

Bayesian Calibration of a RANS Model with a Complex Response Surface - A Case Study with Jet-in-Crossflow Configuration

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

We demonstrate a Bayesian method that can be used to calibrate computationally expensive 3D RANS models with complex response surfaces. Such calibrations, conditioned on experimental data, can yield turbulence model parameters as probability density functions (PDF), concisely capturing the uncertainty in the estimation. Methods such as Markov chain Monte Carlo construct the PDF by sampling, and consequently a quick-running surrogate is used instead of the RANS simulator. The surrogate can be very difficult to design if the model’s response i.e., the dependence of the calibration variable (the observable) on the parameters being estimated is complex. We show how the training data used to construct the surrogate models can also be employed to isolate a promising and physically realistic part of the parameter space, within which the response is well-behaved and easily modeled. We design a classifier, based on treed linear models, to model the “well-behaved region”. This classifier serves as a prior in a Bayesian calibration study aimed at estimating 3 $k - \epsilon$ parameters $\mathbf{C} = (C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ from experimental data of a transonic jet-in-crossflow interaction. The robustness of the calibration is investigated by checking its predictions of variables not included in the calibration data. We also check the limit of applicability of the calibration by testing at an off-calibration point.

Nomenclature

\mathbf{C}	Parameters in the $k - \epsilon$ RANS model to be calibrated
\mathbf{C}_{nom}	Nominal values of \mathbf{C}
$\mathcal{N}(\mu, \sigma^2)$	Normal distribution with mean μ and standard deviation σ
\mathcal{R}	The physically relevant part of the \mathbf{C} parameter space
J	jet-to-crossflow momentum ratio
AIC	Akaike information criterion
GP	Gaussian Process
JIC	Jet in crossflow
JPDF	Joint probability density function
LS	Learning Set
MCMC	Markov chain Monte Carlo
RANS	Reynolds-Averaged Navier-Stokes
TS	Testing Set

^{*}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

I. Introduction

RANS models¹ are routinely used in aerodynamics due to their robustness and speed vis-à-vis other simulations methods e.g., Large Eddy Simulations. Typically, they are used in conjunction with $k - \epsilon$ equations to model the evolution of the turbulent kinetic energy and its dissipation rate. The equations for k and ϵ contain empirical “constitutive laws”, e.g., the eddy viscosity model mapping strain-rates to turbulent stresses, the model for turbulent viscosity in terms of k and ϵ and empirical constants (parameters). The values of these parameters are deemed to be universal and obtained by calibrating to canonical flows e.g., channel and shear flows.^{2,3} We call these the “nominal” values of the parameters. However, there is nothing sacrosanct about these parameters’ values and there have been attempts to tune them^{2,4,5,6,7,8,9,10} for specific flows. The tuned parameters can vary quite substantially from the nominal ones.

Despite their widespread use, the predictive skill of $k - \epsilon$ simulations leave much to be desired. This is especially true for complex turbulent flows, such as transonic jet-in-crossflow interactions. The lack of predictive accuracy has two sources: (1) the structural or model-form error, due to approximations inherent in the RANS $k - \epsilon$ formulation i.e., errors due to missing physics and (2) the use of sub-optimal values of the parameters, simply picked from literature. The latter can be rectified by calibrating the RANS model with experimental data from a flow interaction that is similar to the regime in which the calibrated model will be used.

Our interest lies in devising a principled way for tuning/calibrating a RANS model for jet-in-crossflow (JIC) interactions. JIC is a canonical model for many practical flow interactions e.g., the maneuvering of launch vehicles, mixing of (injected) fuel with an oxidizer, and even the plume from a volcano interacting with the wind.¹¹ JIC interactions are poorly simulated using RANS¹² when using parameters held at their nominal values. As a first step towards improving the predictive accuracy of RANS in JIC, we hypothesized that more relevant parameter values could be obtained by calibrating to a strongly vortical flow. In our previous work,¹³ we tested this hypothesis by designing a Bayesian calibration technique to estimate 3 RANS parameters, $\mathbf{C} = (C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$, from data from an *incompressible*, flow over a square cylinder (FOSC) experiment. The parameters were estimated as a 3D joint PDF (JPDF), capturing the uncertainty in the estimates due to limited data and the inherent shortcomings of RANS. The predictions obtained by running JIC simulations using parameters sampled from the post-calibration parameter PDFs were far more accurate than the one obtained using $\mathbf{C}_{nom} = \{0.09, 1.92, 1.44\}$, the nominal values of the same parameters. The improvement in predictive accuracy was surprising since there is little obvious similarity between a transonic JIC interaction and incompressible FOSSC, beyond the strongly vortical nature of the flow. The main contributions of our previous work were (1) to identify how inadequate \mathbf{C}_{nom} were for JIC interactions, and a more relevant flow interaction that could be used to obtain better \mathbf{C} ; (2) a rigorous method to create surrogate models for Reynolds stresses in FOSSC interactions; (3) developing a joint PDF for \mathbf{C} using MCMC and surrogate models of a relatively complex flow (2D turbulent flow over a square cylinder) and (4) demonstrating the usefulness of ensemble predictions. Using parameters sampled from the post-calibration PDF, we found a parameter set that was very close (about 15% relative errors) to measurements from a JIC experiment.

While our previous work showed the inadequacy of \mathbf{C}_{nom} and a way to overcome it, it nevertheless did not address whether $k - \epsilon$ models for JIC could be improved further. The prediction errors in Ref. 13 contained contributions from the structural error as well as parametric sub-optimality - recall that the joint PDF was obtained by calibrating to FOSSC, not JIC, experimental data. In this work, we seek to isolate the impact of the structural error by calibrating to JIC experimental data. Thereafter, any deviations of RANS predictions from observations can be attributed to structural deficiencies of the $k - \epsilon$ model itself and the inaccuracies in the experimental data. We will investigate whether the structural error affects certain variables more than others. We will also investigate whether the calibration is overly sensitive to the gross flow properties, specifically the strength of the jet relative to the freestream. This is the first step in trying to improve the predictive capability of RANS, and is the first contribution of this work.

As a first step to calibrating to JIC experimental data, we attempted to create surrogate models for JIC interactions in the manner described in Ref. 13. It failed; the response of our calibration variable (vorticity in the cross-plane; described in detail in Sec. III) to variations in \mathbf{C} is too complex and non-stationary (in the statistical sense i.e., the lengthscale of variation of the response changes in the parameter space) to be captured by GP (also called kriging) surrogates or polynomials. Instead, we used experimental data to identify a promising, but irregular sub-domain in the parameter space, which had a higher potential to provide predictive RANS simulations. We call this region \mathcal{R} . We developed polynomial surrogates that are valid only inside \mathcal{R} . We use \mathcal{R} as an improper prior (for Bayesian calibration), with uniform density inside

\mathcal{R} , and zero outside. We develop a classifier, based on treed linear models, to model \mathcal{R} , and use that inside a MCMC calibration of \mathbf{C} . The use of a classifier to model the valid parameter region \mathcal{R} is borrowed from climate modeling¹⁴ but, to date, it has not been used to as a prior in any aerodynamic studies. This new way of using surrogate models, which are valid only in a sub-region of the parameter space, in a Bayesian calibration setting is the second novel contribution of this work.

The paper is structured as follows. In Sec. II, we review literature on calibration of aerodynamic models (both Bayesian and otherwise) and the use of surrogate models for calibration purposes in aerodynamics. We also review treed models. We describe relevant work in JIC simulations. In Sec. III, we formulate our inverse problem for estimating \mathbf{C} (it is a conventional one), identify the need to develop surrogate models, and provide the details of the construction. We also describe in detail the experimental studies that supply us with our calibration data. In Sec. IV, we provide results of the calibration, and check the predictive skill of the calibrated RANS model for variables omitted from the calibration’s likelihood. We also investigate the usefulness of the calibration at off-calibration points. We conclude in Sec. V.

II. Background

II.A. Jet-in-crossflow interactions

JIC is a canonical flow interaction and is used as a model for a host of practical and natural flow phenomena.¹¹ It also plays a central role in the maneuvering of certain finned bodies of revolution e.g., launch vehicles by spin rockets. In such cases, the exhaust from the spin rockets may interact with control surfaces, modifying aerodynamic forces and moments.^{15,16} The problem is strongly vortical. Our interest lies in JIC interactions where the freestream (or crossflow) is transonic, a problem that has been investigated in detail, experimentally;^{17,18,19,20} we will refer to these as the “Beresh experiments”. The jet undergoes a Kelvin-Helmholtz roll-up of the shear layer formed at the boundary of the jet and crossflow, and the flowfield is dominated by a counter-rotating vortex pair (CVP). The CVP tracks the evolution of the jet in the crossflow, and has horseshoe vortices circling it. The CVP and the horseshoe vortices are primarily responsible for modification of the flow in the vicinity of the control surfaces.

The problem of JIC, in general, has also been studied numerically; references can be found in the review paper by Mahesh.¹¹ The particular supersonic jet-in-transonic-crossflow that is the subject of this study has been investigated using $k - \omega$ models,¹² and were also compared with the Beresh experiments. The study found that all $k - \omega$ models overpredicted turbulent intensities inside the jet, resulting in simulated jets “fatter” than their experimental counterpart i.e., the turbulent diffusion was too large in the simulations. Further, the CVP, as captured on the crossplane (a plane perpendicular to the streamwise direction, that slices through the CVP) resided at a point higher than in the experiments, a result at odds with overly strong turbulent diffusion. The study conjectured that turbulent stress were underpredicted in the nearfield of jet, resulting in an erroneous exchange of momentum between the jet and crossflow. The study investigated a number of jet-to-crossflow momentum ratios (J) and jet cant angle, but the general behavior did not change i.e., the lack of accuracy was due to a fundamental deficiency in $k - \omega$ models and was not a consequence of the particular flow configuration. Large Eddy Simulations^{21,22,23} (LES) and Detached Eddy Simulations²⁴ (DES) of the same JIC configuration have also been performed and compared to the Beresh experiments; the mean-flow on the midplane was used as a figure of merit. No comparisons were performed on the crossplane. As expected, their agreement with experimental observations were far better than RANS. While this is encouraging, DES and LES are too computationally intensive to be used in routine design calculations. Thus the need to analyse and reduce RANS prediction errors retains its importance.

II.B. Improving RANS

The accuracy that LES and DES simulations can achieve have allowed them to be used as “numerical experiments” to quantify RANS errors as well as to calibrate RANS models. Note that such studies are limited to (simple) flow configurations where LES or DNS are actually feasible. The motivation of these studies have been to either quantify the uncertainty in RANS predictions or to actually improve RANS via calibration. Both types of studies require one to augment the empirical models in RANS for the creation and dissipation of turbulent kinetic energy with a “correction” term, which is estimated from LES or DNS (Direct Numerical Simulation) data. The “correction”, or structural error is deemed to be spatially variable and is modeled as such. The estimation of the structural error has been investigated in two ways. In the

first case (typically geared towards uncertainty quantification of RANS predictions), the parameters of the RANS model are deemed accurate, and the discrepancy between RANS and LES (or DNS) is computed and its spatial variations modeled in various ways. In Ref. 25, turbulence variables from 2D RANS and LES simulations of transonic flows over a bump were compared to estimate a spatially variable error. The error was then modeled and a random term was introduced to evaluate its impact on flow variables. The model of the error so constructed was tested by performing an ensemble of RANS simulations and adding realizations of the error field by drawing from the PDF. Recently, this approach has been extended to 3D JIC computations.²⁶ The second approach is, quite simply, the calibration of RANS models using LES/DNS data. Such studies have been performed for simple flows e.g., channel flows, flow over a flat plate etc.,^{27,28,29} and involve estimating both flow turbulence model parameters via inverse modeling as well as fitting a structural error model that captures the inadequacy of the calibrated model. The method is very data-intensive i.e., in order to estimate the structural error, the “good” DNS/LES results have to be made available at each point of the RANS mesh. In Ref. 28, parameter estimates and structural error is estimated for the Spalart-Allmaras model using a Bayesian approach, with observational data from incompressible boundary layer flows over a flat plate under both favorable and adverse pressure gradients. The study by Edeling *et al.*³⁰ addressed the same problem, but used a $k - \epsilon$ model instead.

A subsequent study by Edeling *et al.*³¹ addressed the question whether turbulent flow parameters developed for a certain pressure gradient could simply be transferred to a different flow configuration and be predictive there. They were not. They hypothesized that RANS models, once calibrated, could only be used in “similar” flow regimes. They demonstrated this hypothesis in the context of incompressible flat-plate boundary layer flows, at different pressure gradients (called scenarios). They chose a set of scenarios (the training scenarios) and calibrated a collection of models ($k - \epsilon$, $k - \omega$, Spalart-Allmaras, etc.) to estimate both the models’ parameters and their probability i.e., the relative probability of a model’s ability to reproduce observations vis-à-vis others in the collection. These models were subjected to Bayesian Model Averaging over a number of pressure gradients (to weigh models appropriately) and used to predict both the calibration variable (scaled streamwise velocity) and the skin-friction coefficient at pressure gradients that were not included in the training scenarios. Much of the mathematical development in their study involves the choosing and weighing of scenarios that are used for training/calibrating the models.

The studies described above i.e., Refs. 28,30,31 used MCMC to perform the calibration and all inferences were probabilistic. The simplicity of the flow configuration (incompressible, flat-plate boundary layer) was instrumental in reducing computational costs and making the phenomenally expensive MCMC method feasible for calibration purposes. The use of MCMC to calibrate models for more complex flows requires one to replace the flow simulator with a surrogate. In Ref. 32 the authors constructed surrogate models, using kriging, for the velocities and temperatures predicted by LES in the crossplane of a hot jet-in-crossflow interaction. They used the surrogates to estimate optimal values of the constant in the Smagorinsky model and a fourth-order dissipation constant using JIC data obtained from a DNS simulation. In Ref. 13 the authors constructed polynomial surrogate models for the Reynolds stress in the wake of an incompressible 2D flow over a square cylinder (as simulated using RANS) as a function of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ and used Bayesian calibration to estimate a JPDF for them using experimental data. The PDFs were then tested for their predictive skill in a transonic JIC interaction.

II.C. Treed models

The use of simple GP (kriging) and polynomials to construct surrogate models assumes that the mapping between the calibration variable(s) (observables) and the parameters being calibrated is simple and smooth. This assumption often holds true only in a small portion of the parameter space. In most cases, the parameter space is defined using upper and lower bounds on each parameter, and a random combination of parameters may not even be physical i.e., the RANS or LES model may not even run successfully to completion. Further, non-physical parameter combinations do not provide model predictions that are close to experimental observations. In such cases, one could, pragmatically, excise the non-physical portion of the parameter space and focus on constructing a surrogate model that is valid only in \mathcal{R} , the valid, physically realistic part of the parameter space. This approach was followed in Ref. 14, where the authors used a random forest classifier to discriminate between the physical and non-physical parts of the parameter space of a climate model. The classifier was trained on a training set of 10,000 climate model runs.

A second way of constructing a separatrix between physical and non-physical parts of a parameter space is to define a function $\zeta(\mathbf{C})$ that assumes a value of ± 1 in the physical and non-physical parts of the

parameter space respectively. Then the level-set $\zeta(\mathbf{C}) = 0$ defines the separatrix. The problem then reduces to approximating $\zeta(\mathbf{C})$ using a training set of points in the parameter space where the binary function is defined.

Complex, non-stationary functions can be approximated, from a training set of data, using Classification and Regression Trees³³ (CART). CART models are particularly simple if the parameter space is a hyper-cuboid. CART models are built under the assumption that the space of predictors \mathbf{C} can be divided into smaller subdomains where a piecewise constant model would provide an adequate representation of the available data. Subdomains are created by recursively splitting along a dimension and computing the means and variances in the newly created subdomains; an improvement of fit (reduction in variance) results in the split being deemed acceptable. Depending on the function approximated by the training data, very different partitions of the subdomain may result in CART models of very similar accuracies i.e., the partitioning is non-unique. The initial domain is assumed to be the root node in a tree; thereafter, the splits lead to the subdomains becoming the children nodes in a binary tree. The binary tree can change by further splitting of nodes, coalescing of a parent-children pair (provided the child is a leaf node in the binary tree) and by randomly changing the splitting rule of a node. Due to the non-unique nature of partitionings, stochastic search methods are often used to discover multiple trees/CART models that fit the data.

A variation of CART is the treed linear model, where the function in a subdomain is approximated by a linear model (Treed linear models; Ref. 34) or by a kriged model (treed Gaussian Process models; Ref. 35), rather than a constant. Due to the recursive splitting of the domain, complex, non-stationary variations of $\zeta(\mathbf{C})$, including discontinuities, can be easily accommodated. In Ref. 35, the authors modeled lift generated by a rocket booster as a function of Mach number, angle of attack and a sideslip angle. Discontinuous changes in lift in the transonic regime and increase in the angle of attack (due to formation of shocks on aerofoil surfaces) were captured by the recursive nature of subdividing the parameter space. Thus a treed model can serve as a classifier since it can model $\zeta(\mathbf{C})$ and identify which side of a separatrix an arbitrary \mathbf{C} lies on. We will use this approach to model \mathcal{R} , the physically relevant portion of the parameter space.

III. Formulation of the calibration problem

In this section, we formulate a Bayesian inverse problem for estimating $\mathbf{C} = \{C_\mu, C_{\epsilon 2}, C_{\epsilon 1}\}$ from experimental measurements. We identify the need for surrogate modeling and also develop an informative prior.

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 structural error (the inherent inability of the model to reproduce observations due to approximations). 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. If the measurement errors are low, σ^2 provides an estimate of the structural 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

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

where $\|\cdot\|_2$ is the ℓ^2 -norm of the vector. By Bayes’ theorem, the calibrated distribution (or posterior distribution) of (\mathbf{C}, σ^2) can be given as

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

This inverse problem can be solved via sampling i.e., we use a method such as MCMC³⁶ to draw samples of $\{\mathbf{C}, \sigma^2\}$. We reconstruct $P(\mathbf{C}, \sigma^2 | \mathbf{y}_e)$ empirically by plotting histograms or via kernel density estimation.³⁷ In this particular work, we use an adaptive MCMC method called Delayed Rejection Adaptive Metropolis (DRAM; Ref. 38), as implemented in the R³⁹ package FME.⁴⁰ The convergence of the MCMC chain to a stationary distribution $P(\mathbf{C}, \sigma^2 | \mathbf{y}_e)$ i.e., the determination of the sufficiency of $\{\mathbf{C}, \sigma^2\}$ samples to construct

$P(\mathbf{C}, \sigma^2 | \mathbf{y}_e)$ is performed using the Raftery-Lewis method⁴¹ implemented in the R package `mcmcgbssit`.⁴² 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 $\sigma^2 > 5$.

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 variable (described below) on \mathbf{C} . This mapping is expected to hold true (i.e., within an acceptable level of error) in the support of $\Pi_1(\mathbf{C})$. The surrogate model will serve as a computationally inexpensive proxy for SIGMA, our compressible 3D RANS simulator. SIGMA implements a $k - \epsilon$ model (as described in Ref. 43) and uses linear model for the turbulent viscosity. The model uses a turbulent Reynolds number based (y^+ independent) damping functions at the walls making it suitable for large scale parallel computations of complex flows. A Roe-TVD flux scheme with a min-mod limiter for discontinuity capturing has been used for the spatial discretization. Time integration is carried out using a first order point-implicit scheme. The calculations presented herein are initialized using a first-order spatial scheme and relaxed to convergence using the second order scheme. Time marching was carried out in steady state mode (local time stepping) with a gradual CFL ramp to accelerate convergence.

We consider the following bounds on \mathbf{C} : $0.06 \leq C_\mu \leq 0.12, 1.7 \leq C_{\epsilon 2} \leq 2.1, 1.2 \leq C_{\epsilon 1} \leq 1.7$, taken from Ref. 2. While it is tempting to combine these bounds into a uniform distribution in the cuboid \mathcal{C} in $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ space, we shall refrain from doing so since random samples drawn from it may not be physically relevant. While the experiment revealed that the flow is turbulent, but steady, many $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ combinations drawn from \mathcal{C} lead to RANS computations that do not converge to steady state; other combinations lead to converged solutions that bear little resemblance to high Reynolds number transonic flowfields (and consequently the discrepancy $\mathbf{y}_e - \mathbf{y}_m(\mathbf{C})$ is very large). The inability to generate reasonable model predictions \mathbf{y}_m makes it very difficult to construct surrogate models that are accurate proxies for RANS predictions everywhere in \mathcal{C} . Consequently, to circumvent this problem we will develop a prior density that restricts the parameters to a region \mathcal{R} , within which the parameters produce physically realistic flowfields.

III.A. Experimental data

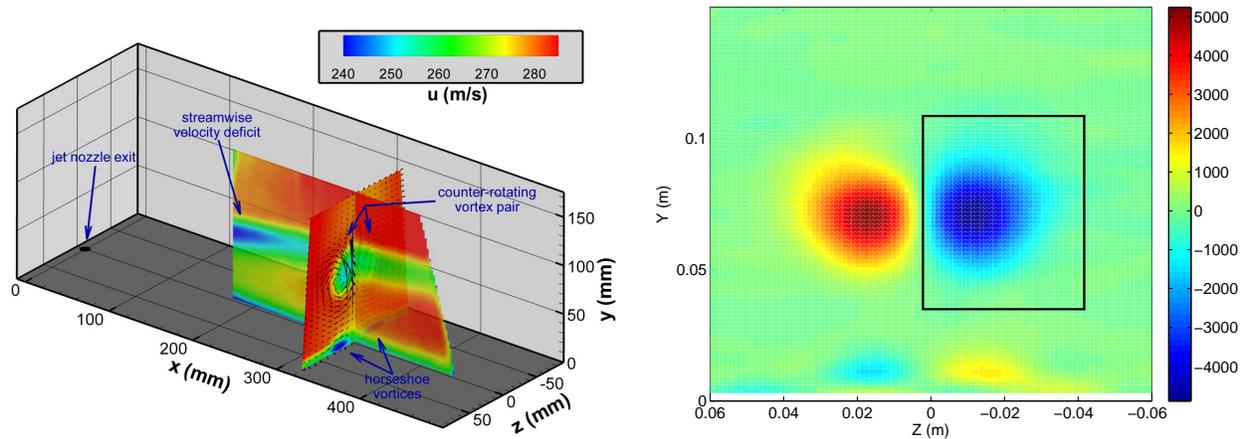


Figure 1. Left: Schematic of the test section showing the orifice where the jet is introduced and the mid- and crossplane where experimental measurements are made. Right: Plot of the experimental streamwise vorticity and the window W ($0 \leq z \leq 0.04, 0.031 \leq y \leq 0.11$) containing one of the vortices. Calibration and validation based on vorticity will be performed inside this box.

The wind tunnel experiments from which we obtain our calibration and validation data are described in Ref. 17, 18, 19. The test section simulated using SIGMA (see schematic in Fig. 1 (left)) has a square cross section, 304.8 mm per side. Turbulent flow is introduced into the test section from the left. A jet, 9.53 mm in diameter and perpendicular to the inflow, is introduced at the floor of the test section. PIV measurements of the flowfield are conducted in the midplane (the plane of symmetry in JIC interaction) as well as a crossplane (perpendicular to the jet), 321.8 mm (33.8 jet diameters) downstream from the center of the jet. Measured

quantities include velocities and the fluctuating turbulent velocities, from which turbulent stresses ($\overline{u'_i u'_j}$) are computed. The studies track the penetration of the jet into the crossflow. This tracking is performed by plotting streamwise and vertical velocity profiles at 5 locations, starting at a distance 200 mm downstream of the center of the jet and thereafter, 50 mm apart. Velocity and turbulent stress distributions in the crossplane are used to demarcate the CVP. Most of the measurements are for $M = 0.8$ freestream flow, with $J = 10.2$. Both midplane and crossplane measurements are also available for $J = 5.6, 16.7$. Measurements are restricted to the midplane for $M = 0.6, 0.7$.

In this study we will use the $M = 0.8, J = 10.2$ experiment as our calibration case. In particular, we will use the experimental velocity measurements on the crossplane to compute the streamwise vorticity. Fig. 1 (right) shows a plot of the experimental crossplane vorticity field, clearly showing the CVP. The window \mathcal{W} ($0 \leq z \leq 0.04$ m, 0.031 m $\leq y \leq 0.11$ m), surrounding one of the vortices, demarcates the region where the vortices are observed in the RANS solution as \mathbf{C} is varied. The window also eliminates the boundary layer seen at the bottom of the figure. Superimposing the RANS grid on the crossplane reveals that the window is covered by a 8×28 grid. We refer to these 224 grid cells as the set \mathcal{P}_f of “probes” where both experimental and model predictions of streamwise vorticity are available for comparison and matching during calibration.

III.B. Designing an informative prior

In this section we devise a way of selecting \mathcal{R} , the part of \mathcal{C} containing $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ values that produce physically realistic flowfields. We draw 2744 (14^3) samples of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ from \mathcal{C} using a space-filling, quasi Monte Carlo method (Halton sequence) and use them to seed 3D RANS JIC simulations. Each simulation (approximately 10 hours on 1024 cores of a PowerPC A2 processor) results in a streamwise vorticity prediction on the crossplane. 100 runs failed and a number of the remaining 2644 runs resulted in non-physical flowfields. On the whole, the variation of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ led to CVPs of varying sizes and strengths; they were also seen to vary in location (above or below the experimentally observed CVP position).

In order to identify \mathcal{R} , we compute the root mean square error (RMSE) between the streamwise vorticity produced by each of the simulations and the experimental counterpart on the crossplane. We retain the $N_s = 525$ runs whose RMSE lie below the 20th percentile. The $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ for these selected runs, which identify \mathcal{R} , are plotted in Fig. 2. The 3D scatterplot clearly show that \mathcal{R} excludes a large portion of \mathcal{C} . Our next step lies in using the training set of N_s $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ points to devise a function in \mathcal{C} that can be used to demarcate \mathcal{R} .

We define a function $\zeta(\mathbf{C})$ in \mathcal{C} that is set to 1 at the N_s points that define \mathcal{R} and -1 at the remaining $2744 - 525 = 2219$ $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ combinations that we had drawn from \mathcal{C} . We define the level set $\zeta(\mathbf{C}) = 0$ as $\partial\mathcal{R}$, the demarcation surface for \mathcal{R} . The technical challenge then reduces to finding a model $\zeta_m(\mathbf{C})$ that approximates $\zeta(\mathbf{C})$ at any arbitrary \mathbf{C} . We use a Bayesian treed linear model³⁴ to construct $\zeta_m(\mathbf{C})$. We use a MCMC method to stochastically generate new trees, followed by newly fitted linear models in the leaf nodes. It allows us to generate an ensemble of trees (and associated linear models), which are then all used to generate a mean prediction for $\zeta_m(\mathbf{C})$, given a specific \mathbf{C} . We use the implementation of Bayesian treed linear models in the R package `tgp`.⁴⁴ In Fig. 3, we plot 2D projections of the most probable tree of subdomains, along with the N_s points in \mathcal{C} where $\zeta(\mathbf{C}) = 1$. $\zeta(\mathbf{C}) = -1$ points are not plotted as there are too many of them; however, they are uniformly distributed in the “empty” space in the figure. It is clear that there are subdomains which predominantly contain $\zeta(\mathbf{C}) = 1$ (or -1); a few contain a mixture. It is in these few nodes that linear models will be useful in predicting $\zeta_m(\mathbf{C})$ at an arbitrary \mathbf{C} ; a positive value of $\zeta_m(\mathbf{C})$ indicates that \mathbf{C} lies inside \mathcal{R} . Thus the Bayesian treed linear model acts as a classifier.

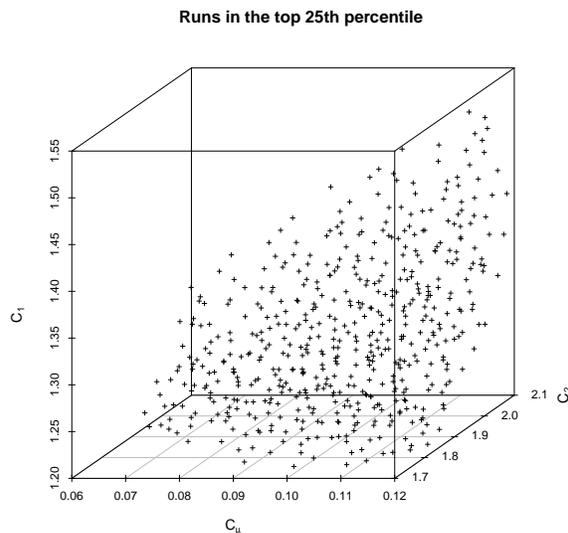


Figure 2. Scatterplot of the set of N_s points in the $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ space that constitute \mathcal{R} .

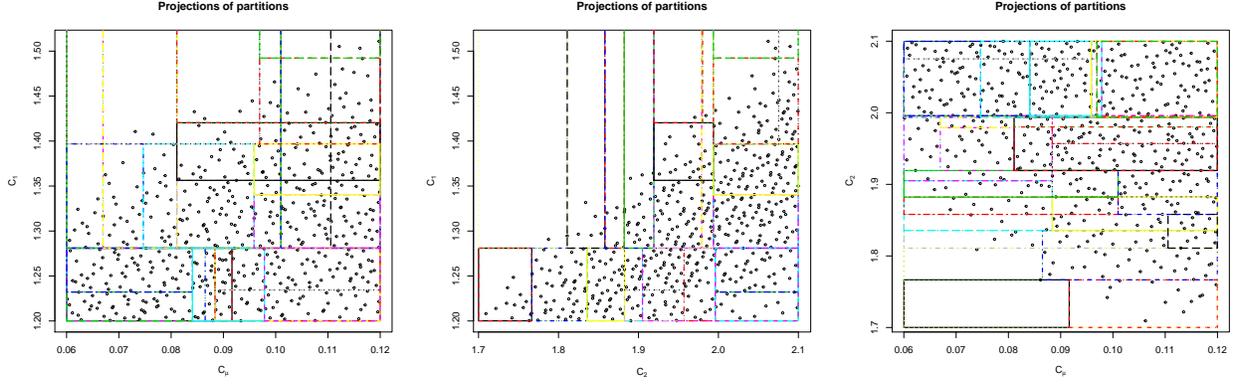


Figure 3. The highest probability or maximum a posteriori partitioning tree of the N_s points in \mathbf{C} that define \mathcal{R} . The points that lie outside \mathcal{R} are used in the calculation of the tree, but are not plotted here because there are too many of them for a clear plot. They inhabit the blank regions of the plots above and are uniformly distributed there. The boxes are the nodes of the tree projected on 2D slices. Data in leaf nodes are approximated with a linear model.

We test the accuracy of $\zeta_m(\mathbf{C})$ using repeated random subsampling validation (a form of cross-validation). Results were averaged over 20 rounds (or splits). We randomly divide the 2744 samples of $\zeta(\mathbf{C})$ into a Learning Set (LS) with 2332 samples (85% of the total) and a Testing Set (TS) with the rest. A treed linear model (classifier) is built using the LS and used to classify/predict at the $(C_\mu, C_{\epsilon_2}, C_{\epsilon_1})$ combinations in the TS. The misclassification rate is then computed. This is repeated 20 times, with different LS/TS pairs, and the mean misclassification rate is found to be around 4%. Thus we find the resulting classifier to be quite accurate. We will use this classifier to define $\partial\mathcal{R}$, the boundary of \mathcal{R} . We model the prior density within \mathcal{R} as a uniform distribution i.e., all $(C_\mu, C_{\epsilon_2}, C_{\epsilon_1})$ combination are equally probable. Outside \mathcal{R} , the prior density is zero.

III.C. Polynomial surrogate models

Due to the large number of model evaluations that are needed for an MCMC method to converge to a stationary posterior distribution, our surrogate models are meant to serve as computationally inexpensive proxies for 3D RANS predictions of streamwise vorticity on the crossplane. Let $y^{(p)}(\mathbf{C})$ be the streamwise vorticity on the crossplane, predicted by RANS at probe $p, p \in \mathcal{P}_f$. We seek to construct an approximation $y_s^{(p)}(\mathbf{C})$ using multivariate polynomials involving $\mathbf{C}, \mathbf{C} \in \mathcal{R}$. Since $y^{(p)}(\mathbf{C})$ is obtained by taking the derivative of the velocity field, it tends to be contaminated with some numerical noise. For probes p in or near the center of the CVP, the large magnitude of the vorticity allows its dependence on \mathbf{C} to be discovered and modeled. The rest of the probes, where numerical noise is substantial, have to be eliminated.

For a given RANS run, we compute the 75th percentile of the vorticity magnitudes obtained at probes $p, p \in \mathcal{P}_f$ and retain the $224/4 = 56$ probes with absolute vorticities above this threshold. The vorticity predicted at these probes are expected to be large enough to be minimally affected by numerical noise. This is repeated for all $\mathbf{C}, \mathbf{C} \in \mathcal{R}$; a different set of probes are selected for each run since the CVP changes its size and location as \mathbf{C} varies. We take a union of these probes to create a set \mathcal{P}_* , containing 108 probes, where we will attempt to create surrogate models. The probes in \mathcal{P}_* are plotted on top of the experimental vorticity field in Fig. 4 (left). Note that \mathcal{P}_* is not the set of probes that we will use in the calibration; as we will see below, we will be able to create surrogate models, of an acceptable quality, only for a subset of them.

The method to create and test the robustness of surrogate models is the same as the one described in Ref. 13. At any probe $p, p \in \mathcal{P}_*$, we postulate that

$$\begin{aligned}
 y^{(p)} &= \sum_{l=0}^{l=3} \sum_{m=0}^{m=3} \sum_{n=0}^{n=3} \alpha_{l,m,n}^{(p)} (C_\mu)^l (C_{\epsilon_2})^m (C_{\epsilon_1})^n + d, \\
 &= y_s^{(p)} + d, \\
 3 &\geq l + m + n,
 \end{aligned} \tag{2}$$

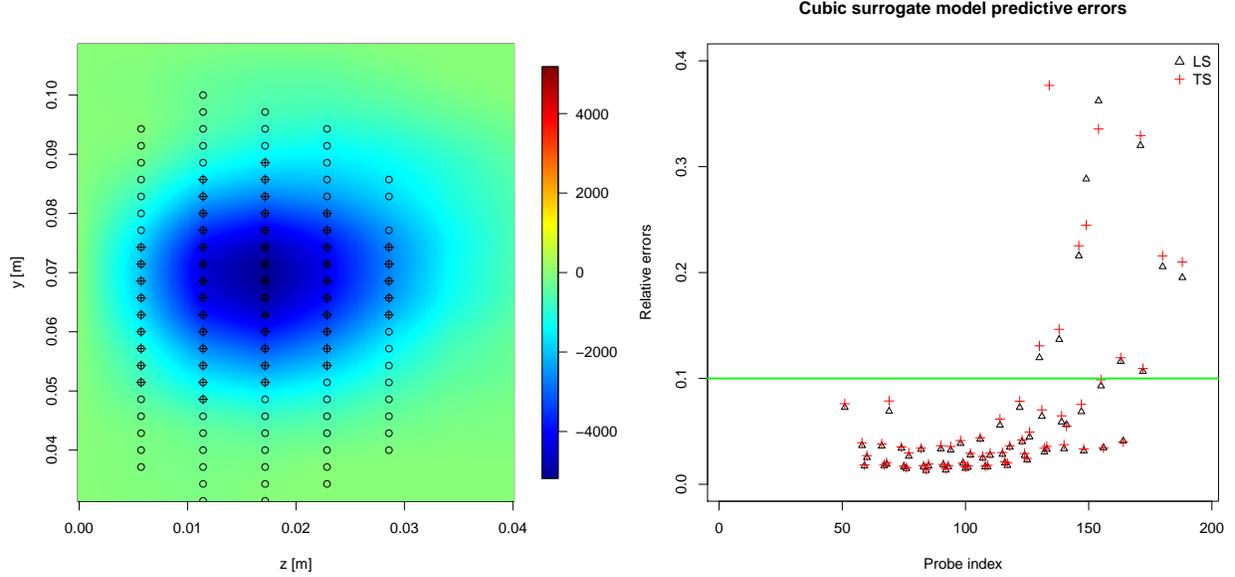


Figure 4. Left: Plot of experimental vorticity inside \mathcal{W} with the probes constituting \mathcal{P}_* overlaid on it and marked with \circ . Right: Comparison of $\overline{\delta_{LS}^{(p)}}$ and $\overline{\delta_{TS}^{(p)}}$ for all 108 probes in \mathcal{P}_* . We note that the two relative errors are of the same magnitudes, indicating that there is little overfitting of the surrogates. We also see that for some probes, the relative errors can be big i.e., surrogates are not necessarily accurate for all probes. Some of the errors also lie outside the range of the vertical axis and are not plotted. The green line at 10% error demarcates the probes that can be used in calibration; the rest of the polynomial surrogates are too inaccurate for any practical use. The probes which lie under the green line are plotted on the right with + symbols.

where $y_s^{(p)}$ is the polynomial approximation (surrogate) of the 3D RANS prediction of the crossplane stream-wise vorticity $y^{(p)}$ and d is the approximation error. In order to estimate the coefficients $\alpha_{l,m,n}^{(p)}$, we create a dataset $\{C_\mu, C_{\epsilon 2}, C_{\epsilon 1}, y^{(p)}\}_i, \mathbf{C}_i \in \mathcal{R}$, and fit the model in Eq. 2 via least-squares fitting. The cubic model so formed is rarely useful in practice due to its tendency to overfit the $\{C_\mu, C_{\epsilon 2}, C_{\epsilon 1}, y^{(p)}\}$ data. We simplify the model by incrementally removing high-order terms and refitting the shrunk model to the same data. The (original) cubic model and its shrunk counterpart are competed using AIC; if the shrunk model results in a lower AIC, it is retained and subjected to the same incremental simplification process. We find that simplification removes terms from Eq. 2 for all probes in \mathcal{P}_* , and in some cases reduces the cubic model to quadratic.

We next test whether the AIC-based simplification procedure yields robust surrogate models. This is done via repeated random subsampling validation, and results are averaged over 100 rounds. The N_s RANS runs are separated into a LS containing 446 (approximately 85% of the runs) randomly selected runs, while the remainder constitute the TS. The simplified polynomial model is constructed using the LS and the relative error for each parameter combination i in the LS, $\delta_{i,LS}^{(p)} = (y_i^{(p)} - y_{s,i}^{(p)})/y_i^{(p)}$, $i \in LS$, evaluated. Here $y_i^{(p)}$ refers to RANS predictions computed using the parameters $\mathbf{C}_i, i \in LS$. The fitted model is also used to evaluate $y_{s,j}^{(p)}$ for the parameter set $j, j \in TS$, and calculate the corresponding relative error i.e., $\delta_{j,TS}^{(p)} = (y_j^{(p)} - y_{s,j}^{(p)})/y_j^{(p)}$, $j \in TS$. These individual relative errors are summarized by their RMS (root mean square) value for the LS and TS. This process is repeated 100 times, using different LS/TS pairs; the mean of the RMS relative errors, $\overline{\delta_{LS}^{(p)}}$ and $\overline{\delta_{TS}^{(p)}}$, are taken as measures of accuracy of the polynomial surrogate. The process is repeated for all the 108 probes in \mathcal{P}_* .

In Fig. 4 (right) we plot $\overline{\delta_{LS}^{(p)}}$ and $\overline{\delta_{TS}^{(p)}}$ for all the probes. We notice that for about half the probes, the two relative errors are very close i.e., the polynomial model fitted to the LS data is equally predictive of the TS. This indicates that the polynomial model does not overfit the LS data; had this not been the case, $\overline{\delta_{LS}^{(p)}}$ would have been substantially smaller than $\overline{\delta_{TS}^{(p)}}$. We also notice that for certain probes the relative error is high i.e., the polynomial model is not an accurate representation of $y^{(p)}$, and should not be used in calibration. Using 10% error (green line in Fig. 4 (right)) as the criterion, we retain 52 out of 108 probes in \mathcal{P}_* where

the polynomial model is deemed to be sufficiently accurate. Further, for these selected probes, the LS and TS relative errors are very close. We refer to this set of 52 probes as \mathcal{P} ; since we can construct accurate surrogate models for all probes $p, p \in \mathcal{P}$, we will use them for calibration purposes. The probes in \mathcal{P} are plotted in Fig. 4 (left) using + signs. They occupy the center of the experimental vorticity field with large vorticity magnitudes. We repeated the same surrogate model construction, but with the models limited to second-order terms rather than cubic terms; we could construct accurate surrogates at on 32 / 108 probes in \mathcal{P}_* . This provides a second verification that the surrogate models do not overfit the RANS data and thus provide an accurate proxy for the 3D RANS simulator.

III.D. The inverse problem

We will solve a slight variation of Eq. 1 to estimate \mathbf{C} from the Beresh experiments

$$P(\mathbf{C}, \sigma^2 | \mathbf{y}_e) \propto \mathcal{L}(\mathbf{y}_e | \mathbf{C}, \sigma^2) \Pi_1(\mathbf{C}) \Pi_2(\sigma^2) \propto \exp\left(-\frac{\|\mathbf{y}_e - \mathbf{y}_s(\mathbf{C})\|_2^2}{\sigma^2}\right) \Pi_1(\mathbf{C}) \Pi_2(\sigma^2), \quad (3)$$

the modification being the substitution of crossplane streamwise vorticity predictions via a 3D RANS model, \mathbf{y}_m , by predictions via surrogates $\mathbf{y}_s = \{y_s^{(p)}\}, p \in \mathcal{P}$. We will use the inverse Gamma prior described above for σ^2 . The prior for $\mathbf{C}, \Pi_1(\mathbf{C})$, is a uniform distribution described over \mathcal{R} . It is defined as an improper prior i.e., $\Pi(\mathbf{C}) = 1, \mathbf{C} \in \mathcal{R}$ and zero otherwise. It is implemented as a classifier, which in turn makes use of Bayesian treed linear models to demarcate \mathcal{R} inside \mathcal{C} . We will estimate the posterior distribution (Eq. 3) using the DRAM algorithm implemented in FME. Post-calibration, we will draw samples of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ from $P(\mathbf{C}, \sigma^2 | \mathbf{y}_e)$ and run RANS simulations to generate an ensemble of crossplane vorticity predictions. Statistical summaries of the vorticity field lying inside the window \mathcal{W} will be used for validation purposes i.e., to check the predictive skill of the of the calibration at $J = 16.7$ when the Mach number is held fixed. We will also check the degree of accuracy with which we can predict the time-averaged (mean) flow on the midplane.

IV. Results

The Bayesian inverse problem defined in Eq. 3 is solved using DRAM and tested for convergence using the Raftery-Lewis method. About 25,000 MCMC steps are required to reach a stationary posterior distribution. The posterior is marginalized and the 1D PDFs for the individual parameters are plotted in Fig. 5; the nominal values are plotted using vertical lines. It is clear that the maximum a posteriori (MAP) values of the parameters vary significantly from the nominal ones, especially for $C_{\epsilon 2}$. An estimate of the model - data mismatch, in the form of σ is also obtained. Comparing σ with the vorticity plotted in Fig. 4 (left), we see that the misfit is quite large. We also performed a calibration, using the same formulation, but using a genetic algorithm; the convergence of the algorithm is in Ref. 45. The best/converged value of the parameters are $C_\mu = 0.105, C_{\epsilon 2} = 2.099, C_{\epsilon 1} = 1.42$. These reflect the values of these parameters that are found to be most probable in Fig. 5.

We next examine the calibrated model's predictive skill at a different value of J . We use the 100 $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1}, \sigma^2)$ samples from the posterior distribution to generate vorticity fields on the crossplane by running the 3D RANS simulator. The vorticity field is then approximated by a point vortex, whose circulation is obtained by integrating the vorticity in \mathcal{W} . The centroid of the vorticity distribution in \mathcal{W}

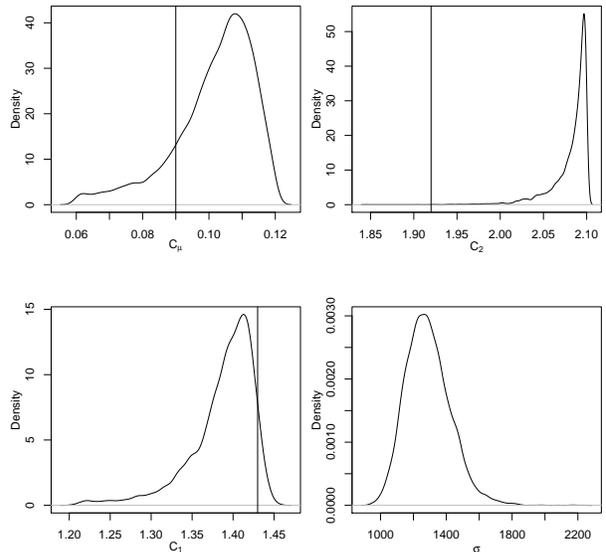


Figure 5. Marginalized PDFs for $C_\mu, C_{\epsilon 2}, C_{\epsilon 1}, \sigma$ computed as a solution to Eq. 3. The nominal values of the parameters are shown using the vertical line. The MAP estimates of all parameters except $C_{\epsilon 1}$ vary significantly from their nominal value.

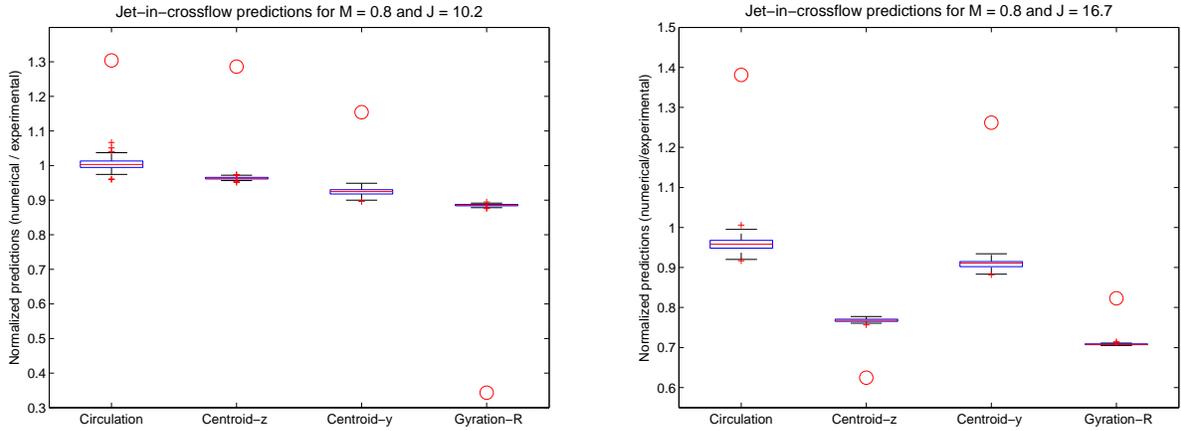


Figure 6. Left: Distribution of circulations, the centroids and radii of gyration generated from 100 samples from the posterior distribution, normalized by their experimental counterparts. Predictions using C_{nom} are plotted with \circ . The results correspond to the $M = 0.8, J = 10.2$ case. Right: Distributions for the $M = 0.8, J = 16.7$ case.

is assigned the position (y, z) of the point vortex. In addition, we compute the radius of gyration of the vorticity distribution inside the \mathcal{W} as a simple proxy for its size. We will refer to these metrics as the “point vortex metrics”. In Fig. 6 (left), we boxplot the “point vortex metrics”, normalized by their experimental counterparts for the $M = 0.8, J = 10.2$ case. The individual circles plot the predictions obtained using C_{nom} . We see that for circulation and the position, our errors are well within 5%; this holds true even for the interquartile range. The radius of gyration is underpredicted by about 10%. Such an accurate representation is somewhat expected since the vorticity field was used to calibrate the model. In Fig. 6 (right), we plot the results for the $M = 0.8, J = 16.7$ case. Note that the model was *not* calibrated for this case and the plot measures the predictive skill of the model at an off-calibration point. We see that the circulation and height of the centroid are accurately captured for the $J = 16.7$ case, but 20% - 30% errors are seen in z and the radius of gyration. The improvement of the predictions, post-calibration, are stark.

The set of runs plotted in Fig. 6, for the $M = 0.8, J = 10.2$ case, allow us to identify an optimal parameter combination (from the 100 points sampled from the posterior). The deviations (from 1) of the ratios plotted in Fig. 6 allow us to define an error (bounded by ± 1) and consequently a mean square error that combines the deviations in circulation, centroidal position and the radius of gyration. The run with the minimum mean squared error yields the best $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ combination $\mathbf{C}_{opt} = \{0.1025, 2.099, 1.416\}$. Note that this value is not very different from the optimal values yielded by the genetic algorithm optimization ($C_\mu = 0.105, C_{\epsilon 2} = 2.099, C_{\epsilon 1} = 1.42$). In Fig. 7, we plot, in the left column, the streamwise vorticity distribution computed on the right half of the crossplane using C_{nom} . In the right column, we plot

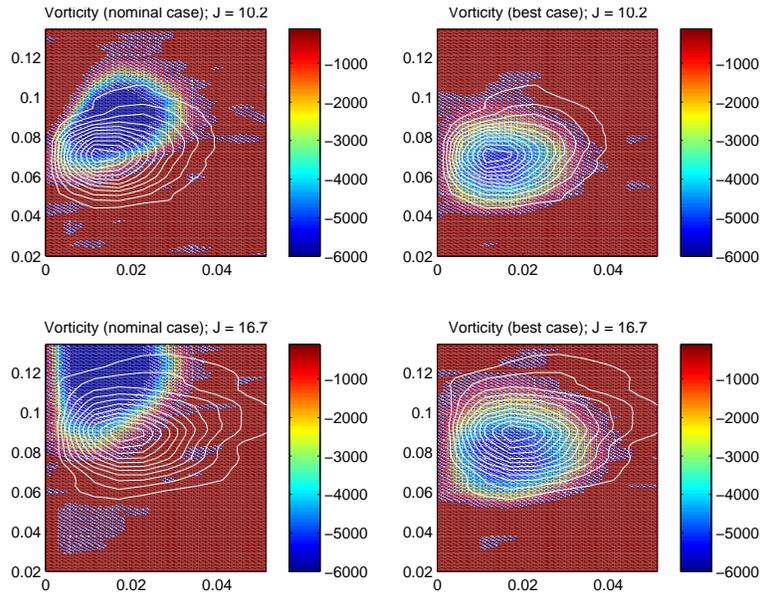


Figure 7. Left column: Streamwise vorticity in the crossplane computed using C_{nom} , overlaid with vorticity, as measured in experiments. $J = 10.2$ and 16.7 for the top and bottom rows. Right column: Crossplane streamwise vorticity predicted using the “optimal” values of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1}), \mathbf{C}_{opt}$. Qualitatively, the agreement is quite good. The improvement over C_{nom} (left column) is stark.

the corresponding vorticity distribution obtained using \mathbf{C}_{opt} . Overlaid on each are contour plots of the vorticity distribution obtained experimentally. The top row contains results for the $M = 0.8, J = 10.2$ case; $J = 16.7$ is plotted in the the bottom row. In general, the vorticity distribution obtained using \mathbf{C}_{opt} (right column) is very close to the experimental values. The improvement over the predictions with \mathbf{C}_{nom} (left column) is substantial. These figures bear out the summaries in Fig. 6, which showed that calibration to the $M = 0.8, J = 10.2$ case improved the predictive skill of RANS at an off-calibration value of J .

In Fig. 8 (top), we plot the streamwise velocity deficit in the midplane at $x/D_j = 21$ for $M = 0.8, J = 10.2$ and 16.7. In Fig. 8 (bottom) we plot the vertical velocity on the midplane at the two flow conditions. The data for the plots were obtained using \mathbf{C}_{nom} , \mathbf{C}_{opt} and the ensemble mean (from the 100 $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ samples drawn from the posterior). These are compared with the experimental values (\circ). We see that the ensemble mean and \mathbf{C}_{opt} predictions always lie closer to the experimental value than \mathbf{C}_{nom} ; in the case of $M = 0.8, J = 10.2$, the agreement between the predicted and experimental values is very close. Predictions using \mathbf{C}_{nom} bear out the inferences drawn from Fig. 6 - the CVP sits far above the experimentally observed CVP, and is much too strong, at both values of J . Post-calibration, the position as well as the strength of the CVP is closer to the experimental values. The vertical velocity in the midplane is largely governed by the CVP. Fig. 6 and Fig. 7 show that the calibration immensely improves our ability to capture the crossplane vorticity (vis-à-vis predictions with \mathbf{C}_{nom}), and improvements in predictions of vertical velocity largely follow. Note that the calibration was performed only on the $M = 0.8, J = 10.2$ case, and that too, using streamwise vorticity on the crossplane; however, the calibration improves the entire flowfield (as seen in the streamwise deficit and vertical velocity profiles) and is applicable also at other values of J .

Post-calibration, the only cause of the discrepancy between experimental observations and numerical predictions is the structural or model-form error in RANS. These errors are caused by the manner in which the turbulent stresses and their evolution are modeled. We use a linear eddy viscosity model, which

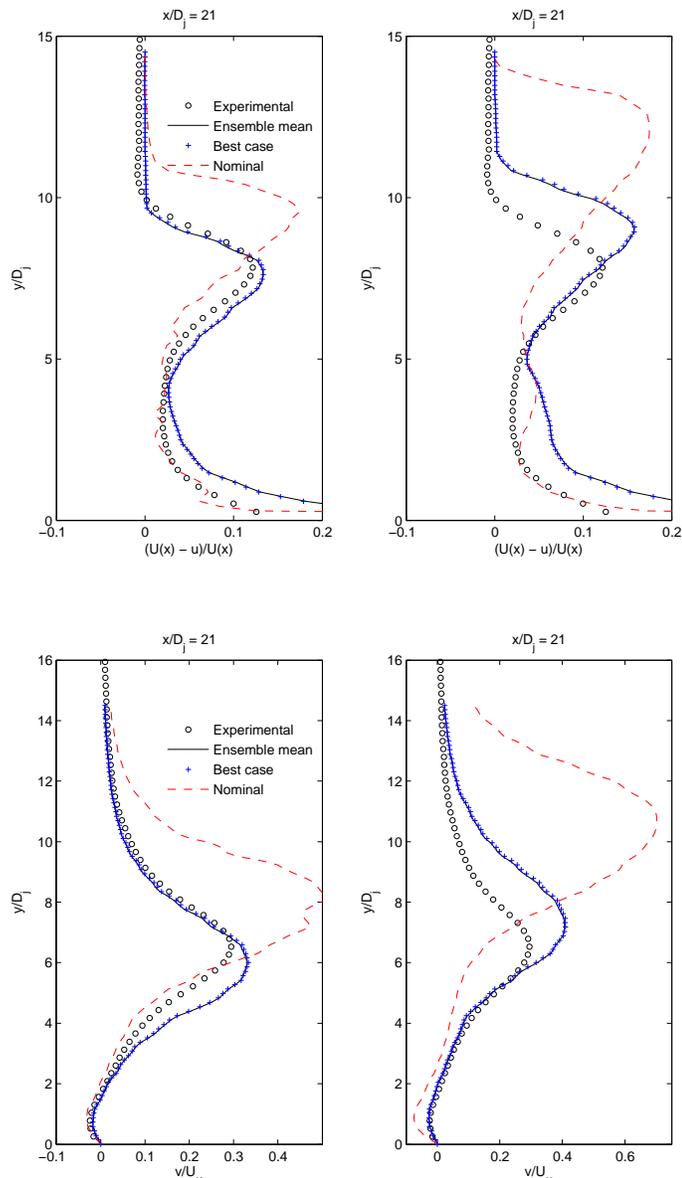


Figure 8. Top row: Streamwise velocity deficit at $x/D_j = 21$, as computed for the $J = 10.2$ and $J = 16.7$. In each subfigure on the left, we plot the streamwise velocity deficit $(U(x) - u)/U(x)$, where $U(x)$ is the bulk (or plug) flow velocity as measured experimentally (\circ), as predicted using \mathbf{C}_{nom} (dashed red line, $- - -$) and \mathbf{C}_{opt} ($+$). We also plot the ensemble mean predictions from the 100 $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ samples from the posterior distribution as the thin black line $-$ (it is very close to the \mathbf{C}_{opt} prediction). Bottom row: We compare the vertical velocity v/U_∞ instead of streamwise velocity deficit, for both values of J .

relates the turbulent stresses τ_{ij} to strain-rates S_{ij} and turbulent kinetic energy k , is stated as

$$\tau_{ij} = \overline{-u'_i u'_j} = \frac{2}{3} k \delta_{ij} - 2\nu_T S_{ij}, \quad k = \sum_i \overline{u_i'^2} \quad \text{and} \quad S_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) - \frac{1}{3} \delta_{ij} \sum_k \frac{\partial u_k}{\partial x_k} \quad (4)$$

and forms a convenient place to quantify and analyze the structural error our RANS model after calibration. In Fig. 9, we compare ensemble mean predictions of turbulent stress obtained from the 100 ($C_\mu, C_{\epsilon 2}, C_{\epsilon 1}$) samples obtained from the posterior distribution against experimental measurements from Ref. 18 for our two tests. We also plot the predictions obtained using \mathbf{C}_{nom} . Calibration to crossplane vorticity (for the $M = 0.8, J = 10.2$ case) results in a loss of agreement with measurements for τ_{22} (under-prediction) and τ_{11} (over-prediction). τ_{12} is grossly underpredicted. In fact, the ensemble mean predictions of τ_{11} and τ_{22} are almost indistinguishable. The predictions with \mathbf{C}_{nom} show the CVP sitting too high. The same behavior is seen at the off-calibration point - there is agreement for τ_{22} but the rest of the turbulent stresses are badly predicted. The agreement between the experimental values and the ensemble mean prediction degrades somewhat the off-calibration point. This behavior pinpoints the source of the structural error.

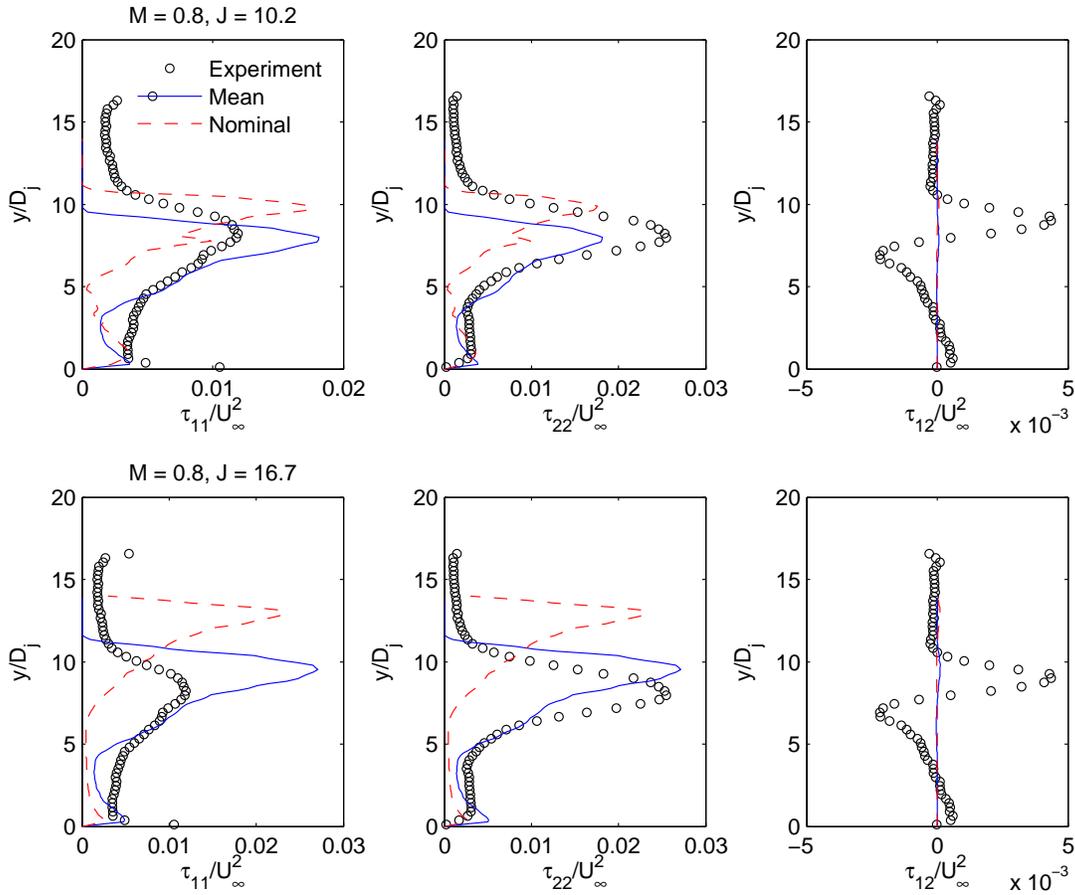


Figure 9. Profiles of $-\tau_{11}/U_\infty^2$ (first column), $-\tau_{22}/U_\infty^2$ (second column) and $-\tau_{12}/U_\infty^2$ (last column) as a function of y/D_j at $x = 200$ mm behind the jet on the mid-plane ($x/D_j = 20.9$). Experimental measurements are plotted with a \circ , the ensemble mean predictions from our 100 ($C_\mu, C_{\epsilon 2}, C_{\epsilon 1}$) values from the posterior are plotted with a solid line (—) and the predictions with \mathbf{C}_{nom} are plotted with a dashed line (---). Top row, we plot results from the calibration case ($M = 0.8, J = 10.2$). The bottom row contains results from the ($M = 0.8, J = 16.7$) test case.

The reason the ensemble mean predictions of τ_{ii} are so close is because the strain rates in the RANS predictions are too small. This leads to a gross under-prediction of τ_{12} using Eq. 4 (see Fig. 9) whereas as for τ_{ii} , k dominates over the term with S_{ij} . Calibration leads to a k that splits the difference between τ_{11} and τ_{22} . At higher values of J , the calibrated parameters over-predict k . This error is systematic and is

observed even at off-calibration points. The linear eddy viscosity model has too simple a form to capture the significant differences in the magnitude of the three stresses and a more complex model is required to bridge this gap. Yet, despite these large errors in the turbulent stress predictions, the large scale structures in the mean flow were only moderately affected, suggesting that the effect of these stresses are somewhat muted.

V. Conclusions

In this study, we have investigated whether RANS models can be calibrated to be predictive in jet-in-transonic-crossflow interactions. We have proceeded under the hypothesis that the nominal values of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1}) = \mathbf{C}_{nom}$ are unsuitable and a better set could be identified via calibration to experimental data. We find that it is possible to do so; further, the calibrated RANS model is predictive for a jet quite different from the one employed in the calibration.

We adopted a Bayesian calibration method and estimated the RANS parameters $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ as a joint PDF, thus quantifying the uncertainty in the estimation. We employed an adaptive Markov chain Monte Carlo (MCMC) method to do so. We also obtained a crude estimate of the data - model mismatch, an amalgamation of measurement and structural errors. Vorticity on the crossplane was used as the observable, and was found to be sufficient for estimating $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$. The calibration used surrogates of the 3D RANS simulator, due to the phenomenal computation cost of using MCMC. Construction of the surrogates posed a formidable challenge. While the bounds on $C_\mu, C_{\epsilon 2}$ and $C_{\epsilon 1}$ are known, their random combinations do not yield physically realistic flows, and in some cases, the RANS simulator will not even yield results. We identified a physically realistic region \mathcal{R} of the parameter space and developed polynomial surrogates that were valid only inside \mathcal{R} . The region \mathcal{R} is complex and we developed a classifier, using treed linear models, to identify whether an arbitrary $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ combination was inside it. \mathcal{R} was also leveraged to serve as an informative, if improper, prior for $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ - prior density was 1 inside \mathcal{R} and zero outside. The classifier, as well as the polynomial surrogates, were used with MCMC to obtain a 4-dimensional joint PDF for $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1}, \sigma^2)$.

The joint PDF for $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1}, \sigma^2)$ was constructed using streamwise vorticity data (from the crossplane) from an experiment conducted at $M = 0.8, J = 10.2$. 100 samples drawn from the posterior were used to run the 3D RANS simulator and predict the vorticity on the crossplane. Prediction errors were $\leq 5\%$, which was expected since vorticity was the calibration variable. However, the same runs yielded very good agreement with the experimentally measured streamwise velocity deficit and vertical velocity on the midplane, providing some corroboration that flowfield away from the crossplane was similar to the experimental one. The improvement over the flowfield computed with the nominal values of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1}), \mathbf{C}_{nom}$ are substantial, and can be seen in all the flow variables that we investigated. We used the joint PDF, computed using $M = 0.8, J = 10.2$ data, to predict the flowfield at $J = 16.7$. Our calibrated RANS model yielded improvements that ranged from substantial to surprising. Thus our hypothesis that the inaccuracy of RANS in jet-in-crossflow interactions was due to the unsuitability of \mathbf{C}_{nom} proved correct; in comparison, the structural error was small.

The structural error, nevertheless, does exist, and post-calibration, is the dominant error term. The structural error causes the degradation of predictive skill when we apply our RANS model at off-calibration points. We isolated the effect of the structural error on the midplane and pinpointed its source to be the simplicity of the linear eddy viscosity model used in our simulations. It predicts certain turbulent stresses well and is utterly ineffective for others. The contributions to the turbulent stresses by mean flow quantities (strain rate in the case of the linear eddy viscosity model) are grossly underpredicted. This error could be rectified by augmenting this closure model, for example, by employing a nonlinear eddy viscosity model. This investigation is left for future work.

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

- ¹Wilcox, D. C., *Turbulence Modeling for CFD*, D C W Industries, 1998.
- ²Poroseva, S. and Iaccarino, G., “Simulating separated flows using the $k - \epsilon$ model,” Tech. rep., Center for Turbulence Research, Annual Research Briefs, 2001.
- ³Kimura, I. and Hosoda, T., “A non-linear $k - \epsilon$ model with realizability for prediction of flows around bluff bodies,” *International Journal for Numerical Methods in Fluids*, Vol. 42, 2003, pp. 813–837.
- ⁴Durbin, P. A., “Near-wall turbulence closure modelling without damping functions,” *Theoretical and Computational Fluid Dynamics*, Vol. 3, 1991, pp. 1–11.
- ⁵Durbin, P. A., “Separated flow computations using the $k - \epsilon - v^2$ model,” *AIAA Journal*, Vol. 33, 1995, pp. 659–664.
- ⁶Yakhot, V. and Orszag, S., “Renormalization group analysis of turbulence,” *Physical Review Letters*, Vol. 57, 1986, pp. 1722–1724.
- ⁷Shih, T. H., Liou, W. W., Shabbir, A., and Yang, Z., “A new $k - \epsilon$ eddy viscosity model for high Reynolds number turbulence flows,” *Computers and Fluids*, Vol. 24, 1995, pp. 227–238.
- ⁸Lumley, J. L., “Computational modeling of turbulent flows,” *Advances in Applied Mechanics*, Vol. 18, 1978, pp. 123–177.
- ⁹Ristorcelli, J. R., Lumley, J. L., and Abid, R., “A rapid-pressure covariance representation consistent with Taylor-Proudman theorem materially-frame-indifferent in the 2D limit,” *Journal of Fluid Mechanics*, Vol. 292, 1995, pp. 111–152.
- ¹⁰Grimaji, S. S., “Pressure-strain correlation modelling of complex turbulent flows,” *Journal of Fluid Mechanics*, Vol. 422, 2000, pp. 91–123.
- ¹¹Mahesh, K., “The interaction of jets with crossflows,” *Annual Reviews of Fluid Mechanics*, Vol. 45, 2013, pp. 379–407.
- ¹²Arunajatesan, S., “Evaluation of two-equation RANS models for simulation of jet-in-crossflow problems,” *50th AIAA Aerospace Sciences Meeting*, 2012.
- ¹³Ray, J., Lefantzi, S., Arunajatesan, S., and Dechant, L., “Bayesian calibration of a $k - \epsilon$ model for predictive jet-in-crossflow simulations,” *Proceedings of the 44th AIAA Fluid Dynamics Conference*, 2014, Paper No. AIAA-2014-2085.
- ¹⁴Sargsyan, K., Safta, C., Najm, H. N., Debusschere, B. J., Ricciuto, D., and Thornton, P., “Dimensionality reduction for complex models via Bayesian compressive sensing,” *International Journal for Uncertainty Quantification*, Vol. 4, No. 1, 2014, pp. 63–93.
- ¹⁵Peterson, C. W., Wolfe, W. P., and Payne, J. L., “Experiments and computations of roll torque induced by vortex-fin interaction,” *42nd Aerospace Sciences Meeting and Exhibit, Reno, Nevada*, 2004.
- ¹⁶Beresh, S. J., Heineck, J. T., Walker, S. M., Shairer, E. T., and Yaste, D. M., “Vortex structure produced by a laterally inclined supersonic jet in transonic crossflow,” *Journal of Propulsion and Power*, Vol. 23, No. 2, 2007, pp. 353–363.
- ¹⁷Beresh, 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.
- ¹⁸Beresh, 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.
- ¹⁹Beresh, 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.
- ²⁰Beresh, 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.
- ²¹Chai, X. and Mahesh, K., “Simulations of high speed turbulent jets in crossflow,” *49th AIAA Aerospace Sciences Meeting, Orlando Florida*, 2011.
- ²²Kawai, S. and Lele, S. K., “Dynamics and mixing of a sonic jet in a supersonic turbulent crossflow,” *Center for Turbulence Research, Annual Research Briefs*, 2009.
- ²³Genin, F. and Menon, S., “Dynamics of sonic jet injection into supersonic crossflow,” *Journal of Turbulence*, Vol. 11, 2010, pp. 1–13.
- ²⁴Rodebaugh, G. P., Brinckman, K. W., and Dash, S. M., “DDES of aeropropulsive flows based on an extended $k - \epsilon$ RANS model,” *51st AIAA Aerospace Sciences Meeting, Grapevine, Texas*, 2013.
- ²⁵Emory, M., Pecnik, R., and Iaccarino, G., “Modeling structural uncertainties in Reynolds-Averaged computations of shock/boundary layer interactions,” *49th AIAA Aerospace Sciences Meeting*, 2011.
- ²⁶Gorlé, C. and Iaccarino, G., “A framework for epistemic uncertainty quantification of turbulent scalar flux models for Reynolds-averaged Navier-Stokes simulations,” *Physics of Fluids*, Vol. 25, 2013, pp. 055105.
- ²⁷Oliver, T. A. and Moser, R. D., “Bayesian uncertainty quantification applied to RANS turbulence models,” *Journal of Physics: Conference Series*, Vol. 318, 2011, pp. 042032.
- ²⁸Cheung, S. H., Oliver, T. A., E. E. Prudencio, S. P., and Moser, R. D., “Bayesian uncertainty analysis with applications to turbulence modeling,” *Reliability Engineering and System Safety*, Vol. 96, 2011, pp. 1137–1149.
- ²⁹Dow, E. and Wang, Q., “Quantification of structural uncertainties in the $k - \omega$ turbulence model,” *52nd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, 2011.
- ³⁰Edeling, 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.
- ³¹Edeling, W. N., Cinnella, P., and Dwight, R. P., “Predictive RANS simulations via Bayesian Model-Scenario Averaging,” *Journal of Computational Physics*, Vol. 275, 2014, pp. 65–91.
- ³²Jouhaud, J.-C., Sagaut, P., Enaux, B., and Laurenceau, J., “Sensitivity analysis and multiobjective optimization for LES numerical parameters,” *Journal of Fluid Engineering*, Vol. 130, 2008, pp. 021401.
- ³³Duda, R. O., Hart, P. E., and Storck, D. G., *Pattern Classification*, Wiley-Interscience, 2000, Second Edition.
- ³⁴Chipman, H. A., George, E. I., and McCulloch, R. E., “Bayesian treed model,” *Machine Learning*, Vol. 48, 2002, pp. 299–320.

- ³⁵Gramacy, R. B. and Lee, H. K. H., “Bayesian treed Gaussian process models with an application to computer modeling,” *Journal of the American Statistical Association*, Vol. 103, No. 483, 2008, pp. 1119–1130.
- ³⁶Gilks, W. R., Richardson, S., and Spiegelhalter, D. J., editors, *Markov Chain Monte Carlo in Practice*, Chapman & Hall, 1996.
- ³⁷Silverman, B. W., *Density Estimation for Statistics and Data Analysis*, Chapman and Hall, 1986.
- ³⁸Haario, H., Laine, M., Mira, A., and Saksman, E., “DRAM-Efficient adaptive MCMC,” *Statistics and Computing*, Vol. 16, No. 4, 2006, pp. 339–354.
- ³⁹R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2012, ISBN 3-900051-07-0.
- ⁴⁰Soetaert, 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.
- ⁴¹Raftery, 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.
- ⁴²Warnes, G. R. and with contributions by Robert Burrows, *mcgibbsit: Warnes and Raftery’s MCGibbsit MCMC diagnostic*, 2011, R package version 1.0.8.
- ⁴³Brinkman, K. W., Calhoun, W. H., and Dash, S. M., “Scalar fluctuation modeling for high-speed aeropropulsive flows,” *AIAA Journal*, Vol. 45, No. 5, 2007, pp. 1036–1046.
- ⁴⁴Gramacy, R. B., “tgp: An R package for Bayesian nonstationary, semiparametric nonlinear regression and design by treed Gaussian process models,” *Journal of Statistical Software*, Vol. 19, No. 9, 2007, pp. 1–46.
- ⁴⁵Lefantzi, S., Ray, J., Arunajatesan, S., and Dechant, L., “Estimation of $k - \epsilon$ parameters using surrogate models and jet-in-crossflow data,” SAND Report SAND2015-0707, Sandia National Laboratories, Livermore, CA 94551-0969, February 2015, Unclassified and unlimited release.