

Highlighting the main differences between neural networks and linear stochastic estimates The cause of reversals in fluid flow

V Krishna Veni¹, I Kalpana²,

Assistant Professor^{1,2},

Mail Id : krishnaveni.v@visvodayata.ac.in, Kalpana.i@visvodayata.ac.in,

Department of Mathematics,

PBR VISVODAYA INSTITUTE OF TECHNOLOGY AND SCIENCE, KAVALI.

Abstract

Neural networks and linear stochastic estimation (LSE) are widely employed as effective methods for fluid-flow regressions. Two classical fluid-flow issues exist: (1) the state estimation in a turbulent channel flow from the wall features, and (2) the estimate of high-order proper orthogonal decomposition coefficients from low-order counterparts for a two-dimensional cylindrical flow. Within the framework of these two issues, we examine their underlying distinctions. In the first problem, LSE is contrasted with a multi-layer perceptron (MLP). In the channel flow example, a nonlinear model called the convolutional neural network (CNN) is employed to manage high-dimensional flows. Because of the nonlinear activation functions, the nonlinear NNs perform better than the linear techniques in both scenarios. Error-curve studies are also used to analyze the response of weights in models and the estimate error. Our study aims to visualize the error-curve resistance against noise while emphasizing the key distinctions between the covered tools for fluid-flow regressions. NNs have shown their enormous promise in data estimation, control, reduced-order modeling, and turbulence modeling at the fundamental level 3–5. Despite the present excitement for the use of NNs, we are forced to rely on linear theories for the time being because of their transparency and generalizability.

INTRODUCTION

If we can learn from the interaction between linear theories and neural networks, we can design more sophisticated NN-based methods for fluidflow analysis. The connection between NNs and linear approaches is of special significance here. Low-dimensionalization was achieved using an autoencoder (AE) in Milano and Koumoutsakos's research. AE based on multi-layer perceptron (MLP) with linear activation functions is comparable to proper orthogonal decomposition (POD)¹⁰, according to Tey's research. Murata et al.¹¹ studied the similarity between an AE based on a convolutional neural network (CNN) and a POD in more recent research. NNs' nonlinear activation functions were shown to have significant strength. Nair & Goza¹² recently examined the MLP, Gappy POD, and linear stochastic estimation (LSE) for the POD coefficient estimate from local sensor measurements of the laminar wake of a flat plate using local sensor data. LSE, MLP, and CNN are used to compare the abilities of linear methods and NNs by comparing two canonical fluid flow regression problems whose complexities are different; 1. estimation of high-order POD coefficients from low-order counterparts of a flow around a cylinder¹³, and 2. state estimation from information on the wall in a turbulent channel flow¹⁴. By concentrating on the influence of biases in NNs, optimization approaches and resilience against noisy inputs, we seek crucial features for their estimates.

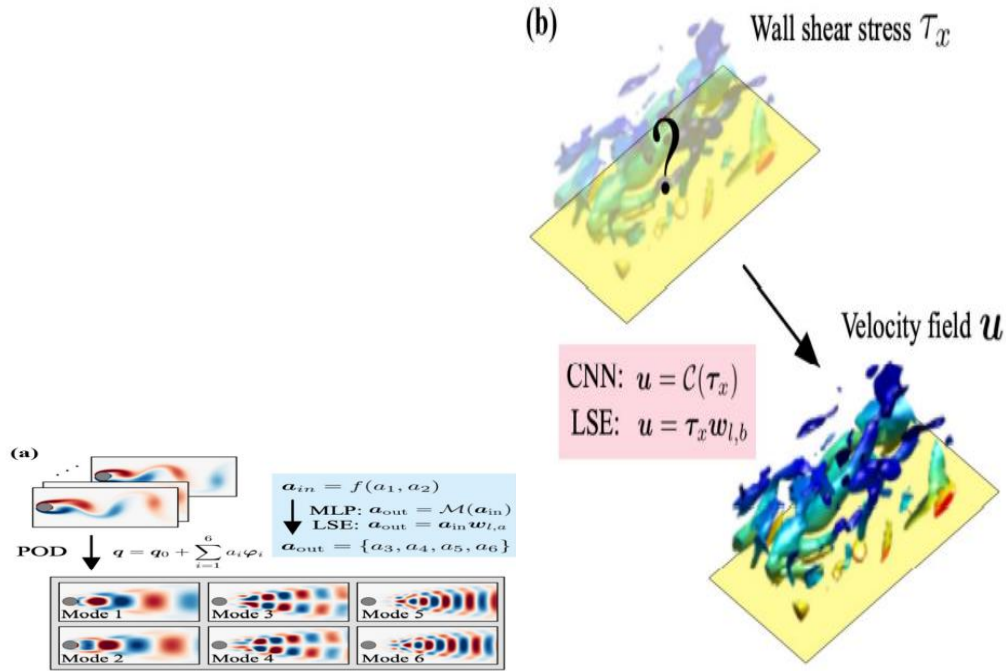


Figure 1 depicts the regressions used in this work for covered fluidflow regressions. Estimation of a flock of birds around a cylinder using POD coefficients A nonlinear activation function is used once biases b have been introduced,

$$q_i^{(l)} = \phi \left(\sum_j W_{ij}^{(l)} q_j^{(l-1)} + b_j^{(l)} \right),$$

allowing for state estimation in a turbulent channel flow.

when the layer index l is used. Back-propagation¹⁸ is used to optimise the weights of all edges W_{ij} to minimise a loss function. E. While the number of output nodes is four in the current MLP, there are four hidden units (3rd to 6th POD modes). Sect. 3.1 provides information on the various circumstances that are taken into account when determining the number of input nodes. It is possible to express the issue in terms of weights w_m (representation of both W and b) within the MLP as follows: $a_{out} = a_3, a_4, \dots, a_5, a_6$, and w_m represents both W and b .

$$w_m = \operatorname{argmin}_{w_m} \|a_{out} - \mathcal{M}(a_{in}; w_m)\|_2.$$

L2 error norm is the loss function we utilise. Notably, the current loss function does not take penalty terms like the Lasso and Ridge penalties into account due to the difficulty and expense of setting the hyperparameter for regularisation¹⁷.

Convolutional neural network

. Due to its completely linked nature, the number of edges in the MLP may expand when dealing with high-dimensional data, such as fluid flows. A convolutional neural network (CNN) [19] has been largely acknowledged as a promising choice for dealing with fluid flow issues. [3, 20]. This work utilises a two- and three-dimensional CNN combination for state estimation. A filter operation is used to extract spatial information from input data in the convolutional layer, which is a basic process in CNNs.

$$q_{ijm}^{(l)} = \phi \left(b_m^{(l)} + \sum_{k=0}^{K-1} \sum_{p=0}^{H-1} \sum_{s=0}^{H-1} h_{pskm}^{(l)} q_{i+p-C, j+s-C, k}^{(l-1)} \right),$$

A convolution layer has K filters, and the bias is $b^{(l)}$ ($m = \text{floor}(H/2)$). For two- or three-dimensional flow fields, this filter technique provides efficient data management. Similar to Eq. (3), the three-dimensional convolution of the CNN is conducted in Eq. (3). The font size H has been set to 5 in the current document. There are two types of activation functions used in the CNN: the filtering activation function and the MLP activation function (Eq. (1)). In order to minimise a loss function, back-propagation is used in the optimization of h filters. Three-dimensional turbulent state u is generated by combining two-dimensional and three-dimensional CNNs (details will be provided in Sect. 3.2). In order to solve the weight optimization issue for the CNN C , we may write it as

$$w_c = \text{argmin}_{w_c} \|u - \mathcal{C}(\tau_{x,wall}; w_c)\|_2.$$

The purpose of this research is to compare the LSE and NN models. To provide a fair comparison, it is best to use the same number of weights for LSE and NNs. Singular value decomposition (SVD) is used to align the number of weights

Layer	Data size	Layer	Data size
Input $\tau_{x,wall}$	(32, 32, 1)	14th 2D Conv.	(32, 32, 16)
1st 2D Conv.	(32, 32, 8)	15th 2D Conv.	(32, 32, 32)
2nd 2D Conv.	(32, 32, 8)	16th 2D Conv.	(32, 32, 32)
3rd 2D Conv.	(32, 32, 8)	Reshape	(32, 32, 32, 1)
4th 2D Conv.	(32, 32, 8)	1st 3D Conv.	(32, 32, 32, 8)
5th 2D Conv.	(32, 32, 8)	2nd 3D Conv.	(32, 32, 32, 8)
6th 2D Conv.	(32, 32, 16)	3rd 3D Conv.	(32, 32, 32, 16)
7th 2D Conv.	(32, 32, 16)	4th 3D Conv.	(32, 32, 32, 16)
8th 2D Conv.	(32, 32, 16)	5th 3D Conv.	(32, 32, 32, 8)
9th 2D Conv.	(32, 32, 16)	6th 3D Conv.	(32, 32, 32, 8)
10th 2D Conv.	(32, 32, 16)	7th 3D Conv.	(32, 32, 32, 8)
11th 2D Conv.	(32, 32, 16)	8th 3D Conv.	(32, 32, 32, 8)
12th 2D Conv.	(32, 32, 16)	9th 3D Conv.	(32, 32, 32, 3)
13th 2D Conv.	(32, 32, 16)	Output u	(32, 32, 32, 3)

in the CNNs and the LSE in order to do this.

Table 1. Structure of 2D-3D CNN C for turbulent channel flow example.

In the next part, we'll go through the LSE's weight distribution. It is possible to describe the operation as $w_l = UVT$, where U is a singular vector and VT is an output vector with the same rank as the input vector, and (nrnk) is a diagonal matrix with singular values on the diagonal. According to our preliminary test, the number of LSE weights based on SVD reduction nwLSE, SVD is 197,632 ($= \text{ninputnrnk} = 1024 \times 193$), which is the current number of LSE

weights. As a result of this SVD-based weight reduction, we now have the 197,000 weights needed to compute the CNN parameters. According to Table 1, the 196 608 weights in our CNN have the following parameters: Estimation using linear stochastic models. Linear stochastic estimation (LSE)^{14,21} is used as a comparison to the neural networks (NNs). $Q = Pw_l$, where n_{data} represents the number of training snapshots, n_{input} represents the number of input attributes, and n_{output} represents the number of output attributes. In this study, we express target data Q w/R $n_{input} n_{output}$ as a linear map with respect to input data P w/R $n_{data} n_{input}$, where $Q = Pw_l$. The linear map w_l may be created by minimization in a manner similar to that used for NN optimization.

$$w_l = \operatorname{argmin}_{w_l} \|Q - w_l P\|_2 = (P^T P)^{-1} P^T Q.$$

For the sake of a fair comparison with NNs, penalization terms are not taken into account in the current loss function. It's important to note, however, that the LSE is analytically optimised by solving Equation (5), whereas the NNs are numerically optimised by back-propagation. Because of the difference in optimization methodologies, we may also compare NNs and LSEs. It is therefore possible to apply the optimal weights w_l to test data.

Results

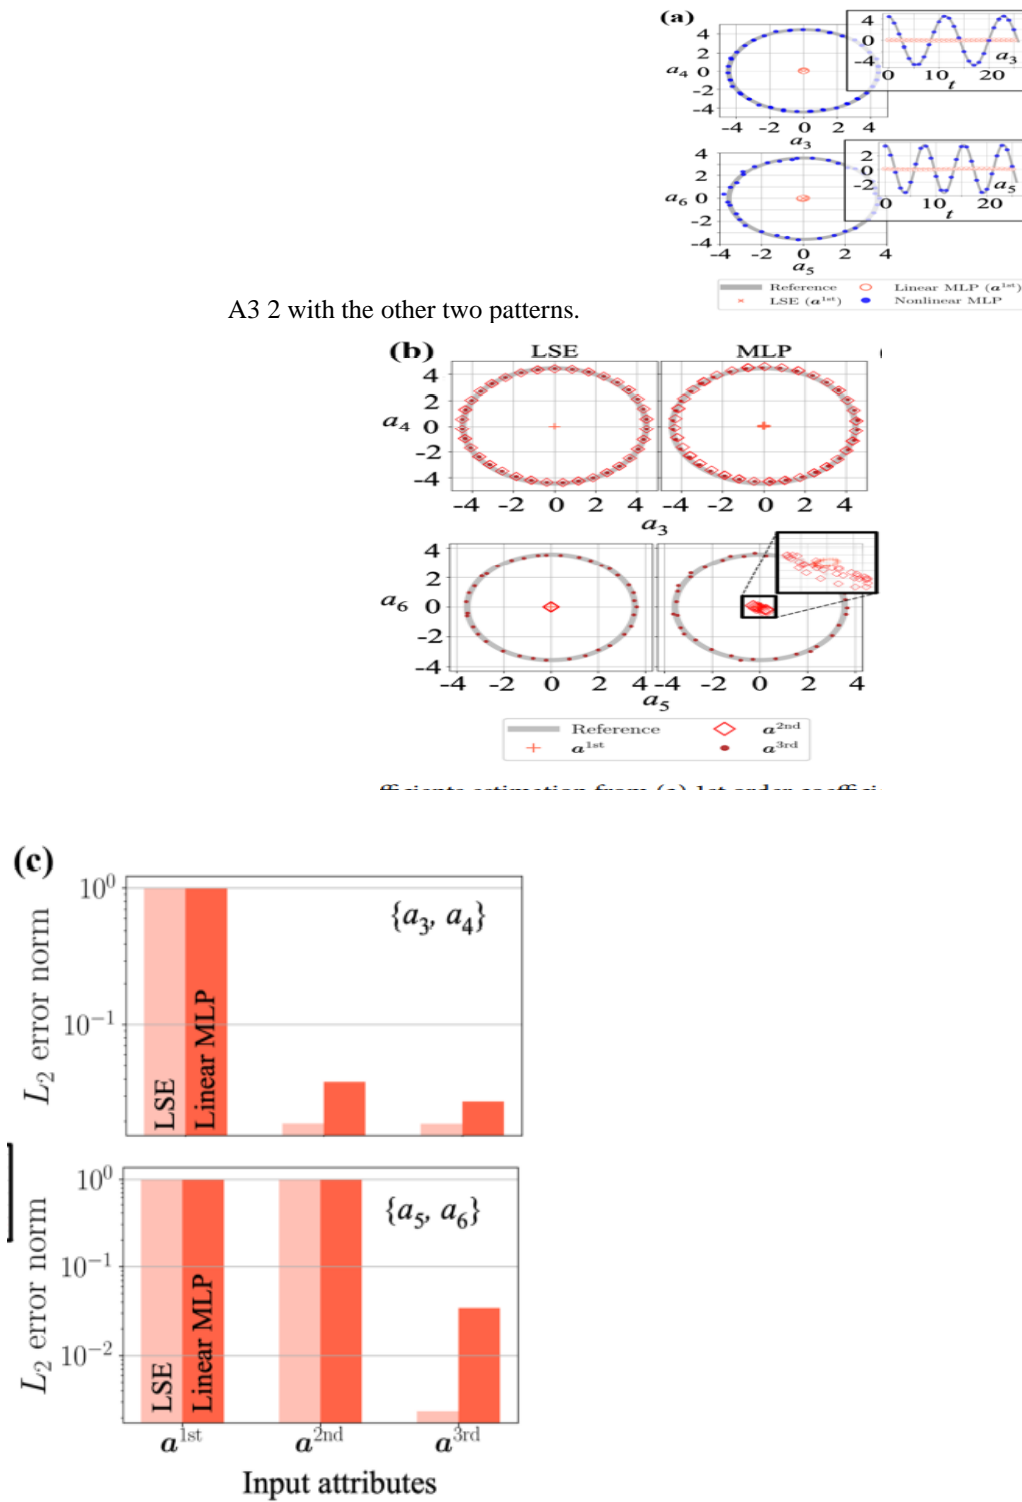
Example 1: The POD coefficient of a two-dimensional cylinder wake at $Re_D = 100$. First, as shown in Fig. 1a, we aim at estimating the two-dimensional high-order POD coefficients of an initial cylinder wake at $Re_D = 100$ from knowledge about its low-order counterparts, a_{in} , such that F1 designates the model for this purpose. We next use this model to estimate $a_{out} = a_{in}$. LSE and MLP serve as model F1 in this experimentation. A two-dimensional direct numerical simulation is used to create flow snapshots (DNS). Incompressible Navier–Stokes equations are the

$$\nabla \cdot \mathbf{u} = 0, \quad \partial_t \mathbf{u} + \nabla \cdot (\mathbf{u}\mathbf{u}) = -\nabla p + Re_D^{-1} \nabla^2 \mathbf{u},$$

governing equations.

The pressure is represented by p , and the velocity vector by \mathbf{u} . It is non-dimensionalized using the fluid density, the free-flow velocity, and the diameter of the cylinder. There are 25.6 and 20.0 units in the computational domain, respectively, and the cylinder's centre lies at $(x, y) = (9, 0)$. Using an immersed boundary method²², the no-slip boundary condition is applied to the cylinder's surface, with a grid spacing of 0.025 and a time step of $2.5 \cdot 10^{-3}$. The

DNS uses $(N_x, N_y) = (N_x, N_y)$ grid points (1024, 800). $(N_x, N_y) = (N_x, N_y) = (N_x, N_y)$ for the POD, which extracts the vorticity field around the cylindrical surface (384, 192). For each POD mode, we take the average of the flow field's temporal average and divide it by one to get the POD basis. Our current MLP and LSE are trained using 5000 snapshots. We do not separate the MLP training data into training and validation for the sake of comparison with LSE. Additional 5000 images are taken into account throughout the analysis as well. Here, we compare the LSE with the linear MLP and the nonlinear MLP using the ReLU activation function²³. TeReLU has a reputation for being a strong contender in the fight against disappearing gradients. Using just a_{1st} with the LSE and linear MLP, we consider three patterns of input $a_{in} = F(a_1, a_2)$: $a_{1st} = A1, A2, A1A2, A2\ 1, A2\ 2, A2\ 1a2\ 2, A3\ 1$, and



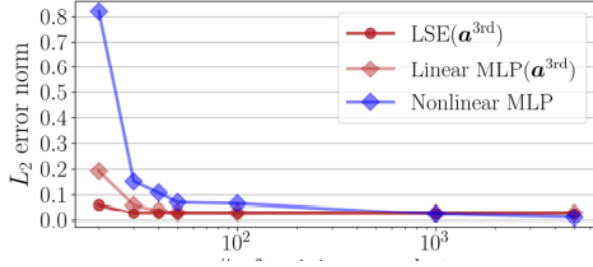


Figure 3. Dependence of the L2 error norm for the POD coefficient estimation on the number of training snapshots.

a nonlinear MLP In this study, we investigate if the nonlinear function in MLP works to capture such nonlinear interactions, and whether the linear models may also be used if the essential combinations of nonlinear terms are given as input. As illustrated in Fig. 2a, the first-order coefficients $a_{1st} = a_1, a_2$ may be used to estimate $a_{out} = a_{3st}$, which is equal to $F_1(a_{1st})$. For both coefficient maps, the nonlinear MLP has a clear advantage over the LSE and the linear MLP. Linear MLP has an L2 error norm of 1.00; non-linear MLP has an L2 error norm of 0.0119; and LSE has an L2 error norm of 1.00. According to this, estimates may benefit from consideration of the nonlinear activation function. As seen in Figs. 2b and 2c, this nonlinearity may be restored even if we solely utilise linear techniques by providing the right inputs, $a_{in} = a_{2nd}$ and a_{3rd} . In order to get an accurate estimate of a_3, a_4 , we can use input up to the second order term a_{2nd} with the LSE and the linear MLP, but for a_5, a_6 we need the third order term a_{3rd} . According to Loiseau et al.13's findings, this tendency is similar to the one seen here. As demonstrated in Fig. 2c, the LSE beats the linear MLP with high-order coefficient inputs. However, as we'll see later, there's a significant difference between the two in terms of noise resilience. The LSE and MLP are then compared in terms of the amount of training data that is readily available. Figure 3 shows the relationship between the L2 error and the amount of training snapshots. For the linear models, we use the third-order coefficients $a_{in} = a_{3rd}$ and the first-order coefficients for the nonlinear MLP, as input for the models. When the number of training snapshots is minimal, the LSE displays its advantage over the MLP models. Because the MLP has a greater degree of flexibility than the LSE, this is the case. In terms of noise resilience, the MLP clearly outperforms the LSE, revealing the linear MLP's inherent advantages over the LSE. Consider the Gaussian white noise denoted by the signal-to-noise (SNR), which is defined as the ratio of the variances of input data and background noise, $2 \text{ data}/2 \text{ noise}$. In Fig. 4, the results for noisy inputs are summarised. We employ the LSE and the linear MLP with $a_{in} = a_{3rd}$ as our linear models. As a point of reference,

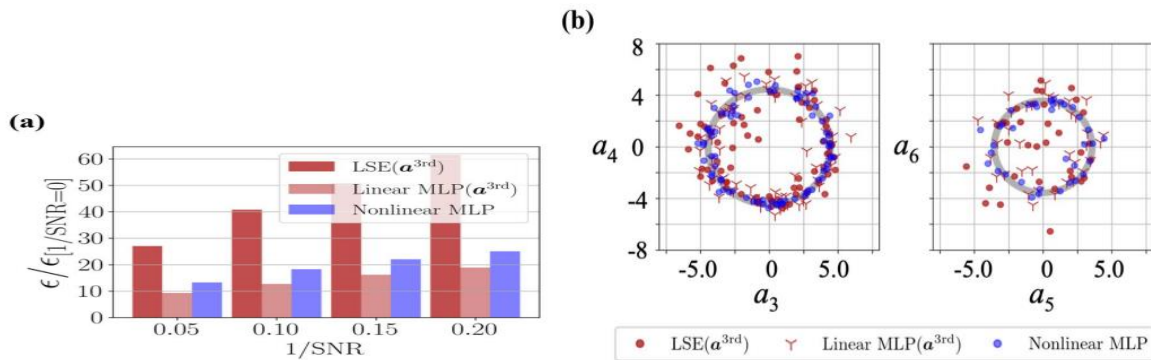
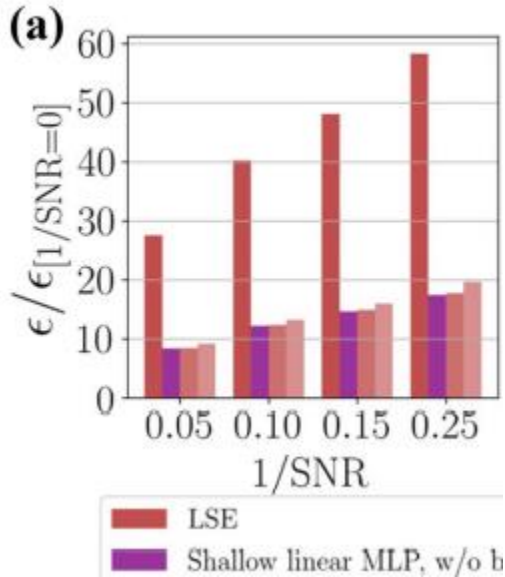


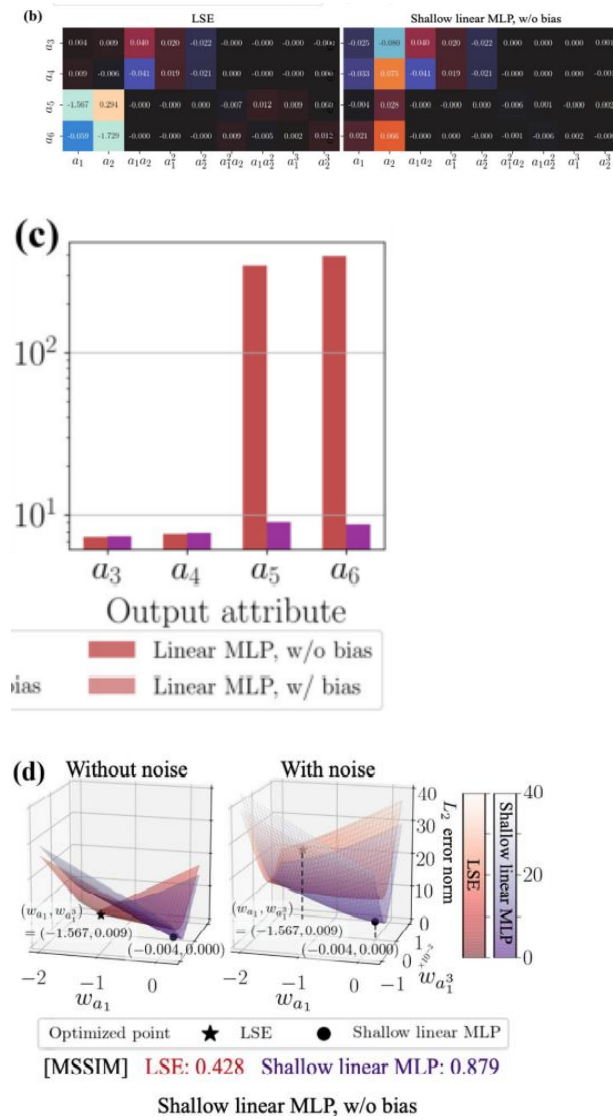
Figure 4: POD coefficients estimation is robust to noisy input. (a) The relationship between the L2 error norm $[\epsilon/\epsilon_1 / \text{SNR}=0]$ increasing $1/\text{SNR}$ and noise magnitude. For linear models, $a_{in} = a_{3rd}$, and for nonlinear MLP models, with $a_{in} = a_{1st}$, we get an a_{out} with $1/\text{SNR} = 0.1$. The nonlinear MLP is likewise monitored with $a_{in} = a_{1st}$. Compared to the covered MLPs, LSE's reaction is far more precise. Biases in MLPs (as described in Eq. (1)) and differences in optimization approaches are to blame for this problem, which may be summarised as follows: In this section, we'll look for the component that has the most impact on noise robustness. The LSE and three kinds of MLPs are taken into account for the study of the primary contribution as previously described. 1. LSE model: the same LSE model as used

earlier. 2. LSE model: This is the same linear MLP model that was used before, but with a bias of M1. Third, the biases are removed from the linear MLP model M1 in order to examine their influence on the bias. When the number of weights is aligned with the LSE, it may be used to evaluate the differences between optimization strategies. Figure 5a shows the relationship between the L2 error norm and the noise level of each model. The covered MLP models have essentially little difference in terms of performance. These results imply that noise resilience is not much influenced by the bias or the number of layers. The shallow linear MLP is still more resilient than the LSE, despite the fact that the model structure is same. Figure 5b shows the weights in the LSE and the shallow linear MLP to illustrate this idea. Second-order term input weights are optimised to the same values as those for the first-order term input, although the first-order term input weights differ significantly. This is a result of a difference in the optimization techniques. Figure 5d depicts an error surface for the inputs of a_1 , 1 and a_3 1 with the output of a_5 , as illustrated. The output a_5 is one of the LSE's most sensitive components to noise, as seen in Fig. 5c. This is why this input-output combination was chosen. To understand why the optimal solutions differ, it may be necessary to look at the techniques employed to optimise them. The presence of noise significantly alters the error-surface form of the LSE, but only marginally alters the shape of the MLP. The MSSIM (mean structural similarity index measure)²⁴ may be used to quantify this difference. If E_r RMN, we apply this to the elevation of each error surface (i.e., without E_r , and with E_r , and with noise E_n , where M , N are the number of samples on each weight axis). MSSIM may be stated numerically as

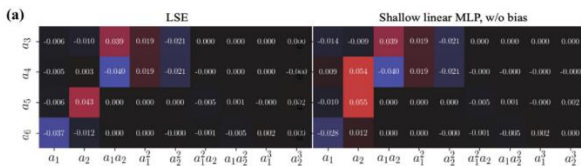
$$\text{MSSIM}(E_r, E_n) = \frac{1}{M'N'} \sum_{i=1}^{M'} \sum_{j=1}^{N'} \text{SSIM}(e_{r,ij}, e_{n,ij}), \quad \text{SSIM}(e_r, e_n) = \frac{(2\mu_r\mu_n + C_1)(2\sigma_m + C_2)}{(\mu_r^2 + \mu_n^2 + C_1)(\sigma_r^2 + \sigma_n^2 + C_2)}.$$

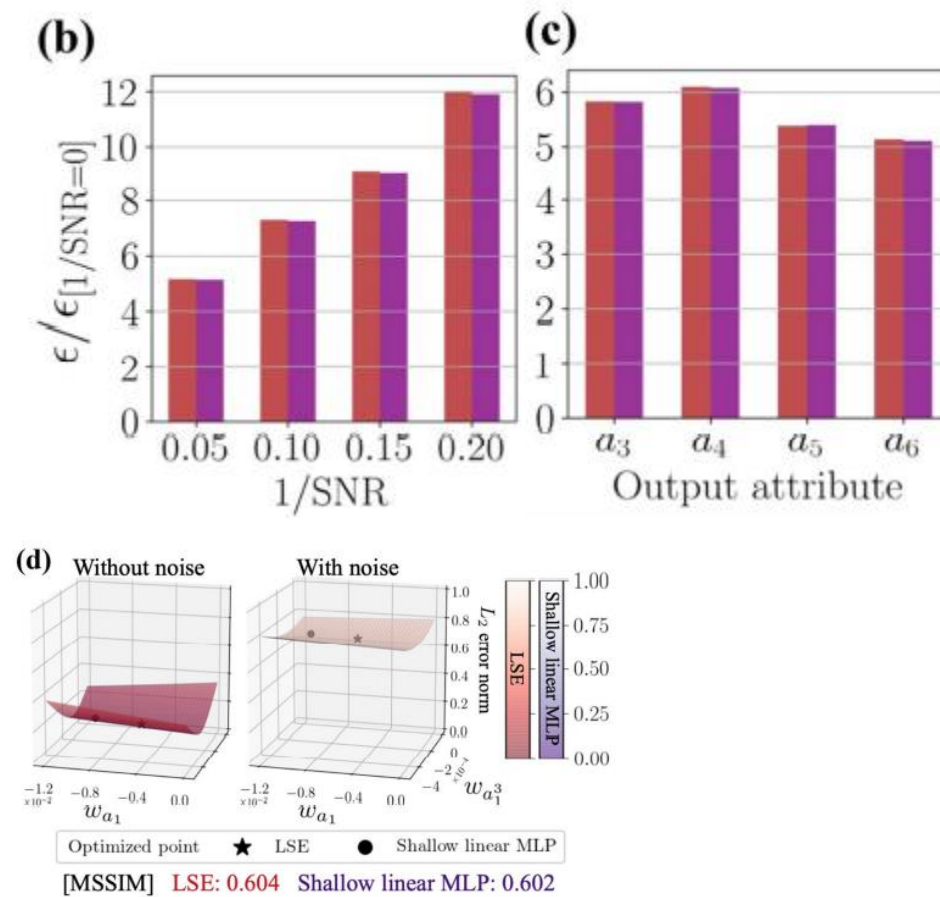
In order to compare two photos, This uses the mean and standard deviation of the two images to determine how similar they are. If you want to get MSSIM, you need to take the average across the whole picture of M' and N' , where M' is equal to $M + m$ and $N + n + 1$, in order to calculate the SSIM in that tiny window of two photos. In accordance with Wang et al.²⁴, we used $C_1, C_2=0.16, 1.44$ as the constant values in Eq. (7). Error surfaces are significantly deformed in LSE, as seen in Fig. 5d, where MSSIM for shallow linear MLPs (0.879) and LSE (0.428) are shown. The LSE's optimal point in Fig. 5d's error space is raised vertically as a result of the substantial distortion of the error surface. According to this, the LSE weights derived analytically ensure the global optimum solution over the training data, but this solution may not be optimal from the standpoint of noise resilience. Although it isn't an exact global best across the training data, the MLP weights are a numerical solution achieved using back-propagation, which makes it a noise-resistant solution. The shallow linear MLP (without bias) and LSE both benefit from adding $\text{SNR} = 0.05$ noisy data to the training data, thus we do so (as described in Fig. 6) to account for the LSE's properties. Fig. 6a illustrates this.



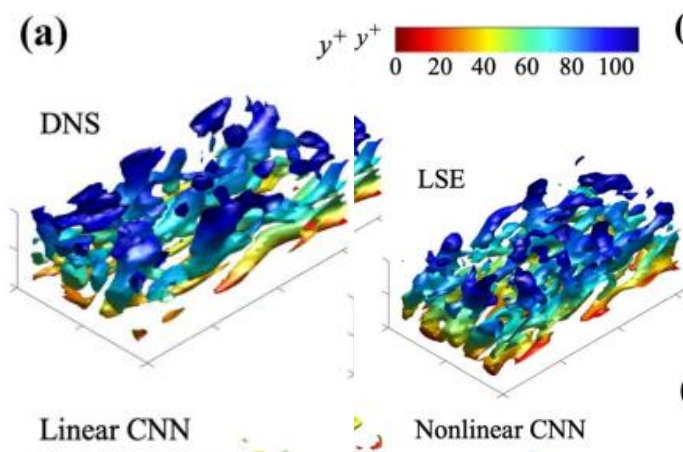


Comparison of LSE and MLP based on bias and optimization approach shown in Figure 5. To increase the L2 error norm from the initial error [1/SNR=0] without the noisy input, we may use the following strategy: (b) LSE and shallow linear MLP weight values. [1/SNR=0] and the output POD coefficients depend on the growing ratio [1/SNR=0] of L2 error norms. d) The error surface surrounding the optimised location is shown graphically.





Using the noisy training data, we can see a comparison between the linear MLP and the LSE in Figure 6. (a) LSE and shallow linear MLP weight values. For example, instead of using the noisy input, we may boost the L2 error norm by a ratio of $[1/SNR=0]$. $[1/SNR=0]$ and the output POD coefficients depend on the growing ratio $[1/SNR=0]$ of L2 error norms. d) The error surface surrounding the optimised location is shown graphically.



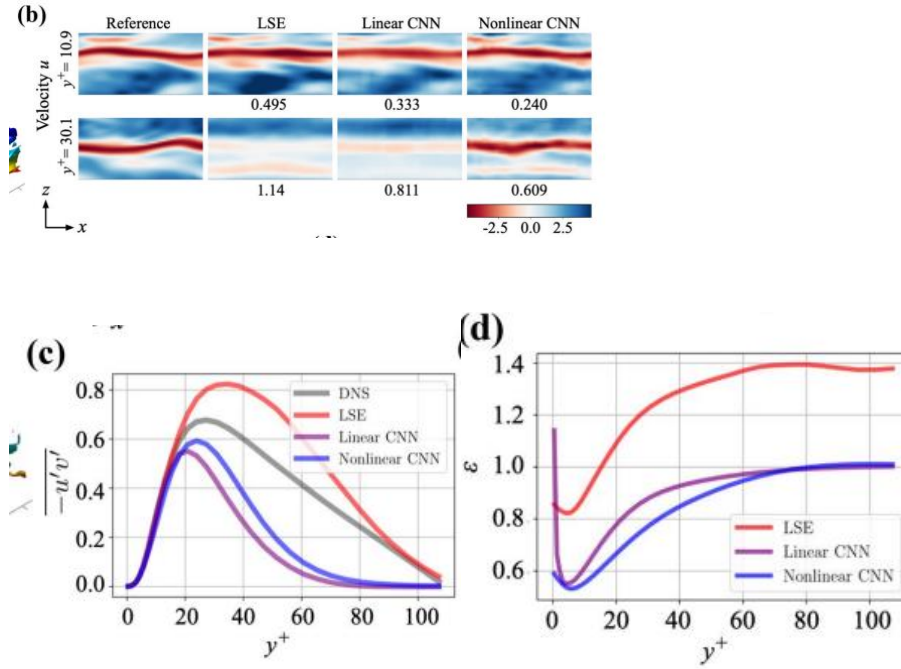


Figure 7 illustrates how streamwise wall-shear stress may be used to estimate turbulent channel flow. Q^+ criteria isosurfaces ($Q^+ = 0.005$). Assuming that (a) is the case, then (b) is the case as well. For each velocity attribute, the Te values beneath the contours indicate the L2 error norm for the corresponding Te . c) Reynolds shear stress $u'v'$. (d) Ensemble L2 error norm dependence on y position for three velocity components.

While there was a difference in weight values in Figure 5, there is almost no difference here. Figs. 6b–d show comparable results for a variety of other studies. According to Tese, training data may be made more robust by include noise. As you can see in Figure 6b, the growing ratio of $[1/SNR=0] = 5$ is not nearly huge since the initial error $[1/SNR=0]$ is tiny ($[1/SNR=0] = 4.49 \times 10^2$ for the LSE, whereas the shallow linear MLP has a little larger error of $4.52 \times 10^2 \times 10^2$).

Example 2: velocity field in a minimal turbulent channel flow at $Re_\tau = 110$. We next compare the LSE and NNs for a more difficult issue, which requires a more thorough analysis. According to the streamwise wall shear stress input x in the least turbulent channel flow at $Re = 110$, it is possible to estimate the velocity field (u) of the channel. Model F2 is an example of a model used in Example 2 and is denoted by the symbol F2. LSE and CNN are the models F2 that we employ. Training data is created by numerically solving the incompressible Navier–Stokes equations with a three-dimensional DNS.

$$\nabla \cdot \mathbf{u} = 0, \quad \partial_t \mathbf{u} + \nabla \cdot (\mathbf{u}\mathbf{u}) = -\nabla p + Re_\tau^{-1} \nabla^2 \mathbf{u},$$

where u and p denote the velocity and pressure velocities, respectively^{25,26}. The channel half-width and the friction velocity u are non-dimensionalized. L_x , L_y and L_z are the computing domains with the number of grid points of $(N_x, N_y \text{ and } N_z) = 0.5. (32, 64, 32)$. There are two distinct grids: one is symmetrical in the x and z dimensions while the other is asymmetrical in the y direction. As the subscript $+$ specifies the number of wall units, $\tau^+ = 0.0385$. The models are trained using 10,000 pictures. We do not separate the training data into training and validation in order to get comparable results to LSE. In addition, we create 2700 new snapshots for the evaluation. Model-estimated flow channels are shown in Fig. 7a. The amount of weights inside the CNN and LSE are almost identical, as mentioned

before in Section 2.2. A CNN with a different filter operation from that of a fully-connected MLP is also taken into account to see whether the same results can be obtained as for the linear MLP in the cylinder example using a linear CNN (i.e., the same CNN structure, but with a linear activation function). The second invariant of the velocity gradient tensor Q^+ of 0.005 is used to illustrate the estimated fields. However, it should be noted that the LSE field only offers turbulent-like structure that is distinct from that seen in the DNS, whilst the DNS field displays qualitatively comparable behaviour to that of the reference LSE data. Figure 7b shows the $x-z$ sectional streamwise velocity distributions to study this site. However, the nonlinear CNN is superior to both the linear CNN and LSE in terms of accuracy, which is surprising given the discovery of Q isosurfaces. At $y^+ = 30.1$, the advantages of the nonlinear technique are most apparent. Fig. 7c also shows the time-averaged Reynolds shear stress $u'v'$. Despite its high error level, the form of the LSE curve seems to be very consistent (albeit exaggerated). A decent reconstruction of LSE in Fig. 7a suggests that the LSE model's predicted fowfield approximates the DNS data in a time-ensemble sense, even if it does not match for each instantaneous field. In Fig. 7d, we examine the estimation's relationship to the y -coordinate. Due to the paucity of data in the vicinity of the wall, it is difficult to determine the velocity field away from the

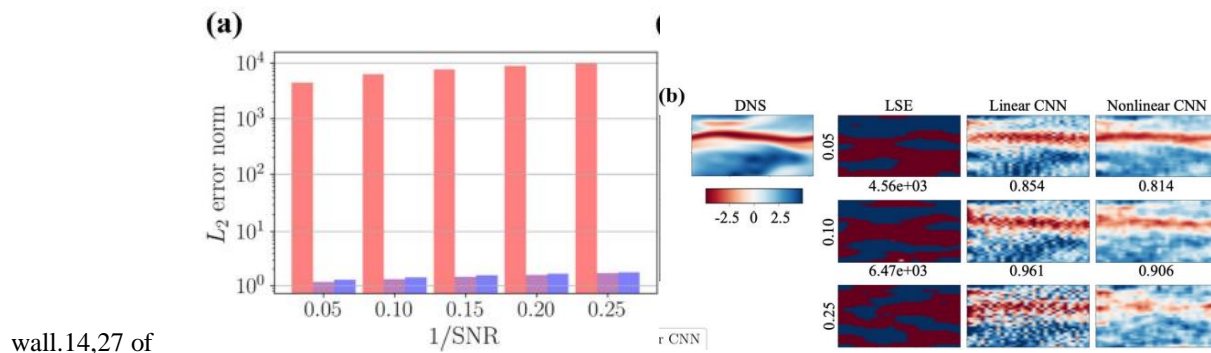


Figure 8 illustrates the robustness of the system in the case of turbulence. The (a) $1/SNR$ on the noise magnitude. Assumed streamwise velocity fluctuation at $y^+ = 10.9$ (b) Contours. L_2 error norm across three velocity components is shown by the Te values shown below. Percentage of SNR (SNR) on the side of the contour is shown by the numbers

Conclusions

The standard fluidflow regression issues were used to study the fundamental differences between NNs and linear methods: A flow is studied in two ways: first, to determine its POD coefficients, then, to determine its state using wall measurements in a turbulent channel flow. SSE, MLP, and a convolutional neural network were evaluated against LSE and SSE with a multi-layer perceptron (CNN). Nonlinear function efficacy may be found in both regression tasks. As a result, we discovered that a linear model might be used as a substitute for a nonlinear model by providing an adequate mix of inputs. A similar discovery has been made in numerous prior studies^{28–30}, thus we may assume that the combination of nonlinear activation functions and adequate inputs can boost the prediction capabilities of the model further. It was also found that linear NNs were more resistant to noise than the LSE, and that visualising the error surface helped to uncover why. Noise resilience may be attributed to the difference in optimization strategies, according to error surfaces. The learning process of nonlinear neural networks (NNs) may be unstable depending on the issue setting since it is based on a gradient approach. Despite this, we discovered that NNs have a number of advantages over linear neural networks. We may not be able to get to a tolerable weight valley, particularly if the issue is multimodal and training data is scarce. In this regard, the LSE may provide us a theoretically sound answer. These features may be combined for further enhancement of NN learning pipelines, such as transfer learning^{32,33} so that the NN may be launched from a reasonable solution and achieve sufficient noise resilience, as shown in our study. For this reason, we may include a loss function connected with physics-based presuppositions in order to minimise the weights used in both neural networks and lasso-style estimation (LSE). Incorporating a physical loss function with both systems might be a significant route toward practical use. 34–3

References

1. Brunton, S. L., Noack, B. R. & Koumoutsakos, P. Machine learning for fluid mechanics. *Annu. Rev. Fluid Mech.* 52, 477–508 (2020).
2. Fukami, K., Fukagata, K. & Taira, K. Assessment of supervised machine learning for fluid flows. *Teor. Comp. Fluid Dyn.* 34, 497–519 (2020).
3. Fukami, K., Fukagata, K. & Taira, K. Super-resolution reconstruction of turbulent flows with machine learning. *J. Fluid Mech.* 870, 106–120 (2019).
4. Fukami, K., Fukagata, K. & Taira, K. Machine-learning-based spatio-temporal super resolution reconstruction of turbulent flows. *J. Fluid Mech.* 909, A9. <https://doi.org/10.1017/jfm.2020.948> (2021).
5. Fukami, K., Maulik, K., Ramachandra, N., Fukagata, K. & Taira, K. Global field reconstruction from sparse sensors with Voronoi tessellation-assisted deep learning. *Nat. Mach. Intell.* 3, 945–951 (2021).
6. Brenner, M. P., Eldredge, J. D. & Freund, J. B. Perspective on machine learning for advancing fluid mechanics. *Phys. Rev. Fluids* 4, 100501. <https://doi.org/10.1103/PhysRevFluids.4.100501> (2019).
7. Maulik, R., Fukami, K., Ramachandra, N., Fukagata, K. & Taira, K. Probabilistic neural networks for fluid flow surrogate modeling and data recovery. *Phys. Rev. Fluids* 5, 104401. <https://doi.org/10.1103/PhysRevFluids.5.104401> (2020).
8. Duraisamy, K., Iaccarino, G. & Xiao, H. Turbulence modeling in the age of data. *Annu. Rev. Fluid. Mech.* 51, 357–377 (2019).
9. Milano, M. & Koumoutsakos, P. Neural network modeling for near wall turbulent flow. *J. Comput. Phys.* 182, 1–26 (2002).
10. Lumley, J. L. The structure of inhomogeneous turbulent flows. In Yaglom, A. M. & Tatarski, V. I. (eds.) *Atmospheric turbulence and radio wave propagation* (Nauka, 1967).
11. Murata, T., Fukami, K. & Fukagata, K. Nonlinear mode decomposition with convolutional neural networks for fluid dynamics. *J. Fluid Mech.* 882, A13. <https://doi.org/10.1017/jfm.2019.822> (2020).
12. Nair, N. J. & Goza, A. Leveraging reduced-order models for state estimation using deep learning. *J. Fluid Mech.* 897. <https://doi.org/10.1017/jfm.2020.409> (2020).
13. Loiseau, J.-C., Brunton, S. L. & Noack, B. R. From the POD-Galerkin method to sparse manifold models. <https://doi.org/10.13140/RG.2.2.27965.31201> (2018).

14. Suzuki, T. & Hasegawa, Y. Estimation of turbulent channel flow at $Re_\tau = 100$ based on the wall measurement using a simple sequential approach. *J. Fluid Mech.* 830, 760–796 (2006).
15. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagation errors. *Nature* 322, 533–536 (1986).
16. Domingos, P. A few useful things to know about machine learning. *Commun. ACM* 55, 78–87 (2012).
17. Nakamura, T., & Fukagata, K. Robust training approach of neural networks for fluid flow state estimations. Preprint at arXiv:2112.02751 (2021).
18. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. Preprint at arXiv:1412.6980 (2014).
19. LeCun, Y., Bottou, L., Bengio, Y. & Haffner, P. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 2278–2324 (1998).
20. Morimoto, M., Fukami, K., Zhang, K., Nair, A. G. & Fukagata, K. Convolutional neural networks for fluid flow analysis: Toward effective metamodeling and low-dimensionalization. *Teor. Comp. Fluid Dyn.* 35, 633–658 (2021).