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

PROGRESSIVE MIXTURE-OF-EXPERTS WITH AUTOENCODER ROUTING FOR CONTINUAL RANS TURBULENCE MODELLING

Haoyu Ji
HEDPS, Center for Applied Physics and Technology,
School of Mechanics and Engineering Science,
Peking University
Beijing 100871, China
haoyu.ji@pku.edu.cn

Yinhang Luo
HEDPS, Center for Applied Physics and Technology,
School of Mechanics and Engineering Science,
Peking University
Beijing 100871, China
lnhang0618@alumni.pku.edu.cn

Hanyu Zhou
HEDPS, Center for Applied Physics and Technology,
School of Mechanics and Engineering Science,
Peking University
Beijing 100871, China
hanyuzhou@stu.pku.edu.cn

Yaomin Zhao
HEDPS, Center for Applied Physics and Technology,
School of Mechanics and Engineering Science,
Peking University
Beijing 100871, China
yaomin.zhao@pku.edu.cn

ABSTRACT

In this work, we propose a novel framework, termed the progressive mixture-of-experts (PMoE), designed to enable continual learning for Reynolds-averaged Navier–Stokes turbulence modelling. The framework employs a modular autoencoder-based router to associate each flow scenario with a specialized turbulence model, referred to as an expert. When a new flow regime cannot be adequately represented by the existing router and expert set, a new expert together with its routing component can be introduced at low cost, without modifying or degrading previously trained ones, thereby naturally avoiding catastrophic forgetting. The framework is applied to a range of flows with distinct physical characteristics, including airfoil wake, channel, periodic hill, and square duct flows. The resulting PMoE model effectively integrates multiple experts and achieves improved predictive accuracy across both seen and unseen test cases that differ in operating conditions or configurations. Owing to sparse activation, model expansion does not incur additional computational cost during inference. The proposed framework therefore provides a scalable pathway towards lifelong-learning turbulence models for industrial computational fluid dynamics.

INTRODUCTION

Reynolds-averaged Navier–Stokes (RANS) remains a widely used approach in industrial computational fluid dynamics due to its balance between computational efficiency and predictive capability. Recent advances in machine learning have introduced a variety of data-driven turbulence modelling strategies, including optimization-based, ensemble-based, and symbolic approaches (Duraisamy *et al.*, 2019; Brunton *et al.*, 2020). These models are typically trained for specific flow scenarios and can be regarded as expert models. However, achieving robust generalization across diverse flow regimes remains a major challenge (Rumsey & Coleman, 2022; Sandberg & Zhao, 2022). Existing efforts to improve generalization can be broadly classified into three categories: parameter calibration, multi-case training, and model aggregation (Bin *et al.*, 2023; Fang *et al.*, 2023; de Zordo-Banliat *et al.*, 2024). While these approaches improve performance within or near training distributions, their applicability to unseen flow regimes remains limited. Furthermore, adapting models to new data often requires joint retraining, which may lead to catastrophic forgetting and practical challenges in evolving applications (French, 1999).

To address these limitations, we propose a progressive mixture-of-experts (PMoE) framework based on the mixture-of-experts (MoE) architecture (Jacobs *et al.*, 1991). By decomposing complex flow physics into multiple expert models and activating only a subset through a router, the framework maintains both computational efficiency and modelling flexibility. A key feature of PMoE is its continual-learning capability, enabling the incorporation of new flow regimes without retraining existing components. The proposed framework is designed to address three key challenges: unsupervised classification of flow regimes using RANS data, continual adaptation through a modular router architecture, and the construction of specialized expert models to capture distinct turbulence characteristics. These aspects form the basis for developing a generalizable turbulence modelling framework across diverse flow scenarios.

METHODOLOGY

The MoE framework decomposes turbulence modelling into a set of expert models $E_1,\ldots,E_n$ coordinated by a router G, which assigns weights to their predictions. The model output is given by $y=\sum_{i=1}^{n} G(x)_i E_i(x)$, where $G(x)_i$ denotes the confidence of expert i. Owing to sparse activation (e.g. top-k selection), the computational cost scales with a single expert. However, standard MoE frameworks lack mechanisms for incorporating new regimes without degrading existing performance, as conventional MLP-based routers are prone to catastrophic forgetting when updated with new data (French, 1999). Continual-learning methods in machine learning aim to address this issue through regularization, replay, or architectural strategies (Wang *et al.*, 2024). Among them, an architecture-based approach (Aljundi *et al.*, 2017) that assigns dedicated modules to different tasks is particularly relevant, motivating

Table 1. Flow features used as input in the framework. The normalized feature q_i is obtained by $q_i = \hat{q}_i/(|\hat{q}_i| + |q_i^*|)$.

Feature (q_i)	Raw feature ($\hat{q}_i$)	Normalization factor (q_i^*)
q_1	$\|W\|$	$\|S\|$
q_2	$\frac{1}{2}(\|W\|^2 - \|S\|^2)$	$\|S\|^2$
q_3	$\sqrt{\frac{\partial P}{\partial x_i}\frac{\partial P}{\partial x_i}}$	$\frac{1}{2}\rho\frac{\partial U_k^2}{\partial x_k}$
q_4	d	L_{ref}
q_5	$\tilde{\nu}$	100ν
q_6	$U_l\frac{\partial P}{\partial x_l}$	$\sqrt{\frac{\partial P}{\partial x_j}\frac{\partial P}{\partial x_j}U_iU_i}$
q_7	$\left\|U_iU_j\frac{\partial U_i}{\partial x_j}\right\|$	$\sqrt{U_lU_lU_i\frac{\partial U_i}{\partial x_j}U_k\frac{\partial U_k}{\partial x_j}}$

the development of a continual-learning MoE framework for turbulence modelling.

As shown in figure 1, the PMoE framework comprises pre-processing of flow features from a baseline RANS calculation, a modular autoencoder-based router for flow classification, and a bank of heterogeneous experts each tailored to a specific regime. By replacing a monolithic router with a collection of independently trained autoencoders, the framework prevents interference between old and new modules during progressive expansion, thereby inherently overcoming catastrophic forgetting.

The input to PMoE consists of $M = 7$ local flow features (table 1) extracted from a baseline Spalart–Allmaras (SA) solution (Spalart & Allmaras, 1992), following prior data-driven RANS approaches (Ling & Templeton, 2015; Wang *et al.*, 2017). A physics-informed mask selects points where either vorticity or strain-rate magnitude exceeds a threshold $\varepsilon = 0.05$ times the field median: $|W_j| > \varepsilon \cdot \mathrm{median}(|W|)$ or $|S_j| > \varepsilon \cdot \mathrm{median}(|S|)$. From the masked region, N points are randomly sampled to form the input matrix $X \in R^{N\times M}$.

The router is composed of modular autoencoders, each trained to reconstruct regime-specific features. An encoder–decoder pair maps $x \to z \to \hat{x}$ by minimizing the reconstruction error $\mathscr{L} = \frac{1}{N}\sum_j \sqrt{\frac{1}{M}\sum_i(\hat{x}_{i,j} - x_{i,j})^2}$. For each component C_k, a threshold T_k is defined from the training error distribution. During inference, the confidence p_k is given by the fraction of points with reconstruction error below T_k. The router selects the expert $E_{\mathscr{K}}$ with the highest confidence ($\mathscr{K} = \arg\max_k p_k$). If $p_{\mathscr{K}} \geq 90\%$, $E_{\mathscr{K}}$ is activated; otherwise, the flow is treated as novel and a new autoencoder–expert pair is introduced.

A key advantage of PMoE is its support for heterogeneous expert formulations. Using the SA baseline transport equation, we integrate three distinct expert types. Expert E_1 recalibrates the damping function f_{v1} via symbolic regression (Weatheritt & Sandberg, 2016; Bin *et al.*, 2023) for wall-attached flows. Expert E_2 employs a FIML-trained (field inversion and machine learning) neural network to introduce a multiplicative correction to the production term:

$$\mathscr{P}^* = \beta(x)\cdot C_{b1}\tilde{S}\tilde{\nu}. \tag{1}$$

Expert E_3 predicts a spatially varying coefficient $C_{\mathrm{cr1}}(x)$ within the quadratic constitutive relation (QCR) framework (Spalart, 2000):

$$\tau_{ij,\mathrm{QCR}} = \tau_{ij}^{linear} - C_{\mathrm{cr1}}(x)\left(O_{ik}\tau_{jk} + O_{jk}\tau_{ik}\right). \tag{2}$$

This modular design decouples architecture from baseline constraints, enabling progressive integration of symbolic and neural experts without manual flow-type labeling.

We employ a curriculum learning strategy wherein datasets are introduced in order of increasing geometric complexity. The sequence progresses from the airfoil wake (free-shear turbulence), through fully developed channel flow (equilibrium wall-bounded flow) and periodic hill (curved-wall separation), to the square duct (corner-induced secondary flow). The router autonomously identifies novel regimes via reconstruction-error confidence, while the introduction order is user-prescribed. Since each router component and expert is trained with all prior modules frozen, the final model is invariant to the presentation order.

The model evolves through stages S_0 to S_3, yielding PMoE-S0 to PMoE-S3. At each stage, $N = 10000$ points with $M = 7$ features (table 1) are sampled from the baseline RANS field. If the maximum confidence $p_{\mathscr{K}}$ falls below $T_{\mathrm{accept}} = 90\%$, a new autoencoder component and specialized expert are trained on the novel regime. The datasets and their training/validation roles are detailed in table 2. Regime distinction is fully automatic. In S0, the model is initialized with the 2DANW wake case, where the SA model serves as expert E_0 and achieves 99.9% router confidence, establishing the baseline latent space. In S1, channel flows at $Re_\tau = 2000$ and 5200 yield 0.0% confidence, triggering expert E_1 via symbolic regression correction of f_{v1}; the resulting PMoE-S1 recognizes channel data with 99.5% confidence. In S2, periodic hill flow ($\alpha = 1.0$, $Re = 5600$) yields 58.4% confidence, below the acceptance threshold, leading to expert E_2 with a FIML-based production correction (equation 1); the confidence increases to 93.5%. In S3, square duct flows at $Re = 2500$ and 5693 produce 56.7% confidence, triggering expert E_3 based on the QCR framework (equation 2), and yielding a final confidence of 97.0%.

RESULTS

The validation of the PMoE framework focuses on three critical performance metrics: the interpretability of the unsupervised routing, the prevention of catastrophic forgetting on previously learned regimes, and the generalization capability to cases with different operating conditions or configurations. All results presented herein utilize the final PMoE-S3 model to demonstrate its cumulative capabilities.

The latent space of each autoencoder is analysed using the Mahalanobis distance. Based on this metric, intra- and inter-cluster dispersions are computed to assess compactness and separation. The intra-cluster dispersion remains approximately 1.8, while inter-to-intra ratios exceed 3.4, indicating well-separated clusters. Although routing is based on reconstruction error, the latent-space structure provides an interpretable proxy for routing behavior, yielding robust classification with over 90% accuracy (table 3).

To verify that the modular PMoE framework avoids catastrophic forgetting, we evaluate the final PMoE-S3 model on the training cases from all four stages, as presented in figure 3. The relative RMSE reduction displayed in the top-left corner

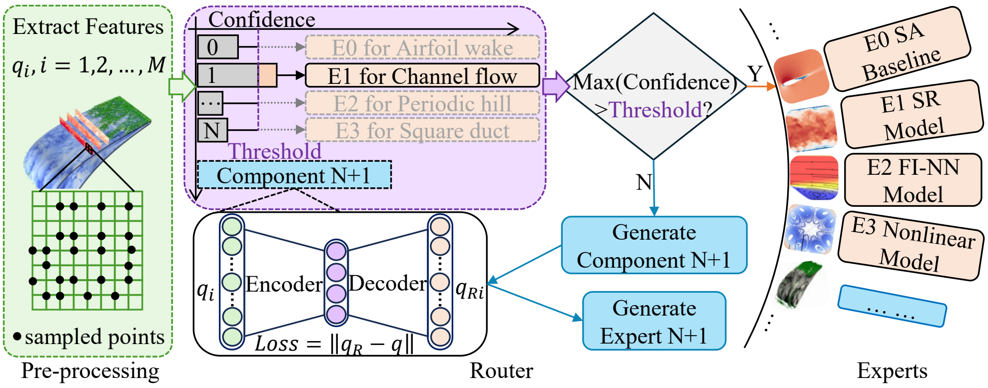

Figure 1. Schematic of the proposed PMoE framework.

Table 2. Summary of flow cases used for training and validation. "T/V" indicates training (T) or validation (V). For each flow regime, references are listed in the caption footnotes.

Regime	Case (Abbreviation)	T/V
Stage 0: Wake[1]	2D Airfoil Near-Wake (2DANW)	T
Stage 1: Channel[2]	$Re_\tau = 2000$ (C2000)	T
	$Re_\tau = 5200$ (C5200)	T
	$Re_\tau = 8000$ (C8000)	V
Stage 2: Periodic hill[3]	$\alpha = 0.8$ (PH0p8)	V
	$\alpha = 1.0$ (PH1p0)	T
	$\alpha = 1.2$ (PH1p2)	V
	$\alpha = 1.5$ (PH1p5)	V
Stage 3: Square duct[4]	$Re = 2500$ (SD2500)	T
	$Re = 3500$ (SD3500)	V
	$Re = 5693$ (SD5693)	T

[1] Nakayama (1985); [2] Lee & Moser (2015); Yamamoto & Tsuji (2018); [3] Xiao *et al.* (2020); [4] Pinelli *et al.* (2010); Vinuesa *et al.* (2018).

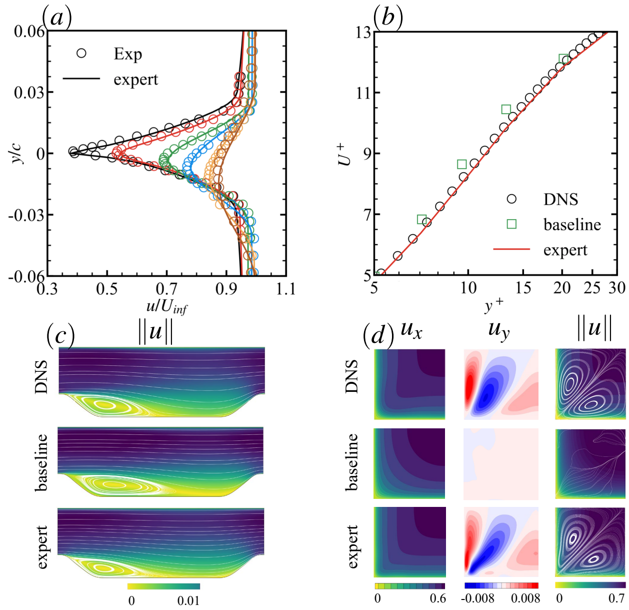

Figure 2. Velocity profiles and contours predicted by the baseline SA model, the expert model and high-fidelity data. (a) S0: 2DANW. (b) S1: C5200. (c) S2: PH1p0. (d) S3: SD2500 and SD5693.

of each panel is defined as

$$\mathscr{R} = (\mathrm{RMSE}_{\mathrm{base}} - \mathrm{RMSE}_{\mathrm{PMoE}})/\mathrm{RMSE}_{\mathrm{base}}, \quad (3)$$

where the root-mean-square error is $\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(u_i - u_i^{\mathrm{ref}})^2}$. Here u_i denotes the model prediction at the i-th reference location for the velocity component shown, and model predictions are interpolated onto reference locations when necessary. For the 2DANW wake (a), predictions remain essentially identical to the baseline SA and agree well with experiments, despite subsequent model expansions, confirming that expert E_0 is not degraded. For channel flows at $Re_\tau = 2000$ and 5200 (b,c) and the periodic hill case ($\alpha = 1.0$) (e), PMoE-S3 preserves the specialized accuracy of experts E_1 and E_2, including accurate reattachment prediction. For square duct flows (d,f–j), which involve corner-induced secondary motions absent in linear eddy-viscosity models, PMoE-S3 yields substantial improvements. Streamwise velocity profiles at $z = 0.3h$ (h is the duct half-height) (d,f) and cross-flow distributions (i) for $Re = 2500$, together with corresponding results at $Re = 5693$ (g,h,j), show consistent gains, demonstrating that expert E_3 captures the nonlinear anisotropy required for secondary flow prediction.

Figure 4 evaluates the robustness of PMoE-S3 on unseen operating conditions. For a channel flow at $Re_\tau = 8000$ and a square duct at $Re = 3500$, both outside the training range, the router correctly assigns the cases to experts E_1 and E_3 (table 3). The predicted velocity profiles (a–e) show excellent agreement with DNS data, indicating that the learned corrections capture the underlying physics rather than overfitting. Further tests on periodic hill flows with varying slopes $\alpha = 0.8, 1.2, 1.5$ (f–h) demonstrate consistent improvements over the baseline SA model, despite training only at $\alpha = 1.0$. While SA overestimates the separation bubble, PMoE-S3 closely matches DNS, showing that expert E_2 provides a generalized separation correction robust to geometric variation. Overall, PMoE-S3 maintains high accuracy across unseen operating conditions while avoiding catastrophic forgetting.

A key practical consideration is computational cost dur-

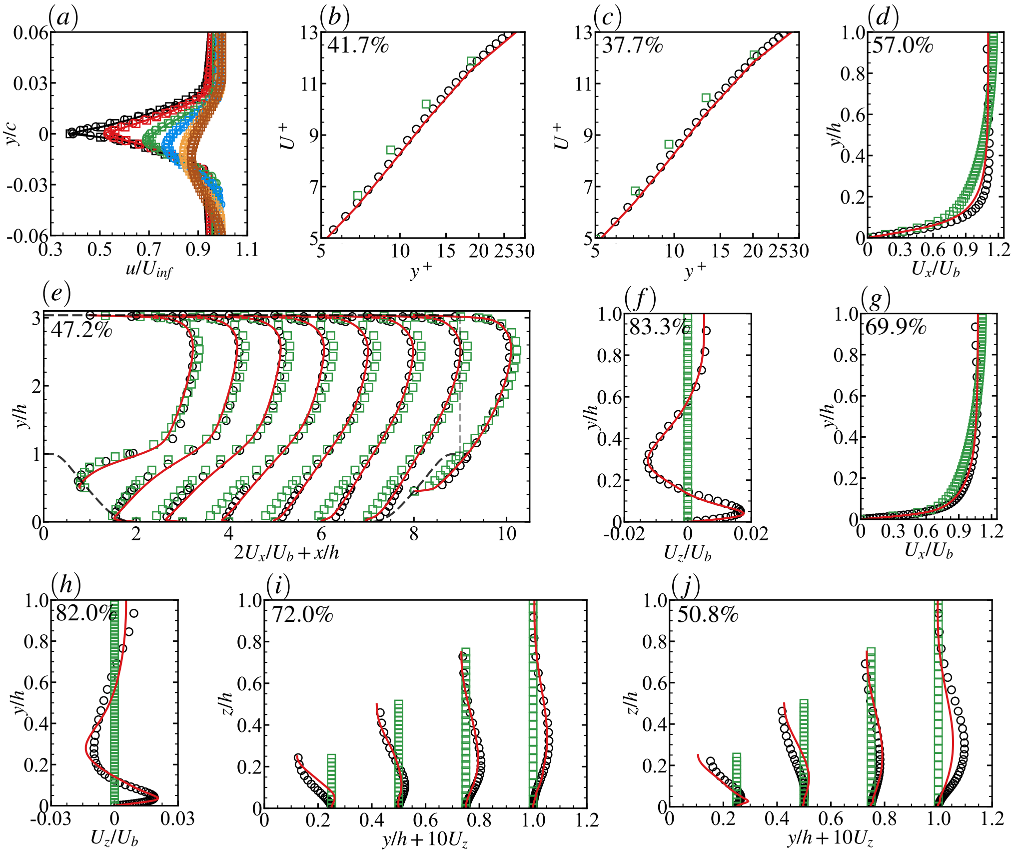

Figure 3. Mean velocity profiles for trained cases. (a) 2DANW (colors distinguish streamwise locations). (b) C2000. (c) C5200. (e) PH1p0. (d, f, i) SD2500. (g, h, j) SD5693. Symbols: ○ DNS/Experimental; baseline SA; —— PMoE-S3. The number shown in the top-left corner of each panel is $\mathscr{R}$ from equation 3. No value is shown in (a) as Expert E_0 coincides with the baseline SA model.

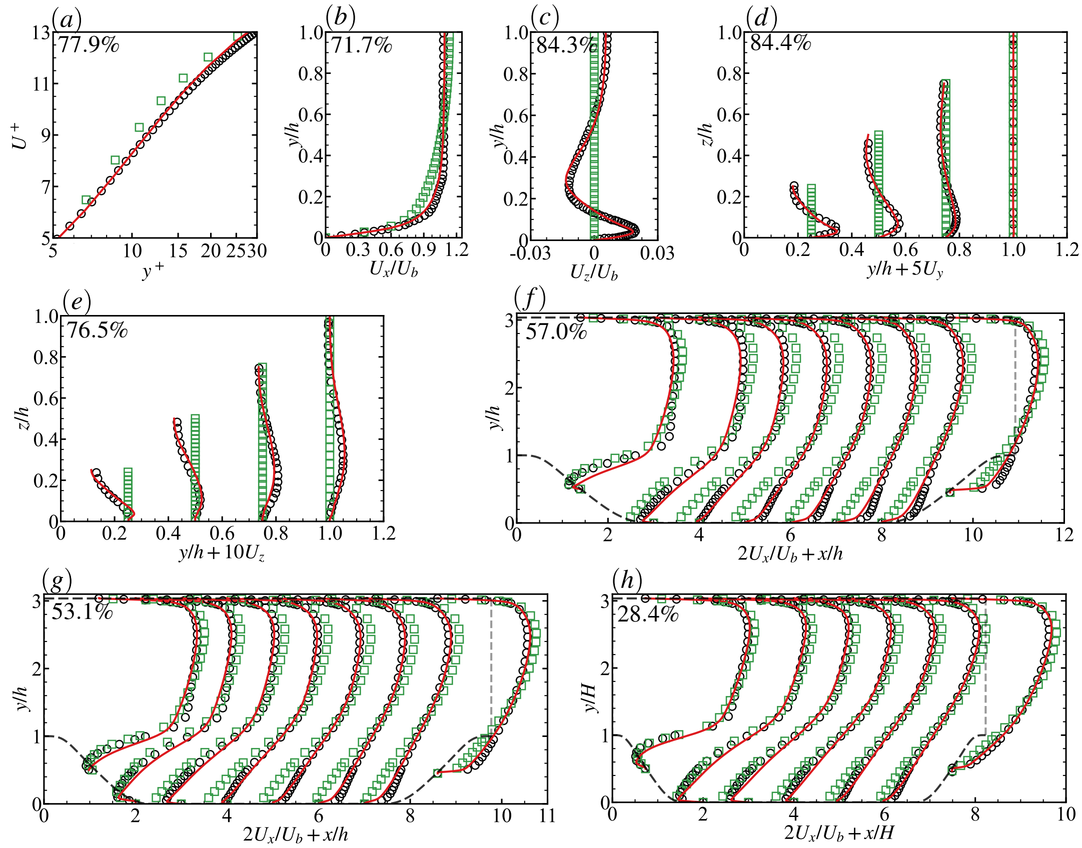

Figure 4. Mean velocity profiles for unseen cases with different operating conditions. (a) C8000. (b–e) SD3500. (f) PH1p5. (g) PH1p2. (h) PH0p8. Symbols denote: ○ DNS/Experimental data; Baseline SA; —— PMoE-S3. The number shown in the top-left corner of each panel is $\mathscr{R}$ from equation 3.

Table 3. Confidence distribution of the PMoE-S3 router over different components for all flow cases.

Case	**Unknown**	C_0	C_1	C_2	C_3
ANW	0.0%	**99.7%**	0.3%	0.0%	0.0%
C2000	0.2%	0.0%	**98.1%**	0.5%	1.2%
C5200	0.0%	0.0%	**99.4%**	0.0%	0.6%
C8000	3.9%	0.0%	**94.5%**	1.3%	0.3%
PH0p8	0.3%	0.0%	5.7%	**91.6%**	2.5%
PH1p0	0.1%	0.0%	5.4%	**91.5%**	3.0%
PH1p2	0.1%	0.0%	4.8%	**92.0%**	3.2%
PH1p5	0.2%	0.0%	4.7%	**91.9%**	3.2%
SD2500	0.4%	0.0%	0.3%	0.0%	**99.3%**
SD3500	0.5%	0.0%	0.0%	0.2%	**99.3%**
SD5693	0.0%	0.0%	2.2%	0.8%	**97.0%**

ing inference. The pipeline consists of baseline SA simulation, autoencoder routing, and expert-corrected RANS. In industrial settings where a baseline solution is preexisting, only routing and correction incur additional cost; if the router selects E_0 (standard SA), no extra simulation is needed. Timing tests are conducted on an Intel® Xeon® Gold 6530 CPU and an NVIDIA® AD102 GPU. Routing time remains below 10 seconds, confirming negligible router overhead. The full-pipeline cost increases by 100%–104% over baseline, reflecting a worst-case scenario of effectively two RANS runs. Crucially, sparse activation ensures that adding experts does not increase per-case cost: expense scales with a single model rather than the ensemble.

The PMoE framework's generalization relies on the router's ability to identify relevant experts via learned feature-space representations and on the experts themselves, which encode physically meaningful corrections rather than case-specific fitting. To evaluate both aspects, we introduce three additional cases: the 2D NASA wall-mounted hump (2DWMH), the curved backward-facing step at $Re = 13,700$ (CBFS), and a rectangular duct at $Re = 2500$ with aspect ratio three (RD2500). These cases share physical mechanisms with the trained regimes—separation under adverse pressure gradients and corner-induced secondary flow—yet introduce substantial geometric and boundary-condition variations. The 2DWMH case further tests the coupling between free-stream, wall-attached, and separated regions within a single external-flow configuration.

The rectangular duct, closely resembling the square duct training case, is assigned to C_3 with 90.8% confidence. CBFS, governed by adverse-pressure-gradient separation akin to the periodic hill, yields 76.7% confidence for C_2. The 2DWMH case, involving separation over a convex hump in an external flow with an upstream attached boundary layer, produces a more mixed representation: 54.7% to C_2 and 35.2% to C_1. This pattern confirms that higher geometric and boundary-condition similarity translates into higher router confidence.

Beyond the router's discrimination, the experts themselves exhibit robust generalization. For CBFS, expert E_2 substantially reduces the separation bubble overprediction and accurately locates reattachment (figure 5a), with reference LES

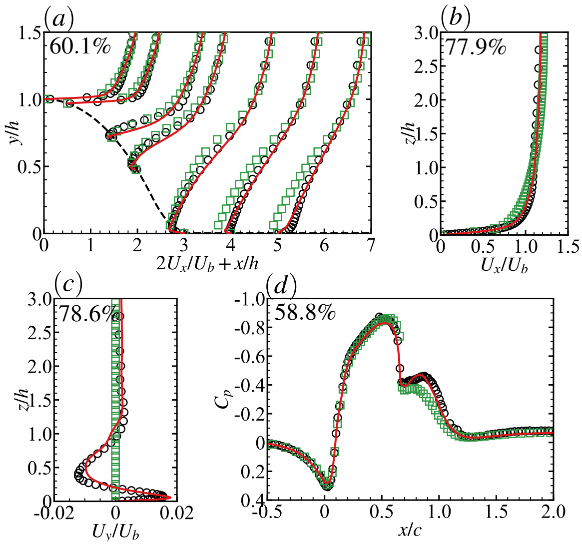

Figure 5. Predictions obtained with the highest-confidence experts for additional test cases. (a) CBFS. (b, c) RD2500. (d) 2DWMH. Symbols denote: ○ DNS/LES/Experimental data; Baseline SA; —— Expert model. The number shown in the top-left corner of each panel is $\mathscr{R}$ from equation 3.

data (Bentaleb *et al.*, 2012). For 2DWMH, despite lower router confidence (below the 90% acceptance threshold), manually activating E_2 improves the pressure coefficient distribution (figure 5d), relative to experiments (Naughton *et al.*, 2006). For RD2500, expert E_3 accurately reproduces corner-induced secondary vortices and improves the streamwise velocity profile taken at one quarter of the duct height from the bottom wall (figures 5(b) and (c)) compared with DNS (Vinuesa *et al.*, 2018), confirming that the QCR-based correction captures essential anisotropy rather than a geometry-specific fit.

Collectively, these results demonstrate that the PMoE framework's generalization stems from the synergy between a physically interpretable router and robust, physics-based experts. The router provides a confidence measure reflecting feature-space similarity, while the experts transfer effectively across geometric variations when the dominant physical mechanism is preserved. Even when confidence falls below the 90% acceptance threshold, the most relevant expert can still improve upon the baseline, as seen for 2DWMH. Conversely, the threshold prevents spurious corrections when no expert is sufficiently relevant. This combination offers a degree of fault tolerance not typical of monolithic data-driven models and suggests a practical extension strategy: when an expert's limits are reached, only that expert and its router component require refinement.

The present PMoE router classifies flows based on features extracted from the baseline SA solution, assuming that the flow can be adequately characterized by the features in table 1 and that the SA model provides a qualitatively correct field. As demonstrated earlier, even notable SA inaccuracies have limited impact on routing outcomes owing to the router's reliance on normalized invariant features; however, a fundamentally incorrect baseline topology may compromise the extracted features and render the framework unreliable. A global top-1 routing strategy is adopted, activating the highest-confidence expert to ensure numerical stability, consistent with industrial practice (Menter & Matyushenko, 2025). This approach performs well for flows similar or moderately related to training regimes and defaults to the baseline SA model for unknown cases. For complex configurations where multiple flow regimes coexist, a single top-1 decision may be insufficient. A region-aware strategy could better match corrections to local physics but introduces challenges with interface discontinuities and solver convergence. A top-k soft blending approach with continuous weights would enable smooth transitions but is currently hindered by the heterogeneous nature of the expert outputs. Adopting a unified expert format, such as a common correction to the Reynolds stress anisotropy tensor, would facilitate blending at the cost of regime-specific flexibility. Resolving this trade-off will be a focus of future work.

CONCLUSIONS

We have proposed a progressive modular mixture-of-experts framework (PMoE) for continual learning in data-driven RANS turbulence modelling. The framework employs a successively trainable autoencoder-based router to progressively introduce specialized experts as new flow regimes emerge. Trained on four representative regimes of airfoil wakes, channel flows, periodic hill flows, and square duct flows, the PMoE-S3 model delivers accurate predictions for both seen and unseen test cases with distinct operating conditions or configurations.

In the PMoE framework, new regimes are incorporated by adding dedicated experts and router components, leaving previously trained modules unchanged. The modular design allows experts of fundamentally different forms to coexist, offering flexibility to integrate heterogeneous turbulence models. The training relies on unsupervised feature extraction, making the framework well suited to industrial workflows where labeled data and repeated retraining are impractical. Owing to sparse activation, inference cost does not scale with the number of experts, which is essential for industrial CFD. The plug-and-play structure enables modular updates and reduces development costs, while its extensibility allows the community to contribute new experts targeting unexplored flow regimes, establishing a living framework that evolves alongside modelling advances.

Future work may extend the expert library to alternative closures, including the k–ω SST model and Reynolds stress models. Region-aware routing strategies will be needed for flows where multiple regimes coexist locally, and the input feature set may require enrichment to accommodate additional flow types beyond those considered here.

REFERENCES

Aljundi, R., Chakravarty, P. & Tuytelaars, T. 2017 Expert gate: Lifelong learning with a network of experts. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 3366–3375.

Bentaleb, Y., Lardeau, S. & Leschziner, M.A. 2012 Large-eddy simulation of turbulent boundary layer separation from a rounded step. *J. Turbul.* (13), N4.

Bin, Y., Huang, G. & Yang, X.I.A. 2023 Data-Enabled recalibration of the Spalart–Allmaras model. *AIAA J.* **61** (11), 4852–4863.

Brunton, S.L., Noack, B.R. & Koumoutsakos, P. 2020 Machine learning for fluid mechanics. *Ann. Rev. Fluid Mech.* **52**, 477–508.

Duraisamy, K., Iaccarino, G. & Xiao, H. 2019 Turbulence modeling in the age of data. *Ann. Rev. Fluid Mech.* **51**, 357–377.

Fang, Y., Zhao, Y., Waschkowski, F., Ooi, A.S.H. & Sandberg, R.D. 2023 Toward more general turbulence models via multicase computational-fluid-dynamics-driven training. *AIAA J.* **61** (5), 2100–2115.

French, R.M. 1999 Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences* **3** (4), 128–135.

Jacobs, R.A., Jordan, M.I., Nowlan, S.J. & Hinton, G.E. 1991 Adaptive mixtures of local experts. *Neural Computation* **3** (1), 79–87.

Lee, M. & Moser, R.D. 2015 Direct numerical simulation of turbulent channel flow up to $Re_\tau = 5200$. *J. Fluid Mech.* **774**, 395–415.

Ling, J. & Templeton, J. 2015 Evaluation of machine learning algorithms for prediction of regions of high Reynolds averaged Navier Stokes uncertainty. *Phys. Fluids* **27** (8), 085103.

Menter, F.R. & Matyushenko, A. 2025 Generalized k- ω (GEKO) two-equation turbulence model. *AIAA J.* **63** (11), 4590–4606.

Nakayama, A. 1985 Characteristics of the flow around conventional and supercritical airfoils. *J. Fluid Mech.* **160**, 155–179.

Naughton, J.W., Viken, S. & Greenblatt, D. 2006 Skin friction measurements on the NASA hump model. *AIAA J.* **44** (6), 1255–1265.

Pinelli, A., Uhlmann, M., Sekimoto, A. & Kawahara, G. 2010 Reynolds number dependence of mean flow structure in square duct turbulence. *J. Fluid Mech.* **644**, 107–122.

Rumsey, C.L. & Coleman, G.N. 2022 NASA symposium on turbulence modeling: Roadblocks, and the potential for machine learning. Technical report. Langley Research Center, Hampton, VA, USA.

Sandberg, R.D. & Zhao, Y. 2022 Machine-learning for turbulence and heat-flux model development: A review of challenges associated with distinct physical phenomena and progress to date. *Int. J. Heat Fluid Flow* **95**, 108983.

Spalart, P.R. 2000 Strategies for turbulence modelling and simulations. *Int. J. Heat Fluid Flow* **21** (3), 252–263.

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

Vinuesa, R., Schlatter, P. & Nagib, H.M. 2018 Secondary flow in turbulent ducts with increasing aspect ratio. *Phys. Rev. Fluids* **3** (5), 054606.

Wang, J., Wu, J. & Xiao, H. 2017 Physics-informed machine learning approach for reconstructing Reynolds stress modeling discrepancies based on DNS data. *Phys. Rev. Fluids* **2**, 034603.

Wang, L., Zhang, X., Su, H. & Zhu, J. 2024 A comprehensive survey of continual learning: Theory, method and application. *IEEE Trans. Pattern Anal. Mach. Intell.* **46** (8), 5362–5383.

Weatheritt, J. & Sandberg, R. 2016 A novel evolutionary algorithm applied to algebraic modifications of the RANS stress–strain relationship. *J. Comp. Phys.* **325**, 22–37.

Xiao, H., Wu, J.-L., Laizet, S. & Duan, L. 2020 Flows over periodic hills of parameterized geometries: A dataset for data-driven turbulence modeling from direct simulations. *Comp. Fluids* **200**, 104431.

Yamamoto, Y. & Tsuji, Y. 2018 Numerical evidence of logarithmic regions in channel flow at $Re_\tau = 8000$. *Phys. Rev. Fluids* **3** (1), 012602.

de Zordo-Banliat, M., Dergham, G., Merle, X. & Cinnella, P. 2024 Space-dependent turbulence model aggregation using machine learning. *J. Comp. Phys.* **497**, 112628.