CWE-8
June 7-11, 2026
London, ON, Canada

Differentiable Hybrid Neural-CFD Modeling of 3D Wall-Bounded Turbulence

Xiantao Fan[a],Fei Ding[b]

[a]Cornell University, Ithaca, NY, USA, xf245@cornell.edu
[b]University of Utah, Salt Lake City, Utah, USA, fei.ding@utah.edu

Summary

Accurate prediction of turbulent wind flows is essential for wind engineering applications, including modeling the atmospheric boundary layer and evaluating wind loads on structures. However, conventional Large-Eddy Simulation (LES) approaches suffer from degraded accuracy on coarse meshes and in complex geometries due to limitations in subgrid-scale (SGS) models and wall treatment. To address these challenges, this work presents a differentiable hybrid neural–CFD framework for high-fidelity simulation of unsteady, wall-bounded turbulence at extremely coarse resolutions. The proposed approach embeds deep neural operators directly into a physics-based CFD solver within a fully differentiable programming environment. This enables end-to-end training through gradient-based optimization. The framework achieves strong a *posteriori* performance and robust generalization without requiring large training datasets. Validation across canonical flows and wind flows around building clusters demonstrates accurate reconstruction of key turbulence statistics at a substantially reduced computational cost, highlighting the method's promise for scalable, data-informed wind engineering simulations.

Keywords: *differentiable physics, CFD-consistent modeling, end-to-end training, wall-bounded turbulence, closure model*

1 INTRODUCTION

Accurate and computationally efficient prediction of turbulent wind flows is fundamental to computational wind engineering. This capability is essential for applications ranging from modeling the atmospheric boundary layer and estimating wind loads on structures, to assessing flow over complex terrain. LES has emerged as a powerful tool for resolving unsteady turbulent structures, yet its practical deployment in large-scale, structure-level computational wind engineering remains limited by the accuracy of SGS models and wall treatments on coarse meshes.

For large-scale simulations, the high computational cost often requires coarse spatial resolution, where traditional SGS closures and wall models fail to capture near-wall turbulence. While recent machine learning enabled turbulence closures have shown promise, their integration into fully coupled CFD solvers poses significant challenges, particularly in achieving stable and robust a *posteriori* performance. Many existing approaches rely on large training datasets, limiting their generalizability and physical consistency.

In this work, we propose a differentiable hybrid neural–CFD framework that overcomes these limitations by embedding deep neural operators directly within a physics-based CFD solver. Implemented in a fully differentiable programming environment, the framework enables end-to-end training through gradient-based optimization of flow-resolved objectives. Importantly, the approach achieves strong generalization with relatively limited training data, enabling the differentiable solver to be effectively applied to civil engineering problems involving both micro-scale and infrastructure-scale aerodynamics. This allows for the accurate evaluation of wind hazards across multiple turbulence scales ranging from localized flow features around buildings to infrastructure-level responses, providing high-fidelity predictions while substantially reducing the computational cost compared to traditional LES simulations.

2 METHODOLOGY

2.1 Problem formulation

Large-eddy simulation (LES) provides a feasible and widely used framework for modeling wall-bounded turbulence. The filtered incompressible Navier–Stokes equations are given by

$$\frac{\partial \bar{u}_i}{\partial t} + \frac{\partial}{\partial x_j}\left(\bar{u}_i \bar{u}_j\right) = -\frac{1}{\rho}\frac{\partial \bar{p}}{\partial x_i} + \nu \frac{\partial}{\partial x_j}\left(\frac{\partial \bar{u}_i}{\partial x_j} + \frac{\partial \bar{u}_j}{\partial x_i}\right) - \frac{\partial \tau_{ij}}{\partial x_j}, \tag{1}$$

where $\bar{u}_i$ is the filtered velocity; $i, j = 1, 2, 3$ correspond to the streamwise (x), wall-normal (y), and spanwise (z) directions, respectively; t denotes time; $\bar{p}$ is the filtered pressure; ν is the kinematic viscosity; ρ is the density; and τ_{ij} represents the subgrid-scale (SGS) stress tensor.

Traditionally, the SGS stress is modeled using a linear eddy-viscosity hypothesis,

$$\tau_{ij} = -2\nu_t \bar{S}_{ij}, \tag{2}$$

where $\bar{S}_{ij}$ is the filtered strain-rate tensor and ν_t is the SGS eddy viscosity, commonly modeled using the Smagorinsky model [1], the Vreman model [2], among others.

Because SGS models alone cannot accurately represent near-wall turbulence, particularly over complex surfaces such as buildings, LES often employs a wall model. A widely used approach is the equilibrium wall function,

$$u^+ = \frac{1}{\kappa}\ln y^+ + C, \tag{3}$$

where $u^+ = U/u_\tau$ is the nondimensional velocity, $y^+ = y u_\tau/\nu$ is the wall coordinate, u_τ is the friction velocity, κ is the von Kármán constant, and C is an empirical constant.

Although these SGS and wall models have been successfully applied to many wall-bounded turbulence problems (e.g., wind engineering), several significant challenges remain. Constant-coefficient SGS models often introduce excessive dissipation, while dynamic SGS models may suffer from numerical instability and increased computational cost. Moreover, SGS models and wall models are typically developed independently, despite their strong coupled influence on accuracy and overall solution fidelity. Their performance is also highly sensitive to numerical discretization choices, and many model parameters require adaptation for non-canonical flow configurations.

These limitations motivate the need to leverage data-driven approaches and machine learning to develop SGS and wall models that are consistent with the underlying CFD solver. The goal of this work is to address these challenges by constructing unified, physics-consistent models that improve accuracy, stability, and generalizability in wall-bounded turbulence simulations.

2.2 Proposed framework

To develop hybrid neural models that are consistent with the underlying CFD solver, an *a posteriori* framework with end-to-end optimization is essential. Building on our differentiable simulation platform, Diff-FlowFSI [3], we propose the modeling framework illustrated in Fig. 1. At each time step, a neural SGS model and a neural wall model are evaluated jointly, allowing their coupled effects to be captured in a physics-consistent manner.

For the neural SGS model, inspired by the Vreman model, we employ a 3D U-Net architecture to predict a global model coefficient C_{ML} from the local flow features, including velocity, pressure, and strain-rate fields. This coefficient is used to compute the spatially varying SGS viscosity according to

$$\nu_t^{\text{Vreman}} = C_{\text{ML}} \sqrt{\frac{B_\beta}{\alpha_{ij}\alpha_{ij}}}, \tag{4}$$

where $\beta_{ij} = \sum_m \Delta_m^2 \alpha_{mi}\alpha_{mj}$ and $\alpha_{ij} = \partial u_j / \partial x_i$.

Simultaneously, a neural wall model estimates the friction velocity u_τ using a 2D U-Net. The inputs consist of velocity and pressure sampled at a matching location above the wall. The wall network is designed as a residual model that predicts the discrepancy between the true friction velocity and that obtained from the classical log-law. This residual formulation improves robustness and generalization. Instead of directly imposing u_τ as a wall boundary condition, our framework converts it into an SGS viscosity at the wall, ensuring a consistent coupling between the wall model and the SGS model. As a result, both neural components influence the solution through the shared variable ν_t, which ultimately modifies the predicted velocity and pressure fields.

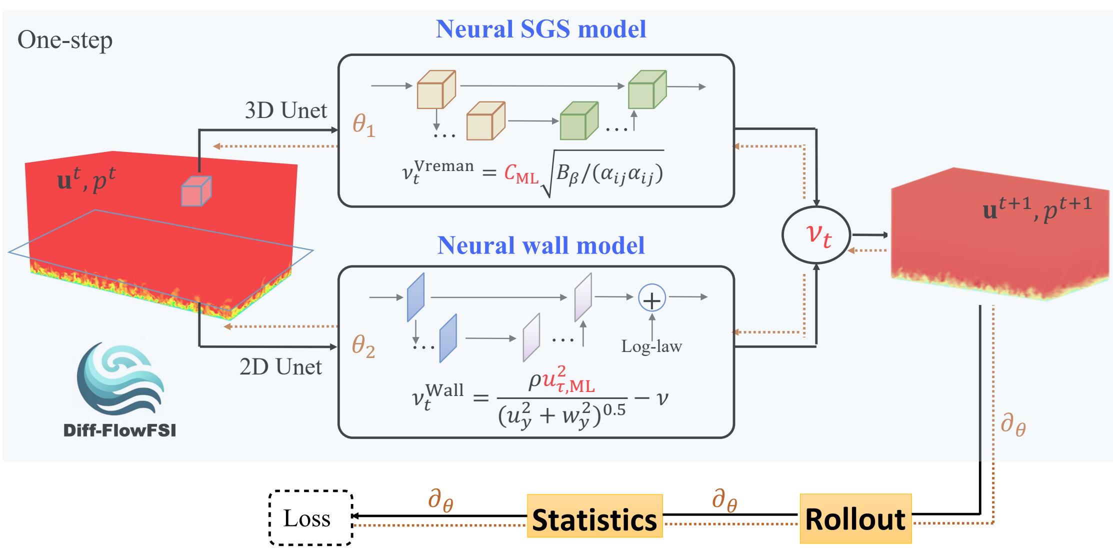

Figure 1: Schematic of the proposed differentiable neural-CFD modeling framework.

Because the one-step module described above is fully differentiable, it can be rolled out over many time steps to obtain converged turbulence statistics, which is crucial for downstream wind-engineering applications. In this work, we roll out the system for $N = 400$ steps for flow with $Re_\theta = 600$ during training, corresponding to approximately two flow-through times. The loss function is defined directly on the statistical quantities:

$$\mathcal{L}(\theta_1, \theta_2) = \|\bar{\hat{u}}_1 - \bar{u}_1\|_2 + \|\hat{u}_i^{\text{rms}} - u_i^{\text{rms}}\|_2 + \|\hat{p}^{\text{rms}} - p^{\text{rms}}\|_2 + \|\hat{u}_\tau - u_\tau\|_2, \tag{5}$$

where $\widehat{\boxed{\cdot}}$ denotes predicted quantities, $\overline{\boxed{\cdot}}$ represents ensemble averages, and $\boxed{\cdot}^{\text{rms}}$ indicates root-mean-square values.

A key advantage of our approach is that it requires only low-order turbulence statistics for training, in contrast to existing models that rely on full instantaneous fields. First- and second-order velocity and pressure statistics are readily obtainable from laboratory measurements, even for non-canonical configurations such as high-Reynolds-number atmospheric boundary layers, flows over complex terrains, or wind-building interactions. Prior knowledge from canonical flows can also be incorporated naturally.

By embedding physics directly into the differentiable solver and optimizing both neural models in an end-to-end *a posteriori* fashion, the proposed framework produces hybrid neural-CFD models that are accurate, robust, and generalizable.

3 *A POSTERIORI* TEST ON UNSEEN FLOW

We conducted an *a posteriori* evaluation on a completely unseen flow configuration at $Re_\theta = 800$, using different initial conditions, domain sizes, and mesh resolutions from those employed in training. The computational settings are summarized in Table 1. Each simulation was integrated for a total of 40 flow-through times, and the final 10 flow-through times were used to compute statistically converged turbulence quantities.

The proposed hybrid neural-CFD model substantially outperforms the baseline approaches across all evaluated metrics. In particular, the energy spectra are predicted with high accuracy despite the use of a coarse computational mesh, a regime in which traditional LES closures typically underpredict spectral energy. These results demonstrate that the proposed model exhibits strong generalizability and robustness when applied to flow conditions outside the training distribution.

Table 1: Computational settings for the testing case.

Domain	Resolution	Δx^+	Δy^+	Δz^+	T
$300 \times 80 \times 160$	$192 \times 80 \times 128$	62	40	50	40

4 CONCLUSIONS

In summary, the proposed hybrid neural-CFD framework demonstrates strong *a posteriori* performance and generalizability when tested on an unseen flow. The model accurately predicts key turbulence statistics even on coarse meshes where traditional SGS models typically fail. By jointly learning SGS and wall-model corrections within a differentiable solver, the framework effectively captures multiscale dynamics that standard closures cannot represent. These results highlight the robustness and practical applicability of the approach for non-canonical wall-bounded turbulence. Overall, our findings suggest a promising pathway toward data-informed, CFD-consistent turbulence modeling for complex wind engineering flows.

ACKNOWLEDGEMENTS

REFERENCES

[1] Joseph Smagorinsky. General circulation experiments with the primitive equations: I. the basic experiment. *Monthly weather review*, 91(3):99–164, 1963.

[2] AW Vreman. An eddy-viscosity subgrid-scale model for turbulent shear flow: Algebraic theory and applications. *Physics of fluids*, 16(10):3670–3681, 2004.

[3] Xiantao Fan, Xin-Yang Liu, Meng Wang, and Jian-Xun Wang. Diff-flowfsi: A gpu-optimized differentiable cfd platform for high-fidelity turbulence and fsi simulations. *Computer Methods in Applied Mechanics and Engineering*, 448:118455, 2026.