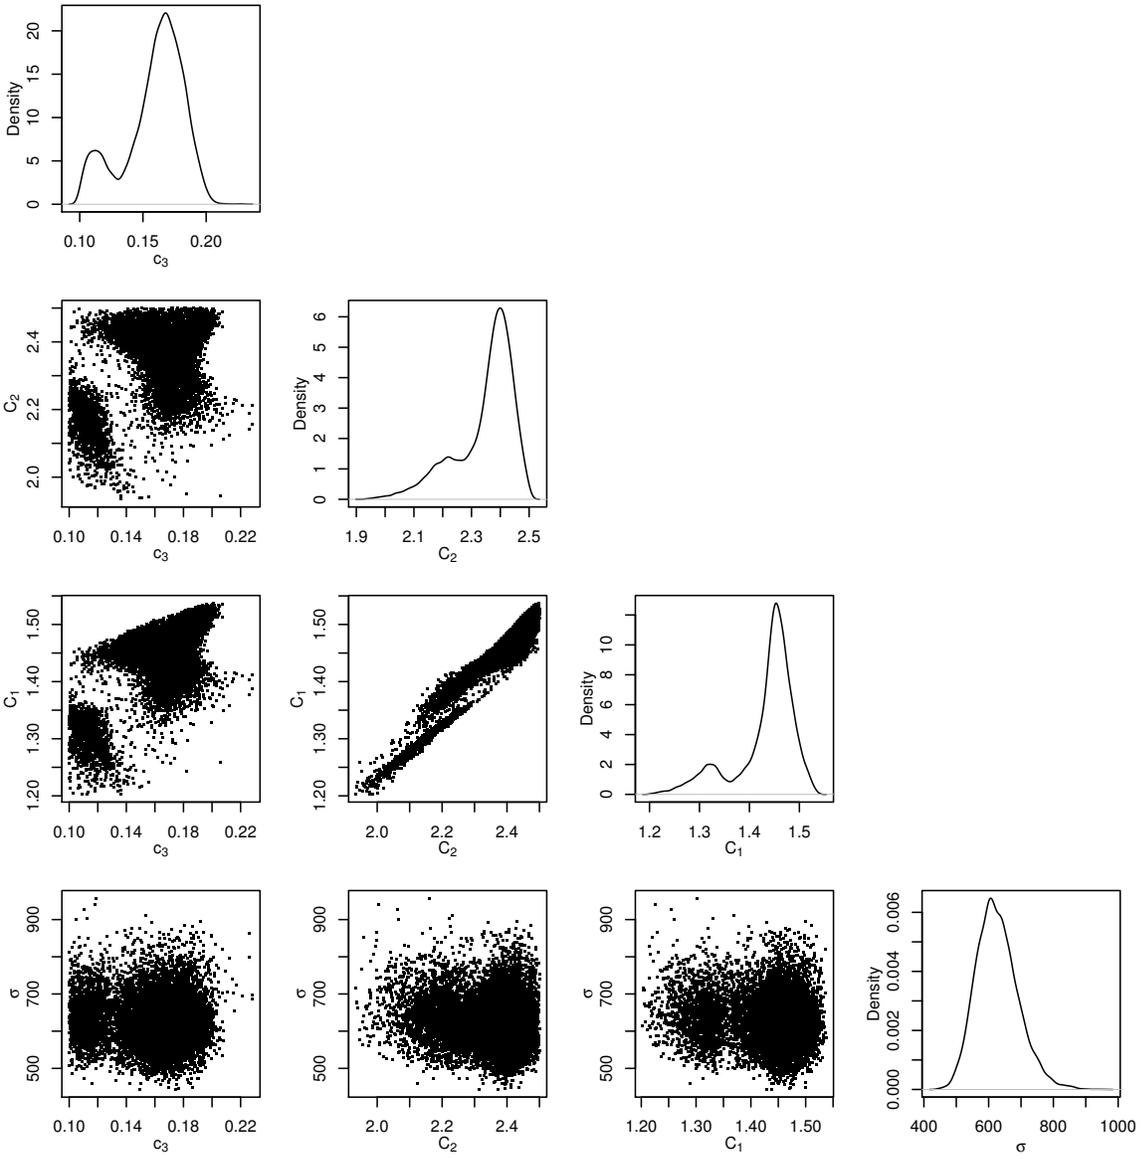
Figure 3: 2D marginalized PDFs of the four-dimensional posterior density in $(c_3, C_{\epsilon 2}, C_{\epsilon 1}, \sigma^2)$ space. Complex correlations are seen between the various parameters. $C_{\epsilon 2}$ and $C_{\epsilon 1}$ show very strong correlation.
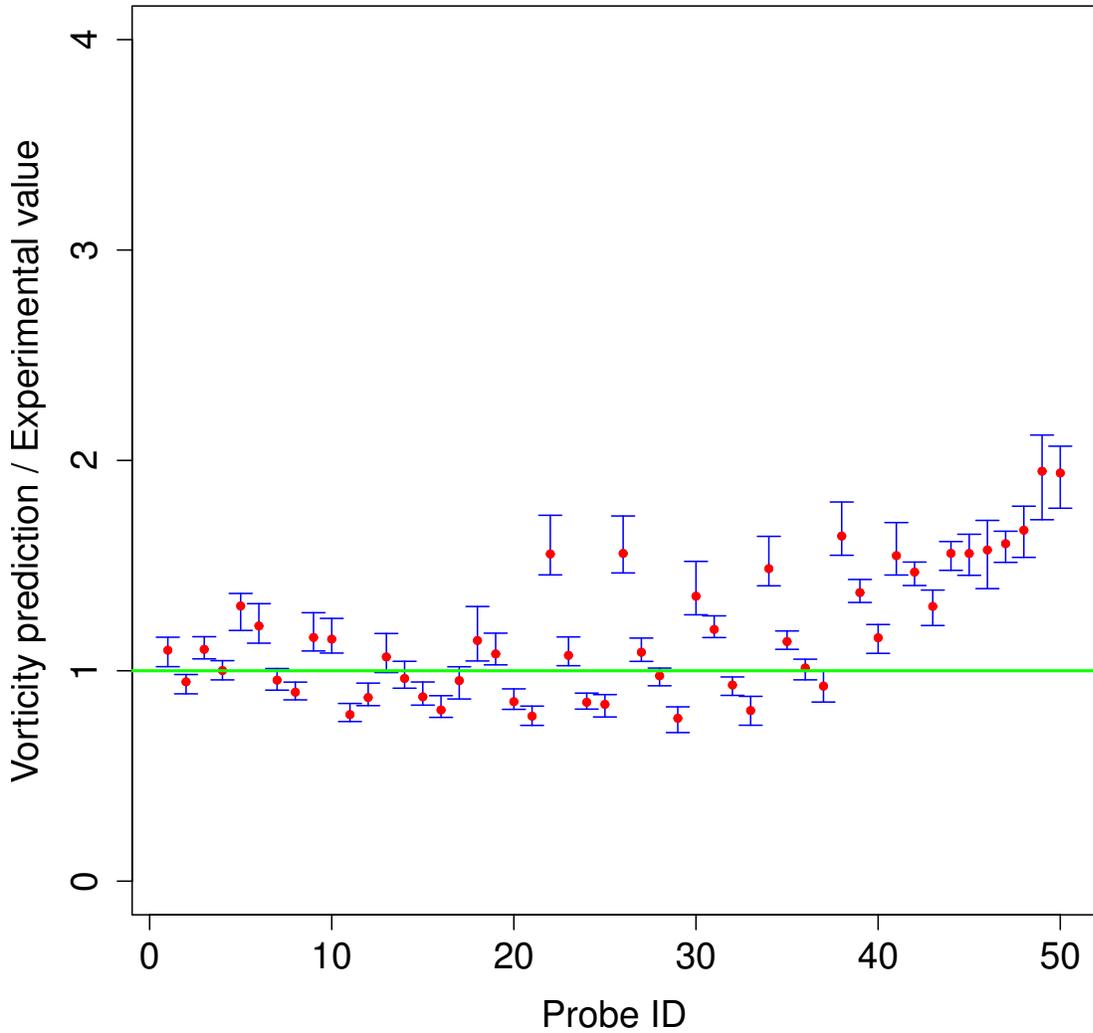
Figure 4: Pushed forward posterior test for the vorticity on the crossplane, computed using surrogate models for the probes with $+$ in Fig. 2 in the main text. The error bars span the 5^{th} - 95^{th} percentile range and the filled $\circ$ are the median prediction. Compared to the variability in Fig. 4 in the main text, the error bars are much smaller, indicating the large role played by ϵ, a composite of measurement and model-form error. Also, the error bars do not bracket the measurements (horizontal line).