International Journal of Heat and Fluid Flow 98 (2022) 109072

Contents lists available at ScienceDirect

International Journal of Heat and Fluid Flow

journal homepage: www.elsevier.com/locate/ijhff

# Robust data-driven turbulence closures for improved heat transfer prediction in complex geometries

James Hammond [a,*], Marco Pietropaoli [b], Francesco Montomoli [a]

[a] Department of Aeronautics, Imperial College London, Exhibition Road, London, SW7 2AZ, UK
[b] TOffeeAM, Winsley Street, London, W1W 8HF, UK

ARTICLE INFO

*Keywords:*
Turbulence modelling
Heat transfer
Data-driven
Machine learning
Complex geometry

ABSTRACT

Data-driven turbulence closures have been used to successfully improve the prediction of first order effects in flows where the test case is similar to the training conditions. In this work, we examine the ability of such closures to improve second order effects, such as heat transfer for complex geometries. First, an upper bound for the available performance of the data-driven model is shown by inserting the Reynolds stresses from the high fidelity training data in a stable manner to a modified Reynolds-Averaged Navier–Stokes (RANS) model. Data-driven closures are then developed from this data set with a focus on robustness and the improvement in the heat transfer effects analysed. The geometries considered are four topology optimised square ducts. Each duct is optimised with varying weights of a multi-objective function to maximise heat transfer and minimise pressure losses. One duct is used in training and the remaining three are set aside for testing the developed closure. For the training case, it is found that simple data-driven closures for the Reynolds stresses alone, are capable of improving prediction of the velocity and temperature fields by 38.2% and 34.8% respectively even in complex geometries. These improvements are largely retained in the testing cases demonstrating the robust generalisation of the developed model for this class of flow.

## 1. Introduction

In recent years there has been growing interest in using data-driven models for CFD simulations. Despite continuing advances in computational power, Reynolds averaged Navier–Stokes (RANS) solutions are still popular in an industrial setting, particularly for wall bounded flows such as gas turbines, heat exchanges and internal flows. Given the well documented shortcomings of such solutions, there has been substantial research effort directed towards improving the accuracy of RANS by leveraging machine learning (ML) methods with high fidelity data sets.

For gas turbines in particular, the internal coolant system is highly three dimensional and there is growing complexity in design to achieve improvements in coolant efficiency. It has been reported that an error of 20 degrees in blade temperature can change the estimated life of the component by up to 33% (Bunker, 2009). The external gas temperature in turbines can reach 2000 K, implying that an error of just 1% can have a substantial impact on real engine applications. It is therefore crucial to have accurate predictions of heat transfer and ML methods can play a crucial role.

Preliminary works were carried out to reduce the epistemic uncertainty of RANS models under the umbrella of uncertainty quantification (Gorlé et al., 2019). Among the first authors to apply dedicated machine learning methods to CFD, Ling and Templeton (2015) introduced a classifier to predict regions of high uncertainty in RANS predicted flows. A significant step forward was made by Parish and Duraisamy (2016) who improved turbulence closures using field inversion and ML. The spatial distribution of discrepancies was inferred for a range of problems and subsequently transformed into a functional form using ML, to be applied as a closure.

Since then, many different machine learning methods have been developed. Most notably the gene expression programming approach of Weatheritt and Sandberg (2016) and the DNNs from Ling et al. (2016). However, they are applied to a relatively small sample of simple training flows with very few examples where three-dimensional flows in complex geometries are considered. A few examples where three-dimensional flows are considered include (Huijing et al., 2021; Hammond et al., 2020). Furthermore, there are limited examples available in literature where second order heat transfer effects are explicitly considered.

One previous work in which data-driven methods have been applied for improving scalar transport came from Weatheritt et al. (2019). They focused on the efficacy of the Gradient-Diffusion Hypothesis (GDH), learning high-order mappings of the scalar flux $\overline{u_i'\theta'}$ from an

* Corresponding author.
*E-mail address:* j.hammond18@imperial.ac.uk (J. Hammond).

https://doi.org/10.1016/j.ijheatfluidflow.2022.109072
Received 28 February 2022; Received in revised form 26 September 2022; Accepted 18 October 2022
Available online 1 November 2022

LES database which could subsequently be applied to a precursor RANS solution. However, as the authors acknowledged, this approach attempts to reproduce second order heat transfer effects when the underlying velocity fields do not match. The scalar flux model is learnt from the LES velocity and then applied in prediction using a RANS velocity field. Given that for incompressible flow the temperature can be considered a passive scalar, it should be sufficient to improve the first order fields, in order to produce an improvement in the second order quantities.

If an improvement in heat transfer is not observed through reconciling the RANS velocity with LES, then attention can turn to the scalar flux closure. Otherwise, the scalar flux model is being tasked with overcoming deficiencies for which it is not responsible. This combined strategy where data-driven models are learnt for both the Reynolds stress and the scalar flux gives the greatest possibility for improved heat transfer prediction but is outside the scope of the present work. Consequently, the focus in this work is on the efficacy of the Reynolds stress closure and subsequent analysis of the improved heat transfer prediction through modification to the momentum equations alone.

Recently Frey Marioni et al. (2021) showed that in industrial applications, even simple networks can accurately reproduce the high fidelity flow field. Moreover, many complex neural networks underperform when applied in an *a posteriori* sense as their high degree of non-linearity can affect solver stability. Hammond et al. (2022a) also provided a method of assessing bounds on the expected performance of data-driven models. They showed that simple closures which are linear with respect to a set of well-chosen input features can provide significant improvements on current state-of-the-art models. Here, this philosophy is continued. Closures of the same class are considered, and it is demonstrated that they can be effective at improving both first and second order effects even in the case of highly complex three-dimensional geometries.

The paper is organised as follows. First, the geometries considered for training and testing purposes are described, then the governing equations are presented for the RANS (baseline and data-driven) and high fidelity frameworks along with the heat transfer model used. The remaining sections follow Fig. 1. An approximation to the upper bound in prediction for the given high fidelity data is uncovered. Simple, robust closures are then developed showing the lower bound that more complex machine learning models should be capable of achieving. These bounds are given in terms of first and second order flow quantities. Finally, the developed closure is applied to the testing geometries to assess how it performs in predictive situations, removed from the training case.

## 2. Training and testing geometries

The geometries considered are obtained as simple representations of an optimised internal coolant system of a gas turbine rotor blade. Whilst it is outside the scope of this paper to describe the optimisation strategy, the geometries are employed as they represent the next generation of coolant system design in modern gas turbines. It has been shown by Hammond et al. (2022b) the critical nature of accurate turbulence modelling in such designs.

A straight square channel with a length–width ratio of 10:1 represents the design domain on which four different optimal geometries are found using the topology optimisation approach described in Pietropaoli et al. (2020). The optimisations are conducted using a multi-objective approach which aims to simultaneously maximise the heat transfer between the fluid and the walls, whilst minimising the pressure losses of the fluid along the channel. An example of one of the optimised designs is presented in Fig. 2. The full domain optimised solid geometry and the corresponding fluid region are visible. A quarter of the full domain is used for simulation given the expected symmetry of the case, with symmetry boundary conditions imposed on the two internal faces highlighted.

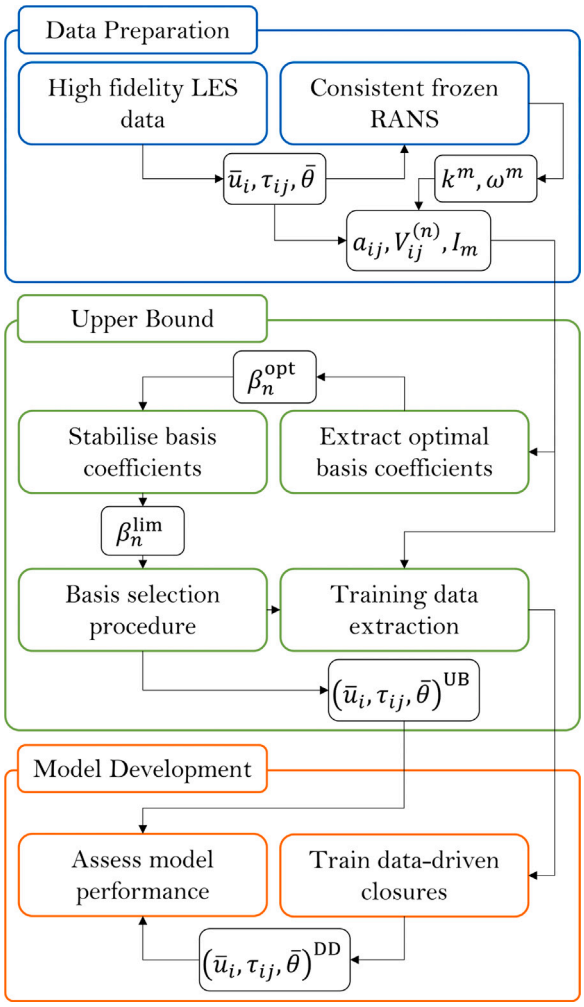


**Fig. 1.** Flow diagram detailing the main steps of the data-driven process.

**Table 1**
Flow parameters.

| Parameter | Value | Units |
|---|---|---|
| $U_\infty$ | 10.25 | $\mathrm{m\,s^{-1}}$ |
| $T_\infty$ | 600 | K |
| $T_{wall}$ | 1000 | K |
| $\rho$ | 0.59 | $\mathrm{kg\,m^{-3}}$ |
| $\nu$ | $5.12 \times 10^{-5}$ | $\mathrm{m^2\,s^{-1}}$ |
| $c_p$ | 0.76 | $\mathrm{kJ\,kg^{-1}\,K^{-1}}$ |
| $k_f$ | $4.66 \times 10^{-2}$ | $\mathrm{W\,m^{-1}\,K^{-1}}$ |

The four geometries, visible in Fig. 3, are obtained using the flow parameters detailed in Table 1 which correspond to real gas turbine working conditions. These conditions combined with a length $h = 0.01$ m relate to an inlet Reynolds number $Re = 5 \times 10^3$. The four geometries show similar patterns. Complex flow turbulators form counter-rotating vortices that increase the contact time between the fluid and the wall, enhancing heat exchange between the fluid and the walls. These complex geometries are not exactly periodic in nature and are instead shaped by the local flow conditions showing the vast superiority in complexity that can be achieved from such optimisation techniques in comparison to human designers. The specific shapes are created to avoid large regions of separation locally, which are inherently linked with high pressure losses.

Variability of the first three square duct designs SD100, SD200 and SD300 is introduced by changing the relative weights of the heat extraction and pressure losses in the multi-objective optimisation

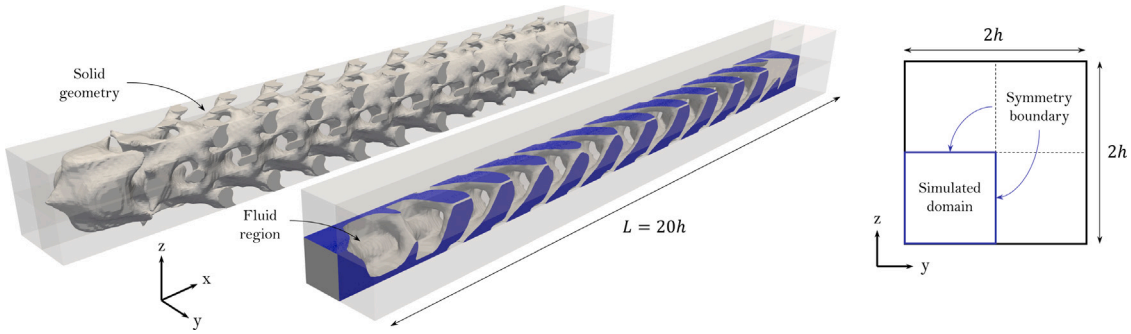


**Fig. 2.** Schematic of the solid optimised geometry and the corresponding simulated fluid region. A quarter of the full domain is simulated, with symmetry conditions imposed on the highlighted faces.

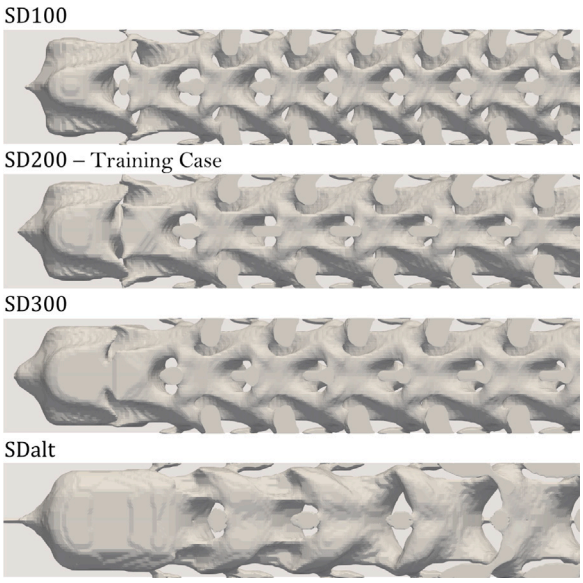


**Fig. 3.** Inlet region and first several ribs of each geometry considered. SD200 is used for training, with SD100, SD300 and SDalt retained for testing.

process. The numbers 100–300 represent the ratio of the heat transfer to pressure losses weighting. It is clear from Fig. 3 that the higher the weighting on heat transfer, the greater the volume of solid geometry that is introduced in the duct. This naturally increases the contact surface area between the fluid and the walls, enhancing heat transfer but presenting a greater blockage effect to the fluid passing along the duct which increases the pressure losses.

The final design SDalt is obtained by deliberately altering the initial optimisation search point to converge to another class of design corresponding to a different local minimum in the optimisation space. In the first three designs, the same rib topology can be observed which acts to create 8 counter rotating vortices along the duct. SDalt has a relatively simpler topology which creates just 4 counter rotating vortices, one in each quarter of the full optimisation domain.

In this work the case SD200 is used for training the data-driven closures as it is the case with mid weighting between pressure drop and heat transfer, with cases SD100, SD300 and SDalt set aside for testing purposes. SD100 and SD300 provide relatively similar cases to the training case and test how well a closure may stand up to changes in optimisation conditions. SDalt aims to provide a test case that is slightly further removed from the training environment by presenting a different topology to see how well the closures perform in a more general setting. If SD100 had been chosen for training, which has optimisation weighting further towards pressure drop than SD200, the data-driven closure would be required to extrapolate further in predicting SD300 which is weighted more towards heat transfer. The same is conversely true if SD300 had been chosen and both situations would come with a likely reduction in efficacy. Whilst it is outside the scope of this work, the sensitivity of the results to the chosen training case would be an interesting study of its own.

Fig. 4 shows the effect more clearly over the first half of each duct. The size of the turbulators clearly reduces from SD100 to SD300 and as a result of the smaller cross-sectional area, Fig. 4(a) shows the fluid velocity increasing. This also has the effect of increasing the level of turbulence in the flow and results in the raised temperature from SD100 to SD300 visible in Fig. 4(b). The differing structure of the flow can also be seen for SDalt which appears to have the lowest level of fluid heating up to the duct midpoint. Fig. 4 highlights the extrapolation that the data-driven closure will have to deal with during testing.

## 3. Governing equations

### 3.1. Baseline RANS framework

The steady RANS equations for an incompressible Newtonian fluid are presented as

$$\frac{\partial \overline{u}_i}{\partial x_i} = 0, \tag{1}$$

$$\overline{u}_j \frac{\partial \overline{u}_i}{\partial x_j} = \frac{\partial}{\partial x_j}\left[-\frac{1}{\rho}\overline{p}\delta_{ij} + \nu\left(\frac{\partial \overline{u}_i}{\partial x_j} + \frac{\partial \overline{u}_j}{\partial x_i}\right) - \tau_{ij}\right], \tag{2}$$

with time averaged velocity $\overline{u}_i$, pressure $\overline{p}$, density $\rho$, kinematic viscosity $\nu$ and Reynolds stress tensor $\tau_{ij}$ which is unclosed and must be modelled. It is conventional to decompose the Reynolds stress into isotropic and anisotropic components as

$$\tau_{ij} = \frac{2}{3}k\delta_{ij} + 2ka_{ij}, \tag{3}$$

with turbulent kinetic energy $k$ and nondimensional anisotropic component $a_{ij}$. It is the anisotropic component that is generally the subject of modelling efforts. Traditionally, turbulence models assume the anisotropic component to be linear with respect to the mean strain rate $S_{ij}$, realised by a proportionality constant known as the eddy (turbulent) viscosity $\nu_t$ such that $2ka_{ij} = -2\nu_t S_{ij}$. This allows the momentum equation Eq. (2) to be rewritten as

$$\overline{u}_j \frac{\partial \overline{u}_i}{\partial x_j} = \frac{\partial}{\partial x_j}\left[-P\delta_{ij} + \nu_{\text{eff}}\left(\frac{\partial \overline{u}_i}{\partial x_j} + \frac{\partial \overline{u}_j}{\partial x_i}\right)\right], \tag{4}$$

where $P = \overline{p}/\rho + 2k/3$ and $\nu_{\text{eff}} = \nu + \nu_t$. These Eddy Viscosity Models (EVMs) require the definition of turbulent scales in order to obtain $\nu_t$, the most common approach being single-point, two-equation models. In this work we adopt the $k-\omega$ SST-2003, two-equation model as a baseline (Menter et al., 2003). Only the relevant components will be presented throughout this work and for the full model, complete with auxiliary expressions and closure coefficients the reader is referred to Menter et al. (2003).

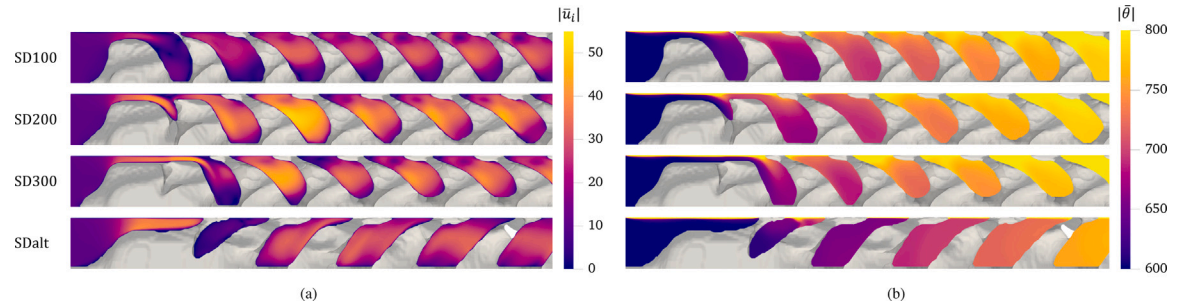


**Fig. 4.** Comparisons of (a) average velocity magnitude and (b) average temperature on the symmetry boundary for the four cases SD100, SD200 (training case), SD300 and SDalt.

### 3.2. Data-driven framework

Data-driven models rely on a more general, nonlinear relationship between $a_{ij}$ and the mean strain $S_{ij}$ and rotation $\Omega_{ij}$ rate tensors and some turbulent timescale $\tau_t$. The relationship, originally presented by Pope (1975), makes use of the Cayley–Hamilton theorem and preserves Galilean invariance. Defining nondimensional forms of the two rate tensors as $s_{ij} = \tau_t S_{ij}$ and $\omega_{ij} = \tau_t \Omega_{ij}$ then $a_{ij} = a_{ij}(s_{ij}, \omega_{ij})$ where

$$a_{ij}(s_{ij}, \omega_{ij}) = \sum_{n=1}^{10} \beta_n V_{ij}^{(n)}(s_{ij}, \omega_{ij}). \tag{5}$$

This constitutes a sum of linearly independent tensor basis $V_{ij}^{(n)}$ for $n = 1, \dots, 10$, each weighted by a coefficient $\beta_n$ which themselves are functions of a set of scalar invariants $I_m$ for $m = 1, \dots, 5$. Whilst in theory all 10 terms are necessary for general mappings of $a_{ij}$ (relying on the assumption that the full functional dependence is in fact $a_{ij} = a_{ij}(s_{ij}, \omega_{ij})$), in cases where the bases are non-degenerate, it has been shown that the first 5 terms alone are required for a local description of $a_{ij}$ (Gatski and Jongen, 2000). Following the aim to develop robust and industrially relevant closures it is considered that these terms, which are of the lowest degree, are also least likely to produce spurious behaviour. Consequently, only these terms which are defined as

$$\begin{aligned} &V_{ij}^{(1)} = s_{ij}, \; V_{ij}^{(2)} = s_{ij}\omega_{ij} - \omega_{ij}s_{ij}, \\ &V_{ij}^{(3)} = s_{ik}s_{kj} - \frac{1}{3}I_1\delta_{ij}, \; V_{ij}^{(4)} = \omega_{ik}\omega_{kj} - \frac{1}{3}I_2\delta_{ij}, \\ &V_{ij}^{(5)} = s_{ik}s_{kl}\omega_{lj} - \omega_{ik}s_{kl}s_{lj}, \end{aligned} \tag{6}$$

are considered in the present work. The corresponding scalar invariants are

$$\begin{aligned} &I_1 = s_{mn}s_{nm}, \; I_2 = w_{mn}w_{nm}, \; I_3 = s_{mn}s_{nk}s_{km}, \\ &I_4 = s_{mn}w_{nk}w_{km}, \; I_5 = s_{mn}s_{nk}w_{kl}w_{lm}. \end{aligned} \tag{7}$$

Data-driven models still require the definition of the turbulent scales in order to define $\tau_t$ so elements of the baseline $k-\omega$ SST model are still required. The turbulent viscosity is defined in the $k-\omega$ SST model as

$$\nu_t = \frac{a_1 k}{\max(a_1\omega, SF_2)}, \tag{8}$$

This definition allows selection of an appropriate turbulent timescale which is $\tau_t = \nu_t/k = a_1/\max(a_1\omega, SF_2)$. It is noted that this definition differs slightly from other works in which the turbulent timescale is defined as $\tau_t = 1/\omega$ (Weatheritt and Sandberg, 2017; Schmelzer et al., 2020). The present definition allows for the switch to a strain based timescale in regions where the ratio of production to dissipation of $k$ is high (Menter, 1994).

Finally, to complete the description of the data-driven framework, the production term, present in both the $k$ and $\omega$ transport equations of the $k-\omega$ SST model, must be given special attention. The production is defined as

$$P_k = \min\left(-2ka_{ij}\frac{\partial \bar{u}_i}{\partial x_j}, 10\beta^*\omega k\right) \tag{9}$$

for incompressible flow, where only the anisotropic component of $\tau_{ij}$ is included as for incompressible flow the isotropic components contribution to $P_k$ is identically zero ($\delta_{ij}\partial_j\bar{u}_i = \partial_i\bar{u}_i \equiv 0$).

### 3.3. Heat transfer model

In order to analyse the second order heat transfer effects, a passive scalar transport equation for time averaged temperature $\bar{\theta}$ is also solved where

$$\bar{u}_j\frac{\partial\bar{\theta}}{\partial x_j} = \frac{\partial}{\partial x_j}\left(\Gamma\frac{\partial\bar{\theta}}{\partial x_j} - \overline{u_j'\theta'}\right), \tag{10}$$

with conductivity $\Gamma = \nu/Pr$ where $Pr = c_p\mu/k_f$ is the Prandtl number calculated from the flow parameters in Table 1. The term $\overline{u_i'\theta'}$ is the correlation between velocity and scalar fluctuations known as the scalar flux. Similar to $\tau_{ij}$, this term is not closed so must be modelled. The standard model for $\overline{u_i'\theta'}$ is the Gradient-Diffusion Hypothesis (GDH) which is given, in a manner analogous to the turbulent viscosity, as

$$\overline{u_i'\theta'} = -\Gamma_t\frac{\partial\bar{\theta}}{\partial x_i} \tag{11}$$

such that the molecular and turbulent conductivities can be combined to give an effective conductivity $\Gamma_{\text{eff}} = \Gamma + \Gamma_t$ and the scalar transport equation may be rewritten

$$\bar{u}_j\frac{\partial\bar{\theta}}{\partial x_j} = \frac{\partial}{\partial x_j}\left(\Gamma_{\text{eff}}\frac{\partial\bar{\theta}}{\partial x_j}\right). \tag{12}$$

The turbulent conductivity is defined as $\Gamma_t = \nu_t/Pr_t$ where the turbulent Prandtl number $Pr_t$ is a constant to be selected for the specific flow and here is set as $Pr_t = 0.85$.

### 3.4. High fidelity framework

The high fidelity database used to train and test the data-driven closures was collected from LES performed on the cases described in Section 2. The LES are solved via the pressure implicit with splitting of operator (PISO) algorithm and use wall-adapting local eddy viscosity (WALE) sub-grid scale (SGS) turbulence model. The meshes for each case are generated using an unstructured polyhedral mesher with boundary layer refinement obtained through the inclusion of prismatic layers along no-slip walls. To ensure a wall $y^+ < 1$, 16 prismatic layers are used with a geometric stretching ratio of 1.15. The resulting meshes each consist of approximately $1 \times 10^7$ elements. Adaptive timestepping is used to prevent the CFL number from rising above 1. The simulations are run for a total of 0.45s which allows for 30 fluid passes through the domain from which to extract mean flow statistics.

Training data is extracted from SD200 whilst the high fidelity data collected for cases SD100, SD300 and SDalt are set aside for testing. For training, the dataset $\mathcal{D}_{\text{train}} = \{a_{ij}, V_{ij}^{(n)}, I_m\}$ must be obtained from the raw high fidelity data. Time averaging the flow history from LES gives the raw high fidelity data $\mathcal{D}_{\text{raw}} = \{\bar{u}_i, \bar{p}, \tau_{ij}\}$. In order to calculate $\mathcal{D}_{\text{train}}$ from $\mathcal{D}_{\text{raw}}$ the turbulent scale quantities $k$ (to calculate $a_{ij}$ from

**Table 2**
Results of upper bound approach for mean square errors $\epsilon(\cdot)$ and normalised errors $\epsilon(\cdot)/\epsilon(\cdot^0)$ where normalised errors are presented with respect to the baseline $k-\omega$ SST.

| $V_{ij}^{(n)}$ | Mean square error | | Normalised error | |
|---|---|---|---|---|
| | $\epsilon(\overline{u}_i)$ | $\epsilon(\overline{\theta})$ | $\epsilon(\overline{u}_i)/\epsilon(\overline{u}_i^0)$ | $\epsilon(\overline{\theta})/\epsilon(\overline{\theta}^0)$ |
| 1 | $1.067\times10^{1}$ | $4.706\times10^{2}$ | 0.4073 | 0.4147 |
| 1, 2 | $1.779\times10^{1}$ | $4.763\times10^{2}$ | 0.6791 | 0.4197 |
| 1, 3 | $1.106\times10^{1}$ | $4.706\times10^{2}$ | 0.4224 | 0.4147 |
| 1, 4 | $9.750\times10^{0}$ | $4.700\times10^{2}$ | 0.3722 | 0.4141 |
| 1, 5 | $1.100\times10^{1}$ | $4.855\times10^{2}$ | 0.4200 | 0.4278 |

Eq. (3)) and $\omega$ (to calculate $V_{ij}^{(n)}$ and $I_m$ from Eqs. (6) and (7) respectively) must first be specified. To facilitate this discussion, quantities calculated directly from high fidelity data will be denoted $\phi^\star$, whilst modelled quantities obtained from solving the RANS turbulent transport equations are denoted $\phi^m$. Researchers recognised quickly that $\omega^\star$, calculated directly from high fidelity data is often inconsistent with the $\omega^m$ obtained from its RANS transport equation (Parneix et al., 1998). Using $\omega^\star$ in training can then significantly affect results when testing developed closures in a RANS environment where all turbulence quantities must be modelled. To remedy this, the frozen training approach has been widely adopted where $\mathcal{D}_{\text{raw}}$, (which are consistent quantities across the high fidelity and RANS environments), is held constant and the $\omega$-transport equation is solved to obtain $\omega^m$ for the given high fidelity mean flow data.

$$\overline{u}_i\frac{\partial\omega}{\partial x_i}=\frac{\gamma}{\nu_t}P_k-\beta\omega^2+\frac{\partial}{\partial x_i}\left[(\nu+\sigma_\omega\nu_t)\frac{\partial\omega}{\partial x_i}\right]+(1-F_1)C\,D_{k\omega}. \tag{13}$$

The result is an $\omega^m$ field that would be obtained, should a RANS calculation admit the exact mean flow observed in $\mathcal{D}_{\text{raw}}$. This leaves the task of appropriately defining $k$ which is now required for Eq. (13) to calculate $\nu_t$, as well as for Eq. (3). A simple choice is to use the high fidelity $k^\star$ calculated by definition as $k^\star=\tau_{ii}^\star/2$. However, it has been observed that whilst $k^\star$ and $k^m$ are often more consistent than their respective $\omega$ counterparts, their consistency cannot be assumed, especially in near wall behaviour. Although $k$ is a quantity that is often of greater interest in analysis than $\omega$ and so its correct prediction is important, the authors consider that inconsistency in the high fidelity and RANS environments is more likely to result in unexpected RANS prediction. Consequently, in the present work, the $k$-transport equation is also solved

$$\overline{u}_i\frac{\partial k}{\partial x_i}=P_k-\beta^*\omega k+\frac{\partial}{\partial x_i}\left[(\nu+\sigma_k\nu_t)\frac{\partial k}{\partial x_i}\right], \tag{14}$$

to obtain the RANS consistent $k^m$ for the raw high fidelity data $\mathcal{D}_{\text{raw}}$. Eqs. (13) and (14) are coupled and must be solved simultaneously to give a consistent frozen RANS solution.

It may be noticed that there is no reference to heat transfer quantities in the preparation of the training data. This is to be expected as the present work explores the possibility of improving heat transfer results through modifications to the equations for first order quantities only. The scalar flux model is present simply as a means of quantifying the improvement in heat transfer prediction, rather than as a means to drive the training of the developed models.

## 4. Model development

### 4.1. Upper bound extraction

Following the definition of an appropriate and consistent training dataset $\mathcal{D}_{\text{train}}$ in Section 3.4, it is prudent to assess the efficacy of the data. This may be achieved by extracting optimal values for the $\beta_n$ coefficients (where some limiting may be required for stability) and substituting them into a RANS calculation of the training case to obtain an estimate for the upper bound in model performance available from the training data on the training case. It should be highlighted that this approach does not constitute a model itself, and is simply the insertion of static fields derived directly from high fidelity data into a RANS calculation. This process also allows the decomposition of the high fidelity Reynolds stresses into basis components. Consequently, an interrogation of which terms play the most significant roles in affecting the mean flow fields can be made, allowing the removal of other less important terms from the model development procedure *a priori*.

In a previous study, Hammond et al. (2022a) applied this approach for statistically 2D flows which is an extension of the work from Wu et al. (2019). Here, the ideas are extended to a much more complex 3D flow. In order to capture the interplay between terms, optimal $\beta_n$ coefficients are extracted for combinations of terms. Rather than solving exactly for $\beta_n$ at each point as is possible with a 2D flow, the $\beta_n$ terms are optimised in a least squares sense. Consider a subset of the first 5 basis defined by $C\subseteq\{1,\dots,5\}$. The $\beta_n$ coefficients for this set can be found by solving the local least squares regression problem

$$\beta_n^{\text{loc}}(x_i)=\underset{\beta_n\ \text{for}\ n\in C}{\arg\min}\left\|a_{ij}^\star-\sum_{n\in C}\beta_n(x_i)V_{ij}^{(n)\star}\right\|_2,\ \text{for}\ n\in C, \tag{15}$$

where $x_i$ is the spatial coordinate of each grid point.

Once the coefficients for a certain combination of terms have been found they must be stably inserted into a RANS calculation to analyse their effects on the mean flow field. To further improve conditioning and reduce the propagation of errors from small discrepancies in the Reynolds stresses the recommended approach of Brener et al. (2021) is followed in which the optimal coefficients are extracted with high fidelity velocity information from the tensor bases, as opposed to recomputing optimal fields based on the RANS predicted tensor bases with each iteration. This often requires limiting of the optimal values as regions of negligible velocity gradient but non-negligible stress force the least squares regression to find large magnitude values of the $\beta_n$ coefficients. In this work, the limits were set across all basis using a bisection method to find limiting values which ensure stability when the optimal coefficients are inserted into a RANS calculation whilst avoiding limiting too severely and damaging the prediction quality. Stable and accurate limits were found to be $-50<\beta_1^{\text{loc}}<0$ and $|\beta_n^{\text{loc}}|<50$ for $n=2,\dots,5$. Combinations of $V_{ij}^{(1)}$ with each of the higher order basis independently were tested to assess the effect of each higher order term. $V_{ij}^{(1)}$ is included in all combinations as it is a linear, viscous term and all calculations become unstable without it. The calculations are run starting with the consistent frozen RANS solution as initial conditions. The continuity, momentum and heat transfer equations are then unfrozen and allowed to converge with fixed $\beta_n^{\text{loc}}$ for $n\in C$ fields.

The results of this analysis are presented in Table 2. The mean square errors $\epsilon(\overline{u}_i)$ and $\epsilon(\overline{\theta})$ give the raw errors with respect to the high fidelity for the mean velocity and temperature fields respectively. The normalised errors present the fraction of the raw mean square error with respect to the raw mean square error of the baseline $k-\omega$ SST simulation. A normalised error of 0, therefore, represents the exact reproduction of the high fidelity data and a normalised error of 1 represents no improvement over the baseline model. First and foremost, it is encouraging to see that all combinations of terms give significantly improved results on the baseline model for both velocity and temperature prediction. Secondly, it is interesting that none of the simulations including higher order terms improve significantly on the model considering $V_{ij}^{(1)}$ in isolation. Terms $V_{ij}^{(3)}$ and $V_{ij}^{(5)}$ are not selected for model development as their inclusion yields very little effect on the mean flows, and even deteriorates prediction compared with considering $V_{ij}^{(1)}$ alone. $V_{ij}^{(2)}$ is also discounted from further analysis from the perspective of model stability. Even when significantly limiting its coefficient $\beta_2$, an oscillatory solution was obtained. Whilst including $V_{ij}^{(4)}$ does give a nominal improvement of 3.5% over using $V_{ij}^{(1)}$ alone, the penalty incurred from the additional model complexity of an extra term, outweighs the observed performance benefits. As

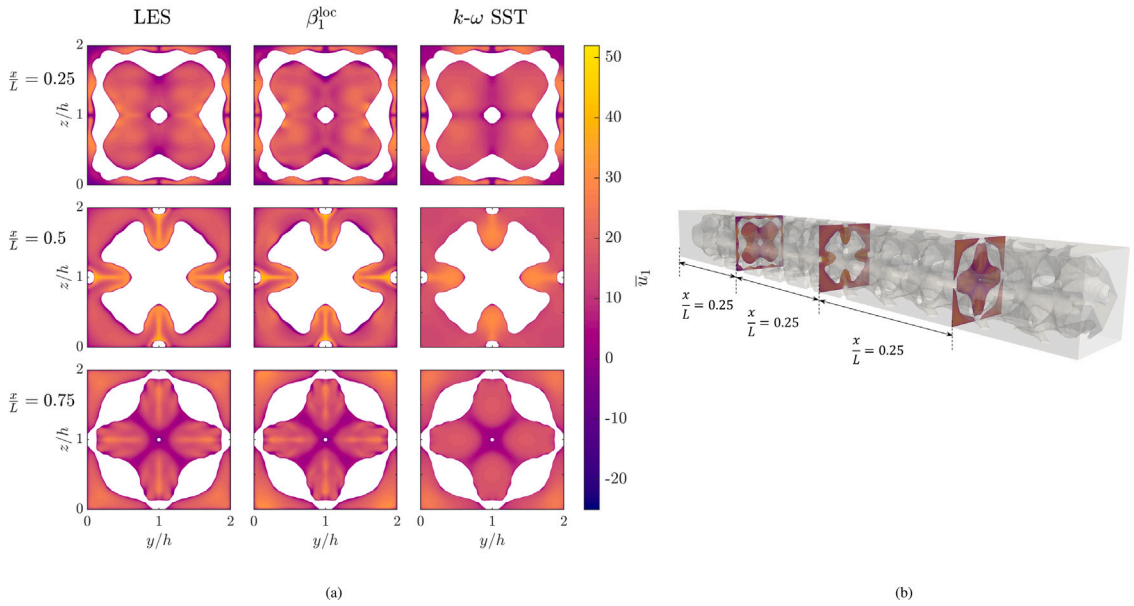


**Fig. 5.** (a) $x$-component of velocity at 3 streamwise locations ($x/L = 0.25, 0.5, 0.75$) as indicated in (b) for: (left) high fidelity LES; (middle) upper bound local regression, $\beta_1^{\text{loc}}$; and (right) baseline $k-\omega$ SST.

**Table 3**
Results of local and linear regressions of $\beta_1$ for mean square errors $\epsilon(\cdot)$ and normalised errors $\epsilon(\cdot)/\epsilon(\cdot^0)$ where normalised errors are presented with respect to the baseline $k-\omega$ SST.

| Model | Mean square error | | Normalised error | |
|---|---|---|---|---|
| | $\epsilon(\overline{u}_i)$ | $\epsilon(\overline{\theta})$ | $\epsilon(\overline{u}_i)/\epsilon(\overline{u}_i^0)$ | $\epsilon(\overline{\theta})/\epsilon(\overline{\theta}^0)$ |
| $\beta_1^{\text{loc}}$ | $1.067 \times 10^1$ | $4.706 \times 10^2$ | 0.4073 | 0.4147 |
| $\beta_1^{\text{lin}}$ | $1.619 \times 10^1$ | $7.403 \times 10^2$ | 0.6181 | 0.6523 |

such, the decision is made to select just $V_{ij}^{(1)}$ in the search for data-driven models and the upper bound of performance with this term is approximated as a 59.3% improvement for velocity prediction and a 58.5% improvement for the temperature field.

A possible explanation as to why the higher order terms have such a small impact on the training case is that the heat exchanger flow contains no significant turbulent core, with substantial influence from the nearest wall along its entire length. The nature of this geometry is likely to restrict the production of phenomena arising from turbulent interactions which may be modelled by the higher order terms. It is therefore likely that the influence of these terms on the present flow is reasonably small and most of the effects can be reasonably captured by considering $V_{ij}^{(1)}$ alone. Fig. 5(a) shows $\overline{u}_1$ contoured over the three sections indicated in Fig. 5(b) ($x/L = 0.25, 0.5, 0.75$) for the high fidelity LES data, the local regression $\beta_1^{\text{loc}}$, and the baseline $k-\omega$ SST. Clearly, including $V_{ij}^{(1)}$ alone has an appreciable impact on the velocity field prediction and the contours much more closely reproduce the high fidelity data than the baseline $k-\omega$ SST model.

### 4.2. Robust closures

Having defined which terms will be considered in the development of the data-driven closures, and having set some expectation of the level of improvement that is possible using these terms, attention can now turn to model development. It has been shown in numerous works that simple closures have the potential to provide very tangible enhancements (Zhao et al., 2020; Schmelzer et al., 2020). In the interest of developing robust closures, the search is restricted to models for $\beta_1$ that are at most linear with respect to the chosen set of input features. Specifically, take $\mathcal{M} \subseteq \{1, \dots, 5\}$ to be the set of indices defining the set of considered input features for a certain model. The input features considered are then defined as $I_m$ for $m \in \mathcal{M}$. Models are searched for in the space $\beta_1^{\text{lin}} = w_m \cdot I_m + b$ where $w_m$ is a vector of weights of length $|\mathcal{M}|$ and $b$ is a bias, such that models for the anisotropic component of the Reynolds stress take the form

$$a_{ij} = (w_m \cdot I_m + b) V_{ij}^{(1)}. \tag{16}$$

To find the values of the weights and bias the least squares multiple linear regression problem is formulated, which can be stated as

$$\langle w_m, b \rangle = \underset{w_m, b \text{ for } m \in \mathcal{M}}{\arg\min} \left\| a_{ij}^{\star} - (w_m \cdot I_m + b) V_{ij}^{(1)\star} \right\|_2. \tag{17}$$

This problem can be solved analytically (or by any other relevant method) to obtain $w_m$ and $b$ for any combination of the input features $I_m$.

Whilst this approach may appear simple it is the opinion of the authors that model complexity should be built up only when required. Many forms of more complex machine learning methods are available to attempt to extract further predictive performance from the training data from Gene Expression Programming (Weatheritt and Sandberg, 2016) and Deep Neural Networks (DNNs) (Ling et al., 2016) to less commonly considered approaches such as Elastic Net Regression (Schmelzer et al., 2020) and Random Forests (Kaandorp and Dwight, 2020). However, such approaches have a much higher chance of developing models which exhibit unpredictable behaviour in predictive cases.

To further explain, for the simplest turbulent flows there may exist a one to one mapping between input features and basis coefficients such that if enough data is used then a very close approximation to the true mapping can be obtained. Such an examples is shown in the turbulent channel flow considered by Taghizadeh et al. (2020). However, as observed by Hammond et al. (2022a), even for a backwards facing step which is still a canonical test case and exhibits only a small increase in complexity, this mapping may no longer be one to one. This implies that even complex DNN models with sufficient data require some level of 'averaging' of the coefficients in regions where the range of the exact coefficients for a given set of input features is large. When this occurs in a DNN it is hard to predict the effect on the flow field, both on its converged state and the dynamics of the network in the convergence process. For example, if the network predicts a value of $\beta_1$ that is lower than the true value, then the flow will struggle to converge in this

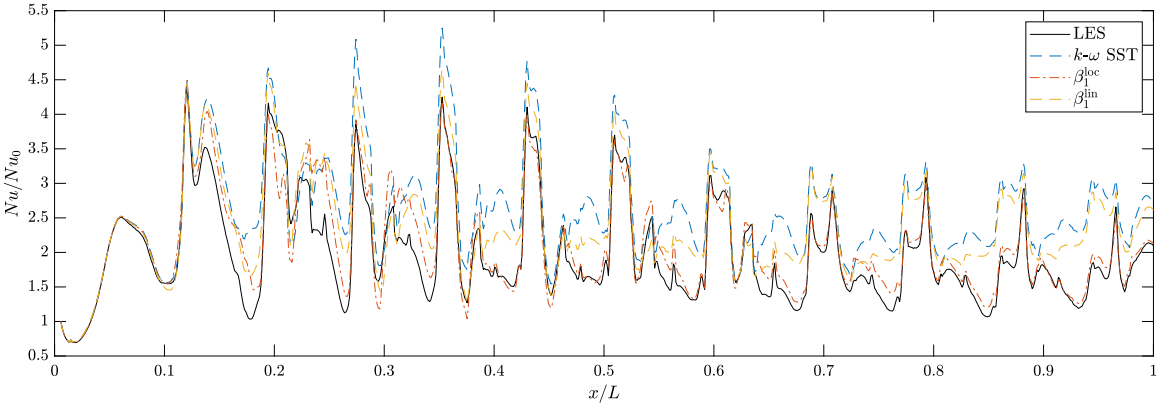


**Fig. 6.** Normalised line averaged Nusselt number plotted from inlet to outlet for the LES, $k-\omega$ SST, $\beta_1^{\text{loc}}$ local regression and $\beta_1^{\text{lin}}$ linear coefficient models.

region when the model is implemented in RANS calculations. During preliminary testing on more complex models, this effect was observed, resulting in spurious oscillations stemming from regions of separation in the flow that were not damped out by the model to obtain a steady solution.

In this work, where second order effects are being considered for the first time, the decision is made to consider only simple linear coefficient closures such as that defined in Eq. (16), with the intention of promoting robustness and setting the benchmark for what more complex models should be capable of achieving.

### 4.3. CFD testing

All possible sets defined by $\mathcal{M}$ were tested and the resulting models were assessed on their *a posteriori* performance in RANS calculations, rather than simply their *a priori* fit to the high fidelity data. Models are run in the finite volume package OpenFOAM. The best model in terms of mean square error for velocity and temperature fields was found as

$$a_{ij} = -0.674 V_{ij}^{(1)}. \tag{18}$$

The mean square and normalised errors of the velocity and temperature are presented in Table 3 along with the results of the upper bound local regression for comparison. Whilst Eq. (18) provides a very simple correction to the anisotropy, equivalent to a 32.6% reduction in the turbulent viscosity, the improvement compared to the baseline $k-\omega$ SST is still considerable for such a complex geometry. The velocity prediction improves by 38.2% over the baseline and the temperature improves by 34.8%. There is a reasonable gap between the performance of the linear coefficient model and the upper bound local regression which more complex machine learning techniques may be capable of exploiting, with a further 21.1% improvement available in the upper bound for the velocity and a further 23.8% for the temperature. However, significantly increasing model complexity to match the upper bound on the training case may come at the cost of overfitting. Consequently, for a simple model whose behaviour should be very predictable, the observed improvements are fairly encouraging.

Figs. 6–8 present further details about the resulting turbulence and heat transfer predictions. In Fig. 6 line averages of local Nusselt number over the length of the heat exchanger for the high fidelity LES, baseline $k-\omega$ SST, upper bound local regression and linear coefficient model, normalised with respect to their inlet values $Nu_0$, are presented. To calculate Nusselt number a Reynolds type analogy is used. First, the wall shear stress $\tau_w$ is extracted directly from each simulation and the skin friction coefficient is calculated as

$$C_f = \frac{\tau_w}{\frac{1}{2}\rho U_\infty^2}. \tag{19}$$

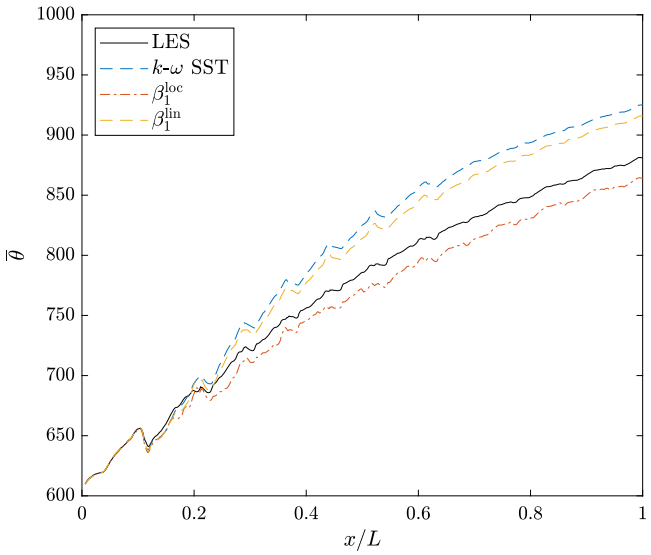


**Fig. 7.** Average temperature at each streamwise section plotted from inlet to outlet for the LES, $k-\omega$ SST, $\beta_1^{\text{loc}}$ local regression and $\beta_1^{\text{lin}}$ linear coefficient models.

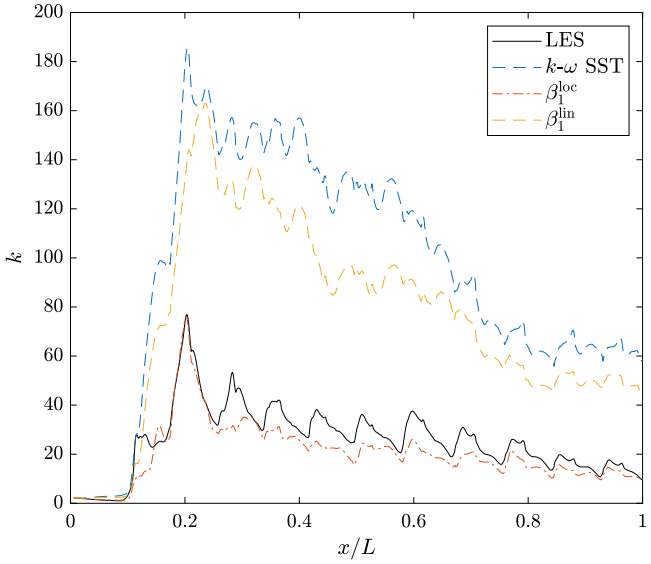


**Fig. 8.** Average turbulent kinetic energy at each streamwise section plotted from inlet to outlet for the LES, $k-\omega$ SST, $\beta_1^{\text{loc}}$ local regression and $\beta_1^{\text{lin}}$ linear coefficient models.

**Table 4**
Results of mean square errors $\epsilon(\cdot)$ and normalised errors $\epsilon(\cdot)/\epsilon(\cdot^0)$ for each of the testing geometries. Normalised errors are presented with respect to the baseline $k-\omega$ SST.

| Test case | Mean square error | | Normalised error | |
|---|---|---|---|---|
| | $\epsilon(\overline{u}_i)$ | $\epsilon(\overline{\theta})$ | $\epsilon(\overline{u}_i)/\epsilon(\overline{u}_i^0)$ | $\epsilon(\overline{\theta})/\epsilon(\overline{\theta}^0)$ |
| SD100 | $1.388\times10^1$ | $7.606\times10^2$ | 0.6745 | 0.5699 |
| SD300 | $1.995\times10^1$ | $6.642\times10^2$ | 0.7431 | 0.6619 |
| SDalt | $3.432\times10^1$ | $6.857\times10^2$ | 0.6733 | 0.6005 |

**Table 5**
Average outlet temperatures for each case.

| Test case | Outlet temperature/K | | |
|---|---|---|---|
| | LES | $k-\omega$ SST | Data-driven |
| SD100 | 872.4 | 928.5 | 914.1 |
| SD200 (training) | 881.0 | 925.1 | 915.8 |
| SD300 | 884.4 | 923.9 | 913.8 |
| SDalt | 881.1 | 926.1 | 913.1 |

By making use of the Chilton–Colburn J-factor analogy between heat and momentum transfer, which states that

$$J_H = \frac{C_f}{2} = J_M = \frac{h}{\rho c_p U_\infty} Pr^{2/3}, \tag{20}$$

to calculate the local heat transfer coefficient $h$, the local Nusselt number $Nu$ can then be evaluated as

$$Nu = \frac{hL}{k_f}. \tag{21}$$

To facilitate this discussion Figs. 7 and 8 show laterally averaged temperature $\theta$, and turbulent kinetic energy $k$, plots for each simulation along the length of the geometry. Over the inlet section up to $x/L = 0.1$ the flows are almost identical. There are then two clear jumps in the level of turbulent kinetic energy which occur between $x/L = 0.1$ and $x/L = 0.2$ where the flow hits the first rib. At this point, the baseline $k-\omega$ SST significantly overpredicts the production of $k$. The result of this is flow with an excessively turbulent velocity profile, where the Reynolds stresses transport fast flow in the freestream to the wall and slow flow at the wall into the freestream, flattening the velocity profile and increasing its gradient at the walls. Large velocity gradients at the wall increase the wall shear stress $\tau_w$ which in turn, drives increased heat transfer at the wall and explains the high $Nu$ values observed in this region for the baseline $k-\omega$ SST. The increased heat transfer then realises itself into higher average fluid temperature as the flow moves downstream. Further downstream the $k-\omega$ SST model is unable to reconcile the overproduction present at the first rib. The level of $k$ remains high throughout the flow, leading again to higher Nusselt numbers and an overprediction of the fluid temperature along the length of the geometry. The average outlet temperature for the LES is 881 K as opposed to 925 K for the $k-\omega$ SST.

Turning attention to the data-driven models, Fig. 8 shows that the upper bound local regression is able to prevent the large overproduction of $k$ after the inlet region. In fact, it slightly overcorrects the effect and underpredicts $k$ along the length of the geometry. Whilst this does not appear to lead to reduced Nusselt numbers further downstream, there are significant sections where the predicted $Nu$ is lower than the LES in the region up to $x/L = 0.5$, particularly immediately following the jump in $k$ at $x/L = 0.2$. This is enough to bring the average temperature prediction below the level of the LES and the average fluid temperature at the outlet is 863 K, representing a 41% improvement in outlet temperature compared to the baseline.

Finally, the data-driven linear coefficient closure only partially remedies the overprediction in $k$ production. However, it does so enough to allow reduced Nusselt numbers along the full length of the geometry and at the outlet the average fluid temperature is 915 K which is a 23% improvement on the baseline prediction. A differential exit temperature of 10 K may not seem significant on first inspection. However, in the safety critical and high-cost industry of aerospace, improvements of this order of magnitude can be significant. As mentioned in Section 1, Bunker (2009) predicted that a change in blade temperature of just 20 K could impact the expected life of the blade by 33% and a work by Rezazadeh Reyhani et al. (2013) concluded that a change in maximum blade temperature of 10 K could reduce the creep life of the blade as much as 40%.

Nevertheless, it is clear that making simple corrections to the turbulence closure alone is sufficient to realise a tangible improvement in the heat transfer prediction, whilst still considering a standard GDH closure for scalar flux $\overline{u_i'\theta'}$, in the scalar transport equation. Further improvement might then be achieved either by considering more complex machine learning techniques for the turbulence closure, or by turning attention to nonlinear or data-driven modelling of $\overline{u_i'\theta'}$.

## 5. Test geometries

To assess how well the constant coefficient data-driven model generalises, it is now subjected to testing on the remaining geometries SD100, SD300 and SDalt. It is reiterated here that as the upper bound local regression is not a model itself but simply a case specific set of static fields, it cannot be applied to these testing geometries.

Table 4 shows the results of the mean square and normalised errors for each of the three test cases. The data-driven model has successfully improved on the baseline $k-\omega$ SST in terms of velocity and temperature prediction for all the test cases. The level of improvement is fairly consistent across all metrics and similar to what was achieved on the training case. Plots of the normalised Nusselt numbers and laterally averaged temperature distributions from inlet to outlet for each of the cases for the LES, baseline $k-\omega$ SST and data-driven models are presented in Figs. 9 and 10 respectively. The overall conclusions are the same as for the training case SD200. Production of turbulence in the baseline $k-\omega$ SST model where the flow hits the first rib is significantly too high and this sets up a flow with excessive turbulence in comparison to high fidelity data along the full length of the duct. These enhanced levels of turbulence drive higher heat transfer between the fluid and the walls of the duct as seen from the elevated Nusselt number predictions. In turn, this leads to an overprediction of the fluid temperature along the duct. The effective reduction in turbulent viscosity of the data-driven model acts to reduce the overprediction of turbulence by the baseline model which corresponds to the improvement observed in both Nusselt number predictions and temperature distribution.

Of the three cases with similar topology (SD100, SD200 and SD300) it is the case with the lowest weighting towards heat transfer SD100, which renders the best improvement over the baseline. This is achieved despite its raw mean square error being the highest. Clearly, this is the case where the baseline model performs the worst, so the greatest improvement is available. As the heat transfer weighting gets higher and the turbulence levels increase, it seems that the baseline model is able to perform better. Table 5, which provides the outlet temperatures for each of the models for each case confirms this, showing a temperature differential between high fidelity data and baseline model of 56.1 K in comparison to 44.1 K for SD200 and 39.5 K for SD300. In Figs. 9 and 10 the overall improvement of the RANS models (both baseline and data-driven) is clear. This to some extent, explains why the improvement of the data-driven closure reduces at higher heat transfer weighting.

Finally, whilst it is difficult to draw direct comparisons of the results for SDalt as the structure of the flow is inherently different to the other cases, it is promising to see the simple data-driven closure holds up well in a predictive case further removed from the training environment.

Fig. 11 shows the average temperature distribution on a cross section at $x/L = 0.7$ for the high fidelity LES, linear coefficient $\beta_1^{\text{lin}}$ model and the baseline $k-\omega$ SST for each of the three testing geometries. The reduction in temperature in the bulk of the flow for the linear coefficient model over the baseline can be visualised in all three cases.

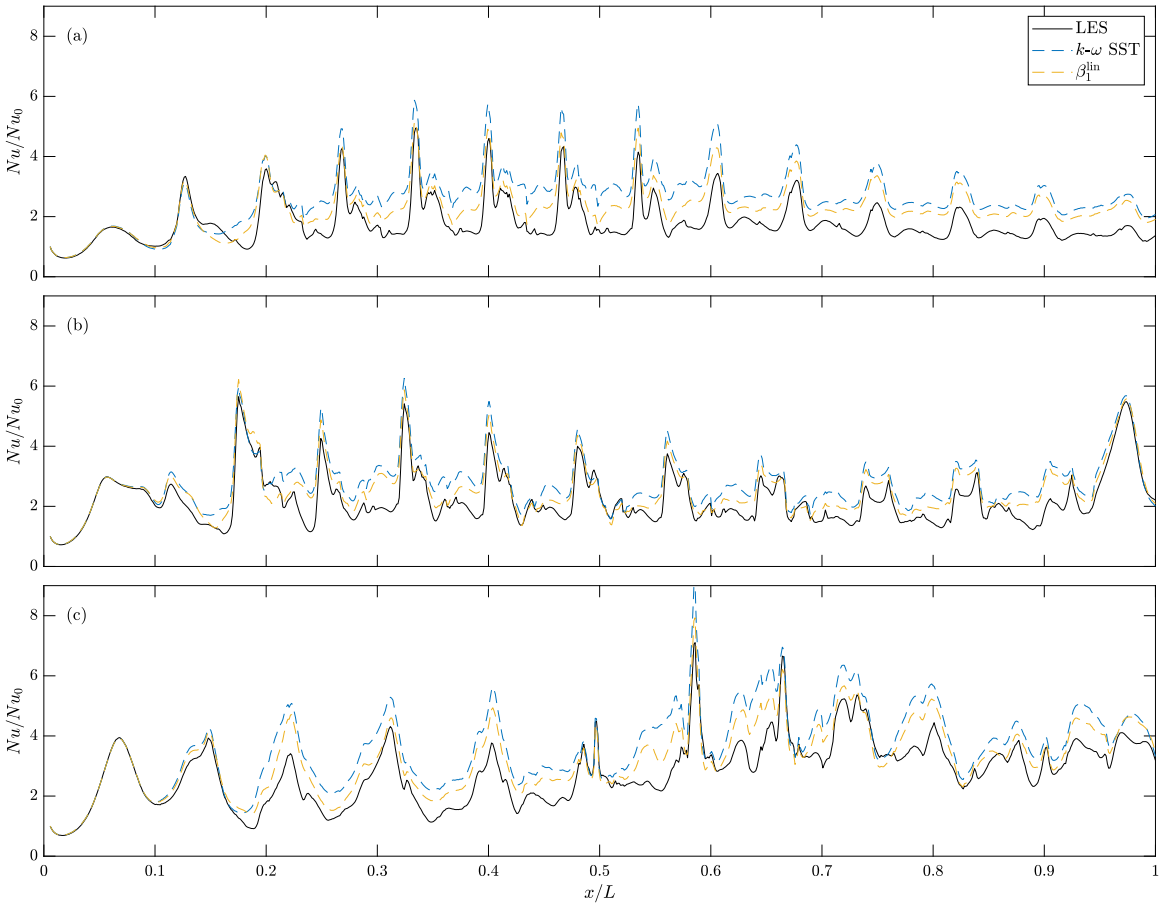


**Fig. 9.** Normalised line averaged Nusselt number plotted from inlet to outlet for the LES, $k-\omega$ SST and $\beta_1^{\text{lin}}$ linear coefficient models for the test cases (a) SD100, (b) SD300 and (c) SDalt.

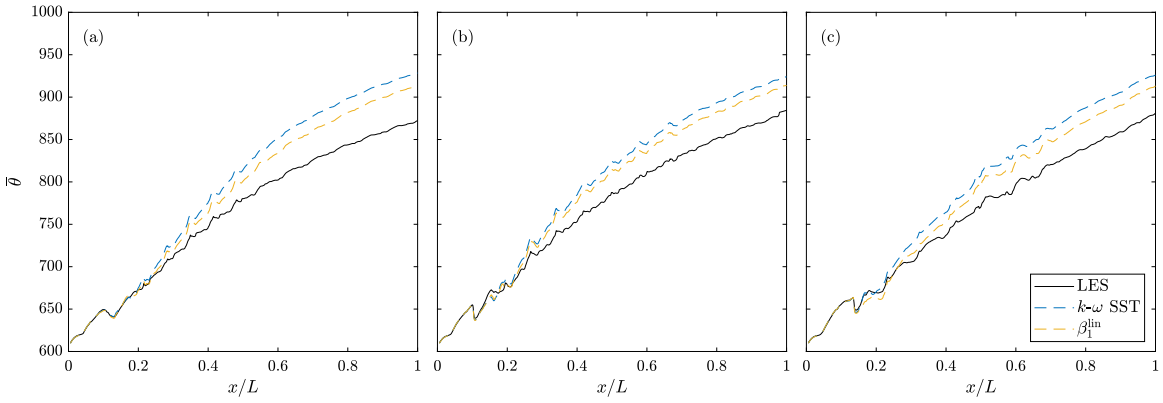


**Fig. 10.** Average temperature at each streamwise section plotted from inlet to outlet for the LES, $k-\omega$ SST and $\beta_1^{\text{lin}}$ linear coefficient models for the test cases (a) SD100, (b) SD300 and (c) SDalt.

However, it is clear there is still some improvement to be made towards the high fidelity data which could in future be exploited by more complex data-driven models.

## 6. Conclusions

In the open literature, there is a lack of work in which data-driven methods are applied to complex geometries representative of real applications, and also which considers heat transfer prediction. In the present study, the geometries used are the result of fluid topology optimisation which is gaining visibility for use in next generation gas turbine design. The geometrical complexity of such designs is very high and closer to those which may be present in real applications. The effect of data-driven closures for the Reynolds stress tensor on heat transfer prediction is also explicitly examined.

First, an approximation to the upper bound on the improvement possible from the available training data is obtained which gives a 59.3% improvement for velocity prediction and a corresponding 58.3% improvement in the temperature field over the baseline $k-\omega$ SST model. Closures are developed in a simple class where basis coefficients are constrained to be linear with respect to the chosen set of input features. Within this class, the model with the best performance improved the velocity prediction by 38.2% and temperature by 34.8%. Similar improvements were obtained for the three test case geometries in both

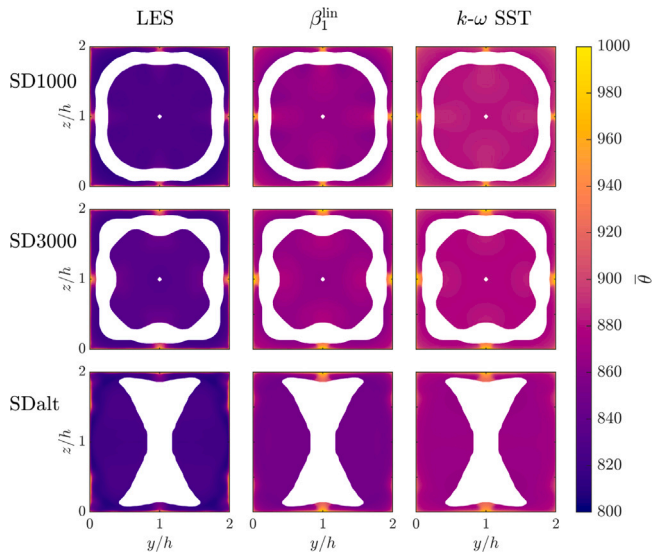


**Fig. 11.** Cross section of temperatures at streamwise location $x/L = 0.75$ for the testing geometries SD100, SD3000 and SDalt for: (left) high fidelity LES; (middle) linear coefficient model, $\beta_1^{\text{lin}}$; and (right) baseline $k-\omega$ SST.

first and second order effects, highlighting the robust nature of the linear coefficient closures.

The study demonstrates the capability of simple data-driven turbulence closures to improve second order heat transfer effects in incompressible flows, where the temperature may be modelled as a passive scalar. Further, the study shows that instead of using a forest of simple simulations to train models, it is possible to improve predictions for the same class of problems by training on a single complex geometry.

## CRediT authorship contribution statement

**James Hammond:** Conceptualization, Methodology, Software, Investigation, Writing – original draft, Writing – review & editing. **Marco Pietropaoli:** Resources, Writing – original draft. **Francesco Montomoli:** Conceptualization, Writing – original draft, Supervision, Funding acquisition.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Data availability

Data will be made available on request.

## Acknowledgements

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC), UK Grant EP/L016230/1. The authors would like to acknowledge Baker Hughes who support this research. In particular, a huge thanks to Prof. Vittorio Michelassi. Further thanks goes to Prof R. D. Sandberg for his continued advice and support.

## References

Brener, B.P., Cruz, M.A., Thompson, R.L., Anjos, R.P., 2021. Conditioning and accurate solutions of Reynolds average Navier-Stokes equations with data-driven turbulence closures. J. Fluid Mech. 915, 1–27. http://dx.doi.org/10.1017/jfm.2021.148.

Bunker, R.S., 2009. The effects of manufacturing tolerances on gas turbine cooling. J. Turbomach. 131 (4), 1–11. http://dx.doi.org/10.1115/1.3072494.

Frey Marioni, Y., Ortiz, E.A.d.T., Cassinelli, A., Montomoli, F., Adami, P., Vazquez, R., 2021. A machine learning approach to improve turbulence modelling from DNS data using neural networks. Int. J. Turbomach. Propul. Power 6 (2), http://dx.doi.org/10.3390/IJTPP6020017.

Gatski, T.B., Jongen, T., 2000. Nonlinear eddy viscosity and algebraic stress models for solving complex turbulent flows. Prog. Aerosp. Sci. 36 (8), 655–682. http://dx.doi.org/10.1016/S0376-0421(00)00012-9.

Gorlé, C., Zeoli, S., Emory, M., Larsson, J., Iaccarino, G., 2019. Epistemic uncertainty quantification for Reynolds-averaged Navier-Stokes modeling of separated flows over streamlined surfaces. Phys. Fluids 31 (3), http://dx.doi.org/10.1063/1.5086341.

Hammond, J., Frey Marioni, Y., Montomoli, F., 2022a. Error Quantification for the Assessment of Data-Driven Turbulence Models. Flow Turbul. Combust. http://dx.doi.org/10.1007/s10494-022-00321-1.

Hammond, J., Montomoli, F., Pietropaoli, M., Sandberg, R.D., Michelassi, V., 2020. Machine Learning for the Development of Data Driven Turbulence Closures in Coolant Systems. In: Turbomachinery Technical Conference and Exposition. Virtual, Online, pp. 1–12.

Hammond, J., Pietropaoli, M., Montomoli, F., 2022b. Topology optimisation of turbulent flow using data-driven modelling. Struct. Multidiscip. Optim. 65 (2), 1–21. http://dx.doi.org/10.1007/s00158-021-03150-4.

Huijing, J.P., Dwight, R.P., Schmelzer, M., 2021. Data-driven RANS closures for three-dimensional flows around bluff bodies. Comput. & Fluids 225, 104997. http://dx.doi.org/10.1016/j.compfluid.2021.104997.

Kaandorp, M.L., Dwight, R.P., 2020. Data-driven modelling of the Reynolds stress tensor using random forests with invariance. Comput. & Fluids 202, 104497. http://dx.doi.org/10.1016/j.compfluid.2020.104497.

Ling, J., Kurzawski, A., Templeton, J., 2016. Reynolds averaged turbulence modelling using deep neural networks with embedded invariance. J. Fluid Mech. 807, 155–166. http://dx.doi.org/10.1017/jfm.2016.615.

Ling, J., Templeton, J., 2015. Evaluation of machine learning algorithms for prediction of regions of high Reynolds averaged Navier Stokes uncertainty. Phys. Fluids 27 (8), http://dx.doi.org/10.1063/1.4927765.

Menter, F.R., 1994. Two-equation eddy-viscosity turbulence models for engineering applications. AIAA J. 32 (8), 1598–1605. http://dx.doi.org/10.2514/3.12149.

Menter, F.R., Kuntz, M., Langtry, R., 2003. Ten Years of Industrial Experience with the SST Turbulence Model Turbulence heat and mass transfer. AIAA J. 4 (July 2014), 625–632.

Parish, E.J., Duraisamy, K., 2016. A paradigm for data-driven predictive modeling using field inversion and machine learning. J. Comput. Phys. 305, 758–774. http://dx.doi.org/10.1016/j.jcp.2015.11.012.

Parneix, S., Laurence, D., Durbin, P.A., 1998. A procedure for using DNS databases. J. Fluids Eng. Trans. ASME 120 (1), 40–47. http://dx.doi.org/10.1115/1.2819658.

Pietropaoli, M., Gaymann, A., Montomoli, F., 2020. Three-dimensional fluid topology optimization and validation of a heat exchanger with turbulent flow. In: Proceedings of the ASME Turbo Expo, Vol. 7A. Virtual, Online, pp. 1–11. http://dx.doi.org/10.1115/GT2020-14479.

Pope, S.B., 1975. A more general effective-viscosity hypothesis. J. Fluid Mech. 72, 331–340. http://dx.doi.org/10.1017/S0022112075003382.

Rezazadeh Reyhani, M., Alizadeh, M., Fathi, A., Khaledi, H., 2013. Turbine blade temperature calculation and life estimation - a sensitivity analysis. Propul. Power Res. 2 (2), 148–161. http://dx.doi.org/10.1016/j.jppr.2013.04.004.

Schmelzer, M., Dwight, R.P., Cinnella, P., 2020. Discovery of Algebraic Reynolds-Stress Models Using Sparse Symbolic Regression. Flow Turbul. Combust. 104, 579–603. http://dx.doi.org/10.1007/s10494-019-00089-x.

Taghizadeh, S., Witherden, F.D., Girimaji, S.S., 2020. Turbulence closure modeling with data-driven techniques: Physical compatibility and consistency considerations. New J. Phys. 22 (9), http://dx.doi.org/10.1088/1367-2630/abadb3, arXiv:2004.03031.

Weatheritt, J., Sandberg, R.D., 2016. A novel evolutionary algorithm applied to algebraic modifications of the RANS stress-strain relationship. J. Comput. Phys. 325, 22–37. http://dx.doi.org/10.1016/j.jcp.2016.08.015.

Weatheritt, J., Sandberg, R.D., 2017. The development of algebraic stress models using a novel evolutionary algorithm. Int. J. Heat Fluid Flow 68 (November), 298–318. http://dx.doi.org/10.1016/j.ijheatfluidflow.2017.09.017.

Weatheritt, J., Zhao, Y., Sandberg, R.D., Mizukami, S., Tanimoto, K., 2019. Data-driven scalar-flux model development with application to jet in cross flow. Int. J. Heat Mass Transfer 147 (118931), http://dx.doi.org/10.1016/j.ijheatmasstransfer.2019.118931.

Wu, J.L., Xiao, H., Sun, R., Wang, Q., 2019. Reynolds-averaged Navier-Stokes equations with explicit data-driven Reynolds stress closure can be ill-conditioned. J. Fluid Mech. 869, 553–586. http://dx.doi.org/10.1017/jfm.2019.205.

Zhao, Y., Akolekar, H.D., Weatheritt, J., Michelassi, V., Sandberg, R.D., 2020. RANS turbulence model development using CFD-driven machine learning. J. Comput. Phys. 411, http://dx.doi.org/10.1016/j.jcp.2020.109413.