

Turbo-RANS: straightforward and efficient Bayesian optimization of turbulence model coefficients

Ryley McConkey, Nikhila Kalia, Eugene Yee and Fue-Sang Lien
*Department of Mechanical and Mechatronics Engineering, University of Waterloo,
Waterloo, Canada*

Received 8 December 2023
Revised 11 March 2024
Accepted 11 April 2024

Abstract

Purpose – Industrial simulations of turbulent flows often rely on Reynolds-averaged Navier-Stokes (RANS) turbulence models, which contain numerous closure coefficients that need to be calibrated. This paper aims to address this issue by proposing a semi-automated calibration of these coefficients using a new framework (referred to as turbo-RANS) based on Bayesian optimization.

Design/methodology/approach – The authors introduce the generalized error and default coefficient preference (GEDCP) objective function, which can be used with integral, sparse or dense reference data for the purpose of calibrating RANS turbulence closure model coefficients. Then, the authors describe a Bayesian optimization-based algorithm for conducting the calibration of these model coefficients. An in-depth hyperparameter tuning study is conducted to recommend efficient settings for the turbo-RANS optimization procedure.

Findings – The authors demonstrate that the performance of the k - ω shear stress transport (SST) and generalized k - ω (GEKO) turbulence models can be efficiently improved via turbo-RANS, for three example cases: predicting the lift coefficient of an airfoil; predicting the velocity and turbulent kinetic energy fields for a separated flow; and, predicting the wall pressure coefficient distribution for flow through a converging-diverging channel.

Originality/value – To the best of the authors' knowledge, this work is the first to propose and provide an open-source black-box calibration procedure for turbulence model coefficients based on Bayesian optimization. The authors propose a data-flexible objective function for the calibration target. The open-source implementation of the turbo-RANS framework includes OpenFOAM, Ansys Fluent, STAR-CCM+ and solver-agnostic templates for user application.

Keywords Turbulence modelling, Bayesian optimization, Calibration, Data-driven methods, Reynolds-averaged Navier Stokes, GEKO model

Paper type Research paper

1. Introduction

Turbulent flows are found in many engineering applications, such as automotive, aerospace, chemical, nuclear and wind energy. Engineering activities such as design, optimization, safety assessment, maintenance and failure investigation in these industries frequently require predicting the behaviour of turbulent flows. Predicting turbulent flow is a necessary, but complex problem, that is often addressed using experimental or computational

This work was funded by the National Sciences and Engineering Research Council of Canada (NSERC) and the Tyler Lewis Clean Energy Research Foundation (TLCERF). The authors thank S. Rezaeiravesh (University of Manchester) for useful discussion regarding the presently proposed techniques, and W. Melek (University of Waterloo) for his ongoing feedback on this work. Authors also thank the four anonymous peer reviews for their careful review and insightful comments.



approaches. While experimental approaches often produce the most physically accurate results, they are expensive and limited in the quantities of interest (information) they can provide. For these reasons, predictions obtained using computational fluid dynamics (CFD) are attractive due to their low cost, capability to rapidly explore a (frequently large) design space and ability to access detailed information about practically all aspects of the flow.

Many industrial CFD simulations rely on solving the Reynolds-averaged Navier-Stokes (RANS) equations, which are more computationally affordable than conducting large-eddy simulation (LES) (Witherden and Jameson, 2017). RANS models have been considered an industrial workhorse for the past several decades. However, there are several well-known deficiencies with the application of these methods. A major source of error in RANS is the turbulence model, which often relies on a single-point, linear-eddy-viscosity closure model [e.g. turbulence closure models such as $k-\varepsilon$ (Lauder and Spalding, 1974), $k-\omega$ (Wilcox, 1988), Spalart-Allmaras (Spalart and Allmaras, 1992)]. The past several decades of experience with RANS have resulted in the development of sophisticated turbulence closure models such as the $k-\omega$ shear stress transport (SST) model (Menter, 1994), the $k-\varepsilon-\phi-f$ model (Laurence *et al.*, 2005) and the recently-developed generalized $k-\omega$ (GEKO) model (Menter and Lechner, 2021). However, structural limitations that arise as a result of the assumptions and approximations underpinning single-point RANS closure models affect their predictive performance in various engineering and industrial applications (Wilcox, 1994; Pope, 2000).

Recently, there has been widespread interest in using data-driven approaches to improve and calibrate RANS models (Brunton *et al.*, 2020; Duraisamy *et al.*, 2019). Because of the ability to automate traditional physics-based and heuristic turbulence modelling, more advanced models can be potentially developed. Machine learning has been used to develop complex non-linear eddy viscosity models, which can provide good approximations for highly-resolved LES and direct numerical simulation (DNS) closure terms for use with the RANS modelling framework for incompressible flows (Ling *et al.*, 2016; McConkey *et al.*, 2022; Kaandorp and Dwight, 2020; Wu *et al.*, 2018) and compressible flows (Grabe *et al.*, 2023; Sciacovelli *et al.*, 2023). Additionally, machine learning can be used to generate LES subgrid closure models (Maulik *et al.*, 2019) and accelerate DNS through the application of super-resolution models (Kochkov *et al.*, 2021). While some approaches aim to use highly complex machine learning models to replace or augment the turbulence closure relationship, another promising approach is using reference data to calibrate closure coefficients in an existing turbulence model. Most RANS turbulence closure models contain a number of tuneable coefficients, which can be calibrated to improve the RANS predictive performance for certain classes of flows. The problem of calibrating the turbulence model coefficients so that a prediction is in better conformance with the reference data can be formulated as an optimization problem. To this purpose, a number of previous studies have proposed the use of genetic algorithms (Gimenez and Bre, 2019; Mauro *et al.*, 2017), Nelder-Mead simplex algorithm (Yoder and Orkwis, 2021), optimal Latin hypercube method (Gu *et al.*, 2012) and the global search method (Barkalov *et al.*, 2022) for addressing this problem. Bayesian optimization was demonstrated to be an effective method for calibrating model coefficients in multiscale dynamical systems by Nadiga *et al.* (2019). As discussed by Diessner *et al.* (2022), the Bayesian optimization methodology is particularly promising for addressing many fluid mechanics problems involving model coefficient calibration because it is efficient at optimizing expensive functions. For example, Morita *et al.* (2022) demonstrated the efficient performance of Bayesian optimization for applications to shape optimization and to optimal icing control problems. Additionally, zonal coefficient optimization can lead to improved predictive performance, as demonstrated by the zonal eddy viscosity approach of (Matai and Durbin, 2019).

Naturally, different optimization problems will require different optimization techniques. Some techniques are better suited for high-dimensional problems, noisy problems or problems where the gradient of the objective function is analytically available. In our view, the primary consideration for the turbulence model calibration problem is the computational cost associated with evaluating the objective function. If we formulate the objective function as an error measure between our reference data and the associated predictions of these data for a given set of model coefficients, then every evaluation of the objective function will require a new CFD calculation using the new turbulence model coefficients. The cost of a CFD calculation varies with the specific problem, but for an industrially relevant flow, a single three-dimensional (3D) RANS calculation can easily require on the order of hundreds to thousands of central processing unit (CPU) hours. Therefore, this optimization problem demands a technique that can carefully and efficiently select new query points (namely, new values for the tuneable model coefficients at which the evaluation of the predictive performance will occur next). Bayesian optimization is a popular choice for optimizing computationally expensive objective functions. For example, Bayesian optimization is commonly used in neural network hyperparameter tuning, where each evaluation of the objective function typically involves an expensive training run (Garnett, 2023).

Bayesian approaches have been applied for calibrating turbulence models in a number of previous studies. In particular, Bayesian inference has received attention because it can provide information on uncertainty in a RANS model prediction. Maruyama *et al.* (2021) applied Bayesian inference to infer model coefficients of the Spalart-Allmaras (SA) model for a transonic flow over a wing. SA model coefficients were also inferred for several transonic airfoil problems in an investigation conducted by Da Ronch *et al.* (2020), who found that a properly calibrated set of model coefficients improved the predictive accuracy on transonic test cases outside of the reference data used for the calibration. Ray *et al.* (2018) developed an analytical model for the optimal coefficients of the k - ε turbulence model for a jet in a cross flow and also performed Bayesian calibration of the turbulence model coefficients. In particular, these investigators hypothesized that for many flows, model-form errors in RANS may be significantly smaller than coefficient calibration errors. A similar investigation, involving the calibration of the SA model for a jet flow with a concomitant determination of the uncertainty in the tuned model coefficients, was undertaken by Li *et al.* (2022). Edeling *et al.* (2014) analysed the error in the Launder-Sharma k - ε turbulence model for wall-bounded flows and applied Bayesian inference for calibrating the model coefficients. An earlier investigation conducted by Oliver and Moser (2011) applied Bayesian uncertainty quantification to four eddy-viscosity models. Recently, Fischer *et al.* (2023) coupled a neural network with Bayesian optimization to predict optimal values of coefficients of the GEKO model (Menter and Lehner, 2021) for a trench-shaped film-cooling flow. Another popular approach is the field inversion machine learning approach proposed by Singh *et al.* (2017), which is often applied to infer local values of the calibration coefficients.

In this work, we introduce the tuning parameters using the Bayesian optimization-RANS (turbo-RANS) platform. More specifically, turbo-RANS is a framework for applying Bayesian optimization to the RANS coefficient calibration problem. Bayesian optimization is particularly suited to problems where an evaluation of the objective function is computationally expensive – this is certainly the case for the turbulence model calibration problem in RANS. We formulate an objective function for use with turbo-RANS that is versatile in terms of the types of reference data that may be used. For three different model coefficient calibration problems, we demonstrate that turbo-RANS can be applied to tune turbulence model coefficients in a computationally

efficient manner. Another contribution of the present work is the provision of an open-source implementation of our proposed calibration framework. We illustrate the application of this framework with several example codes. Templates and source code have been made available in a GitHub repository (McConkey, 2023). We also investigate how the hyperparameters associated with the Bayesian optimization can be chosen and, in the course of this investigation, we provide general recommendations for the settings of these various hyperparameters that lead to a computationally efficient and stable calibration of the RANS model coefficients.

A number of previous studies have applied a Bayesian framework to the turbulence model coefficient calibration problem. However, most of these previous studies have focused on the problem as an inference problem. For example, given a set of RANS calculations and some reference data, Bayesian inference can provide a best estimate plus uncertainty of the turbulence closure coefficients. However, this methodology is computationally prohibitive for use with very expensive RANS computations (forward model predictions). Indeed, the computational demands of the RANS model predictions make it practically impossible as a tool for sampling from the posterior distribution using Markov chain Monte Carlo to produce a sequence of dependent draws from the posterior. In the case where one is under the constraint of a very limited computational budget (typically less than about 100 RANS model calculations), we advocate the application of Bayesian optimization to the problem: given previous knowledge of the problem, what is the next most efficient point to query? The use of the Bayesian optimization approach here results in a naturally computationally efficient optimization procedure.

The Bayesian optimization approach allows the exploration and exploitation trade-off to be tightly controlled during optimization. Moreover, the objective function proposed here is flexible and pragmatic and, as a result, is well-suited for use with Bayesian optimization. Whereas many previous studies used simple objective functions (e.g. root-mean-squared [RMS] error), the objective function proposed here produces a well-posed optimization problem and can simultaneously be used with different types of reference data. We demonstrate that using this objective function for calibration leads generally to an accurate prediction of various quantities of interest derived from the turbulent flow. Finally, we believe that the complexity of the implementation and the ease of use (or lack thereof) are often overlooked factors that limit the widespread use of Bayesian calibration. To address these factors, we present a straightforward methodology and provide an open-source implementation that can be immediately used in a variety of CFD contexts.

This paper is organized as follows. In Section 2, we highlight relevant concepts of Bayesian optimization. We then discuss the formulation of the generalized objective function used in the turbo-RANS framework, providing a rationale for this formulation. The turbo-RANS framework is then described in detail along with how this framework is implemented. In Section 3, we apply turbo-RANS to calibrate coefficients in the k - ω SST and GEKO turbulence closure models for three different types of reference data: namely, integral parameter data, dense DNS data and sparse wall measurement data. In Section 4, we perform an extensive study of how to choose the hyperparameters associated with the Bayesian optimization used in the RANS model calibration. Based on this hyperparameter investigation, we provide recommended settings for the hyperparameters that appear in the Bayesian optimization schema, including those used in the Gaussian process regression. Finally, Section 5 provides a conclusion – here, we contextualize the results obtained in this study and provide suggestions for future work.

2. Background and methodology

This section is organized as follows. Sections 2.1–2.3 provide a high-level overview of Bayesian optimization. Section 2.4 explains the proposed generalized error and default coefficient preference (GEDCP) objective function, which we recommend for optimizing turbulence model coefficients. Section 2.5 provides details and a discussion of the turbo-RANS framework. Implementation of this framework is discussed in Section 2.6.

2.1 Bayesian optimization

We summarize the main ideas underpinning Bayesian optimization and Gaussian process regression as they are applied in the turbo-RANS framework. Bayesian optimization is described in detail in Garnett (2023), and Gaussian process regression is explained in detail in Rasmussen and Williams (2006).

The goal of Bayesian optimization is to maximize an objective function $f(x)$ with respect to a vector of parameters x . Since this objective function can be computationally expensive to compute, we create a surrogate model of the objective function that is significantly less expensive to compute. It is useful for the surrogate model to return both the best estimate $\mu(x)$ of the objective function and the uncertainty $\sigma(x)$ in this best estimate. We introduce a utility function $\alpha(x) \equiv \alpha[\mu(x), \sigma(x)]$ that combines the information embodied in the predicted objective function value and uncertainty in this value into a single scalar-valued function (see Section 2.3). The next x value queried is the point that maximizes the utility function $\alpha(x)$. This utility function allows us to control the trade-off between exploration and exploitation. Exploration is the evaluation of the objective function in new regions far away from the currently available data [i.e. points with high uncertainty $\sigma(x)$]. Exploitation is the evaluation of the objective function in regions where it is expected to be favourable based on past data [i.e. points with a high expected value of $\mu(x)$]. The primary advantage of Bayesian optimization is the ability to control this exploration/exploitation trade-off, which produces an optimization procedure that is efficient in terms of the number of objective function evaluations required to find the optimal value of x .

2.2 Gaussian process regression

In Bayesian optimization, Gaussian process regression is typically used to construct a surrogate model of the objective function. Given a data set D consisting of a set of values of the objective function $f(x)$ at points x ($\{(x_1, f(x_1)), (x_2, f(x_2)), [\dots], (x_n, f(x_n))\}$), we can fit a Gaussian process regression model $GP(x)$ and use it as a surrogate model for $f(x)$. A Gaussian process regression model fits a Gaussian process to $f(x')$ for all x' in the domain of definition of the function. The Gaussian process is *conditioned* on the given data D consisting of $f(x)$ measured at points x , so $f(x') \sim N(\mu(x'), k(x, x'; \theta))$ where $N(\mu(x'), k(x, x'; \theta))$ is the Gaussian process with mean function $\mu(x')$ and $k(x, x'; \theta)$ is the covariance function between the pair of points x and x' . The covariance function between the pair of points depends on the hyperparameters θ . A complete description of Gaussian process regression is provided in the book by Rasmussen and Williams (2006).

In this work, the covariance function $k(x, x')$ (suppressing the vector θ of hyperparameters from the notation) used for Gaussian process regression is the Matern kernel, which has the following form:

$$k(x_i, x_j) = \frac{1}{\Gamma(\nu)2^{\nu-1}} \left(\frac{\sqrt{2\nu}}{l_k} d(x_i, x_j) \right)^\nu K_\nu \left(\frac{\sqrt{2\nu}}{l_k} d(x_i, x_j) \right), \quad (1)$$

where $k(x_i, x_j)$ is the covariance of the points x_i and x_j , ν is a kernel parameter (hyperparameter), $\Gamma(\nu)$ is the gamma function evaluated at ν , l_k (hyperparameter) is the

kernel length scale (or characteristic length scale of the Gaussian process), $d(x_i, x_j)$ is the Euclidean distance between x_i and x_j and $K_\nu(\cdot)$ is the modified Bessel function of order ν . More details are available in the scikit-learn Matern kernel documentation (Pedregosa *et al.*, 2011). The critical hyperparameters of the Matern covariance function (kernel) are ν and l_k , which are discussed below.

The behaviour of the Matern kernel depends on the value of the hyperparameter ν . When $\nu = 1/2$, this kernel reduces to the absolute exponential kernel. When $\nu \rightarrow \infty$, the Matern kernel reduces to the squared exponential kernel. For $\nu = 3/2$ and $5/2$, the associated Gaussian process is once- and twice-differentiable, respectively.

The length scale l_k of the Matern kernel is a hyperparameter that needs to be optimized for a given data set D . In this work, we use scikit-learn for the Gaussian process regression, which automatically fits this hyperparameter to the available data. However, an initial guess for l_k needs to be specified, and tuning this initial guess can lead to a more efficient fitting of the Gaussian process which, in turn, improves the subsequent Bayesian optimization performance (Nogueira, 2014). We reparameterize this initial guess in terms of a relative length scale l . The l_{k,c_i} length scale parameter associated with a model coefficient c_i is determined as follows:

$$l_{k,c_i} = l(c_{i,\text{upper}} - c_{i,\text{lower}}), \tag{2}$$

where l_{k,c_i} is an initial guess for the kernel length scale associated with coefficient c_i . Moreover, $c_{i,\text{upper}}$ and $c_{i,\text{lower}}$ are the upper and lower bounds of the domain of definition for the coefficient c_i , respectively. Also, l_k is the length scale vector with components $(l_{k,c_1}, l_{k,c_2}, \dots, l_{k,c_N})$. This vector defines an anisotropic Matern covariance function (kernel) in the sense that the length scale associated with each model coefficient c_i can have a different value. The motivation behind parameterizing l_k as a vector is to specify a scale-appropriate initial guess for l_k (namely, an appropriate scale for each model coefficient). For example, two model coefficients $c_1 \in [0.5, 1.5]$ and $c_2 \in [0.001, 0.005]$ will require two different length scales. Equation (2) will apply the same relative length scale to both dimensions, resulting in scale-appropriate length scales for each dimension. Finally, we note that the hyperparameter vector θ for the Matern covariance function (kernel) has the explicit form $\theta = (\nu, l_k^T)^T$ where the superscript ‘‘T’’ denotes vector transposition.

2.3 Utility functions

2.3.1 Upper confidence bound utility function. The upper confidence bound (UCB) utility function features a clear exploration/exploitation trade-off that is controlled by the κ hyperparameter. The UCB utility function is given by:

$$\alpha_{\text{UCB}}(x) = \mu(x) + \kappa\sigma(x), \tag{3}$$

where $\alpha_{\text{UCB}}(x)$ is the utility function value at point x , $\mu(x)$ is the expected value of the objective function, and $\sigma(x)$ is the standard deviation at point x (namely, $\sigma(x) = k(x, x)^{1/2}$). Both $\mu(x)$ and $\sigma(x)$ are readily available from the Gaussian process regression of the objective function. Larger values of κ lead to more exploration of points in the parameter space with large uncertainty, and smaller values of κ lead to more exploitation of regions in the parameter space where the expected value $\mu(x)$ is known to be large.

2.3.2 Probability of improvement utility function. The probability of improvement (POI) utility function returns the probability that a given point x will result in an improvement of

the objective function value. The POI hyperparameter ξ specifies a minimum POI. The POI utility function is given by:

$$\alpha_{\text{POI}}(x) = \text{CDF} \left(\frac{\mu(x) - \mu^* - \xi}{\sigma(x)} \right), \quad (4)$$

where $\alpha_{\text{POI}}(x)$ is the utility function value at point x , $\text{CDF}(z)$ is the cumulative distribution function of the standard Gaussian distribution (namely, a Gaussian distribution with zero mean and unit variance), $\mu(x)$ is the expected value of the objective function, μ^* is the best objective function value observed so far and $\sigma(x)$ is the standard deviation at x . Again, $\mu(x)$ and $\sigma(x)$ are immediately available from the Gaussian process regression of the objective function. Larger values of ξ lead to a more explorative behaviour, whereas smaller values of ξ lead to a more exploitative behaviour.

2.3.3 Expected improvement utility function. The expected improvement (EI) utility function improves on the POI utility function. Whereas the POI utility function finds the point with the highest probability of providing any improvement, the EI utility function finds the point with the highest EI. The hyperparameter ξ is also used to control the exploration/exploitation trade-off for the EI utility function. The EI utility function has the following form:

$$\alpha_{\text{EI}}(x) = \text{CDF} \left(\frac{\mu(x) - \mu^* - \xi}{\sigma(x)} \right) (\mu(x) - \mu^{\hat{a}} - \xi) + \text{PDF} \left(\frac{\mu(x) - \mu^* - \xi}{\sigma(x)} \right) \sigma(x), \quad (5)$$

where all variable definitions are the same as those of the POI function (Section 2.3.2), and $\text{PDF}(z)$ is the probability density function for a standard Gaussian distribution. As in the case of the POI utility function, smaller ξ values lead to more exploitation, whereas larger ξ values lead to more exploration, in the parameter space.

2.4 Generalized error and default coefficient preference objective function

Objective function design is important for optimization. In general, the purpose of turbo-RANS is to reduce the error in a RANS prediction by adjusting the turbulence model coefficients. Therefore, the objective function should represent errors in the RANS predictions that ensure a well-posed optimization problem. In view of this, we formulate an objective function for use with turbo-RANS and demonstrate its utility.

We formulate a functional form for an objective function, which will be referred to henceforth as the GEDCP function. This function aims to be both highly versatile in terms of the types of reference data that can be used in the optimization and well-posed in terms of the process of the optimization itself. This function computes an error measure that can be used to assimilate dense reference data, sparse reference data or integral reference data (or any combination of these types of data), making it adaptable generally for any kind of calibration problem. For example, RANS coefficients can be calibrated so that the predicted velocity field matches sparse wind-tunnel velocity data or so that the predicted force coefficients match those obtained from experimental data. For unsteady simulations, the coefficients can be calibrated based on an experimental Strouhal number or on a time-averaged force coefficient. The GEDCP function can also combine different types of reference data into a single objective function. For example, coefficients can be calibrated to simultaneously reduce errors in sparse reference data and in reference data for an integral parameter (e.g. lift or drag coefficient).

An optional term is included in the GEDCP function that embeds a preference for the “default” values of the turbulence model coefficients. In practice, including this term serves two goals: namely, it embeds prior knowledge of the problem into the objective function, and it often results in the optimization problem having a single global optimum. For an example that motivates the first goal, consider the prior knowledge that $C_\mu = 0.09$ is commonly used in linear-eddy viscosity models. This prior knowledge may be an important preference that is relevant for some applications. This default value is based on the correct ratio of $\nu_t \varepsilon / k^2$ in a turbulent channel flow away from the wall (Pope, 2000). Here, ν_t is the eddy viscosity, ε is the dissipation rate and k is the turbulent kinetic energy. If the optimized C_μ deviates from this value, there must be a corresponding significant reduction in the misfit error for this state of affairs to be acceptable. This preference penalty must also increase with increasing deviation from the default value of the coefficient, so if the nominally optimal value for C_μ is very far from 0.09, a large reduction in the fitting error to the reference data must occur for this to be acceptable. The second goal of the preference term is to ensure that the optimization problem has a single (unique) solution. In our experience, an error-based objective function based solely on the misfit between the model predictions and the reference data for calibrating RANS model coefficients often exhibits a saturation in the misfit. In other words, if the coefficient exceeds a certain value, the misfit error remains constant and, as a result, the optimization problem does not have a unique solution. Saturation or asymptotic behaviour in the misfit can arise from the nature of the physics in the problem or from the presence of minimum, maximum, hyperbolic tangent and other functions that are used in various combinations in many turbulence models. With a preference for the default value in a model coefficient, a small slope is introduced into the objective function, which drives coefficients towards the default value of the coefficient if this circumstance arises.

The GEDCP objective function, which depends on the model coefficients $\{c_1, c_2, \dots, c_N\}$ that are to be calibrated against the available reference data, is formulated as follows:

$$f_{\text{GEDCP}}(c_1, c_2, \dots, c_N) = -(\lambda_F E_F + \lambda_I E_I)(1 + \lambda_p \hat{p})|_{c_1, c_2, \dots, c_N}, \quad (6)$$

$$E_F(c_1, c_2, \dots, c_N) = \sum_{\phi \in \Phi} \text{MAPE}(\phi)|_{c_1, c_2, \dots, c_N} = \sum_{\phi \in \Phi} \frac{1}{N_\phi} \sum_{i=1}^{N_\phi} \left| \frac{\phi_i - \tilde{\phi}_i(c_1, c_2, \dots, c_N)}{\phi_i} \right|, \quad (7)$$

$$E_I(c_1, c_2, \dots, c_N) = \sum_{\psi \in \Psi} \left| \frac{\psi - \tilde{\psi}(c_1, c_2, \dots, c_N)}{\psi} \right|, \quad (8)$$

$$\hat{p}(c_1, c_2, \dots, c_N) = \frac{1}{N} \sum_{n=1}^N \left| \frac{c_{n, \text{default}} - c_n}{c_{n, \text{default}}} \right|. \quad (9)$$

The variables and terms in equations (6)–(9) are defined below.

In equation (6), $f_{\text{GEDCP}}(c_1, c_2, \dots, c_N)$ is the objective function which depends on the tuneable coefficients c_1, c_2, \dots, c_N ; E_F is the field error; E_I is the integral parameter error; \hat{p} is the default coefficient preference term; and, λ_F , λ_I and λ_p are regularization parameters that control the amount of preference (relative importance) ascribed to each term. The negative sign in equation (6) converts the problem of minimizing E_F , E_I and \hat{p} to a maximization problem, which is the convention used for Bayesian optimization.

The reference data used for calibration of the model coefficients consist of a set of field values (e.g. velocity, pressure) and integral parameters (e.g. lift coefficient). The set of calibration fields is specified by Φ , and the individual calibration fields are denoted by $\phi \in \Phi$. Individual reference values for a given calibration field are denoted by ϕ_i . In a similar manner, the set of reference integral parameters is specified by Ψ , and the individual integral parameters are denoted by $\psi \in \Psi$. As an example, given calibration data consisting of 15 locations with pressure measurements P , 50 locations with measured velocity components U and V , a reference lift coefficient c_l and a reference drag coefficient c_d , we have the following sets: $\Phi = \{P, U, V\}$, $N_P = 15$, $N_U = N_V = 50$ and $\Psi = \{c_b, c_d\}$.

The field error E_F [equation (7)] is a summation of the mean absolute percentage errors (MAPE) in each field ϕ over the set of fields in the set Φ for which reference data are available for calibration. Here, $i = 1, 2, \dots, N_\phi$ indexes the points where reference data corresponding to the field ϕ are available. Furthermore, ϕ_i is the reference datum associated with the field ϕ at a given location and $\hat{\phi}_i(c_1, c_2, \dots, c_N)$ is the predicted value of the field at this location obtained using the set of model coefficients $\{c_1, c_2, \dots, c_N\}$. This term can accommodate arbitrarily sparse and/or dense reference data and involves the determination of MAPE over all available reference data for a given field. Moreover, the use of MAPE implies that the influence of the reference data is not affected by the scale of ϕ_i . For example, a 25% error at a point with low velocity in the boundary layer is penalized equally to that of a 25% error at a point with high velocity. The use of MAPE here is also motivated so that varying scales of $\phi \in \Phi$ are treated equally; for example, the velocity specified in units of m s^{-1} can have a different scale than that associated with the pressure in units of Pa. This equal treatment would not hold true if a mean-squared error or mean-absolute error were to be used in the formulation of the GEDCP objective function.

The integral parameter error E_I [equation (8)] is a summation of the percentage error in each integral parameter for which reference data are available. Here, Ψ is the set of integral parameters. Again, the percentage error is used so that varying scales associated with different quantities of interest within Ψ are treated equally. For example, consider the calibration of an aerodynamic simulation to match reference data for the lift c_l and drag c_d coefficients. A given prediction using model coefficients c_1, c_2, \dots, c_N will provide an estimate for the lift $\tilde{c}_l(c_1, c_2, \dots, c_N)$ and drag $\tilde{c}_d(c_1, c_2, \dots, c_N)$ coefficients. The use of a relative error implies that even though c_l and c_d can have different orders of magnitude, the percentage errors will be penalized equally.

The default coefficient preference term p is given by equation (9). This term can be interpreted as the mean relative deviation of all model coefficients from their corresponding default values. The use of a relative (rather than an absolute) deviation is again motivated by possible scale differences in the values of the model coefficients.


In equation (6), λ_F , λ_I and λ_p can be adjusted to determine the preference for reference field data, reference integral data and for default values of the model coefficients. We recommend setting $\lambda_F = \lambda_I = 1$, and $\lambda_p = 1/2$. These settings retain some preference for the default values of the model coefficients while ensuring that the objective function is primarily influenced by the error-based terms E_F and E_I . If no reference field data is available (namely, $\Phi = \emptyset$), then $\lambda_F = 0$. Similarly, the absence of reference integral parameter data (namely, $\Psi = \emptyset$) implies $\lambda_I = 0$.

The multiplicative design incorporated in equation (6) is to allow deviations from the default values of the model coefficients, as long as a corresponding significant decrease in misfit (prediction) error is realized. If the optimizer begins to “stray” too far from the default values of the model coefficients with only a marginal or negligible reduction in the prediction error, then the default coefficient preference term will penalize this behaviour.

However, if this departure from the default values of the model coefficients results in a significant prediction error reduction, then the objective function will still reward the proposed departure in the optimization process. For example, in the event that the proposed model coefficients result in perfect predictions of the quantities of interest (namely, $E_F, E_1 \rightarrow 0$), then the default coefficient preference term has no effect. As long as the calibration of the model coefficients results in some error in the fitting process (which is expected for the majority of use cases), then the default coefficient preference term will influence the objective function.

As part of the hyperparameter optimization in this work (see Section 4), the GEDCP loss functions for two example problems (see Sections 3.1–3.2) were computed on a dense grid of RANS predictions associated with a grid of values of the model coefficients. This computation allows for the visualization of the GEDCP loss functions for two typical model coefficient calibration problems.

Figure 1 shows the GEDCP loss function associated with the calibration of the a_1 coefficient of the $k-\omega$ SST turbulence closure model. The calibration of this model coefficient was conducted to improve the predictive performance of the turbulence model for the lift coefficient of an airfoil. More details of this example problem are provided in Section 3.1. Over a wide range of values of a_1 , the GEDCP objective function for this case exhibits a single global optimum point and is a relatively straightforward function to optimize. For this example, turbo-RANS is able to find the optimum value of $a_1 = 0.28$ in approximately nine evaluations of the objective function (see Section 3.1 for more detailed results). In the range of $0.34 < a_1 < 0.37$, a small negative slope can be seen in $f_{\text{GEDCP}}(a_1)$. This small negative slope is due to the presence of the default coefficient preference term in the GEDCP objective function. Without the inclusion of this default coefficient preference term, an objective function that involves only a misfit term would have a zero slope here and, as a result, the optimizer would not be “guided away” from the upper bound for a_1 when exploring this plateau in the objective function.



Note: The objective function is computed for a set of 50 RANS predictions with a_1 taking values in the interval $[0.25, 0.37]$

Source: Author's own work

Figure 1.
Objective function
[equation (16)] for the
airfoil case

A more involved example that justifies the inclusion of the default coefficient preference term in the objective function is exhibited in Figure 2. The GEDCP objective function shown here corresponds to optimizing the two model coefficients a_1 and β^* of the $k-\omega$ SST turbulence closure model to improve the predictive performance of the model for a flow over periodic hills. The misfit error used in this example is a summation over the entire flow field because dense DNS data are available for this case. More details on this case are provided in Section 3.2. Again, a single local maximum over a wide range of values of the model coefficients exists in this case. Without the inclusion of the default coefficient preference term, the objective function would asymptotically saturate (namely, achieve a maximum value) along a line defined by $\beta^* \approx 0.08$ in the range $0.36 < a_1^* < 0.60$. With a pure misfit error-based objective function, in spite of the fact that there is virtually no improvement in the predictive accuracy for increasing values of a_1 , the optimizer in this case nevertheless recommends that the best estimate of $a_1 = 0.60$ (i.e. this value for a_1 maximizes the objective function). This large value for a_1 is not justified (or even reasonable) because there is only a marginal improvement in the predictive performance for values of a_1 greater than about 0.34. The GEDCP objective function addresses this problem through the inclusion of a small preference for the default values of the model coefficients. In so doing, a single (unique) solution is obtained for the optimization problem, resulting in reasonable values for the model coefficients: namely, $\beta^* = 0.076$ and $a_1 = 0.34$. This optimum is found in approximately 20 evaluations of the objective function (see Section 3.2 for more results).

2.5 Turbo-Reynolds-averaged Navier-Stokes framework

Figure 3 shows the turbo-RANS framework, along with details of the probe-and-suggest sub-routines. Various steps of the algorithms incorporated in this framework will be described below.

To initialize the optimization process, information must be provided by the user as to which model coefficients are to be optimized. This information includes the default values for the model coefficients and their lower and upper bounds. Hyperparameters (see


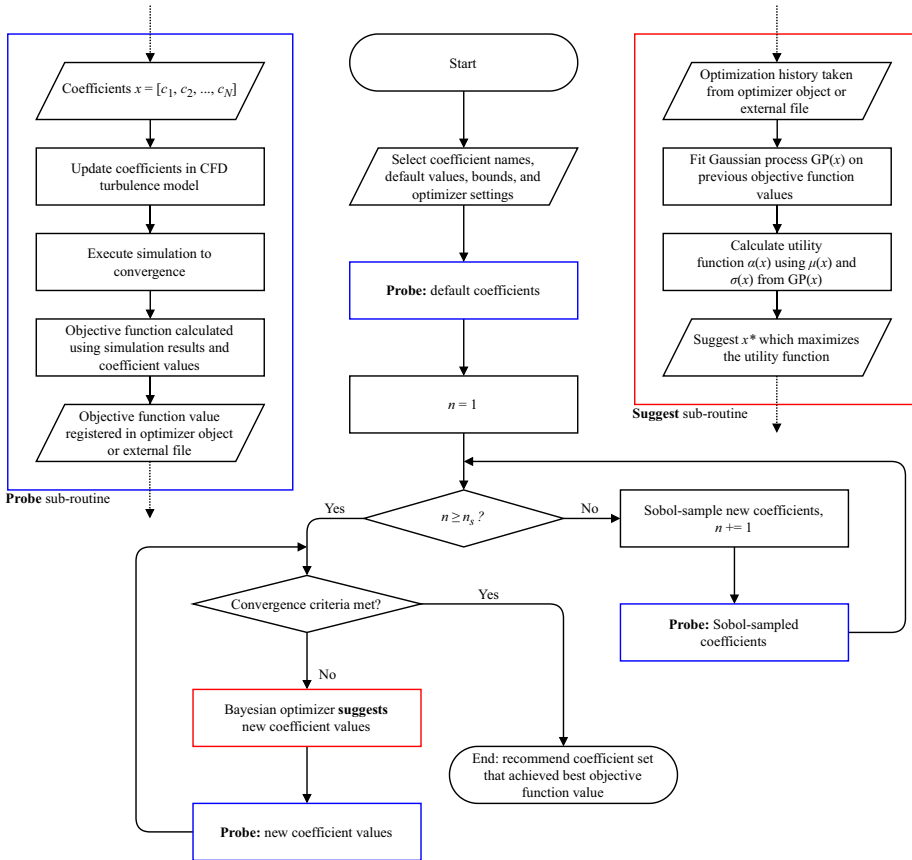


Figure 2.
Objective function
[equation (18)] for the
periodic hills case

Note: The objective function is computed from a set of 900 RANS predictions with $a_1 \in [0.24, 0.60]$ and $\beta^* \in [0.045, 0.140]$
Source: Author's own work



Source: Author's own work

Figure 3. Turbo-RANS framework summarized in an algorithm flowchart

Section 4) used in the optimization, as well as other settings, must also be specified by the user as part of the initialization. It is noted that the hyperparameters can be modified during each iteration of the optimization process. The first point probed is the point corresponding to the default values for the model coefficients. If the chosen optimization problem does not include default values for the model coefficients, then this first default coefficient probe is skipped in the initialization.

After probing the default values for the model coefficients, the sampling loop is executed. In the framework, n_s is the number of sample points that are probed before the Bayesian optimization loop is entered. If n_s is too small, errant behaviour of the Gaussian process regression can occur (see Section 4). At this stage, values of the model coefficients are sampled using Sobol sampling to produce a Sobol sequence of points, which promotes better uniformity of the sampled values compared to a purely random sampling (Sobol, 1967). After a new point is obtained using Sobol sampling, this point is probed. The sampling loop continues until a number of points n_s (prespecified by the user) have been probed.

The Bayesian optimization loop begins after a default point and a sufficient number of additional sample points have been probed. In the Bayesian optimization loop, the optimizer considers the past history and suggests a new value of the model coefficients (namely, a point) to be probed. Following this, the suggested new point is probed. The probe-suggest loop continues until a convergence criterion is satisfied. This convergence criterion is specified by the user. We envision that turbo-RANS loops that involve relatively inexpensive CFD calculations can be stopped when the best objective function value and/or the best set of model coefficients do not provide any further improvements (changes) in the optimization process after a certain number of iterations (early stopping). For problems corresponding to more expensive CFD calculations, an upper bound on the computational budget can be prescribed *a priori* by the user (e.g. repeat the loop until a pre-specified number of CPU hours has been expended in the computation). With either the convergence-based or budget-based stopping criteria, the Bayesian optimization algorithm will result in an efficient optimization in terms of the number of objective function evaluations (namely, CFD calculations). After the stopping criteria are satisfied, the final recommended model coefficient set is the set that achieves the maximum value of the objective function.

Probing a set of values of the model coefficients consists of running a RANS calculation with these coefficients, computing the objective function and registering the objective function value so that it can be used in future iterations. The probe sub-routine in [Figure 3](#) details these steps. The convergence criteria for the RANS calculations are also specified (prescribed) by the user. We recommend the standard practice of evaluating both the residual values of various flow quantities and monitoring an integral parameter to assess convergence. In terms of actual implementation, this process is typically automated using scripts that modify the turbulence model coefficients, execute a RANS calculation and monitor the convergence of this calculation. To reduce computational demands, we recommend using the converged flow fields obtained from the default coefficient probe (namely, the default values of the model coefficients) as initial conditions for future probe RANS calculations. This re-use of the converged flow field from the default coefficient probe greatly reduces the computational costs of the turbo-RANS algorithm, since the flow fields start near convergence. The post-processing of RANS calculations into fields that are compatible with the GEDCP objective function (Section 2.4) is also typically automated. For example, this post-processing can include computing force coefficients or sampling the solution field at the locations where the reference data are available. Details of the objective function registration are provided in Section 2.6, and in the turbo-RANS documentation ([McConkey, 2023](#)).

The suggest sub-routine exhibited in [Figure 3](#) addresses two sub-optimization problems: namely, fitting a Gaussian process using all previously probed points obtained in modelling of the objective function; and, maximizing the utility function. These optimization problems are already implemented in the dependencies for turbo-RANS. The suggest sub-routine executes almost instantaneously, and therefore, the main computational expense in the turbo-RANS algorithm is conducting a RANS calculation with a given set of turbulence model coefficients. The computational time of the suggest sub-routine increases with the number of points previously sampled, as the cost of the Gaussian process regression increases with an increasing number of points. However, in our experience, the computational cost of the suggest sub-routine is negligible compared to the cost of conducting the RANS calculations required for each point probed (undertaken in the probe sub-routine). For large parameter spaces and a large number of previously evaluated points defining the objective function hypersurface (e.g. greater than 100), the Gaussian process regression may perform slowly. However, even in the worst case, we have observed that the cost of the suggest sub-routine remains below 1% of the total cost associated with the RANS calculations required for each probed point.

2.6 Implementation

The core Bayesian optimization package used in turbo-RANS is the Python library Bayesian-optimization (Nogueira, 2014). This library is widely used in many Bayesian optimization applications and includes implementations of the UCB, POI and EI utility functions. The underlying Gaussian process regression library used in Bayesian-optimization is from scikit-learn (Pedregosa *et al.*, 2011). A series of convenient Python interfaces to these two libraries has been provided in the turbo-RANS GitHub repository (McConkey, 2023). The repository includes examples and templates for a wide variety of use cases, including the following:

- two examples based on OpenFOAM: flow over an airfoil (Section 3.1), and flow over periodic hills (Section 3.2);
- one example based on Ansys Fluent: flow through a converging-diverging channel (Section 3.3);
- an example script for STAR-CCM+ users;
- templates for optimizing OpenFOAM simulations;
- solver-agnostic templates for optimizing CFD simulation Python scripts; and
- solver-agnostic templates for optimizing CFD simulation shell scripts.

Full documentation for the code is provided in the GitHub repository (McConkey, 2023). It should be noted that the parallelization of Bayesian optimization is non-trivial. While a RANS calculation with a given set of model coefficients is typically parallelized, the primary turbo-RANS algorithm is currently restricted to running in a sequential manner. This restriction means that only a single RANS calculation is conducted at a time, rather than probing multiple sets of model coefficients to enable multiple RANS calculations to be undertaken simultaneously. It is assumed that the RANS calculation step is by far the most computationally expensive component of the turbo-RANS algorithm since the acquisition of a new point in the Bayesian optimization occurs nearly instantaneously. If desired, the user can opt to implement a parallelized Bayesian optimization by modifying the turbo-RANS example scripts. The code's verbose file input and output structure facilitates implementing parallelized Bayesian optimization but, for now, this functionality is left for a future effort.

3. Demonstrations

3.1 Example 1: Flow over an airfoil

3.1.1 Description. RANS calculations are often used in external aerodynamics, such as in aerospace, wind engineering and automotive engineering. Accurate prediction of integral quantities such as force coefficients is typically desired in these applications. The purpose of this example is to demonstrate how turbo-RANS can be used to calibrate turbulence model coefficients using integral parameter reference data.

3.1.2 Computational setup. We calculate the steady-state, incompressible, turbulent flow over a NACA 0012 airfoil at a 10.12° angle of attack and a chord-based Reynolds number of $Re_c = 6 \times 10^5$. The case details are taken from the NASA turbulence modelling resource "2D NACA 0012 Airfoil Validation" 2DN00 case (Rumsey, 2021). OpenFOAM v2206 was used – the RANS calculations were conducted using the semi-implicit method for pressure-linked equations (SIMPLE) solver (simpleFoam) in OpenFOAM. A second-order upwind scheme was used for the discretization of the convective terms in the momentum equation, and the central difference scheme was used for the discretization of the diffusion terms. A first-order upwind scheme was used for the discretization of the convective terms in the turbulence

transport equations. The geometric agglomerated algebraic multigrid solver was used to obtain the pressure, and the preconditioned bi-conjugate gradient Stab solver was used to solve all the other equations. The full OpenFOAM case setup is available in the turbo-RANS GitHub repository (McConkey, 2023).

The turbulence model used for this case is the k - ω SST model. The OpenFOAM v2206 variant of this model is the updated 2003 Menter version (Menter *et al.*, 2003), with production terms from Menter and Esch (2001) (OpenCFD Ltd, 2019). The following model equations are provided in Menter *et al.* (2003) (k is the turbulent kinetic energy and ω is the specific dissipation rate):

$$\frac{\partial(\rho k)}{\partial t} + \frac{\partial(\rho U_i k)}{\partial x_i} = \tilde{P}_k - \beta^* \rho k \omega + \frac{\partial}{\partial x_i} \left[(\mu + \sigma_k \mu_t) \frac{\partial k}{\partial x_i} \right], \quad (10)$$

$$\frac{\partial(\rho \omega)}{\partial t} + \frac{\partial(\rho U_i \omega)}{\partial x_i} = \alpha \frac{\tilde{P}_k}{\nu_t} - \beta \rho \omega^2 + \frac{\partial}{\partial x_i} \left[(\mu + \sigma_\omega \mu_t) \frac{\partial \omega}{\partial x_i} \right] + 2(1 - F_1) \rho \sigma_{\omega 2} \frac{1}{\omega} \frac{\partial k}{\partial x_i} \frac{\partial \omega}{\partial x_i}, \quad (11)$$

$$\nu_t = \frac{a_1 k}{\max(a_1 \omega, SF_2)}, \quad (12)$$


where the definitions of the blending functions F_1 and F_2 , and limited production term \tilde{P}_k have been omitted as they are not relevant to the present discussion. Given that the coefficients are a blend of the coefficients of the k - ε and k - ω turbulence closure models, the k - ω SST model has ten model coefficients: a_1 , β^* , σ_{k1} , σ_{k2} , α_1 , α_2 , β_1 , β_2 , $\sigma_{\omega 1}$ and $\sigma_{\omega 2}$. Default values for these coefficients are given as: $a_1 = 0.31$, $\beta^* = 0.09$, $\sigma_{k1} = 0.85$, $\sigma_{k2} = 1$, $\alpha_1 = 5/9$, $\alpha_2 = 0.44$, $\beta_1 = 3/40$, $\beta_2 = 0.0828$, $\sigma_{\omega 1} = 0.5$ and $\sigma_{\omega 2} = 0.856$. Here, the subscript one denotes the k - ω model, and the subscript two denotes the k - ε model. The default values of these coefficients arise from the application of various heuristics, such as tuning the behaviour of certain terms to be in conformance with the characteristics of specific types of (usually canonical) flows. For many industrially relevant flows, additional tuning of these model coefficients can increase predictive accuracy. For the airfoil case, we use turbo-RANS to calibrate the a_1 coefficient of the k - ω SST turbulence model for this particular flow.

Figure 4 shows the computational domain and mesh, which were taken from the NASA turbulence modelling resource (Rumsey, 2021). The computational domain is relatively large. More specifically, the domain is $500c$ (c is the chord length) in any direction from the airfoil so as to negate blockage effects. Pressure at the inlet is zero gradient, and all other flow quantities at the inlet are prescribed as follows:

$$U_{\text{inlet}} = (50.68, 0, 9.05) \text{ m s}^{-1}, \quad (13)$$

$$k_{\text{inlet}} = 0.001075 \text{ m}^2 \text{ s}^{-2}, \quad (14)$$

$$\omega_{\text{inlet}} = 0.3279 \text{ s}^{-1}. \quad (15)$$



Notes: (a) Computational domain and boundary names; (b) computational mesh; (c) close-up of computational mesh on the airfoil surface

Source: Author's own work

Figure 4.
Computational setup
for the airfoil example
case

In particular, turbulence quantities at the inlet are based on a 0.052% turbulence intensity and a turbulent length scale of 10% of the chord length. At the outlet plane of the computational domain, all quantities except pressure are zero-normal-gradient, and the gauge pressure is fixed at zero Pa.

The mesh for this case consists of 57,344 structured hexahedral cells, with approximately 129 points along the surface of the airfoil. With the default values for $k\text{-}\omega$ SST model coefficients, the y^+ wall-normal distance from the airfoil surface has values lying between 0.017 and 0.88, with a mean value of 0.30. These values imply that the RANS computation for this case is wall-resolved.

3.1.3 Optimization problem. The purpose of the airfoil example case is to optimize the a_1 model coefficient in the $k\text{-}\omega$ SST turbulence model to improve the accuracy of the model for prediction of the lift coefficient c_l . The objective function for the airfoil case is given by:

$$f_{\text{GEDCP}}(a_1) = -\left(E_1\right)\left(1 + \frac{1}{2}p\right)\Bigg|_{a_1}, \quad (16)$$

$$E_1(a_1) = \left|\frac{c_l - \tilde{c}_l(a_1)}{c_l}\right|. \quad (17)$$


Here, $f_{\text{GEDCP}}(a_1)$ is the GEDCP objective function [equation (6)] with $\lambda = 1$, and $\lambda_b = 1/2$. This objective function primarily aims to minimize the relative error in the prediction of the integral parameter c_l while retaining a minor preference for the default value of the model coefficient; namely, $a_1 = 0.31$. In the present work, we have computed this objective function over a grid of values of a_1 so that it can be visualized (cf. Figure 1). No field quantities are used in the objective function for the airfoil case. The parameter space for the model coefficient for the airfoil case is $a_1 \in [0.25, 0.37]$. The coefficient a_1 was calibrated so that the predictions matched the reference lift coefficient of $c_l = 1.0707$, the latter of which comes from the fully-turbulent experimental measurement for this quantity obtained by Ladson (1988) at a Mach number of 0.15, where the suction and pressure side transition points were fixed at a distance of $0.05c$ from the leading edge, using No. 80-W grit carborundum strips.

3.1.4 Results. The turbo-RANS framework was applied to the airfoil example case, allowing a maximum of ten iterations in the optimization process. The first iteration was the default coefficient run (namely, RANS calculation using the default value for a_1). Subsequently, four additional Sobol-sampled iterations were completed. Therefore, five of these ten iterations correspond to iterations of the Bayesian optimization procedure. Each iteration corresponds to a RANS calculation with a given value for the a_1 model coefficient. The computational time was greatly reduced by using the converged fields from the default coefficient run as initial fields for the remaining RANS calculations involving the new proposals for the value of a_1 . Figure 5 displays the convergence history over the ten iterations.

In nine iterations, the optimization process attained a maximum value of the objective function, and the value of a_1 corresponding to this maximum (optimum) objective function value yielded a predicted value for the lift coefficient of $\tilde{c}_l = 1.0706$. This predicted value for the lift coefficient is in excellent conformance with the experimental value for the lift coefficient of $c_l = 1.0707$. The airfoil example calibration results are summarized in Table 1. While the default coefficient performs reasonably well at predicting the experimental value for the lift coefficient, this example demonstrated that turbo-RANS can efficiently calibrate the model coefficient so that the predicted value of the lift coefficient provides an essentially perfect agreement with the corresponding experimental measurement of this coefficient.

3.2 Example 2: Periodic hills

3.2.1 Description. Predicting the evolution of separated, turbulent flow is a major challenge for RANS turbulence models (Wilcox, 1994). Non-equilibrium effects, non-local



Source: Author's own work

Figure 5. Convergence history for the optimization process for the airfoil example case

Table 1. Summary of the airfoil lift coefficient calibration case

Case	a_1	c_l
Experimental calibration data (Ladson, 1988)	–	1.0707
$k-\omega$ SST, default coefficient	0.31	1.0847
$k-\omega$ SST, turbo-RANS optimized	0.28	1.0706
Source: Author's own work		

effects and history effects are frequently present for this type of flow. Unfortunately, commonly used RANS turbulence models such as the $k-\omega$ SST model often do not accurately predict these flow phenomena. Nevertheless, these models are frequently applied to separated turbulent flows since separation occurs in many industrial flows. A canonical example of separation is turbulent flow over periodic hills. In this flow, a periodic domain is used so that the flow repeatedly separates from the smooth surface of one hill, reattaches, accelerates up a slope of the hill immediately downstream of it and then separates again from the surface of this second hill. When we apply single-point RANS turbulence closure models such as the $k-\omega$ SST model to this type of flow, we can calibrate the model coefficients so that the effects of the poorly captured physics are at least mimicked. Otherwise, the turbulence model will be unable to accurately predict the general characteristics of the flow. In this example, we demonstrate how turbo-RANS can be used to calibrate turbulence model coefficients so that the predicted field quantities are in better conformance with the associated reference data.

3.2.2 Computational setup. We calculate the steady-state, incompressible, turbulent flow over a series of periodic hills at a height-based Reynolds number of $Re_H = 5,600$. DNS data and OpenFOAM meshes are provided for this flow by [Xiao et al. \(2020\)](#). The software, solver and schemes are identical to the airfoil example case (cf. Section 3.1.2). The $k-\omega$ SST turbulence model is also calibrated in this example [equations (10)–(12)].

The computational domain and mesh for this case are taken from [Xiao et al. \(Xiao et al., 2020\)](#). The computational domain is two-dimensional (2D) – the 3D DNS data have been averaged along the spanwise direction. The mesh consists of 14,751 hexahedral cells. Boundary conditions and mesh convergence for the RANS calculation, in this case, are discussed in [McConkey et al. \(2021\)](#). In short, periodic boundary conditions are imposed at the inlet and outlet planes of the computational domain. A no-slip wall boundary condition is imposed on the top and bottom walls of the domain. Along these no-slip walls, $0.01 \leq y^+ \leq 0.80$, with an average value for y^+ along the top and bottom walls of 0.71 and 0.20, respectively. This range of values for y^+ along the top and bottom walls implies that the RANS calculations conducted here are wall-resolved. [Figure 6](#) shows the computational domain and mesh for the periodic hills case.

3.2.3 Optimization problem. For the case of the periodic hills, the GEDCP objective function is used, but flow field quantities are used in the misfit error term of the function. The purpose is to calibrate the a_1 and β^* model coefficients of the $k-\omega$ SST turbulence model so that the predictions of the magnitude of the velocity \bar{U} and k fields are in better conformance with the available DNS reference data. No integral parameters are included in the objective function for this example. In consequence, the objective function assumes the following form:




Figure 6.
Computational setup
for the periodic hills
example case

Notes: (a) Computational domain and boundary names; (b) computational mesh

Source: Author's own work

$$f_{\text{GEDCP}}(a_1, \beta^*) = -\left(E_{\text{F}}\left(1 + \frac{1}{2}b\right)\right)\Bigg|_{a_1, \beta^*}, \tag{18}$$

$$\begin{aligned} E_{\text{F}}(a_1, \beta^*) &= \sum_{\phi \in \{|\vec{U}|, k\}} \text{MAPE}(\phi) \\ &= \sum_{\phi \in \{|\vec{U}|, k\}} \frac{1}{N_{\phi}} \sum_{i=1}^{N_{\phi}} \left| \frac{\phi_i - \tilde{\phi}_i(a_1, \beta^*)}{\phi_i} \right|. \end{aligned} \tag{19}$$

This objective function encapsulates a trade-off between reducing the MAPE in the prediction of $|\vec{U}|$ and k while retaining a minor preference for the default value of the model coefficients: namely, $a_1 = 0.31$ and $\beta^* = 0.09$. The parameter space corresponding to the model coefficients for the case of the periodic hills is $a_1 \in [0.24, 0.60]$ and $\beta^* \in [0.045, 0.14]$. The DNS reference data have been interpolated onto the RANS mesh so that $N_{\phi} = 14,751$. Given that every point in the RANS calculation has corresponding reference data, this example demonstrates how turbo-RANS can be used in an application involving highly dense reference data.

3.2.4 Results. A fixed computational budget of 30 iterations was prescribed for the optimization. The default values of the model coefficients were used for the first iteration. For the next nine iterations, Sobol sampling was used to generate the nine additional points. Subsequently, Bayesian optimization was used for the remaining 20 iterations.

Figure 7 illustrates the convergence of the optimization process for the periodic hills example case. After approximately 20 iterations, the Bayesian optimization has found the maximum (optimum) value (dashed line corresponding to the maximum value obtained using a 30×30 grid search in the parameter space). At the end of the optimization process, Bayesian optimization found model coefficients whose associated objective function value was slightly larger (better score) than the maximum (best score) obtained in the grid search – naturally, this optimal value occurs between the grid points used in the exhaustive grid search.

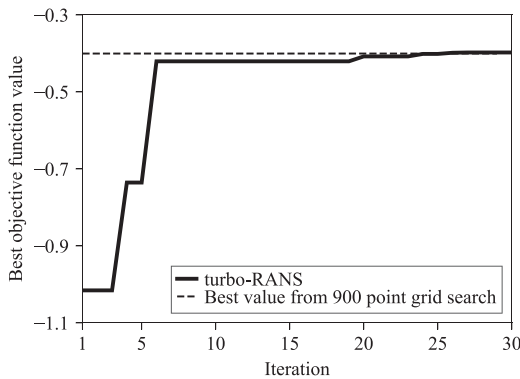




Figure 7. Convergence of the optimization process for the case of the periodic hills

Source : Author's own work

The predicted velocity and turbulent kinetic energy fields obtained using RANS calculations based on the default and optimized model coefficients were compared to the DNS reference data. Figures 8 and 9 display the streamwise U and vertical V components of the velocity along several vertical sampling lines in the computational domain, respectively. The RANS predictions obtained using the turbo-RANS-optimized model coefficients agree very well with the DNS data, demonstrating a significant improvement in comparison to the corresponding RANS predictions obtained using the default values for the model coefficients. The default values of the model coefficients lead to predictions that overestimate the flow separation and, as a result, do not accurately predict the velocity after the occurrence of the separation. In the bulk flow over the separated flow region, the velocity predicted by the RANS models using the default values of the model coefficients is also high compared to the DNS reference data. After calibrating the a_1 and β^* model coefficients, the velocity here agrees well with the DNS reference data. Both the default and optimized RANS predictions exhibit discrepancies in the near-wall quantities directly along the separated bump. Here, the breakdown of assumptions made in modelling the RANS boundary layers is likely responsible for this discrepancy. Overall, the inclusion of the DNS U and V reference data through the GEDCP objective function [equation (18)] leads to a greatly



Source: Author's own work



Source: Author's own work

Figure 8.
Streamwise (U)
component of velocity
along various vertical
sampling lines for the
case of the periodic
hills

Figure 9.
Vertical (V)
component of velocity
along various vertical
sampling lines for the
case of the periodic
hills

improved conformance between the RANS predictions and the DNS reference data for this separated flow, as summarized in Table 2.

Figure 10 compares vertical profiles of the turbulent kinetic energy k along various sampling lines in the computational domain. The k field was used in conjunction with the $|\vec{U}|$ field in the calibration of the model coefficients. While there is a clear improvement in the k estimation in many regions of the flow, there is still a discrepancy between the DNS reference k and the RANS predicted k in the shear zone at the top of the separation region. Further accuracy improvements may be possible through the inclusion of additional model coefficients in the calibration procedure, such as the various coefficients that appear in the k -transport equation [equation (10)] and the ω -transport equation [equation (11)]. Nevertheless, the calibration procedure has generally resulted in an improved prediction of k for this separated flow, as summarized in Table 2.

While the reference data used in the GEDCP objective function were $|\vec{U}|$ and k , it was of interest to determine whether other flow quantities were also predicted more accurately. Even though the wall shear stress was not included as an explicit calibration target, Figure 11 shows that the wall shear stress along the bottom wall is in much better conformance with the DNS data for the wall shear stress (not used in the calibration) after calibration of the model coefficients. In fact, the baseline k - ω SST turbulence model does not predict any reattachment of the flow along the bottom wall before the flow reaches the

Table 2.
Values of MAPE (lower is better) before and after turbo-RANS calibration of the model coefficients in the k - ω SST turbulence model for flow over periodic hills

Coefficients	Mean absolute percentage error (MAPE)	
	k - ω SST, default coefficients	k - ω SST, turbo-RANS
$\{a_1, \beta^*\}$	{0.31, 0.09}	{0.34, 0.076}
$ \vec{U} $	0.522	0.120
k	1.285	0.850

Source: Author's own work




Figure 10.
Turbulent kinetic energy k along various vertical sampling lines for the case of the periodic hills

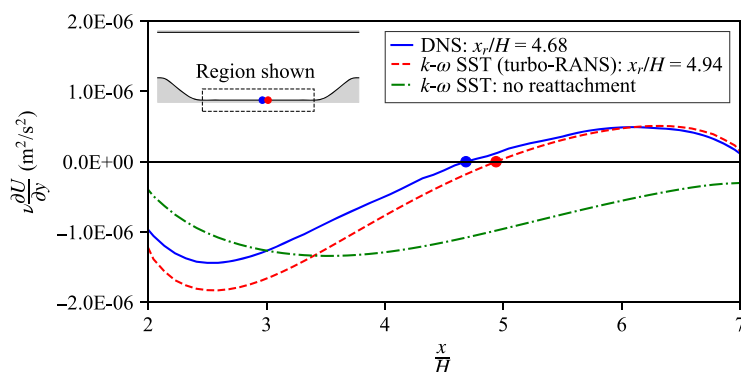
Source: Author's own work

downstream hill. After calibration, the $k-\omega$ SST turbulence model now predicts reattachment, with the reattachment point in very good agreement with the DNS data. Figure 11 demonstrates that calibrating turbulence model coefficients based on a specific set of flow quantities can lead to improvements in the predictive accuracy of other flow quantities not used in the calibration.

3.3 Example 3: Converging-diverging channel

3.3.1 Generalized $k-\omega$ turbulence model. The latest turbulence model that has been developed by Ansys, Inc. is the “generalized $k-\omega$ ” (GEKO) two-equation model. The aim of this model is to provide flexibility for the user to fine-tune different coefficients for a RANS computation. GEKO has several parameters that can be fine-tuned to increase predictive accuracy. There are six parameters that can be tuned by the user. In this paper, the focus is on fine-tuning the model coefficients C_{SEP} and C_{NW} to produce more accurate predictions for flow in a converging-diverging channel. The exact equations by which these GEKO coefficients modify the underlying two-equation turbulence closure model are not provided by Ansys, Inc.; however, general guidance for their use is provided in the Ansys User Manual. The parameter C_{SEP} modifies the characteristics of a separation in the flow. According to the Ansys User Manual, “increasing C_{SEP} reduces eddy viscosity, which leads to more sensitivity to adverse pressure gradients for boundary layers and lower spreading rates for free shear flow” (Menter and Lechner, 2021). A change in the value of C_{SEP} has an influence on all classes of flows. The parameter C_{NW} modifies the near-wall characteristics of the flow that affect the nature of the wall boundary layer. The parameter C_{NW} generally has little or no effect on free shear flows. An increase in the value of C_{NW} leads to a larger wall shear stress. Here, we focus on calibrating global values for the GEKO model coefficients, although local values can also be accommodated through the application of user-defined functions. These coefficients have a specified (recommended) range of values, which is documented in the Ansys User Manual (Menter and Lechner, 2021). In accordance with this documentation, the model coefficients C_{SEP} and C_{NW} should take values in the range summarized in Table 3.

3.3.2 Computational setup. We calculate the steady-state, incompressible, turbulent flow through a converging-diverging channel using the GEKO turbulence closure model in a RANS framework. The DNS Reynolds number for this flow, based on U_{max} and h , is



Source: Author's own work

Figure 11.
Wall shear stress
along the bottom wall
for the case of the
periodic hills

$Re_{U_{\max}} = 12,800$. Here, U_{\max} is the maximum velocity in the fully-developed channel cross-section. The computational domain used for the 2D RANS calculation is shown in Figure 12. The DNS reference data for this flow were obtained from Marquillie *et al.* (2008, 2011). The mesh is adapted from McConkey *et al.* (2021), and contains 610,720 grid cells as displayed in Figure 13. Mesh independence with the selected mesh density near the converging section was demonstrated in McConkey *et al.* (2021).

Boundary conditions and fluid properties for the RANS calculation are prescribed to match those of the DNS reference simulation. The inlet plane of the computational domain is placed at a distance of $20h$ upstream of the converging section, where h is the half-channel height. A Neumann pressure boundary condition is imposed at the inlet plane (namely, the gradient of the pressure in the streamwise direction is zero). The uniform streamwise velocity is prescribed at the inlet with a value of 0.845 m s^{-1} . This value was chosen to match the mass flow rate from the original RANS computation undertaken by McConkey *et al.* (2021). The inlet plane turbulent kinetic energy value is set to $k = 4.28421 \times 10^{-4} \text{ m}^2 \text{ s}^{-2}$, and specific dissipation rate at the inlet plane is $\omega = 0.26993 \text{ s}^{-1}$, in conformance with the computations conducted by


Table 3.

Recommended range of values for the two GEKO coefficients, C_{SEP} and C_{NW}

Parameter	Minimum value	Maximum value	Default value
C_{SEP}	0.7	2.5	1.75
C_{NW}	-2.0	2.0	0.50


Source: Author's own work

Figure 12. Computational domain for the converging-diverging channel



Source: Author's own work

Figure 13. Mesh for the converging-diverging channel (McConkey *et al.*, 2021)



Source : Author's own work

McConkey *et al.* (2021). A fixed gauge pressure of zero Pa is imposed at the outlet plane of the computational domain, where all other flow quantities are zero-gradient. The outlet plane is placed at a distance of $20h$ downstream of the bump. Finally, a no-slip boundary condition was applied on all walls in the computational domain.

The fluid properties of the flow were specified as follows: density $\rho = 1 \text{ kg m}^{-3}$ and kinematic viscosity $\nu = 7.9365 \times 10^{-5} \text{ m}^2 \text{ s}^{-1}$. The SIMPLE-consistent (SIMPLEC) algorithm was used for the pressure-velocity coupling. Convective terms in the momentum equation were discretized with a second-order upwind scheme. The convective terms in the k - and ω -transport equations were discretized with a first-order upwind scheme. The diffusion terms in the momentum and turbulence transport equations were discretized with a second-order central difference scheme.

3.3.3 Optimization problem. For the converging-diverging channel case, we use the GEDCP objective function to calibrate the model coefficients based on sparse reference data. Specifically, we calibrate C_{SEP} and C_{NW} based on a sparse set of measurements of the pressure coefficient along the bottom wall of the channel. The pressure coefficient is determined as follows:

$$C_p = \frac{P - P_0}{\frac{1}{2}\rho U_{\text{max}}^2}. \quad (21)$$

Here, P is the pressure and P_0 is the pressure taken at a point downstream of the converging-diverging channel (approximately at a downstream distance of $12h$ from the start of the converging section).

No integral parameters are included in the objective function. The objective function used here is given by:

$$f_{\text{GEDCP}}(C_{\text{SEP}}, C_{\text{NW}}) = -(E_{\text{F}}) \left(1 + \frac{1}{2}p \right) \Bigg|_{C_{\text{SEP}}, C_{\text{NW}}}, \quad (22)$$

$$E_{\text{F}}(C_{\text{SEP}}, C_{\text{NW}}) = \text{MAPE}(C_p)$$

$$= \frac{1}{N_{C_p}} \sum_{i=1}^{N_{C_p}} \left| \frac{C_{p,i} - \tilde{C}_{p,i}(C_{\text{SEP}}, C_{\text{NW}})}{C_{p,i}} \right|. \quad (23)$$

Ten points along the bottom wall have been selected as sparse wall measurements ($N_{C_p} = 10$) used as the reference data in the calibration of the model coefficients. These points are indicated by the blue scatter points in [Figure 14](#). The misfit error term in the GEDCP objective function is the MAPE in C_p over the ten reference data points. For the upper and lower bounds of C_{SEP} and C_{NW} , we use the Ansys recommended limits of $0.75 \leq C_{\text{SEP}} \leq 2.5$ and $-2 \leq C_{\text{NW}} \leq 2$, as discussed previously in Section 3.3.1.

3.3.4 Results. A Python script was created to modify and execute an Ansys Workbench journal file. A fixed computational budget of 30 iterations was used. The first iteration corresponds to the default values of the model coefficients. Subsequently, nine additional Sobol-sampled iterations are executed. The remaining 20 iterations use the Bayesian optimization schema. After 30 iterations, the recommended values for C_{SEP} and C_{NW} are those that maximize the objective function: $C_{\text{SEP}} = 0.7$, and $C_{\text{NW}} = 1.778$.

Figure 14 compares the DNS data to the predictions obtained using the default and turbo-RANS-optimized GEKO models. After optimization, the streamwise profile of the pressure coefficient is in better conformance with the DNS data. In particular, in the diverging section of the channel, the turbo-RANS-optimized GEKO model better predicts the pressure recovery, in contrast to the default GEKO model which is seen to significantly under-predict the recovery. The optimized model exhibits a slightly greater error in its prediction of the minimum pressure at the top of the converging section (namely, at $x/h \approx 5$), but achieves greater predictive accuracy in every other region of the flow.

3.4 Discussion

The preceding three examples demonstrate the flexibility of the GEDCP objective function to work with various reference data types for the calibration of turbulence closure models. A mixture of integral, dense and sparse reference data were used in these examples. The GEDCP objective function also permits the use of a mixture of reference data types for a given problem. For example, the airfoil case can be modified to include surface pressure coefficient reference data for use in calibration. As demonstrated by the case of the periodic hills, improving predictions of specific quantities of interest often results in improved predictions of other flow quantities (not used as reference data for the calibration). Moreover, when the GEDCP objective function is used as a calibration target in turbo-RANS, a computationally efficient optimization of turbulence model coefficients is possible. Sobol sampling, which occurred until iteration five for the airfoil case and iteration ten for the periodic hills and converging-diverging channel cases, proved to be a highly effective sampling methodology. During the Sobol sampling period, a point close to the optimal value was almost always identified, and therefore, the Bayesian optimization algorithm was well-positioned to determine whether to exploit this region or to explore the more sparsely sampled (albeit more uncertain) regions in the parameter space. For more complex flows, there is no guarantee that the Sobol sampling algorithm will happen to provide a point near the optimal value, and exploration within the Bayesian optimization loop may become more important. Section 4 will investigate and provide recommendations for the number of points to be sampled in the initialization phase of Bayesian optimization.

4. Hyperparameter optimization

Complete details of the hyperparameter optimization study are provided in the supplementary information files. In summary, the purpose of the hyperparameter




Figure 14. Pressure coefficient along the bottom wall for the converging-diverging channel case

Source: Author's own work

optimization procedure was to determine optimal settings for Bayesian optimization for typical turbo-RANS objective functions. The hyperparameter search spaces are summarized in Table 4, with each hyperparameter combination being evaluated using 30 different random seeds. Each evaluation consisted of a “simulated optimization run”, which used a set of pre-determined RANS calculations, rather than conducting a new RANS calculation for each coefficient suggestion. A run was considered convergent based on proximity to the optimal solution from the pre-determined set of RANS calculations. In total, 7,110 turbo-RANS optimization loops were simulated. The recommended hyperparameters for turbo-RANS are summarized in Table 5.

The optimal hyperparameter settings likely change on a case-by-case basis, since the highly customizable GEDCP objective function varies depending on the type of reference data used in the calibration and on the default coefficient preference. Nevertheless, the two relatively different objective functions [equations (16) and (18)] have similar optimal hyperparameters. Moreover, the UCB utility function is relatively insensitive to κ for these typical objective functions. As the objective function increases in complexity and the number of optimized coefficients increases, the optimal value of κ changes slightly but remains close to the recommended default value of $\kappa = 2.576$ in the Bayesian-optimization library. This value of κ represents a reasonable exploration and exploitation trade-off for a general case, and the optimal turbo-RANS setting of $\kappa = 2$ indicates that a slight preference for exploitation is favourable for turbo-RANS. For different flows, settings in the GEDCP

Hyperparameter	Values tested	Notes
Utility function	UCB, EI, POI	See Section 2.3 for more details
κ (UCB)	$0.5 \leq \kappa \leq 6$	Linear spaced interval, increment of $\Delta\kappa = 0.1$ (56 values)
ξ (EI, POI)	$0 \leq \xi \leq 0.5$	Logarithmic spacing, with 56 values
ν	1 3 5 $\frac{1}{2}, \frac{3}{2}, \frac{5}{2}, \dots, \infty$	See Section 2.2 for more details
l	$0.01 \leq l \leq 1.0$	14 values total, more concentrated around 0.1, then every 0.1 to 1.0
n_s	$3 \leq n_s \leq 15$	13 values total, every integer

Source: Author’s own work

Table 4.
Hyperparameter
search space

Hyperparameter	Recommended setting	Notes
Utility function	Upper confidence bound (UCB)	Expected improvement (EI) also performs well, but use $\xi \lesssim 10^{-3}$
κ	2.0	Minor efficiency gains can be made by lowering κ , but at the risk of becoming stuck
ν	5 $\frac{5}{2}$	$\nu = \frac{3}{2}$ also performs well
l	0.1	Since this is just the initial guess for the Gaussian process optimizer, the choice for l is often not significant
n_s	10	For very simple optimization problems (e.g. single variable, under 10 iterations total), n_s can be lowered, at the risk of becoming stuck

Source: Author’s own work

Table 5.
Recommended
hyperparameter
settings

function can depend on the number of coefficients in the problem, so the core optimization problem can vary greatly. While we demonstrated here that the recommended setting (UCB, $\kappa = 2$) performs well for different problems, it may not perform well for all cases. The Bayesian optimization hyperparameters should be re-tuned if the convergence speed is unsatisfactory for a different objective function.

5. Conclusion

This work focused on an important and widespread optimization problem in engineering flow simulation: calibrating turbulence model coefficients. The major contributions of this work are to propose a framework for Bayesian optimization of turbulence model coefficients (turbo-RANS), propose a general objective function, demonstrate the predictive accuracy improvements that can be achieved and recommend hyperparameters for future users of turbo-RANS. We presented the turbo-RANS framework, demonstrated how RANS models can be calibrated using various types of reference data to produce more accurate predictive results and performed a systematic hyperparameter tuning for the Bayesian optimization process.

Even with perfectly calibrated model coefficients, there is an upper limit to the predictive accuracy that can be achieved with RANS. Errors associated with unavoidable assumptions and approximations (model errors) in RANS can severely degrade the accuracy of these predictions. For example, the assumptions of locality can be violated for many real-world flows (Pope, 2000). Duraisamy *et al.* (2019) list four main sources of error in RANS: errors arising from ensemble averaging (L1), errors in the functional representation of the Reynolds stress (L2), errors in the model functional form (L3) and errors in the model coefficients (L4). It is possible that for certain classes of flows, the calibration of model coefficients (which reduces the L4 error) can help offset some L2 and L3 errors. However, this effect is difficult to quantify. Nevertheless, even with perfectly calibrated model coefficients, turbo-RANS is only as good as the “best RANS” solution for a particular problem, which is often limited by unrecoverable L1 and L2 errors. Although turbo-RANS can potentially improve the accuracy of a RANS prediction to some upper bound limit in their predictive accuracy for a given flow, this upper bound limitation in the accuracy will always exist due to implicit modelling assumptions and approximations used in RANS, which may not hold true for many industrially relevant (complex) flows.

Within the wider context of modern machine learning techniques for RANS, we believe that an important future research direction is comparing coefficient calibration to highly sophisticated data-driven algebraic Reynolds stress closure models. Although the process of calibrating model coefficients is comparatively simpler than training a non-linear eddy viscosity model, the resulting performance in the present work appears to compete with these machine learning models. Previous results for the periodic hills case by Wu *et al.* (2018), McConkey *et al.* (2022) and Kaandorp and Dwight (2020) show similar predictive accuracy to the presently calibrated linear-eddy viscosity models. We demonstrated here that simply by tuning model coefficients, there is significant performance to be leveraged from the physics represented in the RANS turbulence closure models. Even further performance can be leveraged from the Bayesian optimization of zonal coefficient values. Currently, Bayesian inference has been applied to infer local turbulence model coefficients, but Bayesian optimization applied to local or zonal model coefficients has not been explored to the authors’ knowledge.

Despite years of optimization methods being applied to turbulence model calibration, and despite the significant accuracy improvements that are possible, these methods are not yet widely used in industrial settings. We believe this may be due to the complicated

implementation and expert background required to perform Bayesian calibration. Therefore, by designing an easy-to-use software package, we hope to help spur industrial interest in this technique. All data and code from the present work can be found in the turbo-RANS GitHub repository (McConkey, 2023). We also include the 900 periodic hills RANS calculations used for the present hyperparameter tuning investigation, as broader investigations concerned with uncertainty quantification may find this data set useful.

Future work will include model coefficient calibration for more challenging flows. The presently proposed GEDCP objective function can be immediately applied to quantify error in an unsteady quantity of interest, such as a Strouhal number, time-averaged force coefficient or RMS force coefficient. We also aim to explore optimal turbo-RANS hyperparameter tuning for more complex 3D flows and investigate calibrating a RANS turbulence model on a mixture of experimental and numerical reference data. Turbo-RANS can also be extended to calibrate coefficients within a Reynolds-stress transport model, or other turbulence closure relations. Ultimately, the algorithm and GEDCP objective function proposed here can be applied to calibrate any coefficient of interest in a RANS model.

Currently, CFD best practices such as a mesh independence study, careful numerical scheme selection, designing meshes appropriate for wall treatment, verification and validation are missing an important step: coefficient calibration. Previously, this would have been a largely manual activity. However, turbo-RANS automates this activity in a highly computationally efficient manner. We envision that, along with validation, calibration can become a more frequently used practice in RANS. With the more widespread use of closure coefficient calibration, future RANS practitioners can potentially generate a coefficient database and make an educated selection of coefficients to use in their specific application based on the nature of the physics present in their flow. To some extent, this is the idea behind the GEKO turbulence closure model (Menter and Lechner, 2021). However, the GEKO model still requires user tuning of the model coefficients. With sufficient reporting and the use of calibrated RANS models, a recommendation engine (e.g. a large language model) can be embedded into CFD software of the future. This model can be trained to recommend turbulence model coefficients given a description of the flow physics present and quantities of interest. Such a model relies on more reporting of calibrated closure coefficients, which we aim to make easier with turbo-RANS.

References

- Barkalov, K., Lebedev, I., Usova, M., Romanova, D., Ryazanov, D. and Strijhak, S. (2022), "Optimization of turbulence model parameters using the global search method combined with machine learning", *Mathematics*, Vol. 10 No. 15, p. 2708, doi: [10.3390/math10152708](https://doi.org/10.3390/math10152708). ISSN: 2227-7390, available at: www.mdpi.com/2227-7390/10/15/2708
- Brunton, S.L., Noack, B.R. and Koumoutsakos, P. (2020), "Machine learning for fluid mechanics". In", *Annual Review of Fluid Mechanics*, Vol. 52 No. 1, pp. 1-32, doi: [10.1146/annurev-fluid-010719-060214](https://doi.org/10.1146/annurev-fluid-010719-060214). issn:0066-4189.
- Da Ronch, A., M., Panzeri, J., Drofelnik, R. and d'Ippolito, R. (2020), "Sensitivity and calibration of turbulence model in the presence of epistemic uncertainties", *CEAS Aeronautical Journal*, Vol. 11 No. 1, pp. 33-47, doi: [10.1007/s13272-019-00389-y](https://doi.org/10.1007/s13272-019-00389-y), available at: <http://link.springer.com/10.1007/s13272-019-00389-y>, ISSN: 1869-5582.
- Diessner, M., O'Connor, A., Wynn, S., Laizet, Y., Guan, K., Wilson, R. and Whalley, R.D. (2022), "Investigating Bayesian optimization for expensive-to-evaluate black box functions: application in fluid dynamics". In", *Frontiers in Applied Mathematics and Statistics*, Vol. 8, p. 1076296, doi: [10.3389/fams.2022.1076296](https://doi.org/10.3389/fams.2022.1076296). ISSN: 2297-468, available at: www.frontiersin.org/articles/10.3389/fams.2022.1076296/full

- Duraisamy, K., Iaccarino, G. and Xiao, H. (2019), "Turbulence modeling in the age of data", *Annual Review of Fluid Mechanics*, Vol. 51 No. 1, pp. 357-377, doi: [10.1146/annurev-fluid-010518-040547](https://doi.org/10.1146/annurev-fluid-010518-040547), ISSN: 0066-4189, available at: www.annualreviews.org/doi/10.1146/annurev-fluid-010518-040547
- Edeling, W.N.P., Cinnella, R.P. and Bijl, H. (2014), "Bayesian estimates of parameter variability in the k-turbulence model", *Journal of Computational Physics*, Vol. 258, pp. 73-94, doi: [10.1016/j.jcp.2013.10.027](https://doi.org/10.1016/j.jcp.2013.10.027), ISSN: 00219991, available at: <https://linkinghub.elsevier.com/retrieve/pii/S0021999113007031>
- Fischer, L., James, D., Jeyaseelan, S. and Pfitzner, M. (2023), "Optimizing the trench shaped film cooling design", In: *International Journal of Heat and Mass Transfer*, Vol. 214, p. 124399, doi: [10.1016/j.ijheatmasstransfer.2023.124399](https://doi.org/10.1016/j.ijheatmasstransfer.2023.124399), ISSN: 00179310, available at: <https://linkinghub.elsevier.com/retrieve/pii/S0017931023005458>
- Garnett, R. (2023), *Bayesian Optimization*, Cambridge University Press, Cambridge.
- Gimenez, J.M. and Bre, F. (2019), "Optimization of RANS turbulence models using genetic algorithms to improve the prediction of wind pressure coefficients on low-rise buildings", In: *Journal of Wind Engineering and Industrial Aerodynamics*, Vol. 193, p. 103978, doi: [10.1016/j.jweia.2019.103978](https://doi.org/10.1016/j.jweia.2019.103978), ISSN: 01676105, available at: <https://linkinghub.elsevier.com/retrieve/pii/S0167610519305781>
- Grabe, C., Jäckel, F., Khurana, P. and Dwight, R.P. (2023), "Data-driven augmentation of a RANS turbulence model for transonic flow prediction", In: *International Journal of Numerical Methods for Heat Fluid Flow*, Vol. 33 No. 4, doi: [10.1108/HFF-08-2022-0488](https://doi.org/10.1108/HFF-08-2022-0488), ISSN: 0961-5539.
- Gu, Z., Song, X., Jiang, Y. and Gong, X. (2012), "Optimization of the realizable k - Turbulence model especially for the simulation of road vehicle", doi: [10.4271/2012-01-0778](https://doi.org/10.4271/2012-01-0778).
- Kaandorp, M.L.A. and Dwight, R.P. (2020), "Data-driven modelling of the Reynolds stress tensor using random forests with invariance", *Computers and Fluids*, Vol. 202, p. 104497, doi: [10.1016/j.compfluid.2020.104497](https://doi.org/10.1016/j.compfluid.2020.104497), ISSN: 00457930, available at: <https://doi.org/10.1016/j.compfluid.2020.104497> <https://linkinghub.elsevier.com/retrieve/pii/S0045793020300700>
- Kochkov, D., Smith, J.A., Alieva, A., Wang, Q., Brenner, M.P. and Hoyer, S. (2021), "Machine learning accelerated computational fluid dynamics", pp. 1-13, available at: <http://arxiv.org/abs/2102.01010>
- Ladson, C.L (1988). "Technical Memorandum 4074: effects of Independent Variation of Mach and Reynolds Numbers on the Low-Speed Aerodynamic Characteristics of the NACA 0012 Airfoil Section", Tech. rep. NASA.
- Lauder, B.E. and Spalding, D.B. (1974), "The numerical computation of turbulent flows", "In", *Computer Methods in Applied Mechanics and Engineering*, Vol. 3 No. 2, pp. 269-289, doi: [10.1016/0045-7825\(74\)90029-2](https://doi.org/10.1016/0045-7825(74)90029-2), ISSN: 00457825, available at: <https://linkinghub.elsevier.com/retrieve/pii/0045782574900292>
- Laurence, D.R., Uribe, J.C. and Utyuzhnikov, S.V. (2005), "A robust formulation of the v_2 -f model", *Flow, Turbulence and Combustion*, Vol. 73 Nos 3/4, pp. 169-185, doi: [10.1007/s10494-005-1974-8](https://doi.org/10.1007/s10494-005-1974-8), ISSN: 13866184.
- Li, J., Chen, S., Cai, F., Wang, S. and Yan, C. (2022), "Bayesian uncertainty analysis of SA turbulence model for supersonic jet interaction simulations", *Chinese Journal of Aeronautics*, Vol. 35 No. 4, pp. 185-201, doi: [10.1016/j.cja.2021.07.039](https://doi.org/10.1016/j.cja.2021.07.039), ISSN: 10009361, available at: <https://linkinghub.elsevier.com/retrieve/pii/S1000936121003332>
- Ling, J., Kurzwski, A. and Templeton, J. (2016), "Reynolds averaged turbulence modelling using deep neural networks with embedded invariance", *Journal of Fluid Mechanics*, Vol. 807, pp. 155-166, doi: [10.1017/jfm.2016.615](https://doi.org/10.1017/jfm.2016.615), ISSN: 14697645.
- McConkey, R. (2023), "turbo-RANS", available at: <https://github.com/rmconke/turbo-rans> (accessed November, 2023).
- McConkey, R., Yee, E. and Lien, F.-S. (2021), "A curated dataset for data-driven turbulence modelling". In", *Scientific Data*, Vol. 8 No. 1, pp. 1-14, doi: [10.1038/s41597-021-01034-2](https://doi.org/10.1038/s41597-021-01034-2), ISSN: 20524463.
- McConkey, R., Yee, E. and Lien, F.-S. (2022), "Deep structured neural networks for turbulence closure modeling", *Physics of Fluids*, Vol. 34 No. 3, p. 35110, doi: [10.1063/5.0083074](https://doi.org/10.1063/5.0083074), ISSN: 1070-6631.

- Marquillie, M., Laval, J.P. and Dolganov, R. (2008), "Direct numerical simulation of a separated channel flow with a smooth profile", *Journal of Turbulence*, Vol. 9, pp. 1-23, doi: [10.1080/14685240701767332](https://doi.org/10.1080/14685240701767332), ISSN: 14685248.
- Marquillie, M., Ehrenstein, U. and Laval, J.P. (2011), "Instability of streaks in wall turbulence with adverse pressure gradient", *Journal of Fluid Mechanics*, Vol. 681, pp. 205-240, doi: [10.1017/jfm.2011.193](https://doi.org/10.1017/jfm.2011.193), ISSN: 14697645.
- Maruyama, D., Bekemeyer, S., Görtz, S., Coggon, S. and Sharma, S. (2021), "Data-driven Bayesian inference of turbulence model closure coefficients incorporating epistemic uncertainty", *Acta Mechanica Sinica*, Vol. 37 No. 12, pp. 1812-1838, doi: [10.1007/s10409-021-01152-5](https://doi.org/10.1007/s10409-021-01152-5), ISSN: 16143116.
- Matai, R. and Durbin, P.A. (2019), "Zonal eddy viscosity models based on machine learning", *Flow, Turbulence and Combustion*, Vol. 103 No. 1, pp. 93-109, doi: [10.1007/s10494-019-00011-5](https://doi.org/10.1007/s10494-019-00011-5), ISSN: 15731987.
- Maulik, R., San, O., Rasheed, A. and Vedula, P. (2019), "Subgrid modelling for two-dimensional turbulence using neural networks", *Journal of Fluid Mechanics*, Vol. 858, pp. 122-144, doi: [10.1017/jfm.2018.770](https://doi.org/10.1017/jfm.2018.770), ISSN: 14697645.
- Mauro, S., Lanzafame, R., Messina, M. and Pirrello, D. (2017), "Transition turbulence model calibration for wind turbine airfoil characterization through the use of a Micro-Genetic algorithm", *In: International Journal of Energy and Environmental Engineering*, Vol. 8 No. 4, pp. 359-374, doi: [10.1007/s40095-017-0248-2](https://doi.org/10.1007/s40095-017-0248-2), ISSN: 2008-9163, available at: <http://link.springer.com/10.1007/s40095-017-0248-2>
- Menter, F.R. (1994), "Two-equation eddy-viscosity turbulence models for engineering applications", *AIAA Journal*, Vol. 32 No. 8, pp. 1598-1605, doi: [10.2514/3.12149](https://doi.org/10.2514/3.12149), ISSN: 00011452.
- Menter, F.R. and Esch, T. (2001), "Elements of industrial heat transfer predictions", 16th Brazilian Congress of Mechanical Engineering.
- Menter, F.R. and Lechner, R. (2021), "Best Practice: generalized k-(GEKO) Two-Equation turbulence modeling in ANSYS CFD".
- Menter, F.R., Kuntz, M. and Langtry, R. (2003), "Ten years of industrial experience with the SST turbulence model", *Turbulence, Heat and Mass Transfer*, Vol. 4, pp. 625-632.
- Morita, Y.S., Rezaeiravesh, N., Tabatabaei, R., Vinuesa, K., Fukagata, P. and Schlatter, P. (2022), "Applying Bayesian optimization with Gaussian process regression to computational fluid dynamics problems", *Journal of Computational Physics*, Vol. 449, p. 110788, doi: [10.1016/j.jcp.2021.110788](https://doi.org/10.1016/j.jcp.2021.110788), ISSN: 10902716.
- Nadiga, B.C., Jiang, D. and Livescu, D. (2019), "Leveraging Bayesian analysis to improve accuracy of approximate models". *Journal of Computational Physics*, Vol. 394, pp. 280-297, doi: [10.1016/j.jcp.2019.05.015](https://doi.org/10.1016/j.jcp.2019.05.015), ISSN: 00219991, available at: <https://linkinghub.elsevier.com/retrieve/pii/S0021999119303407>
- Nogueira, F. (2014), "Bayesian optimization: open source constrained global optimization tool for Python", available at: <https://github.com/fmfn/BayesianOptimization>
- Oliver, T.A. and Moser, R.D. (2011), "Bayesian uncertainty quantification applied to RANS turbulence models", *Journal of Physics: Conference Series*, Vol. 318 No. 4, p. 42032, doi: [10.1088/1742-6596/318/4/042032](https://doi.org/10.1088/1742-6596/318/4/042032), ISSN: 1742-6596, available at: <https://iopscience.iop.org/article/10.1088/1742-6596/318/4/042032>
- OpenCFD Ltd (2019), "OpenFOAM: user guide v2006", available at: www.openfoam.com/documentation/guides/latest/doc/
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M. and Duchesnay, E. (2011), "Scikit-learn: machine learning in python", *Journal of Machine Learning Research*, Vol. 12, pp. 2825-2830.
- Pope, S.B. (2000), "Turbulent flows", ISBN: 9780521598866.

- Rasmussen, C.E. and Williams, C.K.I. (2006), *Gaussian Processes for Machine Learning*, MIT Press, London.
- Ray, J.L., Dechant, S., Lefantzi, J., Ling, S. and Arunajatesan, S. (2018), "Robust Bayesian calibration of a k-model for compressible jet-in-Crossflow simulations", *AIAA Journal*, Vol. 56 No. 12, pp. 4893-4909, doi: [10.2514/1.J057204](https://doi.org/10.2514/1.J057204), ISSN: 0001-1452.
- Rumsey, C. (2021), "NASA Langley research Center turbulence modeling resource", available at: <https://turbmodels.larc.nasa.gov/>
- Sciacovelli, L., Cannici, A., Passiatore, D. and Cinnella, P. (2023), "A priori tests of turbulence models for compressible flows", In: *International Journal of Numerical Methods for Heat Fluid Flow*, doi: [10.1108/HFF-09-2023-0551](https://doi.org/10.1108/HFF-09-2023-0551), ISSN: 0961-5539.
- Singh, A.P., Duraisamy, K. and Zhang, Z.J. (2017), "Augmentation of turbulence models using field inversion and machine learning", 55th AIAA Aerospace Sciences Meeting, *Reston, VT*, American Institute of Aeronautics and Astronautics, ISBN: 978-1-62410-447-3, doi: [10.2514/6.2017-0993](https://doi.org/10.2514/6.2017-0993), available at: <https://arc.aiaa.org/doi/10.2514/6.2017-0993>
- Sobol, I.M. (1967), "On the distribution of points in a cube and the approximate evaluation of integrals", *USSR Computational Mathematics and Mathematical Physics*, Vol. 7 No. 4, pp. 86-112, doi: [10.1016/0041-5553\(67\)90144-9](https://doi.org/10.1016/0041-5553(67)90144-9), ISSN: 00415553, available at: <https://linkinghub.elsevier.com/retrieve/pii/0041555367901449>
- Spalart, P.R. and Allmaras, S.R. (1992), "A one-equation turbulence model for aerodynamic flows", 30th Aerospace Sciences Meeting and Exhibit Vol. 1, *Reston, Virginia*, American Institute of Aeronautics and Astronautics, pp. 5-21, doi: [10.2514/6.1992-439](https://doi.org/10.2514/6.1992-439), available at: <http://arc.aiaa.org/doi/10.2514/6.1992-439>
- Wilcox, D.C. (1988), "Reassessment of the scale-determining equation for advanced turbulence models", *AIAA Journal*, Vol. 26 No. 11, pp. 1299-1310, doi: [10.2514/3.10041](https://doi.org/10.2514/3.10041), ISSN: 00011452.
- Wilcox, D.C. (1994), *Turbulence Modelling for CFD*, 2nd ed., DCW Industries, ISBN: 0-9636051-0-0.
- Witherden, F.D. and Jameson, A. (2017), "Future directions of computational fluid dynamics", 23rd AIAA Computational Fluid Dynamics Conference, pp. 1-16.
- Wu, J., Long, H. and Paterson, E. (2018), "Physics-informed machine learning approach for augmenting turbulence models: a comprehensive framework", *Physical Review Fluids*, Vol. 3 No. 7, pp. 1-28, doi: [10.1103/PhysRevFluids.3.074602](https://doi.org/10.1103/PhysRevFluids.3.074602), ISSN: 2469990X.
- Xiao, H., Jin Long Wu, S., Laizet. and L., Duan. (2020), "Flows over periodic hills of parameterized geometries: a dataset for data-driven turbulence modeling from direct simulations", *Computers and Fluids*, Vol. 200, p. 104431, doi: [10.1016/j.compfluid.2020.104431](https://doi.org/10.1016/j.compfluid.2020.104431), ISSN: 00457930.
- Yoder, D.A. and Orkwis, P.D. (2021), "On the use of optimization techniques for turbulence model calibration", *Computers and Fluids*, Vol. 214 No. May, p. 104752, doi: [10.1016/j.compfluid.2020.104752](https://doi.org/10.1016/j.compfluid.2020.104752), ISSN: 00457930, available at: <https://linkinghub.elsevier.com/retrieve/pii/S0045793020303224>

Corresponding author

Ryley McConkey can be contacted at: rmconke@uwaterloo.ca

For instructions on how to order reprints of this article, please visit our website:

www.emeraldgroupublishing.com/licensing/reprints.htm

Or contact us for further details: permissions@emeraldinsight.com