Quantification of Turbulence Modeling Uncertainties Using Full Field Inversion

Eric Parish[*] and Karthik Duraisamy[†]

Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI 48109.

Abstract

Reynolds Averaged Navier-Stokes (RANS) turbulence models have historically been developed using a combination of theoretical/physical/mathematical arguments, modeler expertise, and empirical data-fitting. As a consequence of the loss of information incurred during the ensemble averaging process, RANS model development has always been data-driven, to an extent, out of necessity. With the profusion of high resolution simulations and experimental methods, there exists a prodigious opportunity to more comprehensively inform closure models with data and improve their utlility in practical predictions of complex turbulent flows. In contrast to inferring model parameters, this work uses inverse modeling to directly obtain corrective, spatially distributed terms that are consistent with the predictive RANS models, thus offering a route to address the issue of model discrepancy which has long plagued turbulence model development. The posterior distributions of the inferred terms are quantified using Bayesian inversion and sampling. Results are presented for turbulent channel flow and flows with wall curvature. While the source terms inferred in this work are problem dependent, an ongoing subset of our work is to use machine learning methodologies to obtain general functional forms of the type that are inferred in this work. This paper specifically focuses on the mechanics of Bayesian inversion and its implications on the quantification of errors and uncertainties in turbulence modeling.

I. Introduction

The growth in computing power over the past several decades has allowed computational modeling to become one of the primary predictive tools in engineering and applied sciences. High fidelity computational fluid dynamics methods, namely direct numerical simulation (DNS) and large eddy simulation (LES), have provided previously unattainable insights into the intricate physical processes involved in a turbulent flow. However, their utility is still limited to low Reynolds number or free shear flows flows due to constraints on computing power and memory requirements. Moreover, the use of such high resolution methods for important design problems will remain impractical for the foreseeable future.[1] As such, these high fidelity methods are not tractable for an industrial design process, where a large number of functional evaluations are required. Instead, the mainstay is to either solve the Reynolds Averaged Navier-Stokes (RANS) equations or use a RANS-type near-wall approximation within an LES. Closure models for RANS are traditionally derived using a combination of theoretical, physical and mathematical arguments along with intuition and empiricism. In this sense, RANS modeling has always been data-driven. Despite the drastic simplifications of the governing equations of motion, RANS models have found utility in several engineering problems. RANS methods, however, have been found to be lacking in accuracy for flows with adverse pressure gradients, separation, curvature/rotations, transitional flow, etc. Such shortcomings severely limit the use of CFD simulations in the aerospace design process.[1]

The development of methodologies to improve turbulence models is an active area of research with many recent efforts pursuing advanced theories.[2–5] On the other hand, the growth in computing power has led to the availability of vast amounts of high-fidelity data and concurrent developments in experimental techniques have made high resolution flow-field measurements possible. The availability of such high-quality data presents a great opportunity for the use of data-driven methods to improve closure models. While the idea to use data-driven methods to assist in model development is obvious, the details on its implementation are not as intuitive. Model discrepancy[6] limits the ability to

---

[*]PhD Candidate, Department of Aerospace Engineering, University of Michigan, Ann Arbor

[†]Assistant Professor, Department of Aerospace Engineering, University of Michigan, Ann Arbor

American Institute of Aeronautics and Astronautics

to parameterize high resolution data in a way that is useful for model development. Instead, methods that produce data directly correlated to the low resolution model need to be considered. Much of the data-driven approaches that have been applied in turbulence modeling have involved the calibration of a set of scalar model coefficients. Yarlanki et al.,[7] Kato and Obayashi,[8] and Edeling et al.[9] all apply modern data-assimilation techniques to infer information about model coefficients. Dow et al.[10] infer a multiplier of the eddy viscosity relationship. From these works, the limitations of tuning model coefficients are becoming clear. Instead, it may be necessary to take a more comprehensive approach to the closure and move beyond the inference of calibration parameters. One such approach could be to infer the full field of discrepancy (for instance, as an additional source term). In this work, we present a subset of our ongoing work[11] that uses full field inversion to infer the spatio-temporal form of model corrections. Once the field variables are inferred over a sufficiently large and diverse set of problems, machine learning techniques are used to relate the corrections to RANS features, such that they can be used in predictive modeling. This paper deals with the inversion aspect of our approach, with a deeper investigation on the mechanics of Bayesian inversion and its implications for uncertainty quantification. In this work, existing turbulence models (Spalart-Allmaras,[12] $k-\omega$,[13] and BL-$\overline{v^2}/k$,[14]) are considered. The inversion is applied to fully developed turbulent channel flow and flows with concave and convex curvature.

II.   Inversion

In our approach, we seek to comprehensively improve closure models by inferring and reconstructing deficiencies in the functional form of turbulence models, rather than estimating model parameters. Bayesian inversion is used to directly obtain functional corrections in the form of additive source terms or multiplicative functions that are consistent with the underlying closure model. As a demonstrative example consider a turbulence model for a transport variable Q,

$$\frac{DQ}{Dt} = P - D + M \tag{1}$$

where P, D, and M could be production, destruction, and diffusion, respectively. A spatial multiplicative field variable β can be introduced into Equation 1 to create a modified turbulence model given by

$$\frac{DQ}{Dt} = \beta(x)P - E + D. \tag{2}$$

Note that the inclusion of β as a multiplier to the production term does not mean that we are restricting the form of the model corrections. It is indeed equivalent to adding an independent additive source term. The purpose of the inversion is to infer β such that solution predicted by the modified turbulence model is closer to a given benchmark solution. It has to be recognized that adding a correction term inherently changes the entire model, as opposed to just one term. Although the resulting model is only valid/accurate for the problem used to generate it, attempts to generalize it using machine learning are a part of our ongoing work.[11, 15]

A Bayesian approach[16] is used to solve the inverse problem. β is considered to be a random function with a certain probability distribution. Using Bayes theorem we can relate the conditional probability distribution of β given benchmark data $\mathbf{d}$ (defined as $q(\beta|\mathbf{d})$) with the likelihood of a solution to be $\mathbf{d}$, given β (defined as $h(\mathbf{d}|\beta)$) , and the prior distribution of β ($p(\beta)$)

$$q(\beta|\mathbf{d}) = \frac{h(\mathbf{d}|\beta)p(\beta)}{c} \tag{3}$$

where $c = \int h(\mathbf{d}|\beta)p(\beta)d\beta$. The computation of the probability distribution using Equation 3 is difficult because of the integration in the denominator, which involves a large number of dimensions. Furthermore, single evaluation of h itself can be expensive; in this work h involves solving the set of RANS equations. However, assumptions regarding the distribution of $\mathbf{d}$ and β can lead to simplifications which can be handled in practice. If we can assume that β and the dataset $\mathbf{d}$ are normally distributed and that the distribution $q(\beta|\mathbf{d})$ is Gaussian, it can then be shown[16] that the problem of finding the distribution of β can be reduced to estimating the *maximum a posteriori* (MAP) solution which is found by solving a deterministic optimization problem

$$\beta_{map} = arg \min \frac{1}{2}\left[\left(\mathbf{d}-h(\beta)\right)^T \mathbf{C_m}^{-1}\left(\mathbf{d}-h(\beta)\right) + \left(\beta - \beta_{prior}\right)^T \mathbf{C}_\beta^{-1}\left(\beta - \beta_{prior}\right)\right] \tag{4}$$

where $\mathbf{C_m}$ and $\mathbf{C}_\beta$ are the observational and prior covariance matrices, respectively. The term being minimized is referred to as the cost function; i.e.

$$\beta_{map} = arg \min \mathfrak{J}. \tag{5}$$

American Institute of Aeronautics and Astronautics

In this work we take **d** to be a subset of observational (high fidelity simulation or experimental) data. Linearizing about the MAP point, the covariance of the posterior is given by the inverse of the Hessian of the cost function[a] $\mathfrak{J}$

$$\mathbf{C}_{\beta_{MAP}} = \mathbf{H}^{-1}\Big|_{\beta_{MAP}} = \left(\frac{d^2\mathfrak{J}}{d\beta_i d\beta_j}\right)^{-1}\Big|_{\beta_{MAP}}. \tag{6}$$

With the MAP solution and covariance of the posterior, realizations of β can be drawn from the posterior distribution. A standard procedure is to first perform a Cholesky decomposition on $\mathbf{C}_{\beta_{MAP}}$ such that

$$\mathbf{R}\mathbf{R}^{\mathbf{T}} = \mathbf{C}_{\beta_{MAP}}. \tag{7}$$

The posterior is then sampled through,

$$\beta = \beta_{MAP} + \mathbf{R}\mathbf{s} \tag{8}$$

where **s** is a vector whose components are independent standard normal variates. The inversion involves three major tasks: (a) calculation of the objective function gradient, (b) the optimization algorithm which utilizes the gradient to minimize the objective function, and (c) computation of the Hessian matrix. It is well-recognized that in non-linear problems, the posterior distribution will not be Gaussian in nature and that Equation 3 is an approximation, we consider it to be necessary step - and not one that is inaccurate - to make our approaches feasible in large dimensional problems.

II.A.   Computational Framework for Inversion

The key computational challenges in the inverse are the high dimensional optimization problem that stems from minimizing Equation 4 and computation of the Hessian, as we are targeting functional inversion and thus the number of dimensions scales with number of mesh points. Here gradient-based methods are used to solve the inverse problem. These methods require derivatives with respect to a large number of parameters, which are efficiently calculated using discrete adjoints. The gradients are determined by solving the adjoint equation

$$\left(\frac{\partial R}{\partial u}\right)^T \psi = -\left(\frac{\partial \mathfrak{J}}{\partial u}\right)^T \tag{9}$$

where $\mathfrak{J}$ is the cost function, R is the residual of the primal equations, and u are the flow variables. The gradient is then computed by

$$G = \frac{d\mathfrak{J}}{d\beta} = \frac{\partial \mathfrak{J}}{\partial \beta} + \psi^T \frac{\partial R}{\partial \beta}. \tag{10}$$

The optimization problem is solved using BFGS,[17] or in problematic cases, using steepest decent. The solution of the adjoint equation requires the computation of the Jacobian of the primal equations. These Jacobians are either computed by automatic differentiation (AD) or analytically. In the case of AD, the engine Tapenade[18] is used for source code generation. The sensitivities obtained from adjoints are validated using finite difference and complex step method.

In this work both an exact and approximate computation of the Hessian are used. The exact Hessian is computed using an adjoint-adjoint method. For a system with M flow variables and N design variables, the Hessian is computed by:[19]

$$H_{ij} = \frac{\partial^2 \mathfrak{J}}{\partial \beta_i \partial \beta_j} + \psi_m \frac{\partial^2 R_m}{\partial \beta_i \partial \beta_j} + \mu_{i,m} \frac{\partial R_m}{\partial \beta_j} + \nu_{i,n} \frac{\partial^2 \mathfrak{J}}{\partial u_n \partial \beta_j} + \nu_{i,n} \psi_m \frac{\partial^2 R_m}{\partial u_n \partial \beta_j} \quad m,n \in [1,M] \tag{11}$$

where,

$$\nu_{i,n} \frac{\partial R_m}{\partial u_n} = -\frac{\partial R_m}{\partial \beta_i} \quad m,n \in [1,M] \quad i \in 1,N] \tag{12}$$

$$\mu_{i,m} \frac{\partial R_m}{\partial u_k} = -\frac{\partial^2 F}{\partial \beta_i \partial u_k} - \psi_m \frac{\partial^2 R_m}{\partial \beta_i \partial u_k} - \nu_{i,n} \frac{\partial^2 \mathfrak{J}}{\partial u_n \partial u_k} - \nu_{i,n} \psi_m \frac{\partial^2 R_m}{\partial u_n \partial u_k} \quad k,m,n \in [1,M] \quad i \in [1,N] \tag{13}$$

Following Ref. 19, Equation 12 is first solved N times for ν. Equation 13 is then solved for the N rows of μ. While more efficient algorithms exist for computing the Hessian, an exact calculation is always expensive. A low-rank

---

[a]evaluated at the MAP point

approximation is additionally used to approximate the Hessian.[16] For a cost function with a diagonal observational covariance matrix and no prior the cost function simplifies to

$$\mathfrak{J} = \frac{1}{2}\left[\sum_i^M \left(\frac{h(\beta)_i - d_i}{\sigma_i}\right)^2\right]. \tag{14}$$

To approximate the Hessian, M scalar valued functions are defined,

$$f_i(\beta) = \frac{h(\beta)_i - d_i}{\sigma_i} \quad \text{for } i = 1, 2, ..., M. \tag{15}$$

The gradient of the non-regularized cost function can computed by

$$\nabla\mathfrak{J}(\beta) = \sum_{i=1}^M \nabla f_i(\beta)^2 = \sum_{i=1}^M 2f_i(\beta)\nabla(f_i(\beta)). \tag{16}$$

Or equivalently

$$\nabla\mathfrak{J}(\beta) = 2\mathbf{J}(\beta)^T\mathbf{F}(\beta) \tag{17}$$

where

$$\mathbf{F} = \begin{bmatrix} f_1(\beta) & f_2(\beta) & ... & f_M(\beta) \end{bmatrix}^T. \tag{18}$$

The Jacobian of the scalar valued functions can then be used to approximate the Hessian by

$$\mathbf{H} \approx \mathbf{J}(\beta)^T\mathbf{J}(\beta). \tag{19}$$

The contribution of the prior is easily computed analytically. It is noted that calculation of the low-rank Hessian involves performing M adjoints for each of the scalar valued functions.

III.   Turbulent Channel Flow Calculations

Planar internal channel flow is considered for the first set of test cases. The Reynolds averaged momentum equation for incompressible fully-developed channel flow is given by

$$0 = -\frac{\partial\overline{p}}{\partial x} + \frac{\partial}{\partial y}\left[\mu\frac{\partial\overline{u}}{\partial y} - \rho\overline{u'v'}\right]. \tag{20}$$

The Reynolds stresses are closed with a turbulence model.

III.A.   Application to the $k - \omega$ Model

First we consider the Wilcox $k - \omega$[13] model which adds two additional model equations to the system. For the planar channel the closure model becomes

$$0 = \nu_t\left(\frac{\partial\overline{u}}{\partial y}\right)^2 - \alpha^* k\omega + \frac{\partial}{\partial y}\left[\left(\nu + \sigma^*\frac{k}{\omega}\right)\frac{\partial k}{\partial y}\right] \tag{21}$$

$$0 = \gamma\left(\frac{\partial\overline{u}}{\partial y}\right)^2 - \alpha\omega^2 + \frac{\partial}{\partial y}\left[\left(\nu + \sigma\frac{k}{\omega}\right)\frac{\partial\omega}{\partial y}\right]. \tag{22}$$

No viscous corrections are used. The inverse is posed by modifying the production term in the turbulent kinetic energy equation to include the spatial multiplier β,

$$0 = \nu_t\left(\frac{\partial\overline{u}}{\partial y}\right)^2\beta(y) - \alpha^* k\omega + \frac{\partial}{\partial y}\left[\left(\nu + \sigma^*\frac{k}{\omega}\right)\frac{\partial k}{\partial y}\right]. \tag{23}$$

The governing equations are discretized using second order second finite differences and solved by introducing a pseudo-time derivative. A geometrically graded mesh is used such that the viscous sublayer is completely resolved.

American Institute of Aeronautics and Astronautics

The first mesh point is placed at approximately $y^+ = 0.05$. A channel flow solver[13] was used as a guideline for the mesh spacing.

The inversion requires the specification of the covariance matrices $\mathbf{C_m}$ and $\mathbf{C_\beta}$. For planar channel flow DNS data are readily available and provide almost perfect observations of the true *mean* flow variables, hence we take $\mathbf{C_m}$ to be diagonal with constant entries of $\sigma_{obs} = 1e - 10$. Such a choice neglects the covariance of the data. Selection of $\mathbf{C_\beta}$ is less straightforward as there is very little knowledge of the prior. Here $\mathbf{C_\beta}$ is taken to be diagonal with constant variance σ_p. The variance is selected such that the DNS solution falls within the locus of the $\pm 2\sigma$ distribution of profiles predicted by the prior model. For all cases it was determined $\sigma_p = 0.5$. At this point we note that this selection corresponds to an uninformed prior. As information about the prior is gained from solving inverse problems, the specification of the prior can become considerably more complex. Note that, with $\sigma_{obs} = 1e - 10$ and $\sigma_p = 0.5$, the dependence on the prior has effectively been eliminated for regions where the inferred function is very sensitive to the data. Specification of a reasonable prior, however, is still important as the overall problem is still is ill-posed.

DNS data from the Jimenez group[20] are used in the inverse. Of this data, the velocity profiles are targeted; i.e. $\mathbf{d} = \mathbf{u}_{DNS}$. For clarity, the entire solution process used is as follows:

1. Sample the prior distribution of β via Equations 7 and 8 with the prior covariance matrix $\mathbf{C_\beta}$.

   (a) For each sample of β solve the governing equations to determine prior distributions for all flow variables.

   (b) Determine if $\mathbf{C_\beta}$ was appropriately chosen by ensuring that the targeted DNS profile falls within the $\pm 2\sigma$ limits of the distribution predicted by the prior turbulence model.

2. Solve the inverse problem with $\mathbf{C_m} = \sigma_{obs}^2 I$ and with $\sigma_{obs} = 1e - 10$.

3. Sample the posterior distribution of β_{MAP} via Equations 7 and 8 with $\mathbf{C_{\beta_{MAP}}}$.

   (a) For each sample of β_{MAP} solve the governing equations to determine the posterior distributions for all flow variables about the MAP point.

A summary of the cases considered is given in Table 1.

Table 1: Summary of channel flow cases.

Case	Re_τ	Grid Points	$\Delta x_e / \Delta x_f$	σ_p	σ_{obs}
1	186.3	201	80	0.5	1×10^{-10}
2	546.7	201	305	0.5	1×10^{-10}
3	933.9	201	600	0.5	1×10^{-10}
4	2004.3	201	1450	0.5	1×10^{-10}
5	4200	201	3340	0.5	1×10^{-10}

The results of the inversion are similar for the different cases. Specifically, we will discuss the entire inference process for $Re_\tau = 934$. As previously noted, the covariance of the prior C_β was selected such that the targeted DNS solution (in this case velocity) fell within the locus of the 2σ limits of the prior distribution. From this process the prior covariance was determined to be $C_\beta = \sigma_p^2 I$ with $\sigma_p = 0.5$. The prior mean is taken to be $\beta_{prior} = 1$, which corresponds to the baseline model. Random draws of β are taken from the multivariate PDF to generate the probability distributions for other flow variables. This prior distribution determined from 100 draws is given in Figure 1.

The goal of the inversion now is twofold; to determine the MAP solution and to determine the posterior distribution. Figure 2 gives a summary of the inversion and compares it to the DNS data that were used in the inverse. 100 draws of the posterior distribution of β we taken to construct the PDFs of other flow variables. Figure 3 additionally shows the posterior PDFs of inferred β at different y^+ locations; since sampling β is trivial 10,000 samples were used. The inversion was able to infer a correction term that decreased the magnitude of the objective function by over two orders of magnitude; as such the match between the targeted velocity profiles is excellent. The inverse also has a positive effect on secondary variables that were not targeted. The turbulent kinetic energy predicted by the inverse model has a more pronounced peak near the wall which is more representative of underlying physics. The turbulent production term in the inverse almost perfectly matches the DNS. It is noted that a composite objective function including variables

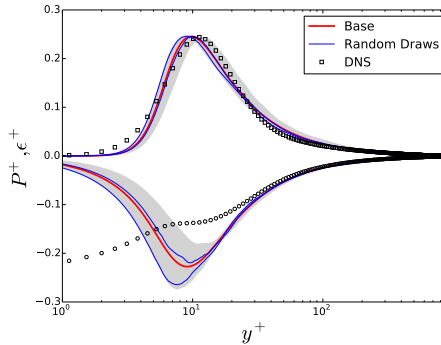American Institute of Aeronautics and Astronautics
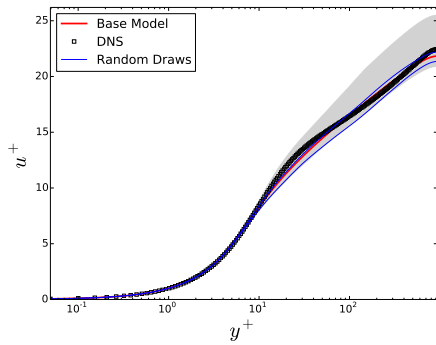
(a) Prior distribution for velocity with $\pm 2\sigma$ limits shaded.
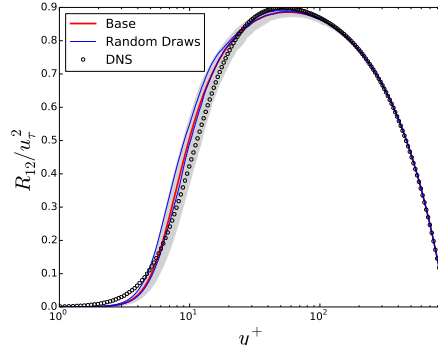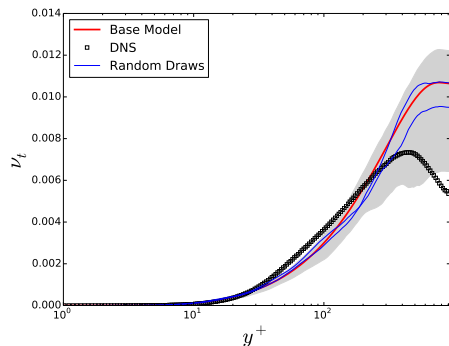

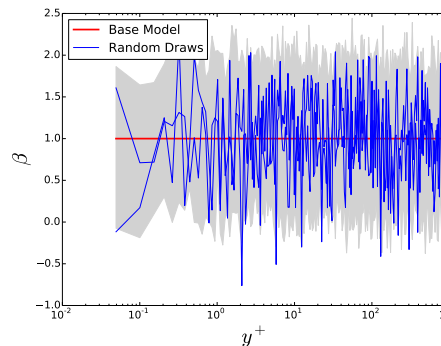(b) Prior distribution for Reynolds stress with $\pm 2\sigma$ limits shaded.
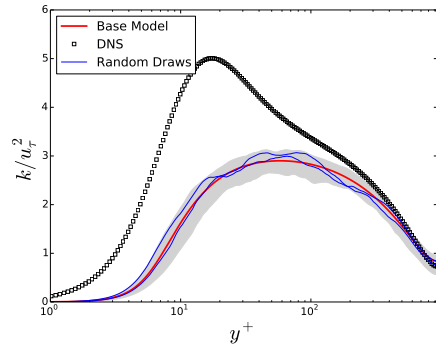
(c) Prior distribution for TKE with $\pm 2\sigma$ limits shaded.

(d) Prior distribution for turbulent energy production and dissipation.

(e) Prior distribution for eddy-viscocity $\pm 2\sigma$ limits shaded compared to effective DNS eddy-viscosity.

(f) Prior distribution for β with $\pm 2\sigma$ limits.

Figure 1: Prior (base $k-\omega$ model) distributions with random draws for planar channel flow at $Re_\tau = 934$ with $C_\beta = \sigma_p^2 I$.

other than just the velocity would greatly improve the prediction of secondary variable such as turbulent kinetic energy near the wall.

The learned β profile undergoes various changes of turbulent production and destruction in the inner layer. Away from the wall in the channel core, β settles below unity, corresponding to an increase in turbulent destruction.

The posterior distributions predicted through the inverse are of interest. First it is seen that the σ limits for β in the viscous sublayer are *identical to the prior*, hence the inverse has provided us with no information about what β should be in the viscous sublayer. This is a result of the ill-posedness of the inverse problem and provides an excellent example of the importance of the specification of the prior. Additionally, this result is completely consistent with what one would expect. In the viscous sublayer turbulent production becomes zero, hence the value of β is of no consequence. In the overlap and outer layers the variance on β collapses to $\sigma = O(10^{-3})$, implying that much information about β can be inferred in these regions. This is further seen in Figure 3. The Gaussian approximation also

American Institute of Aeronautics and Astronautics

appears to be reasonable in this context. The posterior distributions for all other flow variables all but collapse to the MAP solution after inversion. The posterior distribution predicted for velocity is quite small and has a mean standard deviation of $\sigma \approx 1e - 10$, which is of the same order as the observational variance. In summary, the uncertainties on all flow variables in the post-inversion model are greatly decreased. The posterior distribution for β indicates that, for all regions other than the viscous sublayer, the inversion provided substantial information of what β should be.

It is important to note that it is not critical for the DNS profiles of secondary flow variables to fall within the distribution of the model obtained in the inversion. The posterior distribution instead provides information about the inferred β with respect to the effectiveness of the inverse. Such information is useful to the modeler and as inputs to machine learning algorithms to reconstruct discrepancy terms.[11] As an example, an ongoing subset of our work is to parameterize β in terms of flow variables. The information about the posterior distribution of β allows us to realize that β_{MAP} in the viscous sublayer should not be used in machine learning parameterization since no information could be inferred in the inverse.

Figure 4 provides a summary of inferred β for the five cases. The collapse of the data for the cases, which span a wide range of Reynolds numbers, is excellent and coincides with physical arguments. In the inner layer β behaves similarly for all cases and varies with dimensionless wall distance y^+. Away from the wall the effects of outer variables become important and a uniform scaling with y^+ is not expected. For $y^+ > 100$, β begins to scale with y/H and a decent collapse is again observed. While this work focuses on the inversion step, it is worthwhile to show that is is reasonable to believe that the inferred β can be parameterized in terms of RANS features. An important feature in the $k - \omega$ model, and one that is used extensively for low Reynolds number modifications, is the turbulence Reynolds number which is given by

$$Re_T = \frac{k}{\omega \nu}. \tag{24}$$

The turbulence Reynolds number provides a measure of the local strength of the turbulent stresses compared to the laminar stresses. Figure 4 additionally shows inferred β in the inner layer ($y^+ < 300$) as a function of Re_T. The collapse of the data with the local turbulence number is comparable to the collapse with y^+; the key difference being Re_T is a local flow feature and independent of geometry. While the full extent of our work uses machine learning algorithms to obtain a functional form of β, these simple scalings can already be of use to the modeler. Furthermore, the decent collapse of data with dimensionless quantities shows that it is reasonable to believe that machine learning methods can infer a general functional form for the correction.

III.B. Application to the BL-$\overline{v^2}/k$ Model

In the previous section the inversion process was used to infer a velocity-correction multiplier for the baseline $k - \omega$ model without viscous corrections. It is well argued, however, that accurate velocity prediction is not a feature that turbulence models find great difficulty with for channel flows. Near-wall predictions of turbulent kinetic energy and dissipation are more challenging tasks. These near-wall flow features are both non-homogeneous and anisotropic; as such they can not be accurately captured by simple turbulence models that use quasi-homogeneous assumptions unless ad-hoc damping functions are used. We now turn our attention to the BL-$\overline{v^2}/k$ model[14] , which is a variant of the $k - \varepsilon - \overline{v^2} - f$ model family. In addition to k and ε, an additional transport equation is solved for $\phi = \overline{v^2}/k$ and an elliptic equation is solved for a blending parameter α. The BL-$\overline{v^2}/k$ model is as follows:

$$0 = P - \varepsilon_h + \frac{\partial}{\partial y}\left[\left(\frac{\nu}{2} + \frac{\nu_t}{\sigma_k}\right)\frac{\partial k}{\partial y}\right] - C_{\varepsilon 3}(1-\alpha)^3 \frac{k}{\varepsilon_h} 2\nu\nu_t\left(\frac{\partial^2 u}{\partial y^2}\right)^2 \tag{25}$$

$$0 = \frac{C_{\varepsilon 1}P - C_{\varepsilon 2}\varepsilon_h}{T} + \frac{\partial}{\partial y}\left[\left(\frac{\nu}{2} + \frac{\nu_t}{\sigma_{\varepsilon_h}}\right)\frac{\partial \varepsilon_h}{\partial y}\right] \tag{26}$$

$$0 = (1-\alpha^3)f_w + \alpha^3 f_h - P\frac{\phi}{k} + \frac{2}{k}\frac{\nu_t}{\sigma_k}\frac{\partial \phi}{\partial y}\frac{\partial k}{\partial y} + \left[\left(\frac{\nu}{2} + \frac{\nu_t}{\sigma_\phi}\right)\frac{\partial \phi}{\partial y}\right] \tag{27}$$

$$\alpha - L^2\frac{\partial^2 \alpha}{\partial y^2} = 1 \tag{28}$$

$$f_w = -\frac{\varepsilon_h}{2}\frac{\phi}{k} \qquad f_h = -\frac{1}{T}\left(C_1 - 1 + C_2\frac{P}{\varepsilon_h}\right)\left(\phi - \frac{2}{3}\right) \tag{29}$$

American Institute of Aeronautics and Astronautics

$$L = \sqrt{C_L^2 \left(\frac{k^3}{\varepsilon_h^2} + C_\eta^2 \frac{\nu^{3/2}}{\varepsilon_h^{1/2}} \right)} \qquad T = \sqrt{\frac{k^2}{\varepsilon_h^2} + C_T^2 \frac{\nu}{\varepsilon_h}} \qquad T_{lim} = \frac{0.6}{\sqrt{6} C_\mu \phi \sqrt{S_{ij} S_{ij}}} \qquad (30)$$

$$\nu_t = C_\mu \phi k \min(T, T_{lim}) \qquad C_{\varepsilon_2}^* = C_{\varepsilon 2} + \alpha^3 (C_{\varepsilon 4} - C_{\varepsilon 2}) \tanh \left(\left| \frac{1}{\varepsilon_h} \frac{\partial}{\partial_j} \left[\frac{\nu_t}{\sigma_k} \frac{\partial k}{\partial_j} \right] \right| \right) \qquad (31)$$

To improve the performance of the BL-$\overline{v^2}/k$ model, the inverse is posed by adding the corrective multiplier to f_w,

$$f_w = -\frac{\varepsilon_h}{2} \frac{\phi}{k} \beta \qquad (32)$$

Such a choice ensures that β will be active in the near-wall region. Instead of only targeting velocity in the inverse, we now target both velocity and turbulent kinetic energy. Uncorrelated observational and prior covariances are again assumed. With these assumptions the cost function can be written as

$$\mathfrak{J} = \mathfrak{J}_u + \mathfrak{J}_k + \mathfrak{J}_\beta = \sum_{i=0}^{N} \left(\frac{u_i^{RANS} - u_i^{DNS}}{\sigma_u^{obs}} \right)^2 + \left(\frac{k_i^{RANS} - k_i^{DNS}}{\sigma_k^{obs}} \right)^2 + \sum_{j=0}^{M} \left(\frac{\beta_j - \beta_j^{prior}}{\sigma^{prior}} \right)^2. \qquad (33)$$

In Equation 33 it is seen that the variances act as weightings between the different components of the cost function. Since the observational variance of DNS is so low, it is fair to manipulate the variances to obtain a desired weight on each component of the cost function. This is not the case with lower fidelity observations, where the observational variance carries more physical significance. The observational variances were selected such that $\mathfrak{J}_k \approx 2\mathfrak{J}_u$ at the first iteration. A summary of the inversion details are given in Table 2. It is noted that BL-$\overline{v^2}/k$ gives accurate predictions with $y^+ \approx 1$, as such the governing equations are solved on a grid that is coarser than before.

Table 2: Summary of channel flow cases for the BL-$\overline{v^2}/k$.

Case	Re_τ	Grid Points	$\Delta x_e / \Delta x_f$	σ_p	σ_u^{obs}	σ_k^{obs}
6	186.3	61	8	0.5	1×10^{-10}	6×10^{-12}
7	546.7	81	31	0.5	1×10^{-10}	2×10^{-11}
8	933.9	101	31	0.5	1×10^{-10}	7×10^{-11}
9	2004.3	121	61	0.5	1×10^{-10}	3×10^{-10}

We will again discuss the results of the inversion for the moderate Reynolds number of $Re_\tau = 933$, which are given in Figure 5. First, it is observed that the inversion was not able to infer a correction that matched both velocity and turbulent kinetic energy exactly. In general it was observed that matching TKE in the buffer layer led to an under prediction of velocity in the same region. This is seen in the velocity profiles shown in Figure 5a. While the overall velocity prediction is much improved, a slight under prediction is present in the buffer layer. The prediction of TKE is much improved in this region. The posterior solution almost perfectly matches the DNS profile in the viscous sublayer and predicts a much more accurate peak within the buffer layer.

As previously mentioned, the choice to have β be a multiplier on f_w makes it active in the near wall region (since it is introduced through $(1 - \alpha^3)$). The correction β becomes negative at the edge of the viscous sublayer, then grows quickly throughout the buffer layer until it peaks towards the end of the buffer layer. It then rapidly decays throughout the log-layer, becoming negative around $y^+ \approx 30$ and eventually relaxing to $\beta = 1$ in the channel core. In the channel core f_w has a negligible influence on the solution. This is seen in the posterior distribution for β, where the posterior variance is small in the near wall region and increases in the channel core. The posterior distributions for all other variables again collapse to the map solution.

Figure 6 shows a summary of the inferred β profiles for the other cases. The inferred corrections again display similarities within the inner layer, varying with y^+. Since the effect of β is not strong away from the wall, $\beta \approx 1$ in the outer layer. For all cases, the oscillations in β seem to coincide with the viscous, buffer, and overlap layers. This again suggests that these solutions are rooted in the fundamental physics/modeling of wall-bounded flows; opposed to random solutions with no physical meaning.

American Institute of Aeronautics and Astronautics

The decision to have β be a multiplier on f_w is, again, by no means unique. The correction can be applied anywhere. Furthermore, it is possible to infer more than one correction in the inversion. While increasing the dimensionality of the overall correction poses a more daunting challenge to machine learning algorithms at introduces the possibility (and often a probability) that the corrections depends on each other, it needs to be recognized that the process is valid. To demonstrate an alternative setup both the length scale and $C_{\varepsilon2}^*$ coefficient are modified,

$$L = \sqrt{C_L^2\left(\frac{k^3}{\varepsilon_h^2} + C_\eta^2 \frac{v^{3/2}}{\varepsilon_h^{1/2}}\right)\beta_1} \qquad C_{\varepsilon2}^* = \beta_2 C_{\varepsilon2}. \qquad (34)$$

It is noted that in this case we consider the baseline $C_{\varepsilon2}^* = C_{\varepsilon2}$, opposed to the spatially varying value given in Equation 31. As such the baseline model over predicts v_t in the outer layer. A summary of the properties used in the inversion are given in Table 3. It is noted that the decision to have a multiplier on the length scale makes the inverse extremely ill-posed, as such a higher variance for the DNS data is used.

Table 3: Summary of channel flow cases for the BL-$\overline{v^2}/k$ with the setup given in Equation 34.

Case	Re_τ	Grid Points	$\Delta x_e/\Delta x_f$	σ_1^p	σ_2^p	σ_u^{obs}	σ_k^{obs}
6	186.3	61	8	0.1	0.5	1×10^{-7}	9×10^{-9}
7	546.7	81	31	0.1	0.5	1×10^{-7}	2.5×10^{-8}
8	933.9	101	31	0.1	0.5	1×10^{-7}	4.5×10^{-8}
9	2004.3	121	61	0.1	0.5	1×10^{-7}	1.7×10^{-7}

The results of the inversion at $Re_\tau = 933$ are given in Figure 7. It is again seen that, while corrections that exactly match both TKE and velocity could not be inferred, the posterior solutions are much improved. There are several features of interest. First, it is observed that the influence of β_2 is felt across the entire channel. This is a result of $C_{\varepsilon2}^*$ being non-zero everywhere. As such the posterior distribution for β_2 collapses to the MAP solution everywhere. β_2 now contains corrective information that is related to the channel centerline. Second we note that little can be inferred about β_1 outside of the buffer layer. This is credited to β_1 being introduced to the system only through the second derivative of the blending parameter α, which decays to zero in the channel core. Additionally it is believed that β_1 can not be inferred in the viscous sublayer because the turbulent activity is not very high. This result suggests that the turbulent kinetic energy and velocity predictions are insensitive to the length scale definition (in the elliptic equation for α) outside of the buffer layer.

Figure 8 shows a summary of the inferred β profiles for the other cases. These data don't collapse as well as what was seen in previous cases. It is noted that the inversion at $Re_\tau = 180$ did not perform well; the cost function was only decreased by approximately 20%. The differences between different Reynolds numbers could be due to a multitude of reasons, but it is believed that using multiple corrections decreased the uniqueness of the solution.
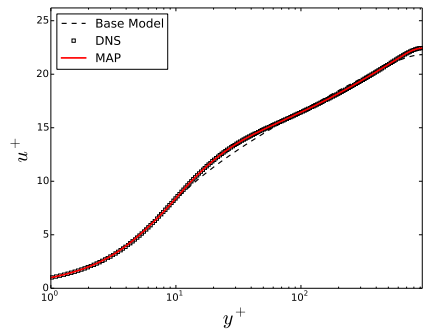
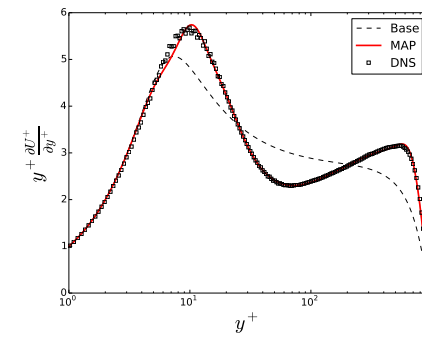American Institute of Aeronautics and Astronautics
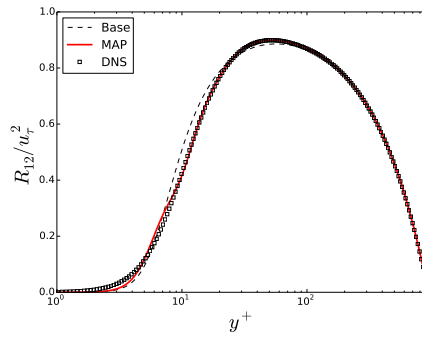
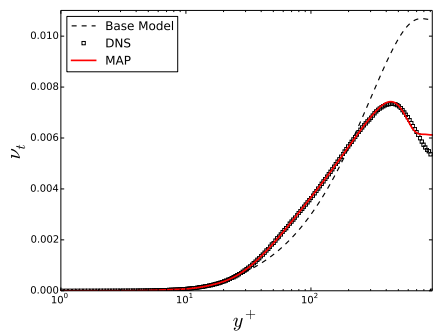(a) Posterior distribution for velocity with $\pm 2\sigma$ limits shaded.

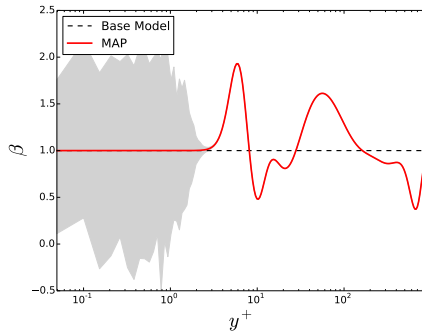(b) Posterior distribution for Reynolds stress with $\pm 2\sigma$ limits shaded.
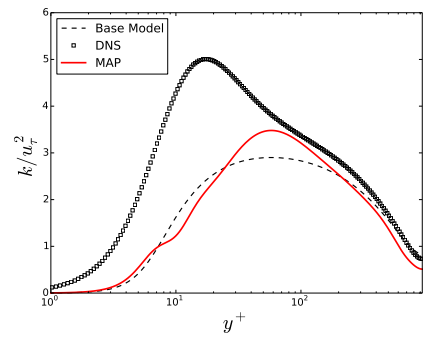
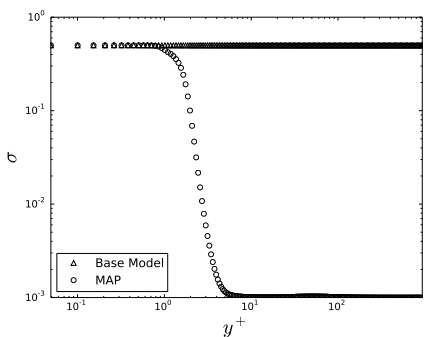(c) Posterior distribution for TKE with $\pm 2\sigma$ limits shaded.
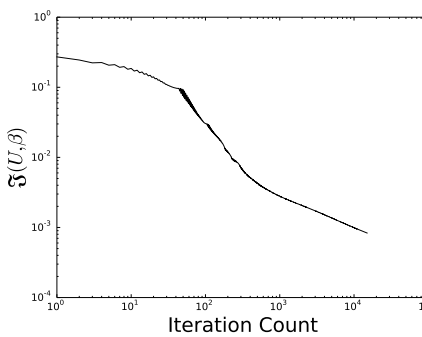
(d) Posterior distribution for $y^+ \frac{dU^+}{dy^+}$.

(e) Posterior distribution for eddy-viscocity $\pm 2\sigma$ limits shaded compared to effective DNS eddy-viscosity.

(f) Posterior distribution for β with $\pm 2\sigma$ limits.

(g) Posterior variance for β.

(h) Iteration history of the convergence of the cost function.

Figure 2: Posterior model distributions for planar channel flow at $Re_\tau = 934$ with $C_\beta = \sigma_p^2 I$.

American Institute of Aeronautics and Astronautics
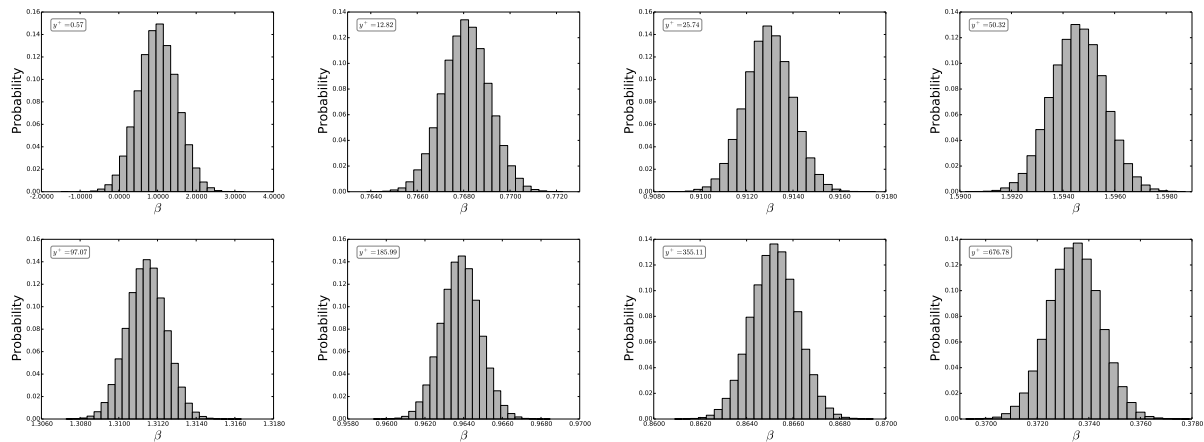
Figure 3: Posterior PDFs of β for planar channel flow at $Re_\tau = 934$ with $C_\beta = \sigma_p^2 I$. 10,000 samples of β were used to construct the PDFs.
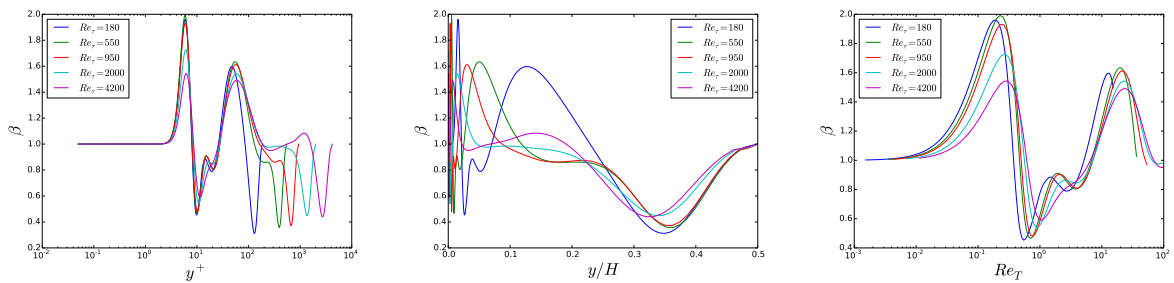

Figure 4: Summary for β's obtained in the inversion. It is seen that in the inner layer β varies with y^+, while in the outer layer β varies with y/H. A decent collapse is also observed with Re_T for data in the inner layer.
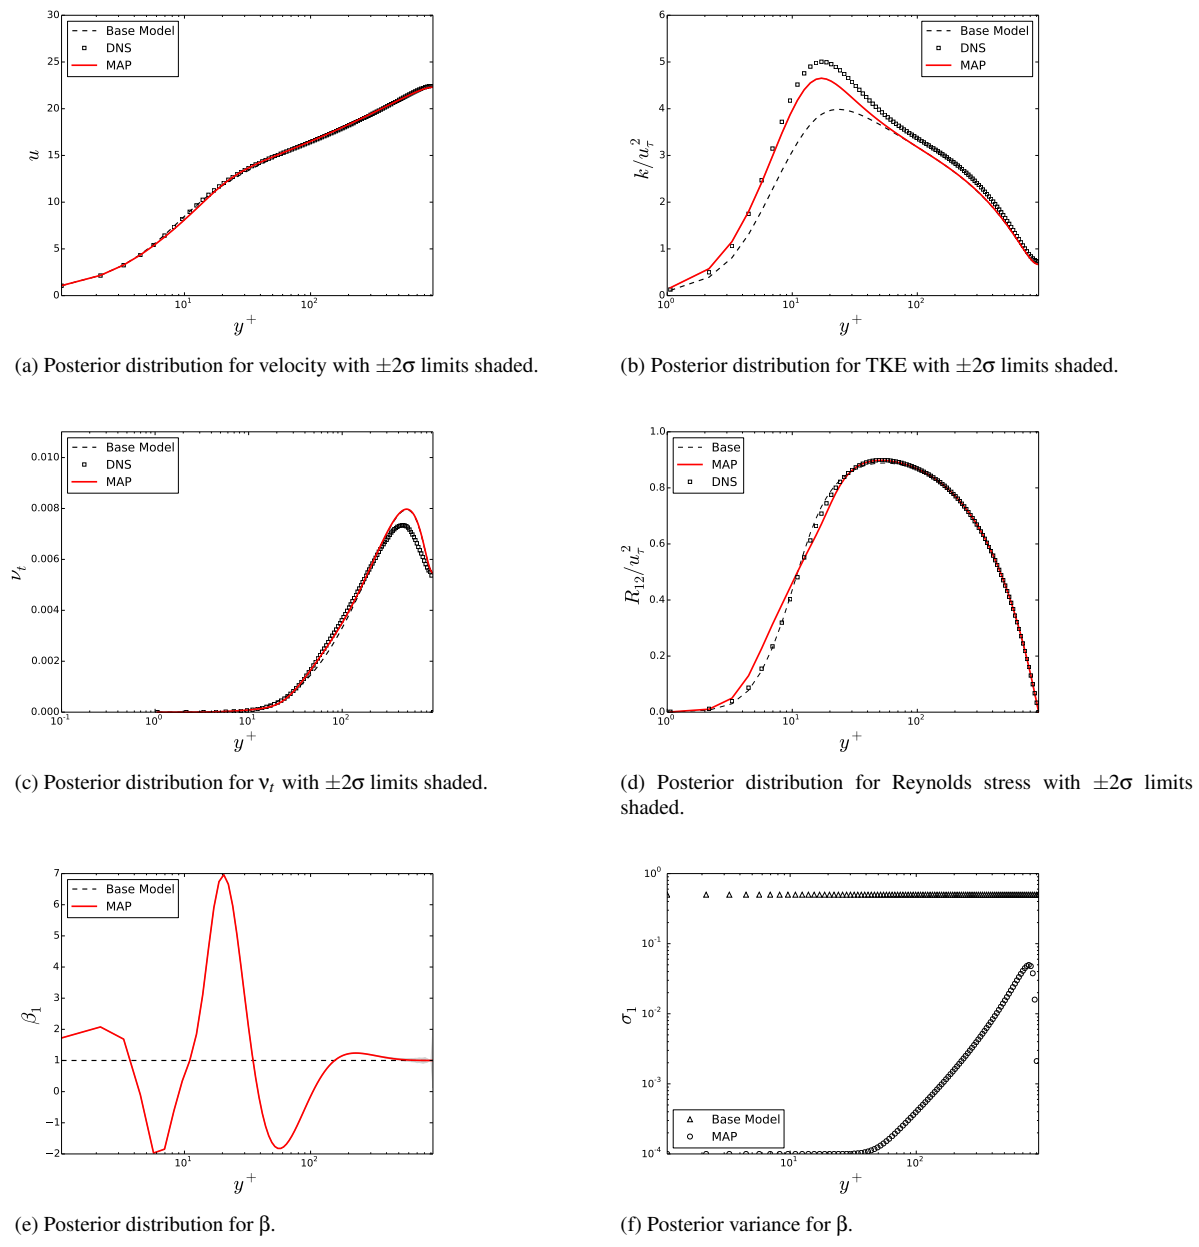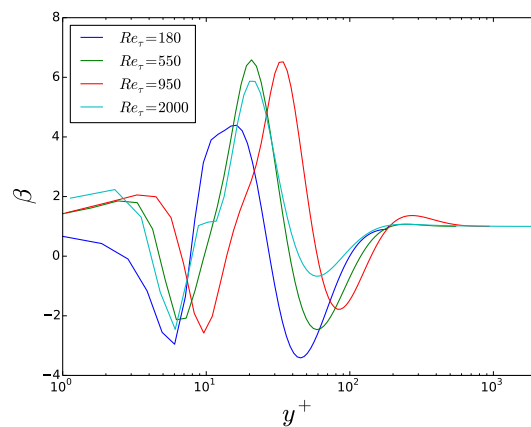
American Institute of Aeronautics and Astronautics

(a) Posterior distribution for velocity with $\pm 2\sigma$ limits shaded.

(b) Posterior distribution for TKE with $\pm 2\sigma$ limits shaded.

(c) Posterior distribution for ν_t with $\pm 2\sigma$ limits shaded.

(d) Posterior distribution for Reynolds stress with $\pm 2\sigma$ limits shaded.

(e) Posterior distribution for β.

(f) Posterior variance for β.

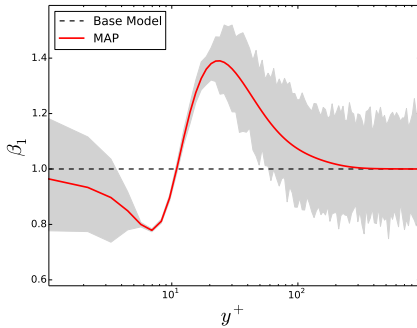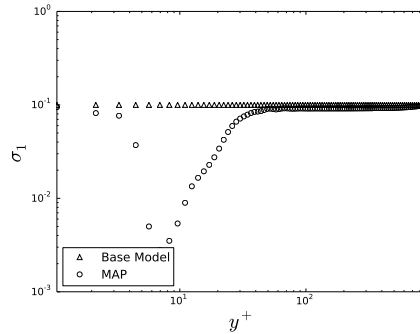Figure 5: Posterior model distributions for planar channel flow at $Re_\tau = 950$ with f_w modifier.

American Institute of Aeronautics and Astronautics

Figure 6: Summary for β's obtained in the inversion with an f_w modifier.
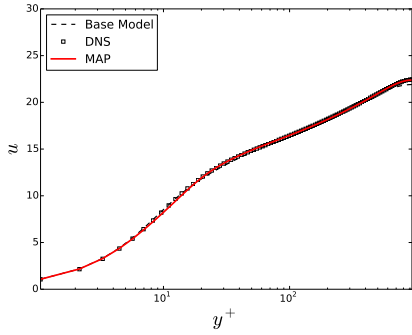
American Institute of Aeronautics and Astronautics
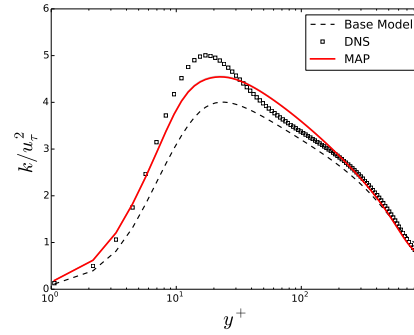
(a) Posterior distribution for velocity with $\pm 2\sigma$ limits shaded.

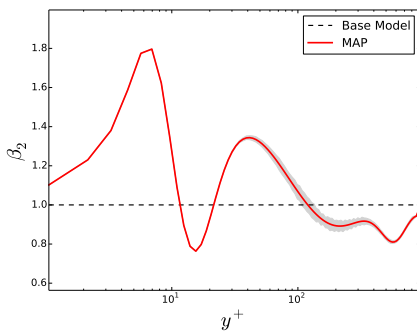
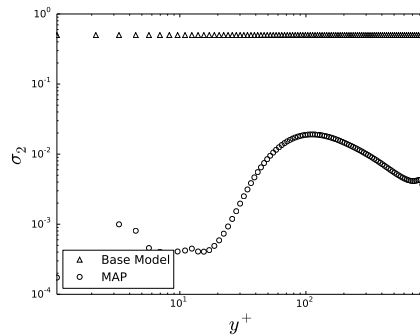(b) Posterior distribution for TKE with $\pm 2\sigma$ limits shaded.

(c) Posterior distribution for β_1.

(d) Posterior variance for β_1.

(e) Posterior distribution for β_2.

(f) Posterior variance for β_2.

Figure 7: Posterior model distributions for planar channel flow at $Re_\tau = 950$ with $C_{\varepsilon 2}^*$ and L modifiers.
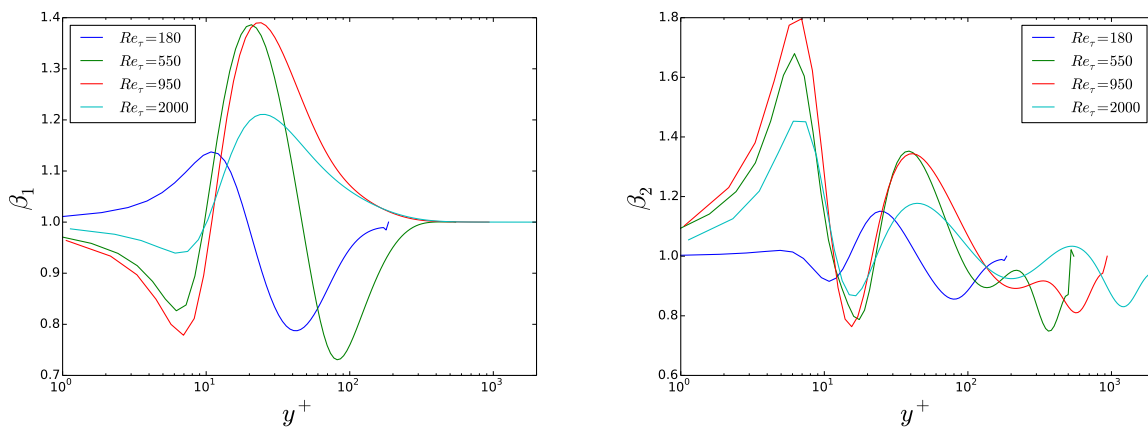
American Institute of Aeronautics and Astronautics

Figure 8: Summary for β's obtained in the inversion. β_1 modifies the length scale and β_2 modifies $C^*_{\varepsilon 2}$.

American Institute of Aeronautics and Astronautics

IV.   Application to Complex Cases

As a more complex test case, we now consider two dimensional flow over concave and convex surfaces. The Spalart Allmaras[12] model will be used as the baseline model, again with a multiplier $\beta(x,y)$ of the production term as the function to be inferred. Specifically, the effect of curvature is studied in (a) a convex channel and (b) a concave channel with a turning angle of 90 degrees (Figure 9). Curvature has a profound effect on the turbulence, acting to destabilise the boundary layer in the concave channel while stabilizing the convex channel boundary layer. While progress has been made in understanding the physical structure of such flows, accurate modeling still remains a challenge. We will use the Spalart-Allmaras model with a correction term $\beta(x,y)$ introduced in the following form

$$`\frac{D\tilde{v}}{Dt} = \beta(x,y)P(\tilde{v}) - D(\tilde{v}) + T(\tilde{v}),$$ (35)

where P, D and T are the production, dissipation and transport terms, respectively. Note that the Spalart Allmaras model with curvature correction (SARC[21]) also introduces corrections in the same form and will thus be used as a basis for comparison.

For both geometries, LES solutions from Reddy & Durbin[22] are used as benchmarks. The RANS grid is generated by taking every fourth grid point in the streamwise direction from the LES grid. The grid inlet is further extended to remove the effect of small differences in the inlet profile. The extended RANS mesh consists of 472 cells in the streamwise and 110 cells in the wall-normal direction. At the inlet, the LES solution results in a boundary layer profile
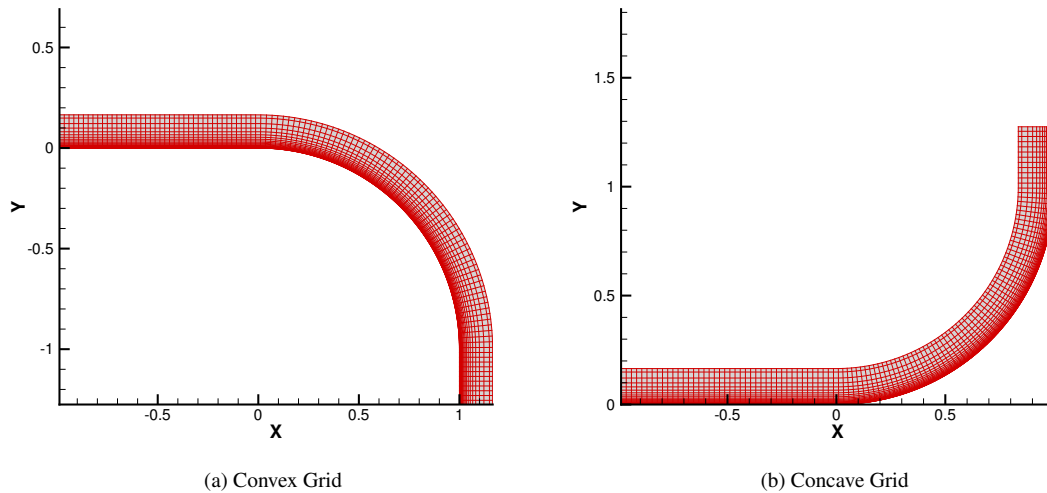

(a) Convex Grid

(b) Concave Grid

Figure 9: RANS domain and grid for convex / concave wall problems.

with $\mathfrak{R}_\theta \approx 2000$, and this is matched by the RANS via a precursor boundary layer simulation. It has to be noted that the LES is incompressible, whereas the RANS was performed at an inlet Mach number of 0.1. The physical variables from the boundary layer simulation are used to set ghost cell values for a characteristic-based inlet boundary condition. At the outlet boundary, pressure is set to match the LES. The lower wall is treated as a viscous wall, whereas the upper wall is inviscid.

IV..1.   Inverse Problem & Results

For both the geometries, the skin friction coefficient on the lower wall C_f is used as the objective function to be matched. The standard deviation for the prior is set to a value of 0.45 based on variations on β in our initial inverse tests. The standard deviation of the observable is set to 2.5E-6, a value based on the variability of $C_{f,LES}$.

Figure 10 compares the skin friction coefficient for the baseline, SARC, and inverse models with LES. The baseline SA model is seen to be highly inaccurate and the SARC, while much improved, still has significant errors. The MAP solution is seen to almost coincide with the benchmark for both cases. To compare the changes to the baseline model, the additional source terms $\delta = P(\beta - 1)$ introduced by the SARC and inversion are compared in Figures 11 and 12. It is interesting to note that the effective corrections are consistent with the expected physica and qualitatively similar

American Institute of Aeronautics and Astronautics

for the SARC and MAP solutions. Figure 13 shows the uncertainty bounds generated via the procedure described in the one dimensional channel flow case.
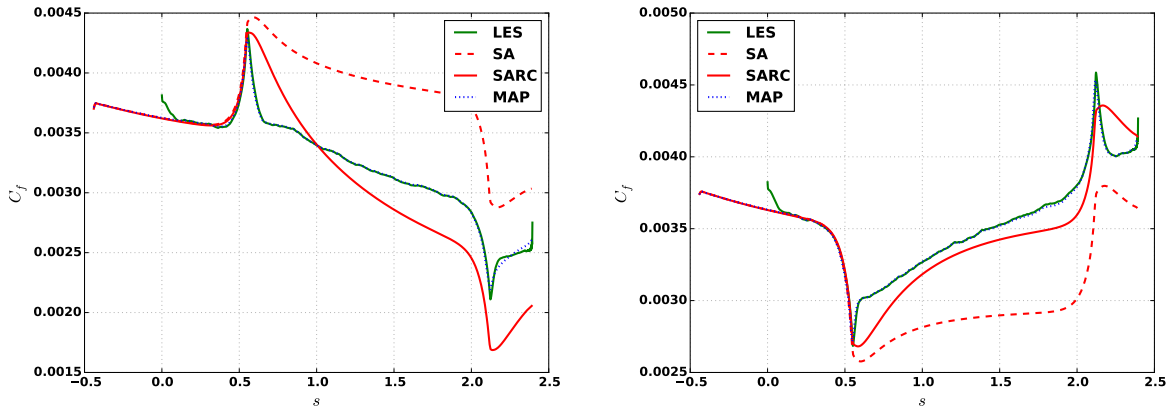

Figure 10: Comparison of skin friction coefficient with LES solution.
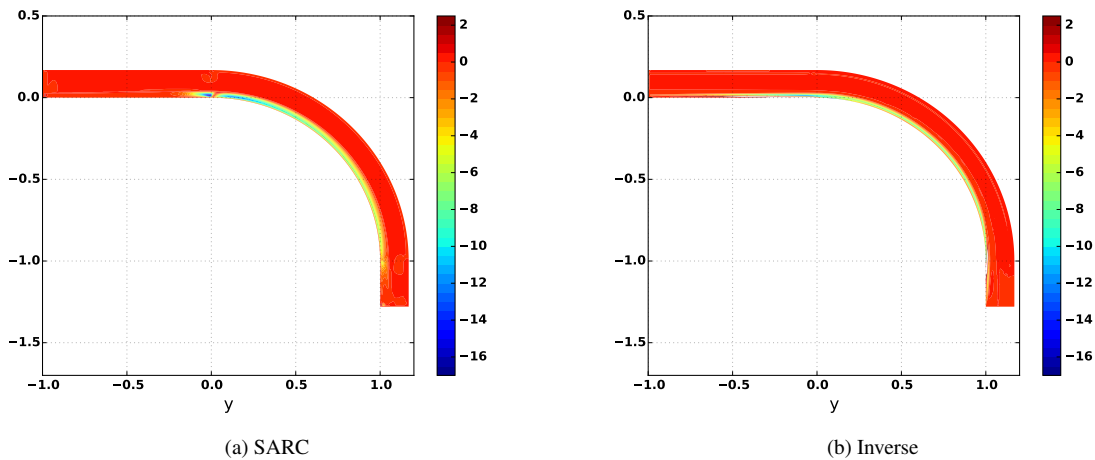

(a) SARC

(b) Inverse

Figure 11: Comparison of additive correction to the baseline SA model for Convex curvature.

V.   Model Prediction

For completeness we now demonstrate how the inversion step can be combined with machine learning algorithms to create a predictive model. More details of the overall methodology can be found in Duraisamy *et al.*[11] By taking the same type of Bayesian approach and making the appropriate assumptions, the process of inferring β_{post} can still be solved via the deterministic optimization problem

$$\beta_{map} = arg \min \frac{1}{2}\left[\left(\mathbf{d} - h(\beta)\right)^T \mathbf{C_m}^{-1}\left(\mathbf{d} - h(\beta)\right) + \left(\beta - \beta_{prior}\right)^T \mathbf{C_\beta}^{-1}\left(\beta - \beta_{prior}\right)\right]. \tag{36}$$

In this problem, the parameter to observable map *h* and observational covariance matrix $\mathbf{C_m}$ are functions of the data used to solve the inverse, the inverse itself, *and* the machine learning algorithm. To make the problem more tractable, the solution process is broken down into multiple steps. As before, an additive correction term is added to the low resolution model. This correction is completely described by $\mathfrak{B}$, which contains β and covariance C_β. Bayesian inversion and machine learning algorithms are applied to the prior correction, $\mathfrak{B}_{prior}$, to produce a machine learned
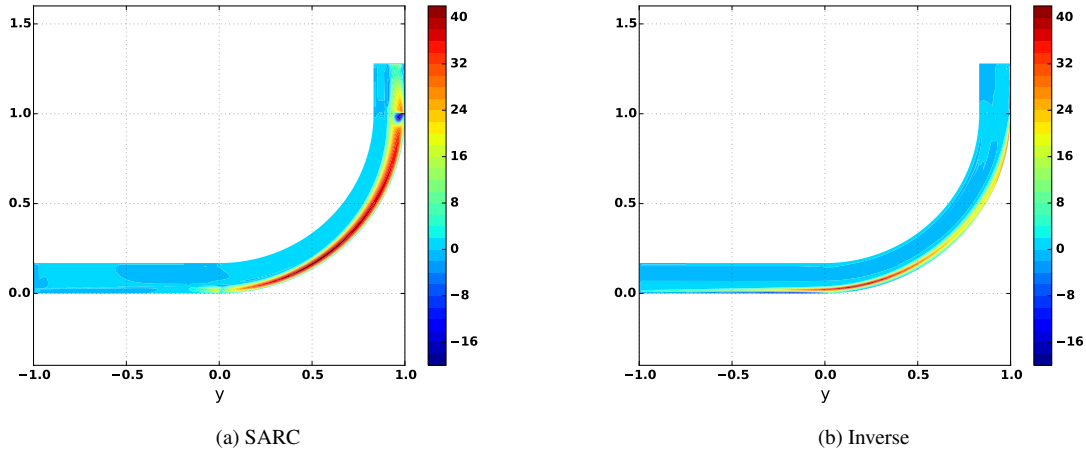
American Institute of Aeronautics and Astronautics

(a) SARC

(b) Inverse

Figure 12: Comparison of additive correction to the baseline SA model for Concave curvature.
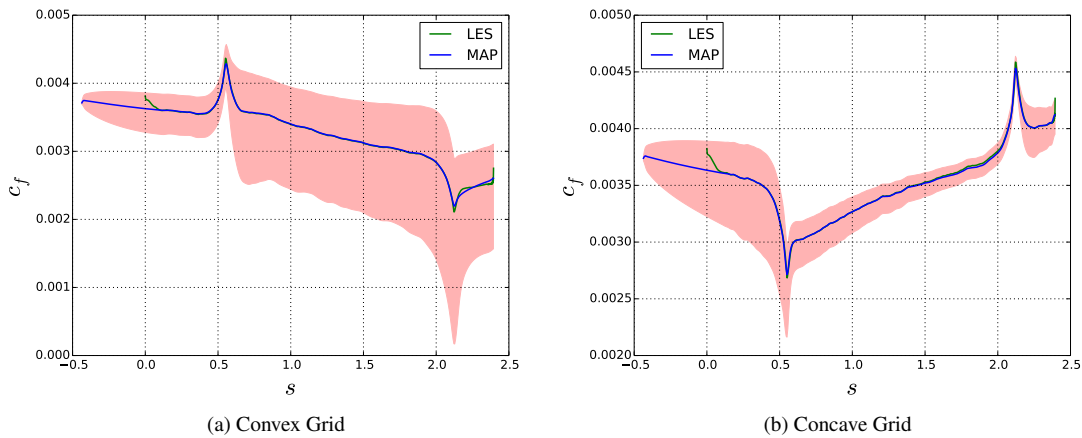

(a) Convex Grid

(b) Concave Grid

Figure 13: Uncertainty bounds for 2D curved channel cases.

(ML) correction

$$\mathfrak{B}_{prior}(\mathbf{U}) \xrightarrow{\textit{Inversion}} \mathfrak{B}_{inverse}(\mathbf{x}) \xrightarrow{\textit{Machine Learning}} \mathfrak{B}(\mathbf{U})_{ML}. \tag{37}$$

By propagating β and C_β through this process, the final posterior correction can then be determined by a second Bayesian inference step that combines the prior and machine learned corrections

$$\begin{matrix} \mathfrak{B}_{prior}(\mathbf{U}) \\ \mathfrak{B}_{ML}(\mathbf{U}) \end{matrix} \rightarrow \mathfrak{B}_{post}(\mathbf{U}). \tag{38}$$

The inversion and machine learning step given in Equation 37 are offline processes, while the final inference step given by Equation 38 is an online process. For the cases considered in this work the final inference step is given by

$$\beta_{post} = arg \min \frac{1}{2}\left[(\beta - \beta_{ML})^T \mathbf{C}_{\beta_{ML}}^{-1} (\beta - \beta_{ML}) + (\beta - \beta_{prior})^T \mathbf{C}_{\beta_{prior}}^{-1} (\beta - \beta_{prior}) \right]. \tag{39}$$

Note that the solution of Equation 39 is trivial compared to the previous optimization problem, since the parameter to observable map h has become the identity matrix. Properly accounting for uncertainty in every step of this process allows for formal quantification of the uncertainty in the posterior distribution of β. It is acknowledged that the implementation of this process is highly-nonlinear, and as a result, the true posterior distribution will not be Gaussian.

To provide a predictive example, planar channel flow is again considered. First, the inverse problem discussed in Section III is solved for the cases given in Table 1, except for $Re_\tau = 2000$. The data generated from the inversion
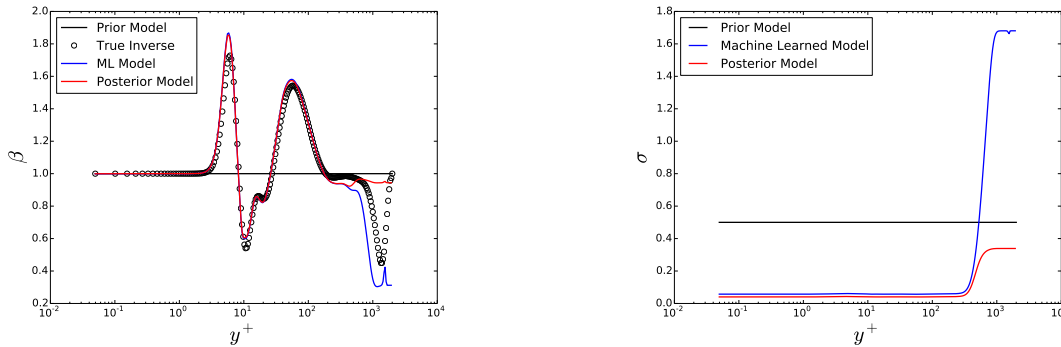
(a) Posterior model correction compared to the prior and machined learned models. The true inverse solution is additionally provided.

(b) Posterior variance predictions for the final posterior model.

Figure 14: Summary of the final Bayesian update step. For regions where the ML model predicts a low training variance, the posterior correction collapses to the ML model. For regions where a high training variance is predicted, the posterior correction is closer to the prior model.

are then passed on to a machine learning algorithm, which in turn obtains a general functional form for β. Details of the machine learning aspect of our work are discussed in a companion paper.[15] The machine learned algorithm is then injected back into the computational solver, where it is queried in a predictive fashion. The output of the machine learned algorithm is then filtered through the final Bayesian inference step, which combines the information gained through the entire inverse/machine learning process with the prior, to produce a final posterior prediction of the corrective source term.

The resulting posterior model is now tested for the $Re_\tau = 2000$. This case was omitted from the training process of the machine learned algorithm. In this example the machine learned model is only queried once with the RANS quantities obtained from the solution of the prior model to obtain $\beta_{ML}(y)$. The posterior correction term β_{post} is then determined by solving Equation 39. Figure 14 shows the results of this final inversion step. It is seen that, for regions where the ML model predicts a low training variance[b], β_{post} assimilates to the ML correction β_{ML}. For regions where a high training variance is predicted, the posterior correction β_{post} reverts to the prior correction β_{prior}. Figure 15 shows the resulting mean profiles of the final posterior model. The results of the implementation are good, especially in the inner layer. The posterior model is able to predict the velocity profile in the buffer layer quite well. The prediction of turbulent production and dissipation are additionally improved from the baseline model. Upon reaching the outer layer, however, the machine learning algorithm is not very accurate in determining the functional form of the training data. As such the posterior model reverts to the prior model, which is seen in Figure 14. The failure to predict the increased destruction of TKE in the channel core leads to a slight under-prediction of velocity. It is interesting to note, however, that this suggests that the prior model accurately predicts the center line velocity due to errors within the inner layer. The model is simply calibrated to negate these errors. Indeed, the turbulent eddy viscosity predicted by the $k - \omega$ model is too high in the channel core. Nonetheless, the error norms on the solution mesh[c] for velocity, turbulent kinetic energy, and Reynolds stress are decreased by $2\%, 2\%$, and 50%, respectively, from the baseline model. Despite limited training data, the inverse/machine learning framework is able to predict an unknown correction term to a good degree of accuracy. Furthermore, it is clear that properly accounting for uncertainties in every step of the process can allow the machine learned and prior models to be blended together in a rigorous way, allowing for the uncertainty of the posterior model to be directly quantified.

VI.   Conclusions and Future Work

The availability of high-fidelity numerical and experimental data presents new opportunities for the use of data-driven methods to improve closure models for turbulence. Here a subset of a predictive modeling framework was presented. Bayesian inference was used to infer full field model corrections for turbulent channel flow and a convex/concave wall bounded flow. In both cases, the accuracy of the (problem dependent) post-inversion model was

---

[b]In this example the machine learning algorithm was not informed of the variance predicted through the inverse.

[c]As previously mentioned, the solution mesh is concentrated in the inner layer.

American Institute of Aeronautics and Astronautics

(a) Velocity profiles predicted by the posterior model.

(b) Turbulent Reynolds stress predicted by the posterior model.

(c) Kinetic energy predicted by the posterior model.

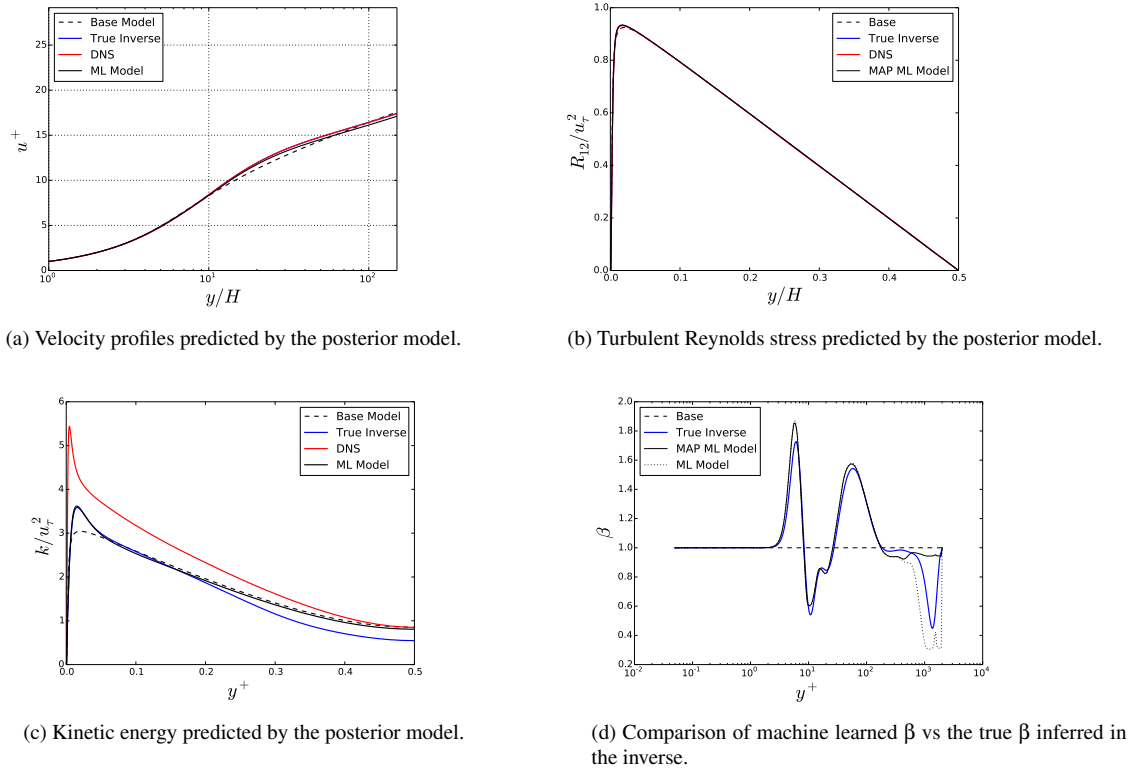(d) Comparison of machine learned β vs the true β inferred in the inverse.

Figure 15: Machine learned $k - \omega$ model compared to the exact inverse model and DNS for planar channel flow at $Re_\tau = 2000$.

greatly increased and modeling uncertainties were quantified. By examining the posterior distribution of the model correction, it is possible to determine regions of the flow where the inversion is effective, as well as regions where no useful model corrections could be inferred. This information is useful for both the modeler as well as in the context of our overall goal of improving turbulence models by coupling inversion results with machine learning algorithms.

Current work is focused on a more rigorous specification of the prior as well as the observational covariance matrices. In the case of an uninformed prior, the posterior distributions are highly dependent on observation information; as such the effects of building in additional physics (such as spatial correlations) are being considered.

VI.A.   Acknowledgments

References

[1]Mavriplis, D., Lurie, E., Gropp, W., Darmofal, D., Alonso, J., Khodadoust, A., and Slotnick, J., "CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences," Tech. Rep. CR-2014-218178, NASA, 2013.

[2]Iaccarino, G., Duraisamy, K., Campos, A., Emory, M., and Ryu, S., "Large eddy simulation of a wing-body junction flow," *Bulletin of the American Phyysical Society*, 2014.

[3]Poroseva, S. and Murman, S., "Velocity/Pressure-Gradient Correlations in a FORANS Approach to Turbulence Modling," *Proc. the AIAA Aviation and Aeronautics Forum and Exposition*, June 2014.

[4]Gerolymos, G., Lo, C., and Vallet, I., "Tensorial Representation of Reynolds-stress pressure-strain redistribution," *Journal of Applied Mechanics*, 2010.

[5]Olsen, M. E., "Prediction of Large Separations with a Third-Order-Moment Model," *53rd AIAA Aerospace Sciences Meeting*, 2015.

[6]Brynjarsdóttir, J. and OHagan, A., "Learning about physical parameters: The importance of model discrepancy," *Institute of Physics Publishing*, 2014.

[7]Yarlanki, S., Rajendran, B., and Hamann, H., "Estimation of turbulence closure coefficients for data centers using machine learning algorithms," *Thermal and Thermomechanical Phenomena in Electronic Systems*, San Diego, CA, 2012, pp. 38 – 42.

[8]Kato, H. and Obayashi, S., "Data Assimilation for Turbulent Flows," *16TH AIAA Non-Determinisitc Approaches Conference*, The American Institute of Aeronautics and Astronautics, 2014.

[9]Edeling, W. N., Cinnella, P., and Dwight, R. P., "Bayesian estimates of parameter variability in the k - e turbulence model," 2014.

[10]Dow, E. and Wang, Q., "Quantification of Structural Uncertainties in the k- ω Turbulence Model," *AIAA Paper*, Vol. 1762, 2011.

[11]Duraisamy, K., Zhang, Z., and Singh, A., "New Approaches in Turbulence and Transition Modeling Using Data-driven Techniques," *AIAA Scitech*, Orlando, Florida, January 2015.

[12]Spalart, P. and Allmaras, S., "A One-Equation Turbulence Model for Aerodynamic Flows," *30th Aerospace Sciences Meeting and Exhibit*, Reno, NV, Jan 1992.

[13]Wilcox, D. C., *Turbulence Modeling for CFD*, Vol. Third Edition, DCW Industries, 2006.

[14]Billard, F. and Laurence, D., "A robust k-e-v2/k elliptic blending turbulence model applied to near-wall, separated and buoyant flows," *International Journal of Heat and Fluid Flow*, 2011.

[15]Zhang, Z. J., Duraisamy, K., and Gumerov, N. A., "Machine Learning Methods for Data-Driven Turbulence Modeling," *AIAA Aviation and Aeronautics Forum and Exposition*, Dallas, Texas, June 2015.

[16]Aster, R., *Parameter Estimation and Inverse Problems*, Elsevier Academic Press, 2005.

[17]Dennis, Jr, J. E. and Moré, J. J., "Quasi-Newton methods, motivation and theory," *SIAM review*, Vol. 19, No. 1, 1977, pp. 46–89.

[18]Hascoët, L. and Pascual, V., "The Tapenade Automatic Differentiation tool: Principles, Model, and Specification," *CM Transactions on Mathematical Software*, 2013.

[19]Caplan, P., "Numerical Computation of Second Derivatives with Applications to Optimization Problems," Unpublished Academic Report, MIT.

[20]Jimenez, J. and Hoyas, S., "Turbulent flucuations above the buffer layer of wall-bounded flows," *Journal of Fluid Mechanics*, Vol. 611, 2008, pp. 215–236.

[21]Spalart, P. R. and Shur, M. L., "On the Sensitization of Turbulence Models to Rotation and Curvature," *Aerospace Science and Technology*, Vol. 1, 1997, pp. 297–302.

[22]Reddy, K. R. and Durbin, P., *Personal communication*, 2015.

American Institute of Aeronautics and Astronautics