SANDIA REPORT

SAND2013-8158

Unlimited Release

Printed September 2013

Tuning a RANS $k-\varepsilon$ Model for Jet-in-Crossflow Simulations

Sophia Lefantzi, Jaideep Ray, Srinivasan Arunajatesan and Lawrence Dechant
Sandia National Laboratories, CA and Sandia National Laboratories, NM

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

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.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.osti.gov/bridge>

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online ordering: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>

Tuning a RANS $k - \epsilon$ Model for Jet-in-Crossflow Simulations

Sophia Lefantzi and Jaideep Ray
Sandia National Laboratories, Livermore CA 94551

Srinivasan Arunajatesan and Lawrence Dechant
Sandia National Laboratories, Albuquerque NM 87185
{slefant,jairay,sarunaj,ljdecha}@sandia.gov

Abstract

We develop a novel calibration approach to address the problem of predictive $k - \epsilon$ RANS simulations of jet-in-crossflow. Our approach is based on the hypothesis that predictive $k - \epsilon$ parameters can be obtained by estimating them from a strongly vortical flow, specifically, flow over a square cylinder. In this study, we estimate three $k - \epsilon$ parameters, C_μ , $C_{\epsilon 2}$ and $C_{\epsilon 1}$ by fitting 2D RANS simulations to experimental data. We use polynomial surrogates of 2D RANS for this purpose. We conduct an ensemble of 2D RANS runs using samples of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ and regress Reynolds stresses to the samples using a simple polynomial. We then use this surrogate of the 2D RANS model to infer a joint distribution for the $k - \epsilon$ parameters by solving a Bayesian inverse problem, conditioned on the experimental data. The calibrated $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ distribution is used to seed an ensemble of 3D jet-in-crossflow simulations. We compare the ensemble's predictions of the flowfield, at two planes, to PIV measurements and estimate the predictive skill of the calibrated 3D RANS model. We also compare it against 3D RANS predictions using the nominal (uncalibrated) values of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$, and find that calibration delivers a significant improvement to the predictive skill of the 3D RANS model. We repeat the calibration using surrogate models based on kriging and find that the calibration, based on these more accurate models, is not much better than those obtained with simple polynomial surrogates. We discuss the reasons for this rather surprising outcome.

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.

Contents

1	Introduction	7
2	Background	9
3	Calibration to a square-cylinder flow	11
3.1	Flow over a square cylinder	11
3.2	Construction of polynomial surrogates for 2D RANS	11
3.3	Bayesian calibration	12
3.4	Checking the calibration	13
4	Jet-in-crossflow predictions	19
5	Calibration with kriging surrogates	25
6	Conclusions	27
	References	28

Figures

1	Above: Flood plot of the streamwise velocities with streamlines seeded just upstream of the cylinder and the location of the probes. The picture is not to scale. The dominant flow is left to right. Below: The location of all 96 probes (open symbols), the 55/96 for which surrogate models can be constructed with $< 10\%$ error (filled symbols) and the 28/96 probes which are included in $\mathcal{P}$ (in rectangles).	14
2	Comparison of $\delta_{LS}^{(p)}$ and $\delta_{TS}^{(p)}$ for all 96 probes. 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.	15
3	Plots of the posterior density for $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$. The prior for each parameter is plotted with a dashed line. The vertical line denotes the nominal value.	16
4	Results from the posterior predictive test, performed with the polynomial surrogates. Results are only for probes in $\mathcal{P}$. The error bars denote the interquartile range (IQR) of the predicted $\overline{u'v'}$, while the circle is the median prediction. We see that most experimental readings lie in the IQR.	17
5	Comparison of predicted quantities from the 2D square-cylinder RANS simulations seeded using the calibrated JPDP. Top: Comparison of the experimental (symbols) $\overline{u'v'}$ with the ensemble predictions (median and IQR) and predictions with C_{nom} (dashed, in red). The predictions using the MAP $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ values are in green. Predictions are for $x/D = 3$ (left) and $x/D = 8$ (right) of probes. Bottom: Comparison of streamwise velocity u at the same locations.	18
6	Above: A schematic of the simulation test case. Measurements are available at $x/D_j = 33.6$ and $z/D_j = 0$. The figure shows the position of the jet and its roll-up into a counter-rotating vortex pair. Below: The streamwise vorticity field at $x/D_j = 33.6$. The box denotes the region of interest $\mathcal{R}$ where we will quantify the vorticity as we test our calibration. Vorticity units in s^{-1}	20
7	Distribution of normalized circulation, the location of the vorticity centroids and the radius of gyration, from JIC runs seeded by $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ samples from the JPDP. The metrics calculated from a JIC simulation using C_{nom} is plotted as a reference.	21
8	Above: Vorticity predictions using C_{nom} with experimental results plotted as contours. Below: Predictions using C_{opt} . The improvement in the vorticity field is quite obvious - the uncalibrated vorticity prediction is too large and is in the wrong position. Vorticity units in s^{-1}	22
9	Above: We plot the streamwise velocity deficit $(U_\infty - u(y))/U_\infty$ at $x/D = 31.5$ and 42.0 on the $z = 0$ plane of symmetry in the top two figures. We plot the experimental values (symbols), the ensemble predictions (median and IQR) and the predictions using C_{nom} (dashed, in red) and C_{opt} (green). Below: We plot $v(y)$ at the same locations in the bottom two figures.	23
10	Marginal posterior distributions for $C_\mu, C_{\epsilon 2}$ and $C_{\epsilon 1}$, developed with surrogates with kriged d , are plotted in red. The black lines are densities obtained using a polynomial surrogate (Sec. 3), whereas the dashed line is the prior.	26

1 Introduction

Jet-in-crossflow (JIC) is a canonical flow problem in fuel-air mixing in scramjet engines [1]. It is also relevant in the flight dynamics of vehicles maneuvered by spin rockets, where the rockets' exhaust modifies the flow over control fins [2]. In most engineering settings, this interaction is simulated using Reynolds-Averaged Navier-Stokes (RANS) equations, in particular, with the $k - \epsilon$ turbulence model. Prior investigations that compared a number of RANS formulations have shown that they are not very predictive for JIC studies [3]. Consequently, Large Eddy Simulations (LES) have been proposed as an alternative [4]. Others have postulated that the $k - \epsilon$ model is fundamentally lacking and have attempted to estimate the $k - \epsilon$ model deficiency by comparing against LES of JIC. The estimated uncertainty is thereafter propagated through to RANS predictions [1]. Both the approaches have their drawbacks in a practical engineering situation. LES are extremely expensive and infeasible in routine engineering design where a huge number of runs have to be made. The second method also requires one to modify existing RANS codes (if one has access to the source) and repeat the entire verification and validation exercise.

In this paper we will demonstrate an alternative approach to predictive JIC simulations, based on a *calibrated* $k - \epsilon$ model. The $k - \epsilon$ RANS model [5, 6] has four independent flow parameters (and one that is kept at a constant ratio to one of the independent ones); if wall models are used, they may add a few more. The parameter values are considered to be *universal* and are evaluated by calibrating to Direct Numerical Simulations and experiments of simple flow configurations e.g., isotropic homogeneous turbulence, channel flow etc. [6, 7]. We call these parameter values the “nominal” ones. There is ample empirical evidence that these parameters are far from being universal [6, 8, 9, 10, 11, 12, 13, 14] and the optimal parameter values can vary substantially (from the nominal ones) depending on the flow configuration. Thus there is no reason to suppose that the nominal values of $k - \epsilon$ model parameters should yield accurate JIC simulations, and calibration (preferably to experimental data) should be a pre-requisite for predictive JIC investigations. However, “wrapping” an optimizer e.g. L-BFGS [15] around a 3D JIC RANS simulation is hardly practical - the model is too computationally expensive for the $O(100)$ sequential evaluations that might be required.

We hypothesize that predictive JIC simulations can be obtained by calibrating a $k - \epsilon$ model to a simpler but related flow problem. The JIC flow is strongly vortical and in this paper, flow over a square cylinder [16, 17] will serve as the calibration problem. The flow will be approximated using a 2D RANS $k - \epsilon$ model, and $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ will serve as the calibration parameters. We will first sample the $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ space to seed 2D RANS simulations. The calibration observable $(\overline{u'v'})$ will be recorded at each experimental probe location and their variation with $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ will be represented using a surrogate model. We will investigate two surrogates: (1) the simpler polynomial model and (2) a more complicated one, based on universal kriging. In case of the polynomial model, for each probe location, we will have a polynomial that captures the variations of $\overline{u'v'}$ as a function of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$. We will then pose a Bayesian inverse problem, predicated on the polynomial surrogates of RANS, and conditioned on experimental data [16, 17]. We will use an adaptive Markov chain Monte Carlo (MCMC) method to infer a joint distribution for $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$. The calibration will be checked against the square cylinder measurements by performing a posterior predictive test i.e., we will attempt to reproduce the experimental results with the calibrated model. Thereafter, we will sample the joint $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ distribution to seed an ensemble of 100 3D JIC simulations. The ensemble results will provide a probabilistic prediction of the flowfield and we will check its accuracy against available JIC experimental data [18, 19].

In Sec. 2 we review recent efforts to improve the predictive skill of RANS and LES JIC models. In Sec. 3, we describe the construction and subsequent simplification of the polynomial models, the solution of the Bayesian inverse problem to estimate $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ as well as a simple estimate of the model deficiency (or structural error) i.e., the inability of the *calibrated* RANS model to reproduce experimental data. These will be performed using the square-cylinder calibration test case. In Sec. 4 we use the calibration (specifically the joint distribution of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$) to estimate the improvement in JIC prediction accuracy, versus the nominal $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ values, by comparing against JIC experiments. In Sec. 5 we calibrate using the kriged models, to check their usefulness versus simple polynomial models. We conclude in Sec. 6.

This page intentionally left blank

2 Background

Jet-in-crossflow (JIC) interaction plays a central role in the control of finned bodies of revolution that are maneuvered by spin rockets; the exhaust from the rockets induce an angle of attack at the fins, significantly modifying the aerodynamic forces and moments [20]. The problem is strongly vortical and has been extensively studied experimentally [18, 19, 21] (henceforth, collectively known as the “Beresh experiments”). The flowfield is dominated by a counter-rotating vortex pair, formed by a Kelvin-Helmholtz roll-up of the shear layer at the boundary between the jet and the crossflow. The vortex pair tracks the spatial evolution of the jet and has horse-shoe vortices wrapped around it. The counter-rotating vortex pairs and horse-shoe vortices are primarily responsible for the modification of the flowfield in the vicinity of the fins [2].

Arunajatesan [3] performed a comparison of various $k - \omega$ RANS models [22] versus experimental JIC results. He found that all the $k - \omega$ models over-predicted the turbulent intensities and Reynolds shear stresses inside the jet, resulting in a wider (“fatter”) jet vis-à-vis experiments i.e., the turbulent diffusion was over-predicted. However, the velocity field on a plane transverse to the streamwise flow (henceforth, the cross-plane) showed that the vortex pair resided at a point higher than the experimental measurements, a result at odds with the hypothesis of overly strong turbulent diffusion. He conjectured that turbulent stresses are underpredicted in the nearfield of the jet, resulting in an erroneous exchange of momentum between the jet and the crossflow. Investigations for various jet-to-crossflow momentum ratios and jet-crossflow cant angle revealed that the general behavior did not change - i.e., the lack of accuracy was characteristic of $k - \omega$ models being investigated rather than the particular flow configuration.

The JIC problem has also been studied using LES [23, 24, 25] and Detached Eddy Simulations (DES, [26]). The simulations were compared against the Beresh experiments, but only on the midplane (the plane of symmetry) and using mean velocities (streamwise and wall-normal) as the metrics of comparison. As expected, the comparison was far superior, vis-à-vis RANS. The performance of these models in capturing the velocity/vorticity field on the cross-plane was not tested. While this is encouraging, LES and DES are too computationally expensive to be used in routine design calculations; however, they can serve as numerical experiments, which are then used to quantify the uncertainty in RANS models, and to calibrate them.

Given the inaccuracy of RANS simulations (and not just in JIC simulations), it is natural to try and augment the empirical models for the creation and dissipation of turbulent kinetic energy with a “correction term”. The correction term is assumed to be spatially variable, is modeled statistically – usually as a multivariate Gaussian field – and estimated from numerical simulations’ data obtained from LES or Direct Numerical Simulations (DNS). This process is known as estimating the “structural” or “model-form” error in RANS models and has been performed for simple flows e.g., channel flow, flow over flat plate etc. [27, 28, 29]. However, 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. Extending the method to use experimental measurements may pose a stiff challenge since data are rarely available in such abundance. An alternative to calibration (which results in a more predictive RANS model) is to simply quantify the error/uncertainty in RANS predictions by comparing against LES data. This was outlined, for 2D flows, in [30], where turbulence variables from LES and RANS were compared to estimate the (spatially variable) difference/error. A statistical model (a probability density distribution) was created for the error. Thereafter, an ensemble of RANS simulations, each with a realization of the “error” field injected into it, were run to create an ensemble of predictions. Recently, this approach was extended to JIC computations [1].

The studies described above assume that the nominal parameters of the RANS model are optimal and the deficiencies of RANS predictions require augmentation of the model. One could also postulate that the model be *calibrated* first, i.e., estimate optimal parameters before embarking on a quest for structural errors. Such an approach, aimed at calibrating a LES model for JIC is described in [4]. The optimization of LES parameters, predicated on experimental data, was performed using surrogates of the LES model, constructed using kriging [31].

In this paper, we will describe a method to *calibrate* RANS models for JIC; the question of structural errors will be addressed in the future if the calibrated RANS model is found to be insufficiently accurate. Thus, conceptually, this work is similar to the LES calibration described in [4]. However, we introduce two novelties:

1. Our calibration is Bayesian i.e., instead of obtaining a point estimate for $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$, we construct a distribution for them. This distribution (henceforth, the “posterior distribution”) can then be used to perform an ensemble of JIC RANS runs to provide flow predictions as well as a quantification of prediction uncertainty.
2. We obtain our distribution of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ by calibrating to a strongly vortical, 2D flow over a square cylinder. 2D RANS simulations are computationally inexpensive (each run takes 4 minutes on 8 cores of a 2.6 GHz Intel Sandy Bridge processor), and our method for calibrating RANS models can be adopted in practical settings.

We demonstrate that the ensemble of predictions, obtained from the posterior distribution, results in far lower prediction errors (when compared against the Beresh experiments) vis-à-vis those obtained with nominal $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ values (henceforth, “nominal predictions”). Further, the uncertainty bounds on the prediction are sufficiently tight such that some variables in the nominal predictions are relegated to being statistical outliers. Finally, we show that samples drawn from the posterior distribution can be used to identify predictive point estimates of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$, for use in JIC simulations.

3 Calibration to a square-cylinder flow

In this section we describe the calibration of a 2D $k - \epsilon$ RANS model to experimental measurements of flow over a square cylinder. Per the discussion in Sec 1, there does not seem to be a universal value for $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$; consequently, we model them as random variables, whose joint probability density function (JPDF) is the object of calibration. The JPDF also allows us to accommodate uncertainty in the estimates due to limited experimental data. We expect that $O(10^4)$ sequential evaluations of the $k - \epsilon$ model will be required to obtain a converged $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ JPDF, which is too computationally expensive even for 2D RANS; consequently, we will construct polynomial surrogates for 2D RANS, before posing and solving the Bayesian calibration problem.

3.1 Flow over a square cylinder

$(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ are calibrated to experimental data from a flow-over-a-square-cylinder experiment conducted in a closed water channel. Details of the experiment are in [16, 17] and we provide a summary below. The channel, with a $39 \text{ cm} \times 56 \text{ cm}$ cross-section, is driven with a constant-head tank, with an average velocity of 0.535 ms^{-1} . The square cylinder, with dimensions of $56 \text{ cm} \times 4 \text{ cm} \times 4 \text{ cm}$ ($W \times D \times D$) is mounted in the middle, resulting in a blockage ratio of 9.75%. The Reynolds number was 21,400 and the flow is statistically steady. Reynolds-averaged shear stresses and fluctuating velocities are measured in the upper half of the vertical mid-plane via laser Doppler velocimetry. Measurements are made at a set of “probe” locations arranged in a grid. The grid of probes begin at the leading edge of the cylinder, though for our paper, we consider only those 96 probes which are in the wake. These are shown in Fig. 1 which also depicts the flowfield.

3.2 Construction of polynomial surrogates for 2D RANS

Let $\overline{u'v'}$ at probe p be denoted by $y^{(p)}$. We postulate that the following cubic polynomial approximation holds:

$$\begin{aligned} y^{(p)} &= \sum_{l=0}^3 \sum_{m=0}^3 \sum_{n=0}^3 \alpha_{l,m,n}^{(p)} (C_\mu)^l (C_{\epsilon 2})^m (C_{\epsilon 1})^n + d, \\ &= y_s^{(p)} + d, \\ 3 &\geq l + m + n, \end{aligned} \tag{1}$$

where $y_s^{(p)}$ is the polynomial approximation of $y^{(p)}$ and d is the approximation error.

In order to complete the model in Eq. 1, we first construct a dataset $\{C_\mu, C_{\epsilon 2}, C_{\epsilon 1}, y^{(p)}\}$ to capture the variation of $y^{(p)}$ with $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$. We use space-filling Halton sequences [32] (as implemented in the `randtoolbox` package [33] in R [34]) to draw 14^3 samples from the $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ parameter space and seed 2D RANS simulations of flow over a square cylinder. The number of samples to use was chosen as a compromise between the need to fully characterize the $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ space and the number of 2D RANS simulations that could be completed in limited time. The values of $y^{(p)}$ at the 96 probes are recorded, resulting in 96 separate $\{C_\mu, C_{\epsilon 2}, C_{\epsilon 1}, y^{(p)}\}$ datasets for the probes. Eq. 1 is regressed to the datasets using least-squares minimization.

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 Akaike Information Criterion (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. 1 for all 96 probes, 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 a 100-fold cross-validation test. The 14^3 RANS runs are separated into a “Learning Set” (LS) containing 2332 (approx-

imately 85% of 14^3) randomly selected runs, while the remainder constitute the “Testing Set” (TS). The simplified polynomial model is fitted to the “Learning Set” and the relative error for each parameter combination i in the LS, $\delta_{i,LS}^{(p)} = (y_i^{(p)} - y_{s,i}^{(p)})/y_i^{(p)}$, $i \in \text{LS}$, evaluated. The fitted model is also used to evaluate $y_{s,j}^{(p)}$ for the parameter set j in the TS, and calculate the corresponding relative error i.e., $\delta_{j,TS}^{(p)} = (y_j^{(p)} - y_{s,j}^{(p)})/y_j^{(p)}$, $j \in \text{TS}$. 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 96 probes.

In Fig. 2 we plot $\overline{\delta_{LS}^{(p)}}$ and $\overline{\delta_{TS}^{(p)}}$ for all the probes. We notice that the two relative errors have the same magnitudes 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. 2) as the criterion, we retain the 55 / 96 probes where the polynomial model is deemed to be sufficiently accurate.

3.3 Bayesian calibration

The polynomial models for $\overline{u'v'}$ at each probe p allow us to estimate $\mathbf{C} = (C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$ given experimental measurements $y_e^{(p)}$. Let $\mathbf{y}_s = \{y_s^{(p)}\}$, $p \in \mathcal{P}$, be the predictions by the polynomial surrogates for a set of probes $\mathcal{P}$, corresponding to $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$. Let $\mathbf{y}_e = \{y_e^{(p)}\}$, $p \in \mathcal{P}$ be the experimental values at the same set of probes. We model the experimental values at each probe p as

$$y_e^{(p)} = y_s^{(p)}(\mathbf{C}) + \delta_m, \quad \delta_m \sim \mathcal{N}(0, \sigma^2), \quad (2)$$

where δ_m is the model-data mismatch and $\mathcal{N}(0, \sigma^2)$ denotes a normal distribution with zero mean and a standard deviation of σ . Then the probability $P(\mathbf{y}_e|\mathbf{C})$ that the experimental data $\mathbf{y}_e$ can be obtained using $\mathbf{C}$ is

$$\begin{aligned} P(\mathbf{y}_e|\mathbf{C}) &= \prod_{p=1}^{|\mathcal{P}|} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(y_e^{(p)} - y_s^{(p)}(\mathbf{C}))^2}{2\sigma^2}\right) \\ &= \frac{1}{(2\pi)^{\frac{|\mathcal{P}|}{2}} \sigma^{2|\mathcal{P}|}} \exp\left(-\frac{\|\mathbf{y}_e - \mathbf{y}_s(\mathbf{C})\|_2^2}{2\sigma^2}\right) \end{aligned} \quad (3)$$

Let $\pi_\mu(C_\mu)$, $\pi_2(C_{\varepsilon 2})$ and $\pi_1(C_{\varepsilon 1})$ be probability density functions (PDFs) reflecting our prior beliefs regarding the variation of turbulence model parameters. Then by Bayes’ rule, the probability density (a.k.a the posterior density) of $\mathbf{C}$ is

$$P(\mathbf{C}|\mathbf{y}_e) \propto P(\mathbf{y}_e|\mathbf{C})\pi_\mu(C_\mu)\pi_2(C_{\varepsilon 2})\pi_1(C_{\varepsilon 1}) \quad (4)$$

We construct a JPDF for $\{C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1}, \sigma\}$, where σ serves as a nuisance variable and is marginalized over to obtain $P(\mathbf{C}|\mathbf{y}_e)$. The JPDF is constructed by sampling from the distribution (Eq. 3) using a Markov chain Monte Carlo (MCMC) technique [35]. Each sample requires one evaluation of the surrogate model for each probe in $\mathcal{P}$ (defined below). For efficiency, we use an adaptive MCMC called Delayed-Rejection Adaptive Metropolis [36], as implemented in the R package FME [37]; about 25,000 samples are required to reach convergence, as measured by the Raftery-and-Lewis metric [38], implemented in mcgibbsit [39]. The prior distributions used in this work are:

$$\begin{aligned} \pi_\mu(C_\mu) &= \mathcal{U}(0.06, 0.12), \\ \pi_2(C_{\varepsilon 2}) &= \mathcal{U}(1.7, 2.1), \\ \pi_1(C_{\varepsilon 1}) &= \mathcal{U}(1.2, 1.7), \end{aligned} \quad (5)$$

where $\mathcal{U}(a, b)$ is the uniform distribution bounded by (a, b) . σ^{-2} uses a conjugate prior from the inverse gamma distribution; see [37] for details. The bounds for $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$ are taken from [6]. The nominal values of $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$ are $\mathbf{C}_{nom} = \{0.09, 1.92, 1.44\}$.

Before we invert for $P(\mathbf{y}_e | \mathbf{C})$, we construct the ProbeSet $\mathcal{P}$. RANS predictions of $\overline{u'v'}$ for $\mathbf{C}_{nom}$ can have large errors and we restrict $\mathcal{P}$ to those probes where

$$\left| \frac{(y^{(p)} - y_s^{(p)}(\mathbf{C}_{nom}))}{y^{(p)}} \right| \leq 0.1 \quad \text{and} \quad \frac{1}{4} \leq \left| \frac{y_s^{(p)}}{y_e^{(p)}} \right| \leq 4.$$

Only 28 / 96 probes are included in $\mathcal{P}$; these are plotted with squares in Fig. 1.

In Fig. 3 we plot the marginalized PDFs for $C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1}$ and σ . The nominal values are plotted as vertical lines. We see that the nominal values are contained in the support of the marginalized PDFs; further, there is no well-defined peak for $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$. This implies a correlated JPDP for $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$, potentially with multiple modes. Further, the PDF for C_μ is not very different from the prior, indicating that $\mathbf{y}_e$ is not very informative about it. The observations are somewhat more informative about $C_{\varepsilon 2}$ and $C_{\varepsilon 1}$. The lack of a distinct mode also implies that estimating $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$ with a gradient-based optimization method could be difficult. In contrast, the PDF for σ has a well-defined peak, and consequently a *maximum a posteriori* (MAP) value.

3.4 Checking the calibration

We first conduct an interim check of the calibration by performing a posterior predictive test (PPT) i.e., we attempt to recreate the observations using the calibrated JPDP, σ and the polynomial surrogates. We draw 100 samples of $\{C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1}, \sigma\}$ from the JPDP, and use Eq. 1 to evaluate $y_e^{(p)}$, its median, the 25th and 75th percentiles. We plot these in Fig. 4 for the probes in $\mathcal{P}$, along with experimental values of $\overline{u'v'}$ from [17] and predictions using $\mathbf{C}_{nom}$. We note that the inter-quartile range (IQR) spanned by the error bars capture the experimental measurements, whereas the medians are not particularly close to them - indeed, they are not very different from the nominal predictions. Thus by adopting a Bayesian approach to account for uncertainty caused by the model-data mismatch (δ_m) and limited experimental data, we have successfully constructed predictive, but computationally inexpensive, surrogate models.

Next, we verify the calibration using RANS simulations of flow over a square cylinder. This allows us to compare model predictions to experiments at locations where we do not have accurate surrogate models. As before, we use the 100 samples of $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$. In Fig. 5 we compare $\overline{u'v'}$ at the column of probes at $x/D = 3$ and $x/D = 8$. RANS predictions are represented using the median and error bars corresponding to the IQR. Note that these predictions do not have δ_m added to them. We also plot the velocity u at the two locations in the bottom of the figure. We see that the turbulent region inside the separation vortex is far smaller than the one observed in experiments - in Fig. 5 (top) we see that in computations, non-zero values of $\overline{u'v'}$ are restricted to a small region. However post-calibration, the experimental values of peak Reynolds shear stress (defined as $\overline{u'v'}$ for incompressible flow, as opposed to the more conventional $\rho \overline{u'v'}$) are contained within the IQR of the numerical simulation results at $x/D = 3.0$. At $x/D = 8.0$, we see that the separation vortex experiences a smaller peak Reynolds stress in the experiment, as compared to numerical predictions. We note that the median predictions of $\overline{u'v'}$ are not too different from those obtained by $\mathbf{C}_{nom}$, and the net effect of Bayesian calibration is to add the credibility intervals. In Fig. 5 (below) we plot the streamwise velocity. We see clearly that the circulation in the vortex is over-predicted by the calibrated numerical simulation - the velocity is negative at $y/D = 0$ and greater than U_∞ at $y/D = 2$ at $x/D = 3$. In contrast, experimentally, the streamwise velocity is never negative and it reaches U_∞ at $y/D = 2$. At $x/D = 8$, the MAP (*maximum a posteriori*, the parameters with the maximum value of probability density) $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$ estimates of $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$ predict a streamwise velocity which is very close to the experimental value (plotted in green).

Figure 1. Above: Flood plot of the streamwise velocities with streamlines seeded just upstream of the cylinder and the location of the probes. The picture is not to scale. The dominant flow is left to right. Below: The location of all 96 probes (open symbols), the 55/96 for which surrogate models can be constructed with $< 10\%$ error (filled symbols) and the 28/96 probes which are included in $\mathcal{P}$ (in rectangles).

Figure 2. Comparison of $\overline{\delta_{LS}^{(p)}}$ and $\overline{\delta_{TS}^{(p)}}$ for all 96 probes. 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.

Figure 3. Plots of the posterior density for $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$. The prior for each parameter is plotted with a dashed line. The vertical line denotes the nominal value.

Figure 4. Results from the posterior predictive test, performed with the polynomial surrogates. Results are only for probes in $\mathcal{P}$. The error bars denote the interquartile range (IQR) of the predicted u/v , while the circle is the median prediction. We see that most experimental readings lie in the IQR.

Figure 5. Comparison of predicted quantities from the 2D square-cylinder RANS simulations seeded using the calibrated JPDF. Top: Comparison of the experimental (symbols) $\overline{u'v'}$ with the ensemble predictions (median and IQR) and predictions with C_{nom} (dashed, in red). The predictions using the MAP ($C_{\mu}, C_{\epsilon 2}, C_{\epsilon 1}$) values are in green. Predictions are for $x/D = 3$ (left) and $x/D = 8$ (right) of probes. Bottom: Comparison of streamwise velocity u at the same locations.

4 Jet-in-crossflow predictions

We finally test the predictive skill of the calibrated JPDF on the JIC problem. The test case is described in detail in [3, 18, 19] and we produce a summary here. A supersonic jet exits from a nozzle into a 12in $\times$ 12in test section of a blow-down wind-tunnel. The tunnel Mach number is 0.8. The nozzle diameter, D_j , is 0.375in, has a design Mach number of 3.73 and is pointed vertically upwards. PIV (particle image velocimetry) measurements of velocities (mean and turbulent) are available along the spanwise symmetry plane and on a plane, perpendicular to the streamwise flow, at a location 33.8 jet diameters downstream (the cross-plane). The tunnel stagnation temperature is 320K and that of the jet is 300K. In the RANS simulations, the inflow boundary is kept 10.66 jet diameters upstream. The coordinate system is anchored at the mid-point of the jet exit. The tunnel side and top walls are 16 and 32 jet diameters away. The top wall is modeled as a slip wall. The tunnel exit is 100 jet diameters downstream and characteristic non-reflective boundary conditions are imposed there. The entire jet nozzle, including the convergent part is meshed and included in the numerical model. The $k - \epsilon$ RANS equations are solved using SIGMA CFD, an in-house research code in the Aerosciences Department at Sandia National Laboratories. The version of the model used corresponds to that in [40]. Briefly, the high Reynolds number form of the $k - \epsilon$ model is augmented with a modified version of the damping functions proposed in [41] for near-wall low Reynolds number regions. These equations are solved in a fully coupled form with the conservation equations for mass, momentum and energy using a conservative second order TVD finite-volume scheme in a block-structured solver. Fig. 6 shows the schematic of the test case, including the location of the jet, its roll-up into a vortex pair and the spanwise cross-plane at $x = 33.6$ where the vorticity can be derived from experimental measurements. We define a rectangular region $\mathcal{R}$, $0.04\text{m} \leq y \leq 0.12\text{m}, 0 \leq z \leq 0.05\text{m}$ (the black box in Fig. 6 (below)), containing one of the vortices, where we will compute the metrics used to compare the calibrated versus nominal predictions.

We seed 100 3D RANS JIC simulations using $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ samples drawn from the calibrated JPDF (the same samples used to make Fig. 3 and 5) and, for each simulation, quantify the vorticity in $\mathcal{R}$ by calculating the circulation, the centroid (z, y) of the vorticity distribution and its radius of gyration (as an approximate measure of the vorticity distribution's spatial extent). These metrics are normalized by their counterparts calculated from experimental measurements [18, 19]. In Fig. 7, we plot the RANS JIC predictions from the JPDF, and compare with those obtained using $\mathbf{C}_{nom}$. The box spans the IQR of the 100 predictions and the red line is the median. It is clear that the circulation is vastly overestimated before calibration; in fact, the calibrated predictions forming the IQR are all more accurate than the nominal predictions. However, there is a clear bias in the circulation - it is overpredicted by the numerical simulations. The centroid of the vorticity distribution also shows a bias - z is underpredicted whereas y is overpredicted. The nominal predictions fall within the IQR. The radius of gyration is also overpredicted, but the median prediction is almost perfect, and not very different from the nominal prediction. Thus the main improvement is in the circulation prediction.

The clustering of the IQRs around 1 raise the question whether there exists a $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ combination for which the normalized predictions are all near 1. We define a goodness-of-fit score as the L_2 norm of the relative errors in circulation, the centroid and the radius of gyration (calculated by subtracting 1 from the normalized values) and search the 100 samples for the one with the minimal score. $\mathbf{C}_{opt} = \{0.117, 1.936, 1.262\}$ yields the smallest relative errors of 11.7%, -11.7%, -8.8% and 13.7% for circulation, the z - and y -centroids and radius of gyration respectively. In Fig. 8 we plot the vorticity distribution predicted using $\mathbf{C}_{nom}$ (above) and $\mathbf{C}_{opt}$ (below). The experimental vorticity distribution is overlaid as contours. The improvement in the vorticity prediction is quite stark.

In Fig. 9 we check the ensemble prediction on the $z = 0$ plane. We plot the streamwise velocity deficit $(U_\infty - u(y))/U_\infty$ and vertical velocities $v(y)$ at $x/D_j = 31.5$ and 42.0. We plot the experimental values (symbols), the ensemble predictions (median and IQR) and the solution corresponding to $\mathbf{C}_{opt}$ (green) and $\mathbf{C}_{nom}$ (red). We see that the velocity deficit is underpredicted by the predictions as well as $\mathbf{C}_{opt}$, but the vertical position of the peak deficit is approximately at the correct height. The improvements in the vertical velocity are more impressive. At both the locations, the vertical velocity predicted by $\mathbf{C}_{opt}$ is very close to the experimental results. $v(y)$, as predicted by $\mathbf{C}_{nom}$ sometimes lies outside the IQR.

Figure 6. Above: A schematic of the simulation test case. Measurements are available at $x/D_j = 33.6$ and $z/D_j = 0$. The figure shows the position of the jet and its roll-up into a counter-rotating vortex pair. Below: The streamwise vorticity field at $x/D_j = 33.6$. The box denotes the region of interest $\mathcal{R}$ where we will quantify the vorticity as we test our calibration. Vorticity units in s^{-1} .

Figure 7. Distribution of normalized circulation, the location of the vorticity centroids and the radius of gyration, from JIC runs seeded by $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ samples from the JPDF. The metrics calculated from a JIC simulation using C_{nom} is plotted as a reference.

Figure 8. Above: Vorticity predictions using C_{nom} with experimental results plotted as contours. Below: Predictions using C_{opt} . The improvement in the vorticity field is quite obvious - the uncalibrated vorticity prediction is too large and is in the wrong position. Vorticity units in s^{-1} .

Figure 9. Above: We plot the streamwise velocity deficit $(U_\infty - u(y))/U_\infty$ at $x/D = 31.5$ and 42.0 on the $z = 0$ plane of symmetry in the top two figures. We plot the experimental values (symbols), the ensemble predictions (median and IQR) and the predictions using C_{nom} (dashed, in red) and C_{opt} (green). Below: We plot $v(y)$ at the same locations in the bottom two figures.

This page intentionally left blank

5 Calibration with kriging surrogates

In this section we perform calibration using not just a model for $y_s^{(p)}$, but also a Gaussian process model for $d(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$ (see Eq. 1). A Gaussian process model (also called kriging) for $d(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$ can be constructed provided its variation in $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$ -space is smooth. In its absence, i.e., if the variation of d resembles white-noise, the estimate for d will be near zero (the mean of white noise).

Kriging models are constructed using $d_i, i = 1 \dots N_d$ defects obtained after the polynomial surrogates $y_s^{(p)}$ are constructed. Then, at any arbitrary point $\mathbf{C}^*$, one obtains an estimate as $d^* = \sum w_i d_i$. The kriging model decides w_i . Kriging is based on the assumption that the value of d at an unknown point should be an average of its known neighbors, perhaps weighted by the distance. There are many forms of kriging [42], but we will use *ordinary kriging* for d . When the ordinary-kriging model for d is augmented by the polynomial model for $y_s^{(p)}$, the composite model is equivalent to *universal kriging* (a Gaussian process model with a complex, polynomial mean).

We assume that $d_i, i = 1 \dots N_d$ are realizations from a Gaussian process D_i defined at points $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})_i$ in the $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$ space. In our case, the mean or expected value $E(D) = 0$ since the defects are the residuals of a least-squares regression to obtain $y_s^{(p)}$. The random variables D_i are correlated among themselves. We make the assumption that the Gaussian process in $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$ space is stationary i.e., the covariance between two arbitrary points $C(D_i, D_j) = f(|\mathbf{C}_i - \mathbf{C}_j|)$, where $|\mathbf{C}_i - \mathbf{C}_j|$ is the distance between the two points in the $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$ space. Thus the covariance between two points is dependent only on the distance between them and independent of their respective locations in $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$ -space. Colloquially, this implies that d is a smooth function without steep gradients. In such a case, the samples d_i can be used to construct a variogram model (in our case, we use Gaussian variograms), and therefore, a modeled covariance matrix. This, in turn, can be used to write a likelihood function for d at an arbitrary point in $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$ -space, using d_i . In our case, we used a fifth of our 14^3 RANS runs ($N_d \approx 600$) to construct the kriged model.

Constructing kriged models, for each of the probes in $\mathcal{P}$ is a time-consuming, manual affair. For each probe, we tested a number of variogram models (e.g., Gaussian, exponential etc). Fitting a particular variogram model required a nonlinear regression. This was repeated 10 times, in a cross-validation procedure, using a Testing Set that was 15% of N_d . This, in turn, was repeated for all probes in $\mathcal{P}$. The evaluation of a kriged model involves a sum over approximately 600 data points; in contrast, our polynomial models have less than 20. The kriged models are therefore a order-of-magnitude more expensive than polynomial ones. Further, the kriged interpolation is effective if d_i are correlated i.e., w_i are mostly non-zero. In many probes, the polynomial model provided an adequate fit, and d , as a function of $(C_\mu, C_{\varepsilon 2}, C_{\varepsilon 1})$, resembled white noise. The kriged models were ineffective in such a case.

The rationale behind modeling d lies in providing a better approximation to $y^{(p)}$ in Eq. 1. Since this is to be used in calibration, modeling d is relevant only if it is comparable (or larger) than the model-data mismatch δ_m (Eq. 2); else, its effect can be subsumed into δ_m . Fig. 3 shows that $\sigma, \delta_{\text{delta}_m} \sim \mathcal{N}(0, \sigma^2)$ is large and it is unclear whether kriging d will make a difference.

In Fig. 10 we plot the distribution of $C_\mu, C_{\varepsilon 2}$ and $C_{\varepsilon 1}$ obtained by augmenting the polynomial models from Sec. 3 with a kriged model for d . Clearly, the results are disappointing; the polynomial models provide an adequate calibration. The reason lies in the distribution for σ , which is virtually unchanged i.e., adding in the defect d made no impact on the model-data mismatch, largely due to RANS's inability to reproduce experimental data. Thus a more sophisticated surrogate, based on kriging, is largely unjustified, even if one makes allowances for the tedious process of constructing them and the order-of-magnitude increase in computational times.

Figure 10. Marginal posterior distributions for C_μ , $C_{\varepsilon 2}$ and $C_{\varepsilon 1}$, developed with surrogates with kriged d , are plotted in red. The black lines are densities obtained using a polynomial surrogate (Sec. 3), whereas the dashed line is the prior.

6 Conclusions

We have developed a method for calibrating computationally expensive 3D RANS models. The method is based on Bayesian calibration to experimental data. The method was demonstrated by evaluating its impact on improving the predictive skill of RANS in JIC simulations. The adoption of a statistical calibration method allows us to accommodate the shortcomings of RANS in many flow configurations, as well as calibration issues due to limited experimental measurements.

In developing this method, we have introduced two new approaches. First, instead of calibrating a large problem as-is [4], we search for a simpler, but related flow, which can be simulated in an inexpensive manner. Since JIC is strongly vortical, we chose flow over a square cylinder as a calibration case. It can be simulated accurately under a 2D approximation, which simplified the calibration process.

Second, we introduce Bayesian calibration or parameter estimation. This method evaluates the model (2D RANS, in this case) $O(10^4)$ times requiring us to create polynomial surrogates. We used existing model simplification methods using AIC and demonstrated, via cross-validation, that the resulting models are robust i.e., they do not overfit. Further, Bayesian calibration allows us to predict in a probabilistic manner using an ensemble, rather than attempt to reproduce experimental data with a single set of model parameters. Thus we can account for the various approximations introduced during calibration.

Bayesian calibration using flow over a square cylinder yielded a JPDF that was surprisingly predictive for JIC. Its predictions were far more accurate than the nominal values of $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$; in fact, the quantification of uncertainty in predictions indicated that some nominal predictions were statistical outliers. Certain samples drawn from the JPDF yielded accuracies in vorticity generation which have never been achieved by RANS simulations of JIC.

We repeated our calibration using a surrogate that contained a kriging model for the defect d in Eq. 1. The model endows enhanced accuracy in $y^{(p)}$ if d is smoothly distributed $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$ -space. We find that this enhanced accuracy does not lead to a calibration that is substantially different from the one obtained using $y_s^{(p)}$. The calibration, in this case, is determined by the ProbeSet $\mathcal{P}$ and the difference between experiment observations and RANS predictions. Since the surrogates do not influence either of these issues, the same $\mathcal{P}$ was used and the slight improvement obtained by modeling d paled into insignificance versus δ_m , the measure of our inability, via RANS, to reproduce experimental observation. Since kriged models are significantly more complex than polynomial ones, we will de-emphasize their use in the future.

Finally, our method is conceptually and practically simple. The majority of the statistical methods used in this work are available as R packages; new development consisted mainly of postulating and testing various model forms (e.g., cubic in $(C_\mu, C_{\epsilon 2}, C_{\epsilon 1})$). This is not difficult, and our method has the potential to be adopted in a practical engineering setting.

To summarize, we have demonstrated that turbulence models can be calibrated in complex flow configurations. The calibration is Bayesian and yields turbulence model parameters as distributions. The calibration improves the accuracy of turbulent flow simulations. The statistical tools and software required to do so are freely available as R packages.

References

- [1] C. Gorié and G. Iaccarino. A framework for epistemic uncertainty quantification of turbulent scalar flux models for Reynolds-averaged Navier-Stokes simulations. *Physics of Fluids*, 25:055105, 2013.
- [2] S. J. Beresh, J. T. Heineck, S. M. Walker, E. T. Schairer, and D. M. Yaste. Planar velocimetry of jet/fin interaction on a full-scale flight vehicle configuration. *AIAA Journal*, 45(8):1827–1840, 2007.
- [3] S. Arunajatesan. Evaluation of two-equation RANS models for simulation of jet-in-crossflow problems. In *50th AIAA Aerospace Sciences Meeting*, 2012.
- [4] J.-C. Jouhaud, P. Sagaut, B. Enaux, and J. Laurenceau. Sensitivity analysis and multiobjective optimization for LES numerical parameters. *Journal of Fluid Engineering*, 130:021401, 2008.
- [5] W. P. Jones and B. E. Launder. The prediction of laminarization with a two-equation model of turbulence. *International Journal of Heat and Mass Transfer*, 15:1–32, 1972.
- [6] S. Poroseva and G. Iaccarino. Simulating separated flows using the $k - \epsilon$ model. Technical report, Center for Turbulence Research, Annual Research Briefs, 2001.
- [7] I. Kimura and T. Hosoda. A non-linear $k - \epsilon$ model with realizability for prediction of flows around bluff bodies. *International Journal for Numerical Methods in Fluids*, 42:813–837, 2003.
- [8] P. A. Durbin. Near-wall turbulence closure modelling without damping functions. *Theoretical and Computational Fluid Dynamics*, 3:1–11, 1991.
- [9] P. A. Durbin. Separated flow computations using the $k - \epsilon - v^2$ model. *AIAA Journal*, 33:659–664, 1995.
- [10] V. Yakhot and S. Orszag. Renormalization group analysis of turbulence. *Physical Review Letters*, 57:1722–1724, 1986.
- [11] T. H. Shih, W. W. Liou, A. Shabbir, and Z. Yang. A new $k - \epsilon$ eddy viscosity model for high reynolds number turbulence flows. *Computers and Fluids*, 24:227–238, 1995.
- [12] J. L. Lumley. Computational modeling of turbulent flows. *Advances in Applied Mechanics*, 18:123–177, 1978.
- [13] J. R. Ristorcelli, J. L. Lumley, and R. Abid. A rapid-pressure covariance representation consistent with taylor-proudman theorem materially-frame-indifferent in the 2d limit. *Journal of Fluid Mechanics*, 292:111–152, 1995.
- [14] S. S. Grimaji. Pressure-strain correlation modelling of complex turbulent flows. *Journal of Fluid Mechanics*, 422:91–123, 2000.
- [15] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific and Statistical Computing*, 16(5):1190–1208, 1995.
- [16] D. A. Lyn and W. Rodi. The flapping shear layer formed by the flow separation from the forward corner of a square cyliner. *Journal of Fluid Mechanics*, 267, 1994.
- [17] D. A. Lyn, S. Einav, W. Rodi, and J. H. Park. A laser Doppler velocimetry study of ensemble averaged characteristics of the turbulent near wake of a square cylinder. *Journal of Fluid Mechanics*, 304:285–319, 1995.
- [18] S. J. Beresh, J. F. Henfling, R. J. Erven, and R. W. Spillers. Crossplane velocimetry of a transverse supersonic jet in a transonic crossflow. *AIAA Journal*, 44(12):3051–3061, 2006.
- [19] S. J. Beresh, J. F. Henfling, R. J. Erven, and R. W. Spillers. Penetration of a transverse supersonic jet into a subsonic compressible crossflow. *AIAA Journal*, 43(2):379–389, 2005.
- [20] C. W. Peterson, W. P. Wolfe, and J. L. Payne. Experiments and computations of roll toque induced by vortex-fin interaction. In *42nd Aerospace Sciences Meeting and Exhibit, Reno, Nevada*, 2004.

- [21] S. J. Beresh, J. T. Heineck, S. M. Walker, E. T. Shairer, and D. M. Yaste. Vortex structure produced by a laterally inclined supersonic jet in transonic crossflow. *Journal of Propulsion and Power*, 23(2):353–363, 2007.
- [22] D. C. Wilcox. *Turbulence Modeling for CFD*. D C W Industries, 1998.
- [23] X. Chai and K. Mahesh. Simulations of high speed turbulent jets in crossflow. In *49th AIAA Aerospace Sciences Meeting, Orlando Florida*, 2011.
- [24] S. Kawai and S. K. Lele. Dynamics and mixing of a sonic jet in a supersonic turbulent crossflow. In *Center for Turbulence Research, Annual Research Briefs*, 2009.
- [25] F. Genin and S. Menon. Dynamics of sonic jet injection into supersonic crossflow. *Journal of Turbulence*, 11:1–13, 2010.
- [26] G. P. Rodebaugh, K. W. Brinckman, and S. M. Dash. DDES of aeropropulsive flows based on an extended $k - \epsilon$ RANS model. In *51st AIAA Aerospace Sciences Meeting, Grapevine, Texas*, 2013.
- [27] T. A. Oliver and R. D. Moser. Bayesian uncertainty quantification applied to RANS turbulence models. *Journal of Physics: Conference Series*, 318:042032, 2011.
- [28] S. H. Cheung, T. A. Oliver, S. Prudhomme E. E. Prudencio, and R. D. Moser. Bayesian uncertainty analysis with applications to turbulence modeling. *Reliability Engineering and System Safety*, 96:1137–1149, 2011.
- [29] E. Dow and Q. Wang. Quantification of structural uncertainties in the $k - \omega$ turbulence model. In *52nd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, 2011.
- [30] M. Emory, R. Pecnik, and G. Iaccarino. Modeling structural uncertainties in Reynolds-Averaged computations of shock/boundary layer interactions. In *49th AIAA Aerospace Sciences Meeting*, 2011.
- [31] J. Laurenceau and P. Sagaut. Building efficient response surfaces of aerodynamic functions with kriging and cokriging. *AIAA Journal*, 46(2):498–507, 2008.
- [32] J. H. Halton. Algorithm 247: Radical-inverse quasi-random point sequence. *Communications of the Association for Computing Machinery*, 7(12):701–702, 1964.
- [33] Christophe Dutang and Petr Savicky. *randtoolbox: Generating and Testing Random Numbers*, 2013. R package version 1.13.
- [34] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2012. ISBN 3-900051-07-0.
- [35] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, editors. *Markov Chain Monte Carlo in Practice*. Chapman & Hall, 1996.
- [36] Heikki Haario, Marko Laine, Antoinietta Mira, and Eero Saksman. DRAM-Efficient adaptive MCMC. *Statistics and Computing*, 16(4):339–354, 2006.
- [37] Karline Soetaert and Thomas Petzoldt. Inverse modelling, sensitivity and monte carlo analysis in R using package FME. *Journal of Statistical Software*, 33(3):1–28, 2010.
- [38] A. Raftery and Steven M. Lewis. Implementing MCMC. In W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, editors, *Markov Chain Monte Carlo in Practice*, pages 115–130. Chapman and Hall, 1996.
- [39] Gregory R. Warnes and with contributions by Robert Burrows. *mcgibbsit: Warnes and Raftery’s MCGibbsit MCMC diagnostic*, 2011. R package version 1.0.8.
- [40] Dash S.M. Brinkman K.W., Calhoon W.H. Scalar fluctuation modeling for high-speed aeropropulsive flows. *AIAA Journal*, 45(5):1036–1046, 2007.
- [41] Gerodimos G. So R.M.C., Sarkar A. and Zhang J. A dissipation rate equation for low reynolds number and near-wall turbulence. *Theoretical and Computational Fluid Dynamics*, 9:47–63, 1997.
- [42] N. Cressie. *Statistics for spatial data*. Wiley, New York, 1993.

DISTRIBUTION:

1	Sophia Lefantzi, 08954	MS 9152
1	Jaideep Ray, 08954	MS 9159
1	Jerry Mcneish, 08954	MS 9159
1	Srinivasan Arunajatesan, 01515	MS 0825
1	Lawrence Dechant, 01515	MS 0825
1	Jeffrey Payne, 01515	MS 0825
1	Mary Gonzales, 08250	MS 9042
1	Technical Library, 08944 (electronic)	MS 0899