

Data-driven subgrid-scale modeling of forced Burgers turbulence using deep learning with generalization to higher Reynolds numbers via transfer learning

Cite as: Phys. Fluids **33**, 031702 (2021); <https://doi.org/10.1063/5.0040286>

Submitted: 12 December 2020 . Accepted: 01 February 2021 . Published Online: 01 March 2021

 Adam Subel,  Ashesh Chattopadhyay,  Yifei Guan (关逸飞), et al.



View Online



Export Citation



CrossMark

ARTICLES YOU MAY BE INTERESTED IN

[Physics guided machine learning using simplified theories](#)

Physics of Fluids **33**, 011701 (2021); <https://doi.org/10.1063/5.0038929>

[Lattice-Boltzmann modeling of a turbulent bluff-body stabilized flame](#)

Physics of Fluids **33**, 031701 (2021); <https://doi.org/10.1063/5.0038089>

[Deep learning methods for super-resolution reconstruction of turbulent flows](#)

Physics of Fluids **32**, 025105 (2020); <https://doi.org/10.1063/1.5140772>

Physics of Fluids

SPECIAL TOPIC: Flow and Acoustics of Unmanned Vehicles

Submit Today!

Data-driven subgrid-scale modeling of forced Burgers turbulence using deep learning with generalization to higher Reynolds numbers via transfer learning

Cite as: Phys. Fluids **33**, 031702 (2021); doi: [10.1063/5.0040286](https://doi.org/10.1063/5.0040286)

Submitted: 12 December 2020 · Accepted: 1 February 2021 ·

Published Online: 1 March 2021



View Online



Export Citation



CrossMark

Adam Subel,¹ Ashesh Chattopadhyay,¹ Yifei Guan (关逸飞),¹ and Pedram Hassanzadeh^{1,2,a)}

AFFILIATIONS

¹Department of Mechanical Engineering, Rice University, Houston, Texas 77005, USA

²Department of Earth, Environmental, and Planetary Sciences, Rice University, Houston, Texas 77005, USA

^{a)} Author to whom correspondence should be addressed: pedram@rice.edu

ABSTRACT

Developing data-driven subgrid-scale (SGS) models for large eddy simulations (LESs) has received substantial attention recently. Despite some success, particularly in *a priori* (offline) tests, challenges have been identified that include numerical instabilities in *a posteriori* (online) tests and generalization (i.e., extrapolation) of trained data-driven SGS models, for example, to higher Reynolds numbers. Here, using the stochastically forced Burgers turbulence as the test-bed, we show that deep neural networks trained using properly pre-conditioned (augmented) data yield stable and accurate *a posteriori* LES models. Furthermore, we show that transfer learning enables accurate/stable generalization to a flow with $10\times$ higher Reynolds number.

Published under license by AIP Publishing. <https://doi.org/10.1063/5.0040286>

Due to their high computational cost, the direct numerical simulation (DNS) of turbulent flows will remain out of reach for many real-world applications in the foreseeable future. As a result, the need for parameterization of subgrid-scale (SGS) processes in coarse-resolution models such as large eddy simulation (LES) continues in various areas of science and engineering.^{1,2} In recent years, there has been substantial interest in applications of deep learning for data-driven modeling of turbulent flows,^{3–13} including for developing data-driven SGS parameterization (DDP) models.^{14–27} In many of these studies, the goal is to learn the relationship between the filtered variables and SGS terms in high-fidelity data (e.g., DNS data), and use this DDP model in LES. *A priori* tests in some of these studies^{17–19,25} have shown that such a non-parametric approach can yield DDP models that capture important physical processes (e.g., energy backscatter^{28,29}) beyond the simple diffusion process that is represented in canonical physics-based SGS models such as Smagorinsky and dynamic Smagorinsky (DSMAG).^{30–32} However, these studies have also reported that *a posteriori* (i.e., online) LES tests, in which the DDP model is coupled to a coarse-resolution Navier-Stokes solver, show numerical instabilities or lead to physically unrealistic flows.^{17–19,25,26} As a remedy, often *ad hoc* post-processing steps of the DDP models'

outputs are introduced, e.g., to remove backscattering or to attenuate the SGS feedback into the numerical solver. Usually, such post-processing steps substantially take away the advantages gained from using deep learning. As a result, numerical instabilities remain a major obstacle to broadening the applications of LES with DDP models.

Another major concern with DDP models is their (in)ability to accurately generalize beyond the flow they are trained for, particularly to flows that have higher Reynolds numbers (Re). However, such extrapolations are known to be challenging for neural networks.^{23,33} Some degree of generalization is essential for building robust and trustworthy LES models with DDP. Furthermore, given that high-fidelity data from often-expensive simulations (e.g., DNS) are needed to train DDP models, some capability to extrapolate to higher Re makes such DDP models much more practically useful.

In this paper, with a particular focus on the issues of stability and generalization, we use a deep artificial neural network (ANN) to develop a DDP model for stochastically forced Burgers turbulence. The forced Burgers equation is³⁴

$$\frac{\partial u}{\partial t} + \frac{1}{2} \frac{\partial(uu)}{\partial x} = \nu \frac{\partial^2 u}{\partial x^2} + F, \quad (1)$$

where u is velocity, $\nu = 1/Re$, and F is a stochastic forcing (defined later). The domain is periodic with length L . Despite being one-dimensional, the presence of strongly nonlinear local regions in the form of shocks, often multiple shocks [Fig. 1(a)], makes Burgers turbulence a complex and challenging system, which has been used as the test-bed in various SGS and reduced-order modeling studies.^{34–40} $F(x, t)$ is defined as³⁴

$$F = \sum_{k=1}^3 \frac{\alpha_k A}{\sqrt{20k\Delta t}} \cos\left(2\pi\left(\frac{kx}{L} + \Phi_k\right)\right), \quad (2)$$

where k , Δt , and A are the wavenumber, time step, and forcing amplitude, respectively. Φ_k and α_k are real, random numbers. To develop the LES model, we spatially filter Eq. (1) to obtain

$$\frac{\partial \bar{u}}{\partial t} + \frac{1}{2} \frac{\partial(\bar{u}\bar{u})}{\partial x} = \nu \frac{\partial^2 \bar{u}}{\partial x^2} + \bar{F} + \Pi, \quad (3)$$

with SGS term

$$\Pi = -\frac{1}{2} \frac{\partial}{\partial x} (\overline{uu} - \bar{u}\bar{u}). \quad (4)$$

Here, we use a box filter;¹ explorations with Gaussian and sharp spectral filters yield the same findings and conclusions. Overbars indicate filtered (and coarse-grained to LES resolution) variables. Note that the difference between F and $\bar{F}$ is negligible. Our aim is to train an ANN




FIG. 1. A sample profile and statistics of the stochastically forced Burgers turbulence (from DNS data at $Re = Re_c$). (a) u showing three distinct shocks. (b) The KE spectrum, showing the inertial range. (c) PSD, as a function of frequency ω , showing chaotic behavior.

to learn Π as a function of $\bar{u}$ in the DNS data, and then use this DDP model as a closure in Eq. (3).

We define a setup, referred to as “control” and indicated with subscripts “c,” with the following parameters (identical to those used by Dolapchiev *et al.*³⁴) $L = 100$, $\nu = 0.02$, and $A = \sqrt{2}/100$. Φ_k and α_k are drawn randomly from $\mathcal{N}(0, 1)$ every $20\Delta t$ to update F . To obtain the DNS data, which are treated as the “truth,” Eq. (1) is integrated using a pseudo-spectral solver with 1024 Fourier modes and time step $\Delta t = 0.01$. (The second-order Adams-Bashforth and Crank-Nicholson methods are used for time integration.) Figure 1 shows a sample profile of $u(x)$, and the kinetic energy (KE) spectrum and power spectral density (PSD) of the flow. To perform LES, Eq. (3) with the DDP model of $\Pi(\bar{u})$ is integrated using the same pseudo-spectral solver but with 128 Fourier modes and time step $20\Delta t$. At this spatial resolution, the filtered velocity field accounts for $\sim 81\%$ of the total KE of the flow (from DNS), conforming with the commonly used ratio.¹ Also, note that the spatial and temporal resolutions of the LES solver are, respectively, $8\times$ and $20\times$ lower than those of the DNS solver, substantially reducing the computational cost.

The schematic of LES with DDP is shown in Fig. 2(a). Next we present the details of the ANN and the training data/procedure. We use a multilayer perceptron ANN⁴¹ to develop the DDP model. This ANN is unidirectional (information only passes in one direction from input to output) and is fully connected between the layers. The ANN is trained, i.e., all learnable parameters of the network (weights and




FIG. 2. (a) The schematic of the LES with DDP model. With normalized $\bar{u}(x, t)$ as input, the trained ANN predicts Π , which is then de-normalized and used in Eq. (3) to compute $\bar{u}(x, t + 20\Delta t)$, and the cycle continues. (b) The pre-conditioning step to augment the training data by adding random shifts in x to produce spatially diverse samples from a relatively small DNS dataset.

biases, collectively represented by θ) are computed, by minimizing the mean square-error $MSE = \sum_{i=1}^M \|\text{ANN}(\tilde{u}_i; \theta) - \tilde{\Pi}_i\|_2^2 / M$. Here, M is the number of training samples, $\|\cdot\|_2$ is the L_2 norm, $\tilde{u}$ and $\tilde{\Pi}$ are calculated from DNS data, and $\tilde{\cdot}$ indicates pre-conditioned (augmented) training data (discussed shortly). The best network architecture, found based on extensive trial and error using MSE , consists of an input layer, six hidden layers with 250 nodes each, and a linear output layer. On all but the final layer, the swish activation function⁴² is used. Overall, the ANN has 394 640 trainable parameters.

Our first attempts to train the DDP model with $M = O(10^5)$ resulted in inaccurate Π terms in *a priori* tests and unstable LES with DDP in *a posteriori* tests. Further analysis showed that the problem is due to the fact that the SGS dynamics and thus the Π terms in Burgers turbulence are highly localized around the shocks,³⁷ which as explained below, leads to overfitting, i.e., poor generalization of ANN (at the same Re) beyond the training set. Shocks are persistent and can remain fairly stationary for many time steps, which can lead to small or near-zero Π terms in some regions of the domain that do not experience shocks throughout the training set. The ANN trained on such a dataset will predict $\Pi \approx 0$ in those regions no matter what the inputted $\tilde{u}$ is during (*a priori* or *a posteriori*) tests. Note that by design, the flow during training could be very different, in terms of the location of shocks and their evolution, from the flow during testing. (Though the training and testing sets have the same Re , the latter is chosen from an independent DNS run or from a time window far from the time window of the training set.) Of course, this overfitting problem can be resolved by using a much larger training set that contains a sufficient number of samples of shock waves occurring in all regions; however, such large training sets are often unavailable. Here, we propose a simple strategy, based on pre-conditioning the training samples, to overcome this problem without the need for a larger dataset.

As shown in Fig. 2(b), a random shift η , drawn from the uniform distribution $\mathcal{U}(0, L)$, is added to x for each input-output pair $(\tilde{u}, \tilde{\Pi})$

$$\tilde{u}(x, t) = \bar{u}(x - \eta, t) \text{ and } \tilde{\Pi}(x, t) = \Pi(x - \eta, t). \quad (5)$$

The periodicity in x is used when $x - \eta < 0$. It should be noted that this type of artificially enhancing the richness of information inside the training set is commonly used in the machine learning community and is called data augmentation.⁴³ For example, in processing of natural images, data augmentation generally involves artificially enhancing the training set by rotating, mirroring, or cropping images. Here, we have exploited the periodicity of x to introduce a physically meaningful augmentation, which allows us to enrich the information of the localized flow and SGS terms around shock waves in the training set without the need for a longer DNS dataset. Finally, as is common practice in machine learning, the input $\tilde{u}$ and output $\tilde{\Pi}$ samples are separately normalized (through removing the mean and dividing by the standard deviation).

The pre-conditioned input-output pairs $(\tilde{u}, \tilde{\Pi})$ are used to train the ANN. As shown next, the DDP model with an ANN trained using augmented data leads to accurate Π terms in *a priori* tests and stable and accurate LES models in *a posteriori* tests without the need for any post-processing of the trained ANN or its output [with the exception of de-normalizing the predicted Π ; see Fig. 2(a)]. We have used $M = 5 \times 10^5$ samples for training and another (independent) 5×10^4 samples for validation from a DNS run at $Re = Re_c$. For testing, we have used data from the same run but $5 \times 10^4 \Delta t$ separated from the

training/validation sets as well as data from two other independent DNS runs at $Re = Re_c$.

We examine the performance of the LES with DDP in a *posteriori* (online) tests to assess both accuracy (of the SGS modeling) and stability of the hybrid model. Given that the numerical solution of Eq. (3) blows up without any SGS modeling (i.e., with $\Pi = 0$), we use a conventional SGS scheme, DSMAG,⁴⁴ as the baseline. Figures 3(a) and 3(b) show the spectrum and the probability density function (PDF) of the Π terms predicted by DDP and DSMAG compared against those of the filtered DNS (FDNS), which is treated as the truth. Both panels show that the statistics of Π predicted by DDP closely follow those of the truth at any k and even at the tails of the PDF. Furthermore, both




FIG. 3. Statistics of the resolved flow $\bar{u}$ and SGS term Π calculated using results from a *posteriori* tests at $Re = Re_c$. The training and testing data are both at $Re = Re_c$. (a) Spectrum of Π , denoted as $\hat{\Pi}(k)$. The spectrum for FDNS agrees with those reported in previous studies of Burgers turbulence.³⁶ (b) PDF of Π . (c) Spectrum of KE. The curl up in KE around the maximum resolved k of LES is a common feature of spectral LES solvers applied to Burgers turbulence.^{38,39,46} In (a)–(c), each curve is produced using 3×10^5 sequential samples that are $20\Delta t$ apart. (d) PDF of $\bar{u}$ computed using a kernel estimator.⁴⁵ Inset panels in (d) show the zoomed-in left and right tails. Shading shows uncertainty as ± 1 standard deviation obtained from bootstrapping three independent LES or DNS runs that are combined (each providing 3×10^5 samples as before). In (b) and (d), σ is the variable's standard deviation.

panels show that DDP outperforms DSMAG in modeling the statistics of the SGS term (Π). The better performance of DDP is clearly seen at high and low k in (a) and beyond ± 1 standard deviation in (b). Note that the difference between the Π 's PDFs from FDNS and DSMAG (DDP) is (is not) statistically significant at 95% confidence level based on both Kolmogorov-Smirnov, KS, and Kullback-Leibler divergence, KL, tests.⁴⁵

To examine the statistics of the resolved flow, Figs. 3(c) and 3(d) show the spectrum of KE and the PDF of $\bar{u}$. Both LES with DDP and LES with DSMAG capture the KE spectrum up to near the maximum resolved k ($= 64$) although DDP does slightly better and agrees with the FDNS' KE spectrum up to $k \approx 60$ while DSMAG does so up to $k \approx 50$. Furthermore, as shown in panel (d), LES with DDP outperforms LES with DSMAG in capturing the PDF's tails, which correspond to shocks. Note that the differences between the PDFs of DDP, FDNS, and DSMAG are not statistically significant (at 95% confidence level) based on the KS or KL test, but that is because such tests mainly assess similarities in the bulk rather than the tails of the PDFs. A closer visual inspection shows that the difference between the tails of the PDFs from FDNS and DDP (DSMAG) is within (outside) the uncertainty range, indicating that DDP (DSMAG) accurately captures (does not capture) the statistics of the rare events.

So far we have discussed the results with LES resolution of 128 Fourier modes, which as mentioned before, conforms with the commonly used criterion for LES resolution based on the KE of the filtered flow. With the lower resolution of 96 modes, the DDP model still leads to a stable LES that outperforms LES with DSMAG. Further lowering the resolution to 64 modes leads to an unstable DDP. While LES with DSMAG is stable (due to the purely diffusive nature of this SGS model), the accuracy is poor. At this resolution, the $\bar{u}$ field only accounts for $\sim 40\%$ of the total KE, and LES (with any SGS model) is not expected to be used at such low resolutions.

In summary, the DDP model that uses an ANN trained with augmented data (from $Re = Re_c$) leads to a stable LES model (with reasonably high enough resolution) in *a posteriori* tests (at $Re = Re_c$) that is more accurate than LES with DSMAG. Next, we examine whether a DDP model trained with augmented data from a given Re can be used for LES of a flow that has higher Re .

Figure 4 shows the statistics of the resolved flow and of Π calculated using results from *a posteriori* tests at $Re = Re_c$ but with a DDP model that uses an ANN trained on data from $Re = Re_c/10$ (see the dashed blue lines). It is clear that this DDP model *does not* generalize as the spectrum and PDF of Π and the spectrum of KE all deviate from those of the FDNS. The results are not surprising as it is known that data-driven models often have difficulty with generalization to a different (especially more complex) system. For example, using a multi-scale Lorenz 96 system, we²³ showed that ANN- and recurrent neural network-based data-driven SGS models do not accurately generalize when the system is forced to become more chaotic. However, we also showed that transfer learning (TL)⁴⁷ provides an effective way for addressing this challenge, at least for a simple chaotic toy model. Below, we show the effectiveness of TL in making DDP generalizable to higher Re in a turbulent flow.

Figure 5 shows the schematic of TL applied to the ANN of a DDP model. In general, the weights of an ANN are randomly initialized and then they are updated through training on M samples from a given data distribution (here, data from a flow with $Re = Re_c/10$).




FIG. 4. Statistics of the resolved flow and SGS term calculated using results from *a posteriori* tests at $Re = Re_c$ but with DDP models mainly trained on data from $Re = Re_c/10$. Each curve is produced using 3×10^5 sequential samples that are $20\Delta t$ apart. The DDP model without transfer learning (TL) uses the ANN trained on $M = 5 \times 10^5$ samples from DNS at $Re = Re_c/10$. The DDP model with TL uses the same ANN but after its last two layers are re-trained with 5×10^4 samples from DNS at $Re = Re_c$ (Fig. 5). (a) Spectrum of Π . (b) PDF of Π . (c) Spectrum of KE.

The test in Fig. 4 showed that this ANN does not accurately work for $Re = Re_c$. The idea of TL is that we re-train this ANN (starting with its current weights rather than random initializations) and update the weights only in the deeper layers using a smaller number of samples (e.g., $M_{TL} = M/10$) from the new data distribution (i.e., the flow with $Re = Re_c$). The underlying idea of TL is that in deep networks, the initial layers learn high-level features, and only the deeper layers learn low-level features that are specific to a particular data distribution.⁴⁷ Thus, for generalization, we only need to re-train the deeper layers, which can be done using a small amount of data from the new distribution.

Figure 4 shows that the DDP model with TL (solid blue lines) accurately generalizes to the flow with Re_c as the spectrum and PDF of Π and spectrum of KE closely match those of FDNS. In fact, the accuracy of the DDP model with TL in Fig. 4 (which only uses $M_{TL} = 5 \times 10^4$ training samples from Re_c) is comparable with the accuracy of the DDP model in Fig. 3 (which uses $M = 5 \times 10^5$ training samples from Re_c). Furthermore, Fig. 6 shows how gradually increasing M_{TL} improves the generalization capability of the DDP model. Finally, a figure in the supplementary material further demonstrates the effect of the number of re-trained layers (as well as M_{TL}), showing that at large enough M_{TL} , re-training more than one layer yields accurate generalization.

In conclusion, we have investigated ANN-based data-driven SGS modeling of Burgers turbulence, with a particular focus on the stability




FIG. 5. Schematic of transfer learning (TL) to develop an accurate DDP model for $Re = Re_c$. Without TL, the ANN in the DDP model is trained, starting with random weights, on $M = 5 \times 10^5$ samples from DNS at $Re = Re_c/10$. This DDP model does not generalize to $Re = Re_c$ (dashed blue lines in Fig. 4). Then, TL is applied: the weights in the first three layers (blue) of this ANN are fixed, and the last two layers (red) are re-trained, starting with the previously computed weights, and using only $M_{TL} = 5 \times 10^4$ samples from DNS at $Re = Re_c$. The DDP model with TL is accurate and stable in *a posteriori* tests at $Re = Re_c$ (solid blue lines in Fig. 4).




FIG. 6. Spectrum of Π in *a posteriori* tests on $Re = Re_c$ as M_{TL} (the number of training samples from Re_c used in TL) is increased. $M_{TL} = 0$ correspond to no TL and the original ANN trained on $M = 5 \times 10^5$ samples from $Re = Re_c/10$. Adding $M_{TL} = 500$ –5000 samples improves the generalization capability of the DDP model to some degree. $M_{TL} = 10^4$ (2% of M) leads to substantial improvements although Π is underestimated at high k while overestimated at low k . Increasing M_{TL} to 5×10^4 (10% of M) further improves the generalization capability and Π that is just slightly underestimated.

of *a posteriori* LES models and generalization to higher Re . We show that developing a DDP model for Burgers turbulence is particularly challenging due to the presence of shocks, which localize the SGS term (Π), resulting in ANNs that overfit in the absence of a large training set. The overfitting ANNs lead to inaccurate/unstable DDP models. To overcome this challenge, we introduce a pre-conditioning step in which, exploiting periodicity, training samples are randomly shifted, thus enriching and augmenting the training set. The DDP model trained on this augmented dataset leads to stable and accurate *a posteriori* LES models. These results suggest that similar data augmentation strategies that exploit symmetries, invariances, and other physical properties (see Xie *et al.*,⁴⁸ Pan and Duraisamy,⁴⁹ and Formentin *et al.*⁵⁰ for more examples) should be considered in developing DDP models for more complex flows when large training sets are unavailable, not only to improve accuracy but also to improve the stability of *a posteriori* LES runs.

We have also found the DDP model not to generalize (i.e., extrapolate) to a flow with $10\times$ higher Re . However, we show, for the first time to the best of our knowledge, the application of TL to making a DDP model generalizable in a turbulent flow. Transfer learning enables the development of DDP models for high- Re flows with most of the training data provided by high-fidelity simulations at lower Re , which is highly appealing for practical purposes because the computational cost of simulating turbulent flows rapidly increases with Re .

In future work, the application of TL and data augmentation to develop accurate, stable, generalizable DDP models for more complex turbulent flows that are 2D and 3D will be investigated.

See the supplementary material for a figure showing the effects of the number of re-trained layers and the number of samples used for re-training on the generalization accuracy.

We thank Romit Maulik, Rambod Mojangani, and Ebrahim Nabizadeh for insightful discussions. We are grateful to two anonymous reviewers for helpful comments and suggestions that improved the quality of the manuscript. This work was supported by an award from the ONR Young Investigator Program, N00014-20-1-2722, and by NSF Grant No. OAC-2005123 (to P.H.). A.C. thanks the Rice University Ken Kennedy Institute for Information Technology for a BP HPC Graduate Fellowship. Computational resources were provided by NSF XSEDE (Allocation No. ATM170020) and by the Rice University Center for Research Computing.

DATA AVAILABILITY

The data that support the findings of this study are openly available in Zenodo at <https://doi.org/10.5281/zenodo.4316338> and in GitHub at https://github.com/envfluids/Burgers_DDP_and_TL.

REFERENCES

- ¹S. B. Pope, *Turbulent Flows* (Cambridge University Press, 2001).
- ²P. Sagaut, *Multiscale and Multiresolution Approaches in Turbulence: LES, DES and Hybrid RANS/LES Methods: Applications and Guidelines* (World Scientific, 2013).
- ³J. Ling, A. Kurzawski, and J. Templeton, "Reynolds averaged turbulence modelling using deep neural networks with embedded invariance," *J. Fluid Mech.* **807**, 155–166 (2016).
- ⁴J. N. Kutz, "Deep learning in fluid dynamics," *J. Fluid Mech.* **814**, 1–4 (2017).
- ⁵S. L. Brunton, B. R. Noack, and P. Koumoutsakos, "Machine learning for fluid mechanics," *Annu. Rev. Fluid Mech.* **52**, 477–508 (2020).

- ⁶J. Pathak, B. Hunt, M. Girvan, Z. Lu, and E. Ott, "Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach," *Phys. Rev. Lett.* **120**, 024102 (2018).
- ⁷J.-L. Wu, K. Kashinath, A. Albert, D. Chirila, H. Xiao *et al.*, "Enforcing statistical constraints in generative adversarial networks for modeling chaotic dynamical systems," *J. Comput. Phys.* **406**, 109209 (2020).
- ⁸A. T. Mohan, D. Tretiak, M. Chertkov, and D. Livescu, "Spatio-temporal deep learning models of 3D turbulence with physics informed diagnostics," *J. Turbul.* **21**, 484–441 (2020).
- ⁹A. Chattopadhyay, P. Hassanzadeh, and D. Subramanian, "Data-driven predictions of a multiscale Lorenz 96 chaotic system using machine-learning methods: Reservoir computing, artificial neural network, and long short-term memory network," *Nonlinear Processes Geophys.* **27**, 373–389 (2020).
- ¹⁰A. Chattopadhyay, E. Nabizadeh, and P. Hassanzadeh, "Analog forecasting of extreme-causing weather patterns using deep learning," *J. Adv. Modeling Earth Syst.* **12**, e2019MS001958 (2020).
- ¹¹M. Raissi, A. Yazdani, and G. E. Karniadakis, "Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations," *Science* **367**, 1026–1030 (2020).
- ¹²H. Eivazi, H. Veisi, M. H. Naderi, and V. Esfahanian, "Deep neural networks for nonlinear model order reduction of unsteady flows," *Phys. Fluids* **32**, 105104 (2020).
- ¹³S. Pandey, J. Schumacher, and K. R. Sreenivasan, "A perspective on machine learning in turbulent flows," *J. Turbul.* **21**, 567–584 (2020).
- ¹⁴S. Pan and K. Duraisamy, "Data-driven discovery of closure models," *SIAM J. Appl. Dyn. Syst.* **17**, 2381–2413 (2018).
- ¹⁵K. Duraisamy, G. Iaccarino, and H. Xiao, "Turbulence modeling in the age of data," *Annu. Rev. Fluid Mech.* **51**, 357–377 (2019).
- ¹⁶C. Xie, J. Wang, H. Li, M. Wan, and S. Chen, "Artificial neural network mixed model for large eddy simulation of compressible isotropic turbulence," *Phys. Fluids* **31**, 085112 (2019).
- ¹⁷R. Maulik, O. San, A. Rasheed, and P. Vedula, "Subgrid modelling for two-dimensional turbulence using neural networks," *J. Fluid Mech.* **858**, 122–144 (2019).
- ¹⁸A. Beck, D. Flad, and C.-D. Munz, "Deep neural networks for data-driven LES closure models," *J. Comput. Phys.* **398**, 108910 (2019).
- ¹⁹Z. Zhou, G. He, S. Wang, and G. Jin, "Subgrid-scale model for large-eddy simulation of isotropic turbulent flows using an artificial neural network," *Comput. Fluids* **195**, 104319 (2019).
- ²⁰T. Bolton and L. Zanna, "Applications of deep learning to ocean data inference and sub-grid parameterisation," *J. Adv. Modeling Earth Syst.* **11**, 376–399 (2019).
- ²¹S. Pawar, O. San, A. Rasheed, and P. Vedula, "A priori analysis on deep learning of subgrid-scale parameterizations for Kraichnan turbulence," *Theor. Comput. Fluid Dyn.* **34**, 429–427 (2020).
- ²²C. Xie, J. Wang, and E. Weinan, "Modeling subgrid-scale forces by spatial artificial neural networks in large eddy simulation of turbulence," *Phys. Rev. Fluids* **5**, 054606 (2020).
- ²³A. Chattopadhyay, A. Subel, and P. Hassanzadeh, "Data-driven super-parameterization using deep learning: Experimentation with multi-scale Lorenz 96 systems and transfer learning," *J. Adv. Modeling Earth Syst.* **21**, e2020MS002084 (2020).
- ²⁴H. Frezat, G. Balarac, J. Le Sommer, R. Fablet, and R. Lguensat, "Physical invariance in neural networks for subgrid-scale scalar flux modeling," *Phys. Rev. Fluids* **6**, 024607 (2021).
- ²⁵L. Zanna and T. Bolton, "Data-driven equation discovery of ocean mesoscale closures," *Geophys. Res. Lett.* **47**, e2020GL088376, <https://doi.org/10.1029/2020GL088376> (2020).
- ²⁶M. Kurz and A. Beck, "A machine learning framework for LES closure terms," [arXiv:2010.03030](https://arxiv.org/abs/2010.03030) (2020).
- ²⁷S. Pawar, S. E. Ahmed, and O. San, "Interface learning in fluid dynamics: Statistical inference of closures within micro-macro-coupling models," *Phys. Fluids* **32**, 091704 (2020).
- ²⁸U. Piomelli, W. H. Cabot, P. Moin, and S. Lee, "Subgrid-scale backscatter in turbulent and transitional flows," *Phys. Fluids A: Fluid Dyn.* **3**, 1766–1771 (1991).
- ²⁹H. T. Hewitt, M. Roberts, P. Mathiot, A. Biastoch, E. Blockley, E. P. Chassignet, B. Fox-Kemper, P. Hyder, D. P. Marshall, E. Popova *et al.*, "Resolving and parameterising the ocean mesoscale in earth system models," *Curr. Clim. Change Rep.* **6**, 137–116 (2020).
- ³⁰J. Smagorinsky, "General circulation experiments with the primitive equations: I. The basic experiment," *Mon. Weather Rev.* **91**, 99–164 (1963).
- ³¹M. Germano, U. Piomelli, P. Moin, and W. H. Cabot, "A dynamic subgrid-scale eddy viscosity model," *Phys. Fluids A: Fluid Dyn.* **3**, 1760–1765 (1991).
- ³²C. Meneveau and T. S. Lund, "The dynamic Smagorinsky model and scale-dependent coefficients in the viscous range of turbulence," *Phys. Fluids* **9**, 3932–3934 (1997).
- ³³D. Krueger, E. Caballero, J.-H. Jacobsen, A. Zhang, J. Binas, R. L. Priol, and A. Courville, "Out-of-distribution generalization via risk extrapolation (REx)," [arXiv:2003.00688](https://arxiv.org/abs/2003.00688) (2020).
- ³⁴S. I. Dolaptchiev, U. Achatz, and I. Timofeyev, "Stochastic closure for local averages in the finite-difference discretization of the forced Burgers equation," *Theor. Comput. Fluid Dyn.* **27**, 297–317 (2013).
- ³⁵S. S. Girimaji, "Spectrum and energy transfer in steady Burgers turbulence," *Phys. Lett. A* **202**, 279–287 (1995).
- ³⁶A. Das and R. D. Moser, "Optimal large-eddy simulation of forced Burgers equation," *Phys. Fluids* **14**, 4344–4351 (2002).
- ³⁷M. Love, "Subgrid modelling studies with Burgers' equation," *J. Fluid Mech.* **100**, 87–110 (1980).
- ³⁸A. LaBryer, P. J. Attar, and P. Vedula, "A framework for large eddy simulation of Burgers turbulence based upon spatial and temporal statistical information," *Phys. Fluids* **27**, 035116 (2015).
- ³⁹R. Maulik and O. San, "Explicit and implicit LES closures for Burgers turbulence," *J. Comput. Appl. Math.* **327**, 12–40 (2018).
- ⁴⁰J. Alcalá and I. Timofeyev, "Subgrid-scale parametrization of unresolved scales in forced Burgers equation using generative adversarial networks (GAN)," [arXiv:2007.06692](https://arxiv.org/abs/2007.06692) (2020).
- ⁴¹I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio, *Deep Learning* (MIT Press, Cambridge, 2016), Vol. 1.
- ⁴²P. Ramachandran, B. Zoph, and Q. V. Le, "Searching for activation functions," [arXiv:1710.05941](https://arxiv.org/abs/1710.05941) (2017).
- ⁴³M. A. Tanner and W. H. Wong, "The calculation of posterior distributions by data augmentation," *J. Am. Stat. Assoc.* **82**, 528–540 (1987).
- ⁴⁴L. Yanan and W. Z. J., "A priori and a posteriori evaluations of sub-grid scale models for the Burgers equation," *Comput. Fluids* **139**, 92–104 (2016).
- ⁴⁵R. Wilcox, *Fundamentals of Modern Statistical Methods: Substantially Improving Power and Accuracy* (Springer Science & Business Media, 2010).
- ⁴⁶C. Bayona, J. Baiges, and R. Codina, "Variational multiscale approximation of the one-dimensional forced Burgers equation: The role of orthogonal subgrid scales in turbulence modeling," *Int. J. Numer. Methods Fluids* **86**, 313–328 (2018).
- ⁴⁷J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *NeurIPS Proceedings of Advances in Neural Information Processing Systems* (2014), pp. 3320–3328.
- ⁴⁸Y. Xie, E. Franz, M. Chu, and N. Thuerey, "tempoGAN: A temporally coherent, volumetric GAN for super-resolution fluid flow," *ACM Trans. Graph.* **37**, 1–15 (2018).
- ⁴⁹S. Pan and K. Duraisamy, "Long-time predictive modeling of nonlinear dynamical systems using neural networks," *Complexity* **2018**, 1.
- ⁵⁰S. Formentin, M. Mazzoleni, M. Scandella, and F. Previdi, "Nonlinear system identification via data augmentation," *Syst. Control Lett.* **128**, 56–63 (2019).