



Science Arts & Métiers (SAM)

is an open access repository that collects the work of Arts et Métiers ParisTech researchers and makes it freely available over the web where possible.

This is an author-deposited version published in: <https://sam.ensam.eu>
Handle ID: <http://hdl.handle.net/10985/15499>

To cite this version :

Martin SCHMELZER, Richard P. DWIGHT, Wouter EDELING - Estimation of Model Error Using Bayesian Model-Scenario Averaging with Maximum a Posteriori-Estimates -

Any correspondence concerning this service should be sent to the repository

Administrator : archiveouverte@ensam.eu



Estimation of model error using Bayesian Model-Scenario Averaging with *Maximum A Posteriori*-estimates

Martin Schmelzer¹, Richard P. Dwight¹, Wouter Edeling², Paola Cinnella³

¹ Delft University of Technology,
Faculty of Aerospace Engineering,
Kluyverweg 2, Delft, The Netherlands
m.schmelzer@tudelft.nl, r.p.dwight@tudelft.nl

² Stanford University, Center for Turbulence Research,
488 Escondido mall, Stanford 94305, USA
wedeling@stanford.edu

³ Arts et Métiers ParisTech,
DynFluid laboratory,
151 Boulevard de l'Hospital, 75013, Paris, France
paola.cinnella@ensam.eu

Abstract. The lack of an universal modeling approach for turbulence in Reynolds-Averaged Navier-Stokes simulations creates the need for quantifying the modelling error without additional validation data. Bayesian Model-Scenario Averaging (BMSA), which exploits the variability on model closure coefficients across several flow scenarios and multiple models, gives a stochastic, a posteriori estimate of a quantity of interest. The full BMSA requires the propagation of the posterior probability distribution of the closure coefficients through a CFD code, which makes the approach infeasible for industrial relevant flow cases. By using *maximum a posteriori* (MAP) estimates on the posterior distribution, we drastically reduce the computational costs. The approach is applied to turbulent flow in a pipe at $Re = 44,000$, over 2D periodic hills at $Re_H = 5600$ and finally over a generic falcon jet test case (Industrial challenge IC-03 of the UMRIDA project).

Keywords: CFD, UQ, RANS, Turbulence Modelling, Bayesian Calibration, Bayesian Scenario-Model Averaging

1 Introduction

The understanding of turbulence is one of the key challenges in classical mechanics. A turbulent flow has a three-dimensional, time-dependent and random velocity field [1], which is composed of a wide range of scales varying from the level close to molecular dissipation of heat to the scales determined by the boundary conditions of the flow domain. Despite great developments in the field of high-fidelity turbulence modelling, such as improvements of the large eddy simulation (LES) approach and hybrid-methods combining Reynolds-averaged Navier-Stokes (RANS) and LES, RANS continues to be the standard approach used to predict a wide range of flows for very complex configurations across virtually all aerospace product categories' [2]. However, using the

less-computationally demanding RANS approach comes at the price of uncertainty due to approximate physical modelling such as turbulence modelling.

In recent years research has focussed on two types of epistemic uncertainties in relation to turbulence modelling, namely parameter uncertainty and model-form uncertainty. Parameter uncertainty arises due to the fact that the closure coefficients of the models are determined via calibration against simple flow configurations using experimental data or scale-resolving simulations. A Bayesian calibration incorporating both the error of the experimental data and model inadequacy revealed how strongly the posterior distributions of the closure coefficients vary even for a simple flow scenario of flow over a flat plate, with different pressure gradients ranging from favourable to strongly adverse [3]. This observation makes the generalisation of the coefficient unjustifiable. Similarly, a Bayesian analysis dealing with the predictive performance of commonly used turbulence models, i.e. Launder-Sharma $k - \epsilon$, Wilcox $k - \omega$, Spalart-Allmaras, Baldwin-Lomax and Stress- ω , showed that no superior model could be identified either for the given flow configurations [3].

Bayesian Model-Scenario Averaging (BMSA) uses a set of different closure models to compute both an ensemble prediction as well as a-posteriori measures of uncertainty due to the choice of closure model. In addition, the variability of the closure coefficients over different calibration scenarios is included to inject uncertainty caused by applying a set of coefficients to a predictive case for which they were not specifically calibrated [4]. The full BMSA approach requires the propagation of various posterior distributions through a CFD code, an expensive endeavour when the underlying flow case is computationally expensive. The costs could be mitigated by using surrogate models per turbulence model in order to propagate the distributions more efficiently [5]. However, also the construction of the surrogate can become expensive for models with many closure coefficients.

In this work we use a major simplification of the full BMSA approach in order to make the technique accessible for industrial flow cases. The reduction of the costs is based on using *maximum a posteriori* (MAP) estimates of the posterior distributions of the closure coefficients, which means that only a single set of closure coefficients per scenario and per model needs to be propagated through the code [6]. Furthermore, the BMSA method was developed using data of flat plate boundary layer flows, for which the Bayesian inference of the posterior closure coefficient was conducted with a cheap boundary-layer code. The method has recently been applied to airfoil and wing cases as examples in order to assess the predictive capabilities of the BMSA framework for wall-bounded flows in external aerodynamics [6], which represents a natural next step in terms of flow complexity. In this work we apply the method to other flow cases such as flow in a turbulent pipe at $Re = 44,000$ and flow over periodic hills at $Re_H = 5600$ in order to assess the robustness of the method also for flow configurations outside of this natural domain. Finally, we also show results for one of the industrial challenges of the UMRIDA project: a generic Falcon Jet (IC-03).

2 Bayesian predictive methodology

2.1 Calibration

The BMSA framework is based on the Bayesian calibration of the coefficients of closure models. Treating calibration as a stochastic problem we obtain posterior probability distributions for the coefficients, which serve as our uncertain estimates of the coefficients under the measurement error of the reference data and the modelling error. The data used for the calibration consists of boundary-layer data-sets from the 1968 AFOSR-IFP-Stanford conference proceedings [7], which are both highly resolved and subject to low measurement noise. The set contains a collection of wide range of favourable and adverse pressure gradients.

A *scenario*, denoted S , is a particular flow set-up, including boundary-conditions, material parameters, and all other physical properties needed to define the flow, with corresponding experimental data \mathbf{z} . The CFD code $m_{\text{CFD}}(S; M, \boldsymbol{\theta})$ takes as arguments the scenario $S \in \mathcal{S} = \{S_1, \dots, S_K\}$, a turbulence model $M \in \mathcal{M} = \{M_1, \dots, M_I\}$ and its closure-coefficients $\boldsymbol{\theta}$. Given an operator $H_z(\cdot)$ that maps the state to the measured quantities \mathbf{z} we can define a statistical model to relate \mathbf{z} and $\boldsymbol{\theta}$:

$$\mathbf{z} = \delta [H_z \circ m_{\text{CFD}}(S; M, \boldsymbol{\theta})] + \varepsilon, \quad (1)$$

where ε , δ are random-variables (RVs) representing measurement noise and multiplicative model error respectively. The error of the measurements is modelled as zero-mean additive Gaussian noise and the RV δ is a representation for the model error, which following the approach of Cheung et. al. [8] specifies δ as a correlated Gaussian process, see also [3]. The Gaussian choice for both ε and δ yields a Gaussian likelihood function, i.e. the pdf describing the probability of observing the data given a realization of $\boldsymbol{\theta}$. Finally, an application of Bayes theorem [9] yields the expression for the posterior distribution of $\boldsymbol{\theta}$

$$p(\boldsymbol{\theta} | z, M, S) = \frac{p(z | \boldsymbol{\theta}, M, S) p(\boldsymbol{\theta} | M, S)}{p(z | M, S)} \propto p(z | \boldsymbol{\theta}, M, S) p(\boldsymbol{\theta} | M, S). \quad (2)$$

Here, $p(\boldsymbol{\theta} | M, S)$ is the prior distribution which represents our knowledge of $\boldsymbol{\theta}$ prior to training M on z . Since the denominator in (2) does not depend on $\boldsymbol{\theta}$ it is omitted from consideration.

To obtain samples from the posterior distribution $p(\boldsymbol{\theta} | z, M, S)$, we employ the Markov-Chain Monte Carlo method [10]. To reach convergence of the Markov-chain, we observed that roughly 40,000 code samples were required [3]. Ordinarily this would constitute an excessive strain on available computational resources in a CFD context. However, as our experimental data consisted of boundary-layer quantities, we were in a position to use a fast boundary layer code. As such, no real computational bottleneck exists during the calibration phase.

2.2 Prediction

Let Δ be a quantity of interest (QoI) in a particular scenario, which may be a scalar, vector, or functional quantity derived from the flow-state, which can be expressed as

$$\Delta \simeq H_{\Delta} \circ m_{\text{CFD}}(\tilde{S}; M, \boldsymbol{\theta}), \quad (3)$$

in which \tilde{S} represents a flow scenario outside of the set of scenarios used within the calibration phase.

We wish to obtain a stochastic estimate of Δ conditional on a set of models \mathcal{M} and a set of training scenarios \mathcal{S} for the predictive scenario $\tilde{S} \notin \mathcal{S}$. The BMSA methodology offers the evaluation of the posterior predictive distribution (ppd) for Δ in case \tilde{S} conditional on all training data:

$$p(\Delta | \tilde{S}, \mathbf{z}) = \sum_{i=1}^I \sum_{k=1}^K \int_{\boldsymbol{\theta}_i} p(\Delta | \tilde{S}, M_i, \boldsymbol{\theta}) p(\boldsymbol{\theta} | S_k, M_i, z_k) \mathbb{P}(M_i | S_k, z_k) \mathbb{P}(S_k) d\boldsymbol{\theta}. \quad (4)$$

The first term on the right-hand side inside the integral represents the probabilistic equivalent of the simulation results for the QoI given the flow scenario \tilde{S} , a turbulence model M_i and closure coefficients $\boldsymbol{\theta}$. The second and third terms are the posterior probability density distribution of the closure coefficients and the posterior model probabilities respectively, which are the output of the previously conducted calibration procedure summarized in Section 2.1 and detailed in Edeling et al. [3,4]. The last term represents the scenario probabilities. The solution of the multi-dimensional integral over the closure coefficients is the expensive part of the method, because it requires as many code calls as the entries in the Markov-chain, in this case 40,000.

Therefore to obtain a practical estimate of $p(\Delta | \tilde{S}, \mathbf{z})$ we propose to approximate the marginal posterior probability distributions $p(\boldsymbol{\theta} | S_k, M_i, z_k)$ with Dirac- δ functions at their *maximum a posteriori* (MAP) values⁴

$$\boldsymbol{\theta}_{i,k}^{\text{MAP}} := \underset{\boldsymbol{\theta} \in \boldsymbol{\theta}_i}{\operatorname{argmax}} p(\boldsymbol{\theta} | S_k, M_i, z_k) \quad (5)$$

so that

$$p(\boldsymbol{\theta} | S_k, M_i, z_k) \simeq \delta(\boldsymbol{\theta} - \boldsymbol{\theta}_{i,k}^{\text{MAP}}). \quad (6)$$

The effect of this approximation is to neglect the effect of *within-model within-scenario* variance on the ppd. Thus the ppd variance will be reduced, but still include the effect of multiple models and scenarios. Note that if perfectly plentiful data were available in the training scenarios (and the models were able to fit the data exactly for some values of the closure coefficients), then $p(\boldsymbol{\theta} | S_k, M_i, z_k)$ would be δ -functions. So one way to think of this approximation is as neglecting the effect of imperfect training.

Substituting (6) into (4) leads to an approximation of the posterior predictive distribution $p(\Delta | \mathbf{z}) \simeq \hat{p}(\Delta | \mathbf{z})$

$$\begin{aligned} \hat{p}(\Delta | \tilde{S}, \mathbf{z}) &= \sum_{i=1}^I \sum_{k=1}^K \mathbb{P}(M_i | S_k, z_k) \mathbb{P}(S_k) \int_{\boldsymbol{\theta}_i} p(\Delta | \tilde{S}, M_i, \boldsymbol{\theta}) \delta(\boldsymbol{\theta} - \boldsymbol{\theta}_{i,k}^{\text{MAP}}) d\boldsymbol{\theta} \\ &= \sum_{i=1}^I \sum_{k=1}^K \mathbb{P}(M_i | S_k, z_k) \mathbb{P}(S_k) p(\Delta | \tilde{S}, M_i, \boldsymbol{\theta}_{i,k}^{\text{MAP}}) \\ &\stackrel{(a)}{=} \sum_{i=1}^I \sum_{k=1}^K \mathbb{P}(M_i | S_k, z_k) \mathbb{P}(S_k) \delta(\Delta - m_{\text{CFD}}(\tilde{S}; M_i, \boldsymbol{\theta}_{i,k}^{\text{MAP}})). \end{aligned} \quad (7)$$

⁴The MAP estimates are available online at [11]

Equality (a) follows from the fact that the prediction of m_{CFD} is deterministic for deterministic $\boldsymbol{\theta}$. The approximate ppd is therefore a weighted-sum of $I \times K$ δ -functions, one at each prediction of $m_{\text{CFD}}(\tilde{S})$ for each model, and each scenario's MAP-estimate of $\boldsymbol{\theta}$. The cost of evaluating the ppd is $I \times K$ runs of $m_{\text{CFD}}(\tilde{S})$.

The first moment of $\hat{p}(\Delta | \tilde{S}, \mathbf{z})$ can be derived directly from (7). The expectation is

$$\begin{aligned} \mathbb{E}[\Delta | \tilde{S}, \mathbf{z}] &= \int \Delta \cdot \hat{p}(\Delta | \tilde{S}, \mathbf{z}) d\Delta \\ &\stackrel{(a)}{=} \sum_{i=1}^I \sum_{k=1}^K \mathbb{P}(M_i | S_k, z_k) \mathbb{P}(S_k) \int \Delta \cdot \delta \left[\Delta - m_{\text{CFD}} \left(\tilde{S}; M_i, \boldsymbol{\theta}_{i,k}^{\text{MAP}} \right) \right] d\Delta \\ &\stackrel{(b)}{=} \sum_{i=1}^I \sum_{k=1}^K \mathbb{P}(M_i | S_k, z_k) \mathbb{P}(S_k) m_{\text{CFD}} \left(\tilde{S}; M_i, \boldsymbol{\theta}_{i,k}^{\text{MAP}} \right), \end{aligned} \quad (8)$$

where (a) follows from (7) and (b) is the integral-identity encoding the statement that the mean of a deterministic quantity is the quantity itself.

Unlike the posterior model probabilities $\mathbb{P}(M_i | S_k, z_k)$, which are informed using the reference data during the training phase [4], the scenario probabilities $\mathbb{P}(S_k)$ are of a predictive nature. If reference data for the scenario \tilde{S} would be available, $\mathbb{P}(S_k)$ could be optimized accordingly. However, in a predictive setting this data isn't always available. Therefore, $\mathbb{P}(S_k)$ is defined based on model agreement per scenario S_k : If the models show a high level of agreement regarding the value of Δ under a specific scenario S_k , this scenario receives a higher weight compared to other scenarios, in which the models rather disagree. For that we exploit the principle, that if \tilde{S} is similar to S_k the models are expected to give accordingly similar predictions, because $\boldsymbol{\theta}_{i,k}^{\text{MAP}}$ has been calibrated under the same data z_k . As introduced in [4] this principle is modelled by

$$\mathbb{P}(S_k) := \frac{\xi_k^{-p}}{\sum_{j=1}^K \xi_j^{-p}}, \quad \xi_k = \sum_{i=1}^I \|m_{\text{CFD}} \left(\tilde{S}; M_i, \boldsymbol{\theta}_{i,k}^{\text{MAP}} \right) - \mathbb{E}[\Delta | \tilde{S}, z_k]\|_2, \quad (9)$$

in which p serves as a tuning parameter scaling the weighting procedure, i.e. for $p = 0$ all scenarios are equally weighted but for $p \rightarrow \infty$ a single scenario is preferred.

Given all terms in (7) we now want to deduce uncertainty estimates from the ppd. If the full range of $\mathbb{P}(M_i | S_k, z_k) \mathbb{P}(S_k)$ is used, the uncertainty bound is largest and determined by the extreme predictions of the entire ensemble. However, a more reasonable approach is to draw samples from the ppd and use percentiles as min/max levels to obtain a confidence interval of the ppd leading potentially to tight uncertainty bounds. This is done in the following.

3 Incompressible pipe flow at $Re = 44,000$

We apply the BMSA method to turbulent flow in a straight pipe at $Re = 44,000$ with validation data from a DNS simulation [12]. This flow case is, together with turbulent flow over a flat plate and channel flow, one of the three canonical wall-bounded flow types [1]. The main difference between flow in a pipe or channel and over a flat plate, lies in the development of the boundary layer. Assuming an uniform inflow into a circular

straight pipe a boundary layer at the wall develops and its thickness increases downstream similarly to the flat plate case. Forced by the geometry of the pipe the boundary layers from all sides will eventually merge in the middle of the pipe. The distance from the entrance of the pipe to the point, where the boundary layer fills the entire diameter, is known as the entry length. Beyond the entry length the flow is homogeneous in stream-wise direction. In contrast the boundary layer thickness of a flat plate flow with zero pressure gradient is not bounded. However, for all these wall-bounded flow types the physics close to the wall can be assumed to be similar. Therefore, with the application of BMSA based on flat plate flow to pipe flow is a proof-of-concept in order to assess the predictive capabilities of BMSA for wall-bounded flows in general.

3.1 Simulation setup

The validation of the BMSA method is based on a comparison with averaged velocity profiles from a DNS simulation [12]. Therefore, the steady-state RANS simulation mimics the time-averaged DNS simulation by using appropriate boundary conditions given in Table 1. In the DNS simulation the mass flow rate is kept constant via a time-varying adjustment of the pressure gradient in the stream-wise direction [12]. For the RANS simulation the mass flow rate is kept constant via a fixed inlet velocity and the pressure gradient is set to zero at the inlet. At the outlet both the velocity gradients are set to zero and pressure is set to zero. In this way the pressure drop is computed according to the fixed mass flow. For the forward-simulations $m_{CFD}(M_i, \boldsymbol{\theta}_{i,k}^{MAP})$ the results of simulations with nominal values of the coefficients $m_{CFD}(M_i, \boldsymbol{\theta}_{i,k}^o)$ are used as an initial condition in order to reduce iteration counts.⁵

Table 1. The Inlet conditions of velocity U , pressure P/ρ and turbulent quantities.[13]

	\mathbf{U} [m/s]	P/ρ [m ² /s ²]	k	ε	ω	$\tilde{\nu}$
Inlet	$(1, 0, 0)^T$	$\mathbf{n} \cdot \nabla P = 0$	0.00375	0.00835	25.0516	0.00027
Outlet	$\mathbf{n} \cdot \nabla \mathbf{U} = 0$	$(0, 0, 0)^T$	$\mathbf{n} \cdot \nabla k = 0$	$\mathbf{n} \cdot \nabla \varepsilon = 0$	$\mathbf{n} \cdot \nabla \omega = 0$	$\mathbf{n} \cdot \nabla \tilde{\nu} = 0$
Wall	$(0, 0, 0)^T$	$(0, 0, 0)^T$	$(0, 0, 0)^T$	$(0, 0, 0)^T$	8367238.00803	$(0, 0, 0)^T$

The geometry of a straight pipe is determined by its diameter and length. While the diameter D is set in order to achieve the target Reynolds number, the length of the pipe needs to be chosen sufficiently longer than the entry length in order to overcome all effects at the inlet, such as the usage of different closure coefficients or the uniform inflow velocity profile. The entry length is estimated by an empirical relation based on the Reynolds number [14]

$$L_{urb} = 1.359 \cdot D \cdot Re_D^{1/4} \Rightarrow L_{urb} \sim 20D. \quad (10)$$

⁵The template-cases for each turbulence model are available on GitHub: <https://github.com/shmlzr/UQOpenFOAM>

For all simulations a conservative value $L = 200D$ was chosen. Due to the symmetry of the case a wedge-shaped mesh with symmetric boundary conditions orthogonal to the stream-wise direction was chosen[15]. Different meshes were initially studied, with a total number of 2900, 5900 and 9900 cells respectively. The differences for the velocity profiles between the meshes were negligible, but in order to keep the error small the finest mesh was chosen for all forward simulations. In order to study the effect of the turbulence modelling and not additional modelling due to wall-functions the thickness of the smallest cells close to the wall was defined according to $y^+ \leq 1.0$ and Low- Re turbulence models were used.

3.2 BMSA-prediction for radial velocity profile





Fig. 1. BMSA prediction using different smart weighting tuning parameter p .

Applying the BMSA framework to the set of simulations leads to the expectation of the radial velocity $\mathbb{E}(U/U_{bulk}|\mathbf{z})$ as shown in Figures 1(a) to 1(c) for different tun-



(a) L^2 error norm of difference between DNS (b) Scenario-weights for several p -values. De- data and BMSA expectation against p -value. scription of scenarios in [3].

Fig. 2.

ing parameters of the smart-weighting method $p = [0, 1, 2]$. For the confidence interval 10000 samples were drawn from the posterior predictive mass function given in eq. 7 for each x/R position. The 5th and 95th percentile of the population was used as the lower and upper limit respectively, containing 90% of the samples. For $p = 0$, i.e. uniform weighting of the scenarios, the impact of outliers on the confidence interval is large leading to an over-prediction of the modelling error. Outliers are caused by sensitivities of the models with respect to certain coefficients leading to predictions very different compared to the ensemble mean. The effect is already reduced for $p = 1$, for which the confidence interval is covering a tight bound around the expectation. The shape matches with the shape of the majority of the ensemble of forward simulations. For increasing p the confidence interval becomes tighter, since the smart-weighting consecutively increases the weight of $S_k = 1300$ as shown in Figure 2(b) and decreases the influence of the others, reducing the overall variance. Interestingly, the MAP estimates for scenario $S_k = 1300$ are acquired based on a flow over a flat plate with a moderate favourable pressure gradient and a boundary layer close to an equilibrium state [4], which is very similar to the conditions for the flow in a pipe. A comparison of the DNS data and the expectation for different p via the L^2 -norm

$$\|\mathbb{E}[U/U_{bulk}|\tilde{S}, \mathbf{z}] - U_o\|_2 \quad (11)$$

given in Figure 2(a), where the high-resolution DNS data is interpolated by a cubic spline to enable a point-by-point comparison, shows a minimal error for $0 < p < 2$, but also an increase of the error for higher values of p . Note that due to the use of MAP estimates, the variance is missing the *within-model within-scenario* contribution. Therefore, a reduction of the variance with smart-weighting should be done with care in order to not reduce the already underestimated variance even further. Thus, the purpose of the smart-weighting for the BMSA-with-MAP-approach should be to exclude outliers, which can be achieved with relatively low p values. In such a way an optimum

of the width of the confidence interval and the correctness of the expectation can be achieved.




Fig. 3. BMSA prediction of pressure drop including confidence interval for several tuning parameter p of smart-weighting method.

Table 2. Pressure drop over pipe length for Launder-Sharma $k - \varepsilon$, Wilcox (2006) $k - \omega$ and Spalart-Allmaras using nominal closure coefficient [4]. The corresponding Darcy friction coefficient f_{Darcy} is calculated using the Darcy-Weisbach equation. The empirical value for f_{Darcy} is based on the Colebrook equation for $Re = 44,000$ and the corresponding $\Delta P/\rho$ is calculated using the Darcy-Weisbach equation.

Model	$\Delta P/\rho [\frac{m^2}{s^2}]$	f_{Darcy}
Wilcox(2006) $k - \omega$	2.1427	0.02144
Launder-Sharma $k - \varepsilon$	2.0478	0.02049
Spalart-Allmaras	2.2074	0.02208
Empirical value	2.1499	0.021499

3.3 BMSA-prediction for pressure drop

Another relevant quantity of interest for the turbulent pipe flow is the pressure drop ΔP over the pipe length. Table 2 gives the values for the pressure drop $\Delta P/\rho$ normalised by the fluid density ρ and the corresponding Darcy friction coefficient f_{Darcy}

$$f_{Darcy} = \frac{2 \Delta P D}{\rho L U^2} \quad (12)$$

per model. The empirical reference value for the pressure drop is calculated from the Colebrook-White equation, which is an approximation of the empirical Moody-diagram,

$$\frac{1}{\sqrt{f}} = -2\log_{10}\left(\frac{2.51}{Re\sqrt{f}}\right) \quad (13)$$

with zero roughness at $Re = 44,000$. For several tuning parameters of the smart-weighting methods, i.e. $p \in [0, 8]$, the BMSA expectation of the pressure drop $\mathbb{E}[\Delta P | \tilde{S}, \mathbf{z}]$ is shown in Figure 3 as well as the empirical reference data calculated based on the Colebrook-White equation. Up to $p = 4$ the confidence interval still covers a range including also the empirical reference value, which suggests that the method successfully gives a reasonable uncertainty range. For larger tuning parameters the confidence interval reduces drastically and puts all weight on one single scenario, so that the confidence interval is not visible any more. In line with the results for the velocity profile (see Section 3.2), also for the pressure drop $p=2$ is a reasonable choice. The expectation consistently underestimates the true value and is therefore worse than the baseline predictions, see Table 2. However, since the BMSA method wasn't trained on pressure drop data this is not unexpected and the main gain of BMSA are the confidence bounds instead of replacing the baseline prediction with the point-estimate of the expectation.

4 Incompressible flow over periodic hills at $Re_H = 5600$

The periodic hill test case is based on the channel flow case, but with a lower surface modified by a series of periodically arranged hills, and deals with flow separation on the curved surface of the hills and reattachment on the flat plate region between the hills [16]. Being computationally relatively cheap but still challenging in terms of flow physics it has been used in several workshops as a benchmark test case [17]. We use DNS data of the mean flow field from Breuer et al. [16] for $Re_H = 5600$, defined using the streamwise bulk velocity U_b between the hills crest and the upper surface and the hill height H . The test case is known to be especially challenging for linear eddy-viscosity models, which aren't able to predict the mean effect of the unsteady fluctuation of the separation and reattachment points correctly [18]. Thus, the application of BMSA to this challenging flow configuration is a test of the limits of the framework utilizing linear eddy-viscosity models being calibrated for flat plate boundary layer flows.

4.1 Simulation setup

In order to mimic the periodicity of the hill-geometry, periodic boundary conditions were applied at the inlet and outlet and no-slip conditions at the walls. A volume forcing is applied to each cell, which maintains a bulk velocity of $U_b = 1.0$ between the hill's crest and the upper surface. The functional form of the lower surface is defined according to the ERCOFTAC test case description⁶. In order to study the effect of the turbulence modelling and not additional modelling due to wall-functions the thickness of the smallest cells close to the wall were defined according to $y^+ \leq 1.0$ and Low- Re turbulence models were used: Spalart-Allmaras, Launder-Sharma $k - \epsilon$ and Wilcox

⁶Underlying flow regime 3-30, 2D Periodic Hill Flow: <http://qnet-ercoftac.cfms.org.uk>

(2006) $k - \omega$ (for details of these models see [4]). Furthermore, a mesh convergence study was conducted for each model using baseline coefficients and the meshes in table 3 were used for the application of the BMSA method.

Table 3. Mesh size per closure model for flow of periodic hills at $Re_H = 5600$.

	mesh ($n_x \times n_y$)
Spalart-Allmaras	100×110
Launder-Sharma $k - \epsilon$	150×140
Wilcox (2006) $k - \omega$	100×110

Similar to the pipe flow case in Section 3 the results of the forward-simulations $m_{CFD}(M_i, \boldsymbol{\theta}_{i,k}^{MAP})$ with nominal values of the coefficients $m_{CFD}(M_i, \boldsymbol{\theta}_{i,k}^o)$ are used as an initial condition in order to reduce iteration counts. For S_{14} the simulation using Wilcox (2006) $k - \omega$ didn't converge so that this scenario was excluded from the set leading to 13 used scenarios in total.

4.2 BMSA prediction for velocity field

The baseline simulations for the three models show the expected behaviour as reported in the literature for a larger Re -number [18]: while both Spalart-Allmaras and the Wilcox (2006) $k - \omega$ over-predict the size of the recirculation zone, characterised by a zero velocity component close to the wall, the Launder-Sharma $k - \epsilon$ under-predicts this flow feature.

The BMSA approach using MAP estimates of the posterior probability distributions outputs both an expectation of the velocity $\mathbb{E}[\Delta | \tilde{\mathcal{S}}, \mathbf{z}]$, for which $\Delta = U_x$, and a confidence interval. For the latter two types were chosen: one obtained by sampling from the posterior predictive distribution and using the range from the 1.0 to the 99.0 percentile, i.e. 98% confidence interval, and another one by using the min/max range of the ppd. We have chosen both a larger confidence interval and the entire range of the distribution in order to take the expected low performance of the linear eddy-viscosity models for this test case into account. Especially the latter can be seen as an approach trying to envelop the true process by all models and for all scenarios. The smart-weighting technique to obtain the scenario probabilities as described in Section 2.2 was applied along the y/H -direction for each streamwise location for the 98% confidence interval. In that manner the predictive similarity for each scenario per model is evaluated locally for the streamwise direction x .

The BMSA expectation $\mathbb{E}[\Delta | \tilde{\mathcal{S}}, \mathbf{z}]$, as shown in Figure 4, over-predicts the recirculation zone, gives similar velocity profiles for $x = 3.0$ and 4.0 close to the lower surface, but is completely off for other locations. Especially, for $x = 0.0$ the expectation doesn't capture the local maximum of the velocity close to $y/H = 1.0$, i.e. at the hills crest, and over-predicts the velocity within the channel for every other location. Throughout the different locations the expectation shows large differences compared to the DNS for the upper part of the velocity profile.

The entire range of the posterior predictive distribution (ppd) and the 98% confidence interval show large differences over the entire domain as shown in Figure 4. The ppd captures the DNS data for $x \geq 3$, but doesn't capture the local velocity maxima in the area of the free shear layer in the leeward region of $x = 0$ to 2.0. Interestingly for $x \leq 3.0$ when the range of the ppd shrinks locally in y/H -direction also the DNS is still inside. However, this pattern is not the same for every y/H -position, e.g. at $x = 2.0$ the expectation and the DNS match for $0.5 \leq y/H \leq 1$, but the error is high. For the relatively large confidence value 98% the intervals shrink drastically, so that the DNS data is only captured in a small lower band $y/H \leq 1.0$ for $3 \leq x \leq 7$ and for the part of the velocity profile in the middle of the channel.

Interestingly, in the reattachment region between the two hills, where the case is similar to a flat plate, the BMSA approach in the present form with a confidence of 98% is able to capture the DNS data. However, the interval for the upper surface is always negligible but the DNS data doesn't match with the expectation.

The large effect of the min/max profiles, which determine the min/max range of the ppd, to capture the DNS data proves the aforementioned fact that the used linear eddy-viscosity models suffer from restrictions which inhibit the reproduction of the true flow state for this test case. However, posterior probability distributions of the closure coefficients of the models inferred on other flow scenarios more equal to the one here might lead to a different picture.

5 Industrial Test-case: Generic Dassault Falcon

The final test-case we consider is the generic Falcon jet – a business jet including tail, engine nacelles and winglets – at transonic conditions (Industrial challenge IC-03 of UMRIDA). The geometry is far more complex than any considered so far, with a resulting spectrum of primary and secondary flows. We provide this case as a demonstration of the applicability of our approach to problems of relevance in the aerospace industry - unfortunately without a reference solution. LES is impractically expensive here, and experimental data is not available. We attempt to justify the results of BMSA based on the expected response of the simulation to the closure modelling.

The CFD code used is Petrov-Galerkin finite-element RANS solver *AETHER* used within Dassault Aviation. As a preliminary step it was verified that the closure models implemented in *AETHER* were identical to the models for which closure coefficients were calibrated. The implementations of Spalart-Allmaras and Wilcox $k - \omega$ were found to be sufficiently similar to the calibrated models, and in addition are regarded by Dassault Aviation as suitable for this test-case. Of the 14 sets of closure coefficients computed with each model, only 5 cases were able to be successfully converged with $k - \omega$, whereas all S-A cases converged without issue. BMSA requires a minimum of 2 models per scenario, and as such only 5 scenarios could be considered. Thus the spectrum of model results is substantially more limited than for both previous test-cases.

The resulting uncertainty in the pressure distribution at a cut on the wing (with $p = 0$) can be seen in Figures 5(a) and 5(b). The uncertainty is concentrated around the shock on the suction side, and even there is barely visible. Increasing p reduces the uncertainty




Fig. 4. BMSA for streamwise velocity component U_x at several x -locations, $\mathbb{E}[\Delta|\tilde{S}, \mathbf{z}]$ (red), full posterior predictive distribution (ppd) range (shaded grey) and confidence interval of 98% (shaded blue) using $p = 2$. DNS data of Breuer et al. [16] (black).

further. The very low variance can be attributed to two effects: (a) the real lack of sensitivity of the pressure distribution to the turbulent boundary layer in an essentially attached flow, for which an inviscid solution is already satisfactory, and (b) the limited range of models and coefficients used in the study. Indeed by eliminating exactly those coefficients that cause convergence problems (for $k - \omega$), we are introducing additional bias into the sampling, likely biasing against extreme results. The only alternative – of including unconverged solutions – is even less attractive however.






Fig. 5. Pressure distribution at 30% span of the Falcon wing, $y = 3848\text{mm}$.

More insight can be gained by looking at force coefficients: Figure 6 shows total, pressure and viscous drag coefficients, all evaluated by surface integration, with mean and 95% confidence intervals, as a function of p . While pressure- and viscous-drag have approximately the same magnitude, the viscous part completely dominates the uncertainty - with $C_{D,p}$ varying at most 3 counts, and $C_{D,v}$ varying between 15 and 25 counts, depending on p . Thus the uncertainty in total drag is driven entirely by $C_{D,v}$. Given the observed lack of uncertainty in C_p -profiles, and the high sensitivity of friction coefficients to closure modelling, this is not surprising – though it does indicate that varying the closure model does not appear have a significant effect on separation behaviour in this case. Examining the relationship of uncertainty against p reveals which quantities are dominated by model-differences (where uncertainty is approximately constant with p), and for which quantities scenario differences are significant. Here $C_{D,p}$ belongs to the former class, and $C_{D,v}$ to the latter.


In summary, the limited number of turbulence models and range of coefficients lead to limited uncertainty in this case, for quantities of interest related to the pressure, which is essentially governed by the inviscid flow behaviour. Significant uncertainty is observed for viscous quantities such as friction drag. Resolving this in practical applications is the subject of ongoing work.



(a) total



(b) pressure



(c) viscous

Fig. 6. Total, pressure and viscous drag coefficients for the full Falcon.

6 Conclusion

The BMSA approach based on MAP estimates of the posterior probability distributions of the closure coefficient has been applied to three different test cases.

For the flow in a straight pipe at $Re = 44,000$ the results for the velocity profile and the pressure drop are in line with the validation data for this case, which shows in general that the approach can successfully be applied to other test cases outside of the set used for calibration.

The application of BMSA to the periodic hill test case at $Re_H = 5600$ proves that this test case is challenging for the linear eddy-viscosity assumption and also shows the limits of the BMSA framework in the present form, provided that all models in the chosen set employ the Boussinesq hypothesis. However, the resulting ppd is able to capture most of the validation data successfully. A more suitable test case will be evaluated for BMSA in the future, such as the flow over a backward-facing step, for which the separation is forced by the geometry and not result of the simulation itself. Furthermore, the question of how to incorporate model-form error within a stochastic framework for UQ purposes is the topic of ongoing research.

Finally, the Falcon Jet test case (IC-03) revealed the open problem of how to deal with many non-converged solutions within the BMSA methodology and the successive artificial reduction of uncertainty by excluding these simulations from the set of simulations, which is also a topic for further research.

Acknowledgement

We thank Dassault Aviation and especially Gilbert Roge for the collaboration on the test case of the generic falcon jet.

References

1. Pope, S.B.: Turbulent Flows. Cambridge University Press (2000)
2. Slotnick, J., Khodadoust, A.: CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences. Technical report, NASA - National Aeronautics and Space Administration, Langley Research Center Hampton, Virginia 23681-2199 (2014)
3. Edeling, W.N., Cinnella, P., Dwight, R.P., Bijl, H.: Bayesian estimates of parameter variability in the k-epsilon turbulence model. *Journal of Computational Physics* **258** (2014) 73–94
4. Edeling, W.N., Cinnella, P., Dwight, R.P.: Predictive rans simulations via bayesian model-scenario averaging. *Journal of Computational Physics* **275** (2014) 65–91
5. Edeling, W.N.: Quantification of modelling uncertainties in turbulent flow simulations. PhD thesis, Technical University Delft, Aerospace Department (1 2015)
6. Edeling, W.N., Schmelzer, M., Dwight, R.P., Cinnella, P.: Pragmatic estimation of model error of rans turbulence models for industrial relevant flow cases. In preparation.
7. Coles, D., Hirst, E.: Computation of turbulent boundary layers. In: Proceedings of AFOSR-IFP Stanford Conference. Volume 2. (1968)
8. Cheung, S.H., Oliver, T.A., Prudencio, E.E., Prudhomme, S., Moser, R.D.: Bayesian uncertainty analysis with applications to turbulence modeling. *Reliability Engineering & System Safety* **96**(9) (2011) 1137–1149

9. Gelman, A., Carlin, J., Stern, H., Rubin, D.: Bayesian data analysis. Volume 2. Chapman & Hall/CRC Boca Raton, FL, USA (2014)
10. Hastings, W.K.: Monte carlo sampling methods using markov chains and their applications. *Biometrika* **57**(1) (1970) 97–109
11. Edeling, W.N., Cinnella, P., Dwight, R.: MCMC traces of posterior distributions of closure coefficients for 6 turbulence models [online database]. https://www.researchgate.net/publication/308609583_MCMC_traces_of_posterior_distributions_of_closure_coefficients_for_6_turbulence_models_ASCII_format (2016) [Retrieved 25-September-2016].
12. Wu, X., Moin, P.: A direct numerical simulation study on the mean velocity characteristics in turbulent pipe flow. *Journal of Fluid Mechanics* **608** (2008) 81–112
13. Schwarze, R.: CFD Modellierung. Grundlagen und Anwendungen bei Strömungsprozessen. Springer Vieweg (2013) In German.
14. Venkateshan, S.: Mechanical Measurements. 2 edn. John Wiley and Sons Ltd (2015)
15. OpenFOAM: OpenFOAM User Guide. OpenFOAM Foundation. (February 2014)
16. Breuer, M., Peller, N., Rapp, C., Manhart, M.: Flow over periodic hills - Numerical and experimental study in a wide range of Reynolds numbers. *Computers and Fluids* **38**(2) (2009) 433–457
17. Manceau, R.: Report on the 10th joint ERCOFTAC (SIG-15)/ IAHR / QNET-CFD Workshop on Refined Turbulence Modelling. ERCOFTAC Series **57** (2002)
18. Jakirlic, S.: Extended excerpt related to the test case: Flow over a periodical arrangement of 2D hills. Technical Report June (2012)