14th International Symposium on Turbulence and Shear Flow Phenomena (TSFP14)
Heidelberg, Germany, July 28–31, 2026

A FORMALLY LOG(RE) SOLUTION TO THE ENGINEERING TURBULENCE PROBLEM

Jiaqi Li
Mechanical Engineering
The Pennsylvania State University
University Park, PA 16802, USA
jfl5663@psu.edu

Robert F. Kunz
Mechanical Engineering
The Pennsylvania State University
University Park, PA 16802, USA
rfk102@psu.edu

George Huang
Mechanical and Materials Engineering
Wright State University
Dayton, OH 45435, USA
george.huang@wright.edu

Xiang I.A. Yang
Mechanical Engineering
The Pennsylvania State University
University Park, PA 16802, USA
xzy48@psu.edu

ABSTRACT

In fluid engineering, the turbulence problem is the long-standing challenge of obtaining accurate predictions of engineering quantities at affordable computational cost. Viewed through the lens of computational complexity, a practical algorithm requires cost growth no worse than $O(N)$, where N denotes the problem size. For turbulent flows, the problem size may be approximated by the number of dynamically relevant scales and, therefore, by the Reynolds number Re. On this basis, the cost of a practical method should scale no worse than $O(Re)$. We propose a multi-fidelity, physics-constrained, data-driven framework that formally satisfies this criterion. For demonstration, we augment the Spalart–Allmaras (SA) model through field inversion and machine learning (FIML) using a constrained formulation that preserves the law of the wall. The model is trained at low Reynolds number, where high-fidelity data are affordable, and is then deployed at higher Reynolds numbers. Relative to the target Reynolds number, the high-fidelity calculation contributes a cost of $O(Re^0)$, the FIML stage also contributes $O(Re^0)$, and the low-fidelity RANS prediction scales as $O(\log(Re))$. The overall formal cost scaling is therefore $O(\log(Re))$. The framework is assessed in plane channel flow and in the periodic hill with slope factor $\alpha = 1$. In channel flow, a model trained at $Re_\tau = 1000$ corrects the wake-layer error of the baseline SA model and retains this improvement when extrapolated to $Re_\tau = 5200$. In the periodic hill, a model trained at $Re_b = 5600$ is tested at $Re_b = 10595$, 19000, and 37000. The constrained formulation preserves the correct separation and recovery behavior as Reynolds number increases. For comparison, we include results from conventional unconstrained FIML and from the baseline model. Across all tested Reynolds numbers, the constrained model yields the lowest root-mean-square error and exhibits nearly Reynolds-number-independent error, indicating robust extrapolation.

INTRODUCTION

Turbulence remains one of the central unresolved problems in physics and engineering. In fluid engineering, the practical difficulty lies in predicting quantities such as drag and lift at high Reynolds numbers with both useful accuracy and acceptable computational cost. The core difficulty arises from the vast separation of dynamically relevant scales that must be represented.

Computational fluid dynamics (CFD) provides a hierarchy of modeling approaches, each reflecting a different compromise between accuracy and cost depending on which scales are resolved and which are modeled. Direct numerical simulation (DNS), which resolves all scales down to the Kolmogorov length, is regarded as the gold standard. However, even with optimized grids and algorithms, its cost scales as $Re_L^{2.91}$ (Yang & Griffin, 2021), rendering it impractical beyond laboratory Reynolds numbers. Large-eddy simulation (LES) reduces this cost by modeling the smallest scales while resolving the inertial range, but still scales as $Re_L^{2.72}$. Wall-modeled LES (WMLES) and hybrid RANS/LES methods further reduce the cost to approximately $Re_L^{1.14}$. Despite their improved efficiency, these turbulence-resolving approaches remain too expensive for routine industrial applications at, for example, $Re \sim 10^9$ (Slotnick *et al.*, 2014). At the other end of the spectrum, Reynolds-averaged Navier–Stokes (RANS) models resolve only the mean flow while modeling all turbulent fluctuations and are therefore far more affordable than DNS and LES. Their weakness is accuracy: because there is no universal mapping between mean flow and turbulence, widely used closures such as Spalart–Allmaras (SA) and k–ω SST show limited performance in flows with strong separation and other non-equilibrium effects (Spalart, 2015; Durbin, 2018).

To place this accuracy–cost trade-off in a broader context, it is useful to borrow the language of computational complexity. In computer science, algorithmic cost is assessed not only by the expense of solving a single instance, but by how that expense grows with problem size N. Algorithms whose cost grows linearly or sublinearly remain practical as N increases, whereas superlinear growth quickly becomes prohibitive in repeated use. A similar viewpoint may be adopted for turbulent flows. Because the Reynolds number measures the separation between the largest and smallest dynamically relevant scales, it may be interpreted as an effective measure of problem size. Cost scalings expressed in terms of Re can therefore be viewed as statements of algorithmic complexity. From this perspective, the limitation of turbulence-resolving methods becomes immediate: their superlinear cost growth renders them impractical at engineering Reynolds numbers. RANS, by contrast, lies on fundamentally different computational footing because

its cost depends only weakly on *Re*. A practical route forward is therefore preserving the affordability of RANS while recovering some of the accuracy of high-fidelity methods.

The emergence of machine learning (ML) has opened new opportunities for turbulence modeling (Duraisamy *et al.*, 2019). Early developments primarily aimed to improve predictions for specific flow configurations (Wang *et al.*, 2017; Wu *et al.*, 2018; Parish & Duraisamy, 2016), whereas more recent efforts have focused on generalization across geometries and flow conditions (Rumsey *et al.*, 2022; Srivastava, 2022; Srivastava *et al.*, 2024). Among the available approaches, field inversion and machine learning (FIML) framework is especially attractive because it augments an existing closure in a model-consistent manner (Singh *et al.*, 2017; Holland *et al.*, 2019; Rumsey *et al.*, 2022). In the FIML framework, a baseline turbulence model is first recalibrated through field inversion so that it reproduces reference data, and a regression model is then trained to learn the relationship between local flow features and the inferred correction. Another closely related paradigm is progressive machine learning, which seeks to improve robustness and extrapolative capability without erasing successful calibrations already embedded in the baseline model (Bin *et al.*, 2022, 2024*a*,*b*; Yang *et al.*, 2025).

Despite these advances, most prior work in data-driven turbulence modeling has emphasized generalization in configuration space rather than extrapolation in Reynolds number. For engineering utility, however, Reynolds-number extrapolation is at least as important. Existing attempts at *Re* extrapolation (Vadrot *et al.*, 2023) have shown promise in canonical flows such as channels and boundary layers, but their relevance to complex separated flows remains limited. This gap motivates the present study.

In this paper, we introduce a multi-fidelity, physics-constrained, data-driven turbulence-modeling framework developed specifically for extrapolation across Reynolds numbers. The framework combines three ingredients: the efficiency of RANS, whose cost depends only weakly on Reynolds number; the accuracy of high-fidelity data obtained at moderate *Re*; and the extrapolative capability of physics-constrained machine learning that preserves calibration with respect to the law of the wall. For demonstration, we augment the SA model using both constrained and unconstrained FIML formulations and assess the resulting models in plane channel flow and in the periodic hill (Breuer *et al.*, 2009; Rapp & Manhart, 2011). The goal is not universal cross-geometry generalization, but reliable Reynolds-number extrapolation within a flow family while retaining the favorable cost structure of RANS.

RANS COST SCALING

This section establishes the cost scaling for RANS of boundary-layer flows. To the authors' knowledge, an explicit derivation form is not commonly stated. Following Yang & Griffin (2021) and Choi & Moin (2012), we consider a flow with a dominant streamwise direction x and a wall-normal direction y, and assume that the grid is aligned with the mean-flow direction. This setting is representative of common RANS practice in wall-bounded flows.

Because RANS resolves only the mean flow, the grid spacing should be determined by the gradients of the mean velocity rather than by the smallest turbulent scales. A convenient estimate is

$$\Delta_x = \frac{c_x U_\infty}{|\partial U/\partial x|} \approx \frac{c_x U_\infty}{|du_\tau/dx|}, \qquad \Delta_y = \frac{c_y u_\tau}{\partial U/\partial y}, \tag{1}$$

where U_∞ is a characteristic outer velocity, u_τ is the friction velocity, and c_x and c_y are constants. The spanwise resolution is taken to be uniform, with a fixed number of points N_z, as is commonly done for, e.g., bodies of revolution.

Under these assumptions, the total number of grid points is

$$N = N_z \iint \frac{dx\,dy}{\Delta_x \Delta_y}. \tag{2}$$

Substituting Eq. (1) into Eq. (2) gives

$$N = \frac{N_z}{c_x c_y U_\infty} \iint \frac{|du_\tau/dx|}{u_\tau} \frac{\partial U}{\partial y}\, dx\,dy. \tag{3}$$

Integrating first in the wall-normal direction yields

$$\int_0^\delta \frac{\partial U}{\partial y}\, dy = U(\delta) - U(0) \sim U_\infty, \tag{4}$$

so that Eq. (3) reduces to

$$N = \frac{N_z}{c_x c_y} \int \frac{|du_\tau/dx|}{u_\tau}\, dx \sim \ln\left(\frac{u_{\tau,1}}{u_{\tau,2}}\right). \tag{5}$$

To express this result in terms of the Reynolds number, we use the standard relation

$$u_\tau = U_\infty \sqrt{\frac{C_f}{2}}, \tag{6}$$

together with the flat-plate turbulent boundary-layer estimate

$$C_f \sim Re_x^{-1/7}. \tag{7}$$

It then follows that

$$N \sim \ln(Re). \tag{8}$$

Equation (8) shows that the RANS grid-point requirement grows only logarithmically with Reynolds number. If the cost of the steady solver scales linearly with the number of grid points, then the overall computational cost of RANS also scales as

$$\mathrm{Cost}_{\mathrm{RANS}} \sim N \sim \log(Re). \tag{9}$$

SA MODEL AND LAW OF THE WALL CONSTRAINT

This section summarizes the baseline SA model and the law-of-the-wall (LoW)-constrained modification used in the present work, following Spalart & Allmaras (1992) and Bin *et al.* (2024*b*). The transport equation for the modified eddy viscosity $\tilde{\nu}$ is

$$\begin{aligned}\frac{\mathrm{D}\tilde{\nu}}{\mathrm{D}t} &= c_{b1}\tilde{S}\tilde{\nu} - c_{w1} f_w \left(\frac{\tilde{\nu}}{d}\right)^2 \\ &\quad + \frac{1}{\sigma}\left[\frac{\partial}{\partial x_j}\left((\nu+\tilde{\nu})\frac{\partial \tilde{\nu}}{\partial x_j}\right) + c_{b2}\frac{\partial \tilde{\nu}}{\partial x_j}\frac{\partial \tilde{\nu}}{\partial x_j}\right],\end{aligned} \tag{10}$$

where $\mathrm{D}/\mathrm{D}t$ denotes the material derivative. The eddy viscosity is given by $\nu_t = \tilde{\nu} f_{v1}$, and the anisotropic part of the Reynolds stress tensor is modeled with the Boussinesq hypothesis as $R^d_{ij} = 2\nu_t S_{ij}$; the modified strain-rate magnitude is

$$\tilde{S} = \Omega + \frac{\tilde{\nu}}{\kappa^2 d^2} f_{v2}, \tag{11}$$

where Ω is the vorticity magnitude, d is the distance to the nearest wall, and $\kappa = 0.41$ is the von Kármán constant; the auxiliary functions are

$$f_{v1} = \left[1 - \exp\left(-\frac{\chi}{\kappa A}\right)\right]^2, \quad f_{v2} = 1 - \frac{\chi}{1 + \chi f_{v1}}, \tag{12}$$

with $\chi = \tilde{\nu}/\nu$; the model coefficients are $A = 17$, $c_{b2} = 0.622$, and $\sigma = 2/3$. The present form of f_{v1} differs from that in the original SA model (Spalart & Allmaras, 1992). It enforces the correct near-wall scaling, $f_{v1} \sim y^2$, while remaining consistent with the original SA formulation away from the wall. Lastly, the function f_w is defined as

$$f_w = g\left(\frac{1 + 2^6}{g^6 + 2^6}\right)^{1/6}, \qquad g = r + 0.3(r^6 - r), \tag{13}$$

where $r = \min\left(\tilde{\nu}/(\tilde{S}\kappa^2 d^2), 10\right)$.

In the present study, the coefficient c_{b1} is allowed to vary in space, following Singh *et al.* (2017), Rumsey *et al.* (2022), and Srivastava *et al.* (2024). Two formulations are considered: a conventional unconstrained FIML approach and a constrained variant. In the unconstrained formulation, coefficients such as c_{w1} remain fixed at their baseline values as c_{b1} varies. In the constrained formulation, the variations of c_{b1} and c_{w1} are coupled so that the model remains consistent with the LoW. This preserves the correct near-wall asymptotic behavior during both field inversion and machine learning. Additional details are given in Bin *et al.* (2024*b*) and are not repeated here.

MULTI-FIDELITY FRAMEWORK

The proposed framework consists of three stages. First, a high-fidelity dataset is generated at a low Reynolds number, where such calculations remain affordable. Second, FIML is used to augment the baseline turbulence model, which in the present study is the SA model. The model correction is introduced through a spatially varying c_{b1}. In the field-inversion step, an optimization problem is solved to infer the correction field that enables the augmented RANS solution to match the high-fidelity target data. In the subsequent machine-learning step, this inferred field is replaced by an explicit functional relation between local nondimensional flow features and the correction. Once trained, this relation can be evaluated directly at a new Reynolds number without further optimization, access to high-fidelity data, or case-by-case retuning. In the present implementation, field inversion is carried out with DAFoam (He *et al.*, 2020), the regression model is a fully connected feedforward neural network, and the learned correction is deployed in OpenFOAM (Weller *et al.*, 1998). Third, the augmented closure is evaluated at higher Reynolds numbers using RANS.

Because the high-fidelity simulation and the learning stage are performed only once at low Reynolds number, their cost does not scale with the target Reynolds number. The online prediction cost therefore inherits the favorable Reynolds-number scaling of RANS. In the language of algorithmic complexity, the overall cost is

$$\text{Cost} = O(Re^0) + O(Re^0) + O(\log Re) = O(\log Re). \tag{14}$$

Two formulations are considered: an unconstrained formulation and a constrained one. In the unconstrained formulation, all coefficients other than c_{b1} retain their baseline values. In the constrained formulation, the coupling between c_{b1} and the associated destruction term is modified so that the law of the wall is preserved. This distinction is important for Reynolds-number extrapolation. As Reynolds number increases, the main structural change in many wall-bounded flows is the compression of the near-wall layer. A useful data-driven correction should therefore improve the outer-flow behavior without destroying the near-wall asymptotic calibration of the baseline model. The underlying assumption is that dominant nonequilibrium outer-layer physics, such as pressure-gradient and separation effects, can be learned at moderate Reynolds number and then transferred to higher Reynolds numbers, provided that the near-wall asymptotics are respected. This assumption is reasonable for many aerodynamic and hydrodynamic flows, in which increasing Reynolds number primarily thins the inner layer rather than introducing entirely new outer-layer physics.

Different feature sets are used for the two flow configurations considered here. For plane channel flow, the single feature $|P|/(|P| + |D|)$ is used, where P and D denote the production and destruction terms in the SA model. This feature isolates the wake-layer imbalance responsible for the baseline-model error. For the periodic hill, two features are used: $1/(f_w + 1)$, which reflects the pressure-gradient effect, and $|P|/(|P| + |D|)$, which captures the balance between mean-shear production and near-wall destruction. The channel model is trained at $Re_\tau = 1000$ using the DNS data of Graham *et al.* (2016), and the periodic-hill model is trained at $Re_b = 5600$.

RESULTS

Plane channel flow

Plane channel flow provides a simple test of whether the framework can learn a correction at one Reynolds number and retain it at another. The baseline SA model reproduces the inner-layer behavior well, but it overpredicts the velocity deficit in the wake. Figure 1 shows that the constrained ML model reduces this wake-layer error at the training condition, $Re_\tau = 1000$, and retains the improvement when extrapolated to $Re_\tau = 5200$. This result is consistent with the expectation that the wake-layer correction should remain transferable when the near-wall scaling is preserved. The channel case therefore serves as a calibration check: if the framework cannot extrapolate in this comparatively simple flow, it is unlikely to succeed in a separated configuration.

Periodic hill

The periodic hill provides a more demanding test because it involves separation, shear-layer development, and recovery, all of which vary with Reynolds number (Breuer *et al.*, 2009; Rapp & Manhart, 2011). Here we compare the baseline SA model, the unconstrained ML model, and the constrained ML model.

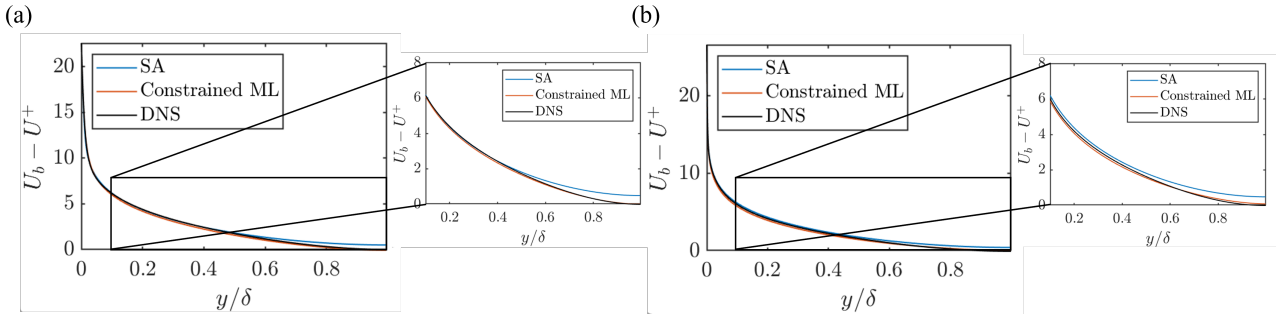

Figure 1. Velocity-defect profiles in plane channel flow at (a) $Re_\tau = 1000$ and (b) $Re_\tau = 5200$.

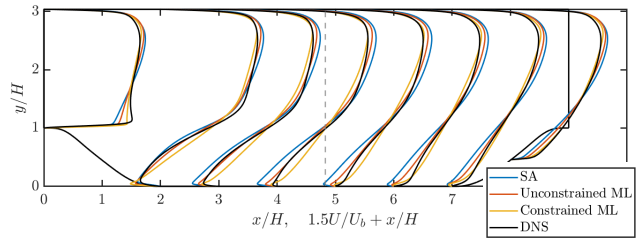

Figure 2. Streamwise velocity profiles in the periodic hill at the training condition $Re_b = 5600$. The dashed line marks the reattachment location.

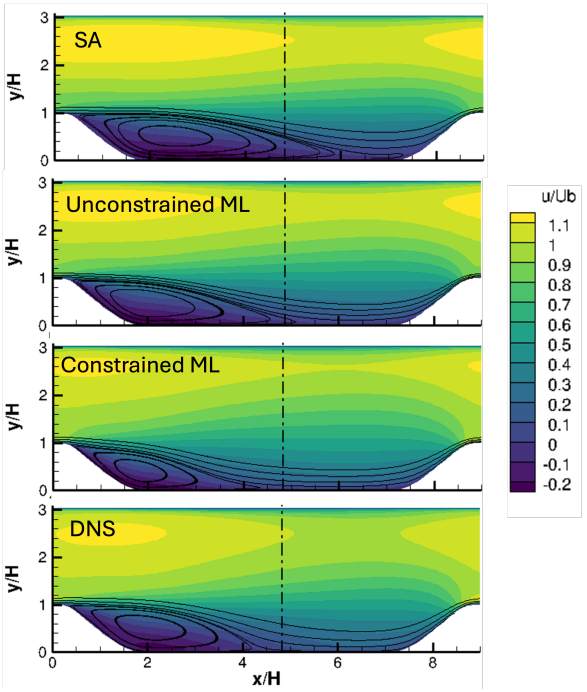

Figure 3. Contours of streamwise velocity and streamlines at the training condition $Re_b = 5600$. The dashed lines mark the reattachment locations.

Figures 2 and 3 show the results at the training condition. Both machine-learned models improve the recirculation and recovery regions relative to the baseline SA model. The velocity profiles show better agreement with the reference data in and downstream of the separated region, and the contour plots indicate a more realistic recirculation bubble and recovery pattern. The difference between the constrained and unconstrained formulations remains modest at the training Reynolds number. This is expected because both models have sufficient flexibility to fit the low-Reynolds-number target case. At this stage, the training case mainly shows that both learned corrections can extract useful information from the reference data.

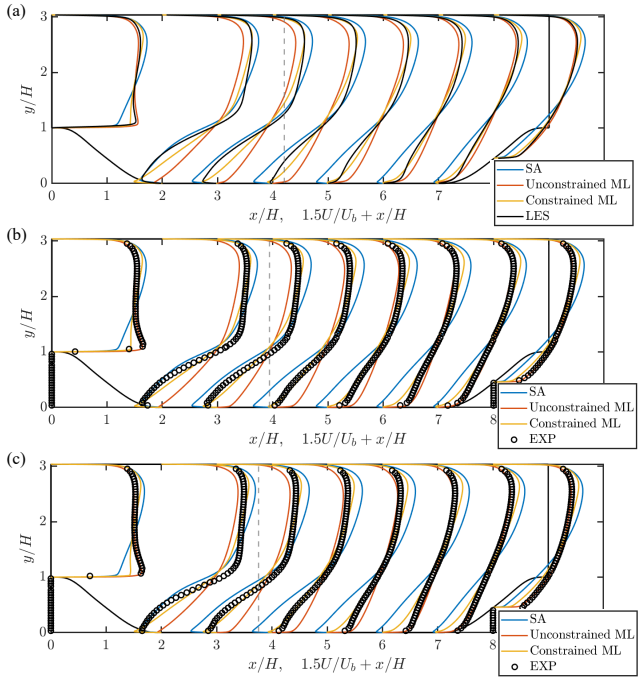

Figure 4. Velocity profiles at test Reynolds numbers $Re_b =$ 10595, 19000, and 37000. Dashed lines mark the reattachment locations.

The difference between the two ML formulations becomes much more pronounced under extrapolation. Figures 4 and 5 show the results at $Re_b = 10595$, 19000, and 37000. The constrained ML model preserves the size of the separation bubble and predicts the downstream recovery much more accurately than either the baseline SA model or the unconstrained ML model. This behavior is visible in both the profile comparisons and the contour plots. The predicted reattachment location remains much closer to the reference value, and the overall mean-flow development downstream of reattachment is captured more faithfully.

In contrast, the unconstrained ML model predicts an unrealistically short separation region as Reynolds number increases. Its performance at the training condition therefore does not translate into robust Reynolds-number extrapolation. This trend indicates that fitting the training case is not sufficient by itself; the augmentation must also preserve the near-

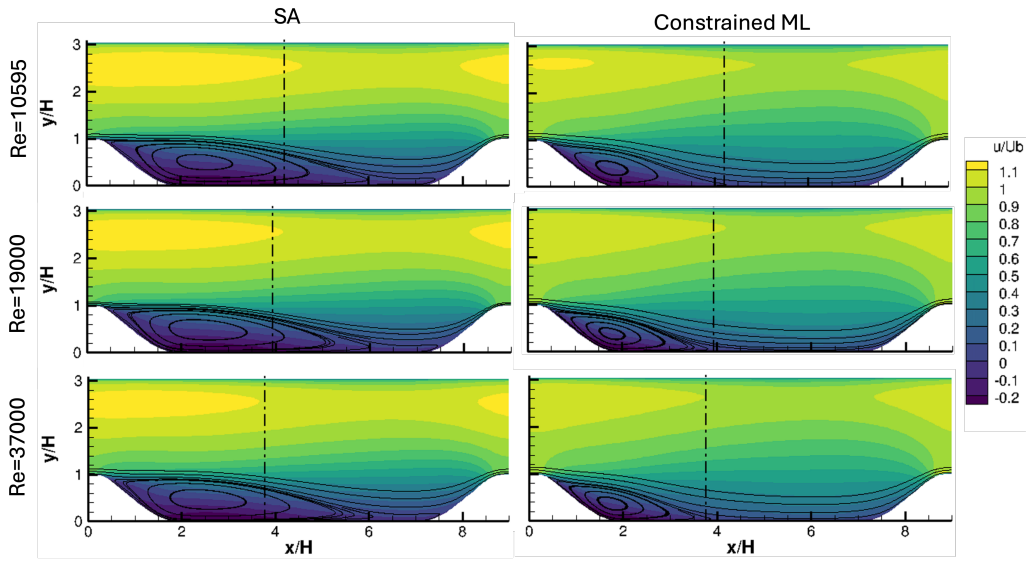

Figure 5. Velocity contours and streamlines at test Reynolds numbers: (a) $Re_b = 10595$, (b) $Re_b = 19000$, and (c) $Re_b = 37000$. Dashed lines mark the reattachment locations.

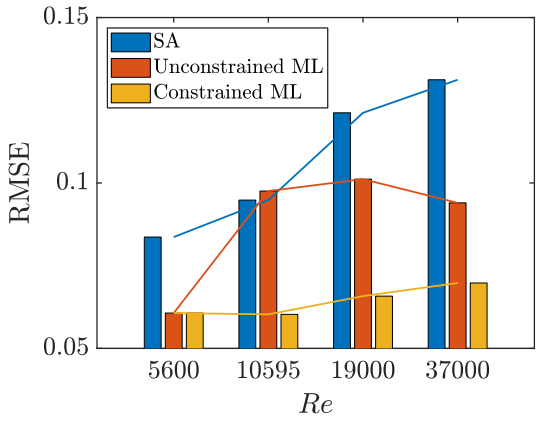

Figure 6. RMSE of the periodic-hill velocity profiles relative to experiment.

wall structure required for extrapolation. The baseline SA model also deteriorates systematically with Reynolds number: the experimentally observed reattachment location shifts with Re, whereas the modeled reattachment location remains comparatively insensitive to it.

To quantify these trends, we compute the root-mean-square error (RMSE) of the streamwise velocity profiles relative to experiment:

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(U_i - U_{i,\mathrm{ref}}\right)^2}, \tag{15}$$

where the average is taken over the locations for which experimental data are available. The results are shown in Figure 6. The baseline SA model exhibits a monotonic increase in error with Reynolds number. The unconstrained ML model reduces the error relative to the baseline, but its error also grows with Reynolds number. By contrast, the constrained ML model yields the lowest RMSE at every test case and, importantly, maintains nearly constant error from $Re_b = 5600$ to $Re_b = 37000$. This behavior strongly suggests that the law-of-the-wall constraint is essential for stable Reynolds-number extrapolation in separated flows.

DISCUSSION AND CONCLUDING REMARKS

The present results suggest that the engineering turbulence problem should be viewed as a balance between accuracy and computational complexity. The goal is not to eliminate modeling or to require scale-resolving simulation for every high-Reynolds-number prediction. Rather, the goal is to combine a computationally practical model with a calibration strategy that transfers the most useful high-fidelity information to the regime of interest. In the present framework, the expensive steps—high-fidelity simulation and field inversion—are performed at low or moderate Reynolds number, while predictions at the target high Reynolds number retain the cost structure of RANS.

This perspective also clarifies why Reynolds-number extrapolation differs from cross-geometry generalization. Cross-geometry generalization requires the model to accommodate different dominant mechanisms and flow organizations. Reynolds-number extrapolation, by contrast, often preserves the large-scale structure of the flow while changing the thickness of the inner layer and the degree of scale separation. For this reason, preserving the law of the wall is especially important. The periodic-hill results support this view: both learned models improve the training case, but only the constrained model remains robust as Reynolds number increases.

The framework also suggests a practical workflow for engineering applications. One may train on a canonical or reduced configuration for which high-fidelity data are feasible, use field inversion to expose the dominant model discrepancy, and then learn a constrained correction for deployment over a family of Reynolds numbers. This strategy is not universally applicable, because increasing Reynolds number may introduce new physics not represented in the training case. Nevertheless, for many attached and separated vehicle-relevant flows, this requirement is more realistic than demanding universal cross-geometry extrapolation from a single dataset.

The present results should therefore be viewed as an existence proof rather than a universal closure. For both a canonical attached flow and a separated flow, the proposed frame-

work retains RANS-level online cost, preserves the correct wall scaling, and substantially improves extrapolative behavior through data-enabled augmentation.

In summary, this work examined Reynolds-number extrapolation in data-driven turbulence modeling using a physics-constrained FIML framework. The learned correction reduces the wake-layer bias of the SA model in plane channel flow and remains effective when extrapolated. In the periodic hill, both constrained and unconstrained ML improve the training case, but only the constrained formulation remains reliable as Reynolds number increases. Across all periodic-hill cases, the constrained model gives the lowest RMSE and nearly Reynolds-number-independent error. These results show that, for Reynolds-number extrapolation, model consistency matters as much as data fitting. More broadly, they suggest that computational complexity, physics-based closure design, and machine learning must be considered together if a practical $\log(Re)$ solution is to be achieved.

ACKNOWLEDGMENTS

Yang acknowledges AFOSR Grant No. FA9550-23-1-0272 with Dr. Gregg Abate as the technical monitor.

REFERENCES

Bin, Yuanwei, Chen, Lihua, Huang, George & Yang, Xiang I A 2022 Progressive, extrapolative machine learning for near-wall turbulence modeling. *Phys. Rev. Fluid* **7** (8), 084610.

Bin, Yuanwei, Hu, Xiaohan, Li, Jiaqi, Grauer, Samuel J & Yang, Xiang I. A. 2024*a* Constrained re-calibration of two-equation Reynolds-averaged Navier–Stokes models. *Theor. Appl. Mech. Lett.* **14** (2), 100503.

Bin, Yuanwei, Huang, George, Kunz, Robert & Yang, Xiang I. A. 2024*b* Constrained recalibration of Reynolds-averaged Navier–Stokes models. *AIAA J.* **62** (4), 1434–1446.

Breuer, Michael, Peller, Nikolaus, Rapp, Ch & Manhart, Michael 2009 Flow over periodic hills–numerical and experimental study in a wide range of Reynolds numbers. *Computers & Fluids* **38** (2), 433–457.

Choi, Haecheon & Moin, Parviz 2012 Grid-point requirements for large eddy simulation: Chapman's estimates revisited. *Phys. Fluids* **24** (1).

Duraisamy, Karthik, Iaccarino, Gianluca & Xiao, Heng 2019 Turbulence modeling in the age of data. *Annu. Rev. Fluid Mech.* **51**, 357–377.

Durbin, Paul A 2018 Some recent developments in turbulence closure modeling. *Annu. Rev. Fluid Mech.* **50**, 77–103.

Graham, J, Kanov, K, Yang, X I. A., Lee, M, Malaya, N, Lalescu, CC, Burns, R, Eyink, G, Szalay, A, Moser, RD *et al.* 2016 A web services accessible database of turbulent channel flow and its use for testing a new integral wall model for les. *J. Turbul.* **17** (2), 181–215.

He, Ping, Mader, Charles A, Martins, Joaquim RRA & Maki, Kevin J 2020 Dafoam: An open-source adjoint framework for multidisciplinary design optimization with openfoam. *AIAA Journal* **58** (3), 1304–1319.

Holland, Jonathan R, Baeder, James D & Duraisamy, Karthik 2019 Towards integrated field inversion and machine learning with embedded neural networks for RANS modeling. In *AIAA Scitech forum*, p. 1884.

Parish, Eric J & Duraisamy, Karthik 2016 A paradigm for data-driven predictive modeling using field inversion and machine learning. *J. Comput. Phys.* **305**, 758–774.

Rapp, Ch & Manhart, M 2011 Flow over periodic hills: an experimental study. *Experiments in fluids* **51** (1), 247–269.

Rumsey, Christopher L, Coleman, Gary N & Wang, Li 2022 In search of data-driven improvements to RANS models applied to separated flows. In *AIAA Scitech 2022 Forum*, p. 0937.

Singh, Anand Pratap, Medida, Shivaji & Duraisamy, Karthik 2017 Machine-learning-augmented predictive modeling of turbulent separated flows over airfoils. *AIAA J.* **55** (7), 2215–2227.

Slotnick, Jeffrey P, Khodadoust, Abdollah, Alonso, Juan, Darmofal, David, Gropp, William, Lurie, Elizabeth & Mavriplis, Dimitri J 2014 CFD vision 2030 study: a path to revolutionary computational aerosciences. *Tech. Rep.*. NASA.

Spalart, Philippe & Allmaras, Steven 1992 A one-equation turbulence model for aerodynamic flows. In *30th Aerospace Sciences Meeting and Exhibit*, p. 439.

Spalart, Philippe R 2015 Philosophies and fallacies in turbulence modeling. *Prog. Aerosp. Sci.* **74**, 1–15.

Srivastava, Vishal 2022 Generalizable data-driven model augmentations using learning and inference assisted by feature-space engineering. PhD thesis, University of Michigan.

Srivastava, Vishal, Rumsey, Christopher L, Coleman, Gary N & Wang, Li 2024 On generalizably improving RANS predictions of flow separation and reattachment. In *AIAA SCITECH 2024 Forum*, p. 2520.

Vadrot, Aurélien, Yang, Xiang I. A., Bae, H Jane & Abkar, Mahdi 2023 Log-law recovery through reinforcement-learning wall model for large eddy simulation. *Phys. Fluids* **35** (5).

Wang, Jianxun, Wu, Jinlong & Xiao, Heng 2017 Physics-informed machine learning approach for reconstructing Reynolds stress modeling discrepancies based on dns data. *Phys. Rev. Fluid* **2** (3), 034603.

Weller, Henry G, Tabor, Gavin, Jasak, Hrvoje & Fureby, Christer 1998 A tensorial approach to computational continuum mechanics using object-oriented techniques. *Comput. phys.* **12** (6), 620–631.

Wu, JinLong, Xiao, Heng & Paterson, Eric 2018 Physics-informed machine learning approach for augmenting turbulence models: A comprehensive framework. *Phys. Rev. Fluid* **3** (7), 074602.

Yang, Xiang I. A. & Griffin, Kevin P 2021 Grid-point and time-step requirements for direct numerical simulation and large-eddy simulation. *Physics of Fluids* **33** (1).

Yang, Zhongxin, Shan, Xianglin, Yang, Xiang I. A. & Zhang, Weiwei 2025 Data-enabled discovery of specific and generalisable turbulence closures. *J. Fluid Mech.* **1016**, R1.