

# Large-Eddy Simulation of Flow over Boeing Gaussian Bump Using Multi-Agent Reinforcement Learning Wall Model

Di Zhou \*

*California Institute of Technology, Pasadena, CA 91125, USA.*

Michael P. Whitmore †

*Stanford University, Palo Alto, CA 94305, USA.*

Kevin P. Griffin ‡

*National Renewable Energy Laboratory, Golden, CO 80401, USA.*

H. Jane Bae §

*California Institute of Technology, Pasadena, CA 91125, USA.*

We develop a wall model for large-eddy simulation (LES) that takes into account various pressure-gradient effects using multi-agent reinforcement learning. The model is trained using low-Reynolds-number flow over periodic hills with agents distributed on the wall at various computational grid points. It utilizes a wall eddy-viscosity formulation as the boundary condition to apply the modeled wall shear stress. Each agent receives states based on local instantaneous flow quantities at an off-wall location, computes a reward based on the estimated wall-shear stress, and provides an action to update the wall eddy viscosity at each time step. The trained wall model is validated in wall-modeled LES of flow over periodic hills at higher Reynolds numbers, and the results show the effectiveness of the model on flow with pressure gradients. The analysis of the trained model indicates that the model is capable of distinguishing between the various pressure gradient regimes present in the flow. To further assess the robustness of the developed wall model, simulations of flow over the Boeing Gaussian bump are conducted at a Reynolds number of  $2 \times 10^6$ , based on the free-stream velocity and the bump width. The results of mean skin friction and pressure on the bump surface, as well as the velocity statistics of the flow field, are compared to those obtained from equilibrium wall model (EQWM) simulations and published experimental data sets. The developed wall model is found to successfully capture the acceleration and deceleration of the turbulent boundary layer on the bump surface, providing better predictions of skin friction near the bump peak and exhibiting comparable performance to the EQWM with respect to the wall pressure and velocity field. We also conclude that the subgrid-scale model is crucial to the accurate prediction of the flow field, in particular the prediction of separation.

## I. Introduction

LARGE-EDDY simulation is an essential technology for the simulation of turbulent flows. The basic premise of LES is that energy-containing and dynamically important eddies must be resolved everywhere in the domain. However, this requirement is hard to meet in the near-wall region, as the stress-producing eddies become progressively smaller. Because of the cost involved in resolving the near-wall region, routine use of wall-resolved LES (WRLES) is far from being an industry standard, where short turnaround times are needed to explore high-dimensional design spaces. Consequently, most industrial computational fluid dynamics analyses still rely on cheaper but arguably lower-fidelity Reynolds-Averaged Navier-Stokes (RANS) tools. This has motivated the development of the wall-modeled LES (WMLES) approach, which uses LES to predict the turbulence in the outer region of the boundary layer but utilizes a reduced-order model in a relatively coarse grid in the near-wall region to account for the effect of the energetic

---

\*Postdoctoral Scholar Research Associate in Aerospace, Graduate Aerospace Laboratories, Member AIAA

†PhD Candidate, Center for Turbulence Research, Member AIAA

‡Postdoctoral Researcher, National Renewable Energy Laboratory, Member AIAA

§Assistant professor, Graduate Aerospace Laboratories, Senior Member AIAA

near-wall eddies, thus drastically reducing the grid resolution requirement and simultaneously increasing the maximum allowable time step size. Because of such characteristics, WMLES has been anticipated as the next step to enable the increased use of high-fidelity LES in realistic engineering and geophysical applications.

The most popular and well-known WMLES approach is the so-called RANS-based wall modeling [1–4], which computes the wall-shear stress using the RANS equations. To account for the effects of nonlinear advection and pressure gradient, the unsteady three-dimensional RANS equations are solved [2, 4]. However, these models assume explicitly or implicitly a particular flow state close to the wall (e.g. fully-developed turbulence in equilibrium over a flat plate) and/or rely on RANS parametrization which is manually tuned for various pressure-gradient effects. To remove the RANS legacy in wall modeling, Bose and Moin [5] and Bae et al. [6] proposed a dynamic wall model using slip wall boundary conditions for all three velocity components. Although these models are free of *a priori* specified coefficients and add negligible additional cost compared to the traditional wall models, they were found to be sensitive to numerics of the flow solver and the subgrid-scale (SGS) model, which hinders the application of these models in WMLES. The recent rise of machine learning has prompted supervised learning as an attractive tool for discovering robust wall models that automatically adjust for different conditions, such as variations in the pressure gradient. Zhou et al. [7] proposed a data-driven wall model that accounts for pressure-gradient effects based on the supervised learning approach. While the trained model performed well in *a priori* testing for a single time step, the model broke down in *a posteriori* testing due to integrated errors that could not be taken into account via supervised learning.

Recently, Bae and Koumoutsakos [8] demonstrated the efficacy of multi-agent reinforcement learning (MARL) as a model development tool for wall models in a canonical channel and zero-pressure-gradient boundary layer flows. Reinforcement learning (RL) is a learning framework with foundations on dynamic programming [9] and has been used in the applications of flow control [10, 11] and SGS model development [12]. In the study [8], a series of RL agents are distributed along the computational grid points, with each agent receiving local states and rewards and then providing local actions at each time step. The MARL wall model performs as well as the RANS-based equilibrium wall-shear stress model (EQWM) [13, 14], which has been tuned for this particular flow configuration. However, The MARL model is able to achieve these results by training on moderate Reynolds number flows with a reward function only based on the mean wall-shear stress. Furthermore, instead of relying on *a priori* knowledge or RANS parametrization to perfectly recover the wall boundary condition computed from filtered DNS data, RL can develop novel models that are optimized to accurately reproduce the flow quantities of interest. This is achieved by discovering the dominant patterns in the flow physics, which enables the model to generalize beyond the specific conditions used for training. Therefore, the models are trained *in situ* with WMLES and do not require any higher-fidelity flow fields. More recently, Vadrot et al. [15] further improved the capability of the MARL wall model proposed by Bae and Koumoutsakos [8] via normalizing one of the model states in a novel way and it was found that the improved model is able to accurately predict the log law in a channel flow for a wider range of Reynolds numbers.

In the present study, we extend the methodology of [8] for pressure-gradient flows, retraining a wall model on low-Reynolds-number flow over periodic hills and testing on higher-Reynolds-number flows. Our first objective is to develop a wall model for LES based on MARL that is robust to pressure-gradient effects in a data-efficient way. Another objective of this study is to evaluate the performance of the trained wall model in predicting the flow over a three-dimensional tapered Gaussian bump [16] (referred to as the Boeing Gaussian bump). The flow configuration is a canonical case of smooth-body separation of a turbulent boundary layer (TBL) subject to pressure-gradient and surface-curvature effects. As a widely studied flow configuration, extensive experimental data exist [17–21] for validating computational fluid dynamics codes, and various computational approaches, including Reynolds-averaged Navier-Stokes (RANS) methods [17], direct numerical simulation (DNS) [22, 23], hybrid LES–DNS [24], WMLES [25–27] and detached-eddy simulations (DES) [28] have been evaluated and compared for the flow around the bump. The WMLES studies [25–27] showed that the performance of the classical EQWM is still less than satisfactory, particularly for predicting the separation bubble. In the present study, the detailed comparison of simulations with the EQWM and the MARL wall model are conducted in terms of the mean skin friction as well as the pressure on the bump surface, and the prediction of the velocity field.

The remainder of this paper is organized as follows. In Section II, the details of wall model training based on the flow over periodic hills are introduced. In addition, the testing results of the developed wall model for flow over periodic hills at the Reynolds number of training and higher Reynolds numbers are presented. Section III outlines the simulation setup for the flow over the Boeing Gaussian bump, and demonstrates the results of mean skin friction, mean pressure coefficients on the bump surface, and mean velocity in the flow field. Additionally, the simulations using the developed wall model are compared to those using the EQWM and experimental data. Finally, Section IV summarizes the conclusions drawn from this study.

**Table 1** Simulation cases in comparison to reference data, including mesh size, mean separation location  $x_{\text{sep}}$  and mean reattachment location  $x_{\text{rea}}$  at  $Re_H = 10595$ .

Case	Mesh size ( $N_x \times N_y \times N_z$ )	$x_{\text{sep}}/H$	$x_{\text{rea}}/H$
RLWM	$128 \times 64 \times 64$	0.29	4.57
RLWM, coarse mesh	$64 \times 32 \times 32$	0.43	3.66
EQWM	$128 \times 64 \times 64$	0.57	3.05
DNS [36]	$896 \times 448 \times 448$	0.20	4.51
WRLES [37]	$512 \times 256 \times 256$	-0.11	4.31

## II. Model Development and Validation

### A. Methodology

#### 1. Flow simulation

For the flow solver, we utilize a finite-volume, unstructured-mesh LES code [29]. The spatially-filtered incompressible Navier-Stokes equations are solved with second-order accuracy using cell-based, low-dissipative, and energy-conservative spatial discretization and a fully-implicit, fractional-step time-advancement method with the Crank–Nicholson scheme. The Poisson equation for pressure is solved using the bi-conjugate gradient stabilized method (BiCGStab) proposed by Van der Vorst [30]. The subgrid-scale stress is modeled using the dynamic Smagorinsky model (DSM) [31, 32].

For training the wall model, we prefer a flow configuration that (i) has widely available wall-shear stress profiles for several Reynolds numbers and (ii) does not require tuning of the inlet profile or other boundary conditions. The flow over periodically arranged hills in a channel as proposed by Mellen et al. [33] has well-defined boundary conditions, can be computed at affordable costs, and nevertheless inherits all the features of a flow separating from a curved surface and reattaching on a flat plate. Furthermore, the periodic-hill channel does not require configuring the inlet boundary condition for different grid resolutions and wall models, which is necessary for non-periodic flows. This configuration has become a popular benchmark test case for validating computational fluid dynamics codes since numerous experimental and high-fidelity numerical references exist [34–37] and provide extensive data on a wide range of Reynolds numbers covering  $700 \leq Re_H \leq 37000$ , where  $Re_H$  is the Reynolds number based on the hill height  $H$  and the bulk velocity  $U_b$  at the top of the hill.

The periodic-hill channel flow configuration has the dimensions of  $9H \times 3.035H \times 4.5H$  in streamwise ( $x$ ), vertical ( $y$ ), and spanwise ( $z$ ) directions, respectively. In the simulations of the present study, periodic boundary conditions are applied on streamwise and spanwise boundaries, and the EQWM is employed at the top wall. To maintain constant bulk velocity in time, the flow is driven by a time-varying body force following the control algorithm proposed by Balakumar et al. [38]. Two meshes with different densities are used in the present study, and the details of the meshes are listed in Table 1. The meshes are evenly spaced in  $z$  direction, and approximately uniform in both  $x$  and  $y$  directions. A Courant–Friedrichs–Lewy (CFL) number of 1 is used for all simulations.

#### 2. Reinforcement learning architecture

The MARL architecture of wall-model training in the present study is based on the one proposed by Bae and Koumoutsakos [8]. During model training, the agents distributed above the wall receive states based on local instantaneous flow information and a reward based on the estimated wall-shear stress, then provide local actions to update the wall boundary condition at each time step. The agents infer a single optimized policy through their repeated interactions with the flow field to maximize their cumulative long-term rewards.

In order to utilize MARL as a model development tool for wall models, an RL toolbox, smarties [39], is coupled with the aforementioned unstructured flow solver. The RL tool is an open-source C++ library and is optimized for high CPU-level efficiency through fine-grained multi-threading, strict control of cache-locality, and computation-communication overlap. It implements many established deep RL algorithms [39] as well as methods that have been systematically tested for data-driven turbulence modeling [12]. On every step, it asynchronously obtains on-policy data by sampling

the environment and computes the updates by sampling from the replay memory, which advances the gradient steps. Upon completion of all tasks, it applies the gradient update and proceeds to the next step. The coupling between the RL tool and the flow solver has been validated using the same training configuration as the preliminary study of Bae and Koumoutsakos [8].

### 3. Training of the wall model

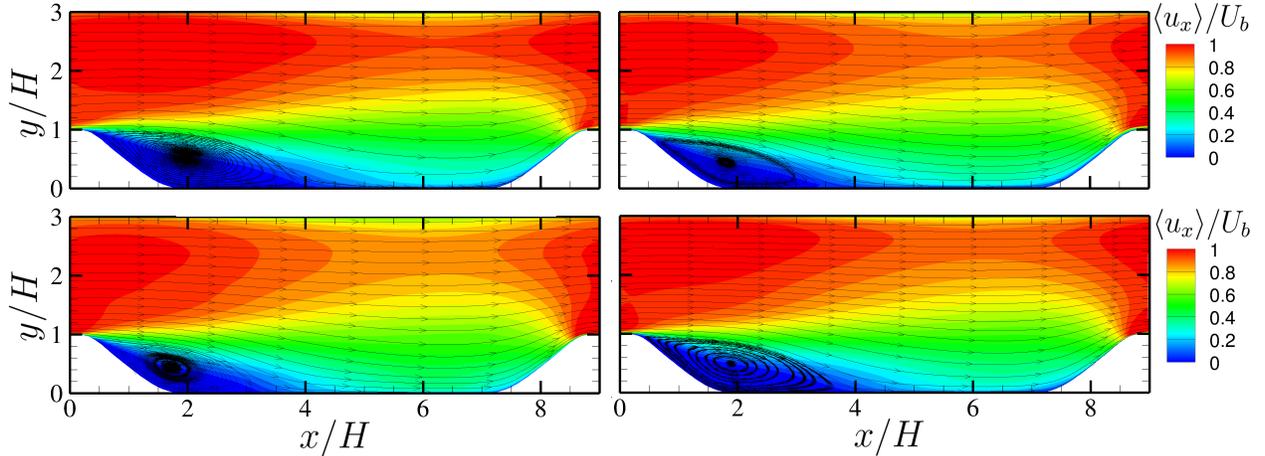
The RL-based wall model (RLWM) training is conducted using the LES of periodic-hill channel flow at  $Re_H = 10595$  with a baseline mesh ( $N_x \times N_y \times N_z = 128 \times 64 \times 64$ ). A total of 512 agents are uniformly distributed along the bottom wall, and the wall-normal locations of the agents  $h_m$  are randomly selected between  $0.01H$  and  $0.09H$  at each agent location.

Considering the existence of log law for inner-scaled mean velocity profile ( $\langle u \rangle^+ = (1/\kappa) \ln y^+ + B$ , where  $\kappa$  the von Kármán constant and  $B$  is the intercept constant) in the near-wall region of turbulent flows, Bae and Koumoutsakos [8] developed a wall model for flat-plate channel flow based on two instantaneous log-law variables  $(1/\kappa_m) = (\partial u / \partial y)^+ h_m^+$  and  $B_m = u^+ - (1/\kappa_m) \ln(h_m^+)$  where the subscript  $m$  denotes that the variable is a modeled quantity, the superscript  $+$  denotes the inner scaling based on friction velocity and kinematic viscosity, and  $h_m$  is the wall-normal location of RL agent. Note that the inner scaling for the instantaneous quantities for the training are computed using the modeled  $u_\tau$  rather than the true  $u_\tau$ . Recently, Vadrot et al. [15] improved the performance and effectiveness of the model at different Reynolds numbers by using  $(1/\kappa_m - 1/\kappa_{\text{ref}}) \ln(h_m^+)$  and  $B_m$  as the model states, where  $\kappa_{\text{ref}}$  is the reference von Kármán constant. On the foundation of these studies, for the current RL wall model, we set the local flow quantities  $(1/\kappa_m - 1/\kappa_{\text{ref}}) \ln(h_m^*)$  and  $B_m = u_s^* - (1/\kappa_m) \ln(h_m^*)$  as the first two model states, where  $1/\kappa_m = (\partial u_s / \partial n)^* h_m^*$ , the subscript  $s$  denotes the wall-parallel direction pointing towards the positive  $x$ -direction,  $n$  denotes the wall-normal direction pointing towards the interior,  $u_s$  is the wall-parallel velocity and  $\kappa_{\text{ref}} = 0.41$ . In order to get better adaptability for flows with different pressure gradients and separations, we also select the turbulence strain rate  $S_{12}^* = (\partial u_s^* / \partial n + \partial u_n^* / \partial s) / 2$  and the local wall-parallel pressure gradient parameter  $(\partial p / \partial s)^* h_m^*$  as the third and fourth model states, where  $u_n$  is the local wall-normal velocity. To increase the applicability of the wall model for a wide range of flow parameters, the states are nondimensionalized using kinematic viscosity  $\nu$  and the composite friction velocity  $u_{\tau p} = (u_\tau^2 + u_p^2)^{1/2}$  introduced by Manhart et al. [40], where  $u_p = |(\nu/\rho)(\partial p_w / \partial s)|^{1/3}$ ,  $\rho$  is the density of the fluid, the subscript  $w$  denotes quantities evaluated at the wall,  $p_w$  is the pressure on the bottom wall and  $u_\tau$  is the friction velocity based on the modeled wall-shear stress. Quantities nondimensionalized by  $\nu$  and  $u_{\tau p}$  are denoted by superscript  $*$ .

It is well known that the method of wall-shear stress application in WMLES affects the mean quantities [38]. A previous study [41] tested two formulations of the boundary condition, namely the wall-shear stress and the wall eddy-viscosity formulations, to determine the appropriate action for the wall model. Zhou et al. [41] found that the wall eddy-viscosity formulation is a more robust method for WMLES of separated flows, and thus we utilize this method for the current wall model. Specifically, in the model, each agent acts to adjust the local wall eddy viscosity  $\nu_{t,w}$  at each time step through a multiplication factor  $\nu_{t,w}(t_{i+1}) = a \nu_{t,w}(t_i)$ , where  $a \in [1 - \alpha \Delta T U_b / \Delta x, 1 + \alpha \Delta T U_b / \Delta x]$ ,  $\Delta T$  is the time step of the simulation,  $\Delta x$  is the grid size in  $x$  direction, and  $\alpha$  is a constant which is selected to be  $10^{-3}$ . The local wall-shear stress can be calculated by the formula  $\tau_w = \rho \nu (1 + \nu_{t,w}^*) (\partial u_s / \partial n)_w$ . The reward  $r$  is calculated based on  $r(t_i) = [|\tau_w^{\text{ref}} - \tau_w(t_i)| - |\tau_w^{\text{ref}} - \tau_w(t_{i-1})|] / \tau_{w,rms}^{\text{ref}}$  at each location, where  $\tau_w^{\text{ref}}$  and  $\tau_{w,rms}^{\text{ref}}$  are the mean and root-mean-square wall-shear stress from the reference simulation. The reward is proportional to the improvement in the modeled wall-shear stress compared to the one obtained in the previous time step, and an extra reward of 0.1 is added when the modeled  $\tau_w$  is within 10% of the reference value.

During model training, each iteration is initialized with the normalized wall eddy viscosity  $\nu_{t,w}^*$  that is randomly selected from  $(0, 10]$ . To generate the initial condition for training, the simulation is started from a flow field generated by the EQWM and run with the given initial  $\nu_{t,w}$  for 20 flow-through times (FTTs) to remove numerical artifacts. Each iteration of the model training simulation is conducted for 5 FTTs, and to increase the data efficiency, the  $\nu_{t,w}$  is updated every 100 time steps.

The model policy is learned based on a neural network with two hidden layers of 128 units each with a Softsign activation function. The parameters of the neural network are identical to those used in [8]. The model training is advanced for 1.6 million policy gradient steps.



**Fig. 1** Contours of the mean velocity in  $x$  direction and the streamlines at  $Re_H = 10595$ : RLWM (top left); RLWM, coarse mesh (top right); EQWM (bottom left); WRLES (bottom right) [37].

## B. Validation

### 1. Testing for flow over periodic hills at $Re_H = 10595$

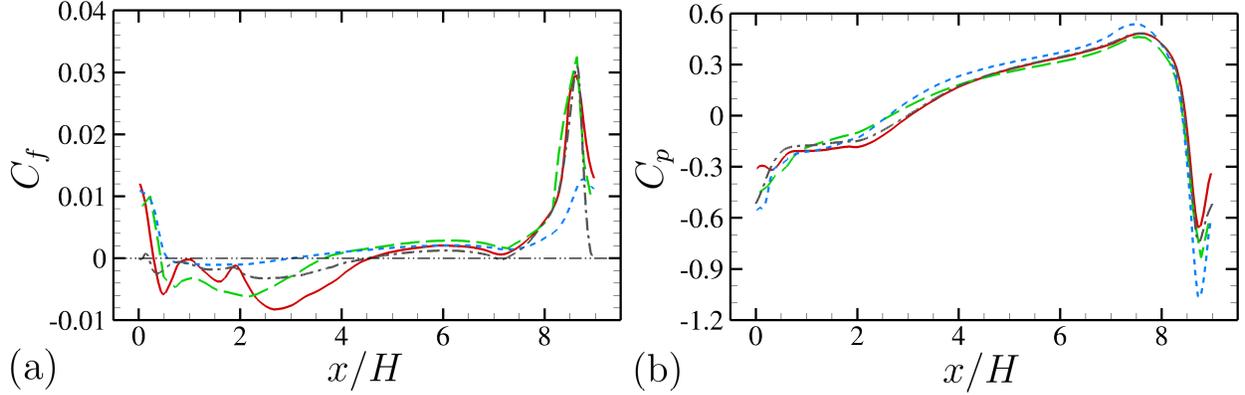
To evaluate the performance of the trained RLWM, two simulations for the periodic-hill channel flow at  $Re_H = 10595$  are carried out using meshes with different resolutions. The details of the simulation cases are listed in Table 1. The number of agents on the bottom wall is consistent with the number of mesh cells on the wall. The wall-normal matching location of the agents was chosen to be within the first off-wall cell. The wall eddy viscosity  $\nu_{t,w}$  is updated based on the model action at every time step. All simulations are run for about 50 FFTs after transients. The flow statistics of all simulations are averaged over spanwise direction and time. For comparison, a WMLES using the EQWM that applies the traditional wall-shear stress boundary condition is conducted with the baseline mesh, and results from two high-fidelity reference databases for this flow [36, 37] are included.

Figure 1 shows the contours of the mean velocity in  $x$  direction and the mean-flow streamlines. The flow separates on the leeward side of the hill due to a strong adverse pressure gradient (APG), and a shear layer is generated near the top of the hill. The flow reattaches in the middle section of the channel, and as the flow approaches the windward side of the downstream hill, it subjects to a strong favorable pressure gradient (FPG) and accelerates rapidly. The simulations with the RLWM successfully capture the separation bubble on the leeward side of the hill and yield more accurate results than the EQWM (see Table 1 for quantitative comparison).

The predictions of the mean skin friction  $C_f$  and the mean pressure coefficient  $C_p$  are shown in Fig. 2. The mean skin friction is defined as  $C_f = \langle \tau_w \rangle / (0.5\rho U_b^2)$ , where  $\langle \cdot \rangle$  denotes averaging in time and spanwise direction, and the positive direction of  $\tau_w$  points toward the opposite direction of bulk flow. The mean pressure coefficient is defined as  $C_p = (\langle p_w \rangle - \langle p_{\text{ref}} \rangle) / (0.5\rho U_b^2)$ , where the pressure at  $x/H = 0$  on the top wall is chosen as reference pressure  $p_{\text{ref}}$  [36]. Regarding  $C_f$ , the results from the RLWM simulations are in reasonable agreement with the DNS data, with large deviations found only near the top of the hill on the leeward side where the skin friction rapidly decreases from its maximum value to a negative value. However, the results are better than the EQWM simulation which largely under-predicts the skin friction on the windward side of the hill. Furthermore, the mean locations of the separation and reattachment points (listed in Table 1) are better predicted by the RLWM, consistent with the streamline shown in Fig. 1. Additionally, it should be noted that the separation bubble size shrinks as the mesh resolution decreases from the baseline; more details about the velocity field will be discussed later in this section.

All simulations capture the qualitative trend of the mean  $C_p$  on the bottom wall including the APG and FPG regimes, but large deviations among the simulation cases are visible near the top of the hill ( $x/H \geq 8.5$  or  $x/H \leq 0.5$ ) where the pressure sees a sudden change from strong FPG to strong APG and the flow separation emerges. Overall, the RLWM provides more accurate predictions of  $C_f$  and  $C_p$  than the EQWM.

Quantitative comparison of mean velocity and Reynolds stress components at five streamwise locations ( $x/H = 0.05, 2, 4, 6$  and  $8$ ) are shown in Fig. 3. The results of the mean velocity profiles and Reynolds stress profiles from the RLWM simulations agree reasonably well with each other and with the reference DNS data. The discrepancies are



**Fig. 2** (a) Mean skin friction coefficient and (b) mean pressure coefficient along the bottom wall at  $Re_H = 10595$ . Lines indicate —, RLWM; - - -, RLWM, coarse mesh; - - -, EQWM; ····, DNS [36]; - - -,  $C_f = 0$ .

visible for the locations on the hill where the pressure gradient is strong. It should be mentioned that the prediction of the velocity field not only depends on the wall boundary conditions but also on the SGS model, and the varying performance of the SGS model in different meshes may contribute to the inconsistency of those two RLWM simulations. On the other hand, prediction from the EQWM simulation is less ideal, particularly for the streamwise Reynolds stress  $\langle u'_x u'_x \rangle / U_b^2$ .

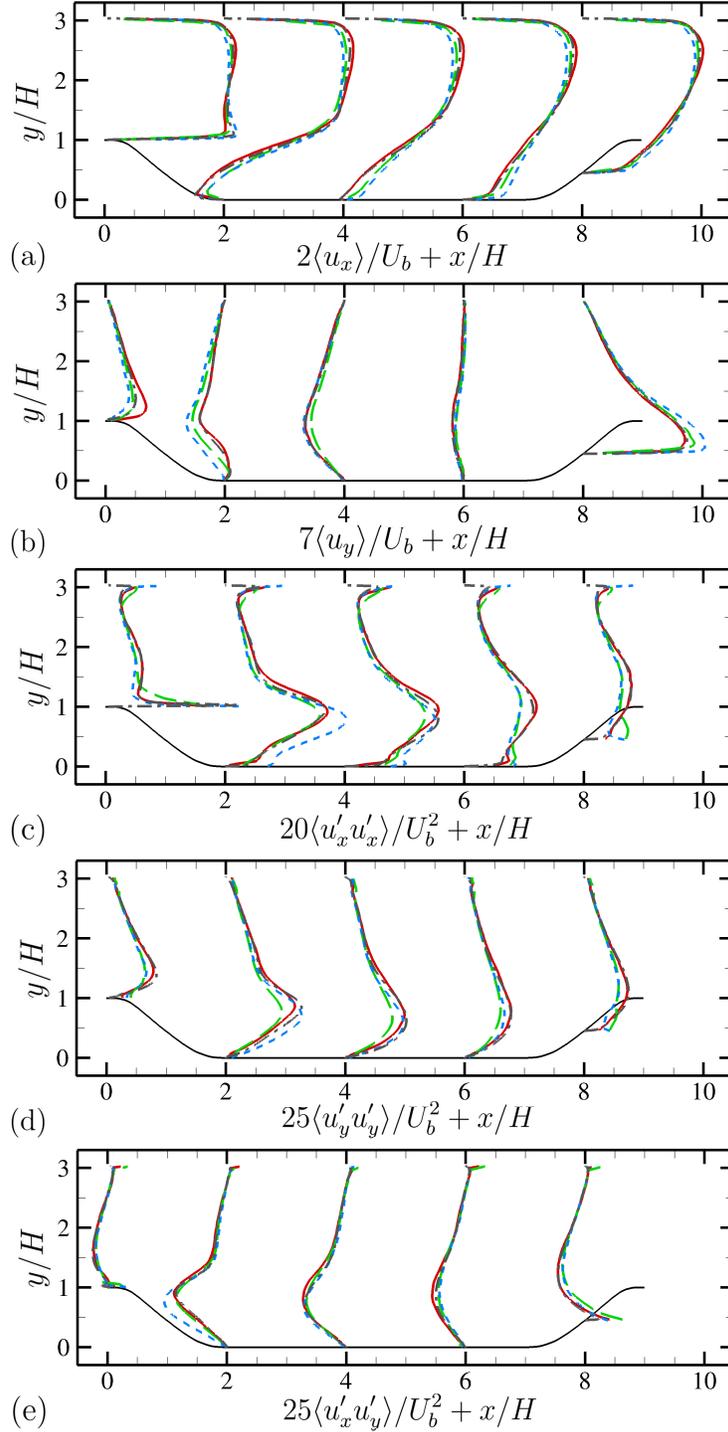
To better understand the mechanism of the trained model, we examine the state-action maps, which are the probability density functions (PDFs) of the likelihood that the model takes a particular action conditioned on the occurrence of positive rewards. Figure 4 shows the maps based on instantaneous states and actions at three streamwise positions  $x/H = 0.1, 2$  and  $8.5$  which are located near the top of the hill on its leeward side, within the separation bubble, and the windward side of the hill, respectively. Overall, the action contour lines for increasing and decreasing  $v_{t,w}$  are well separated, which illustrates the model is able to distinguish flow states and provide appropriate actions.

## 2. Testing for flow over periodic hills at higher Reynolds numbers

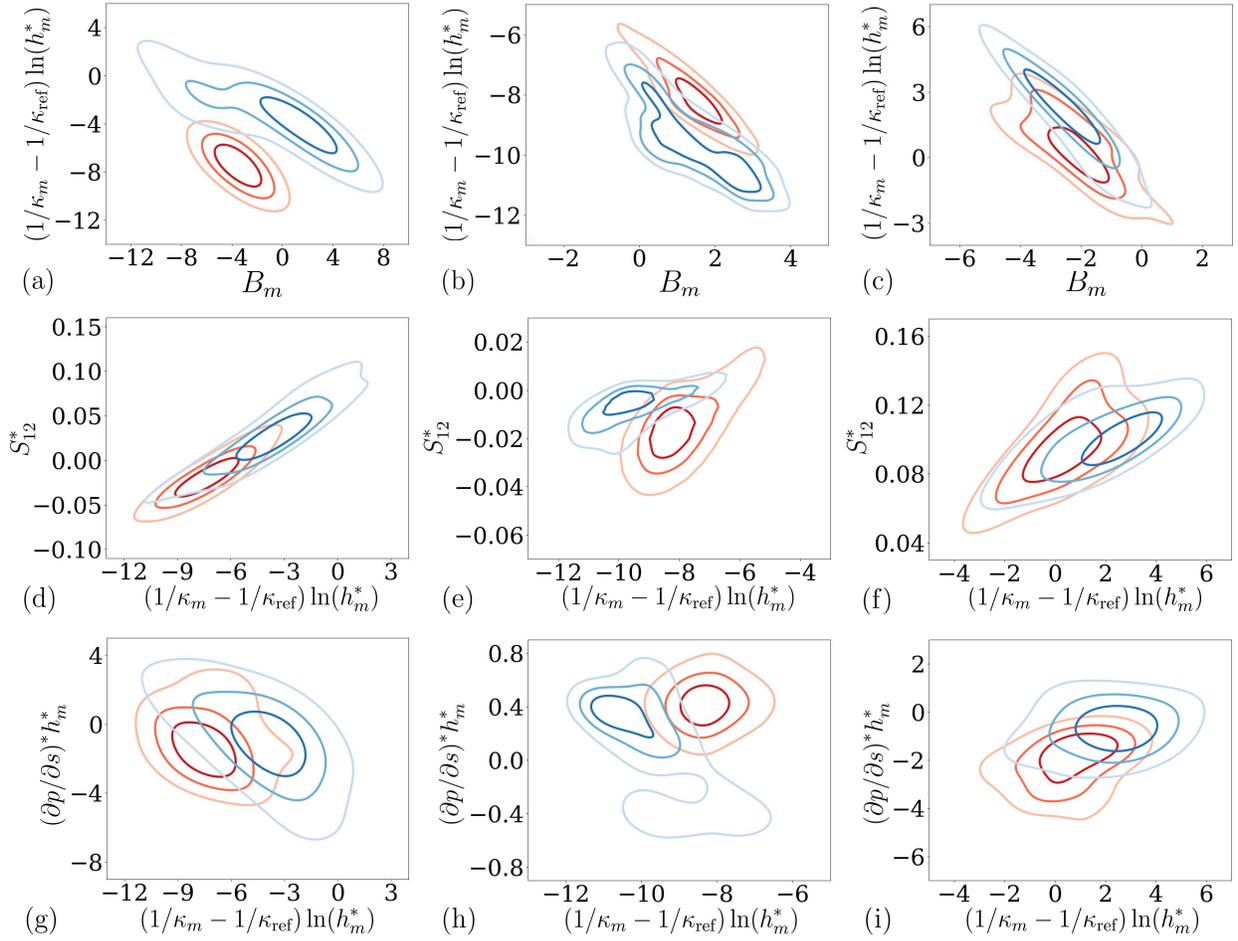
In this section, the RLWM is applied to a WMLES of periodic-hill channel flow at  $Re_H = 19000$  and  $37000$ . The simulations are conducted by using the baseline mesh ( $128 \times 64 \times 64$ ) and the coarse mesh ( $64 \times 32 \times 32$ ), and the implementation of the RLWM is similar to the simulations at  $Re_H = 10595$ . All simulations are run for about 50 FTTs after transients. The results from the EQWM with baseline mesh and the WRLES [37] are included for comparison.

The mean skin friction coefficients along the bottom wall at  $Re_H = 19000$  and  $37000$  are shown in Figs. 5(a,b). The distributions of  $C_f$  at higher Reynolds numbers have a similar shape as the one shown in Fig. 2(a). As the Reynolds number increases, the maximum value of the  $C_f$  on the windward side of the hill decreases. The results from the two RLWM simulations are in reasonable agreement with each other and with the WRLES results. Regarding the separation point, the predicted locations from the RLWM simulations are further downstream than that of the WRLES. The predicted reattachment location from the RLWM simulations is further upstream. Compared to the EQWM simulation, where the separation bubble size is largely under-predicted and the maximum value of  $C_f$  on the windward side of the hill is much smaller than the value of WRLES, the RLWM is more accurate.

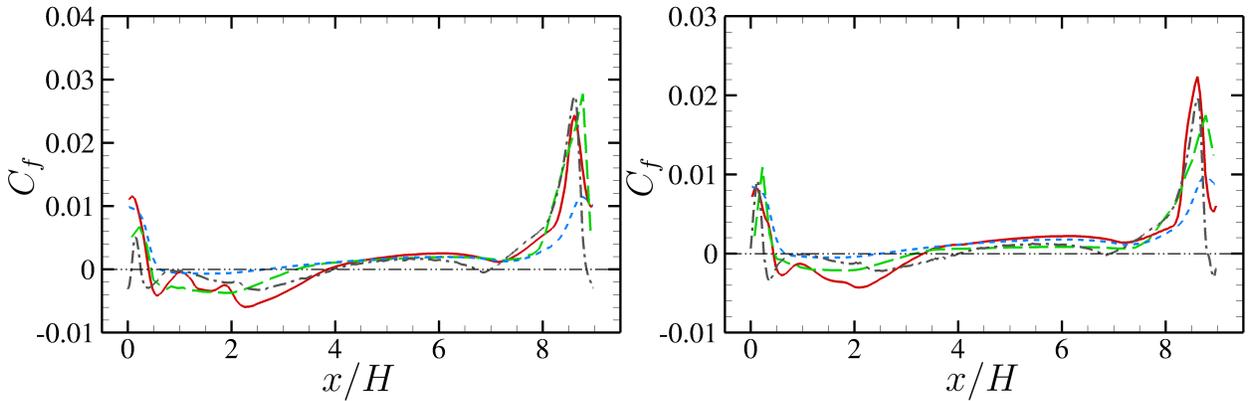
Figure 6 shows the profiles of streamwise components of mean velocity and Reynolds stress at five streamwise stations ( $x/H = 0.05, 2, 4, 6$  and  $8$ ) for  $Re_H = 19000$  and  $37000$ . The deviations among the profiles from the RLWM simulations and the WRLES profiles grow as the Reynolds number increases, particularly for the coarse-mesh RLWM simulation, which implies the degradation of performance in the RLWM. Moreover, as the Reynolds number increases to  $Re_H = 37000$ , the EQWM simulation shows comparable results to the RLWM simulation even though the prediction of  $C_f$  is inaccurate, as shown in Fig. 5(b). It should be noted that the impact of the SGS model on the velocity field away from the wall is greater than the wall model, and this could be an indication that a better SGS model is necessary.



**Fig. 3** Mean velocity and Reynolds stress components at  $Re_H = 10595$ : (a) streamwise velocity; (b) vertical velocity; (c) streamwise Reynolds stress; (d) vertical Reynolds stress; (e) Reynolds shear stress. Lines indicate —, RLWM; - - -, RLWM, coarse mesh; ···, EQWM; - · - ·, DNS [36].

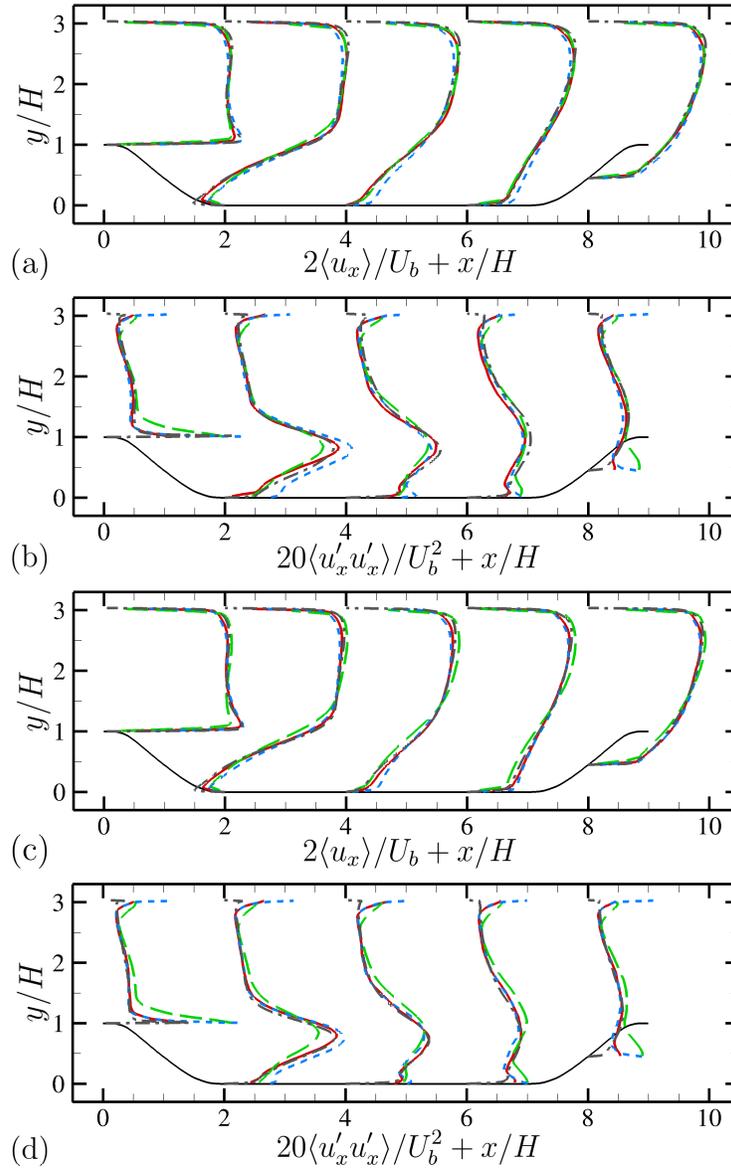


**Fig. 4** PDFs of states for the RLWM conditioned to events with  $r > 0.1$  and  $a < 0.9995$  (blue) or  $a > 1.0005$  (red) at (a,d,g)  $x/H = 0.1$ , (b,e,h)  $x/H = 2$  and (c,f,i)  $x/H = 8.3$  at  $Re_H = 10595$ : PDFs of states (a–c)  $B_m$  and  $(1/\kappa_m - 1/\kappa_{ref}) \ln(h_m^*)$ , (d–f)  $(1/\kappa_m - 1/\kappa_{ref}) \ln(h_m^*)$  and  $S_{12}^*$  and (g–i)  $(1/\kappa_m - 1/\kappa_{ref}) \ln(h_m^*)$  and  $(\partial p/\partial s)^* h_m^*$ . Contour levels are 25, 50, 75% of the maximum value.

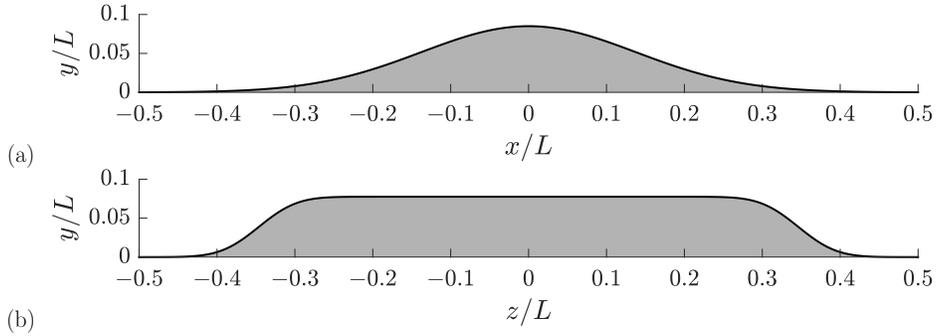


**Fig. 5** Mean skin friction coefficient along the bottom wall at (a)  $Re_H = 19000$  and (b)  $Re_H = 37000$ . Lines indicate —, RLWM; - - -, RLWM, coarse mesh; - - -, EQWM; ···, WRLES [37]; - - -,  $C_f = 0$ .





**Fig. 6** Mean velocity and Reynolds stress components at (a,b)  $Re_H = 19000$  and (c,d)  $Re_H = 37000$ : (a,c) mean streamwise velocity; (b,d) streamwise Reynolds stress. Lines indicate —, RLWM; - - -, RLWM, coarse mesh; ····, EQWM; - · - ·, WRLES [37].



**Fig. 7** Cross-sections of the bump geometry showing (a) a slice along the centerline ( $z/L = 0$ ) and (b) a slice along the span ( $x/L = 0$ ).

### III. Simulations for Flow over Boeing Gaussian Bump

The geometry of the Boeing Gaussian bump is given by the analytic function

$$y = f(x, z) = \frac{h}{2} e^{-(x/x_0)^2} \left\{ 1 + \operatorname{erf} \left[ \left( \frac{L}{2} - 2z_0 - |z| \right) / z_0 \right] \right\}, \quad (1)$$

where  $x$ ,  $y$  and  $z$  are the streamwise (free-stream-aligned), vertical and spanwise coordinates, respectively, and  $f$  is the surface representing the geometry of the wall-mounted bump. The length scale  $L$ , referred to as the bump width, is used to express the other scales of the bump, where  $h = 0.085L$  is the maximum height of the bump,  $x_0 = 0.195L$ , and  $z_0 = 0.06L$ . The axial and spanwise cross-sections are shown in Fig. 7. Additionally, the length scale  $L$  is used to define the Reynolds number  $Re_L = U_\infty L / \nu_\infty$ , where  $(\cdot)|_\infty$  denotes quantities at free-stream conditions. This geometry has been extensively studied experimentally by [17–21]. The experimental geometry includes side and top walls, as the Gaussian bump is wall-mounted on a splitter plate inside of a wind tunnel. In this work, the geometry with side and top walls is considered. Particularly, the case of  $Re_L = 2 \times 10^6$  is simulated, and the WMLES results are compared with the experimental measurements.

#### A. Computational details

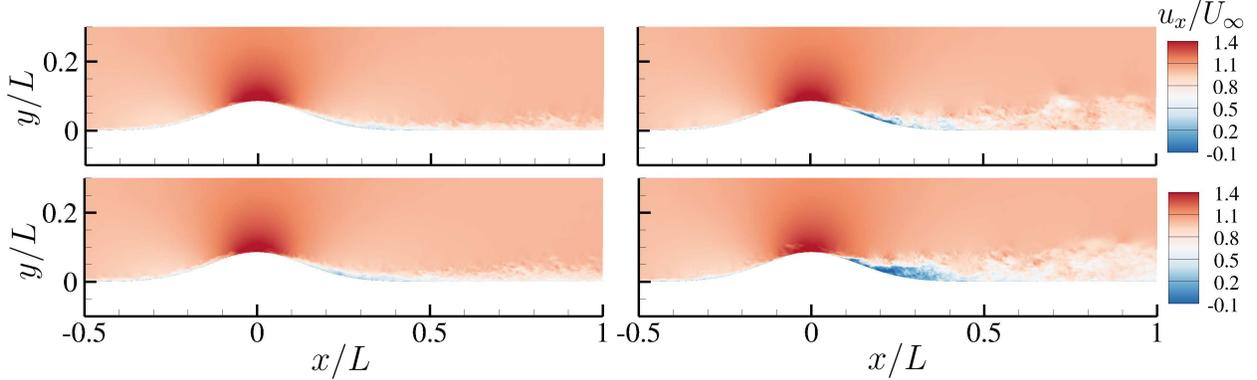
The present simulations employ a rectangular computational domain with the dimension of  $L_x \times L_y \times L_z = 2.5L \times 0.5L \times 0.5L$ , which has the same blockage ratio as in the wind tunnel experiment [18–21]. The origin of the coordinate system in the domain locates at the base of the bump peak as shown in Fig. 8. Because symmetry exists with respect to the center plane at  $z/L = 0$  in the geometry, the simulation domain only covers half of the entire bump span with a symmetry boundary condition applied at  $z/L = 0$ . Moreover, the simulations have a plug flow inlet at  $x/L = -1$ , with the side at  $z/L = 0.5$  and top boundary conditions at  $y/L = 0.5$  treated as inviscid walls to approximate the wind tunnel. The outlet is placed at  $x/L = 1.5$  with a convective outflow boundary condition. Note that preliminary investigations suggest insensitivity to the inlet and tunnel wall boundary conditions [26, 27].

The simulations are conducted using the aforementioned unstructured-mesh, high-fidelity incompressible finite-volume flow solver that is coupled with the RL toolbox, smarties. In order to investigate the influence of SGS models on the simulation results, the DSM, which is used in the aforementioned wall model training, and the anisotropic minimum-dissipation (AMD) model [42] are both used in separate simulations. In addition to the developed RL wall model, we also used the EQWM in the simulations for comparison. The EQWM is based on the Neumann/no-penetration velocity boundary condition that is traditionally employed with wall-stress models [43], and the center of the second off-wall cell is selected as the matching location. Additionally, a CFL number of 2 is used in all simulations.

To study the effect of mesh resolution on the simulation results, three computational meshes with increasing resolutions in each direction are considered. These meshes consist of structured-mesh blocks covering the entire bump surface and the flat wall surfaces in both the upstream and downstream, and unstructured-mesh blocks elsewhere. The wall-normal dimension of the structured-mesh blocks is equal to  $0.2L$ , which is larger than the thickness of the TBL downstream of the bump. The parameters of the computational meshes are provided in Table 2. The outer unstructured mesh has the resolution of  $\Delta/L = 0.03$  and the control volumes are refined gradually towards the bottom wall. It should

**Table 2** Mesh parameters of the simulations for the flow over Boeing Gaussian bump at  $Re_L = 2 \times 10^6$ .

Mesh	$N_{CV}$	$\min \Delta_x/L$	$\min \Delta_y/L$	$\min \Delta_z/L$
Coarse	6 mil.	$3.8 \times 10^{-3}$	$2.6 \times 10^{-3}$	$2.6 \times 10^{-3}$
Medium	37 mil.	$1.9 \times 10^{-3}$	$1.3 \times 10^{-3}$	$1.3 \times 10^{-3}$
Fine	100 mil.	$1.3 \times 10^{-3}$	$9 \times 10^{-4}$	$9 \times 10^{-4}$



**Fig. 8** Instantaneous streamwise velocity  $u_x/U_\infty$  in an  $x-y$  plane at  $z/L = 0$  from medium-mesh WMLES with RLWM and DSM (top left), RLWM and AMD model (top right), EQWM and DSM (bottom left), and EQWM and AMD model (bottom right).

be noted that the thickness of the TBL at  $x/L = -0.4$  is approximately resolved by 3 cells in the coarse mesh, 7 cells in the medium mesh, and 11 cells in the fine mesh, respectively.

In the simulations using the RLWM, the number of agents above the bottom surface is consistent with the number of wall cells in the mesh. The wall-normal matching location of the agents was chosen to be within the first off-wall cell. The wall eddy viscosity  $\nu_{t,w}$  is updated at each time step based on the model action. To eliminate numerical artifacts, all simulations are run for 1.5 FTTs at first to pass the transient process. After that, these simulations are run for another 1.5 FTTs to collect flow statistics.

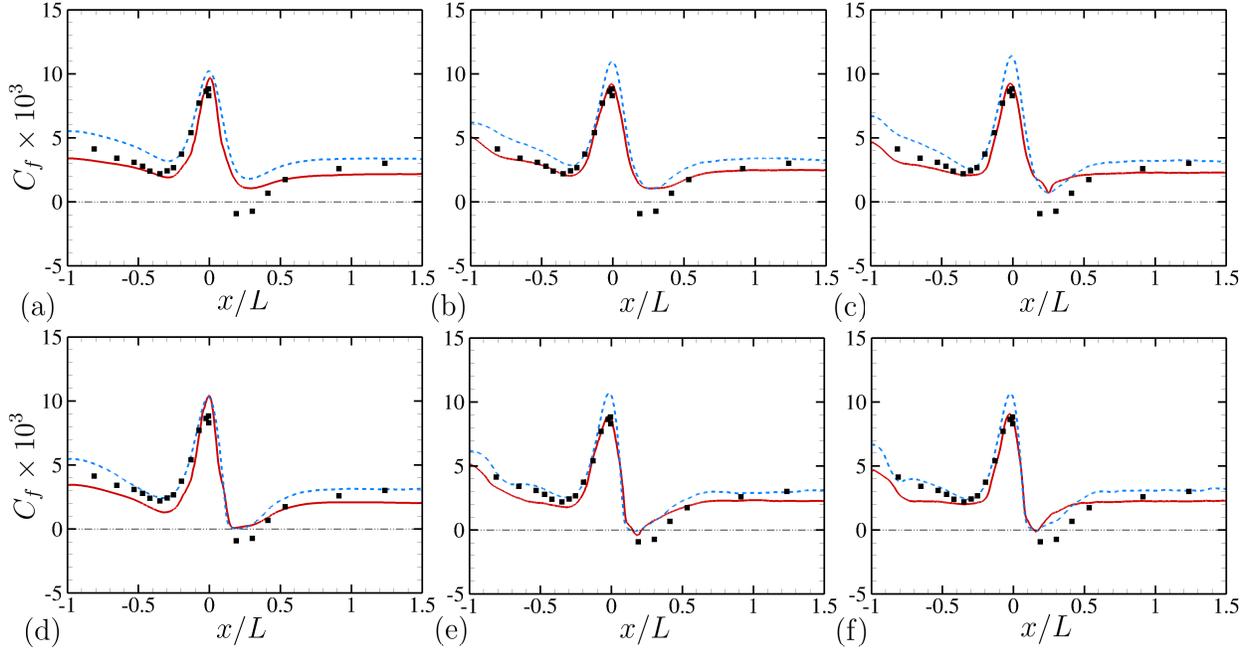
## B. Results and Discussion

Figure 8 shows contours of the instantaneous streamwise velocity  $u_x/U_\infty$  in an  $x-y$  plane at  $z/L = 0$  obtained from the medium-mesh simulations with different SGS models and wall models. The flow gradually accelerates on the windward side of the bump. The flow velocity reaches its maximum value at the bump peak. Downstream of the peak, the flow decelerates over the leeward side of the bump, and the boundary layer thickens rapidly. Note that the flow is attached over the bump surface in the simulations with DSM, but the flow is separated on the leeward side of the bump in the simulations with the AMD model, which is more consistent with the experimental observations [17–21]. The results explicitly demonstrate the strong effect of SGS models on the WMLES. In terms of the effect from wall models, it is comparatively small; however, from the simulations with the AMD model, it can be found that the separation is stronger when the EQWM is used, which shows the influence of wall models on the flow field is not negligible. Detailed quantitative comparison of mean velocity from different simulations will be discussed later.

Additional quantities of interest for the bump-flow simulations include the mean pressure coefficient,  $C_p$ , and the mean skin friction coefficient,  $C_f$ , which are defined as

$$C_p = \frac{\overline{p_w} - p_\infty}{\frac{1}{2} \rho_\infty U_\infty^2} \quad \text{and} \quad C_f = \frac{\overline{\tau_w}}{\frac{1}{2} \rho_\infty U_\infty^2}, \quad (2)$$

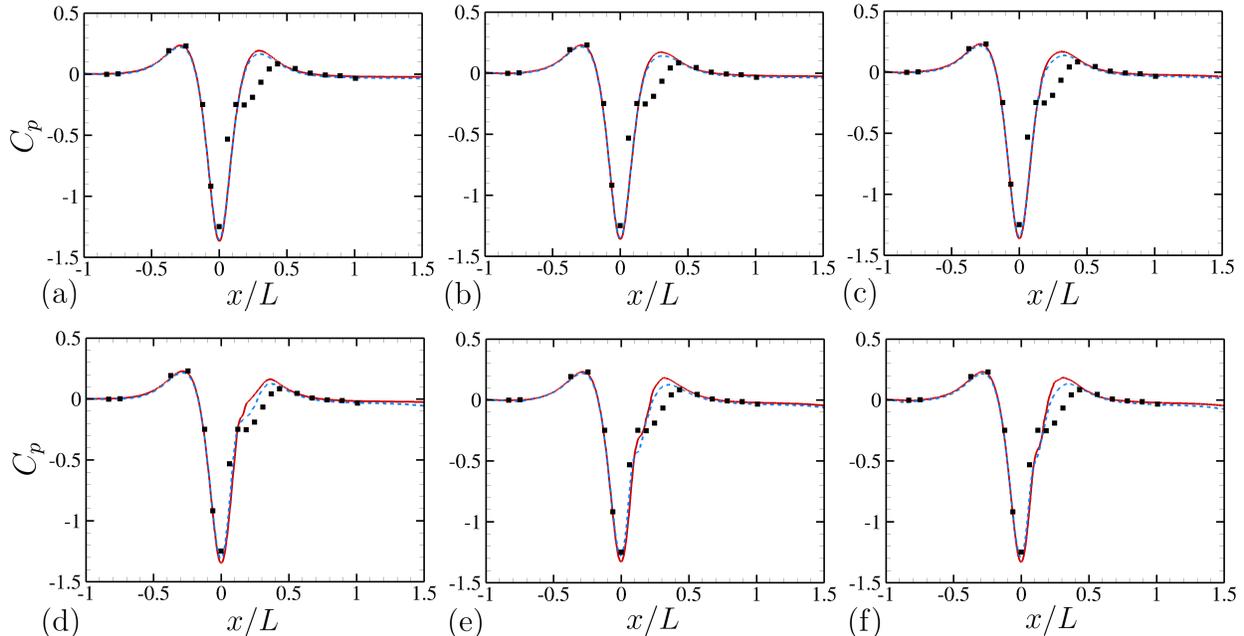
where  $\overline{(\cdot)}$  denotes the averaging in time, and the mean pressure at the inlet near the top boundary is chosen as the reference free-stream pressure  $p_\infty$ . Figure 9 shows the distributions of the mean skin-friction coefficient  $C_f$  over the



**Fig. 9** The distributions of mean skin friction coefficient along the bump surface at  $z/L = 0$  from the WMLES using (a,b,c) DSM and (d,e,f) AMD model with (a,d) coarse mesh, (b,e) medium mesh and (c,f) fine mesh. Lines indicate —, RLWM; ---, EQWM; -·-,  $C_f = 0$ ; ■, experimental measurements at  $Re_L = 2 \times 10^6$  [19].

bump surface. For comparison, the data from the experiment using the same geometry at the same Reynolds number from Gray et al. [19] is also shown. In the experiments, a large separation bubble appears downstream of the bump peak, which is not accurately captured in all simulations. A small separation bubble can be observed on the leeward side of the bump in the simulations with the AMD model. On the contrary, using the DSM, there is no separation in the flow field. The huge contrast in the results of  $C_f$  is consistent with that observed in the instantaneous velocity contours (see Fig. 8). Moreover, upstream of the bump, the  $C_f$  results from the simulations with different SGS models also exhibit noticeable differences, although the difference is reducing as the mesh resolutions increase. Considering the performance of different wall models, the simulations using the EQWM typically overpredict  $C_f$  near the bump peak. However, the developed RLWM provides more accurate predictions for this region, and as the mesh resolutions increase, the  $C_f$  converges to a value that agrees well with the experimental data. Upstream of the bump, the performance of the current RLWM is less optimal, which is expected as the model training does not encompass laminar or transitional flows. On the leeward side of the bump, the predictions of  $C_f$  from both wall models are comparable, but further downstream the RLWM provides lower skin friction than the EQWM.

Figure 10 shows the distribution of  $C_p$  along the bump surface at  $z/L = 0$ . Results from all simulations with different computational meshes are included along with the experimental measurements of Williams et al. [17] at slightly lower Reynolds number ( $Re_L = 1.87 \times 10^6$ ). The  $C_p$  distributions illustrate strong FPG immediately upstream of the bump peak. Downstream of the peak, the flow is first subjected to a very strong APG then followed by a mild FPG. Comparison of simulation results with experimental data reveals significant differences on the leeward side of the bump, which is expected due to the imprecise capture of the separation bubble in the simulations. This region is also the most sensitive to the SGS models employed. Furthermore, simulations with the DSM exhibit closely matching curves, suggesting that variations in wall modeling and mesh resolution do not significantly affect  $C_p$  predictions in these simulations. On the other hand, increasing the mesh resolution results in noticeable changes in  $C_p$  on the leeward side of the bump, which leads to a gradual decline in prediction accuracy. In terms of wall model performance, the RLWM and EQWM provide similar  $C_p$  predictions, with only minor differences observed on the leeward side of the bump. In addition, the distributions of  $C_p$  along the span at the bump peak are shown in Fig. 11. The results from different simulations agree well among themselves and with the experimental data [17] in the outer-span region ( $x/L > 0.25$ ). The differences among different simulations and the experiment become more pronounced near the center of the bump,



**Fig. 10** The distributions of mean pressure coefficient along the bump surface at  $z/L = 0$  from the WMLES using (a,b,c) DSM and (d,e,f) AMD model with (a,d) coarse mesh, (b,e) medium mesh and (c,f) fine mesh. Lines indicate —, RLWM; ---, EQWM; ■, experimental measurements at  $Re_L = 1.87 \times 10^6$  [17].

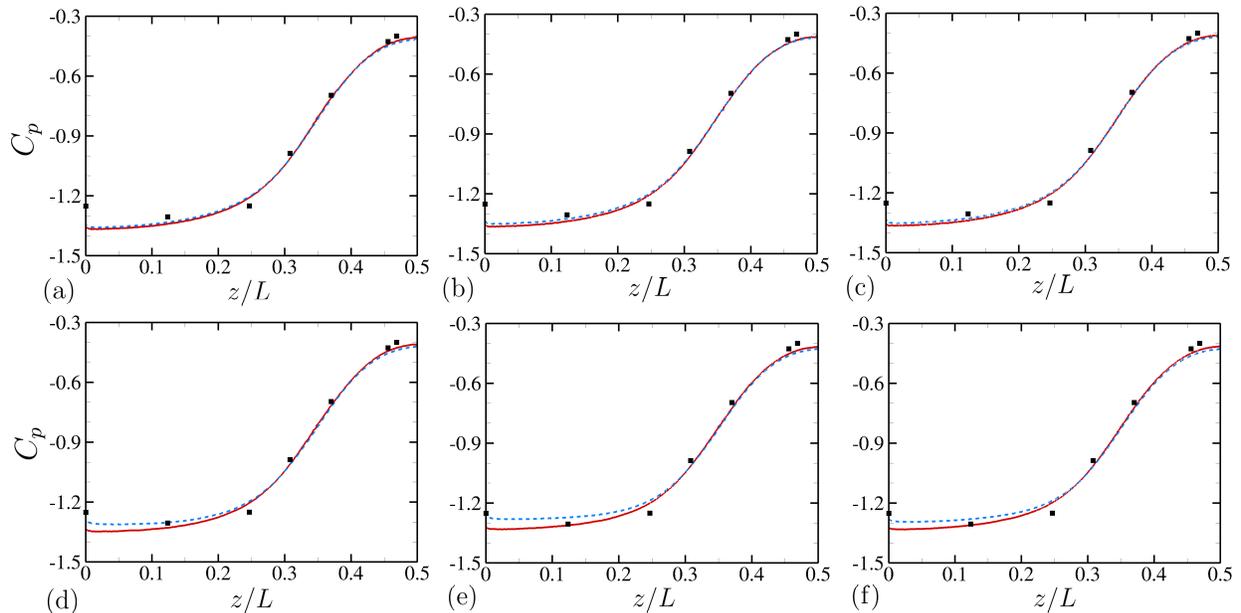
where the separation bubble appears in the downstream region.

In Fig. 12, profiles of the mean streamwise velocity are depicted at four streamwise stations on the leeward side of the bump. These results quantitatively show the flow deceleration and boundary-layer thickening on the bump surface. Based on the comparison in the figure, the simulation results exhibit significant differences from the experiment conducted by Gray et al. [20], with the simulations predicting a thinner TBL downstream of the bump. This finding is consistent with the earlier observation of the instantaneous flow field. The simulations with different SGS models and wall models exhibit similar results near the bump peak, but significant differences emerge further downstream. Specifically, the use of the AMD model leads to the appearance of a separation bubble in the flow field and a significantly thicker TBL. Considering the results from different wall models, their differences are less significant than that due to the use of different SGS models.

While data-driven turbulence models often have limited applicability across different flow configurations [44], the current RLWM, originally trained on the low-Reynolds-number periodic-hill channel flow, shows promise in simulating the flow over the Boeing Gaussian bump, which has a different geometry as well as higher Reynolds number and is likely to exhibit different flow physics. Specifically, the RLWM provides improved predictions for the skin friction near the peak of the bump and performs comparably to the EQWM with respect to the wall pressure and velocity field. Despite the use of different wall models, the simulations exhibit significant sensitivity to SGS models. Notably, the current RLWM was developed using DSM. As such, further research is needed to explore the impact of SGS models on the training of MARL-based wall models and on the applicability of the trained wall models.

#### IV. Summary

In this work, a wall model that can adapt to various pressure-gradient effects is developed for turbulent flow over periodic hills using multi-agent reinforcement learning. The model behaves as a control policy for wall eddy viscosity to predict the correct wall-shear stress. The optimized policy of the wall model is learned through the training process based on LES of low-Reynolds-number flow over periodic hills with cooperating agents using the recovery of the correct wall-shear stress as a reward. The developed wall model is first validated in the LES of the periodic-hill configuration



**Fig. 11** The distributions of mean pressure coefficient along the span at the bump peak ( $x/L = 0$ ) from the WMLES using (a,b,c) DSM and (d,e,f) AMD model with (a,d) coarse mesh, (b,e) medium mesh and (c,f) fine mesh. Lines indicate —, RLWM; - - -, EQWM; ■, experimental measurements at  $Re_L = 1.87 \times 10^6$  [17].

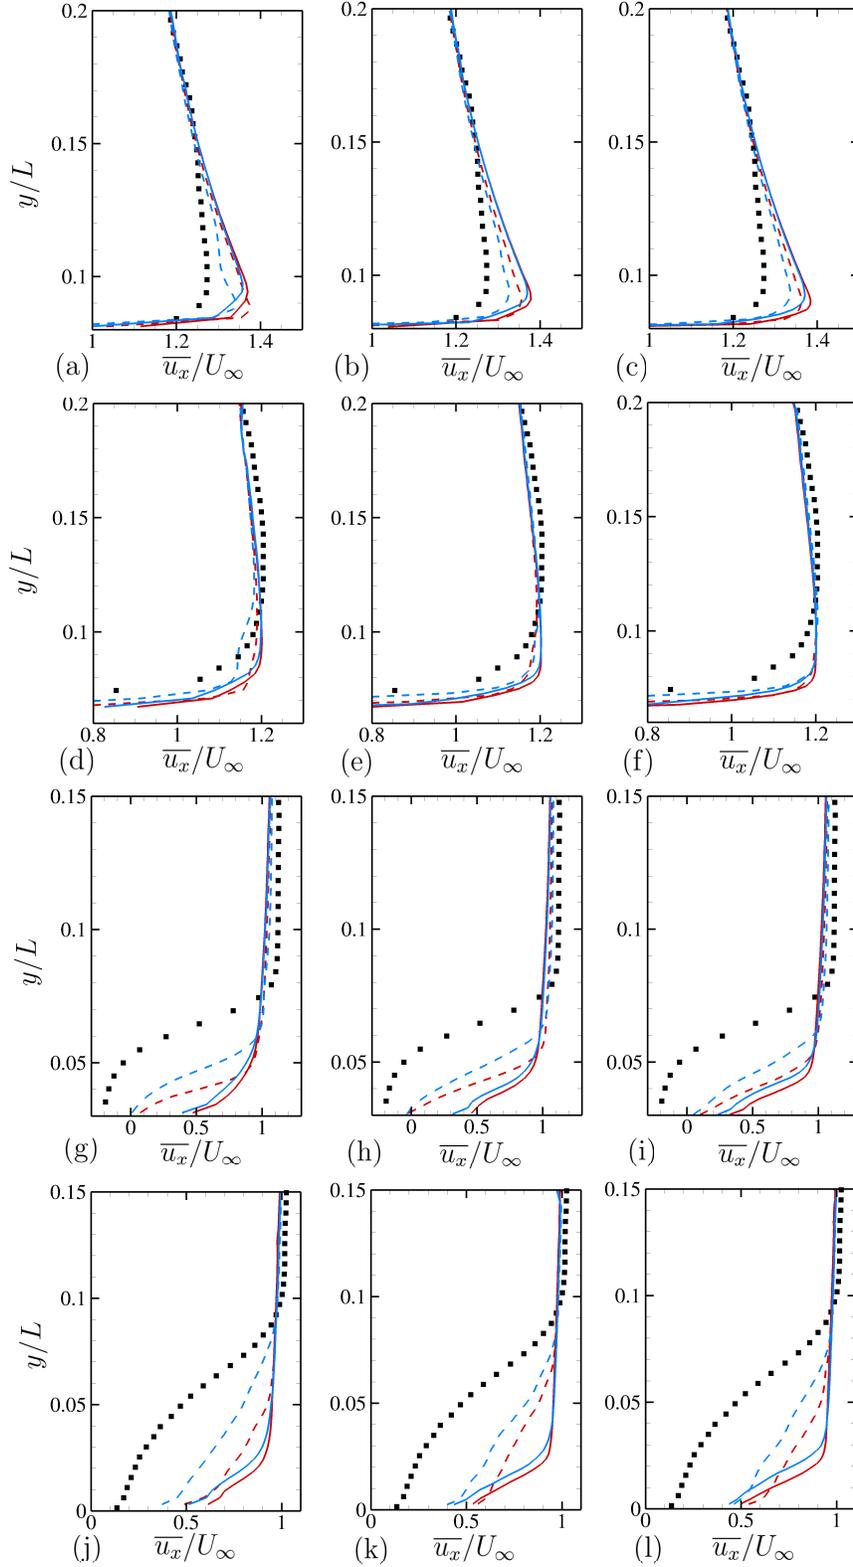
at the same Reynolds number of model training. The wall model provides good predictions of mean wall-shear stress, mean wall pressure, and the mean velocity as well as Reynolds stress in the flow field. The test results also show that the developed model outperforms the EQWM. The performance of the developed model is further evaluated at two higher Reynolds numbers ( $Re_H = 19000$  and  $37000$ ). Good predictions are obtained for the mean wall-shear stress and velocity statistics.

To further investigate the generalizability and robustness of the developed wall model, simulations of flow over the Boeing Gaussian bump are conducted using two different SGS models. The flow geometry is consistent with the experiments [18–21], and the Reynolds number based on the free-stream velocity and the width of the bump is  $2 \times 10^6$ . The results of mean skin friction and pressure on the bump surface and the velocity statistics of the flow field are compared to those from the EQWM simulations and published experimental data sets. It shows that the developed wall model successfully captures the acceleration and deceleration of the TBL on the bump surface. The wall model provides better predictions of skin friction near the bump peak and exhibits comparable performance to the EQWM with respect to the wall pressure and velocity field. In addition, it is found that the SGS model has a significant impact on the simulation results, which is a topic requiring further study.

On the foundation of the test simulations for various flow configurations, it can be concluded that the current wall model, developed based on the training with low-Reynolds-number periodic-hill channel flow, has gained valuable insights into the key physics of complex flows with diverse pressure gradients. However, further improvements can be expected by training the wall model over a wider range of Reynolds numbers and flow geometries using the multi-agent reinforcement learning framework. This would enable the model to capture more complex features of the flow field and enhance its reliability and performance in a broader range of applications.

## Acknowledgments

Authors D. Z. and H. J. B. acknowledge support from NSF under Grant No.2152705 and XSEDE computing resources from PSC under project PHY210119. M. P. W. acknowledges funding from NASA Transformational Tools and Technologies grant #80NSSC20M0201. K. P. G. is supported by the Exascale Computing Project (grant 17-SC-20-SC), a collaborative effort of two US Department of Energy organizations (Office of Science and the National Nuclear Security Administration) responsible for the planning and preparation of a capable exascale ecosystem, including



**Fig. 12** Profiles of mean velocity  $\bar{u}_x/U_\infty$  on an  $x - y$  plane of  $z/L = 0$  at (a,b,c)  $x/D = 0.05$ , (d,e,f)  $x/D = 0.1$ , (g,h,i)  $x/D = 0.2$ , and (j,k,l)  $x/D = 0.4$  from WMLES with coarse mesh (left), medium mesh (center), and fine mesh (right). Solid lines indicate simulations with the DSM, and dashed lines indicate simulations with the AMD model. Red lines denote the simulations with the RLWM, and blue lines denote the simulations with the EQWM. Symbol (■) represents experimental measurements at  $Re_L = 2 \times 10^6$  [20].

software, applications, hardware, advanced system engineering, and early testbed platforms, in support of the nation's exascale computing imperative. This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

## References

- [1] Balaras, E., Benocci, C., and Piomelli, U., "Two-layer approximate boundary conditions for large-eddy simulations," *AIAA J.*, Vol. 34, No. 6, 1996, pp. 1111–1119.
- [2] Wang, M., and Moin, P., "Dynamic wall modeling for large-eddy simulation of complex turbulent flows," *Phys. Fluids*, Vol. 14, No. 7, 2002, pp. 2043–2051.
- [3] Chung, D., and Pullin, D. I., "Large-eddy simulation and wall modelling of turbulent channel flow," *J. Fluid Mech.*, Vol. 631, 2009, pp. 281–309.
- [4] Park, G. I., and Moin, P., "An improved dynamic non-equilibrium wall-model for large eddy simulation," *Phys. Fluids*, Vol. 26, No. 1, 2014, pp. 37–48.
- [5] Bose, S. T., and Moin, P., "A dynamic slip boundary condition for wall-modeled large-eddy simulation," *Phys. Fluids*, Vol. 26, No. 1, 2014, p. 015104.
- [6] Bae, H. J., Lozano-Durán, A., Bose, S. T., and Moin, P., "Dynamic slip wall model for large-eddy simulation," *J. Fluid Mech.*, Vol. 859, 2019, pp. 400–432.
- [7] Zhou, Z., He, G., and Yang, X., "Wall model based on neural networks for LES of turbulent flows over periodic hills," *Phys. Rev. Fluids*, Vol. 6, No. 5, 2021, p. 054610.
- [8] Bae, H. J., and Koumoutsakos, P., "Scientific multi-agent reinforcement learning for wall-models of turbulent flows," *Nat. Commun.*, Vol. 13, 2022, pp. 1–9.
- [9] Bertsekas, D. P., *Reinforcement Learning and Optimal Control*, Athena Scientific, Nashua, NH, USA, 2019.
- [10] Gazzola, M., Hejazialhosseini, B., and Koumoutsakos, P., "Reinforcement learning and wavelet adapted vortex methods for simulations of self-propelled swimmers," *SIAM J. Sci. Comput.*, Vol. 36, No. 3, 2014, pp. B622–B639.
- [11] Novati, G., Verma, S., Alexeev, D., Rossinelli, D., Van Rees, W. M., and Koumoutsakos, P., "Synchronisation through learning for two self-propelled swimmers," *Bioinspir. Biomim.*, Vol. 12, 2017, p. 036001.
- [12] Novati, G., de Laroussilhe, H. L., and Koumoutsakos, P., "Automating turbulence modelling by multi-agent reinforcement learning," *Nat. Mach. Intell.*, Vol. 3, 2021, pp. 87–96.
- [13] Cabot, W., and Moin, P., "Approximate wall boundary conditions in the large-eddy simulation of high Reynolds number flow," *Flow, Turbulence and Combustion*, Vol. 63, 2000, pp. 269–291.
- [14] Kawai, S., and Larsson, J., "Wall-modeling in large eddy simulation: Length scales, grid resolution, and accuracy," *Phys. Fluids*, Vol. 24, No. 1, 2012, p. 015105.
- [15] Vadrot, A., Yang, X. I. A., Bae, H. J., and Abkar, M., "Log-law recovery through reinforcement-learning wall model for large-eddy simulation," *Phys. Fluids*, accepted, 2023.
- [16] Slotnick, J. P., "Integrated CFD validation experiments for prediction of turbulent separated flows for subsonic transport aircraft," *NATO Science and Technology Organization, Meeting Proceedings RDP, STO-MP-AVT-307*, 2019.
- [17] Williams, O., Samuell, M., Sarwas, E. S., Robbins, M., and Ferrante, A., "Experimental study of a CFD validation test case for turbulent separated flows," *AIAA Scitech 2020 Forum*, 2020, p. 0092.
- [18] Gray, P. D., Gluzman, I., Thomas, F., Corke, T., Lakebrink, M., and Mejia, K., "A new validation experiment for smooth-body separation," *AIAA Aviation 2021 Forum*, 2021, p. 2810.



- [19] Gray, P. D., Gluzman, I., Thomas, F. O., and Corke, T. C., “Experimental characterization of smooth body flow separation over wall-mounted Gaussian bump,” *AIAA Scitech 2022 Forum*, 2022, p. 1209.
- [20] Gray, P. D., Gluzman, I., Thomas, F. O., Corke, T. C., Lakebrink, M. T., and Mejia, K., “Benchmark characterization of separated flow over smooth Gaussian bump,” *AIAA Aviation 2022 Forum*, 2022, p. 3342.
- [21] Gluzman, I., Gray, P., Mejia, K., Corke, T. C., and Thomas, F. O., “A simplified photogrammetry procedure in oil-film interferometry for accurate skin-friction measurement over arbitrary geometries,” *Exp. Fluids*, Vol. 63, No. 7, 2022, pp. 1–14.
- [22] Balin, R., and Jansen, K. E., “Direct numerical simulation of a turbulent boundary layer over a bump with strong pressure gradients,” *J. Fluid Mech.*, Vol. 918, 2021.
- [23] Uzun, A., and Malik, M. R., “High-fidelity simulation of turbulent flow past Gaussian bump,” *AIAA J.*, Vol. 60, No. 4, 2022, pp. 2130–2149.
- [24] Wright, J. R., Balin, R., Jansen, K. E., and Evans, J. A., “Unstructured LES\_DNS of a turbulent boundary layer over a Gaussian bump,” *AIAA Scitech 2021 Forum*, 2021, p. 1746.
- [25] Iyer, P. S., and Malik, M. R., “Wall-modeled LES of flow over a Gaussian bump,” *AIAA Scitech 2021 Forum*, 2021, p. 1438.
- [26] Whitmore, M. P., Griffin, K. P., Bose, S. T., and Moin, P., “Large-eddy simulation of a Gaussian bump with slip-wall boundary conditions,” *Center for Turbulence Research Annual Research Briefs*, 2021, pp. 45–58.
- [27] Agrawal, R., Whitmore, M. P., Griffin, K. P., Bose, S. T., and Moin, P., “Non-Boussinesq subgrid-scale model with dynamic tensorial coefficients,” *arXiv preprint arXiv:2202.05502*, 2022.
- [28] Balin, R., Jansen, K. E., and Spalart, P. R., “Wall-modeled LES of flow over a Gaussian bump with strong pressure gradients and separation,” *AIAA Aviation 2020 Forum*, 2020, p. 3012.
- [29] You, D., Ham, F., and Moin, P., “Discrete conservation principles in large-eddy simulation with application to separation control over an airfoil,” *Phys. Fluids*, Vol. 20, 2008, p. 101515.
- [30] Van der Vorst, H. A., “Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems,” *SIAM J. Sci. Comput.*, Vol. 13, 1992, pp. 631–644.
- [31] Germano, M., Piomelli, U., Moin, P., and Cabot, W. H., “A dynamic subgrid-scale eddy viscosity model,” *Phys. Fluids*, Vol. 3, 1991, pp. 1760–1765.
- [32] Lilly, D. K., “A proposed modification of the Germano subgrid-scale closure method,” *Phys. Fluids*, Vol. 4, 1992, pp. 633–635.
- [33] Mellen, C. P., Fröhlich, J., and Rodi, W., “Large eddy simulation of the flow over periodic hills,” *Proceedings of the 16th IMACS world congress*, Lausanne, Switzerland, 2000, pp. 21–25.
- [34] Fröhlich, J., Mellen, C. P., Rodi, W., Temmerman, L., and Leschziner, M. A., “Highly resolved large-eddy simulation of separated flow in a channel with streamwise periodic constrictions,” *J. Fluid Mech.*, Vol. 526, 2005, pp. 19–66.
- [35] Rapp, C., and Manhart, M., “Flow over periodic hills: an experimental study,” *Exp. Fluids*, Vol. 51, No. 1, 2011, pp. 247–269.
- [36] Krank, B., Kronbichler, M., and Wall, W. A., “Direct numerical simulation of flow over periodic hills up to  $Re_H = 10,595$ ,” *Flow Turbul. Combust.*, Vol. 101, 2018, pp. 521–551.
- [37] Gloerfelt, X., and Cinnella, P., “Large eddy simulation requirements for the flow over periodic hills,” *Flow Turbul. Combust.*, Vol. 103, 2019, pp. 55–91.
- [38] Balakumar, P., Park, G. I., and Pierce, B., “DNS, LES, and wall-modeled LES of separating flow over periodic hills,” *Proceedings of the CTR Summer Program*, 2014, pp. 407–415.
- [39] Novati, G., and Koumoutsakos, P., “Remember and forget for experience replay,” *Proceedings of the 36th International Conference on Machine Learning*, PMLR, 2019, pp. 4851–4860.
- [40] Manhart, M., Peller, N., and Brun, C., “Near-wall scaling for turbulent boundary layers with adverse pressure gradient A priori tests on DNS of channel flow with periodic hill constrictions and DNS of separating boundary layer,” *Comput. Fluid Dyn.*, Vol. 22, 2008, pp. 243–260.
- [41] Zhou, D., Whitmore, M. P., Griffin, K. P., and Bae, H. J., “Multi-agent reinforcement learning for wall modeling in LES of flow over periodic hills,” *Proceedings of the CTR Summer Program*, 2022, pp. 25–34.

- [42] Rozema, W., Bae, H. J., Moin, P., and Verstappen, R., “Minimum-dissipation models for large-eddy simulation,” *Phys. Fluids*, Vol. 27, No. 8, 2015, p. 085107.
- [43] Kawai, S., and Larsson, J., “Wall-modeling in large eddy simulation: Length scales, grid resolution, and accuracy,” *Phys. Fluids*, Vol. 24, 2012, p. 015105.
- [44] Duraisamy, K., Iaccarino, G., and Xiao, H., “Turbulence modeling in the age of data,” *Annual review of fluid mechanics*, Vol. 51, 2019, pp. 357–377.