Sharper predictions: The role of loss functions for enhanced turbulent-flow sensing

A. G. Balasubramanian [*]
FLOW, Dept. Engineering Mechanics, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden and Swedish e-Science Research Centre (SeRC), SE-100 44 Stockholm, Sweden

A. Cremades
Instituto Universitario de Matemática Pura y Aplicada, Universitat Politécnica de Valéncia, Valencia 46022, Spain

R. Vinuesa
Department of Aerospace Engineering, University of Michigan, Ann Arbor, Michigan 48109, USA

O. Tammisola [†]
FLOW, Dept. Engineering Mechanics, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden and Swedish e-Science Research Centre (SeRC), SE-100 44 Stockholm, Sweden

(Received 15 December 2025; accepted 9 March 2026; published 16 April 2026)

Accurate estimation of near-wall turbulence from limited surface measurements is critical for various practical and experimental applications; however, it remains challenging due to the complex and multiscale characteristics inherent to turbulent flows. Recent developments in data-driven methods have improved performance over traditional linear approaches but are often limited by loss functions that primarily penalize pointwise errors. In this study, we consider a composite, spectrally informed loss function that augments the conventional mean-squared error with terms that explicitly promote statistical consistency of fluctuation levels and spectral energy distributions. This loss function is evaluated using a baseline convolutional network model applied to direct numerical simulation datasets of turbulent open-channel flow at friction Reynolds numbers, $\mathrm{Re}_\tau = 180$ and 550. We show that the considered loss function substantially reduces reconstruction errors of near-wall fluctuation velocity fields obtained from the network model using wall-shear stress and wall-pressure data as inputs. The proposed method recovers up to a threefold improvement in reconstruction accuracy relative to baseline employing mean-squared loss and retains energy content at small-scale structures. We further show that reconstruction accuracy is only mildly affected when the input wall signals are contaminated with 5–100% Gaussian noise and retains energy content in small flow structures when predicting from coarse wall inputs, indicating that the trained models generalize robustly to noisy, experimentally relevant conditions. Our results demonstrate that properly designed loss functions can improve reconstruction accuracy, highlighting the importance of training strategy in

[*]Contact author: argb@mech.kth.se
[†]Contact author: outi@mech.kth.se

achieving efficient, high-performance neural-network models for practical nonintrusive sensing applications. The present study also provides a foundation for future extensions with advanced architectures, such as generative or diffusion-based models to further improve reconstruction of turbulent flow fields in complex settings.

DOI: 10.1103/26js-tpg4

I. INTRODUCTION

Turbulence remains one of the unsolved problems in classical physics. Although the Navier–Stokes equations provide a mathematical description of fluid motion, the inertial effects give rise to chaotic, multiscale dynamics that makes turbulence prediction extremely challenging. The resulting wide range of interacting length and timescales poses substantial difficulties for both numerical simulations and experimental measurements. Direct numerical simulations (DNSs) can provide accurate datasets, but their computational cost scales steeply with the Reynolds number. Likewise, wall-bounded turbulence introduces additional complexity due to the steep gradients and small structures that emerge in the near-wall region, making accurate flow characterization particularly challenging in practical and experimental settings.

Recent advances in machine learning (ML), particularly in deep learning, have revolutionized turbulence research by offering powerful data-driven alternatives to traditional physics-based modeling [1–3]. Deep neural networks (DNNs) can approximate highly nonlinear relationships, enabling the modeling of complex flow dynamics that are otherwise difficult to capture through conventional analytical or empirical methods [4–6]. ML-based approaches have demonstrated success in diverse areas such as turbulence closure modeling [1,7] for Reynolds-averaged Navier-Stokes [8–10] and large-eddy simulations [11–14], turbulent flow reconstruction [15,16], superresolution [17–19], and low-order turbulence modeling [20–23]. Physics-informed neural networks, in particular, have provided a framework for embedding governing physical laws into learning algorithms, enhancing physical consistency and generalization [10,24,25]. These approaches have enabled superresolution of noisy or sparse flow data [26,27], and reconstruction of physically consistent fields in experimental and numerical settings [28].

Beyond simulation and modeling, ML methods have greatly expanded capabilities in prediction and control of turbulent flows. Temporal evolution of wall-bounded turbulence has been modeled using recurrent neural networks and autoencoder-based architectures [29–31], enabling accurate forecasting of near-wall dynamics. Deep reinforcement learning and control frameworks have also been used to design real-time strategies for drag reduction and flow manipulation [32–35]. Moreover, interpretable deep-learning methods [36] have begun to provide alternative physical insights into turbulence, such as identifying dynamically significant coherent structures and quantifying their relative role in momentum transfer [3,37]. These developments collectively demonstrate how ML can bridge data-driven learning with fundamental turbulence physics.

In wall-bounded turbulence, deep learning has been particularly successful for nonintrusive sensing (see Fig. 1 for illustration), where the goal is to infer off-wall flow fields from limited wall measurements. Fully convolutional networks have been used to predict instantaneous velocity fluctuations from wall-shear stress and wall-pressure distributions [15], significantly outperforming linear techniques such as linear stochastic estimation and extended proper orthogonal decomposition. Generative adversarial networks (GANs) have been further employed for superresolution and enhancement of turbulent flow fields [18], improving reconstruction accuracy even with coarse or noisy wall inputs. An inverse problem, which is to predict wall quantities from off-wall velocity fields, has also been explored [16] toward wall modeling. Lagemann *et al.* [38] proposed a deep-learning framework that predicts wall-shear stress distributions from velocity fields in the logarithmic layer, achieving zero-shot generalization to experimental data. Such approaches complement nonintrusive sensing by offering data-driven strategies for wall modeling and under-

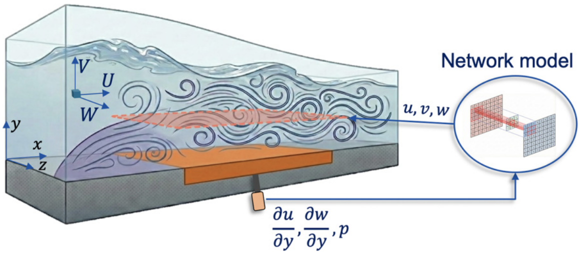

FIG. 1. Illustration of nonintrusive sensing in turbulent flows using network models. Here, the wall measurements such as wall-shear rate components ($\partial u/\partial y$, $\partial w/\partial y$) and wall-pressure (p) data are used to predict the near-wall velocity fluctuation fields (u, v, w).

standing inner–outer layer interactions in wall turbulence. Together, these developments underline the growing synergy between ML and turbulence physics, setting the foundation for more accurate, generalizable, and interpretable flow-reconstruction methodologies.

Despite these advances, accurately estimating near-wall turbulent structures from limited wall data remains an open challenge. The nonlinear coupling between wall-bounded features, such as wall-shear stress, wall pressure, and the velocity fluctuations away from the wall, demands network architectures capable of preserving both local and global flow dynamics. Furthermore, the use of conventional mean squared error (MSE) loss functions often fails to enforce the spectral fidelity [39,40] of the reconstructed turbulence, leading to errors in small-scale structures and flow statistics.

In this context, the present study considers an efficient, loss-function–based framework for the reconstruction of near-wall velocity fluctuation fields from wall-based measurements. Rather than focusing on network complexity, we design a composite loss (CL) that augments standard MSE with spectral and statistical constraints to better capture the multiscale nature of turbulence and reduce prediction errors. This framework is tested using fully convolutional networks as a representative baseline architecture, but it is, by construction, applicable to a broad class of neural-network models including generative and diffusion-based models. Using DNS datasets of turbulent open-channel flows at two Reynolds numbers, we systematically assess the performance of this composite-loss formulation in predicting near-wall velocity fluctuations from wall-shear and wall-pressure data.

The present study highlights that properly designed loss functions can improve reconstruction accuracy while maintaining architectural simplicity (without introducing huge amount of trainable parameters) and computational efficiency. In addition, the trained networks exhibit improved performance across Reynolds numbers, highlighting the importance of targeted training strategies [41] over network complexity. Moreover, the modified loss function generalizes well when the inputs are contaminated with Gaussian noise, mimicking the measurement uncertainty typical of experimental campaigns. Further to account for the limited spatial resolution of wall measurements in practical setups, we also evaluate the methodology using spatially coarsened wall inputs where models trained with the spectral loss (SL) is shown to consistently outperform those trained with a pure MSE loss. Overall, the evaluated framework provides a promising foundation for nonintrusive turbulence sensing and for integrating data-driven models into experimental diagnostics and flow-control applications.

II. METHODOLOGY

A. Datasets

The neural networks employed in this study are trained using datasets obtained from direct numerical simulations of a turbulent open-channel flow. The simulations were performed using SIMSON, a pseudospectral solver [42] where nonlinear terms in the Navier–Stokes equations are evaluated in the physical space using the 3/2-rule (for dealiasing) and transformed to the physical space using fast Fourier transform. Time advancement is carried out using a second-order Crank-Nicolson scheme for the linear terms and a third-order Runge-Kutta scheme for the nonlinear terms, with a constant time step Δt. More details on numerical simulations can be found in Guastoni *et al.* [15].

The computational domain Ω considered in this work is given by $L_x \times L_y \times L_z = 4\pi h \times h \times 2\pi h$, where h is the channel height and L_x, L_y, L_z correspond, respectively, to the length of the computational domain in streamwise (x), wall-normal (y), and spanwise (z) directions. Instantaneous velocity fields in streamwise, wall-normal, and spanwise directions are, respectively, denoted as $\boldsymbol{U}$, $\boldsymbol{V}$, $\boldsymbol{W}$ with the velocity field $\boldsymbol{U}$ decomposed into $\langle \boldsymbol{U} \rangle + \boldsymbol{u}$, denoting the ensemble mean and corresponding fluctuations. Periodic boundary conditions are imposed in the streamwise and spanwise directions, allowing the use of Fourier representations in the wall-parallel directions. In the wall-normal direction, where the flow is nonperiodic, Chebyshev polynomials $T_n(\xi)$ are used on the interval $-1 \leqslant \xi \leqslant 1$. The no-slip condition $\mathbf{U} = 0$ is applied at $y = 0$ and symmetry boundary condition at $y = h$. Two different friction Reynolds numbers are considered in this work at $\mathrm{Re}_\tau = 180$ and $\mathrm{Re}_\tau = 550$ [$\mathrm{Re}_\tau := hu_\tau/\nu$, based on the open-channel height h and friction velocity $u_\tau := \sqrt{\tau_w/\rho}$, with τ_w corresponding to wall-shear stress and ρ, ν indicating, respectively, the density and kinematic viscosity of the fluid]. The resolution of the domain for $\mathrm{Re}_\tau = 180$ is $N_x \times N_y \times N_z = 192 \times 65 \times 192$, and for $\mathrm{Re}_\tau = 550$ it is $= 512 \times 193 \times 512$, where N_x, N_y, N_z correspond to the number of grid points in the streamwise, wall-normal, and spanwise directions, respectively.

Training and testing samples were collected only after the DNS reached a statistically stationary state [43]. Velocity fields were sampled at several inner-scaled wall-normal locations corresponding to $y^+ = 30, 50, 100$. Note that "+" denotes the scaling with friction velocity u_τ and the fluid kinematic viscosity ν. Additionally, the wall-shear stress components in the streamwise and spanwise directions, as well as the wall pressure ($y^+ = 0$), were recorded. Throughout this work, all the sampled fields are ensemble-mean subtracted and their corresponding fluctuation components are considered for the inputs and predictions. To generate statistically independent datasets, multiple DNS runs were initiated with different random seeds for the initial velocity field. The flow fields were stored at uniform temporal intervals of $\Delta t_s^+ = 5.08$ for $\mathrm{Re}_\tau = 180$ and $\Delta t_s^+ = 1.49$ for $\mathrm{Re}_\tau = 550$. In total, approximately 25 200 and 19 920 snapshots were collected for training the networks at $\mathrm{Re}_\tau = 180$ and $\mathrm{Re}_\tau = 550$, respectively. For validation, around 4000 samples were used and the total number of samples in the test dataset N_t is taken to be 4000, which is found sufficient to achieve statistical convergence in turbulence statistics. Using fewer test samples may yield comparable instantaneous prediction accuracy (in terms of mean squared error) but could result in deviations in turbulence statistics such as the root-mean-square (RMS) profiles of velocity fluctuations due to inadequate representation of the underlying stochastic process.

B. Fully Convolutional Neural Networks

Artificial neural networks are computational models inspired by biological neural systems, consisting of interconnected nodes (neurons) that process and transmit information. Deep neural networks, composed of multiple layers of neurons, can capture highly nonlinear relationships between input and output data. A common form of DNNs, the multilayer perceptron, processes flattened input data through fully connected layers, but this approach becomes computationally expensive and inefficient for large, spatially correlated datasets such as velocity fields.

In contrast, convolutional neural networks (CNNs) employ convolutional layers that exploit the spatial structure of the input. A convolution operation between an input field $\mathbf{I} \in \mathbb{R}^{d_1 \times d_2}$ and a learnable kernel $\mathbf{K} \in \mathbb{R}^{k_1 \times k_2}$ produces a feature map $\mathbf{F}$:

$$F_{i,j} = \sum_m \sum_n I_{i-m,j-n} K_{m,n}, \tag{1}$$

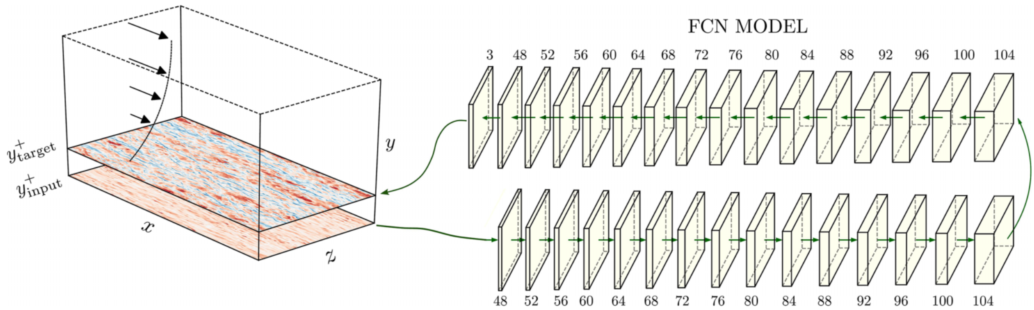

FIG. 2. Workflow representation of fully convolutional network model. Left: Computational domain for the open-channel flow. Right: FCN model with the corresponding number of kernels in each hidden layer is indicated.

where the kernel size and number of kernels determine the layer's receptive field and feature extraction capacity. This operation efficiently identifies spatial patterns such as coherent flow structures, while drastically reducing computational cost compared to dense connections.

In this study, a fully convolutional network (FCN) is used similar to Guastoni *et al.* [15], where all layers are convolutional, making it suitable for mapping two-dimensional (2D) input fields (e.g., wall quantities) to two-dimensional outputs (velocity fluctuations). Nonlinear activation function, specifically, the Rectified Linear Unit (ReLU), is applied after each convolution to enable the network to model complex dependencies.

Neural networks learn the approximation of the conditional distribution of the output given the input by minimizing the error between predictions and reference values (ground truth), using training data sampled from the underlying joint distribution. The training process is carried out using gradient-descent optimization, where the backpropagation algorithm computes gradients to iteratively update the network parameters. In essence, training the network model involves adapting the model parameters to improve predictive accuracy based on sample observations. When corresponding target data are available, this approach is categorized as supervised learning. Each input–output pair constitutes a sample, and a collection of such samples forms a dataset, which is typically divided into training, validation, and testing subsets. The training data is used to optimize the network parameters, the validation data to monitor the generalization performance during training, and the test data to evaluate the predictive accuracy. In this work, the training dataset is divided into a 4:1 ratio between training and validation samples.

The architecture of a neural network is determined by its structure and the associated hyperparameters, such as kernel size, number of layers, and learning rate. These hyperparameters are fixed before training and not updated during optimization. In this study, we employ the FCN architecture introduced by Guastoni *et al.* [15], with additional hidden layers to enhance its ability to capture complex turbulent features as used by Balasubramanian *et al.* [16] (see Fig. 2 for representation of FCN). The total number of trainable parameters is 1 626 843. The selected hyperparameters, summarized in Table I, were kept consistent across all cases to isolate the influence of the training strategies and only the loss function is varied in this study. The focus of this work is not architectural optimization but demonstrating how suitable training methodologies such as the use of spectral loss functions, which can substantially improve performance in data-driven, nonintrusive turbulence sensing applications.

It is important to note that the input fields $\mathbf{I}$ used in the network have dimensions of 256×256, which are larger than 192×192 as sampled from the DNS data for $\text{Re}_\tau = 180$. This increase is due to periodic padding applied to the inputs [15] (see also [44] for discussion on effects of padding operation in the CNN model). As valid convolutional operations inherently reduce the spatial dimensions of feature maps, padding is required to preserve the desired output size. The

TABLE I. Hyperparameters for the network.

Parameters	Values
Kernel size, $k_1 \times k_2$	3×3
Weight initialization	Gaussian
Optimization algorithm	Adam
Batch size, N_b	32 (for $\mathrm{Re}_\tau = 180$) 8 (for $\mathrm{Re}_\tau = 550$)
Activation function	ReLU
Learning rate drop a	0.5
Initial learning rate α_o	0.001
Epoch drop b	40

receptive field of the network, which corresponds to the portion of the input influencing a single output element, determines the extent of this padding. For the FCN model, the receptive field spans 63×63 points. Consequently, the padded input size for $\mathrm{Re}_\tau = 180$ increases by approximately 77%, whereas for $\mathrm{Re}_\tau = 550$, the increase is about 26.5% with respect to size of the DNS fields. In this work, we apply periodic padding only to the input fields and use valid convolutions in all subsequent layers. For unit stride and dilation, this is equivalent (on the retained interior domain) to applying periodic padding at every convolutional layer, since the required padding width equals the effective receptive-field radius of the network model for the considered kernel size.

Batch normalization is applied between consecutive layers to stabilize the training process by normalizing the input distributions for each layer. Within the FCN, multiple feature maps are generated at each stage and passed collectively to subsequent layers, enabling the network to learn complex spatial correlations between input and output fields. The network is trained using the Adam optimization algorithm [45], incorporating a scheduled exponential learning rate decay expressed as $\alpha = \alpha_o a^{\lfloor \mathrm{epoch}/b \rfloor}$, where α_o is the initial learning rate and a, b are tunable hyperparameters. The network models are implemented using the TensorFlow/Keras framework. More details on the obtained learning curves and associated discussions are found in Appendix A. The input and predicted fluctuation velocity fields are normalized by their respective RMS values to ensure that the learning process is independent of the magnitude of the fluctuations at a particular target wall-normal position.

1. Mean-squared error as loss function

Following the approach employed in previous studies, including the baseline FCN framework of Guastoni *et al.* [15], the loss function used for training is the mean-squared error (denoted MSL) between the instantaneous FCN prediction and the reference DNS field, defined as

$$\mathcal{L}_{\mathrm{MSL}} = \frac{1}{N_b N_x N_z} \sum_{k=1}^{N_b} \sum_{i=1}^{N_x} \sum_{j=1}^{N_z} \left\| \mathbf{u}^k_{\mathrm{MSL}}(i,j) - \mathbf{u}^k_{\mathrm{DNS}}(i,j) \right\|_2^2, \qquad \mathbf{u} = [\boldsymbol{u}, \boldsymbol{v}, \boldsymbol{w}]^T. \tag{2}$$

where averaging over the streamwise and spanwise directions is represented by the indices i and j, which will be dropped in the following sections for brevity. The DNS fields are denoted by $\mathbf{u}_{\mathrm{DNS}}$ and the corresponding predictions from the FCN, when employing MSE as the loss function, are denoted by $\mathbf{u}_{\mathrm{MSL}}$. Here, $\|\cdot\|_2$ denotes the Euclidean norm in component space, i.e., $\|\mathbf{a}\|_2^2 = a_u^2 + a_v^2 + a_w^2$ for $\mathbf{a} = [a_u, a_v, a_w]^T$.

The MSE loss [as introduced in Eqs. (2)] penalizes pointwise differences between predicted and reference fields. This form of loss encourages the network model to minimize the difference in turbulence kinetic energy and thereby to capture large-scale flow features first and progressively refine small-scale structures as training proceeds [46]. However, as noted in recent studies on turbulence reconstruction using convolutional networks [15,47], MSE-based objectives can bias the learning process toward minimizing pixel-level differences, often resulting in smoothed predictions and reduced accuracy in capturing energetic small-scale fluctuations.

2. *Composite loss function*

To address these limitations, we augment the MSE objective with additional physically motivated terms, resulting in the CL function:

$$\mathcal{L}_{\mathrm{CL}} = \alpha \frac{1}{N_b} \sum_{k=1}^{N_b} \left\| \mathbf{u}_{\mathrm{CL}}^k - \mathbf{u}_{\mathrm{DNS}}^k \right\|_2^2 + \beta(\tilde{u}_{\mathrm{RMS,CL}} - \tilde{u}_{\mathrm{RMS,DNS}})^2 + \gamma(1 - \rho(\mathbf{u}_{\mathrm{CL}}, \mathbf{u}_{\mathrm{DNS}}))^2. \tag{3}$$

Here, α, β, and γ are hyperparameters controlling the relative importance of each term, and $\mathbf{u}_{\mathrm{CL}}$ denotes the FCN prediction obtained when training with the composite loss. The batch RMS is defined as

$$\tilde{u}_{\mathrm{RMS,DNS}} = \sqrt{\frac{1}{N_b} \sum_{k=1}^{N_b} \left\| \mathbf{u}_{\mathrm{DNS}}^k \right\|_2^2}, \qquad X \in \{\mathrm{DNS}, \mathrm{CL}\}, \tag{4}$$

and the batch correlation coefficient is defined as

$$\rho(\mathbf{u}_{\mathrm{CL}}, \mathbf{u}_{\mathrm{DNS}}) = \frac{\frac{1}{N_b} \sum_{k=1}^{N_b} \left(\mathbf{u}_{\mathrm{DNS}}^k - \bar{\mathbf{u}}_{\mathrm{DNS}}\right) \cdot \left(\mathbf{u}_{\mathrm{CL}}^k - \bar{\mathbf{u}}_{\mathrm{CL}}\right)}{\tilde{u}_{\mathrm{RMS,DNS}}\, \tilde{u}_{\mathrm{RMS,CL}}}, \quad \bar{\mathbf{u}}_X = \frac{1}{N_b} \sum_{k=1}^{N_b} \mathbf{u}_X^k. \tag{5}$$

The first term in Eq. (3) corresponds to the conventional MSE contribution, ensuring accurate spatial reconstruction at the grid-point level. The second term enforces consistency between the overall amplitude of the predicted and reference velocity fluctuations, thereby constraining the energy content of the reconstructed field. The third term accounts for the spatial alignment between the predicted and DNS structures by maximizing their correlation coefficient. This term effectively penalizes phase errors and misalignments in flow structures.

The inclusion of amplitude and correlation-based penalties promotes the preservation of both the statistical and structural integrity of the turbulent flow, complementing the MSE-driven spatial fidelity. Such a composite loss function balances the competing objectives of local accuracy, global energy consistency, and structural coherence. In this study, although the loss coefficients were test-tuned to achieve lower errors in reconstructed fields, it was found that even just setting $\alpha = \beta = \gamma = 1$ is sufficient to obtain good improvements in results in comparison to MSE as loss function (see Appendix B for more details). Since only minor improvements [around 2% relative points in reconstruction accuracy as measured by E_{RMS} for prediction of streamwise velocity fluctuations at $y^+ = 50$; see Eq. (9) for definition of E_{RMS}] were obtained by tuning α, β, γ, and we report the results obtained with equal weight to all loss terms.

3. *Spectral loss function*

In the preceding section, we formulated a composite loss function combining mean-squared error, fluctuation amplitude matching, and structural alignment terms. Although such a formulation already encourages the network to preserve both pointwise accuracy and global statistical consistency, it remains fundamentally tied to spatial-domain errors. To further enhance the reconstruction, especially at smaller scales and in preserving sharp gradients, we include additional spectral-domain loss terms, inspired by image-processing practices in machine learning, to guide the network in the frequency domain.

Rather than measuring error purely in pixel (or spatial) difference, one can transform both predicted and reference fields into the frequency (Fourier) domain [$\hat{u} = \mathcal{F}(u)$, where $\mathcal{F}$ corresponds to discrete Fourier transform operator] and compare the amplitude or phase of each Fourier mode. Yan *et al.* [48] argues that taking the squared difference between Fourier transforms is mathematically equivalent to a spatial MSE loss (for periodic signals and where $\mathcal{F}$ is unitary), but that comparing amplitude spectra (magnitude of the Fourier coefficients) is distinct and can help preserve high-frequency content (fine structures) while being more shift invariant and sensitive to blur. In other words, using SL encourages the network to reproduce the correct energy distribution across spatial frequencies, which complements spatial-domain supervision. Recent works in super-resolution and image reconstruction have leveraged such frequency-domain guidance to ameliorate blurring artifacts and to boost preservation of texture and edges [49]. In the context of turbulent flow reconstruction, the high-wave-number (small-scale) components are especially challenging to capture [15]. Because convolutional networks and MSE-based losses tend to favor smooth solutions, adding a spectral loss term can force the network to more faithfully reconstruct energy across a broad range of wavenumbers.

We further extend the composite loss [Eq. (3)] by adding two additional terms. The first is a gradient-consistency penalty, $\mathcal{G}(\mathbf{u}_{\mathrm{SL}}, \mathbf{u}_{\mathrm{DNS}})$, which discourages discrepancies in spatial derivatives in physical space [50,51]. The second is a spectral-amplitude discrepancy term, $\mathcal{H}(\mathbf{u}_{\mathrm{SL}}, \mathbf{u}_{\mathrm{DNS}})$, which promotes agreement between the magnitudes of the Fourier coefficients of the predicted and reference velocity-fluctuation fields. The FCN predictions obtained when training with this extended objective are denoted by $\mathbf{u}_{\mathrm{SL}}$, and we refer to the resulting objective as the SL:

$$\mathcal{L}_{\mathrm{SL}} = \alpha \frac{1}{N_b} \sum_{k=1}^{N_b} \left\| \mathbf{u}_{\mathrm{SL}}^k - \mathbf{u}_{\mathrm{DNS}}^k \right\|_2^2 + \beta(\tilde{u}_{\mathrm{RMS,SL}} - \tilde{u}_{\mathrm{RMS,DNS}})^2 + \gamma(1 - \rho(\mathbf{u}_{\mathrm{SL}}, \mathbf{u}_{\mathrm{DNS}}))^2 + \delta\, \mathcal{G}(\mathbf{u}_{\mathrm{SL}}, \mathbf{u}_{\mathrm{DNS}}) + \zeta\, \mathcal{H}(\mathbf{u}_{\mathrm{SL}}, \mathbf{u}_{\mathrm{DNS}}). \tag{6}$$

Here, δ and ζ are additional hyperparameters controlling the relative importance of the gradient and spectral terms, respectively. The gradient-consistency term is defined as

$$\mathcal{G}(\mathbf{u}_{\mathrm{SL}}, \mathbf{u}_{\mathrm{DNS}}) = \frac{1}{N_b} \sum_{k=1}^{N_b} \left\| \nabla \mathbf{u}_{\mathrm{SL}}^k - \nabla \mathbf{u}_{\mathrm{DNS}}^k \right\|_F^2, \tag{7}$$

where ∇ denotes a finite-difference operator applied in the streamwise and spanwise directions, and $\| \cdot \|_F$ is the Frobenius norm (sum of squares over components and derivative directions). The spectral-amplitude term is defined as

$$\mathcal{H}(\mathbf{u}_{\mathrm{SL}}, \mathbf{u}_{\mathrm{DNS}}) = \frac{1}{N_b} \sum_{k=1}^{N_b} \left\| \left|\mathcal{F}(\mathbf{u}_{\mathrm{SL}}^k)\right| - \left|\mathcal{F}(\mathbf{u}_{\mathrm{DNS}}^k)\right| \right\|_1, \tag{8}$$

where $\mathcal{F}(\cdot)$ denotes the two-dimensional Fourier transform in the homogeneous directions and $|\cdot|$ indicates the elementwise magnitude of the Fourier coefficients.

In summary, by adding gradient and spectral discrepancy terms, the network model is encouraged not only to match field values at individual points but also to honor the frequency distribution and gradient structure of turbulence. This combination helps produce reconstructions with faithful fine-scale dynamics and improved energy spectra, which are essential for nonintrusive turbulence sensing.

The network model as shown in Fig. 2 is trained to predict the velocity fluctuation fields at $y^+ = 30, 50, 100$. The training of this model is run for at least 75 epochs, with each epoch consisting of the update of the model parameters using all the data in the training set. More details on training network models are provided in Appendix A.

The performance of the network models are evaluated in terms of the relative percentage error in predicting the corresponding root-mean-squared quantities between the true (DNS) fields from the test dataset and the fields predicted by the FCN (indicated by E_{RMS}), and is given by

$$E_{\mathrm{RMS}}(c_{\mathrm{pred,RMS}})\ [\%] = 100 \cdot \frac{|c_{\mathrm{pred,RMS}} - c_{\mathrm{DNS,RMS}}|}{c_{\mathrm{DNS,RMS}}},\ \ c \in \{u, v, w\}, \tag{9}$$

whereas the instantaneous correlation coefficient between the predicted and the DNS fields is defined as

$$R(c_{\mathrm{pred}}) = \frac{\langle c_{\mathrm{pred}} c_{\mathrm{DNS}} \rangle_{x,z,t}}{c_{\mathrm{pred,RMS}} c_{\mathrm{DNS,RMS}}},\ \ c \in \{u, v, w\}, \tag{10}$$

with $\langle \cdot \rangle_{x,z,t}$ corresponding to the average in space (x, z) and time (t denotes the N_t samples in the test dataset) and subscript $\cdot_{\mathrm{RMS}}$ refering to root-mean-squared quantities. Here, c denotes a velocity component (u, v, w), and pred refers to MSL/CL/SL predictions, and so $u_{\mathrm{MSL,RMS}}$, $u_{\mathrm{CL,RMS}}$, $u_{\mathrm{SL,RMS}}$, respectively, correspond to the RMS of the velocity components obtained from FCN employing MSE, CL, and SL as loss functions with the test dataset.

Note that the performance metrics reported in this study are obtained from the mean of at least three different network models to include the effects of stochasticity introduced by the random initialization of kernel weights in FCN and random sampling of minibatches during the training process. The instantaneous correlation coefficient between the predicted and DNS fields averaged over the samples in the test dataset is also highlighted. To evaluate the distribution of energy across the various scales in the flow, a comparison of the premultiplied 2D power-spectral density $k_z k_x \phi_{ij}(\lambda_x^+, \lambda_z^+)$ between DNS fields and the predictions is performed. Here, ϕ_{ij} is the power-spectral density obtained for the quantity ij and k_x, k_z, respectively, denoting the wavenumbers in streamwise and spanwise directions, with the corresponding wavelengths in viscous units denoted by λ_x^+ and λ_z^+.

C. Predictions with noisy inputs

To assess the robustness and generalization capacity of our neural-network models, we consider an additional testing scenario in which the input wall-based fields $\mathbf{I}$ are contaminated with additive Gaussian noise with an amplitude of 5%, 25%, and 100% of their respective RMS values. Specifically, the noisy input is constructed as $\tilde{\mathbf{I}} = \mathbf{I} + \sigma_{\mathrm{noise}}\mathbf{n}$, where $\mathbf{n}$ has the same dimensions as $\mathbf{I}$ and its entries are sampled i.i.d. from a standard normal distribution, i.e., $n_i \sim \mathcal{N}(0, 1)$. For a 5% noise level, we set $\sigma_{\mathrm{noise}} = 0.05\,\mathbf{I}_{\mathrm{RMS}}$ (and analogously for the other noise amplitudes). The target fields are the reference DNS velocity fluctuations at $y^+ = 50$. The resulting prediction quality is compared across models trained with different loss functions, using the noise-free case $\sigma_{\mathrm{noise}} = 0$ as a baseline. This test provides insight not only into performance under idealized inputs but also into the stability of the predictions under input perturbations, as expected in experimental settings where measurement noise and input uncertainty are typically present.

D. Predictions with coarse inputs

To assess the impact of limited spatial resolution in wall measurements, we consider a scenario in which the wall-based input fields, $\mathbf{I}(x, z)$, are systematically coarsened. Starting from fully resolved wall-shear stress and wall-pressure fields discretized on a grid (N_x, N_z), we construct coarsened inputs $\mathbf{I}^{(F)}$ by subsampling over nonoverlapping $F \times F$ blocks, leading to an effective input resolution $(N_x/F, N_z/F)$. For up-sampling of the coarsened input fields using FCN as depicted in Fig. 2, a single bilinear 2D up-sampling layer of size (F, F) is introduced before the first convolutional layer when $F \neq 1$. In the present work, we examine coarsening factors $F = 2, 4, 8$, which we refer to as cases F2, F4, and F8, respectively. For each of these cases, the network is provided with the coarse wall fields $\mathbf{I}^{(F)}$ while the target off-wall velocity fluctuations at $y^+ = 50$ are retained at the original DNS resolution. For $\mathrm{Re}_\tau = 180$, the coarse fields at $F = 8$ corresponds to predicting the DNS resolution $(N_x, N_y) = (192, 192)$ velocity fluctuation fields at $y^+ = 50$ using input fields that are of resolution $(N_x, N_y) = (24, 24)$, and for $\mathrm{Re}_\tau = 550$ the predicted fields are of shape $(N_x, N_z) = (512, 512)$ with input fields corresponding to $(N_x, N_z) = (64, 64)$ at $F = 8$. The resulting predictions are evaluated using the relative error in prediction of RMS quantities and energy distribution across scales with respect to the DNS reference.

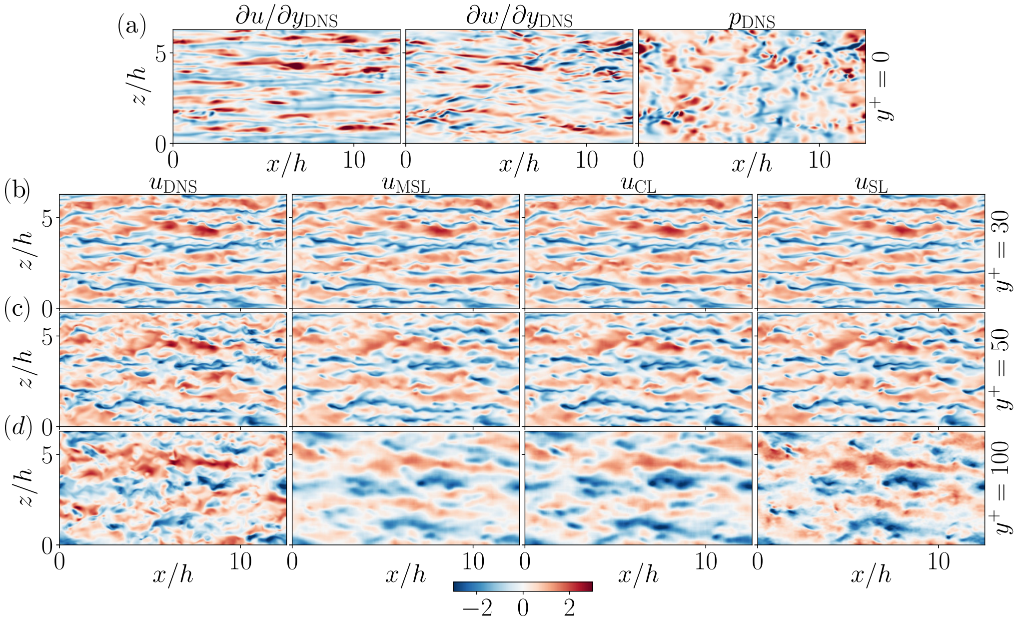

FIG. 3. Sample instantaneous [(a) left to right] RMS-normalized streamwise wall shear, spanwise wall shear, and wall pressure fluctuations at $\mathrm{Re}_\tau = 180$, which are provided as inputs to the FCN. Corresponding instantaneous streamwise velocity-fluctuation fields at (b) $y^+ = 30$, (c) $y^+ = 50$, (d) $y^+ = 100$ (left to right) reference DNS fields, predictions from FCN using MSE (MSL), CL, and SL. The fields are scaled with the corresponding RMS values.

III. RESULTS AND DISCUSSION

A. Comparison of predictions with different loss functions

The predictions from the trained neural networks employing different loss functions are compared against reference DNS fields. First, we assess the performance qualitatively using instantaneous fields, then quantitatively using turbulence statistics and two-dimensional power spectral densities. Moreover, the impact of different loss functions on the reconstruction quality is emphasized.

1. Qualitative analysis

We examine instantaneous velocity-fluctuation predictions at multiple wall-normal targets, using wall-shear rates (streamwise and spanwise) and wall pressure as inputs. In Fig. 3, we show representative snapshots of the streamwise fluctuation field at $\mathrm{Re}_\tau = 180$ from DNS and FCN models trained with MSE loss (denoted MSL), composite loss (CL), and spectral-augmented loss (SL), and the wall-normal and spanwise fields are presented in Appendix C. All predicted fields exhibit strong structural agreement with the DNS ground truth, though the match degrades progressively with wall-normal distance (also due to the decreasing correlation between wall inputs and off-wall velocities with increasing wall-normal position). The MSL predictions capture large-scale features well but tend to be smooth and lack sharpness. The CL formulation, which enforces amplitude and correlation constraints, improves the correspondence in fluctuation magnitudes, especially for the predictions farther from the wall. Most strikingly, the predictions from FCN employing spectral loss recovers more small-scale detail that closely aligns with the DNS structures. This trend is consistent

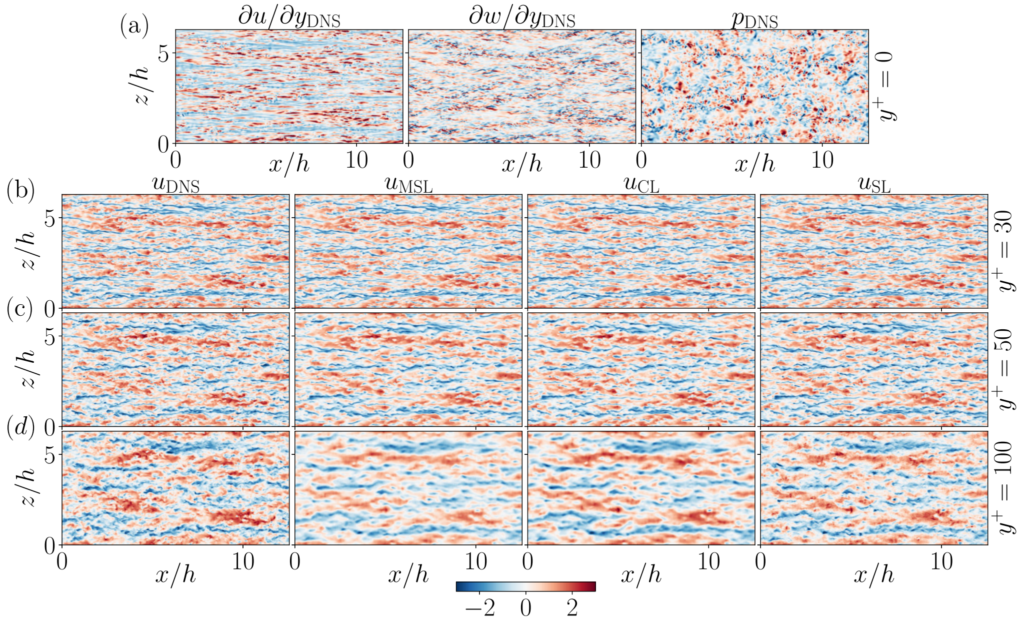

FIG. 4. Sample instantaneous flow fields at $\mathrm{Re}_\tau = 550$. (a) Inputs to the FCN (left to right): RMS-normalized fluctuations in streamwise wall shear, spanwise wall shear, and wall pressure. (b)–(d) Corresponding streamwise velocity-fluctuation fields at $y^+ = 30, 50, 100$, respectively, and each row compares (left to right) the reference DNS data, with predictions from the FCN using MSE (MSL), CL, and SL loss functions. The fields are scaled with the corresponding RMS values.

in the $\mathrm{Re}_\tau = 550$ case as well (see Fig. 4), demonstrating that the advantages of spectral-informed loss generalize across Reynolds numbers.

2. Comparison of MSE

Table II shows the RMS-normalized mean-squared error (where the mean-squared error values are divided by the respective square of RMS values to render it as a nondimensional quantity) for each velocity component (u, v, w) at target wall-normal positions $y^+ = 30, 50, 100$ for $\mathrm{Re}_\tau = 180$. The results compare predictions using three different loss functions: the baseline FCN (with MSE loss function) denoted as MSL, the CL, and the spectra-augmented loss (SL).

For the near-wall region ($y^+ = 30$), the RMS-normalized MSE is roughly halved when using CL or SL relative to MSL for the prediction of streamwise velocity fluctuations. The prediction of wall-normal and spanwise velocity fluctuations at $y^+ = 30$ also show a decreased RMS-normalized MSE with CL or SL in comparison to MSL. This demonstrates the significant benefit of the additional loss terms in capturing fine near-wall fluctuations using wall inputs. At larger wall-normal positions, the errors across the three loss functions are nearly comparable, suggesting that for large wall-normal separation between input and output fields, enforcing amplitude and structural constraints or spectral penalties does not consistently reduce MSE. These observations support the idea that the augmented loss formulations (CL, SL) bring the most benefit close to the wall in terms of reducing the discrepancy in the turbulent kinetic energy in the predicted fields, where the input–output relationship is stronger and fine-scale structures are relatively dominant compared to larger wall-normal distances.

For $\mathrm{Re}_\tau = 550$, Table III presents the RMS-normalized MSE values for MSL, CL, and SL predictions across several wall-normal positions. In this regime, we observe only moderate improvements

TABLE II. RMS-normalized MSE for predictions at $\mathrm{Re}_\tau = 180$ under different loss functions at different wall-normal positions. The average obtained with three different FCN models (with different random initializations of network weights during training) is reported along with the model-to-model variations which quantifies the stochastic variability of the optimization process during training. While the model-to-model variations (obtained as a standard-deviation of results from three different FCN models) provide an indication of training sensitivity, the metric may not be robust given the small number of trials; a more extensive statistical analysis becomes necessary, which is not considered within the scope of present work.

y^+	Component	MSL	CL	SL
30		0.1111 (± 0.001)	0.050 (± 0.001)	0.040 (± 0.001)
50	u	0.1584 (± 0.007)	0.150 (± 0.008)	0.150 (± 0.001)
100		0.5094 (± 0.025)	0.515 (± 0.028)	0.518 (± 0.017)
30		0.0942 (± 0.001)	0.070 (± 0.001)	0.070 (± 0.001)
50	v	0.2458 (± 0.003)	0.240 (± 0.004)	0.240 (± 0.001)
100		0.6643 (± 0.009)	0.664 (± 0.010)	0.633 (± 0.009)
30		0.0718 (± 0.002)	0.060 (± 0.002)	0.060 (± 0.001)
50	w	0.2289 (± 0.009)	0.230 (± 0.010)	0.230 (± 0.003)
100		0.6376 (± 0.009)	0.663 (± 0.010)	0.642 (± 0.013)

(or nearly identical performance) of CL or SL relative to the baseline FCN, especially near the wall. However, at larger wall-normal distances, the error under SL occasionally increases slightly compared to MSL.

As Re_τ increases, the turbulence becomes richer in small-scale, high-wave-number features. The spectral loss pushes the network to reconstruct more of these fine details. In doing so, it may introduce minor deviations at the grid-point level in less-correlated regions, thereby slightly raising the pointwise MSE despite improving spectral fidelity. SL tends to prioritize matching energy across frequencies, which may conflict with minimizing pointwise errors causing local increase in MSE while still improving the predicted energy content and coherence of flow structures.

Thus, the slight increase in MSE at larger y^+ for SL does not necessarily imply poorer physical reconstruction. Rather, it reflects a trade-off by encouraging faithful high-frequency recovery, SL can sometimes deviate locally in pixel error in weakly constrained regions, while still maintaining or improving global spectral consistency. These observations suggest that toward minimizing pointwise errors between predictions and DNS reference, SL is most beneficial closer to the wall, where the input-output coupling is strong and fine-scale structures are prominent.

TABLE III. RMS-normalized MSE for predictions under different loss functions for $\mathrm{Re}_\tau = 550$ at different target wall-normal positions. The average obtained with three different FCN models is reported along with the model-to-model variations.

y^+	Component	MSL	CL	SL
30		0.0695 (± 0.001)	0.0733 (± 0.001)	0.0633 (± 0.002)
50	u	0.1734 (± 0.003)	0.1573 (± 0.002)	0.1601 (± 0.004)
100		0.3952 (± 0.007)	0.4035 (± 0.006)	0.4835 (± 0.005)
30		0.1236 (± 0.002)	0.1281 (± 0.002)	0.1105 (± 0.001)
50	v	0.3074 (± 0.005)	0.2774 (± 0.004)	0.2980 (± 0.008)
100		0.6421 (± 0.007)	0.6731 (± 0.006)	0.8244 (± 0.009)
30		0.0898 (± 0.002)	0.0932 (± 0.001)	0.0815 (± 0.001)
50	w	0.2416 (± 0.005)	0.2190 (± 0.003)	0.2318 (± 0.006)
100		0.6179 (± 0.011)	0.6484 (± 0.012)	0.7256 (± 0.011)

TABLE IV. Relative percentage error in predicting RMS quantities: A comparison between results obtained by Guastoni *et al.* [15] and the present FCN with different loss functions at $\mathrm{Re}_\tau = 180$. Each entry corresponds to the average of the results obtained with three different network models and the model-to-model variation is also indicated.

y^+	Component	Guastoni *et al.* [15]	MSL	CL	SL
	u	6.79 (± 1.31)	2.52 (± 0.33)	1.35 (± 0.12)	1.11 (± 0.52)
30	v	10.18 (± 1.67)	2.97 (± 0.78)	1.43 (± 1.00)	1.45 (± 1.07)
	w	9.65 (± 1.07)	2.24 (± 0.99)	1.32 (± 0.97)	1.37 (± 0.85)
	u	21.47 (± 1.97)	8.14 (± 0.95)	2.72 (± 1.01)	2.42 (± 0.92)
50	v	26.66 (± 0.76)	14.03 (± 1.29)	5.22 (± 1.54)	5.72 (± 0.78)
	w	25.60 (± 1.21)	13.47 (± 0.88)	5.24 (± 0.46)	5.69 (± 0.92)
	u	50.82 (± 2.19)	24.30 (± 3.11)	8.36 (± 3.35)	8.81 (± 0.90)
100	v	59.05 (± 1.61)	36.09 (± 2.64)	16.27 (± 3.08)	13.36 (± 0.61)
	w	59.59 (± 1.31)	40.45 (± 3.10)	17.96 (± 3.74)	12.25 (± 0.53)

3. Comparison of errors in predicting turbulence statistics

To quantitatively assess the fidelity of FCN in retrieving turbulence statistics, we compute for each wall-normal position the relative percentage error in predicted RMS velocity fluctuations relative to DNS denoted by E_{RMS}. In assessing the quality of statistics obtained from FCN, we also include the results of Guastoni *et al.* [15] where the introduced FCN has a fewer number of hidden layers compared to that used in the present work, although the total number of trainable parameters remain similar.

The following Table IV compares the performance obtained by Guastoni *et al.* [15] with the present MSL, CL, and SL models at several y^+ positions for $\mathrm{Re}_\tau = 180$:

The FCN in the present study, trained with MSE loss already yields drastic reductions in RMS error compared to Guastoni *et al.* [15], a benefit we attribute to the increased representational capacity of our architecture, which can combine more abstract features from the input fields. For predictions close to the wall (e.g. $y^+ = 30$), the error in RMS statistics decreases by a factor of about 3. Even at $y^+ = 100$, error reductions of approximately a factor of two are attained.

When comparing the FCN model with MSL and CL variant, we observe consistently large gains in RMS accuracy across all wall-normal positions. These gains primarily arise from better amplitude matching of the velocity fluctuations, which the additional terms in CL explicitly promote. In turn, the SL model remains competitive with CL and marginally outperforms CL in many cases. In particular, for the wall-normal (v) and spanwise (w) components farther from the wall, SL often matches or exceeds the performance of CL and the pure MSE baseline in capturing RMS statistics and is characterized by lower model-to-model variations.

At higher Reynolds number ($\mathrm{Re}_\tau = 550$), as observed from Table V the benefits of our deeper architecture persist and we achieve at least a twofold reduction in relative RMS error over the baseline FCN [15]. The CL model further reduces RMS errors by roughly 50 percentage points in many cases, and the SL model also delivers consistent improvements across all target wall-normal positions and velocities. These results underscore the effectiveness of the CL and SL approaches in lowering E_{RMS} and producing more accurate statistical reconstructions.

Overall, the deeper FCN architecture delivers significant improvements in prediction of RMS statistics relative to FCN introduced by Guastoni *et al.* [15] and enhanced loss functions (CL,SL) yield good improvements over MSL for different target wall-normal positions. The CL already yields a strong improvement by incorporating amplitude and correlation constraints, lowering errors by factors of 2–4 in many cases. The SL model offers further gains, especially at large target y^+, by better preserving small-scale energy and structural alignment. This is particularly evident at $\mathrm{Re}_\tau = 550$, where SL tends to outperform CL more markedly.

TABLE V. Relative percentage error in predicting RMS quantities: A comparison between results obtained by Guastoni *et al.* [15] and the present FCN with different loss functions at $\mathrm{Re}_\tau = 550$. Each entry corresponds to the average of the results obtained with three different network models and the model-to-model variation is also indicated.

y^+	Component	L. Guastoni *et al.* (2021) [15]	MSL	CL	SL
	u	8.10 ($\pm$ 0.62)	2.48 ($\pm$ 0.56)	2.36 ($\pm$ 0.41)	1.61 ($\pm$ 1.16)
30	v	11.21 ($\pm$ 1.41)	6.64 ($\pm$ 0.23)	2.32 ($\pm$ 0.11)	2.96 ($\pm$ 0.92)
	w	9.03 ($\pm$ 0.31)	5.44 ($\pm$ 0.58)	2.41 ($\pm$ 0.46)	2.56 ($\pm$ 0.82)
	u	15.33 ($\pm$ 0.22)	7.31 ($\pm$ 0.32)	5.82 ($\pm$ 0.33)	3.71 ($\pm$ 0.26)
50	v	24.20 ($\pm$ 1.38)	15.97 ($\pm$ 1.75)	10.78 ($\pm$ 1.58)	4.81 ($\pm$ 0.62)
	w	21.21 ($\pm$ 1.27)	13.22 ($\pm$ 0.97)	8.21 ($\pm$ 1.31)	4.52 ($\pm$ 0.64)
	u	33.06 ($\pm$ 0.30)	19.11 ($\pm$ 1.78)	11.36 ($\pm$ 1.47)	8.05 ($\pm$ 0.83)
100	v	50.82 ($\pm$ 0.26)	38.28 ($\pm$ 2.46)	19.85 ($\pm$ 1.89)	14.55 ($\pm$ 1.01)
	w	51.83 ($\pm$ 0.38)	29.03 ($\pm$ 2.71)	17.32 ($\pm$ 1.61)	12.16 ($\pm$ 1.02)

These comparisons underscore that the dual role of architectural capacity and loss design with deeper networks provide the representational space needed to capture complex multiscale mappings, while advanced loss formulations help steer the training toward physically meaningful reconstructions rather than pure pixel-level minimization.

4. Comparison of energy distribution across scales

The distribution of energy in the predicted and DNS data across different scales are compared through the premultiplied two-dimensional power spectral density plots as illustrated in Fig. 5.

From Fig. 5, it is evident that the model using spectra-augmented loss (SL) more faithfully recovers the small-scale energy content than the models trained with pure MSE or CLs. Between MSL and CL, there is little difference in their ability to resolve fine scales and the improvements mainly reflect better matching of fluctuation amplitudes and featurewise correlation rather than enhanced high-frequency recovery. In contrast, because the spectral loss explicitly enforces agreement across all frequency modes as present in reference DNS fields, the FCN model employing SL recovers finer structures with more accurate amplitudes.

It is worth emphasizing that the Fourier coefficients aggregate information globally over the field, so a spectral loss term complements local, pixelwise error measures. The spectral penalty term helps reduce overfitting to local noise or spurious deviations while tolerating small spatial shifts or phase misalignment. In experimental settings where noise or measurement errors may exist, this robustness to slight misalignments would be of particular benefit. This is further investigated in Sec. III B.

Across wall-normal positions, we observe that with increasing y^+ (see $y^+ = 50$ in Fig. 5), the contours (at 10%, 50%, and 90% of maximum spectral density) from the SL predictions align more closely with the DNS reference than those from FCN or CL. At larger wall-normal distances, some discrepancy remains in the large energy containing features which are not reproduced as faithfully in any of the trials. However, SL still outperforms FCN and CL in reproducing those large scales better. Importantly, these improvements are not restricted to the streamwise component u but carry over to v and w fluctuations as well.

A similar trend is observed at $\mathrm{Re}_\tau = 550$ (see Fig. 6). The spectral recovery of fine-scale features by the SL-trained model is markedly superior, and the improvements persist even at larger wall-normal positions.

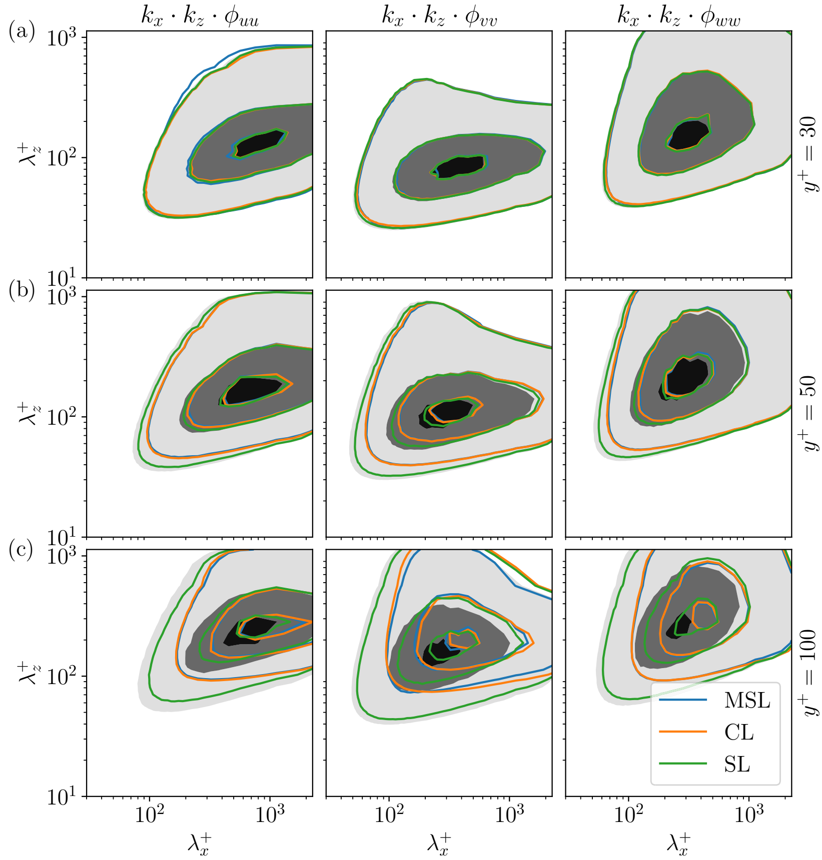

FIG. 5. Premultiplied two-dimensional power-spectral densities of (left column) the streamwise, (middle column) wall-normal, and (right column) spanwise velocity components at target wall-normal position of (a) $y^+ = 30$, (b) $y^+ = 50$, and (c) $y^+ = 100$ for $\mathrm{Re}_\tau = 180$. The contour levels contain 10%, 50%, and 90% of the maximum power-spectral density. Shaded contours refer to DNS, while contour lines correspond to the spectra obtained from predictions with (blue) FCN using MSE loss function, (orange) FCN with CL, and (green) FCN with SL.

B. Predictions with noisy inputs

In the present work, we performed predictions in which network inputs **I** (streamwise and spanwise wall-shear rates and wall pressure) were introduced with additive Gaussian noise of amplitude of 5%, 25%, and 100% of respective DNS RMS values. The network models are retrained to predict the target velocity fluctuation fields with noisy inputs. This serves to test FCN model's robustness to input uncertainty, which is important when considering experimental applications. We perform the analysis with predictions obtained at $y^+ = 50$ for $\mathrm{Re}_\tau = 180, 550$. The obtained relative error in predicting RMS values in the test dataset is provided in Table VI. A sample instantaneous qualitative field highlighting the effects of noise in the inputs and the corresponding predictions of streamwise velocity fluctuation fields at $y^+ = 50$ are plotted in Fig. 7. The distribution of energy in the predicted fields using noisy inputs and reference DNS data across different scales are compared through the spectral analysis as illustrated in Fig. 8.

The results show that when introducing noise to the input fields, the SL-trained network model maintains strong performance in capturing fine-scale structures in contrast to the network model trained with MSE as loss functions. Comparing the results obtained with DNS inputs (refer to Tables

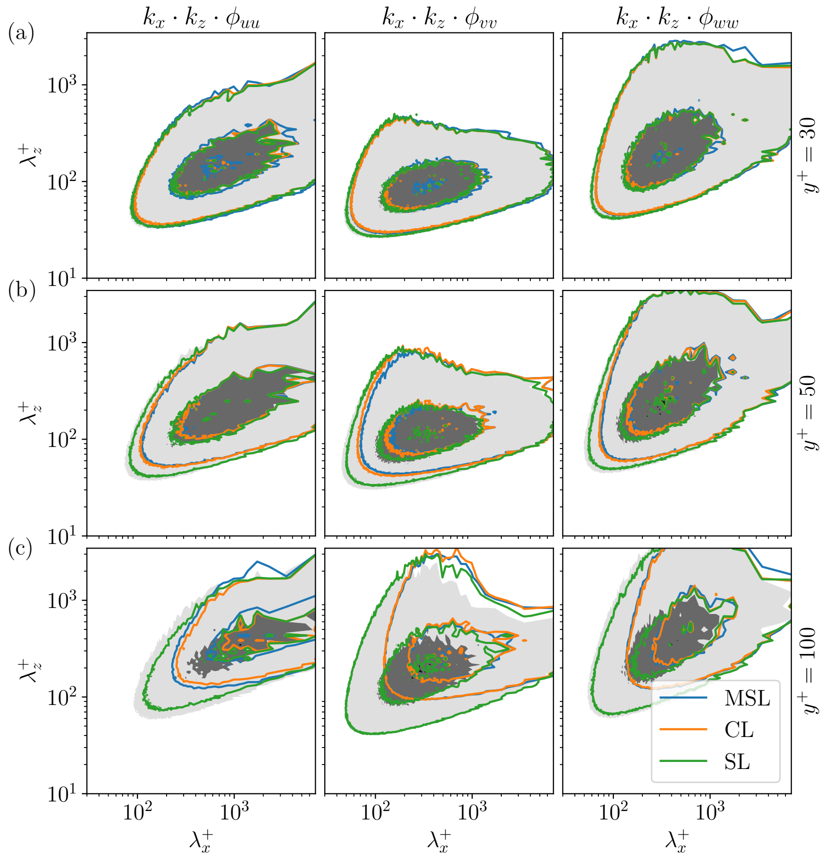

FIG. 6. Premultiplied two-dimensional power-spectral densities of (left column) the streamwise, (middle column) wall-normal, and (right column) spanwise velocity components at target wall-normal position of (a) $y^+ = 30$, (b) $y^+ = 50$, and (c) $y^+ = 100$ for $\mathrm{Re}_\tau = 550$. The contour levels contain 10%, 50%, and 90% of the maximum power-spectral density. Shaded contours refer to DNS, while contour lines correspond to the spectra obtained from predictions with (blue) MSE loss function, (orange) CL, and (green) SL.

IV and V), we observe the spanwise and wall-normal velocity fluctuations to be predicted with near similar accuracy. However, the most significant degradation occurs in the streamwise component, where the MSE-trained network error nearly doubles with noise while the SL-trained error also increases with a near similar amount up to 4% for prediction at $y^+ = 50$. The qualitative fields as plotted in Fig. 7 also highlight that FCN models trained with MSE or SL reproduce the large scale structures at $y^+ = 50$, with SL capturing the fine-scale structures better than the MSL counterpart.

Further, from the 2D power spectral density plot (refer to Fig. 8), we observe the FCN model trained with MSE to lack certain small scales in streamwise and wall-normal fluctuations $\mathrm{Re}_\tau = 180$ when the input fields are introduced with noise. However, for 5% noise in the input fields, we do not observe drastic differences in the energy content in different scales in the predictions.

These findings suggest two key aspects where the spectral-loss model exhibits superior generalization and robustness: Even when wall-based inputs are introduced with noise, it maintains low RMS errors, indicating that the FCN trained with SL as the loss function has learned more stable and noise-tolerant feature representations. Second, since experimental wall-fields will invariably

TABLE VI. Comparison of relative errors in prediction of RMS quantities (E_{RMS} [%]) for velocity components (u, v, w) at $y^+ = 50$, employing FCNs trained with mean-squared loss (MSL), composite loss (CL), and spectral loss (SL) using noisy wall-input fields. Reported values are averaged over three network models initialized with different random weights and the respective model-to-model variations are also indicated.

Re_τ	Model type	$\sigma_{\mathrm{noise}}/\mathbf{I}_{\mathrm{RMS}}$	u	v	w
180	MSL	0.05	12.36 (± 0.97)	19.13 (± 1.12)	17.63 (± 0.92)
180	CL	0.05	7.18 (± 1.12)	13.41 (± 0.92)	11.94 (± 1.04)
180	SL	0.05	5.48 (± 0.59)	9.05 (± 0.65)	8.40 (± 0.87)
180	MSL	0.25	17.15 (± 1.02)	25.69 (± 1.34)	24.24 (± 1.26)
180	CL	0.25	11.62 (± 0.95)	12.71 (± 1.11)	16.37 (± 1.13)
180	SL	0.25	7.0 (± 0.97)	9.42 (± 0.91)	10.51 (± 0.94)
180	MSL	1.0	22.90 (± 1.12)	33.79 (± 1.62)	32.47 (± 1.33)
180	CL	1.0	14.87 (± 1.26)	21.64 (± 1.31)	22.44 (± 1.27)
180	SL	1.0	9.99 (± 1.25)	13.29 (± 1.84)	13.16 (± 1.78)
550	MSL	0.05	11.03 (± 0.49)	20.68 (± 1.29)	15.40 (± 0.83)
550	CL	0.05	8.92 (± 0.48)	14.73 (± 1.46)	12.42 (± 1.31)
550	SL	0.05	5.75 (± 0.62)	7.90 (± 0.72)	7.31 (± 0.74)
550	MSL	0.25	12.19 (± 0.62)	23.64 (± 1.89)	17.52 (± 1.26)
550	CL	0.25	10.76 (± 0.78)	16.55 (± 1.25)	15.27 (± 1.31)
550	SL	0.25	6.80 (± 0.81)	8.95 (± 0.94)	7.87 (± 0.96)
550	MSL	1.0	21.03 (± 1.04)	34.15 (± 1.52)	28.83 (± 1.48)
550	CL	1.0	13.62 (± 1.19)	21.97 (± 1.98)	20.35 (± 2.13)
550	SL	1.0	7.48 (± 1.21)	12.67 (± 1.64)	11.50 (± 1.47)

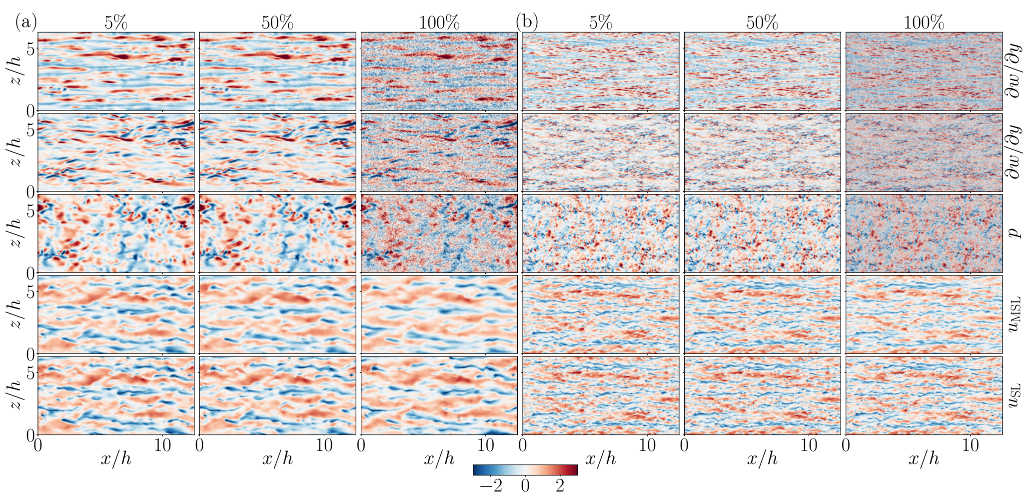

FIG. 7. Instantaneous fields at (a) $\mathrm{Re}_\tau = 180$ and (b) $\mathrm{Re}_\tau = 550$. The first, second, and third columns in (a) and (b) correspond to the effect with the amplitude of noise level corresponding to 5%, 25%, and 100% of the respective RMS values of the input fields. The first three rows, respectively, correspond to the input fields of streamwise wall shear, spanwise wall shear, and wall pressure fields. The fourth and fifth rows, respectively, correspond to the predicted streamwise velocity fluctuation field at $y^+ = 50$ from FCN trained with MSE and SL. Refer to Figs. 3 and 4 for the reference DNS velocity fluctuation fields at $y^+ = 50$. The fields are normalized by the respective RMS values.

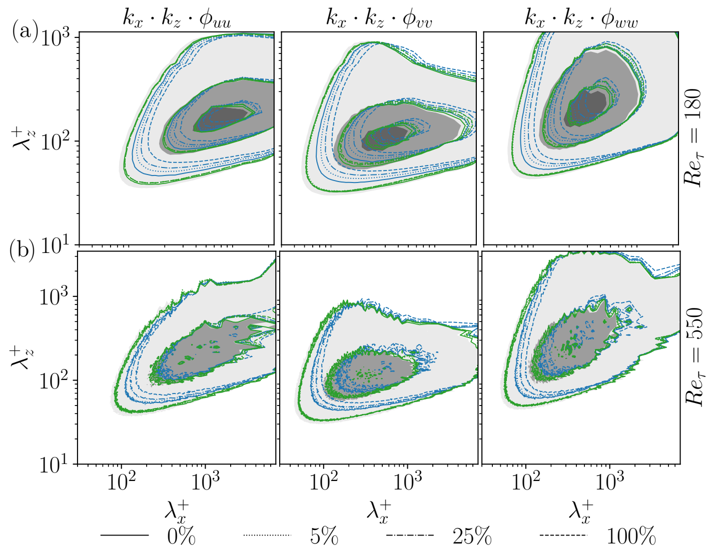

FIG. 8. Premultiplied two-dimensional power-spectral densities of (left column) the streamwise, (middle column) wall-normal, and (right column) spanwise velocity components at target wall-normal position of $y^+ = 50$ for (a) $\mathrm{Re}_\tau = 180$ and (b) $\mathrm{Re}_\tau = 550$. The contour levels contain 10%, 50%, and 90% of the maximum power-spectral density. Shaded contours refer to DNS, while contour lines correspond to the spectra obtained from predictions with (blue) FCN using MSE loss function and (green) FCN with SL. The contour line styles indicate the amplitude of noise level with respect to the input RMS.

include possible measurement noise due to sensor error, calibration drift, signals from other physical processes, using an FCN trained with spectral loss highlights its suitability for practical deployment.

C. Prediction with coarse inputs

In addition to the test with noisy inputs, we examine the sensitivity of the reconstruction to the spatial resolution of the wall-based measurements. To mimic the resolution limitations typically encountered in experimental campaigns, the input fields $\mathbf{I}^{(F)}$ are coarsened with factors $F = 2, 4, 8$ (denoted as F2, F4, and F8, respectively). The network models are retrained to predict the DNS resolution target velocity fluctuation fields at $y^+ = 50$ using coarse inputs for $\mathrm{Re}_\tau = 180, 550$. The obtained relative error in predicting RMS values in the test dataset is provided in Table VII and an instantaneous qualitative prediction with $F = 8$ is provided in Fig. 9. The distribution of energy in the predicted fields at $y^+ = 50$ using noisy inputs and reference DNS data across different scales are compared through the spectral analysis as illustrated in Fig. 10.

From the qualitative comparisons in Fig. 9, the predictions obtained with MSL using coarse input fields at $F = 8$ appear diffused and fail to capture well the amplitude of turbulent fluctuations, whereas SL-based predictions exhibit sharper structures and visibly closer agreement with the reference DNS fields. This improvement is quantitatively confirmed by the results in Table VII, where the relative percentage errors in key turbulence statistics obtained with SL are at least a factor of 2 smaller than those obtained with MSL across the range of tested coarsening factors. For SL, the errors almost increase linearly with coarsening wall inputs ($F > 1$) for $\mathrm{Re}_\tau = 180$, and for $\mathrm{Re}_\tau = 550$ it remains nearly linear with respect to F up to $F = 4$ before tending to saturate at around $F = 8$ with about 15% relative difference in RMS statistics with respect to DNS. This suggests that models trained with SL are able to retain much of its predictive capability even when only low-resolution wall inputs are provided, a scenario relevant for practical experimental configurations. The two-dimensional premultiplied spectra further indicates that SL preserves the small-scale

TABLE VII. Relative error (E_{RMS} [%]) in RMS velocity components (u, v, w) at $y^+ = 50$ for FCNs trained with mean-squared loss (MSL) and spectral loss (SL), using wall-input fields coarsened by factors $F = 2, 4, 8$. Results for predictions with DNS inputs ($F = 1$) are also reported. Errors correspond to deviations in RMS quantities, averaged over three independently initialized network models and the model-to-model variations are also indicated.

Re_τ	F	u_{MSL}	v_{MSL}	w_{MSL}	u_{SL}	v_{SL}	w_{SL}
180	1	8.14 (± 0.95)	14.03 (± 1.29)	13.47 (± 0.88)	2.42 (± 0.92)	5.72 (± 0.78)	5.69 (± 0.92)
180	2	11.27 (± 1.17)	17.07 (± 1.27)	17.09 (± 1.04)	4.16 (± 1.99)	6.42 (± 1.11)	6.68 (± 1.25)
180	4	18.48 (± 1.16)	27.25 (± 1.37)	23.86 (± 1.47)	6.85 (± 1.12)	10.01 (± 1.72)	9.96 (± 1.70)
180	8	34.10 (± 1.67)	46.02 (± 1.95)	48.14 (± 1.78)	11.39 (± 1.64)	15.74 (± 1.86)	16.68 (± 1.89)
550	1	7.31 (± 0.32)	15.97 (± 1.75)	13.22 (± 0.97)	3.71 (± 0.26)	4.81 (± 0.62)	4.52 (± 0.64)
550	2	10.77 (± 1.09)	18.01 (± 1.23)	15.01 (± 1.16)	5.73 (± 1.57)	8.17 (± 1.26)	6.69 (± 1.59)
550	4	18.66 (± 1.71)	26.15 (± 1.91)	22.35 (± 1.48)	7.96 (± 1.79)	10.46 (± 1.91)	9.68 (± 1.83)
550	8	32.02 (± 1.97)	52.51 (± 2.79)	20.75 (± 2.04)	12.17 (± 1.89)	15.88 (± 2.11)	14.75 (± 1.89)

energy distribution at $y^+ = 50$ for both $\mathrm{Re}_\tau = 180$ and 550, whereas MSL underestimates the energy content at high wavenumbers and consequently results in fields that are less consistent with DNS, as also observed from Fig. 9. These findings highlight the advantages of using SL with coarse-wall inputs demonstrating the enhanced potential of SL-based training for nonintrusive sensing in wall-bounded turbulent flows. While the present implementation is demonstrated with an FCN, the same SL framework can be extended to more advanced architectures, such as GANs or diffusion models, which may further improve the recovery of small-scale turbulence.

A practical limitation of the considered spectral loss is its dependence on well-converged statistics and spectral information of the target fields, which are readily available from DNS but may be difficult to estimate in experimental settings. This constraint is particularly restrictive when only a limited number of data realizations can be acquired or when the measurement domain is small. A possibility to extend the present methodology in an applied setting is to calibrate the spectral components of the loss on high-fidelity numerical data and subsequently adapt the models to experimental dataset via transfer learning and setting some loss terms in SL to be zero [e.g., $\delta = \zeta = 0$ in Eq. (6)], depending on the availability of respective data. Alternatively, predefined parametric or reduced-order models for the spectra can be considered as a function of Re_τ and y^+

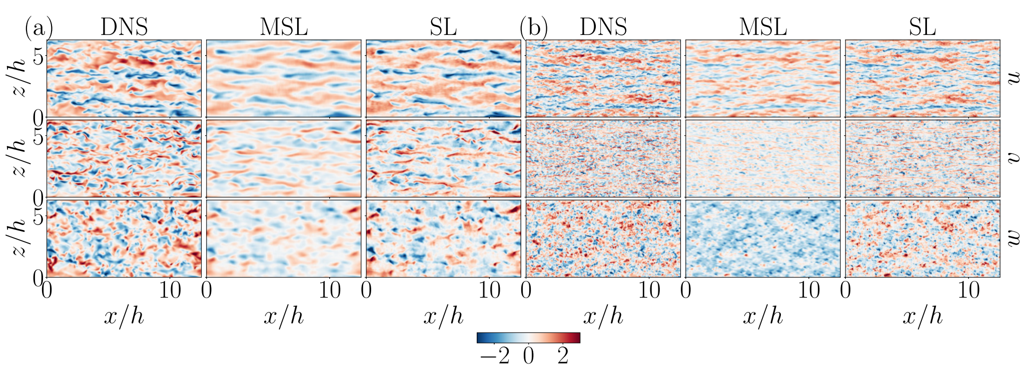

FIG. 9. Sample instantaneous (top to bottom) streamwise, wall-normal, and spanwise velocity fluctuation fields at $y^+ = 50$ for (a) $\mathrm{Re}_\tau = 180$ and (b) $\mathrm{Re}_\tau = 550$. In (a) and (b), the target reference DNS fields and the predictions employing MSL and SL, where inputs are coarsened with $F = 8$, are, respectively, shown from left to right. The fields are scaled with the corresponding RMS values.

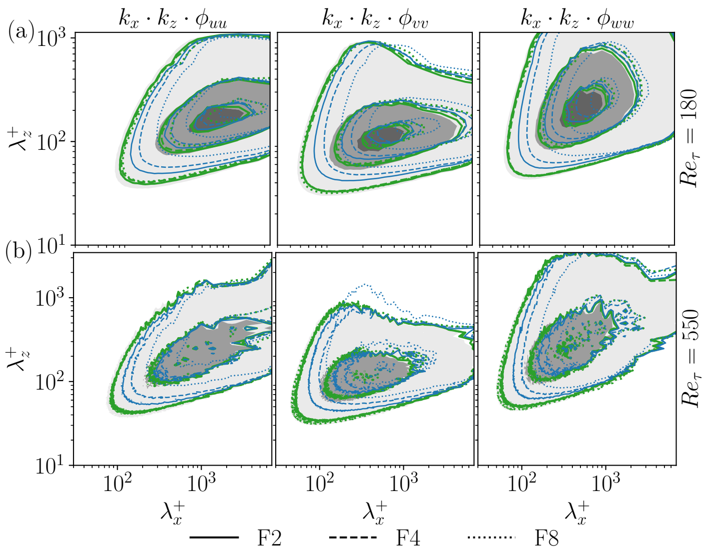

FIG. 10. Premultiplied two-dimensional power-spectral densities of (left column) the streamwise, (middle column) wall-normal, and (right column) spanwise velocity components at target wall-normal position of $y^+ = 50$ at (a) $\mathrm{Re}_\tau = 180$ and (b) $\mathrm{Re}_\tau = 550$. The contour levels contain 10%, 50%, and 90% of the maximum power-spectral density. Shaded contours refer to DNS, while contour lines correspond to the spectra obtained from predictions with coarse inputs at (solid) $F = 2$, (dashed) $F = 4$, and (dotted) $F = 8$. The blue contour lines correspond to the predictions employing MSE as loss function and green contour lines denote the predictions with SL.

for training the network models, and comparisons can be performed in a low-dimensional set of spectral metrics that could be extracted from an experimental dataset.

IV. SUMMARY AND CONCLUSIONS

This study demonstrates that the spectra-informed loss function substantially enhances the accuracy of nonintrusive sensing for wall-bounded turbulence. By optimizing not only the mean-squared error but also the statistical fidelity of the fluctuations, correlations, and spectral energy distribution, the considered approach achieves superior reconstruction of turbulent structures across a range of wall-normal positions and Reynolds numbers ($\mathrm{Re}_\tau = 180$ and 550). Relative to a conventional MSE baseline, the FCN model recovers significantly more small-scale content, underscoring that the loss function formulation is a critical determinant of prediction fidelity. From an application standpoint, the framework exhibits notable robustness, maintaining low prediction errors and increased retrieval of fine-scale structures when the network models are trained with noisy inputs and continues to outperform MSE-based models when predicting from low-resolution wall input fields, thereby demonstrating its suitability for realistic experimental scenarios characterized by measurement noise and limited sensor resolution. At the same time, the considered loss function requires access to sufficiently converged statistical information, including spectra, which may be challenging to obtain from experimental datasets. This limitation can be overcome by adopting hybrid training strategies in which the spectral components of the loss are calibrated on a high-fidelity numerical dataset (DNS) and subsequently adapted to experiments through transfer learning (where contribution from some loss terms can be set to zero depending on data availability) or employing reduced-order spectral models (when low-dimensional spectral metrics can be estimated from experimental data).

In conclusion, this work illustrates the enhanced potential of deep learning for instantaneous flow-field estimation in wall turbulence, with direct utility for flow control, state estimation, and

TABLE VIII. RMS-normalized mean-squared error (MSE) and relative percentage error in RMS statistics (E_{RMS}) for the velocity components (u, v, w) at $\mathrm{Re}_\tau = 550$ and target wall-normal location of $y^+ = 30, 100$ obtained with the training dataset.

Re_τ	y^+	Type	MSE(u)	MSE(v)	MSE(w)	$E_{\mathrm{RMS}}(u)$ [%]	$E_{\mathrm{RMS}}(v)$ [%]	$E_{\mathrm{RMS}}(w)$ [%]
550	30	MSL	0.0674 (± 0.006)	0.1204 (± 0.003)	0.0875 (± 0.001)	3.44 (± 0.81)	6.55 (± 0.93)	5.68 (± 0.83)
550	30	CL	0.0714 (± 0.007)	0.1244 (± 0.002)	0.0906 (± 0.001)	2.58 (± 0.50)	2.31 (± 0.13)	2.64 (± 0.56)
550	30	SL	0.0609 (± 0.002)	0.1062 (± 0.003)	0.0784 (± 0.001)	1.84 (± 1.45)	3.01 (± 1.10)	2.88 (± 1.04)
550	100	MSL	0.3856 (± 0.008)	0.6355 (± 0.007)	0.6092 (± 0.013)	19.79 (± 1.84)	38.05 (± 2.13)	28.95 (± 1.84)
550	100	CL	0.3754 (± 0.008)	0.6402 (± 0.006)	0.5960 (± 0.012)	12.29 (± 1.75)	19.42 (± 2.11)	16.10 (± 1.66)
550	100	SL	0.4604 (± 0.010)	0.8016 (± 0.013)	0.7130 (± 0.016)	8.82 (± 0.96)	14.36 (± 1.16)	11.48 (± 2.44)

TABLE IX. Variation of RMS-normalized mean-squared error and relative percentage error (E_{RMS} [%]) in RMS velocity components (u, v, w) at $\mathrm{Re}_\tau = 550$ and target wall-normal location of $y^+ = 50$ with respect to loss coefficients in SL [see Eq. (6)]. Model-averaged results over three different network initializations are provided along with the corresponding model-to-model variations.

α	β	γ	δ	ζ	MSE(u)	MSE(v)	MSE(w)	$E_{\mathrm{RMS}}(u)$ [%]	$E_{\mathrm{RMS}}(v)$ [%]	$E_{\mathrm{RMS}}(w)$ [%]
0.1	1	1	1	1	0.1629 (± 0.004)	0.2962 (± 0.006)	0.2291 (± 0.004)	3.31 (± 0.49)	4.35 (± 0.88)	4.0 (± 0.84)
1	0.1	1	1	1	0.1608 (± 0.003)	0.2923 (± 0.007)	0.2257 (± 0.002)	3.13 (± 0.59)	4.39 (± 0.78)	5.43 (± 0.71)
1	1	0.1	1	1	0.1941 (± 0.004)	1.1865 (± 0.012)	0.2676 (± 0.006)	2.47 (± 0.34)	4.91 (± 1.12)	4.53 (± 0.82)
1	1	1	0.1	1	0.1582 (± 0.004)	0.2909 (± 0.007)	0.2277 (± 0.003)	3.37 (± 0.43)	5.24 (± 0.75)	3.55 (± 0.66)
1	1	1	1	0.1	0.1717 (± 0.004)	0.3099 (± 0.008)	0.2426 (± 0.004)	5.0 (± 0.51)	6.41 (± 0.81)	5.83 (± 0.92)
0.1	0.1	0.1	1	1	0.1981 (± 0.005)	1.1505 (± 0.010)	0.2648 (± 0.005)	4.21 (± 0.49)	5.6 (± 0.93)	4.65 (± 0.83)
1	1	1	0.1	0.1	0.1572 (± 0.005)	0.2886 (± 0.007)	0.2276 (± 0.003)	2.43 (± 0.61)	5.27 (± 0.77)	2.92 (± 0.72)
0.1	0.1	0.1	0.1	1	0.1763 (± 0.004)	0.3121 (± 0.008)	0.2449 (± 0.004)	1.69 (± 0.31)	4.95 (± 0.68)	4.35 (± 0.51)
1	1	1	1	1	0.1601 (± 0.004)	0.2980 (± 0.008)	0.2318 (± 0.006)	3.71 (± 0.26)	4.81 (± 0.62)	4.52 (± 0.64)
0.5	0.5	0.5	1	1	0.1594 (± 0.003)	0.2893 (± 0.008)	0.2251 (± 0.006)	4.66 (± 0.35)	6.99 (± 0.84)	6.62 (± 0.86)
0.1	0.1	0.1	0.1	2	0.2375 (± 0.008)	1.3304 (± 0.012)	1.5394 (± 0.011)	3.07 (± 0.67)	6.12 (± 1.19)	5.36 (± 1.02)
1	0.5	1	1	1	0.1577 (± 0.004)	0.2877 (± 0.007)	0.2255 (± 0.006)	3.63 (± 0.56)	6.8 (± 0.98)	4.72 (± 0.79)

model-based design. As a next step, the present framework, trained on high-fidelity simulation data, could be transferred to experimental measurements of wall quantities, thereby assessing its performance under realistic measurement constraints. Beyond reconstruction from combined wall-shear and wall-pressure information, an important extension is to quantify how much of the off-wall dynamics can be recovered using wall pressure alone, which is often more accessible experimentally. The present methodology also extends itself naturally to superresolution tasks, where coarse or sparsely sampled wall measurements could be mapped to high-resolution volumetric fields, further enhancing its utility for practical sensing configurations. Moreover, because the considered loss-function design is largely independent of the underlying architecture, it can be straightforwardly integrated into more advanced generative models, such as GANs or diffusion-based networks [52], to further improve the recovery of small-scale turbulent structures. Additional extensions of this work includes the deployment of the framework at higher Reynolds numbers and in more complex wall-bounded configurations relevant to engineering applications. Further, studies like [47] have highlighted that lack of small-scale features in the velocity fluctuations predicted from wall inputs leads to increased errors in prediction of polymeric stresses in viscoelastic drag reducing flows and use of spectra-informed loss functions complemented by superresolution strategies can be a viable option for prediction of stress fields in complex fluids flow.

ACKNOWLEDGMENTS

This work is supported by the funding provided by the European Research Council Grant No. 2019-StG-852529, MUCUS and the Swedish Research Council through Grant No 2021-04820. The present work was performed with resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), funded by the Swedish Research Council through Grant Agreement No. 2022-06725. Some text in this article have been modified with the assistance of the large language model developed by Mistral AI, accessed via Le Chat interface (November 2025 version), where the tool was primarily used to identify suitable synonyms and to improve the readability of certain parts of the text. Part of Fig. 1 was generated with Gemini AI (December 2025 version).

DATA AVAILABILITY

The trained FCN models along with the weights will be made available on GitHub [53]. The dataset for training the network models is available upon request from the authors.

APPENDIX A: TRAINING OF NETWORK MODELS

All network models were trained for a minimum of 75 epochs (after which the stopping criteria is activated) and a maximum of 250 epochs. We also used an early-stopping patience of 20 epochs, i.e., training was terminated if the validation loss did not improve for 20 consecutive epochs. The required number of training epochs decreased with increasing target wall-normal position for each Reynolds number considered. For predictions at $y^+ = 30$, the models typically trained for $\sim$230 epochs, whereas for $y^+ = 100$ the training typically stopped at $\sim$160 epochs. The learning curves for the FCN trained with MSL and SL to predict velocity fluctuations at $y^+ = 50$ and $y^+ = 100$ for $\mathrm{Re}_\tau = 550$ are shown in Figs. 11 and 12. Overall, Fig. 11 shows that the mean-squared errors in the velocity components decreased by roughly a factor of 5 for the target fields at $y^+ = 50$ and by about a factor of 2 for the velocity fluctuations at $y^+ = 100$. In contrast, the loss reduction for the FCN trained with SL (Fig. 12) is smaller than that observed with MSE as loss function (Fig. 11), which we attribute to the added difficulty of jointly minimizing multiple error terms in SL. For the FCN models trained with MSE loss, training stopped at 164 and 155 epochs for targets at $y^+ = 50$ and $y^+ = 100$, respectively. With SL, training stopped at approximately 155

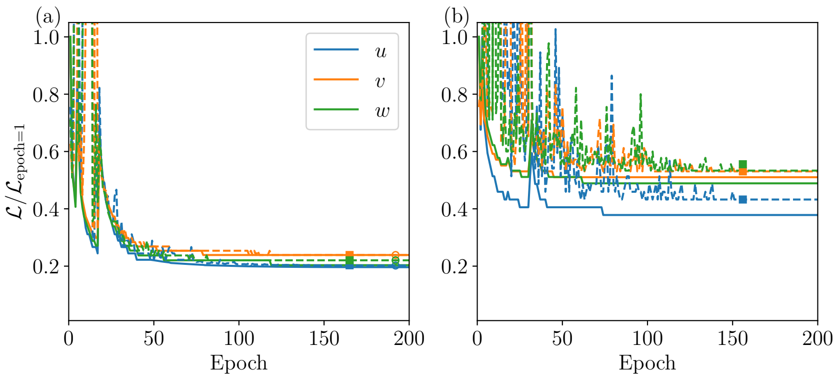

FIG. 11. Variation of loss [Eqs. (2)] for each of the velocity components with respect to epoch for target wall-normal location (a) $y^+ = 50$ and (b) $y^+ = 100$ at $\mathrm{Re}_\tau = 550$. The solid lines correspond to training loss and the dashed lines correspond to validation loss. The total training or validation loss is the sum of the respective losses for three velocity components. The loss values are normalized with the respective component values obtained at the first epoch. The training is stopped at the epoch marked with a colored square marker and the open circle marker corresponds to the epoch position at which the next improvement is obtained.

and 176 epochs for $y^+ = 50$ and $y^+ = 100$, respectively, based on the same stopping criteria. To assess whether further gains were possible beyond these stopping points, training was restarted from the saved checkpoints. The next improvement in the validation loss (summed across the three velocity components) occurred at 184 and 198 epochs for $y^+ = 50$ and $y^+ = 100$, respectively, when using SL. For MSE, the next improvement was observed at 194 and 210 epochs for $y^+ = 50$ and $y^+ = 100$, respectively. However, the relative improvement in the predicted RMS statistics between (i) the checkpoint selected by the stopping criterion and (ii) the checkpoint corresponding to the next validation-loss improvement was less than 0.1% for all velocity components at both wall-normal locations for both MSE and SL.

To assess potential overfitting across models trained with MSL, CL, and SL loss functions, we report the RMS-normalized mean-squared error and the relative percentage error in the predicted

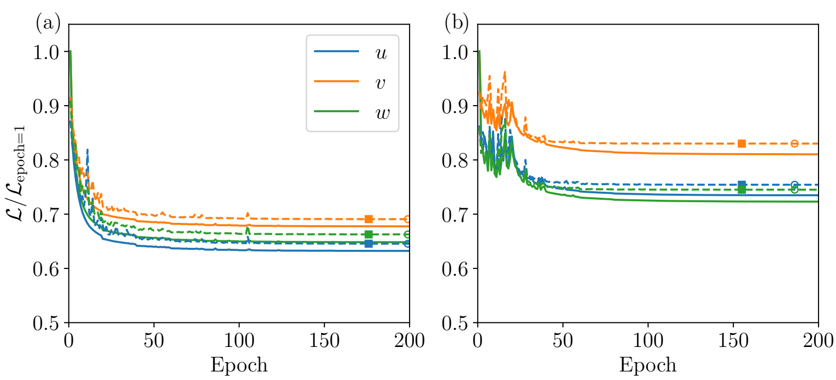

FIG. 12. Variation of loss [Eq. (6)] for each of the velocity components with respect to epoch for target wall-normal location (a) $y^+ = 50$ and (b) $y^+ = 100$ at $\mathrm{Re}_\tau = 550$. The solid lines correspond to training loss and the dashed lines correspond to validation loss. The total training or validation loss is the sum of the respective losses for three velocity components. The loss values are normalized with the respective component values obtained at the first epoch. The training is stopped at the epoch marked with a colored square marker and the open circle marker corresponds to the epoch position at which the next improvement is obtained.

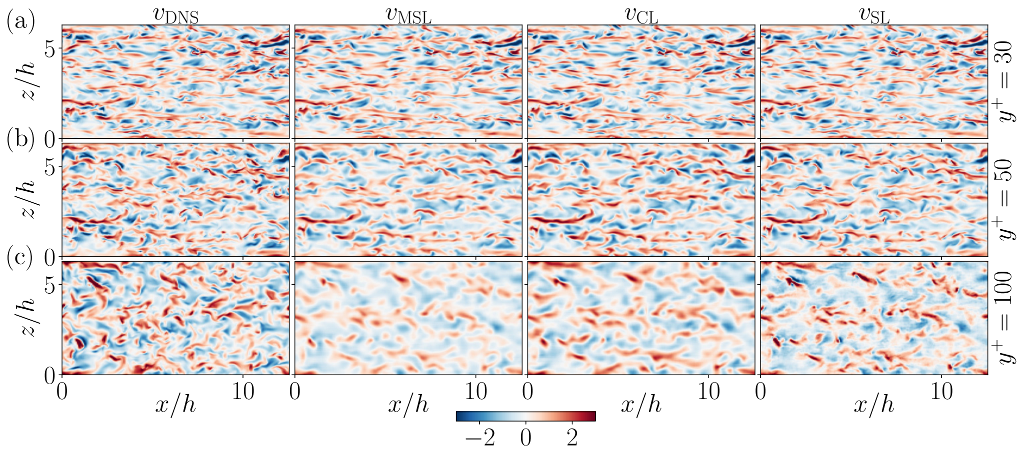

FIG. 13. Sample instantaneous wall-normal velocity fluctuation fields at (a) $y^+ = 30$, (b) $y^+ = 50$, (c) $y^+ = 100$ for $\mathrm{Re}_\tau = 180$. Each row compares (left to right) the DNS reference field with the predictions from FCN using MSE, CL, and SL. The fields are scaled with the corresponding RMS values.

RMS statistics obtained on the training dataset for $\mathrm{Re}_\tau = 550$ at $y^+ = 30$ and $y^+ = 100$ in Table VIII. These values are directly compared against the corresponding test dataset results reported in Tables III and V.

Overall, the errors obtained with the training dataset are comparable in magnitude to those obtained on the test set for both wall-normal locations. This behavior is indicative of generalization and suggests that the models did not memorize the training samples. The results also suggest that the early-stopping strategy in the employed model with different loss functions were sufficient to prevent overfitting for the considered dataset. While the present comparison is restricted to $\mathrm{Re}_\tau = 550$ and two target locations, the close similarity between training and test errors indicates that the FCN models trained with MSL, CL, and SL exhibit reasonable generalization for the range of conditions examined in this work.

APPENDIX B: TUNING LOSS COEFFICIENTS

This section summarizes the tuning trials conducted to assess the sensitivity of the FCN trained with SL to the choice of loss coefficients. Performance is evaluated using the RMS-normalized mean-squared error and the relative percentage error in the predicted RMS statistics. We performed a manual hyperparameter search over the loss weights using a small, ablation-style sweep over a discrete set of candidate combinations as provided in Table IX.

Table IX shows that, for streamwise predictions at $y^+ = 50$, the relative RMS error varies from 1.69% to 5% across the tested coefficient sets, while the unity-weight baseline yields a value of 3.71%. Importantly, comparing with Table V also indicates that the improvement obtained using SL from MSE as the loss function is larger than the sensitivity of loss coefficients. The spread in performance across Table IX suggests that additional gains may be achievable through more extensive tuning of the loss weights. To probe transferability across wall-normal locations, we additionally tested $\alpha = \beta = \gamma = 1$, $\delta = \zeta = 0.1$, alongside the unity-weight baseline, for predictions at $y^+ = 100$. For all three velocity components, the configuration $\alpha = \beta = \gamma = 1$, $\delta = \zeta = 0.1$ produced RMS errors that were at least one percentage point higher than those obtained with unity weights. This observation highlights that the best-performing coefficients can depend on both the Reynolds number and the target wall-normal position. A broader exploration of the coefficient space (e.g., larger sweeps or systematic searches) was not pursued due to the high computational cost of

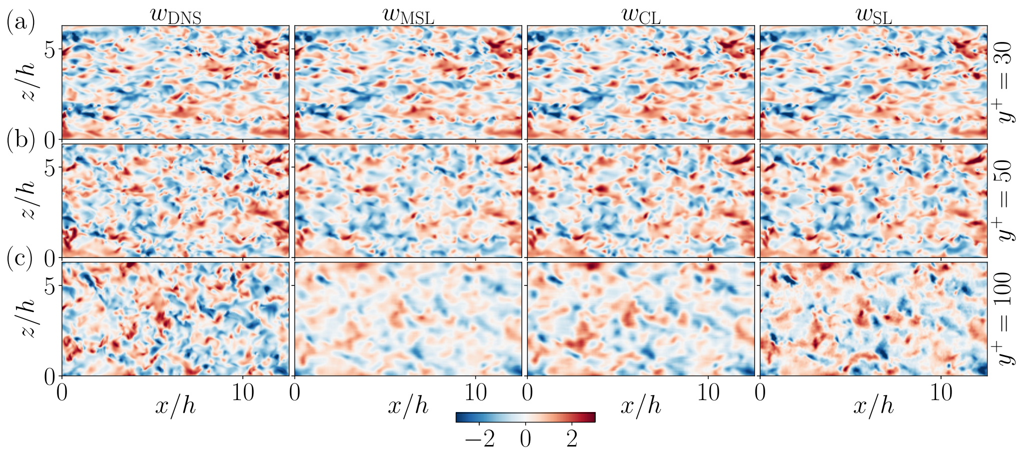

FIG. 14. Sample instantaneous spanwise velocity fluctuation fields at (a) $y^+ = 30$, (b) $y^+ = 50$, (c) $y^+ = 100$ for $\mathrm{Re}_\tau = 180$. Each row compares (left to right) the DNS reference field with the predictions from FCN using MSE, CL, and SL. The fields are scaled with the corresponding RMS values.

training each model. Accordingly, unless stated otherwise, all training and evaluations in this work use equal weights for the loss terms in both CL and SL.

APPENDIX C: QUALITATIVE FIELDS OF WALL-NORMAL AND SPANWISE VELOCITY FLUCTUATIONS

This section contains the wall-normal and spanwise velocity fluctuation fields corresponding to the same instant as the input wall-fields and the streamwise velocity fluctuations that are presented in Figs. 3 and 4. Figure 13 shows the wall-normal velocity fluctuation fields at $\mathrm{Re}_\tau = 180$ and Fig. 14 plots the spanwise velocity fluctuation fields at $\mathrm{Re}_\tau = 180$. Figure 15 depicts the wall-

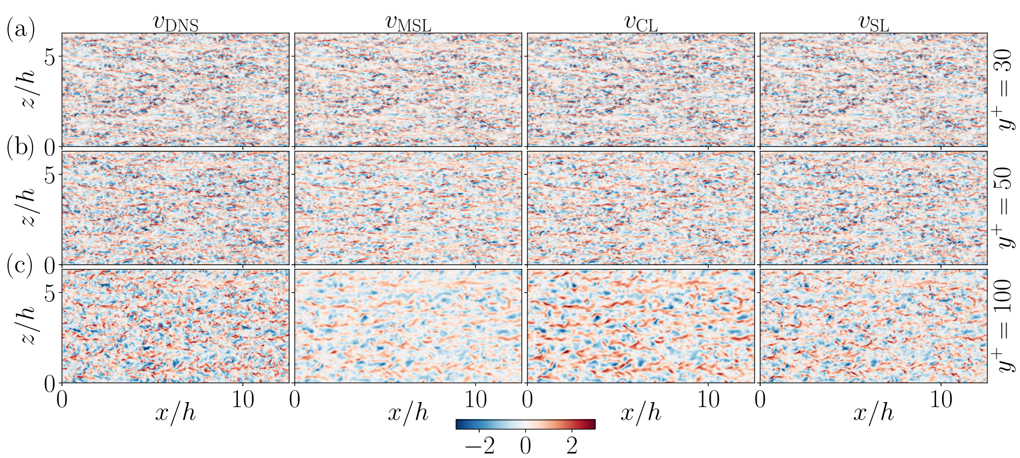

FIG. 15. Sample instantaneous wall-normal velocity fluctuation fields at (a) $y^+ = 30$, (b) $y^+ = 50$, (c) $y^+ = 100$ for $\mathrm{Re}_\tau = 550$. Each row compares (left to right) the DNS reference field with the predictions from FCN using MSE, CL, and SL. The fields are scaled with the corresponding RMS values.

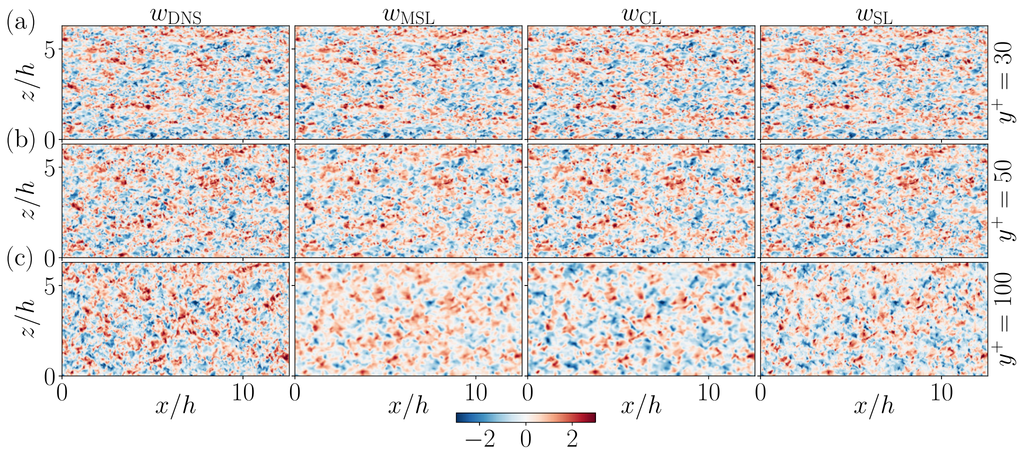

FIG. 16. Sample instantaneous spanwise velocity fluctuation fields at (a) $y^+ = 30$, (b) $y^+ = 50$, (c) $y^+ = 100$ for $\mathrm{Re}_\tau = 550$. Each row compares (left to right) the DNS reference field with the predictions from FCN using MSE, CL, and SL. The fields are scaled with the corresponding RMS values.

normal velocity fluctuations at $\mathrm{Re}_\tau = 550$ and Fig. 16 shows the spanwise velocity fluctuations at $\mathrm{Re}_\tau = 550$.

[1] K. Duraisamy, G. Iaccarino, and H. Xiao, Turbulence modeling in the age of data, Annu. Rev. Fluid Mech. **51**, 357 (2019).

[2] R. Vinuesa, Perspectives on predicting and controlling turbulent flows through deep learning, Phys. Fluids **36**, 031401 (2024).

[3] A. Cremades, S. Hoyas, and R. Vinuesa, Classically studied coherent structures only paint a partial picture of wall-bounded turbulence, Nat. Commun. **16**, 10189 (2025).

[4] J. Kutz, Deep learning in fluid dynamics, J. Fluid Mech. **814**, 1 (2017).

[5] S. Brunton, B. Noack, and P. Koumoutsakos, Machine learning for fluid mechanics, Annu. Rev. Fluid Mech. **52**, 477 (2020).

[6] R. Vinuesa and S. Brunton, Emerging trends in machine learning for computational fluid dynamics, Comput. Sci. Eng. **24**, 33 (2023).

[7] K. Duraisamy, Perspectives on machine learning-augmented Reynolds-averaged and large eddy simulation models of turbulence, Phys. Rev. Fluids **6**, 050504 (2021).

[8] J. Ling, A. Kurzawski, and J. Templeton, Reynolds averaged turbulence modelling using deep neural networks with embedded invariance, J. Fluid Mech. **807**, 155 (2016).

[9] Y. Zhao, H. Akolekar, J. Weatheritt, V. Michelassi, and R. Sandberg, RANS turbulence model development using CFD-driven machine learning, J. Comput. Phys. **411**, 109413 (2020).

[10] H. Eivazi, M. Tahani, P. Schlatter, and R. Vinuesa, Physics-informed neural networks for solving Reynolds-averaged Navier–Stokes equations, Phys. Fluids **34**, 075117 (2022).

[11] Z. Wang, K. Luo, D. Li, J. Tan, and J. Fan, Investigations of data-driven closure for subgrid-scale stress in large-eddy simulation, Phys. Fluids **30**, 125101 (2018).

[12] H. Bae and P. Koumoutsakos, Scientific multi-agent reinforcement learning for wall-models of turbulent flows, Nat. Commun. **13**, 1443 (2022).

[13] J. Sirignano and J. MacArt, Deep learning closure models for large-eddy simulation of flows around bluff bodies, J. Fluid Mech. **966**, A26 (2023).

[14] M. Kurz, P. Offenhäuser, and A. Beck, Deep reinforcement learning for turbulence modeling in large eddy simulations, Int. J. Heat Fluid Flow **99**, 109094 (2023).

[15] L. Guastoni, A. Güemes, A. Ianiro, S. Discetti, P. Schlatter, H. Azizpour, and R. Vinuesa, Convolutional-network models to predict wall-bounded turbulence from wall quantities, J. Fluid Mech. **928**, A27 (2021).

[16] A. Balasubramanian, L. Guastoni, P. Schlatter, H. Azizpour, and R. Vinuesa, Predicting the wall-shear stress and wall pressure through convolutional neural networks, Int. J. Heat Fluid Flow **103**, 109200 (2023).

[17] K. Fukami, K. Fukagata, and K. Taira, Super-resolution reconstruction of turbulent flows with machine learning, J. Fluid Mech. **870**, 106 (2019).

[18] A. Güemes, S. Discetti, A. Ianiro, B. Sirmacek, H. Azizpour, and R. Vinuesa, From coarse wall measurements to turbulent velocity fields through deep learning, Phys. Fluids **33**, 075121 (2021).

[19] K. Fukami and K. Taira, Single-snapshot machine learning for super-resolution of turbulence, J. Fluid Mech. **1001**, A32 (2024).

[20] A. Lozano-Durán, H. Bae, and M. Encinar, Causality of energy-containing eddies in wall turbulence, J. Fluid Mech. **882**, A2 (2020).

[21] D. Floryan and M. Graham, Data-driven discovery of intrinsic dynamics, Nat. Mach. Intell. **4**, 1113 (2022).

[22] A. Linot and M. Graham, Dynamics of a data-driven low-dimensional model of turbulent minimal Couette flow, J. Fluid Mech. **973**, A42 (2023).

[23] A. Solera-Rico, C. Sanmiguel Vila, M. Gómez-López, Y. Wang, A. Almashjary, S. Dawson, and R. Vinuesa, β-variational autoencoders and transformers for reduced-order modelling of fluid flows, Nat. Commun. **15**, 1361 (2024).

[24] M. Raissi, A. Yazdani, and G. E. Karniadakis, Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations, Science **367**, 1026 (2020).

[25] R. Vinuesa and S. Brunton, Enhancing computational fluid dynamics with machine learning, Nat. Comput. Sci. **2**, 358 (2022).

[26] H. Gao, L. Sun, and J.-X. Wang, Super-resolution and denoising of fluid flow using physics-informed convolutional neural networks without high-resolution labels, Phys. Fluids **33**, 073603 (2021).

[27] H. Eivazi, Y. Wang, and R. Vinuesa, Physics-informed deep-learning applications to experimental fluid mechanics, Meas. Sci. Technol. **35**, 075303 (2024).

[28] R. Vinuesa, S. Brunton, and B. McKeon, The transformative potential of machine learning for experiments in fluid mechanics, Nat. Rev. Phys. **5**, 536 (2023).

[29] P. A. Srinivasan, L. Guastoni, H. Azizpour, P. Schlatter, and R. Vinuesa, Predictions of turbulent shear flows using deep neural networks, Phys. Rev. Fluids **4**, 054603 (2019).

[30] H. Eivazi, L. Guastoni, P. Schlatter, H. Azizpour, and R. Vinuesa, Recurrent neural networks and Koopman-based frameworks for temporal predictions in a low-order model of turbulence, Int. J. Heat Fluid Flow **90**, 108816 (2021).

[31] G. Borrelli, L. Guastoni, H. Eivazi, P. Schlatter, and R. Vinuesa, Predicting the temporal dynamics of turbulent channels through deep learning, Int. J. Heat Fluid Flow **96**, 109010 (2022).

[32] J. Rabault, M. Kuchta, A. Jensen, U. Réglade, and N. Cerardi, Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control, J. Fluid Mech. **865**, 281 (2019).

[33] J. Park and H. Choi, Machine-learning-based feedback control for drag reduction in a turbulent channel flow, J. Fluid Mech. **904**, A24 (2020).

[34] C. Vignon, J. Rabault, and R. Vinuesa, Recent advances in applying deep reinforcement learning for flow control: Perspectives and future directions, Phys. Fluids **35**, 031301 (2023).

[35] L. Guastoni, J. Rabault, P. Schlatter, H. Azizpour, and R. Vinuesa, Deep reinforcement learning for turbulent drag reduction in channel flows, EPJ E **46**, 27 (2023).

[36] A. Cremades, S. Hoyas, and R. Vinuesa, Additive-feature-attribution methods: A review on explainable artificial intelligence for fluid dynamics and heat transfer, Int. J. Heat Fluid Flow **112**, 109662 (2025).

[37] A. Cremades, S. Hoyas, R. Deshpande, P. Quintero, M. Lellep, W. Lee, J. Monty, N. Hutchins, M. Linkmann, I. Marusic, and R. Vinuesa, Identifying regions of importance in wall-bounded turbulence through explainable deep learning, Nat. Commun. **15**, 3864 (2024).

[38] E. Lagemann, J. Roeb, S. Brunton, and C. Lagemann, A deep learning approach to wall-shear stress quantification: From numerical training to zero-shot experimental application, J. Fluid Mech. **1014**, A1 (2025).

[39] N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F. Hamprecht, Y. Bengio, and A. Courville, On the spectral bias of neural networks, in *Proceedings of the 36th International Conference on Machine Learning*, edited by K. Chaudhuri and R. Salakhutdinov (PMLR, 2019), Vol. 97, pp. 5301–5310.

[40] M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. Barron, and R. Ng, Fourier features let networks learn high frequency functions in low dimensional domains, in *Advances in Neural Information Processing Systems*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Curran Associates, Inc., 2020), Vol. 33, pp. 7537–7547.

[41] K. Taira, G. Rigas, and K. Fukami, Machine learning in fluid dynamics: A critical assessment, Phys. Rev. Fluids **10**, 090701 (2025).

[42] M. Chevalier, A. Lundbladh, and D. Henningson, SIMSON: A pseudo-spectral solver for incompressible boundary layer flow, KTH Tech. Rep. Trita-Mek 7 (KTH, 2007).

[43] S. Hoyas, R. Vinuesa, P. Schmid, and H. Nagib, Sensitivity study of resolution and convergence requirements for the extended overlap region in wall-bounded turbulence, Phys. Rev. Fluids **9**, L082601 (2024).

[44] M. Morimoto, K. Fukami, K. Zhang, A. Nair, and K. Fukagata, Convolutional neural networks for fluid flow analysis: Toward effective metamodeling and low dimensionalization, Theor. Comput. Fluid Dyn. **35**, 633 (2021).

[45] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, arXiv:1412.6980.

[46] Z.-Q. Xu, Y. Zhang, and Y. Xiao, Training behavior of deep neural network in frequency domain, in *Neural Information Processing*, edited by T. Gedeon, K. Wong, and M. Lee, Lecture Notes in Computer Science, Vol. 11953 (Springer International Publishing, Cham, 2019), pp. 264–274.

[47] A. Balasubramanian, R. Vinuesa, and O. Tammisola, Prediction of flow and polymeric stresses in a viscoelastic turbulent channel flow using convolutional neural networks, J. Fluid Mech. **1009**, A36 (2025).

[48] C. Yan, S. Foo, V. Trinh, D. Yeung, K. Wong, and W. Wong, Fourier amplitude and correlation loss: Beyond using L2 loss for skillful precipitation nowcasting, *Advances in Neural Information Processing Systems*, edited by A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (Curran Associates, Inc., 2024), Vol. 37, pp. 100007–100041.

[49] D. Fuoli, L. Van Gool, and R. Timofte, Fourier space losses for efficient perceptual image super-resolution, in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (IEEE, 2021), pp. 2360–2369.

[50] M. Bode, M. Gauding, K. Kleinheinz, and H. Pitsch, Deep learning at scale for subgrid modeling in turbulent flows: regression and reconstruction, in *High Performance Computing*, edited by M. Weiland, G. Juckeland, S. Alam, and H. Jagode, Lecture Notes in Computer Science, Vol. 11887 (Springer, Cham, 2019), pp. 541–560.

[51] K. Fukami, K. Fukagata, and K. Taira, Super-resolution analysis via machine learning: A survey for fluid flows, Theor. Comput. Fluid Dyn. **37**, 421 (2023).

[52] L. Guastoni and R. Vinuesa, A new perspective on the simulation of stochastic problems in fluid mechanics with diffusion models: Fluid dynamics, Nat. Mach. Intell. **7**, 816 (2025).

[53] A. Balasubramanian, Code repository of Sharper predictions: The role of loss functions for enhanced turbulent-flow sensing, https://github.com/KTH-Complex-fluids-group/2026_Sharper-predictions-The-role-of-loss-functions-for-enhanced-turbulent-flow-sensing (2026).