
Neural Ideal Large Eddy Simulation: Modeling Turbulence with Neural Stochastic Differential Equations

Anudhyan Boral
Google Research
Mountain View, CA 94043, USA
anudhyan@google.com

Zhong Yi Wan
Google Research
Mountain View, CA 94043, USA
wanzy@google.com

Leonardo Zepeda-Núñez
Google Research
Mountain View, CA 94043, USA
lzepedanunez@google.com

James Lottes
Google Research
Mountain View, CA 94043, USA
jlottes@google.com

Qing Wang
Google Research
Mountain View, CA 94043, USA
wqing@google.com

Yi-fan Chen
Google Research
Mountain View, CA 94043, USA
yifanchen@google.com

John Roberts Anderson
Google Research
Mountain View, CA 94043, USA
janders@google.com

Fei Sha
Google Research
Mountain View, CA 94043, USA
fsha@google.com

Abstract

We introduce a data-driven learning framework that assimilates two powerful ideas: ideal large eddy simulation (LES) from turbulence closure modeling and neural stochastic differential equations (SDE) for stochastic modeling. The ideal LES models the LES flow by treating each full-order trajectory as a random realization of the underlying dynamics, as such, the effect of small-scales is marginalized to obtain the deterministic evolution of the LES state. However, ideal LES is analytically intractable. In our work, we use a latent neural SDE to model the evolution of the stochastic process and an encoder-decoder pair for transforming between the latent space and the desired ideal flow field. This stands in sharp contrast to other types of neural parameterization of closure models where each trajectory is treated as a deterministic realization of the dynamics. We show the effectiveness of our approach (niLES – neural ideal LES) on two challenging chaotic dynamical systems: Kolmogorov flow at a Reynolds number of 20,000 and flow past a cylinder at Reynolds number 500. Compared to competing methods, our method can handle non-uniform geometries using unstructured meshes seamlessly. In particular, niLES leads to trajectories with more accurate statistics and enhances stability, particularly for long-horizon rollouts. (Source codes and datasets will be made publicly available.)

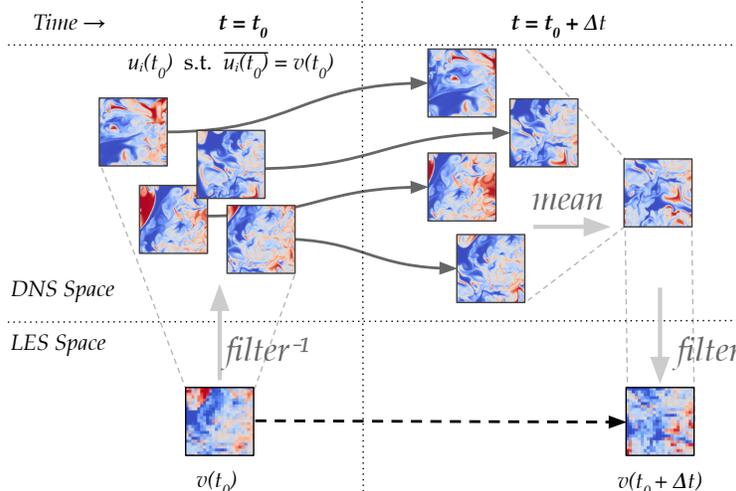


Figure 1: A cartoon of ideal LES. The LES field is lifted into the space of real turbulent fields by applying the multi-valued inverse of the filtering operator. Then the turbulent DNS fields are evolved continuously according to the N-S equations. Finally, at the end of the LES time step, the mean of the DNS fields is filtered to obtain the new LES field. Although the ideal LES is the *ideal* LES evolution, it is analytically intractable. DNS image from [54].

1 Introduction

Multiscale physical systems are ubiquitous and play major roles in science and engineering [5, 7, 16, 25, 57, 74, 38]. The main difficulty of simulating such systems is the need to numerically resolve strongly interacting scales that are usually order of magnitude apart. One prime example of such problems is turbulent flows, in which a fluid flow becomes chaotic under the influence of its own inertia. As such, high-fidelity simulations of such flows would require solving non-linear partial differential equations (PDEs) at very fine discretization, which is often prohibitive for downstream applications due to the high computational cost. Nonetheless, such direct numerical simulations (DNS) are regarded as gold-standard [62, 43].

For many applications where the primary interest lies in the large (spatial) scale features of the flows, solving coarse-grained PDEs is a favorable choice. However, due to the so-called back-scattering effect, the energy and the dynamics of the small-scales can have a significant influence on the behavior of the large ones [51]. Therefore, coarsening the discretization scheme *alone* results in a highly biased (and often incorrect) characterization of the large-scale dynamics. To address this issue, current approaches incorporate the interaction between the resolved and unresolved scales by employing statistical models based on the physical properties of fluids. These mathematical models are commonly known as *closure models*. Closure models in widely used approaches such as Reynolds-Averaged Navier-Stokes (RANS) and Large Eddy Simulation (LES) [68] have achieved considerable success. But they are difficult to derive for complex configurations of geometry and boundary conditions, and inherently limited in terms of accuracy [26, 67, 41].

An increasing amount of recent works have shown that machine learning (ML) has the potential to overcome many of these limitations by constructing data-driven closure models [55, 77, 47, 56]. Specifically, those ML-based models correct the aforementioned bias by comparing their resulting trajectories to coarsened DNS trajectories as ground-truth. Despite their empirical success and advantage over classical (analytical) ones, ML-based closure models suffer from several deficiencies, particularly when extrapolating beyond the training regime [9, 71, 79]. They are unstable when rolling the dynamics out to long horizons, and debugging such black box models is challenging. This raises the question: *how do we incorporate the inductive bias of the desired LES field to make a good architectural choice of the learning model?*

In this paper, we propose a new ML-based framework for designing closure models. We synergize two powerful ideas: ideal LES [50] and generative modeling via neural stochastic differential equations (NSDE) [80, 52, 44, 45]. The architecture of our closure model closely mirrors the desiderata of ideal

LES for marginalizing small-scale effects of inherent stochasticity in turbulent flows. To the best of our knowledge, our work is the first to apply probabilistic generative models for closure modeling.

Ideal LES seeks a flow field of large-scale features such that the flow is optimally consistent with an ensemble of DNS trajectories that are filtered to preserve large-scale features [50]. The optimality, in terms of minimum squared error, leads to the conditional expectation of the filtered DNS flows in the ensemble. Ideal LES stems from the observation that turbulent flows, while in principle deterministic, are of stochastic nature as small perturbations can build up exponentially fast and after a characteristic *Lyapunov* time scale the perturbations *erases* the signal from the initial condition. Thus, while deterministically capturing the trajectory of the large-scales from a single filtered DNS trajectory can be infeasible due to the chaotic behavior and loss of information through discretization, it may be possible to predict the *statistics* reliably over many possible realizations of DNS trajectories (with the same large-scale features) by marginalizing out the random effect. Fig 1 illustrates the main idea and Section 2 provides a detailed description. Unfortunately, ideal LES is not analytically tractable because both the distribution of the ensemble as well and the desired LES field are unknown.

To tackle this challenge, our main idea is to leverage i) the expressiveness of neural stochastic differential equations (NSDE) [52, 83] to model the unknown distribution of the ensemble as well as its time evolution, and ii) the power in Transformer-based encoder-decoders to map between the LES field to-be-learned and a latent space where the ensemble of the DNS flows presumably resides. The resulting closure model, which we term as neural ideal LES (niLES), incorporates the inductive bias of ideal LES by estimating the target conditional expectations via a Monte Carlo estimation of the statistics from the NSDE-learned distributions, which are themselves estimated from DNS data.

We investigate the effectiveness of niLES in modeling the Kolmogorov flow, which is a classical example of turbulent flow. We compare our framework to other methods that use neural networks to learn closure models, assuming that each trajectory is a deterministic realization. We demonstrate that niLES leads to more accurate trajectories and statistics of the dynamics, which remain stable even when rolled out to long horizons. This demonstrates the benefits of learning samples as statistics (i.e., conditional expectations) from ensembles.

2 Background

We provide a succinct introduction to key concepts necessary to develop our methodology: closure modeling for turbulent flows via large eddy simulation (LES) and neural stochastic different equations (NSDE). We exemplify these concepts with the Navier-Stokes (N-S) equations.

Navier-Stokes and direct numerical simulation (DNS) We consider the N-S equations for incompressible fluids without external forcing. In dimensionless form, the N-S equations are:

$$\partial_t u + (u \cdot \nabla)u = -\nabla p + \nu \nabla^2 u \quad \text{with} \quad \nabla \cdot u = 0 \quad (1)$$

where $u = u(x, t)$ and $p = p(x, t)$ are the velocity and pressure of a fluid at a spatial point x in the domain $\Omega \subset \mathbb{R}^d$ at time t ; ν is the kinematic viscosity, reciprocal of the Reynolds number Re , which characterizes the degree of turbulence of the flow. We may eliminate pressure p on the right-hand-side; e.g. by taking the divergence of the momentum equation and using the fact that velocity is divergence-free. Hence, we rewrite the N-S equations compactly as

$$\partial_t u = \mathcal{R}^{\text{NS}}(u; \nu), \quad (2)$$

We impose boundary conditions on $\partial\Omega$ and initial conditions $u(x, 0) = u_0(x)$, $x \in \Omega$ in N-S equations throughout the manuscript. We sometimes omit the implicit ν on the right hand side. To solve numerically, DNS will discretize the equation on a fine grid \mathcal{G} , such that all the scales are adequately resolved. It is important to note that as ν becomes small, the inertial effects becomes dominant, thus requiring a refinement of the grid (and time-step due to the Courant–Friedrichs–Lewy (CFL) condition [21]). This rapidly increases the computational cost [62].

LES and closure modeling For many applications where the primary interests are large-scale features, the LES methodology balances computational cost and accuracy [76, 72]. It uses a coarse grid $\bar{\mathcal{G}}$ which has much fewer degrees-of-freedom than the fine grid \mathcal{G} . To represent the dynamics with respect to $\bar{\mathcal{G}}$, we define a filtering operator $(\bar{\cdot}) : \mathbb{R}^{|\mathcal{G}| \times d} \mapsto \mathbb{R}^{|\bar{\mathcal{G}}| \times d}$ which is commonly implemented

using low-pass (in spatial frequencies) filters [12]. Applying the filter to Eq. (2), we have

$$\partial_t \bar{u} = \mathcal{R}_c^{\text{NS}}(\bar{u}; \nu) + \mathcal{R}^{\text{closure}}(\bar{u}, u) \quad (3)$$

where \bar{u} is the LES field, $\mathcal{R}_c^{\text{NS}}$ has the same form as \mathcal{R}^{NS} with \bar{u} as input, u is the DNS field in Eq. (2), and $\mathcal{R}^{\text{closure}}(\bar{u}, u) : \mathbb{R}^{|\mathcal{G}|\times d} \times \mathbb{R}^{|\mathcal{G}|\times d} \mapsto \mathbb{R}^{|\mathcal{G}|\times d} = \nabla \cdot (\bar{u} \bar{u} - \overline{u u})$ is the closure term. It represents the collective effect of unresolved subgrid scales, which are smaller than the resolved scales in $\bar{\mathcal{G}}$.

However, as u is unknown to the LES solver, the closure term needs to be approximated by functions of \bar{u} . How to model such terms has been the subject of a large amount of literature (see Section 5). Traditionally, those models are mathematical ones; deriving and analyzing them is highly challenging for complex cases and entails understanding the physics of the fluids.

Learning-based closure modeling One emerging trend is to leverage machine learning (ML) tools to learn a data-driven closure model [47, 77] to parameterize the closure term,

$$\mathcal{R}^{\text{closure}}(\bar{u}, u) \approx \mathcal{M}(\bar{u}; \theta) \quad (4)$$

where θ is the parameter of the learning model (often, a neural network). With a DNS field as a ground-truth, the goal is to adjust the θ such that the approximating LES field

$$\partial_t \tilde{u} = \mathcal{R}_c^{\text{NS}}(\tilde{u}; \nu) + \mathcal{M}(\tilde{u}; \theta) \quad (5)$$

matches the filtered DNS field \bar{u} . This is often achieved through the empirical risk minimization framework in ML:

$$\theta^* = \arg \min_{\theta} \sum_i \|\tilde{u}_i - \bar{u}_i\|_2^2 \quad (6)$$

where i indexes the trajectories in the training dataset, each being a DNS field from a simulation run with a different condition.

Despite their success, learning-based models have also their own drawbacks. Among them, *how to choose the learning architecture for parameterizing \tilde{u}* is more of an art than a science. Our work aims to shed light on this question by advocating designing the architecture to incorporate the inductive bias of designed LES fields. In what follows, we describe ideal LES, which motivates our work. To give a preview, our probabilistic ML framing matches very well the formulation of ideal LES in extracting statistics from turbulent flows of inherent stochastic nature.

Ideal LES It has long been observed that while chaotic systems, such as turbulent flows, can be in principle deterministic, they are stochastic in nature due to the fast growth (of errors) with respect to even small perturbation. This has led to many seminal works that treat the effect of small scales stochastically [68]. Thus, instead of viewing each DNS field as a deterministic and distinctive realization, one should consider an ensemble of DNS fields. Furthermore, since the filtering operator $(\bar{\cdot})$ is fundamentally lossy [75], filtering multiple DNS fields could result in the same LES state. Ideal LES identifies an evolution of LES field such that the dynamics is consistent with the dynamics of its corresponding (many) DNS fields. Formally, let the initial distribution $\pi_{t_0}(u)$ over the (unfiltered) turbulent fields to be fixed but unknown. By evolving forward the initial distribution according to Eq. (2), we obtain the stochastic process $\pi_t(u)$.

The evolution of the ideal LES field v is obtained from the time derivatives of the set of unfiltered turbulent fields whose large scale features are the same as v [50]:

$$\frac{\partial v}{\partial t} = \mathbb{E}_{\pi_t} \left[\frac{\partial u}{\partial t} \Big| \bar{u} = v \right] \quad (7)$$

Fig 1 illustrates the conceptual framing of the ideal LES. We can gain additional intuition by observing that the field v also attains the minimum mean squared error, matching its velocity $\partial v / \partial t$ to that of the filtered field $\partial \bar{u} / \partial t$.

It is difficult to obtain the set $\{u \mid \bar{u} = v\}$ which is required to compute the velocity field (and infer v). Thus, despite its conceptual appeal, ideal LES is analytic intractable. We will show how to derive a data-driven closure model inspired by the ideal LES, using the tool of NSDE described below.

NSDE NSDE extends the classical stochastic differential equations by using neural-network parameterized drift and diffusion terms [45, 80, 52]. It has been widely used as a data-driven model for stochastic dynamical systems. Concretely, let time $t \in [0, 1]$, Z_t the latent state and X_t the observed variable. NSDE defines the following generative process of data via a latent Markov process:

$$Z_0 \sim p_0(\cdot), p(Z_t) \sim dZ_t = h_\theta(Z_t, t)dt + g_\theta(Z_t, t) \circ dW_t, X_t \sim p(X_t|Z_t) \quad (8)$$

where $p_0(\cdot)$ is the distribution for the initial state. W_t is the Wiener process and \circ denotes the Stratonovich stochastic integral. The Markov process $\{Z_t\}_{t \in [0,1]}$ provides a probabilistic prior for the dynamics, to be inferred from observations $\{X_t\}_{t \in [0,1]}$. Note that the observation model $p(X_t|Z_t)$ only depends on the state at time t . $h(\cdot)$ and $g(\cdot)$ are the drift and diffusion terms, expressed by two neural networks with parameters θ and ϕ .

Given observation data $x = \{x_t\}$, learning θ and ϕ is achieved by maximizing the Evidence Lower Bound (ELBO), under a variational posterior distribution which is also an SDE [52]

$$q(Z_t|x) \sim dZ_t = h_\phi(Z_t, t, x)dt + g_\theta(Z_t, t) \circ dW_t \quad (9)$$

Note that the variational posterior has the same diffusion as the prior. This is required to ensure a finite ELBO, which is given by

$$\mathcal{L} = \mathbb{E}_q \left\{ \int_0^1 \log(x_t|z_t)dt - \int_0^1 \frac{1}{2} \left(\frac{h_\phi(Z_t, t) - h_\theta(Z_t, t, x)}{g_\theta(Z_t, t)} \right)^2 dt \right\} \quad (10)$$

Note that both the original SDE parameters θ and ϕ and the variational parameter ψ are jointly optimized to maximize \mathcal{L} . For a detailed exposition of the subject, please refer to [45]. In the next section, we will describe how NSDE is used to characterize the stochastic turbulent flow fields.

3 Methodology

We propose a neural SDE based closure model that implements ideal LES. The generative latent dynamical process in neural SDE provides a natural setting for modeling the unknown distribution of the DNS flow ensemble that is crucial for ideal LES to reproduce long-term statistics.

Setup We are given a set of filtered DNS trajectories $\{\bar{u}_i\}$. Each \bar{u}_i is a sequence of ‘‘snapshots’’ indexed by time t and spans a temporal interval \mathcal{T}_i over the domain Ω . We use `VARIABLE(t)` to denote the time t snapshot of the `VARIABLE` (such as $\bar{u}_i(t)$ or $v(t)$). Those trajectories are treated as ‘‘ground-truth’’ and we would like to derive a data-driven closure model $\mathcal{M}(v; \Theta)$ in the form of Eq. (3) to evolve the LES state v

$$\partial_t v = \mathcal{R}_c^{\text{NS}}(v) + \mathcal{M}(v), \quad (11)$$

where we have dropped the model’s parameters Θ for notation simplicity. Our goal is to identify the optimal $\mathcal{M}(v)$ such that the trajectories of Eq. (11) have the same long-term statistics as Eq. (2). Following ideal LES, we would render Eq. (11) equivalent to Eq. (7) by implementing the ideal $\mathcal{M}(v)$ as

$$\mathcal{M}(v(t)) = \mathbb{E}_{\pi_t} [\bar{\partial}_t u | \bar{u} = v(t)] - \mathcal{R}_c^{\text{NS}}(v(t)). \quad (12)$$

Let $\bar{p}_t(u; v(t))$ denote the density of $\pi_t(u)$ restricted to the set $\{u | \bar{u} = v(t)\}$:

$$\bar{p}_t(u; v(t)) \propto \delta(\bar{u} = v(t))\pi_t(u). \quad (13)$$

We can thus rewrite the closure model as

$$\mathcal{M}(v(t)) = \int [\bar{\partial}_t u - \mathcal{R}_c^{\text{NS}}(v(t))] \bar{p}(u; v(t)) du = \int f(u; v(t)) \bar{p}(u; v(t)) du \quad (14)$$

where $f(u)$ denotes the fluctuation exerted on the large scales v if the *true* DNS trajectory was u . This shows that the ideal $\mathcal{M}(v(t))$ should compute the mean effect of the small-scale fluctuations $f(u)$ by integrating over all possible DNS trajectories u that are consistent with the large scales of $v(t)$. However, just as ideal LES is not analytically tractable, so is this ideal closure model. Specifically, while the term $\mathcal{R}_c^{\text{NS}}(v(t))$ can be easily computed using a numerical solver on the coarse grid, the remaining terms are not easily computed. In particular, $\bar{\partial}_t u$ would require a DNS solver, thus defeating the purpose of seeking a closure model. An approximation to Eq. (14) is needed.

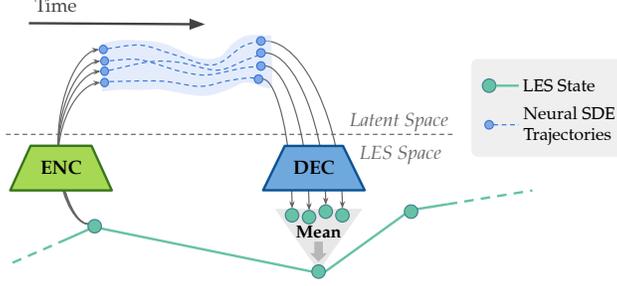


Figure 2: Schematic of our modeling approach motivated from ideal LES (cf. Figure 1 for structural correspondence). The evolution in the low-dimensional latent space follows trajectories of a data-driven Neural SDE, mirroring the fine temporal resolution of the DNS trajectories. The final states of the latent state at the next time step are decoded into the LES space and averaged over to obtain the mean correction due to small-scale fluctuations.

3.1 Main idea

Consider a trajectory segment from t_0 to t_1 where we are told only $\bar{u}(t_0)$ and $\bar{u}(t_1)$, *how can we discover all valid DNS trajectories between these two times to compute the closure model without incurring the cost of actually computing DNS?* Our main idea is to use a probabilistic generative model to generate a representation of those trajectories at a fraction of the cost. One can certainly learn from a corpus of DNS trajectories but the corpus is still too costly to acquire. Instead, we leverage the observation that we do not need the *full* details of the DNS trajectories in order to approximate $M(v(t))$ well — the expectation itself is a low-pass filtering operation. Thus, we proceed by constructing a (parameterized) stochastic dynamic process to emulate the *unknown* trajectories and collect the necessary statistics. As long as the constructed process is differentiable, we can optimize it end to end by matching the resulting dynamics of using the closure model to the ground-truth.

We sketch the main components below. The stochastic process is instantiated in the latent state space of a neural SDE with an encoder whose output defines the initial state distribution controlled by the desired LES state at time t_0 . The desired statistics, *i.e.*, the mean effect of small-scales, is computed in Monte Carlo estimation via a parameterized nonlinear mapping called *decoder*. The resulting correction by the closure model is then compared to the desired LES state at time t_1 , driving learning signals to optimize all parameters. See Fig. 2 for an illustration with details in Appendix D.

Encoder The encoder defines the distribution of the initial latent state variable in the NSDE, denoted by $Z_0 \in \mathbb{R}^L$. Concretely, $\mathcal{E} : \mathbb{R}^{\bar{G} \times d} \mapsto \mathbb{R}^{L+L(L+1)/2}$ maps from $v(t_0)$ in the LES space to the mean and the covariance matrix of a multidimensional Gaussian in the latent space: $\mathcal{E}(v(t_0)) = (\mu_{z_0}, \Sigma_{z_0})$. This distribution is used to sample K initial latent states in \mathbb{R}^L : $Z_0^{(i)} \sim \mathcal{N}(\mu_{z_0}, \Sigma_{z_0})$ ($1 \leq i \leq K$).

Latent space evolution The latent stochastic process evolves according to time $\tau \in [0, 1]$. This is an important distinction from the time for the LES field. Since we are interested in extracting statistics from DNS with finer temporal resolution, τ represents a faster (time-stepping) evolving process whose beginning and end map to the physical time t_0 and t_1 . (In our implementation, $\Delta t = t_1 - t_0$, the time step of the coarse solver, is greater than $\Delta \tau$, the time step for the NSDE.)

The latent variable Z_τ evolves according to the Neural SDE:

$$dZ_\tau = h_\phi(Z_\tau, \tau)d\tau + g_\theta(Z_\tau, \tau) \circ dW_\tau \quad (15)$$

where W_τ is the Wiener process on the interval $[0, 1]$. We obtain trajectories $\{Z_\tau^{(i)}\}_{\tau \in [0, 1]}$ sampled from the NSDE, and in particular, we obtain an ensemble of $\{Z_1^{(i)}\}_{i=1}^K$.

Decoder The decoder $\mathcal{D} : \mathbb{R}^L \mapsto \mathbb{R}^{\bar{G} \times d}$ maps each of the K ensemble members $\{Z_1^{(i)}\}_{i=1}^K$ from latent state back into the LES space. So we can compute the Monte-Carlo approximation of

Algorithm 1 Compute $\mathcal{M}(v(t_0))$

Input: LES state $v(t_0)$, Closure model parameters Θ

1. Encode $v(t_0)$ to a distribution on the latent state: $(\mu_{z_0}, \sigma_{z_0}) = \mathcal{E}_\Theta(v(t_0))$.
 2. Sample K initial latent states in \mathbb{R}^L ; $Z_0^{(i)} \sim \mathcal{N}(\mu_{z_0}, \sigma_{z_0})$ ($1 \leq i \leq K$).
 3. Sample trajectories $Z_\tau^{(i)}$ ($\tau \in (0, 1]$) by solving Eq. (15) with initial conditions $Z_0^{(i)}$.
 4. Decode the final latent states $Z_1^{(i)}$ to the LES space: $\mathbf{x}^{(i)} = \mathcal{D}_\Theta(Z_1^{(i)})$ ($1 \leq i \leq K$).
 5. Take the empirical mean of the samples: $v' = (\frac{1}{K}) \sum_{i=1}^K \mathbf{x}^{(i)}$.
- Output:** $v'/(t_1 - t_0)$.
-

$\mathcal{M}(v(t_0))(t_1 - t_0)$ (cf. Eq. (14) for the definition) as

$$\int_{t_0}^{t_1} dt \int f(u; v(t)) \bar{p}(u; v(t)) du \approx \int dw f_w(w; v) p_w(w) \approx \frac{1}{K} \sum_{i=1}^K \mathcal{D}(Z_1^{(i)}), \quad (16)$$

where w is the spatio-temporal lifted version of u , i.e., the space of DNS trajectories in space and time, and f_w absorbs both f in the closure Eq. (14) and the implicit conditioning in Eq. (13), and p_w is a prior for the trajectories. This notational change allows us to approximate the integral directly by a Monte-Carlo approximation on the lifted variables, i.e., the distribution of trajectories (as shown in Fig. 2) which are modeled using the NSDE in Eq. (15). See Algo. 1 for the calculation steps.

We stress that while we motivate our design via mimicking DNS trajectories, Z is low-dimensional and is not replicating the real DNS field u . However, it is possible that with enough training data, Z might discover the low-dimensional solution manifold in u .

Training objective The NSDE is differentiable. Thus, with the data-driven closure model Eq. (16), we can apply end-to-end learning to match the dynamics with the closure model to ground-truths. Concretely, let $v(t_0)$ be $\bar{u}(t_0)$ and we numerically integrate Eq. (11)

$$v(t_1) \approx v(t_0) + \int_{t_0}^{t_1} \mathcal{R}_c^{\text{NS}}(v(t)) dt + \mathcal{M}(v(t_0))(t_1 - t_0) \quad (17)$$

This gives rise to the likelihood of the (observed) data in NSDE. Specifically, following the VAE setup of [52], we have

$$-\log p(\bar{u}(t_1)|Z) = \mathcal{L}^{\text{recon}}(v(t_1), \bar{u}(t_1)) = (2\sigma^2)^{-1} \|v(t_1) - \bar{u}(t_1)\|^2 \quad (18)$$

where σ is a scalar posterior scale term. The training objective $\mathcal{L}(v, u) = \mathcal{L}^{\text{recon}}(v, u) + \text{KL}^{\text{NSDE}}$ includes a KL divergence term associated with the neural SDE. See Appendix D for more details on the additional term.

To endow more stability to the training, following [81] the training loss incorporates S time steps ($S > 1$) of rollouts of the LES state:

$$\mathcal{L}^{(S)}(v, \bar{u}) = \sum_{k=1}^S \mathcal{L}(v(t_0 + k\Delta t), \bar{u}(t_0 + k\Delta t)) \quad (19)$$

For a dataset with multiple trajectories, we just sum the loss for each one of them and optimize through stochastic gradient descent.

3.2 Implementation details

We describe details in the Appendix G for the Transformer-based encoder and decoder and Appendix D for implementing NSDE with a Transformer parameterizing the drift and diffusion terms.

4 Experimental results

We showcase the advantage of our approach using two instances of a chaotic Navier-Stokes flow: 2D Kolmogorov flow, [47] at a Reynolds number of 20,000 and flow past a circular cylinder at a Reynolds number of 500.

4.1 Setup

Dataset generation The reference data for Kolmogorov flow consists of an ensemble of trajectories generated by randomly perturbing an initial condition. The flow past cylinder dataset is generated as a single long, chaotic trajectory, split up into 14 equal segments with an initial subsegment thrown away to avoid correlation between the segments. Each trajectory is generated by solving the NS equations directly using a high-order spectral finite element discretization in space and time-stepping is done via a 3rd order backward differentiation formula (BDF3) [23], subject to the appropriate Courant–Friedrichs–Lewy (CFL) condition [21]. These DNS calculations were performed on a mesh with 2304 elements for Kolmogorov flow and 292 elements for the cylinder flow with each element having a polynomial order of 8 and 13 respectively. The DNS trajectories were then sampled to a coarse grid with order 4 and 5 respectively for the two datasets. For more details see Appendix B.

Benchmark methods For all LES methods, including niLES, we use a $10\times$ larger time step than the DNS simulation. We compare our method against a classical implicit LES at polynomial order 4 using a high order spectral element solver using the filtering procedure of [28]. As reported in [14], the implicit LES for high order spectral methods is comparable to Smagorinsky subgrid-scale eddy-viscosity models. We also train an encoder-decoder deterministic NN-based closure model using a similar transformer backend as our niLES model. This follows the procedure from prior works [56, 81, 47, 24] in training deterministic NN-based closure models. Additionally, we train a deterministic encoder-processor-decoder network with a neural ODE processor [19] with strictly more parameters than the niLES model.

Metrics For measuring the short-term accuracy we used root-mean-squared-error (RMSE) after unrolling the LES forward. Due to chaotic divergence, the measure becomes meaningless after about 1000 steps. For assessing the long-term ability to capture the statistics of DNS we use turbulent kinetic energy (TKE) spectra. The TKE spectra is obtained by taking the 2D Fourier transform of the velocity field, after interpolating it to a uniform grid. Then we obtain the kinetic energy in the Fourier domain and we integrate it in concentric annuli along the 2D wavenumber k .

4.2 Main Results

Long-term turbulent statistics We summarize in Figs. 3 and 5. In the short-term regime, both niLES and deterministic methods achieve higher accuracy than the implicit LES, even when unrolling for hundreds of steps. Beyond this time frame, the chaotic divergence leads to exponential accumulation of pointwise errors, until the LES states are fully decorrelated with the DNS trajectory. At this point, we must resort to statistical measures to gauge the quality of the rollout.

When considering reconstruction errors from very short-term rollouts, such as when the model is unrolled only 8 times, identical to the training setup, the deterministic approaches and niLES yield similar reconstruction errors. Beyond this regime, and going up to 1000s of rollout steps, the probabilistic approach of niLES provides superior generalization ability even when measuring a deterministic reconstruction loss.

Furthermore, niLES captures long-term turbulent statistics significantly better than the other two approaches, particularly in the high wavenumber regime. The deterministic NN is not stable for long term rollouts due to energy buildup in the small scales, which eventually leads the simulation to blow up.

Inference costs The cost of the inference for our method as well as the baselines are summarized in Table 1. niLES uses four SDE samples for both training and inference, and each SDE sample is resolved using 16 uniformly-spaced time steps corresponding to a single LES time step. The inference cost scales linearly with both the number of samples and the temporal resolution. However, niLES achieves much lower inference cost than the DNS solver while having similarly accurate turbulent statistics for the resolved field. The deterministic NN model has slightly lower inference cost than niLES, since it can forego the SDE solves. While the implicit LES is the fastest method and is long-term stable, it cannot capture the statistics especially in the high wavenumber regime.

Limitations A drawback of our current approach stems from using transformers for the encoder-decoder phase of our model, which might result in poor scaling with increasing number of mesh

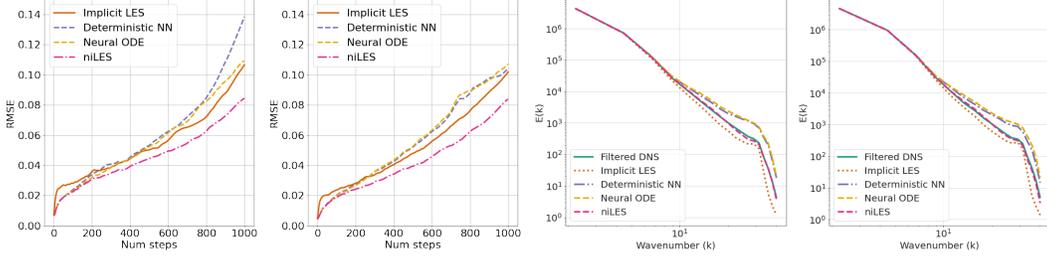


Figure 3: Root mean squared error (RMSE) over the first 1000 steps (first two columns) and the turbulent kinetic energy (TKE) spectrum $E(k)$ averaged over the first 2500 steps (right two columns) of two independent test trajectories unseen during training or validation. niLES has an improved ability to capture the long term statistics accurately compared to both implicit LES and deterministic NN. The energy buildup in the small scales (large wavenumber) in the deterministic NN model eventually leads to unstable trajectories.

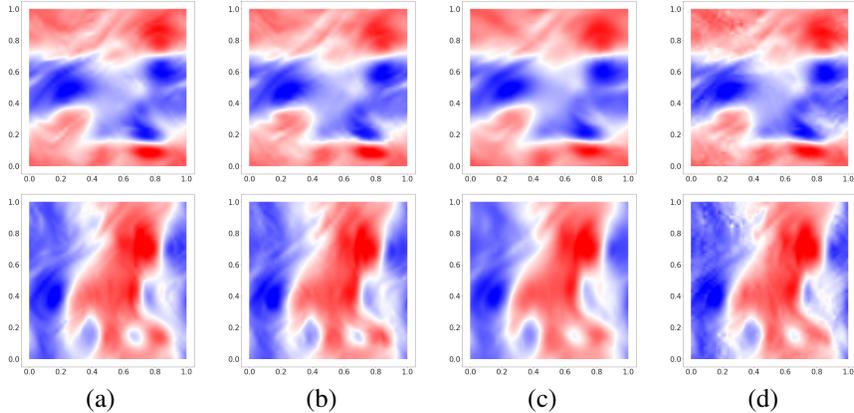


Figure 4: Comparison between rollout predictions after 800 LES steps on a held-out trajectory. Velocities in the x (top row) and y (bottom row) directions respectively. Snapshots of filtered DNS (reference) (a), niLES (b), implicit LES (c) and deterministic NN models (d). The niLES captures several finer scale features of the flow consistent with the reference filtered DNS trajectory. The implicit LES has an overall smoothing effect and some turbulent structures are not captured. The deterministic NN LES shows artifacts which indicate instability.

elements. Alternative architectures which can still handle interacting mesh elements should be explored. Additionally, the expressibility of the latent space in which the solution manifold needs to be embedded can affect the performance of our algorithm, and requires further study.

5 Related work

The relevant literature on closure modeling is extensive as it is one of the most classical topics in computational methods for science and engineering [68]. In recent years, there has also been an explosion of machine learning methods for turbulence modeling [9, 25] and multi-scaled systems in general. We loosely divide the related works into four categories, placing particular emphasis on the treatment of effects caused by unresolved (typically small-scaled) variables.

Classical turbulence methods primarily relies on phenomenological arguments to derive an *eddy viscosity* term [48], which is added to the physical viscosity and accounts for the dissipation of energy from large to small scales. The term may be static [4], time-dependent [76, 31] or multi-scale [39, 40].

Data-driven surrogates often do not model the closure in an explicit way. However, by learning the dynamics directly from data at finite resolution, the effects of unresolved variables and scales are expected to be captured implicitly and embedded in the machine learning models. A variety of architectures have been explored, including ones based on multi-scaled convolutional neural networks [70, 82, 77], transformers [11], graph neural networks [73, 49] and operator learning [65].

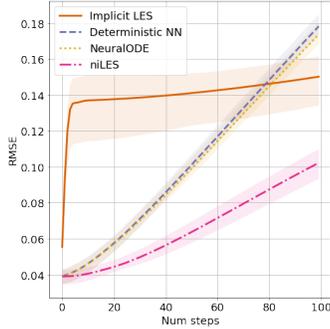


Figure 5: Average RMSE over 10 testcases on the cylinder wake dataset. Shaded regions depict the spread (minimum and maximum) among the testcases.

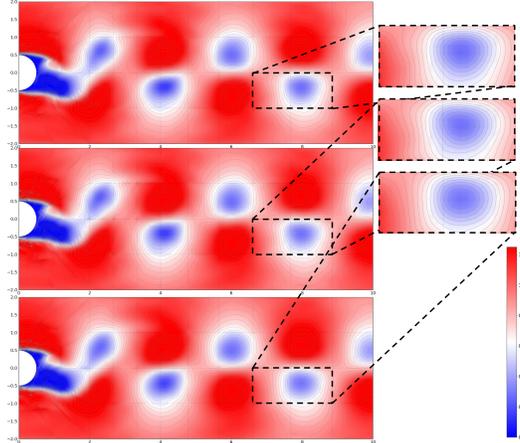


Figure 6: Comparison between rollout predictions of the horizontal component of velocity in the cylinder wake dataset after 100 LES steps on a held-out testcase. From top to bottom: snapshots of (a) filtered DNS (reference), (b) niLES, and (c) ILES. The niLES captures correct shapes of the vortices when compared to ILES. The latter has an excessive smoothing effect and the vortices appear more rounded.

Table 1: Inference times in wall clock seconds to evolve the state for one second of simulation time for the Kolmogorov dataset

| Method | DNS | Implicit LES | Deterministic NN | niLES (Ours) |
|--------------------|-------|--------------|------------------|--------------|
| Inference time [s] | 15600 | 8.85 | 11.4 | 23.8 |

Hybrid physics-ML contains a rich set of recent methods to combine classical numerical schemes and deep learning models [61, 6, 47, 56, 24, 81, 58, 35]. The former is expected to provide a reasonable baseline, while the latter specializes in capturing the interactions between modeled and unmodeled variables that accurately represent high-resolution data. This yields cost-effective, low-resolution methods that achieve comparable accuracy to more expensive simulations.

Probabilistic turbulence modeling seeks to represent small-scaled turbulence as stochastic processes [1, 32, 33, 20, 36]. Compared to their deterministic counterparts, these models are better equipped to model the backscattering of energy from small to large scales (i.e. opposite to the direction of energy flow in eddy viscosity), which is rare in occurrence but often associated with events of great practical interest and importance (e.g. rogue waves, extreme wildfires, etc.).

Our proposed method is inspired by methods in the last two categories. The closure model we seek includes the coarse solver in the loop while using a probabilistic generative model to emulate the stochastic process underlying the turbulent flow fields. This gives rise to long-term statistics that are accurate as the inexpensive neural SDE provides a good surrogate for the ensemble of the flows.

6 Conclusion

Due to chaotic divergence, it is infeasible to predict the state of the system with pointwise accuracy. Fortunately, in many systems of interest, the presence of an attractor allows a much lower-dimensional model to capture the essential statistics of the system. However, the time evolution of the attractor is not known, which makes building effective models challenging.

In this work we have argued that taking a probabilistic viewpoint is useful when modeling chaotic systems over long time windows. Our work has shown modeling the evolution with a neural SDE is beneficial in preserving long-term statistics and this line of thinking is likely fruitful.

References

- [1] Ronald J Adrian. On the role of conditional averages in turbulence theory. *Turbulence in Liquids*, pages 323–332, 01 1977.
- [2] Ronald J Adrian. *Stochastic estimation of the structure of turbulent fields*. Springer, 1996.
- [3] Ronald J Adrian, BG Jones, MK Chung, Yassin Hassan, CK Nithianandan, and AT-C Tung. Approximation of turbulent conditional averages by stochastic estimation. *Physics of Fluids A: Fluid Dynamics*, 1(6):992–998, 1989.
- [4] Nadine Aubry, Philip Holmes, John L Lumley, and Emily Stone. The dynamics of coherent structures in the wall region of a turbulent boundary layer. *Journal of fluid Mechanics*, 192:115–173, 1988.
- [5] Gary S Ayton, Will G Noid, and Gregory A Voth. Multiscale modeling of biomolecular systems: in serial and in parallel. *Current opinion in structural biology*, 17(2):192–198, 2007.
- [6] Yohai Bar-Sinai, Stephan Hoyer, Jason Hickey, and Michael P Brenner. Learning data-driven discretizations for partial differential equations. *Proceedings of the National Academy of Sciences*, 116(31):15344–15349, 2019.
- [7] Peter Bauer, Alan Thorpe, and Gilbert Brunet. The quiet revolution of numerical weather prediction. *Nature*, 525(7567):47–55, 2015.
- [8] Andrea Beck, David Flad, and Claus-Dieter Munz. Deep neural networks for data-driven les closure models. *Journal of Computational Physics*, 398:108910, 2019.
- [9] Andrea Beck and Marius Kurz. A perspective on machine learning methods in turbulence modeling. *GAMM-Mitteilungen*, 44(1):e202100002, 2021.
- [10] G Berkooz. An observation on probability density equations, or, when do simulations reproduce statistics? *Nonlinearity*, 7(2):313, 1994.
- [11] Kaifeng Bi, Lingxi Xie, Hengheng Zhang, Xin Chen, Xiaotao Gu, and Qi Tian. Pangu-weather: A 3d high-resolution model for fast and accurate global weather forecast. *arXiv preprint arXiv:2211.02556*, 2022.
- [12] Hugh M Blackburn and S Schmidt. Spectral element filtering techniques for large eddy simulation with dynamic estimation. *Journal of Computational Physics*, 186(2):610–629, 2003.
- [13] Guido Boffetta and Robert E Ecke. Two-dimensional turbulence. *Annual review of fluid mechanics*, 44:427–451, 2012.
- [14] Christoph Bosshard, Michel O Deville, Abdelouahab Dehbi, Emmanuel Leriche, et al. Udns or les, that is the question. *Open Journal of Fluid Dynamics*, 5(04):339, 2015.
- [15] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, et al. Jax: composable transformations of python+ numpy programs. 2018.
- [16] P Bradshaw. Turbulence modeling with application to turbomachinery. *Progress in Aerospace Sciences*, 32(6):575–624, 1996.
- [17] Johannes Brandstetter, Daniel Worrall, and Max Welling. Message passing neural pde solvers. *arXiv preprint arXiv:2202.03376*, 2022.
- [18] Ashesh Chattopadhyay, Jaideep Pathak, Ebrahim Nabizadeh, Wahid Bhimji, and Pedram Hassanzadeh. Long-term stability and generalization of observationally-constrained stochastic data-driven models for geophysical turbulence. *Environmental Data Science*, 2:e1, 2023.
- [19] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- [20] Carlo Cintolesi and Etienne Mémin. Stochastic modelling of turbulent flows for numerical simulations. *Fluids*, 5(3):108, 2020.
- [21] R. Courant, K. Friedrichs, and H. Lewy. On the partial difference equations of mathematical physics. *IBM Journal of Research and Development*, 11(2):215–234, 1967.

- [22] Daan Crommelin and Eric Vanden-Eijnden. Subgrid-scale parameterization with conditional markov chains. *Journal of the Atmospheric Sciences*, 65(8):2661–2675, 2008.
- [23] M. O. Deville, P. F. Fischer, and E. H. Mund. *High-Order Methods for Incompressible Fluid Flow*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, 2002.
- [24] Gideon Dresdner, Dmitrii Kochkov, Peter Norgaard, Leonardo Zepeda-Núñez, Jamie A Smith, Michael P Brenner, and Stephan Hoyer. Learning to correct spectral methods for simulating turbulent flows. *arXiv preprint arXiv:2207.00556*, 2022.
- [25] Karthik Duraisamy, Gianluca Iaccarino, and Heng Xiao. Turbulence modeling in the age of data. *Annual review of fluid mechanics*, 51:357–377, 2019.
- [26] Paul A Durbin. Some recent developments in turbulence closure modeling. *Annu. Rev. Fluid Mech.*, 50, 2018.
- [27] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6824–6835, 2021.
- [28] Paul Fischer and Julia Mullen. Filter-based stabilization of spectral element methods. *Comptes Rendus de l’Académie des Sciences-Series I-Mathematics*, 332(3):265–270, 2001.
- [29] Kai Fukami, Koji Fukagata, and Kunihiko Taira. Super-resolution reconstruction of turbulent flows with machine learning. *Journal of Fluid Mechanics*, 870:106–120, 2019.
- [30] Kai Fukami, Koji Fukagata, and Kunihiko Taira. Machine-learning-based spatio-temporal super resolution reconstruction of turbulent flows. *Journal of Fluid Mechanics*, 909:A9, 2021.
- [31] Massimo Germano, Ugo Piomelli, Parviz Moin, and William H Cabot. A dynamic subgrid-scale eddy viscosity model. *Physics of Fluids A: Fluid Dynamics*, 3(7):1760–1765, 1991.
- [32] Ian Grooms and Andrew J Majda. Efficient stochastic superparameterization for geophysical turbulence. *Proceedings of the National Academy of Sciences*, 110(12):4464–4469, 2013.
- [33] Ian Grooms and Andrew J Majda. Stochastic superparameterization in quasigeostrophic turbulence. *Journal of Computational Physics*, 271:78–98, 2014.
- [34] Yifei Guan, Ashesh Chattopadhyay, Adam Subel, and Pedram Hassanzadeh. Stable a posteriori les of 2d turbulence using convolutional neural networks: Backscattering analysis and generalization to higher re via transfer learning. *Journal of Computational Physics*, 458:111090, 2022.
- [35] Yifei Guan, Adam Subel, Ashesh Chattopadhyay, and Pedram Hassanzadeh. Learning physics-constrained subgrid-scale closures in the small-data regime for stable and accurate les. *Physica D: Nonlinear Phenomena*, 443:133568, 2023.
- [36] Arthur P Guillaumin and Laure Zanna. Stochastic-deep learning parameterization of ocean momentum forcing. *Journal of Advances in Modeling Earth Systems*, 13(9):e2021MS002534, 2021.
- [37] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *International conference on learning representations*, 2017.
- [38] Thomas JR Hughes, Gonzalo R Feijóo, Luca Mazzei, and Jean-Baptiste Quinicy. The variational multiscale method—a paradigm for computational mechanics. *Computer methods in applied mechanics and engineering*, 166(1-2):3–24, 1998.
- [39] Thomas JR Hughes, Luca Mazzei, and Kenneth E Jansen. Large eddy simulation and the variational multiscale method. *Computing and visualization in science*, 3:47–59, 2000.
- [40] Thomas JR Hughes, Assad A Oberai, and Luca Mazzei. Large eddy simulation of turbulent channel flows by the variational multiscale method. *Physics of fluids*, 13(6):1784–1799, 2001.
- [41] Javier Jimenez and Robert D Moser. Large-eddy simulations: where are we and what can we expect? *AIAA journal*, 38(4):605–612, 2000.
- [42] Myeongseok Kang, Youngmin Jeon, and Donghyun You. Neural-network-based mixed subgrid-scale model for turbulent flow. *Journal of Fluid Mechanics*, 962:A38, 2023.

- [43] Petr Karnakov, Sergey Litvinov, and Petros Koumoutsakos. Computing foaming flows across scales: From breaking waves to microfluidics. *Science Advances*, 8(5):eabm0590, 2022.
- [44] Patrick Kidger. On neural differential equations. *arXiv preprint arXiv:2202.02435*, 2022.
- [45] Patrick Kidger, James Foster, Xuechen Li, and Terry J Lyons. Neural sdes as infinite-dimensional gans. In *International Conference on Machine Learning*, pages 5453–5463. PMLR, 2021.
- [46] Patrick Kidger, James Foster, Xuechen Chen Li, and Terry Lyons. Efficient and accurate gradients for neural sdes. *Advances in Neural Information Processing Systems*, 34:18747–18761, 2021.
- [47] Dmitrii Kochkov, Jamie A Smith, Ayya Alieva, Qing Wang, Michael P Brenner, and Stephan Hoyer. Machine learning–accelerated computational fluid dynamics. *Proceedings of the National Academy of Sciences*, 118(21):e2101784118, 2021.
- [48] Andrei Nikolaevich Kolmogorov. The local structure of turbulence in incompressible viscous fluid for very large reynolds numbers. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, 434(1890):9–13, 1991.
- [49] Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Alexander Pritzel, Suman Ravuri, Timo Ewalds, Ferran Alet, Zach Eaton-Rosen, et al. Graphcast: Learning skillful medium-range global weather forecasting. *arXiv preprint arXiv:2212.12794*, 2022.
- [50] Jacob A Langford and Robert D Moser. Optimal les formulations for isotropic turbulence. *Journal of fluid mechanics*, 398:321–346, 1999.
- [51] CE Leith. Stochastic backscatter in a subgrid-scale model: Plane shear mixing layer. *Physics of Fluids A: Fluid Dynamics*, 2(3):297–299, 1990.
- [52] Xuechen Li, Ting-Kam Leonard Wong, Ricky TQ Chen, and David K Duvenaud. Scalable gradients and variational inference for stochastic differential equations. In *Symposium on Advances in Approximate Bayesian Inference*, pages 1–28. PMLR, 2020.
- [53] Yanghao Li, Chao-Yuan Wu, Haoqi Fan, Karttikeya Mangalam, Bo Xiong, Jitendra Malik, and Christoph Feichtenhofer. Mvitv2: Improved multiscale vision transformers for classification and detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4804–4814, 2022.
- [54] Yi Li, Eric Perlman, Minping Wan, Yunke Yang, Charles Meneveau, Randal Burns, Shiyi Chen, Alexander Szalay, and Gregory Eyink. A public turbulence database cluster and applications to study lagrangian evolution of velocity increments in turbulence. *Journal of Turbulence*, 9(9):N31, 2008.
- [55] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv:2010.08895 [cs, math]*, May 2021. arXiv: 2010.08895.
- [56] Björn List, Li-Wei Chen, and Nils Thuerey. Learned turbulence modelling with differentiable fluid solvers: physics-based loss functions and optimisation horizons. *Journal of Fluid Mechanics*, 949:A25, 2022.
- [57] Peter Lynch. *The emergence of numerical weather prediction: Richardson’s dream*. Cambridge University Press, 2006.
- [58] Jonathan F MacArt, Justin Sirignano, and Jonathan B Freund. Embedded training of neural-network subgrid-scale turbulence models. *Physical Review Fluids*, 6(5):050502, 2021.
- [59] Romit Maulik, Kai Fukami, Nesar Ramachandra, Koji Fukagata, and Kunihiko Taira. Probabilistic neural networks for fluid flow surrogate modeling and data recovery. *Physical Review Fluids*, 5(10):104401, 2020.
- [60] Romit Maulik, Omer San, Jamey D Jacob, and Christopher Crick. Sub-grid scale model classification and blending through deep learning. *Journal of Fluid Mechanics*, 870:784–812, 2019.
- [61] Siddhartha Mishra. A machine learning framework for data driven acceleration of computations of differential equations. *arXiv preprint arXiv:1807.09519*, 2018.
- [62] Steven A. Orszag. Analytical theories of turbulence. *Journal of Fluid Mechanics*, 41(2):363–386, 1970.
- [63] Shaowu Pan and Karthik Duraisamy. Data-driven discovery of closure models. *SIAM Journal on Applied Dynamical Systems*, 17(4):2381–2413, 2018.

- [64] Jonghwan Park and Haecheon Choi. Toward neural-network-based large eddy simulation: Application to turbulent channel flow. *Journal of Fluid Mechanics*, 914:A16, 2021.
- [65] Jaideep Pathak, Shashank Subramanian, Peter Harrington, Sanjeev Raja, Ashesh Chattopadhyay, Morteza Mardani, Thorsten Kurth, David Hall, Zongyi Li, Kamyar Azizzadenesheli, et al. Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. *arXiv preprint arXiv:2202.11214*, 2022.
- [66] Stephen B Pope. Pdf methods for turbulent reactive flows. *Progress in energy and combustion science*, 11(2):119–192, 1985.
- [67] Stephen B Pope. Ten questions concerning the large-eddy simulation of turbulent flows. *New Journal of Physics*, 6:35–35, March 2004.
- [68] Stephen B Pope and Stephen B Pope. *Turbulent flows*. Cambridge university press, 2000.
- [69] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *International Conference on Machine Learning*, pages 5301–5310. PMLR, 2019.
- [70] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [71] Andrew Ross, Ziwei Li, Pavel Perezhogin, Carlos Fernandez-Granda, and Laure Zanna. Benchmarking of machine learning ocean subgrid parameterizations in an idealized model. *Journal of Advances in Modeling Earth Systems*, 15(1):e2022MS003258, 2023.
- [72] Pierre Sagaut. *Large eddy simulation for incompressible flows: an introduction*. Springer Science & Business Media, 2005.
- [73] Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *International conference on machine learning*, pages 8459–8468. PMLR, 2020.
- [74] Tapio Schneider, João Teixeira, Christopher S Bretherton, Florent Brient, Kyle G Pressel, Christoph Schär, and A Pier Siebesma. Climate goals and computing the future of clouds. *Nature Climate Change*, 7(1):3–5, 2017.
- [75] Claude Elwood Shannon. Communication in the presence of noise. *Proceedings of the Institute of Radio Engineers*, 37(1):10–21, 1949.
- [76] Joseph Smagorinsky. General circulation experiments with the primitive equations: I. the basic experiment. *Monthly weather review*, 91(3):99–164, 1963.
- [77] Kimberly Stachenfeld, Drummond B Fielding, Dmitrii Kochkov, Miles Cranmer, Tobias Pfaff, Jonathan Godwin, Can Cui, Shirley Ho, Peter Battaglia, and Alvaro Sanchez-Gonzalez. Learned coarse models for efficient turbulence simulation. *arXiv preprint arXiv:2112.15275*, 2021.
- [78] Adam Subel, Ashesh Chattopadhyay, Yifei Guan, and Pedram Hassanzadeh. Data-driven subgrid-scale modeling of forced burgers turbulence using deep learning with generalization to higher reynolds numbers via transfer learning. *Physics of Fluids*, 33(3):031702, 2021.
- [79] Salar Taghizadeh, Freddie D Witherden, and Sharath S Girimaji. Turbulence closure modeling with data-driven techniques: physical compatibility and consistency considerations. *New Journal of Physics*, 22(9):093023, 2020.
- [80] Belinda Tzen and Maxim Raginsky. Neural stochastic differential equations: Deep latent gaussian models in the diffusion limit. *arXiv preprint arXiv:1905.09883*, 2019.
- [81] Kiwon Um, Robert Brand, Yun Raymond Fei, Philipp Holl, and Nils Thuerey. Solver-in-the-loop: Learning from differentiable physics to interact with iterative pde-solvers. *Advances in Neural Information Processing Systems*, 33:6111–6122, 2020.
- [82] Rui Wang, Karthik Kashinath, Mustafa Mustafa, Adrian Albert, and Rose Yu. Towards physics-informed deep learning for turbulent flow prediction. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1457–1466, 2020.
- [83] Winnie Xu, Ricky TQ Chen, Xuechen Li, and David Duvenaud. Infinitely deep bayesian neural networks with stochastic differential equations. In *International Conference on Artificial Intelligence and Statistics*, pages 721–738. PMLR, 2022.