

## Embedded training of neural-network subgrid-scale turbulence models

Jonathan F. MacArt<sup>1,\*</sup>, Justin Sirignano<sup>2,3,†</sup> and Jonathan B. Freund<sup>4,‡</sup>

<sup>1</sup>*Department of Aerospace and Mechanical Engineering, University of Notre Dame, Notre Dame, Indiana 46556, USA*

<sup>2</sup>*Mathematical Institute, University of Oxford, Oxford OX2 6GG, United Kingdom*

<sup>3</sup>*Department of Industrial & Systems Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA*

<sup>4</sup>*Department of Aerospace Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA*



(Received 3 July 2020; accepted 6 January 2021; published 12 May 2021)

The weights of a deep neural-network model are optimized in conjunction with the governing flow equations to provide a model for subgrid-scale stresses in a temporally developing plane turbulent jet at Reynolds number  $Re_0 = 6000$ . The objective function for training is first based on the instantaneous filtered velocity fields from a corresponding direct numerical simulation, and the training is by a stochastic gradient descent method, which uses the adjoint Navier-Stokes equations to provide the end-to-end sensitivities of the model weights to the velocity fields. In-sample and out-of-sample testing on multiple dual-jet configurations show that its required mesh density in each coordinate direction for prediction of mean flow, Reynolds stresses, and spectra is half that needed by the dynamic Smagorinsky model for comparable accuracy. The same neural-network model trained directly to match filtered subgrid-scale stresses, without the constraint of being embedded within the flow equations during the training, fails to provide a qualitatively correct prediction. The coupled formulation is generalized to train based only on mean-flow and Reynolds stresses, which are more readily available in experiments. The mean-flow training provides a robust model, which is important, though a somewhat less accurate prediction for the same coarse meshes, as might be anticipated due to the reduced information available for training in this case. The anticipated advantage of the formulation is that the inclusion of resolved physics in the training increases its capacity to extrapolate. This is assessed for the case of passive scalar transport, for which it outperforms established models due to improved mixing predictions.

DOI: [10.1103/PhysRevFluids.6.050502](https://doi.org/10.1103/PhysRevFluids.6.050502)

### I. INTRODUCTION

The attractiveness of subgrid-scale models that reduce the resolution needs of large-eddy simulation (LES) methods is well understood [1]. However, the required resolution to achieve sufficiently universal subgrid-scale behavior, which simplifies modeling, remains higher than might be desired [2] and, even for finely resolved LES, the modeling challenge is not necessarily simple [3]. Flexibility to produce diverse subgrid-scale behaviors is, in a sense, at odds with cleanly expressed models, which are necessarily simpler than the diverse range of phenomenology that they seek to represent.

\*jmacart@nd.edu

†justin.sirignano@maths.ox.ac.uk

‡jbfreund@illinois.edu

Models with little empiricism, perhaps most notably dynamic models [4,5], are obviously attractive for their anticipated convergence and concomitant portability, and are expected to see continued use indefinitely, though models with greater parametric flexibility might afford opportunities for further relaxing mesh resolution needs. We examine such an approach here and introduce subgrid-scale models with many parameters. The hope is that the greater empiricism will reduce the resolution requirements yet still provide useful extrapolation to new configurations, beyond the training data. Maybe most importantly for motivation, we also seek procedures that are sufficiently flexible to include additional subgrid-scale physical effects, beyond that expressed in the flow equations and possibly beyond that which has been well described to date by governing equations. The attractiveness of the dynamic procedure depends upon the turbulence physics that underlie its spectral extrapolation procedure. Cases with coupled physics, such as turbulent combustion, do not lend themselves directly to decoupled subgrid-scale predictions. We therefore also seek to use the additional empiricism to increase flexibility in the proposed models to facilitate incorporation of observations that lack as complete of a mathematical description as turbulence.

The approach we take uses the formulations, optimization procedures, and efficient implementations available for deep neural networks. As for any large-eddy simulation, the  $\bar{\cdot}$  filtered equations will be solved numerically,

$$\frac{\partial \bar{u}_i}{\partial t} + \bar{u}_j \frac{\partial \bar{u}_i}{\partial x_j} = -\frac{1}{\rho} \frac{\partial \bar{p}}{\partial x_i} + \nu \frac{\partial^2 \bar{u}_i}{\partial x_j^2} + \frac{\partial h_{ij}(\bar{\mathbf{u}}; \bar{\theta})}{\partial x_j}, \quad (1)$$

$$\frac{\partial \bar{u}_j}{\partial x_j} = 0, \quad (2)$$

with the closure  $h_{ij} = h_{ij}(\bar{\mathbf{u}}; \bar{\theta})$ , with model parameters  $\bar{\theta}$ , dependent on the velocity and its derivatives. The dimension of  $\bar{\theta}$  is small for established subgrid-scale models and nominally zero for models based on a dynamic procedure. For large-eddy simulation, the simplest approach, here referred to as the direct approach, is to minimize

$$L(\bar{\theta}) = D(\tau_{ij}^{\text{SGS}}(\bar{\mathbf{u}}_e, \overline{\mathbf{u}_e^\top \mathbf{u}_e}), h_{ij}(\bar{\mathbf{u}}_e; \bar{\theta})), \quad (3)$$

where  $D$  is an appropriate distance

$$D(\mathbf{a}, \mathbf{b}) = \int_0^T \int_{\Omega} \|\mathbf{a} - \mathbf{b}\| \, d\mathbf{x} \, dt \quad (4)$$

and  $\tau_{ij}^{\text{SGS}}$  and  $\mathbf{u}_e$  are *a priori* trusted data, likely from a direct numerical simulation (DNS) or possibly from advanced experimental methods. Accelerating optimization of  $\bar{\theta}$  by using the gradient

$$\frac{\partial L}{\partial \bar{\theta}} = \frac{\partial L}{\partial \mathbf{h}} \cdot \frac{\partial \mathbf{h}}{\partial \bar{\theta}} \quad (5)$$

is straightforward, with the center dot operator here representing an appropriate inner product as needed. The first factor in (5) follows directly from the selected measure of mismatch (3), and the second is the usual sensitivity gradient available from a back-propagation training algorithm, as is available in any neural-network software [6]. One drawback is the availability of  $\tau_{ij}^{\text{SGS}}$ , which is a complex quantity in a turbulent flow. Even when  $\tau_{ij}^{\text{SGS}}$  and  $\mathbf{u}_e$  are available from advanced diagnostics or direct numerical simulation, it is alarming that the optimization problem (3) ignores the governing equations (1) and (2), so the trained model is at most indirectly constrained by the known and resolved physics expressed in the unclosed governing equations, which risks diminishing its capacity for extrapolative prediction.

Such a direct approach, using a variant of (3), has been employed with some success in Reynolds-averaged Navier-Stokes (RANS) modeling to close the mismatch with trusted data [7]. Of course, in Reynolds-averaged descriptions, the mismatch is relatively simple to infer from mean flow fields. In contrast, for large-eddy simulation, it is time dependent, three dimensional, and smaller scale, so

it is difficult to design a training regimen that would with confidence close the mismatch [8,9]. Still more challengingly, it is unclear how to anticipate the limits of model validity: How far from the training data will the model successfully extrapolate? For this question, how even to define distance is not obvious. Training on the mismatch (the direct approach) has the additional challenge that it requires data specifically aligned with that mismatch. Such correspondingly rich data sets likely do not exist in many machine-learning applications to fluid mechanics [10].

The approach we pursue is designed to both reduce the need for training data and remove the requirement of obtaining the challenging data for  $\tau_{ij}^{sgs}$ . To constrain the models by the physics in (1) and (2), we minimize

$$L(\vec{\theta}) = D(F(\bar{\mathbf{u}}_e), F(\bar{\mathbf{u}})), \quad (6)$$

where  $F$  is a function of the flow field,  $\bar{\mathbf{u}}_e$  indicates that  $F(\bar{\mathbf{u}}_e)$  is derived from a trusted flow field, even if  $\bar{\mathbf{u}}_e$  is not necessarily known in detail, and  $\bar{\mathbf{u}}$  is the *a posteriori* solution of (1) and (2). Thus, minimization of (6) is a PDE-constrained optimization problem, unlike (3). If  $\bar{\mathbf{u}}_e$  is available, at least in some  $\Omega$  part of the domain, then  $F(\bar{\mathbf{u}}) = \bar{\mathbf{u}}$  can be simply used, though any convenient and effective quantity of interest can, in principle, be selected.

Unfortunately, this indirect loss function (6) greatly increases the training challenge because any new set of parameters  $\vec{\theta}$  requires a new flow solution to determine  $\bar{\mathbf{u}}$  and thus  $F(\bar{\mathbf{u}})$ . Gradient-based acceleration of the optimization, in particular, is more challenging to implement due to the dimension of  $\vec{\theta}$ . The gradient of (6) for the discretized form of (1) satisfies

$$\frac{\partial L}{\partial \vec{\theta}} = \frac{\partial L}{\partial \bar{\mathbf{u}}} \cdot \frac{\partial \bar{\mathbf{u}}}{\partial \mathbf{h}} \cdot \frac{\partial \mathbf{h}}{\partial \vec{\theta}}, \quad (7)$$

the middle factor of which describes the change in the flow solution due to changes in the closure term  $\mathbf{h}$ . The center dot is a dot product, and the dimension of  $\frac{\partial L}{\partial \bar{\mathbf{u}}}$  matches the number of space and time degrees of freedom in the discretization. For the continuous case of (1), (7) would be replaced by functional derivatives that satisfy an adjoint partial differential equation (PDE). Adjoint-based methods are well understood to efficiently provide such high-dimensional gradient information, and we employ them here. Most of the details of the numerical implementation are available elsewhere [11], where they were demonstrated for the limited case of  $F(\bar{\mathbf{u}}) = \bar{\mathbf{u}}$  for isotropic turbulence. A similar approach has also been demonstrated in two-dimensional aerodynamics applications [12], based on a similarly motivated framework previously demonstrated on illustrative model applications [13]. The same challenges also motivate recent efforts to achieve more-robust learning via more sophisticated network models, rather than building the governing equations directly into the learning procedure [14]. Another approach uses reverse-mode differentiable programming to implement adjoint PDEs [15], similar to algorithmic differentiation [16], and could be useful in solving (7) for complex systems.

Building on the success for isotropic turbulence [11], several temporally developing planar turbulent shear flows with Reynolds number 6000 are used to assess the method for the more complex case of free shear flow, including its scalar mixing, and multiple training regimens. The numerical simulation methods are summarized in Sec. II. The turbulent flows considered are introduced in Sec. III. The neural-network model is standard; for completeness, its particulars are summarized in Sec. IV. We also consider nontrivial  $F$  in (6) and portability of trained  $\mathbf{h}$  to other flows (Sec. V). The form of the trained  $\mathbf{h}$  is analyzed in Sec. VI, and how it better represents turbulence on relatively coarse meshes is analyzed in Sec. VII. The extrapolative capacity of the method is demonstrated in Sec. VIII, where a model trained on subgrid-scale stresses is shown to be more effective than established models when applied directly, without additional training, to scalar transport. Finally, in Sec. IX we consider the implications of the mesh resolution. Conclusions are discussed in the context of future directions in Sec. X.

TABLE I. Conditions for the initial mean velocity profiles (13) for case A and (18) for cases B–D.

Case	Configuration	$Re_0 = U_{j0}H_0/\nu$	$H_1/H_0$	$U_{j1}/U_{j0}$	$(L, H, W)/H_0$	Uniform mesh
A	Single jet		–	–		
B	Dual jets	6000	1.0	1.0	16,20,12	$1024 \times 1280 \times 768$
C	Dual asym.		2.0	+0.5		
D	Dual antipara.		2.0	–0.5		

## II. FLOW SIMULATIONS

### A. Governing equations

To provide the direct numerical simulation training and testing data, the incompressible-fluid flow equations, including transport of a passive scalar  $\xi$ ,

$$\frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial x_j} = -\frac{1}{\rho} \frac{\partial p}{\partial x_i} + \nu \frac{\partial^2 u_i}{\partial x_j^2}, \quad (8)$$

$$\frac{\partial u_j}{\partial x_j} = 0, \quad (9)$$

$$\frac{\partial \xi}{\partial t} + u_j \frac{\partial \xi}{\partial x_j} = D \frac{\partial^2 \xi}{\partial x_j^2}, \quad (10)$$

are discretized with sufficient resolution such that all turbulence scales are sufficiently resolved.

For the large-eddy simulations, (8) is replaced with either

$$\frac{\partial \bar{u}_i}{\partial t} + \bar{u}_j \frac{\partial \bar{u}_i}{\partial x_j} = -\frac{1}{\rho} \frac{\partial \bar{p}}{\partial x_i} + \nu \frac{\partial^2 \bar{u}_i}{\partial x_j^2} - \frac{\partial \tau_{ij}^r}{\partial x_j}, \quad (11)$$

for comparison with established  $\tau_{ij}^r$  closure approaches, or (1), with the fitted neural-network closure  $h_{ij}$ . For the anisotropic residual stress  $\tau_{ij}^r = \tau_{ij}^{\text{SGS}} - \tau_{kk}^{\text{SGS}} \delta_{ij}/3$ , we consider the Smagorinsky model [17,18] with constant coefficient  $C_s = 0.18$ , the dynamic Smagorinsky model [4,5], and the model of Clark *et al.* [19]. Differently from a previous demonstration [11], where  $\mathbf{h}$  modeled  $\nabla \cdot \bar{\tau}^r$ , here  $h_{ij}$  models  $\tau_{ij}^r$  directly for smoother optimization and greater modeling flexibility (Sec. VE). Moreover, it enables us to compare the functional form of  $h_{ij}$  to traditional models (Sec. VI).

The filter underlying the large-eddy simulation formulation is

$$\bar{\phi}(\mathbf{x}, t) = \int_{\Omega} G(\mathbf{y}, \mathbf{x}) \phi(\mathbf{x} - \mathbf{y}, t) d\mathbf{y}, \quad (12)$$

though this is only explicitly used for comparison with direct numerical simulations and for generating large-eddy simulation initial conditions. For this,  $G(\mathbf{y}, \mathbf{x})$  has unit support in a  $\bar{\Delta}$  cube centered on  $\mathbf{x}$  [19].

### B. Numerical discretization

Second-order centered differences are used on a staggered mesh [20], except for convective terms in the scalar equation (10), which were discretized using a third-order weighted essentially nonoscillatory (WENO) scheme [21]. Time integration is by a fractional-step method [22] and a linearized trapezoid method in an alternating-direction implicit (ADI) framework for the diffusive terms, as implemented by Desjardins *et al.* [23]. Mesh densities for the direct numerical simulation cases listed in Table I were selected so that the Kolmogorov scale  $\eta = (\nu^3/\varepsilon)^{1/4}$ , based on the calculated viscous dissipation rate  $\varepsilon$ , is resolved with uniform mesh spacing  $\Delta x_{\text{DNS}}$  such that  $\Delta x_{\text{DNS}}/\eta < 1.8$  in the turbulence regions, which matches successful resolutions of turbulent jets




FIG. 1. One-dimensional energy spectra from direct numerical simulation of a temporally evolving jet (case A; see Table I) at the centerline ( $y/H_0 = 0$ ) and in the shear layer ( $y/H_0 = 2$ , shown offset by a factor of 10), at time  $t = 62.5t_{j0}$ : (a) streamwise and (b) spanwise spectra. Dashed lines are the filtered direct numerical simulation data, with vertical lines indicating the  $\Delta_{16}$  filter-cutoff wave number. Reference  $\sim \kappa^{-5/3}$  lines are shown.

at similar Reynolds number [24–28]. Spectra in Fig. 1 confirm that several decades of turbulence energy decay are resolved. The time step  $\Delta t_{\text{DNS}}$  was fixed, with initial Courant–Friedrichs–Lewy (CFL) number about 0.37 that decreased to about 0.18 by the end of the simulation. *A posteriori* large-eddy simulations used the same numerical methods together with subgrid-scale models or the deep-learning model discussed subsequently in Sec. IV. Analogous filtered spectra in Fig. 1 with  $\Delta_{16}$  show a relatively high energy at the cutoff wave number. This substantially increases the modeling challenge, as is shown in Sec. III C.

The adjoint-based neural-network model training, which is described in Sec. IV B, uses the same numerical methods for the forward solution with the exception of time integration, which is by the fourth-order Runge–Kutta method for simplicity; thus the temporal adjoint solution is approximate. Since the overall order of accuracy is limited by the pressure-projection step, it is comparable to that of the forward solution. We further note that the adjoint solution is computed over only short time horizons (Sec. IV B), for which the slightly different energy-conservation properties of the forward and adjoint time integrators do not become appreciable. An exact (in space only) discrete adjoint was solved.

### III. CONFIGURATIONS

Free-shear-flow turbulence from multiple temporally developing planar jets, some with multiple streams, will be used for training and testing. These are introduced in the following subsections.

#### A. Turbulent plane jets

The baseline turbulent plane jet has a rounded top-hat initial streamwise mean velocity

$$\hat{u}(y) = \frac{U_{j0}}{2} \left[ \tanh \frac{y/H_0 + 1/2}{\delta_L} - \tanh \frac{y/H_0 - 1/2}{\delta_L} \right] \quad (13)$$

of width  $\approx H_0$  bounded by shear layers of thickness  $\approx \delta_L = 0.1H_0$ . The Reynolds number is  $\text{Re}_0 = U_{j0}H_0/\nu = 6000$ .

To seed transition to turbulence, spatially periodic fluctuations in the  $x$ - $z$  plane are added to the streamwise and cross-stream velocities:

$$\hat{u}(y, z) = a \hat{u}(y) \cos \left[ \frac{16\pi z}{L} \right] + \hat{u}(y), \quad (14)$$

$$\hat{w}(x, y) = a \hat{u}(y) \cos \left[ \frac{16\pi x}{W} \right], \quad (15)$$

where  $a = 0.1$ , and  $L$  and  $W$  are the domain sizes in the periodic  $x$  and  $z$  directions, respectively (see Table I). Three-dimensional random perturbations are superimposed on top of these at each mesh point:

$$u(x, y, z) = \{1 + a[r_x(x, y, z) - 0.5]\} \hat{u}(y, z), \quad (16)$$

$$w(x, y, z) = \{1 + a[r_z(x, y, z) - 0.5]\} \hat{w}(x, y), \quad (17)$$

where  $r_x$  and  $r_z \in [0, 1]$  are uniformly distributed. The initial cross-stream velocity is  $v(x, y, z) = 0$ . This single-jet baseline case is case A in Table I.

Additional dual-jet cases, which are anticipated to involve entrainment and merging processes not seen in the single jet [29], are used for out-of-sample tests. These additional processes occur over significantly longer duration than the development of approximate self-similarity in the single jet. Three specific dual-jet cases are considered: symmetric parallel jets (case B), asymmetric parallel jets (case C), and asymmetric antiparallel jets (case D). These have initial streamwise mean velocity profile

$$\begin{aligned} \hat{u}(y) = & \frac{U_{j0}}{2} \left[ \tanh \frac{(y + H_{s,0})/H_0 + 1/2}{\delta_L} - \tanh \frac{(y + H_{s,0})/H_0 - 1/2}{\delta_L} \right] \\ & + \frac{U_{j1}}{2} \left[ \tanh \frac{(y - H_{s,1})/H_1 + 1/2}{\delta_{L,1}} - \tanh \frac{(y - H_{s,1})/H_1 - 1/2}{\delta_{L,1}} \right], \end{aligned} \quad (18)$$

where  $H_s = \min(H_0, H_1)$ ,  $H_{s,0} = (H_s + H_0)/2$ ,  $H_{s,1} = (H_s + H_1)/2$ , and  $\delta_{L,1} = (H_1/H_0)\delta_L$ . The primary and secondary jets are set to have the same Reynolds number ( $\text{Re}_0 = \text{Re}_1$ ) based on their speed and width. Perturbations are superimposed following the same long-wavelength and random perturbations as for the single jet. Specific  $H_1/H_0$  and  $U_{j1}/U_{j0}$  ratios are listed in Table I.

## B. Jet evolution

Reynolds averages are over the periodic streamwise ( $x$ ) and transverse ( $z$ ) directions and are functions of the cross-stream ( $y$ ) coordinate and time ( $t$ ):

$$\langle \phi \rangle(y, t) = \frac{1}{L_x L_z} \int_0^{L_x} \int_{-L_z/2}^{L_z/2} \phi(x, y, z, t) dz dx, \quad (19)$$

with discrete analog

$$\langle \phi \rangle(y_j, t^n) = \frac{1}{N_x N_z} \sum_{i=1}^{N_x} \sum_{k=1}^{N_z} \phi(x_i, y_j, z_k, t^n), \quad (20)$$

where  $N_x$  and  $N_z$  are the number of mesh points in the  $x$  and  $z$  directions. Subsequently, presented large-eddy simulation results were also ensemble averaged in time using at least two realizations initialized with different random perturbations sampled from the same distributions.

The evolving centerline velocity  $U_{cl}(t) \equiv \langle u \rangle(0, t)$  and half-width  $y_{1/2}(t)$ , defined by  $\langle u \rangle(y_{1/2}, t) = U_{cl}(t)/2$ , are shown in Fig. 2 for the single jet (case A). There is an adjustment from the initial non-self-similar pluglike flow to an approximately linear spreading with  $\sim t^{-1/2}$  centerline velocity decay for  $t \gtrsim 15t_{j0}$ ,  $t_{j0} = H_0/U_{j0}$ , consistent with an onset of similarity. Spectra

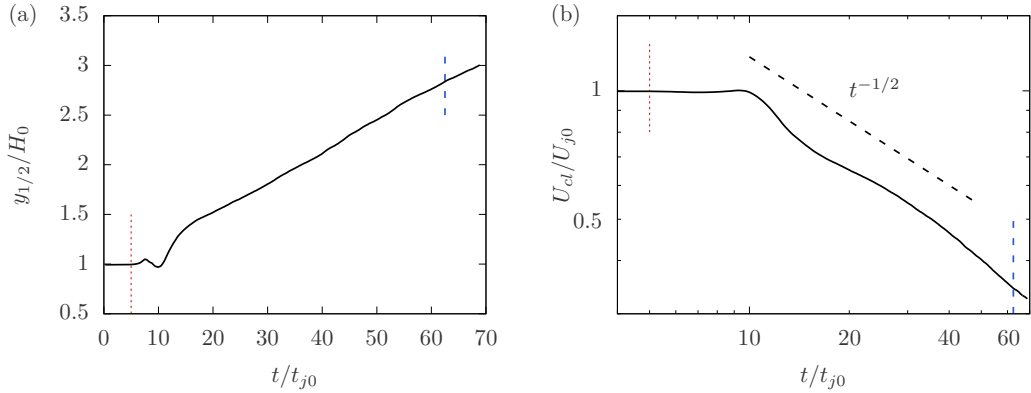


FIG. 2. (a) Half-width and (b) centerline velocity of the case A jet direct numerical simulation. Subsequently discussed large-eddy simulations are initialized at  $t = 5t_{j0}$ , indicated by dotted lines. Comparisons for this case are carried out at  $t = 62.5t_{j0}$ , indicated by dashed lines.

at  $t = 62.5t_{j0}$ , such as are used subsequently for model evaluation, are shown in Fig. 1. The vorticity magnitude at  $t = 62.5t_{j0}$ , as visualized in Fig. 3(a), shows a range of large- and small-scale turbulence.

The large-eddy simulations are initialized by filtering the direct numerical simulation fields at  $t = 5t_{j0}$ . As seen in Fig. 2, this is prior to the period of approximate self-similarity, which presumably increases the difficulty of matching statistics at later times. This initialization is also more representative of realistic LES, for which fully developed filtered-DNS initial conditions are typically unavailable. Primary comparisons for the single-jet case are made at  $t = 62.5t_{j0}$ , which is approximately four times the time for first achieving approximate self-similarity based on the

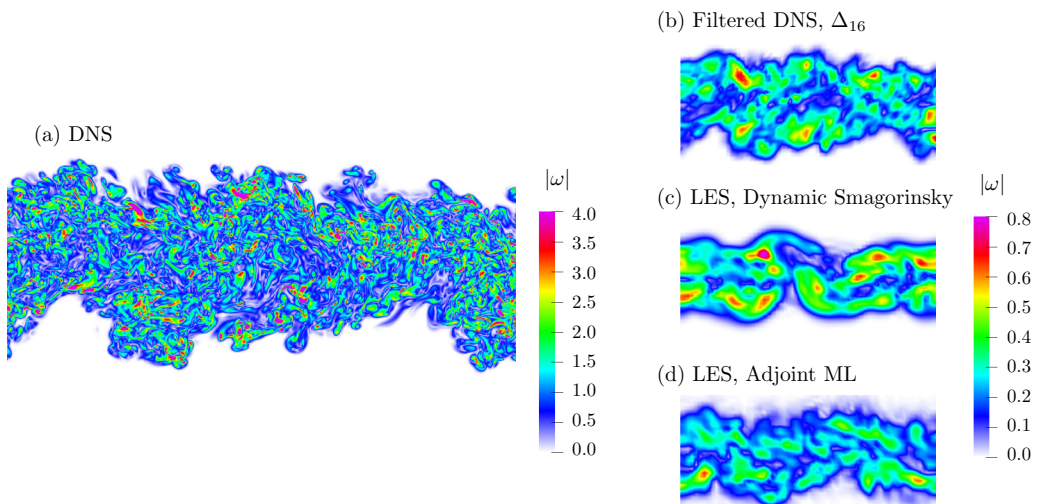


FIG. 3. Vorticity magnitude at  $z = 0$  for  $t = 62.5t_{j0}$ : (a) direct numerical simulation case A ( $1024 \times 1280 \times 768$  mesh points); (b) the same field as (a) filtered using  $\bar{\Delta}/\Delta_{\text{DNS}} = 16$  for an effective mesh of  $64 \times 80 \times 48$  points; (c) large-eddy simulation using the dynamic Smagorinsky subgrid-scale model, initialized at  $t = 5t_{j0}$  (see Sec. III B), on  $64 \times 80 \times 48$  mesh points; and (d) the analogous  $64 \times 80 \times 48$  simulation on using the adjoint-trained deep-learning model (Sec. IV).

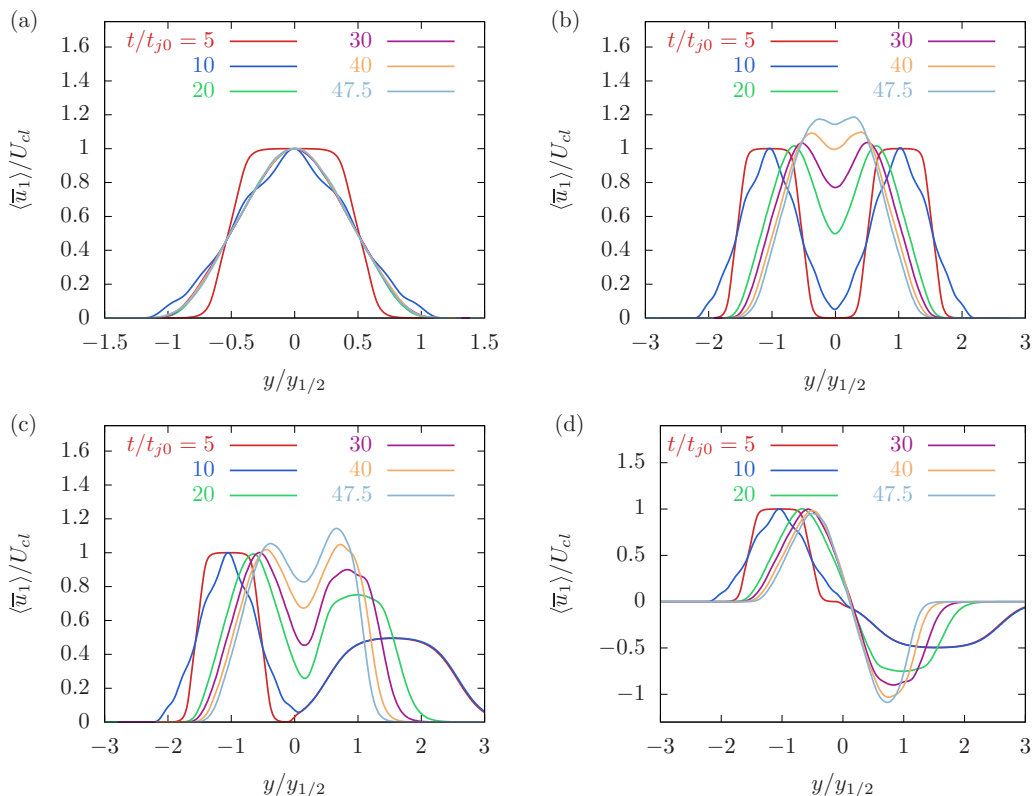


FIG. 4. Mean streamwise velocity for (a) single jet (case A), (b) symmetric parallel jets (case B), (c) asymmetric parallel jets (case C), and (d) asymmetric antiparallel jets (case D), all normalized by case A similarity variables. Profiles are shown starting at  $t = 5t_{j0}$ .

mean velocity. At this time, the jet has grown to approximately 2.75 times its original width. As an example, the vorticity magnitude of filtered DNS data is compared to large-eddy simulations using the dynamic Smagorinsky and adjoint-trained machine-learning model in Figs. 3(b)–3(d). On this coarse grid, the learned model, described in Sec. IV, produces turbulence that qualitatively better matches the filtered direct numerical simulation. Quantitative comparisons of these and other model predictions are presented in Sec. V.

Mean velocities for all cases are shown in Fig. 4. As expected, the single-jet profiles exhibit approximately self-similar evolution for  $t \gtrsim 15t_{j0}$ . However, the dual-jet cases do not, so a successful model will need to reproduce a presumably more challenging non-self-similar evolution. Subsequent quantitative comparisons for the dual-jet cases are at  $t = 50t_{j0}$  (case B) and  $t = 42.5t_{j0}$  (cases C and D).

### C. Baseline subgrid-scale models

The deep-learning subgrid-scale model is analyzed and compared primarily against a standard dynamic Smagorinsky model [4,5]. Baseline performance of the dynamic Smagorinsky model is established by comparing with filtered direct numerical simulations. Since there is not explicit filtering during the large-eddy simulation, the nominal filter width and mesh resolution are equivalent:  $\bar{\Delta} \equiv \Delta_N = N\Delta_{\text{DNS}}$  [18].

Figure 5 shows mean streamwise velocity predictions of the single-jet case (A) for  $\Delta_{N=16,8,4,2}$ . The  $\Delta_{16}$  case overpredicts the centerline velocity and underpredicts the jet spreading, though, as






FIG. 5. Direct numerical simulation, filtered DNS, and dynamic Smagorinsky (DS) [4] large-eddy simulation mean streamwise velocity  $\bar{u}$  for the single jet (case A) at  $t = 62.5t_{j0}$  for filter sizes indicated:  $\Delta_{16,8,4,2}$ .

expected, agreement improves with finer meshes  $\Delta_{8,4,2}$ . This increases the separation between the dynamic-model test-filter scale and the large turbulence scales, thereby improving the realism of the isotropy assumption. The DNS mean filtered velocity is also shown in Fig. 5; this deviates only slightly from the unfiltered mean velocity.

Budgets of turbulence kinetic energy provide a measure of the role of the models' resolved-to-subgrid scale energy transfer versus viscous dissipation. Upon filtering, energy may be decomposed into resolved  $E_f$  and residual (subgrid-scale)  $k_r$  components as

$$\bar{E} \equiv \frac{1}{2} \overline{u_i u_i} = \underbrace{\frac{1}{2} \bar{u}_i \bar{u}_i}_{E_f} + \underbrace{\frac{1}{2} (\overline{u_i u_i} - \bar{u}_i \bar{u}_i)}_{k_r}. \quad (21)$$

A gauge of model importance can be inferred from the relative amounts viscous dissipation of  $E_f$  by the resolved (filtered) velocity field

$$\varepsilon_f \equiv 2\nu \bar{S}_{ij} \bar{S}_{ij} \quad (22)$$

and the production rate of residual kinetic energy

$$P_r \equiv -\tau_{ij}^r \bar{S}_{ij}. \quad (23)$$

For both of these,  $\bar{S}_{ij}$  is the filtered strain rate. Both  $\varepsilon_f$  and  $P_r$  represent sink terms when positive.

Figure 6 compares the large-eddy simulation and filtered direct numerical simulation  $\varepsilon_f$  and  $P_r$  for the baseline jet. Ideally,  $\varepsilon_f$  and  $P_r$  should match the filtered direct numerical simulation values for the same filter size. As filter size increases,  $\varepsilon_f$  is expected to decrease and  $P_r$  to increase, and both do exhibit these trends. However, the dynamic Smagorinsky model increasingly overpredicts both  $\varepsilon_f$  and  $P_r$  with increasing filter size. This means that there is excessive removal of kinetic energy from the resolved scales. For  $\Delta_4$ ,  $P_r$  is a small fraction of  $\varepsilon_f$ , so the model for  $\tau_{ij}^r$  plays only a minor role. The  $\Delta_4$  filtered  $\varepsilon_f^{\text{DNS}}$  is about half the unfiltered dissipation  $\varepsilon^{\text{DNS}}$  [Fig. 6(a)], so it dominates the evolution of the resolved kinetic energy. However, the balance is different for  $\Delta_{16}$ . Here, the subgrid-scale production dominates the resolved dissipation, making the model essential. Such a challenge affords opportunities for improvement, so we focus on  $\Delta_{16}$  for subsequent deep-learning modeling.

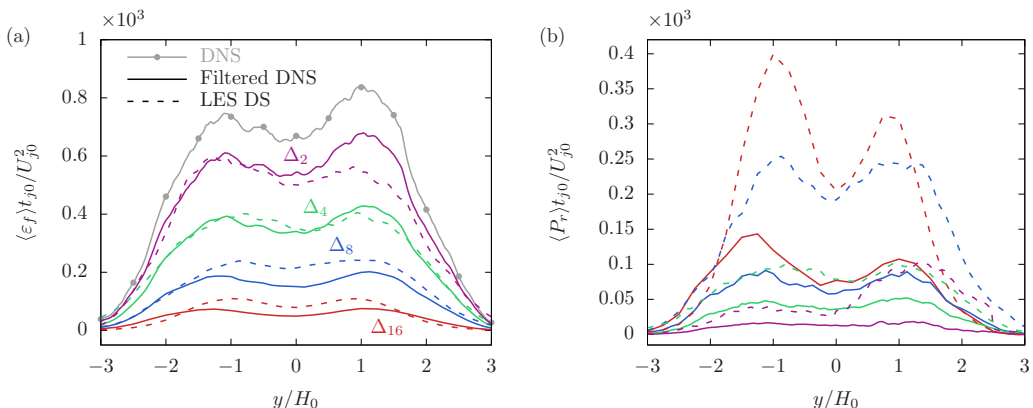


FIG. 6. Evaluation of the dynamic Smagorinsky (DS) model for the baseline single jet (case A): (a) resolved viscous dissipation  $\varepsilon_f$  (22) and (b) residual kinetic energy production rate  $P_r$  (23) at  $t = 62.5t_{j0}$  for filter sizes indicated.

#### IV. NEURAL-NETWORK MODEL AND TRAINING

The deep-learning model architecture, hyperparameters, and training algorithm are described in this section. Section IV A presents the neural-network model. Section IV B discusses the training algorithm.

##### A. Model architecture

We use a deep neural network  $\mathbf{G}(\vec{\phi}(\mathbf{x}, t); \vec{\theta})$ , where  $\vec{\phi}$  represents space- and time-dependent input data from the running simulation, with the following architecture:

$$\begin{aligned}
 H^1 &= \sigma(W^1 \vec{\phi} + b^1), \\
 H^2 &= \sigma(W^2 H^1 + b^2), \\
 H^3 &= G^1 \odot H^2 \quad \text{with} \quad G^1 = \sigma(W^5 z + b^5), \\
 H^4 &= \sigma(W^3 H^3 + b^3), \\
 H^5 &= G^2 \odot H^4 \quad \text{with} \quad G^2 = \sigma(W^6 z + b^6), \\
 \mathbf{G}(\vec{\phi}; \vec{\theta}) &= W^4 H^5 + b^4,
 \end{aligned} \tag{24}$$

where  $\sigma$  is a  $\tanh(\dots)$  element-wise nonlinearity,  $\odot$  denotes element-wise multiplication, and the parameters are  $\theta = \{W^1, W^2, W^3, W^4, W^5, W^6, b^1, b^2, b^3, b^4, b^5, b^6\}$ . The gated units  $G^1$  and  $G^2$  are chosen for their general capability to represent stiff, nonlinear behavior. Each layer includes  $N_H = 100$  hidden units, all initialized using a standard Xavier initialization [30]. Inputs to the neural network are normalized by the same set of constants for all cases.

In (1),  $\mathbf{h}(\mathbf{x}, t; \vec{\theta}) = \Psi(y)\mathbf{G}(\vec{\phi}(\mathbf{x}, t); \vec{\theta})$ , where in the implementation  $\Psi$  simply sets the neural network to zero at the  $y = \pm H/2$  boundary mesh points. We note that the discrete stress tensor  $\mathbf{h}^{i,j,k}$  is not constrained to be exactly symmetric on the staggered mesh at finite resolution. Similarly, the symmetric subgrid-scale stress components (e.g.,  $\tau_{12}^r$  versus  $\tau_{21}^r$ ) appear in different equations ( $x$  versus  $y$  momentum), which are not exactly constrained by the same symmetries as the (unfiltered) Navier-Stokes equations [2]. Thus, further model flexibility in the finite-resolution limit may be achieved by relaxing the symmetry requirement. The same argument applies for corresponding tensor invariants; they will not hold for finite resolution, so are not imposed on the network. To

be clear, as for any well-posed subgrid-scale model, symmetry and invariance properties will be achieved with increasing resolution. Symmetry is revisited in Sec. [VE](#) and invariants in Sec. [VII](#).

The output  $h_{ll}^{i,j,k}$  is at the cell center  $\mathbf{x} = (i\Delta + \frac{\Delta}{2}, j\Delta + \frac{\Delta}{2}, k\Delta + \frac{\Delta}{2})$ . For  $l \neq m$ , the output  $h_{lm}^{i,j,k}$  is at the cell corners

$$\mathbf{x} = \begin{cases} (i\Delta, j\Delta, k\Delta + \frac{\Delta}{2}), & \text{if } (l, m) = (1, 2) \\ (i\Delta, j\Delta + \frac{\Delta}{2}, k\Delta), & \text{if } (l, m) = (1, 3) \\ (i\Delta + \frac{\Delta}{2}, j\Delta, k\Delta), & \text{if } (l, m) = (2, 3). \end{cases} \quad (25)$$

The derivatives  $\frac{\partial h_{ij}}{\partial x_j}$  are then calculated at the cell faces using central differences with mesh size  $\Delta$ , for example,

$$\begin{aligned} \frac{\partial h_{11}}{\partial x} \left( i\Delta, j\Delta + \frac{\Delta}{2}, k\Delta + \frac{\Delta}{2} \right) &= \frac{h_{11}^{i,j,k} - h_{11}^{i-1,j,k}}{\Delta}, \\ \frac{\partial h_{12}}{\partial y} \left( i\Delta, j\Delta + \frac{\Delta}{2}, k\Delta + \frac{\Delta}{2} \right) &= \frac{h_{12}^{i,j+1,k} - h_{12}^{i,j,k}}{\Delta}, \\ \frac{\partial h_{13}}{\partial y} \left( i\Delta, j\Delta + \frac{\Delta}{2}, k\Delta + \frac{\Delta}{2} \right) &= \frac{h_{13}^{i,j,k+1} - h_{13}^{i,j,k}}{\Delta}. \end{aligned} \quad (26)$$

Models were implemented and tested using two different sets of input variables. We first tested using  $\vec{\phi} = (\bar{\mathbf{u}}, \partial_{x_i} \bar{\mathbf{u}}, \partial_{x_i x_j} \bar{\mathbf{u}})$  as an input, although this closure model is not Galilean invariant. For the simpler case of isotropic turbulence, for which this approach was first tested [[11](#)], the model appeared to learn Galilean invariance. However, for the jets here, not surprisingly, it was found to be unstable in *a posteriori* simulations, at least for the amount of training evaluated. Therefore, a Galilean invariant form with  $\vec{\phi} = (\partial_{x_i} \bar{\mathbf{u}}, \partial_{x_i x_j} \bar{\mathbf{u}})$  was used here.

The input  $\vec{\phi}$  to (24) for a large-eddy simulation grid cell  $(i, j, k)$  includes the variables  $\vec{\phi}^{i,j,k}$  for  $(i, j, k)$  as well as the six adjacent cells. Velocity-gradient inputs are calculated using the same stencils as the convective and viscous operators. Thus, on-diagonal derivatives  $\bar{u}_{i,l}$  are evaluated at cell centers, and off-diagonal derivatives  $\bar{u}_{l,m}$ ,  $m \neq l$ , are evaluated at cell faces. As an example,

$$\begin{aligned} \bar{u}_{1,x} \left( i\Delta + \frac{\Delta}{2}, j\Delta + \frac{\Delta}{2}, k\Delta + \frac{\Delta}{2} \right) &= \frac{\bar{u}_1(i\Delta + \Delta, j\Delta + \frac{\Delta}{2}, k\Delta + \frac{\Delta}{2}) - \bar{u}_1(i\Delta, j\Delta + \frac{\Delta}{2}, k\Delta + \frac{\Delta}{2})}{\Delta}, \\ \bar{u}_{1,y} \left( i\Delta, j\Delta, k\Delta + \frac{\Delta}{2} \right) &= \frac{\bar{u}_1(i\Delta, j\Delta + \frac{\Delta}{2}, k\Delta + \frac{\Delta}{2}) - \bar{u}_1(i\Delta, j\Delta - \frac{\Delta}{2}, k\Delta + \frac{\Delta}{2})}{\Delta}, \\ \bar{u}_{1,z} \left( i\Delta, j\Delta + \frac{\Delta}{2}, k\Delta \right) &= \frac{\bar{u}_1(i\Delta, j\Delta + \frac{\Delta}{2}, k\Delta + \frac{\Delta}{2}) - \bar{u}_1(i\Delta, j\Delta + \frac{\Delta}{2}, k\Delta - \frac{\Delta}{2})}{\Delta}. \end{aligned} \quad (27)$$

## B. Training algorithm

A stochastic gradient descent-type scheme is employed, in which the adjoint partial differential equation is solved on random time subintervals  $m = 1, \dots, M$  during each optimization iteration, each of which can be evaluated in parallel. The training algorithm is summarized below. We first consider the case in which we seek to minimize (6) as discussed in Sec. [I](#) with  $F(\bar{\mathbf{u}}) = \bar{\mathbf{u}}(\mathbf{x}, t)$ . The extension to training on mean statistics is described in Sec. [IV C](#).

- (i) Initialize the parameters  $\vec{\theta}^0$  with samples from the Xavier distribution [[30](#)].
- (ii) For training iteration  $k = 1, 2, \dots, K$ :
  - (a) Each subiteration  $m = 1, \dots, M$  selects a random time  $t^m$ ; the corresponding filtered direct numerical simulation data for  $[t^m, t^m + T_s]$  is taken as the training target:  $F(\bar{\mathbf{u}}_e) = \bar{\mathbf{u}}^{\text{DNS}}$ .

- (b) Equations (1) and (2) are solved over  $[t^m, t^m + T_s]$  with the filtered direct numerical simulation data at  $t^m$  projected onto a divergence-free manifold as the initial condition [11].
- (c) For each  $m$ , the adjoint is solved on  $[t^m + T_s, t^m]$  with objective function, rewritten from (6),

$$L^m(\theta) = \int_{\Omega} \|\bar{u}^{\text{DNS}}(t^m + T_s, x) - \bar{u}(t^m + T_s, x)\| dx. \quad (28)$$

- (d) The adjoint solution provides the gradient  $\nabla_{\bar{\theta}} L^m(\theta)$  for  $m = 1, \dots, M$ .
- (e) These  $M$  gradients are averaged as  $G = \frac{1}{M} \sum_{m=1}^M \nabla_{\theta} L^m(\theta)$ .
- (f) The parameter  $\bar{\theta}$  for the deep-learning large-eddy simulation model is updated using the RMSprop algorithm [31] with the gradient  $G$ :

$$\begin{aligned} \bar{r}^{k+1} &= \rho \bar{r}^k + (1 - \rho)G \odot G, \\ \bar{\theta}^{k+1} &= \bar{\theta}^k - \frac{\alpha^k}{\sqrt{\bar{r}^{k+1}} + \epsilon} \odot G, \end{aligned} \quad (29)$$

where  $\alpha^k$  is the learning rate, and we take  $\rho = 0.99$  and  $\epsilon = 10^{-8}$ . The  $\sqrt{\bar{r}^{k+1}}$  operation is element-wise. The learning rate  $\alpha^k$  is decreased by a factor of  $\frac{1}{2}$  every  $K$  iterations.

Taking  $K = 250$  and  $M = 12$  was found to be effective, though of course the optimal choice of such hyperparameters will generally depend upon the application and the model architecture. The time segments used had  $T_s = 5\Delta t_{\text{LES}}$ , with the large-eddy simulation time step  $\Delta t_{\text{LES}} = 10\Delta t_{\text{DNS}}$ .

### C. Training on the Reynolds-averaged statistics

As discussed in Sec. I, direct numerical simulation data, or correspondingly resolved experimental measurements, are not available in all cases. Therefore, as a relevant example of limited availability of data, we train a model only on the mean velocity and resolved Reynolds stresses:  $F(\bar{\mathbf{u}}) = (\langle \bar{\mathbf{u}} \rangle; \langle \bar{\mathbf{u}} \bar{\mathbf{u}}' \rangle)(y, t)$ . We obtain target data  $F(\bar{\mathbf{u}}_e) = (\langle \bar{\mathbf{u}}^{\text{DNS}} \rangle; \langle \bar{\mathbf{u}} \bar{\mathbf{u}}' \rangle^{\text{DNS}})$  from the filtered case A, though in general measured data may be substituted without modification. In this case, the objective function becomes

$$L(\theta) = \frac{1}{L_y} \frac{1}{(T_2 - T_1)} \int_{T_1}^{T_2} \int_0^{L_y} \sum_{i=1}^3 \frac{c_1}{U_{j0}^2} \left( \langle \bar{u}_i^{\text{DNS}} - \bar{u}_i \rangle^2 + \frac{c_1}{U_{j0}^2} \sum_{j=i}^3 \langle \bar{u}_i^{\text{DNS}} \bar{u}_j^{\text{DNS}} - \bar{u}_i \bar{u}_j \rangle^2 \right) dy dt, \quad (30)$$

with  $c_1 = U_{j0}^2 = 467.4$  found to be effective for the current application. Training this objective function is computationally more challenging than with (28) since (30) requires solving the flow and its adjoint on the full  $t \in [T_1 = 5, T_2 = 67.5]t_{j0}$  case A time horizon for each optimization iteration. Despite solving over a relatively long time length, the adjoint equation remains well behaved.

## V. MODEL PREDICTIONS

The objective functions, training data, and testing data for all cases are listed in Table II. We start in Sec. VA with an *a priori*  $\tau_{ij}^{\text{SGS}}$  model, trained directly based on the subgrid-scale stress mismatch (3). This simpler approach is demonstrated to be poor compared with the adjoint-based training in the following sections: in-sample results for the baseline case A jet (Sec. VB) and out-of-sample analysis of cases B, C, and D (Sec. VC). The performance for the averaged objective function (30) is assessed in Sec. VD.

### A. *A priori* training

The model of Sec. IVA is trained to minimize the *a priori* mismatch between the modeled and filtered direct numerical simulation  $\tau_{ij}^{\text{SGS}}$ . This is the simple, direct approach: training requires only

TABLE II. Objective functions, training data, and testing data for primary cases studied.

Objective function ( $F$ )	Training	Testing	Section
$\tau_{ij}^{\text{SGS}}(\bar{\mathbf{u}})$ ( <i>a priori</i> ML)	Jet (13)	Jet (13)	VA
$F(\bar{\mathbf{u}}) = \bar{\mathbf{u}}(\mathbf{x}, t)$ (filtered velocity)	Jet (13)	Jet (13) Dual jets (18) Dual asymmetric (18) Dual antiparallel (18)	VB VC VC
$F(\bar{\mathbf{u}}) = (\langle \bar{\mathbf{u}} \rangle; \langle \bar{\mathbf{u}}' \bar{\mathbf{u}}' \rangle)_{(y,t)}$ ( $x$ and $z$ averaged)	Jet (13)	Jet (13) Dual jets (18) Dual asymmetric (18) Dual antiparallel (18)	VD VD VD
$F(\bar{\mathbf{u}}) = \bar{\mathbf{u}}(\mathbf{x}, t)$ (filtered velocity; $h_{ij} = h_{ji}$ )	Jet (13)	Jet (13) Dual jets (18) Dual asymmetric (18) Dual antiparallel (18)	VE VE VE

5.5 h on 24 Nvidia K20X GPU nodes. The *a priori*  $\tau_{ij}^{\text{SGS}}$  agreement in Fig. 7(a) is good, as expected since deep neural networks are well understood to fit data well. However, this fit is unconstrained by how the model interacts with the governing equations, so it is unsurprising that the *a posteriori* performance in a large-eddy simulation is poor [Fig. 7(b)], even for the mean velocity. Simply minimizing the mismatch is insufficient, even for this relatively simple free-shear flow. In Fig. 7(b), the adjoint-trained neural network with  $F(\bar{\mathbf{u}}) = \bar{\mathbf{u}}$ , which is discussed subsequently, performs far better for the same network model.

### B. Single-jet case (in sample)

In this case, the model is trained on filtered instantaneous velocity fields [ $F(\bar{\mathbf{u}}) = \bar{\mathbf{u}}(\mathbf{x}, t)$ ] from case A and tested on the same configuration. In Fig. 8(a), we see that for  $\Delta_{16}$  the deep-learning model nearly perfectly reproduces the mean velocity, doing better than the dynamic Smagorinsky (DS) model for the coarse  $\Delta_{16}$  case and comparably to it for  $\Delta_8$ . In Fig. 8(b), it also better tracks the

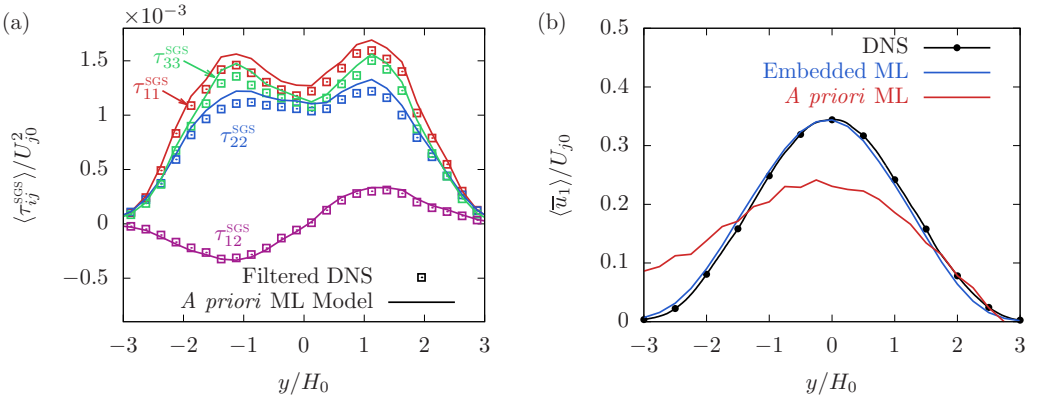


FIG. 7. The *a priori* trained deep-learning model for case A for  $\Delta_{16}$  at  $t = 62.5t_{j0}$ : (a) residual stress agreement from filtered direct numerical simulations and the *a priori* trained model, and (b) predicted mean streamwise velocity.

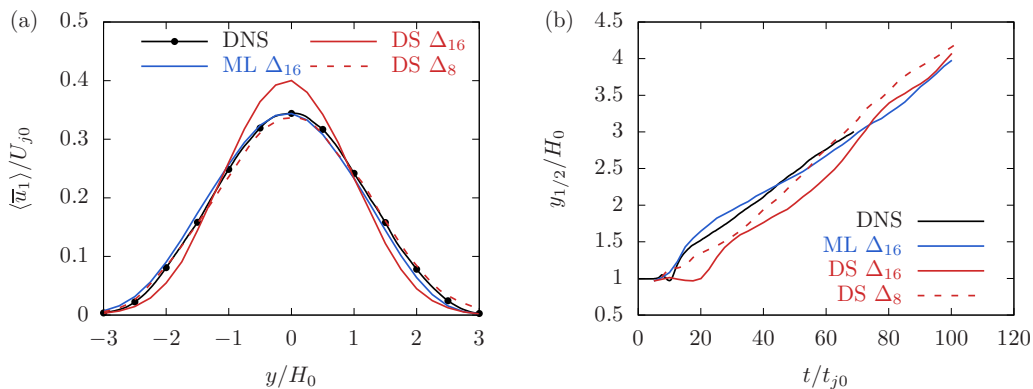


FIG. 8. Single-jet (case A) in-sample comparison for learning (ML) and dynamics Smagorinsky (DS) models: (a) mean streamwise velocity  $\bar{u}_1$  at  $t = 62.5t_{j0}$  and (b) half-width  $y_{1/2}$  evolution for the indicated filter sizes. The direct numerical simulation data are included for comparison.

evolution of  $y_{1/2}$ : for both  $\Delta_8$  and  $\Delta_{16}$ , the dynamic Smagorinsky model underpredicts jet spreading at early times.

Resolved Reynolds stress components  $R_{ij} \equiv \langle (\bar{u}_i - \langle \bar{u}_i \rangle)(\bar{u}_j - \langle \bar{u}_j \rangle) \rangle$  are compared in Fig. 9. As expected, based on the mean-flow evolution, the trained model significantly outperforms dynamic Smagorinsky at this coarse resolution. The complete Reynolds stress from the direct numerical simulation, prior to filtering to match the large-eddy simulation model, is also shown to quantify the subgrid-scale contribution to the turbulence stresses. One-dimensional energy spectra in Fig. 10 confirm that the deep-learning model more accurately reproduces the spectral energy distribution than the dynamic Smagorinsky model, and this is particularly evident at high wave numbers.

### C. Dual-jet cases (out of sample)

Out-of-sample mean-flow results are shown in Fig. 11 for all three dual-jet cases of Table I. The corresponding resolved viscous dissipation rates are shown in Fig. 12.

For the symmetric parallel jets, case B, the deep-learning model on  $\Delta_{16}$  has accuracy comparable to the dynamic Smagorinsky model on  $\Delta_8$  and predicts the resolved viscous dissipation rate. It thus predicts the merging process for which it was not trained. Conversely, for this resolution, the dynamic Smagorinsky model on the coarse  $\Delta_{16}$  mesh substantially underpredicts the jet merging, showing two distinct peaks of mean velocity. Similarly to case A, this corresponds to overprediction of  $\varepsilon_f$  near the high-shear regions.

The overall performance for the asymmetric parallel jets (case C) is similarly good: the coarse-grid deep-learning model more accurately reproduces the jet spreading than the dynamic Smagorinsky model for both  $\Delta_{16}$  and  $\Delta_8$  [Fig. 11(b)]. Of the three models, the deep-learning model also most accurately reproduces the peak velocity of the lower jet ( $y \approx -H_0$ ), though at the same time it is less accurate for the peak velocity of the slower-speed upper jet ( $y \approx 2H_0$ ). This might be because the lower jet more closely matches the training case A, whereas the upper jet is slower ( $U_1 = U_0/2$ ) and wider ( $H_1 = 2H_0$ ), as seen in Table I and Fig. 4. These trends are repeated in Fig. 12(b) for the resolved dissipation rate.

The deep-learning model similarly outperforms the dynamic Smagorinsky model in the asymmetric, antiparallel jets (case D) [Figs. 11(c) and 12(c)] for both  $\Delta_{16}$  and  $\Delta_8$ . Unlike previous cases, this configuration has much larger  $\varepsilon_f$  in the high-shear, interjet region ( $y \approx 0$ ) than in the jets themselves. The deep-learning model performs especially well in this region.

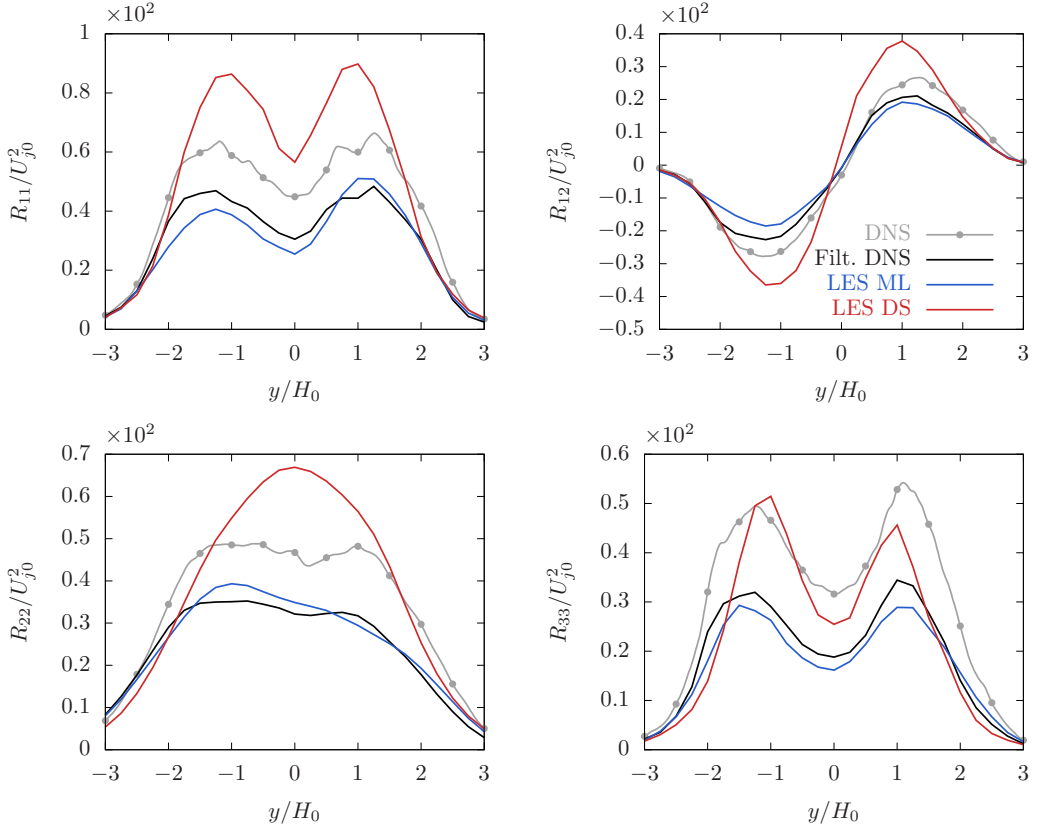


FIG. 9. Resolved Reynolds stress components as labeled in the single-jet case A at  $t = 62.5t_{j0}$  for  $\Delta_{16}$ : the new machine-learning model (ML), dynamic Smagorinsky (DS), and the filtered direct numerical simulation (DNS). The unfiltered Reynolds stress components are also shown for reference.




FIG. 10. (a) Streamwise and (b) spanwise energy spectra for the single-jet case A with  $\Delta_{16}$  at  $t = 62.5t_{j0}$ : filtered (DNS), machine-learning model (ML), and dynamic Smagorinsky (DS). The straight lines have  $\kappa^{-5/3}$  slope. Spectra are shown at the centerline ( $y/H_0 = 0$ ) and at  $y/H_0 = 1$ , which corresponds to the cross-stream location of maximum resolved kinetic energy, is shown offset by a factor of 10.

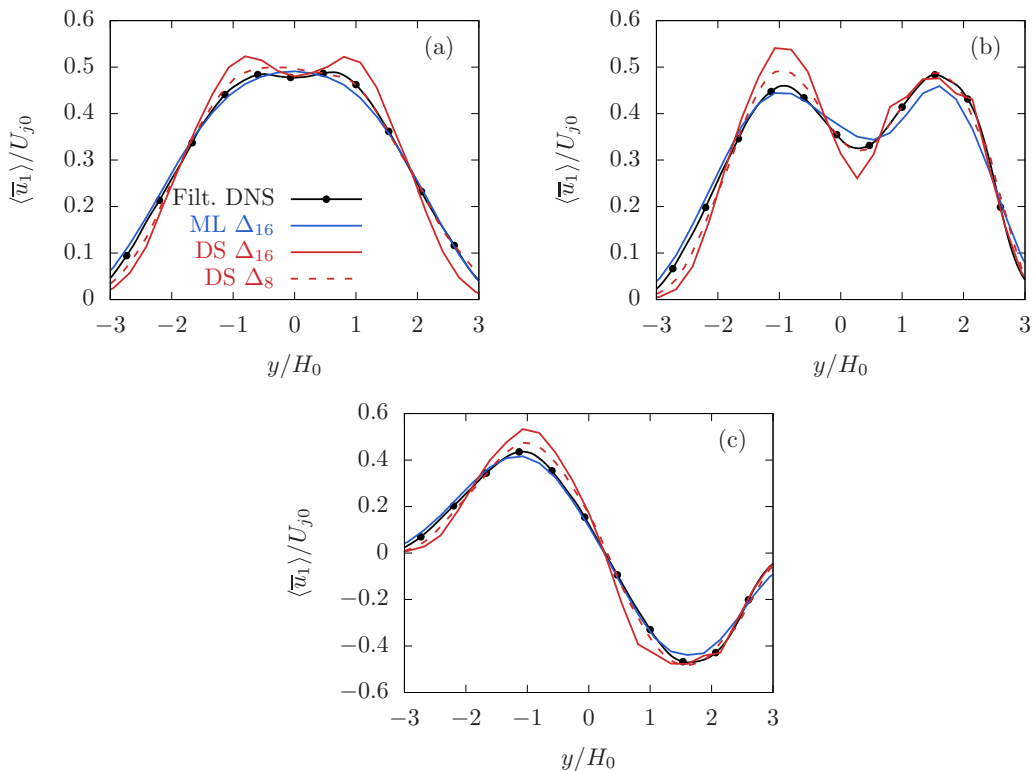


FIG. 11. Out-of-sample mean streamwise velocity  $\bar{u}_1$  for  $\Delta_{16}$  and the deep-learning (ML) and dynamic Smagorinsky (DS) models: (a) case B at  $t = 50t_{j0}$ , (b) case C at  $t = 42.5t_{j0}$ , and (c) case D at  $t = 42.5t_{j0}$ . The  $\Delta_8$  case is shown for reference.

One-dimensional energy spectra in Fig. 13 show that the deep-learning model performs particularly well at high wave numbers. This might follow from the ability of the deep-learning model to compensate for discretization errors, as considered for isotropic turbulence [11]. As can be seen in Fig. 13, the deep-learning model nearly eliminates unphysical near-cutoff wave-number dissipation.

#### D. Training for mean statistics

It is, of course, expected that the reduced information content in this case will make training subgrid-scale turbulence models more challenging. We take  $F(\bar{\mathbf{u}}) = (\langle \bar{\mathbf{u}} \rangle; \langle \bar{\mathbf{u}}' \bar{\mathbf{u}}' \rangle)_{(y,t)}$ , leading to the loss function (30) developed in Sec. IV C trained on the single-jet case A. This corresponds to the third set of tests in Table II. The same neural-network architecture (Sec. IV A) and training algorithm (Sec. IV B) are used as for the previous demonstrations.

The results for the mean flow are shown in Fig. 14. In all cases, the model produces a reasonable prediction, unlike the *a priori* training on the mismatch of Sec. V A. And, as expected, it performs well in sample, which is expected since it was trained specifically to match its Reynolds stress data. For the out-of-sample symmetric dual jet (case B), it also performs comparably to the model trained on the filtered data. However, with its less rich training set, it is not expected to extrapolate as well as the model trained with the richer three-dimensional filtered data. For the asymmetric out-of-sample jets (cases C and D), it only closely matches the data on the sides of the dual jets that are similar to the case A training data ( $y < 0$ ; see Table I and Fig. 4). On the other sides, the model overpredicts spreading, though not catastrophically, despite the huge reduction of training data size: the mean flow statistics have  $N_x N_z$  times fewer degrees of freedom per (20) than the instantaneous velocity.



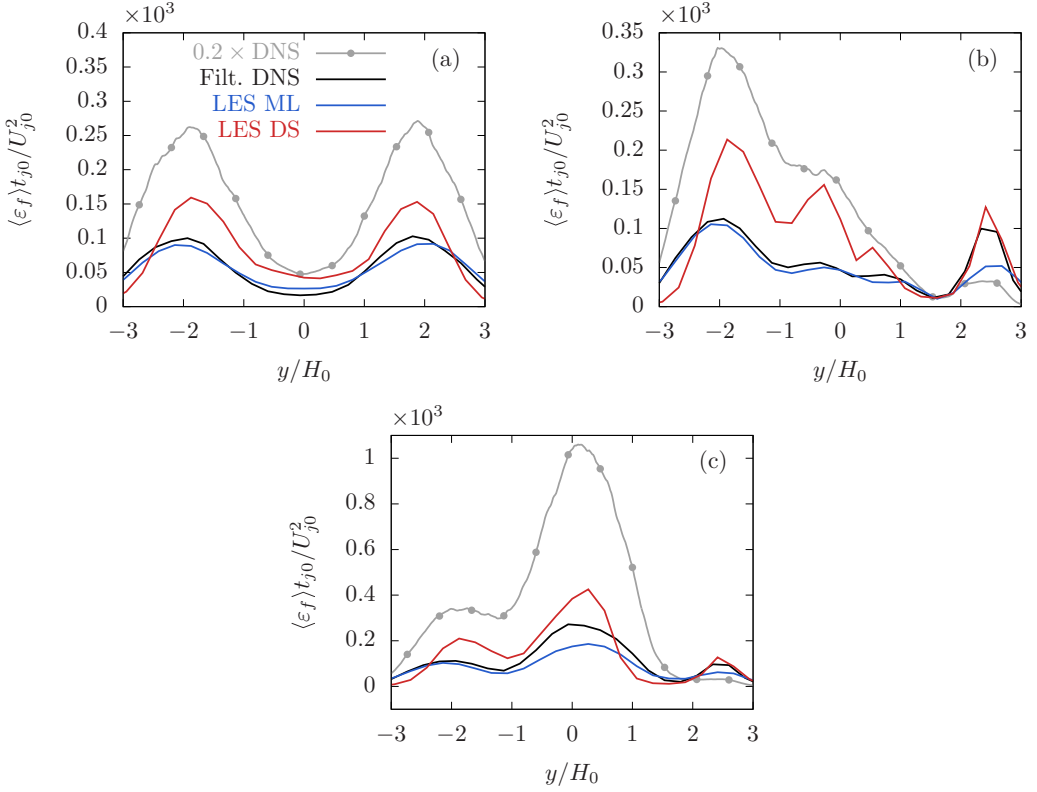


FIG. 12. Out-of-sample resolved viscous dissipation  $\varepsilon_f$  for the deep-learning (ML) and dynamic Smagorinsky (DS) models for  $\Delta_{16}$ : (a) case B at  $t = 50t_{j_0}$ , (b) case C at  $t = 42.5t_{j_0}$ , and (c) case D at  $t = 42.5t_{j_0}$ . The unfiltered direct numerical simulation viscous dissipation rate  $\times 0.2$  is shown for comparison.

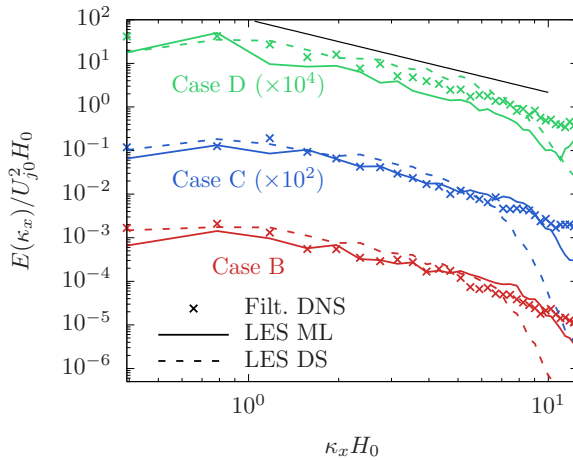


FIG. 13. Out-of-sample one-dimensional streamwise energy spectra at the point of peak resolved turbulence kinetic energy for  $\Delta_{16}$  at  $t = 50t_{j_0}$  for case B, at  $t = 42.5t_{j_0}$  for case C, and at  $t = 42.5t_{j_0}$  for case D. The straight line has  $\kappa^{-5/3}$  slope.




FIG. 14. Mean-flow development for training with (30) mean statistics (see Table II) for  $\Delta_{16}$ : in-sample (a) case A at  $t = 62.5t_{j_0}$  and out-of-sample (b) case B at  $t = 50t_{j_0}$ , (c) case C at  $t = 42.5t_{j_0}$ , and (d) case D at  $t = 42.5t_{j_0}$ . Results for the baseline instantaneous field training (ML  $\bar{\mathbf{u}}$  Asym.) and its symmetrically constrained variant of Sec. VE (ML  $\bar{\mathbf{u}}$  Sym.) are also shown.

It might be anticipated that expanded training data sets are likely necessary for this approach to be more accurate. Additional implications of the mean-statistics training are discussed in Sec. VII.

### E. A symmetric constraint

As introduced in Sec. IV A, the models discussed up to this point were not constrained to produce symmetric stress outputs, which is not a rigid constraint for typical discretizations of the governing equations. Corresponding results for a model that is trained to match the filtered velocity from case A [ $F(\bar{\mathbf{u}}) = \bar{\mathbf{u}}$ ], but now exactly constrained to produce symmetric stresses ( $h_{ij} = h_{ji}$ ), are shown Fig. 14 for both in-sample and out-of-sample cases. Its performance is markedly worse than either previous model in most circumstances. This highlights performance gained by relaxing the symmetry constraint, which is not consistent with the discretized form of the governing equations.

## VI. THE $\mathbf{h}$ CLOSURE

With the trained model working effectively in testing cases, especially when the training is based on the filtered instantaneous data, it is natural to analyze its form. As for any deep-learning model, it is difficult to infer the mode of its operation, which in turn risks hindrance of its physical interpretation for flows [10]. However, we can make some assessment of it. Here, we develop simplified functional forms based on  $\mathbf{h}$  from the trained neural network by calibrating candidate




FIG. 15. Taylor-series expansion  $\mathbf{g}$  from (32) fitted to the deep-learning (ML) model outputs  $\mathbf{h}$  for the  $\tau_{12}$  and  $\tau_{21}$  nonsymmetric model components.

functional forms  $\mathbf{g}$  that minimize

$$J(\vec{\phi}) = D(\mathbf{h}(\bar{\mathbf{u}}; \vec{\theta}), \mathbf{g}(\bar{\mathbf{u}}; \vec{\phi})), \quad (31)$$

as determined by a linear least-squares fit. Each candidate closure  $\mathbf{g}$  has a low-dimensional set of parameters  $\vec{\phi}$  that are fitted to the  $\mathbf{h}$  trained for  $F(\bar{\mathbf{u}}) = \bar{\mathbf{u}}$  based on evaluation at all mesh and time points in case A. Attempts to interpret the fitted models are then made, and several of these  $\mathbf{h}$ -fitted models are implemented for *a posteriori* tests.

We start with a Taylor series expansion in filter width including terms up to  $O(\bar{\Delta}^2)$ , which thus has 39 parameters per  $\tau_{ij}^r$  component (351 total):

$$g_{ij}(\bar{\mathbf{u}}) = a_{ijk}\bar{u}_k + \bar{\Delta}b_{ijkl}\frac{\partial\bar{u}_k}{\partial x_l} + \bar{\Delta}^2c_{ijklm}\frac{\partial^2\bar{u}_k}{\partial x_l\partial x_m}. \quad (32)$$

This expansion for two subgrid-scale stress components is shown in Fig. 15. We focus on components for which the deep-learning model learns asymmetries (discussed in Sec. IV), with  $\tau_{12}^r$  and  $\tau_{21}^r$  shown in Fig. 15 as an example, though we made analogous observations for all other components. The Taylor-series model (32) is able to reproduce the asymmetries learned by  $\mathbf{h}$  and, roughly, its variances. Other model forms, such as those discussed subsequently based on the Smagorinsky and Clark models, are unable to do so. There seem to be two aspects to the relative success of (32): (1) assigning separate parameters to asymmetric velocity gradient tensors (and not using the symmetric strain-rate tensor  $S_{ij}$ ) and (2) including all velocity derivatives in  $g_{ij}$ , not just those with the same  $i, j$  indices. One would expect that asymmetries in decoupled closures (for example, those based on the Boussinesq hypothesis) could be learned simply by including the asymmetric velocity derivatives. However, the improvements of including non-index-matching velocity derivatives reinforces the idea that  $\mathbf{h}$  can learn inter-PDE couplings, consistently with the overall discretization of the equations.

We can also infer properties of  $\mathbf{h}$  by fitting it with the functional forms of standard low-degree-of-freedom turbulence models. We first consider the constant-coefficient Smagorinsky model [17,18]

$$\tau_{ij}^{r,\text{Smag}} = -2C_s^2\bar{\Delta}^2|\bar{S}|(\bar{S}_{ij} - \frac{1}{3}\bar{S}_{kk}\delta_{ij}), \quad (33)$$




FIG. 16. Performance on case A of baseline,  $\mathbf{h}$ -fitted, and  $\bar{\mathbf{u}}$ -fitted models at  $t = 62.5t_{j0}$  and for  $\Delta_{16}$ : (a) the constant-coefficient Smagorinsky model ( $C_s = 0.163$ :  $\mathbf{h}$  fit;  $C_s = 0.118$ :  $\bar{\tau}_{\text{DNS}}^r$  fit), and (b) the Clark model (see text).

where  $|\bar{S}| = (2\bar{S}_{ij}\bar{S}_{ij})^{1/2}$  is the filtered strain-rate magnitude. The fitting parameter is  $\vec{\phi} = [C_s^2]$ . With all components of the model denoted  $\bar{\tau}^{r,\text{Smag}} = \mathbf{a}(\bar{\mathbf{u}})\vec{\phi}$ , we can obtain  $C_s$  by minimizing (31). For case A, this yields  $C_s = 0.163$  for  $\mathbf{h}(\bar{\mathbf{u}}; \vec{\theta})$  and  $C_s = 0.118$  for  $\bar{\tau}_{\text{DNS}}^r$ , both of which are comparable to common values. Simulation results for these candidate coefficients are shown in Fig. 16(a). The coefficient learned from  $\mathbf{h}$ ,  $C_s = 0.163$ , most accurately reproduces the mean flow of the three. The learned coefficients ( $C_s = 0.163$  and  $0.118$ ) also more correctly reproduce jet spreading (not shown), but not as accurately as the full neural-network model.

We also consider the model of Clark *et al.* [19], which is obtained via a Taylor-series expansion of  $\bar{\tau}^r$  in filter size  $\bar{\Delta}$ :

$$\begin{aligned} \tau_{ij}^{r,\text{Clark}} &= \tau_{ij}^{r,\text{Smag}} + \tau_{ij}^{r,\text{grad}} \\ &= \tau_{ij}^{r,\text{Smag}} + \frac{1}{12} \left( \frac{\bar{\Delta}^2}{\Delta} \frac{\partial \bar{u}_i}{\partial x_k} \frac{\partial \bar{u}_j}{\partial x_k} - \frac{\bar{\Delta}^2}{3} \frac{\partial \bar{u}_k}{\partial x_k} \frac{\partial \bar{u}_k}{\partial x_k} \delta_{ij} \right). \end{aligned} \quad (34)$$

The second term on its own is known as the gradient model, which, while accurate in *a priori* evaluations because it provides a good approximation for filter-scale perturbations in an ideally filtered DNS field, can cause numerical instabilities in *a posteriori* evaluation [32].

To fit coefficients with  $\vec{\phi} = [C_s^2, C_g]^T$ , we rewrite (34) as

$$\bar{\tau}^{r,\text{Clark}} = C_s^2 \mathbf{a}_{\text{Smag}}(\bar{\mathbf{u}}) + C_g \mathbf{a}_{\text{grad}}(\bar{\mathbf{u}}). \quad (35)$$

With  $\mathbf{a}(\bar{\mathbf{u}}) = [\mathbf{a}_{\text{Smag}}, \mathbf{a}_{\text{grad}}]$ , the model is  $\bar{\tau}^{r,\text{Clark}} = \mathbf{a}(\bar{\mathbf{u}})\vec{\phi}$ . Fitting to  $\mathbf{h}(\bar{\mathbf{u}})$  yields  $C_s = 0.151$  and  $C_g = 0.422$ , and fitting to  $\bar{\tau}_{\text{DNS}}^r$  yields  $C_s = 0.041$  and  $C_g = 1.288$ . Of these, only the fit to  $\mathbf{h}(\bar{\mathbf{u}})$  appears reasonable based on typical values. This is a surprise since one might expect that the coefficients could be learned directly from the filtered data. However, the same assumption is the reason *a priori* evaluation often incorrectly predicts a model's *a posteriori* performance. *A posteriori* results are shown in Fig. 16(b). With  $\mathbf{h}$ -learned coefficients, results are comparable to the unity-coefficient model, and, indeed, the DNS-fitted coefficients are inaccurate.

These results suggest that high-degree-of-freedom models can be interpreted in the context of prespecified functional forms. In particular, we observe that asymmetries in the modeled subgrid-scale stress can be codified using asymmetric velocity gradient tensors, and that  $\mathbf{h}$  might learn inter-PDE couplings between the different velocity components. However, low-degree-of-freedom attempts to interpret highly parametrized models will often be found wanting, and this can be seen in attempts to fit standard turbulence-model parameters to the neural-network form. While these can

produce better-performing fitted (or at least not seriously performance degrading) parameters, they do not match the *a posteriori* accuracy of the full model.

## VII. PREDICTED TURBULENCE STRUCTURE

The agreement of the deep-learning model prediction with the usual statistical measures (mean, turbulence stresses, energy spectrum) suggests an improved representation of realistic turbulence at the resolved scales. Whether or not this agreement translates to an improved representation of the turbulence structure, which would be important for extrapolation beyond the training metrics, is quantified with barycentric maps. These track the eigenvalue invariants of the normalized resolved Reynolds stress anisotropy  $a_{ij} = R_{ij}/R_{kk} - \delta_{ij}/3$  [33] and as such quantify resolved turbulence componentiality. Although an obvious downside for neural-network models is the lack of proof about realizability and boundedness, barycentric maps offer an empirical assessment of this: the lines connecting the vertices of such maps are physical bounds, and remaining within them is not guaranteed. One could envision that a deep-learning model might, despite a PDE constraint, sacrifice physical realism for a good fit to data; this is tested in this section. The degree to which a deep-learning closure recovers the ideal turbulence structure supports its extrapolative capacity and its additional adherence to the physics beyond reproduction of the training data.

Barycentric maps for the case A resolved turbulence and  $\Delta_{16}$  machine-learning model are shown in Figs. 17(a) and 17(b). All trajectories remain within the physical bounds, and this is also the case for the dual-jet configurations (not shown). The trajectories within the maps show the variation of componentiality in  $y$ . Key features of the direct numerical simulation [Fig. 17(a)] include a relative two-component condition in the outer shear layers ( $|y| \approx 4H_0$ ), predominantly plane strain (the line bridging three- and two-componentiality) within the shear layers, and three-component turbulence near the centerline ( $|y| \lesssim H_0$ ). Also shown are invariant maps of *a posteriori* large-eddy simulation using the dynamic Smagorinsky model, the  $F(\bar{\mathbf{u}}) = (\langle \mathbf{u} \rangle; \langle \bar{\mathbf{u}}' \bar{\mathbf{u}}' \rangle)$  closure, and the  $\mathbf{h}(\bar{\mathbf{u}})$ -fitted Clark model. Both the  $\Delta_{16}$   $\mathbf{h}(\bar{\mathbf{u}})$  states [Fig. 17(b)] and  $\Delta_8$  dynamic Smagorinsky states [Fig. 17(d)] show the general features of the direct numerical simulation, consistent with the ability of these models to reproduce other turbulence statistics. However, the coarse  $\Delta_{16}$  dynamic Smagorinsky model [Fig. 17(c)] fails to achieve the same three-componentiality for  $|y| \rightarrow 0$ .

Trajectories obtained using  $\mathbf{h}$  trained with only mean statistics [ $F(\bar{\mathbf{u}}) = (\langle \mathbf{u} \rangle; \langle \bar{\mathbf{u}}' \bar{\mathbf{u}}' \rangle)$ ] in Fig. 17(e) only recover the turbulence structure for  $|y| \lesssim H_0$ . This case fails to recover the two-component limiting conditions in the outer shear layers ( $|y|/H_0 \approx 4$ ), instead trending toward one-component turbulence. More training data from richer flows might remedy this [7].

Finally, invariant trajectories of the  $\mathbf{h}(\bar{\mathbf{u}})$ -fitted Clark model from Sec. VI are shown in Fig. 17(f). Despite its apparently good mean-flow prediction [Fig. 16(b)], the turbulence produced bears little resemblance to the direct numerical simulation data. Thus, this model is not expected to be extrapolative, efforts for which will likely benefit from better faithful representation of the turbulence structure in addition to any particular statistics. For example, subfilter models for reacting flows will almost certainly benefit from the capability to encode these dynamics [34,35]. As a step toward assessing, out-of-sample tests of the  $\mathbf{h}$ -fitted Clark model are presented along with the machine-learning model for scalar mixing in the following section.

## VIII. MIXING OF A PASSIVE SCALAR

In Sec. VA, it was shown that the high-dimensional neural-network model could, not surprisingly, fit the filtered direct numerical simulation subgrid-scale stresses through a direct mismatch loss function, though this fitted stress failed to provide a viable mode when subsequently used in a simulation. This exemplifies the risk of fitting without physical constraints, which guard against overfitting and unreliability for extrapolative prediction. Optimizing the model as embedded in the governing equation (Secs. VB and VC) provided a model that reproduced turbulence statistics for a class of similar flows, demonstrating some capacity for reliable extrapolation. Here, we further

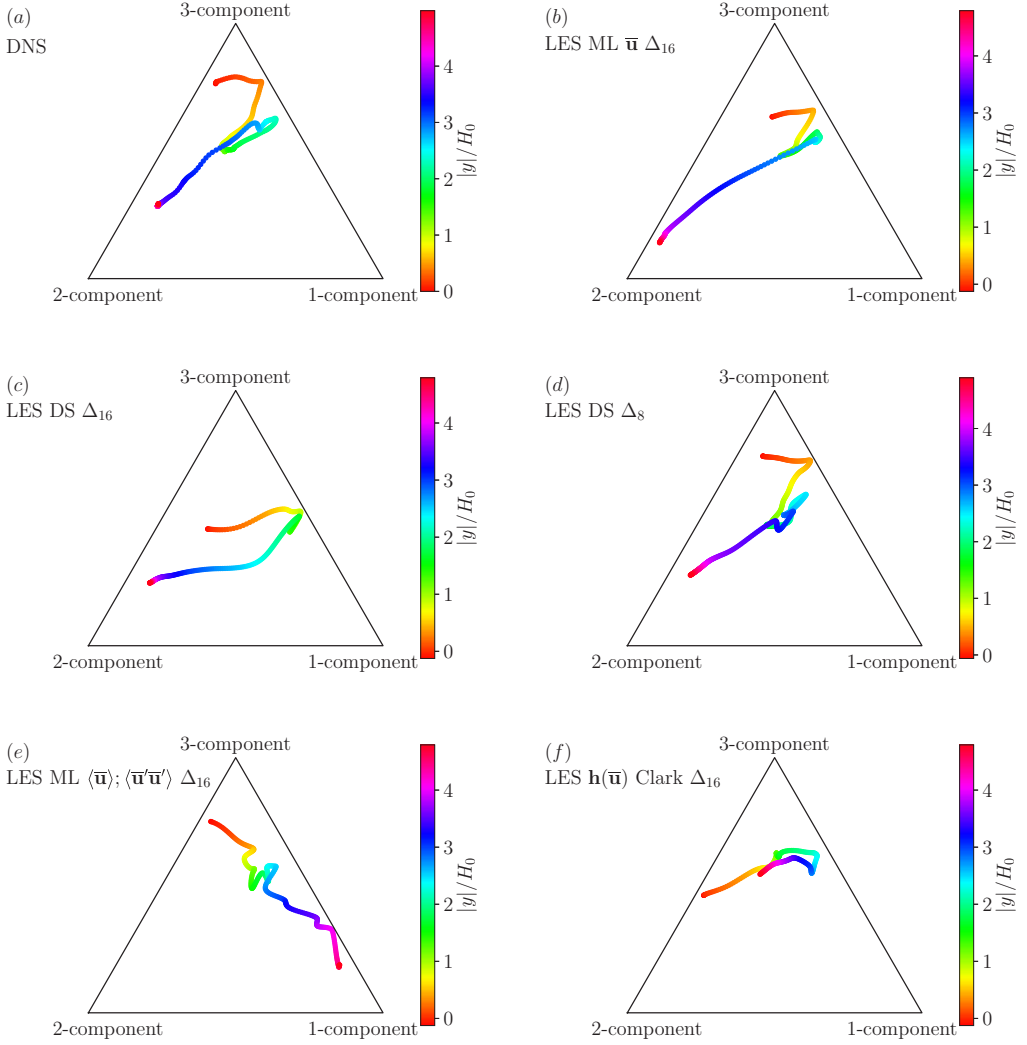


FIG. 17. Barycentric maps of normalized resolved Reynolds stress invariants at various cross-stream locations  $|y|/H_0$  for case A at  $t = 62.5t_{j_0}$ , with data averaged  $\pm y$ : (a) direct numerical simulation, (b) machine-learning model with  $\mathbf{h}(\bar{\mathbf{u}})$  (see Table II), (c), (d) the dynamic Smagorinsky model for  $\Delta_{16,8}$ , (e) training on averaged statistics  $F(\bar{\mathbf{u}}) = (\langle \mathbf{u} \rangle; \langle \bar{\mathbf{u}}' \bar{\mathbf{u}}' \rangle)$  (see Table II), and (f) the Clark model fitted to  $\mathbf{h}(\bar{\mathbf{u}})$  (see Sec. VI).

evaluate the extrapolative capacity of the model by assessing passive scalar  $\xi$  mixing predictions. Filtering (10) yields

$$\frac{\partial \bar{\xi}}{\partial t} + \bar{u}_j \frac{\partial \bar{\xi}}{\partial x_j} = D \frac{\partial^2 \bar{\xi}}{\partial x_j^2} - \frac{\partial}{\partial x_j} (\overline{u_j \xi} - \bar{u}_j \bar{\xi}), \quad (36)$$

where  $D = 0$ . Only numerical diffusion was active in both the direct numerical simulation and the large-eddy simulation, and then only for the scalar advection. The second term on the right-hand side of (36) is the divergence of the unclosed residual scalar flux  $F_j^\xi = \overline{u_j \xi} - \bar{u}_j \bar{\xi}$ . We model  $F_j^\xi$  using the gradient-diffusion hypothesis, in which the eddy diffusivity is obtained using a scalar analog of the dynamic Smagorinsky model [4].

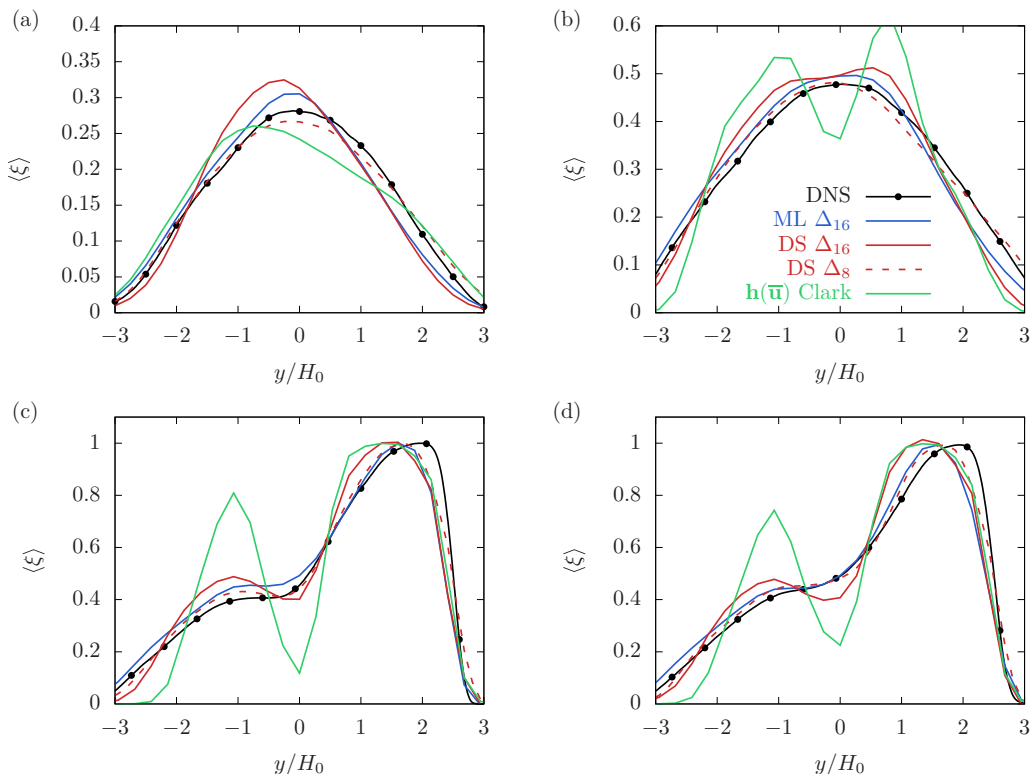


FIG. 18. Mean passive scalar for the deep-learning (ML), dynamic Smagorinsky (DS), and  $\mathbf{h}$ -fitted Clark model for  $\bar{\tau}^r$  for  $\Delta_{16}$  for (a) case A at  $t = 62.5t_{j_0}$ , (b) case B at  $t = 50t_{j_0}$ , (c) case C at  $t = 42.5t_{j_0}$ , and (d) case D at  $t = 42.5t_{j_0}$ . For comparison,  $\Delta_8$  DS results are also shown.

The predicted mean scalar profiles for the four jet cases, simulated with  $\bar{\tau}^r$  obtained from either the deep-learning model, the dynamic Smagorinsky model, or the  $\mathbf{h}$ -fitted Clark model, are evaluated in Fig. 18. The deep-learning model outperforms the others in all four cases. This is particularly evident in the double-jet cases, for which the same  $\Delta_{16}$  using dynamic Smagorinsky underpredicts jet spreading and merging (see Fig. 11). This result is comparably accurate to dynamic-Smagorinsky predictions for  $\Delta_8$ . Conversely, the  $\mathbf{h}$ -fitted Clark model, with coefficients obtained for case A, performs poorly for the out-of-sample dual-jet cases.

The results in this section provide evidence of potentially improved predictive accuracy of coupled-PDE simulations using the deep-learning model. However, we note that the more complicated case of two-way-coupled PDEs has yet to be investigated. Variable-density, nonreacting scalar mixing, with a simple equation of state, is an obvious next step in analyzing two-way couplings.

## IX. MESH DEPENDENCE

We finally consider the effect of different grid resolutions on the learned closure  $\mathbf{h}$ . The *a priori* convergence of  $\mathbf{h}$  outputs with grid refinement is first shown. Then, *a posteriori* performance of the model on finer grids, for which the model was not explicitly trained, is assessed.

### A. Convergence with grid refinement

The preceding sections center around implicitly filtered large-eddy simulation  $\Delta_{\text{LES}}/\Delta_{\text{DNS}} = \bar{\Delta}/\Delta_{\text{DNS}}$ , where  $\Delta_{\text{DNS}}$  and  $\Delta_{\text{LES}}$  distinguish the fine direct numerical simulation mesh and large-eddy




FIG. 19. Relative difference  $\varepsilon_{\Delta x}$  from (37) showing mesh dependence of  $\tau_{12}^r$  for case A with  $\bar{\Delta} = \Delta_{16}$  training for the indicated models (Sec. VB). The corresponding convergence of  $\partial u / \partial y$  is shown for reference.

simulation meshes. For constant  $\bar{\Delta}$ , we now consider *a priori* convergence of subgrid-scale closure approximations with refinement of  $\Delta_{\text{LES}}$ . This enables estimation of discretization errors and the degree to which  $\mathbf{h}$  deviates from the expected convergence of the differential approximation of the underlying numerical scheme.

For constant  $\bar{\Delta}$ , standard models for  $\bar{\tau}^r$  converge as  $O(\Delta_{\text{LES}}^p)$  for order- $p$  spatial discretizations. Here,  $p = 2$ . The deep-learning model is not constrained by this: as described in Sec. IV A, although its inputs are the velocity first and second derivatives at adjacent mesh points, it is nonlinear and so not constrained to show second-order convergence, at least for the resolutions we test. In theory, it can compensate for coarse-grid truncation error in *a posteriori* large-eddy simulation, for which numerical errors appear as additional, unclosed terms when compared to trusted target data [11].

Velocity gradients are computed from filtered direct numerical simulation data for  $\bar{\Delta} = \Delta_{16}$ . The finite-difference approximations are evaluated for  $\Delta_{\text{LES}} / \Delta_{\text{DNS}} = 16, 8, 4, 2, 1$  using the same staggered variable placement as the solver (see Sec. II B).

To quantify the mesh convergence of the different closures, a relative difference [36] is computed as

$$\varepsilon_{\Delta x} \equiv \ell_2(\bar{\tau}_{\Delta x}^r - \bar{\tau}_{\Delta x/2}^r), \quad (37)$$

where  $\ell_2$  is the 2-norm and  $\bar{\tau}_{\Delta x}^r$  and  $\bar{\tau}_{\Delta x/2}^r$  are closure approximations evaluated for mesh spacings  $\Delta_{\text{LES}}$  and  $\Delta_{\text{LES}}/2$ .

Behavior is shown in Fig. 19 for both the deep-learning and Smagorinsky models. Also shown, for reference, is the cross-stream  $u$ -velocity gradient evaluated on  $\Delta_{\text{LES}}$ . Both the velocity gradients and the Smagorinsky model  $\bar{\tau}^r$  appear to follow  $O(\Delta_{\text{LES}}^2)$ , as expected. Conversely, the deep-learning model only approaches linear  $O(\Delta_{\text{LES}})$  behavior over the available range of mesh densities. This is presumably a consequence of the more complex and nonlinear properties that provide its excellent *a posteriori* performance.

In the sense of optimal LES [3], discrepancies between the formal convergence of unclosed LES solutions (or trusted DNS data) and the ideal closed LES solution are expected. The ideal LES closure, which achieves the upper limit of accuracy with respect to trusted data, must approximate Navier-Stokes physics that the filter operator (12) maps to its null space. These are, of course, not representable in an LES solution space even with explicit filtering [3].

Figure 19 may be understood in this context by considering an optimal closure model

$$\bar{\tau}_{\Delta x}^r = \mathbf{P}\bar{\tau}_{\text{ideal}}^r + \varepsilon_{\text{P}}(\Delta x) + c\Delta x^p, \quad (38)$$






FIG. 20. Predicted streamwise mean velocity for the single-jet case (A) for grid-filter ratios  $\bar{\Delta}/\Delta_{\text{DNS}}$ . (a) Instantaneous-velocity training; (b) mean-flow and Reynolds stress training. Both models were trained only on  $\Delta_{16}$ .

where  $\tau_{\text{ideal}}^r$  is the ideal LES closure model,  $\mathbf{P}$  is a (nonunique) projection to the truncated  $\Delta x$  LES solution space,  $\varepsilon_{\mathbf{P}}(\Delta x)$  is the dynamical error in the Navier-Stokes solution due to the projection, and  $c\Delta x^p$  is the spatial discretization error of the model inputs. Since  $\mathbf{P}$  is linear and invertible, the relative difference (37) eliminates  $\mathbf{P}\bar{\tau}_{\text{ideal}}^r$  leaving only the error terms. Traditional LES closures do not model  $\varepsilon_{\mathbf{P}}(\Delta x)$  and so converge as  $O(\Delta x^p)$ . Conversely, the *a posteriori* trained (6) deep-learning model does not, even though its inputs are  $O(\Delta x^p)$  accurate, which suggests that it at least partially accounts for  $\varepsilon_{\mathbf{P}}(\Delta x)$ .

### B. Performance on out-of-sample meshes

The model is trained for a particular coarse resolution, for which, at least for established models, discretization errors are known to couple with model errors. The performance of the model suggests that it also compensates for discretization errors, which was assessed previously for isotropic turbulence [11]. Hence, it comes with no *a priori* assurance that it will be as accurate away from resolutions for which it was trained. We assess this by applying the  $\mathbf{h}$  closure to finer grids than the training  $\Delta_{16}$ . Figure 20 shows the performance over a range of  $\Delta_N$  using both instantaneous-velocity training and mean-statistics training. The instantaneous-velocity-trained model in Fig. 20(a) reproduces the mean profile on  $\Delta_4$  nearly as well as for the  $\Delta_{16}$  trained case, but not so well on  $\Delta_8$ . The mean-flow training is more robust, as seen in Fig. 20(b). However, the implicit filter definition used for *a posteriori* LES is still a potential source of error. The mean-flow training [Fig. 20(b)] is presumably less sensitive to discrepancies between this and the *a priori* explicit filter definition, although further testing is certainly warranted.

Finally, we note that a deep-learning model trained only for isotropic turbulence [11] performed reasonably well when tested out of sample on the jet case A, although it was not as accurate as the deep-learning models in Fig. 20. The model is at least stable and reasonably accurate when applied to different grids and potentially to different flows. However, its overall extrapolative capacity, for example, to different flow phenomena, is likely reduced.

## X. ADDITIONAL DISCUSSION

In some sense, it is unsurprising that a high-degree-of-freedom model can outperform established models with few if any parameters, at least for the modest degree of out-of-sample extrapolation we attempted here. However, determining viable parameters for this is challenging. The adjoint-based PDE-coupled training procedures used here seem to perform well for finding parameters sufficient

for robust and extrapolative models. The predictive agreement with trusted data for coarse meshes is striking. Still, given the number of parameters required and the challenging optimization procedures needed, this should not be taken as an unqualified improvement. Indeed, a fair way to view results is recognition of just how well dynamic Smagorinsky, as an example of established models, does with neither significant parametrization nor a complicated training regimen. In closing, we revisit some of the motivations beyond the basic goal of increasing accuracy for fixed mesh resolution.

One particular hope is that this framework will facilitate representation of turbulence when it is coupled with additional physics, especially in cases where this coupling lacks models of the fidelity available for incompressible turbulence. The final demonstration, where the model is shown to also extrapolate to represent scalar mixing, is an example of this. Similarly, training of a robust model based only on mean statistics provides an avenue for predictions even when part of the physics (e.g., flame dynamics) might be insufficiently described to craft an explicit model. Generalizing high-degree-of-freedom models to out-of-sample flow conditions or new physics couplings is a main target for this procedure.

Even for inert incompressible turbulence, without the complexity of additional physics, we also anticipate potential synergies between the highly parametrized approach we introduce and established models and modeling approaches. Our efforts to interpret the successful neural-network model by fitting the parameters of low-degree-of-freedom models are a first step toward using the successful working representations to improve conventional models. A related opportunity is the use of *a posteriori* successful neural-network models to aid coefficient discovery in moderately parametrized models, based on physical reasoning, such as those that aim to codify nonlinearly coupled multiphysical interactions. We anticipate a fruitful middle ground between fitting accuracy and generalization where the two modeling approaches can complement and augment each other.

An outstanding issue, which is shared with all turbulence subgrid-scale modeling, is the interplay of numerical errors associated with the coarse mesh resolution and the model closures. As trained through solutions of the discrete flow equations and their adjoint, the current deep-learning models account for errors introduced by the coarse discretizations, which is important for its predictive success. A potential drawback is that it risks linking the trained model with the specific mesh spacing (and discretization method) for which it is trained. In the present case, we demonstrated that this is not a severe limitation: the model is robustly predictive for different meshes, though accuracy is indeed reduced. Looking forward, it can be anticipated that systematically training for a range of mesh resolutions might provide improved predictions. We can also speculate that the training on the case A jet, which spreads by over a factor of 2 during the training period, might imbue the learned model with some robustness in this regard since it exposes the model to different nominal resolutions through the evolution of the flow. For example, if the development were exactly self-similar, then the model requirements for the  $\Delta_{16}$  mesh at some later time would match those for the  $\Delta_8$  mesh at the earlier time when the jet has half the later thickness. Systematic studies of extrapolative capacity to different resolutions, mesh configurations, numerical methods, and flows are all warranted.

Finally, it is important to address a general challenge faced when applying adjoint methods to turbulence. It is well understood that the chaos of turbulence will, for long enough times, lead to divergence of sensitivity fields, which diminishes the utility of computed sensitivities. This is not a challenge for the short-time adjoint solutions used in most of the present demonstrations, and we anticipate that many applications of this method will not suffer from this because of the character of the problem: a subgrid-scale model to provide a stress correction for the local instantaneous resolved-scale turbulence. Since the large turbulence scales are relatively local in space and time, the sensitivity will not accumulate significantly. In general, however, this might limit applications with limited data availability, especially if it is sensitive to turbulence that is significantly distant in space and time. Further work will be required to understand how the Lyapunov divergence of turbulence might hinder the adjoint-based training. We do note that the adjoint equations remained sufficiently well behaved when we trained and simulated over the entire time history of the jet flow when training for mean statistics.

## ACKNOWLEDGMENTS

This material is based in part upon work supported by the Department of Energy, National Nuclear Security Administration, under Award No. DE-NA0002374. This research is part of the Blue Waters sustained-petascale computing project, which is supported by the National Science Foundation (Awards No. OCI-0725070 and No. ACI-1238993) and the State of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications.

- 
- [1] C. Meneveau and J. Katz, Scale-invariance and turbulence models for large-eddy simulation, *Annu. Rev. Fluid Mech.* **32**, 1 (2000).
  - [2] S. Ghosal, Mathematical and physical constraints on large-eddy simulation of turbulence, *AIAA J.* **37**, 425 (1999).
  - [3] J. A. Langford and R. D. Moser, Optimal LES formulations for isotropic turbulence, *J. Fluid Mech.* **398**, 321 (1999).
  - [4] M. Germano, U. Piomelli, P. Moin, and W. H. Cabot, A dynamic subgrid-scale eddy viscosity model, *Phys. Fluids* **3**, 1760 (1991).
  - [5] D. K. Lilly, A proposed modification of the Germano subgrid-scale closure method, *Phys. Fluids* **4**, 633 (1992).
  - [6] B. Steiner, Z. DeVito, S. Chintala, S. Gross, A. Paszke, F. Massa, A. Lerer, G. Chanan, Z. Lin, E. Yang, A. Desmaison, A. Tejani, A. Kopf, J. Bradbury, L. Antiga, M. Raison, N. Gimelshein, S. Chilamkurthy, T. Killeen, L. Fang, and J. Bai, Pytorch: An imperative style, high-performance deep learning library, in *Advances in Neural Information Processing Systems 32* (Curran Associates, Inc., 2019), p. 8024.
  - [7] J. Ling, A. Kurzawski, and J. Templeton, Reynolds averaged turbulence modeling using deep neural networks with embedded invariance, *J. Fluid Mech.* **807**, 155 (2016).
  - [8] M. Gamahara and Y. Hattori, Searching for turbulence models by artificial neural network, *Phys. Rev. Fluids* **2**, 054604 (2017).
  - [9] A. Beck, D. Flad, and C. D. Munz, Deep neural networks for data-driven LES closure models, *J. Comput. Phys.* **398**, 108910 (2019).
  - [10] M. P. Brenner, J. D. Eldredge, and J. B. Freund, A perspective on machine learning for advancing fluid mechanics, *Phys. Rev. Fluids* **4**, 100501 (2019).
  - [11] J. Sirignano, J. F. MacArt, and J. B. Freund, DPM: A deep learning PDE augmentation method with application to large-eddy simulation, *J. Comput. Phys.* **423**, 109811 (2020).
  - [12] J. R. Holland, J. D. Baeder, and K. Duraisamy, Towards integrated field inversion and machine learning with embedded neural networks for RANS modeling, in *AIAA SciTech Forum* (AIAA, Reston, VA, 2019), p. 1884.
  - [13] E. J. Parish and K. Duraisamy, A paradigm for data-driven predictive modeling using field inversion and machine learning, *J. Comput. Phys.* **305**, 758 (2016).
  - [14] G. Novati, H. L. de Laroussilhe, and P. Koumoutsakos, Automating Turbulence Modeling by Multi-Agent Reinforcement Learning, *Nat. Machine Intelligence* **3**, 87 (2021).
  - [15] C. Rackauckas, Y. Ma, J. Martensen, C. Warner, K. Zubov, R. Supekar, D. Skinner, A. Ramadhan, and A. Edelman, Universal Differential Equations for Scientific Machine Learning, [arXiv:2001.04385](https://arxiv.org/abs/2001.04385).
  - [16] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, Automatic differentiation in machine learning: A survey, *J. Mach. Learn. Res.* **18**, 1 (2018).
  - [17] J. Smagorinsky, General circulation experiments with the primitive equations I. The basic experiment, *Mon. Weather Rev.* **91**, 99 (1963).
  - [18] R. S. Rogallo and P. Moin, Numerical simulation of turbulent flows, *Annu. Rev. Fluid Mech.* **16**, 99 (1984).
  - [19] R. A. Clark, J. H. Ferziger, and W. C. Reynolds, Evaluation of subgrid-scale models using a fully simulated turbulent flow, *J. Fluid Mech.* **91**, 1 (1979).

- [20] F. H. Harlow and J. E. Welch, Numerical calculation of time-dependent viscous incompressible flow of fluid with free surface, *Phys. Fluids* **8**, 2182 (1965).
- [21] G.-S. Jiang and C.-W. Shu, Efficient implementation of weighted ENO schemes, *J. Comput. Phys.* **126**, 202 (1996).
- [22] J. Kim and P. Moin, Application of a fractional-step method to incompressible Navier–Stokes equations, *J. Comput. Phys.* **59**, 308 (1985).
- [23] O. Desjardins, G. Blanquart, G. Balarac, and H. Pitsch, High order conservative finite difference scheme for variable density low Mach number turbulent flows, *J. Comput. Phys.* **227**, 7125 (2008).
- [24] J. B. Freund, S. K. Lele, and P. Moin, Compressibility effects in a turbulent annular mixing layer. Part 2. Mixing of a passive scalar, *J. Fluid Mech.* **421**, 269 (2000).
- [25] C. Pantano and S. Sarkar, A study of compressibility effects in the high-speed turbulent shear layer using direct simulation, *J. Fluid Mech.* **451**, 329 (2002).
- [26] R. Knaus and C. Pantano, On the effect of heat release in turbulence spectra of non-premixed reacting shear layers, *J. Fluid Mech.* **626**, 67 (2009).
- [27] E. R. Hawkes, O. Chatakonda, H. Kolla, A. R. Kerstein, and J. H. Chen, A petascale direct numerical simulation study of the modeling of flame wrinkling for large-eddy simulations in intense turbulence, *Combust. Flame* **159**, 2690 (2012).
- [28] J. F. MacArt, T. Grenga, and M. E. Mueller, Effects of combustion heat release on velocity and scalar statistics in turbulent premixed jet flames at low and high Karlovitz numbers, *Combust. Flame* **191**, 468 (2018).
- [29] D. R. Miller and E. W. Comings, Force-momentum fields in a dual-jet flow, *J. Fluid Mech.* **7**, 237 (1960).
- [30] X. Glorot and Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics* (Microtome, Brookline, MA, 2010), p. 249.
- [31] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016).
- [32] B. Vreman, B. J. Geurts, and H. Kuerten, Large-eddy simulation of the temporal mixing layer using the Clark model, *Theor. Comput. Fluid Dyn.* **8**, 309 (1996).
- [33] S. Banerjee, R. Krahl, F. Durst, and C. Zenger, Presentation of anisotropy properties of turbulence, invariants versus eigenvalue approaches, *J. Turbul.* **8**, N32 (2007).
- [34] J. O’Brien, C. A. Z. Towery, P. E. Hamlington, M. Ihme, A. Y. Poludnenko, and J. Urzay, The cross-scale physical-space transfer of kinetic energy in turbulent premixed flames, *Proc. Combust. Inst.* **36**, 1967 (2017).
- [35] J. F. MacArt and M. E. Mueller, Scaling and modeling of heat-release effects on subfilter turbulence in premixed combustion, *Center for Turbulence Research Proceedings of the Summer Program 2018* (Stanford University Press, Stanford, CA, 2018), p. 299.
- [36] J. F. MacArt and M. E. Mueller, Semi-implicit iterative methods for low Mach number turbulent reacting flows: Operator splitting versus approximate factorization, *J. Comput. Phys.* **326**, 569 (2016).