Sensing coherent structures from the wall

by

Alejandro Güemes Jiménez

in partial fulfillment of the requirements for the degree of Doctor in
Programa Interuniversitario en Mecánica de Fluidos

Universidad Carlos III de Madrid

Advisors:
Stefano Discetti
Andrea Ianiro

Tutor:
Andrea Ianiro

Leganés, July 2021

Esta tesis se distribuye bajo licencia "Creative Commons **Reconocimiento – No Comercial – Sin Obra Derivada**".

Cover: numerical flow visualization of a turbulent channel flow.

Esta tesis está dedicada a mis padres,
por el apoyo y cariño constante que me han dado.

Acknowledgements

In the first place, I would like to dedicate my first words of gratitude to my supervisors, Andrea and Stefano. I cannot thank you enough your guidance during these 4 years. I will not be able to pay you back for the effort you have made to turn a stubborn person into someone flexible who can face the problems of the world of research. The transformation that I have undergone during these four years is incredible. Everything I have learned from you, both inside and outside the lab, is a gift. Now this stage of my life is coming to an end, but I hope to continue working with you for many years because you have a lot of science inside you are wonderful people. I couldn't have had better supervisors.

My gratitude also extends to Marco, because you were the person who taught me to work in a laboratory. The endless hours in the lab turning the laser on and off must have had a serious effect on me. I just hope that one day we can remove the panels article from our to-do list and we can move to new projects. You should also know that the difficulty of my thesis would have been much greater if I had not had you there with your experience in the water tunnel and to deal with Stefano and, above all, with Andrea. Nor can I forget the other members of the rest of the colleagues from the Experimental Aerodynamics and Propulsion Lab, especially Carlos and Rodrigo. Without the beers when we left the laboratory, I would not have lasted until here.

I would like to thank the people from KTH. Thanks to Ricardo, Hossein and Philipp ,for letting me participate in their projects and giving me the opportunity to use their resources. Luca deserves a special mention, with whom I have been able to work side by side these last two years. I wish you the best for the future.

I cannot forget the Bioengineering and Aerospace engineering department's professors. I would like to mention especially Manolo, Oscar, Manuel Sanjurjo and Manuel Soler. I have learnt a lot from you, thank you very much.

A mis amigos Andrea, Andrés, Dani, Gonzalo, Joaquín, Marta, Nere, Proven, Salme, Sergio, Sisa y Tamara, compañeros de viaje durante todos estos años de estudio, les reservo un lugar especial. Las horas interminables en la biblioteca o en las salas de ordenadores del edificio siete al final tuvieron su recompensa. A los compañeros de Boeing, con especial cariño para Javi, Pablo, Miguel y

David, que me dieron la oportunidad de poder trabajar con ellos, y a Quique, que ha mantenido la confianza en mí.

A toda mi familia, de la que siempre he recibido un inmenso cariño. Para mi, no hay mejores personas de las que estar rodeado, os quiero. A Sandra, por ser mi apoyo durante estos últimos años. Espero poder estar a tu lado cuando me necesites de la misma forma que has estado tu ahí para mi. Te quiero.

Finalmente, a papá y mamá. Los logros de esta tesis son más vuestros que míos. No hay palabras que hagan justicia al gran esfuerzo que habéis hecho todos estos años para educarme, formarme y, sobre todo, quererme. Muchísimas gracias, os quiero.

Sensing coherent structures from the wall

Alejandro Güemes Jiménez
Bioengineering and Aerospace Engineering Department, Universidad Carlos III de Madrid
Avenida de la Universidad, 30, Leganés, Spain

Abstract

Any device or application in which a solid body and a fluid are in relative motion has to deal with the presence of drag forces. Counteract such forces in an efficient and environmentally-friendly way has become very important in recent years, in view of the sustainable-development challenges imposed by global authorities. The interaction between any engineering device and a surrounding fluid occurs most of the time in a turbulent regime. Despite the chaotic appearance of any turbulent flow, a close inspection reveals the presence of patterns that maintain their structure over a period of time. These patterns, known as *coherent structures*, are of utmost importance, since they are a *low-dimensional expression* of a complex highly-dimensional dynamical system that can be used to target affordable control strategies. Taking into account that control models rely on proper identification of the flow state, it is necessary to acquire this information in a way that is sufficiently robust for real applications. For example, the experimental tools normally used in wind tunnels (for instance, velocity field measurement techniques like Particle Image Velocimetry) cannot be used in an airborne flight. It becomes necessary to acquire this information from sensors that can be embedded in the device itself. The present thesis attempts to provide a technique based on deep neural networks that allows to reconstruct turbulent flows and their energy-carrying coherent structures from sensors embedded in the wall. The choice of neural networks, characterized by their ability to approximate non-linear relations, is based on the development of new algorithms that has occurred in recent years due to the availability of new computational resources.

The first part the dissertation aims to develop a model that combines the ability of prediction in nonlinear systems of deep neural networks with the compression capabilities of proper orthogonal decomposition (POD). For that purpose, a direct numerical simulation (DNS) of a turbulent channel flow of $Re_\tau = 1000$ is used. Using wall-shear-stress information, the proposed architecture extracts information into feature maps through a series of convolutional neural networks, which are later converted in POD modes coefficients by using fully-connected neural networks. Since the first POD modes can be ascribed to the most energetic structures in any flow, the proposed networks aim to provide an estimation of the first 10 POD modes. This approach allows to reconstruct wall-parallel flow fields containing the most energetic structures in the flow. This approach is compared with a baseline linear method, proving that neural

networks are able to trace better the nonlinear effect of turbulent flows in the wall-normal direction.

The next step is to improve this model to reconstruct as much flow scales as possible. For such purpose, a DNS of turbulent open-channel flow is used, with $Re_\tau = [180 - 550]$. The aforementioned network is transformed into a fully-convolutional network, which allows to exploit better the spatial organization of the POD modes and retrieve a significant larger amount of flow energy, surpassing 90%. The results show that the prediction quality decays with increasing wall distance. This is because the imprint of the flow structures at the wall is related in a linear way for all scales that have a size of the order of the distance to the wall. In contrast, the smaller scales do not impose their imprint directly on the wall, and the nonlinear relationship between the two attenuates with distance. For these results, the energy spectrum of the flow has also been evaluated, showing that the accuracy loss of the predictions begins with the smallest scales. However, the predictions are capable of recovering the largest structures of the flow, which are the ones that transport the bulk of energy. Furthermore, this network is compared with a network that does not exploit the compressibility advantages of POD. The results are overall similar, although the network based on POD appear slightly more robust for increasing distance from the wall.

While the previous networks have used fully-resolved wall information coming from DNS simulation, n real applications it is often difficult to embed such a fine mesh of sensors in the wall. Among the different types of architecture that comprise neural networks, generative adversarial networks (GANs) have stood out for their ability to recover the small-scale details of low-resolution images. In fact, they have already been proposed as a tool to recover resolution in flow measurements. In this thesis the use of GANs to retrieve high-resolution wall information is presented. GANs performance is evaluated for different information loss ratios, showing that even in extreme cases GANs are capable of recovering the most energetic scales of the flow footprint at the wall.

Finally, as the use of GANs has been proven successful for both wall and flow fields, it is proposed to use them to make a direct estimation of the flow from wall quantities. The results offer a significant improvement in flow reconstruction when compared to the previous results. This improvement can be ascribed to the complexity of GAN architectures, which allows to optimize better the amount of information fed into the network without suffering vanishing-gradient problems. The reconstruction performances are also tested with low-resolution input data. Although the predictions get worse when moving away from the wall also in this case, the rate is lower than in the previous cases. This allows to have reconstructions that capture the coherent structures away from the wall.

Key words: wall turbulence, coherent structures, deep learning.

Detección de estructuras coherentes desde la pared

Alejandro Güemes Jiménez
Departamento de Bioingeniería e Ingeniería Aeroespacial, Universidad Carlos III de Madrid
Avenida de la Universidad, 30, Leganés, España

Resumen

Cualquier dispositivo o aplicación en la que un cuerpo sólido y un fluido estén en movimiento relativo tiene que lidiar con la presencia de fuerzas de arrastre. A la vista de los desafíos de desarrollo sostenible que han impuesto las autoridades mundiales, contrarrestar estas fuerzas de una manera eficiente y amistosa con el medio ambiente se ha convertido en un tema de suma importancia. La interacción entre un dispositivo de carácter industrial y el fluido que lo rodea ocurre la mayor parte del tiempo en el régimen turbulento. A pesar de la apariencia caótica que tienen los flujos turbulentos, una inspección de cerca revela la presencia de patrones que mantienen su estructura durante algún tiempo. Estos patrones, conocidos como *estructuras coherentes*, son de suma importancia, ya que son una *expresión de baja dimensión* de un sistema dinámico complejo de alta dimensionalidad que puede usarse para desarrollar estrategias de control asequibles. Teniendo en cuenta que los modelos de control se basan en la identificación adecuada del estado del flujo, es necesario adquirir esta información de una manera que sea lo suficientemente robusta para aplicaciones reales. Por ejemplo, las herramientas experimentales que se utilizan normalmente en los túneles de viento, como la velocimetría de imágenes de partículas para la medición del campo de velocidad, no se pueden utilizar durante el vuelo de un avión regular. Es necesario adquirir esta información a partir de sensores que se puedan incrustar en el propio dispositivo. La presente tesis intenta proporcionar una técnica basada en redes neuronales profundas que permita reconstruir flujos turbulentos y sus estructuras coherentes portadoras de energía a partir de sensores embebidos en la pared. La elección de las redes neuronales, caracterizadas por su capacidad para aproximar relaciones no lineales, se basa en el desarrollo de nuevos algoritmos que han aparecido en los últimos años, gracias a los avances de los nuevos recursos computacionales.

La primera parte de esta disertación tiene como objetivo desarrollar un modelo que combine la capacidad de predicción en sistemas no lineales de redes neuronales profundas con las capacidades de compresión de la Descomposición Modal Ortogonal (Proper Orthogonal Decomposition, POD). Para ello, se utiliza una simulación numérica directa (Direct Numerical Simulation, DNS) de un flujo de canal turbulento, con $Re_\tau = 1000$. Midiendo los esfuerzos de cortadura en la pared, la arquitectura propuesta extrae información en mapas de características a través de una serie de redes neuronales convolucionales, que luego se convierten en coeficientes de modos POD mediante el uso de redes neuronales completamente conectadas. Dado que los primeros modos POD se

pueden atribuir a las estructuras más energéticas en cualquier flujo, las redes propuestas tienen como objetivo proporcionar una estimación de los primeros 10 modos POD. Este enfoque permite reconstruir campos de flujo paralelos a la pared que contienen sus estructuras más energéticas. Este método propuesto se compara con un método lineal de referencia, lo que demuestra que las redes neuronales pueden rastrear mejor el efecto no lineal de los flujos turbulentos en la dirección normal de la pared.

El siguiente paso es mejorar este modelo para reconstruir tantas escalas del flujo como sea posible. Para ello se utiliza un DNS de flujo turbulento en canal abierto, con $Re_\tau = [180 - 550]$. La mencionada red se transforma en una red totalmente convolucional, lo que permite aprovechar mejor la organización espacial de los modos POD y recuperar una cantidad significativamente mayor de energía de flujo, superando el 90%. Los resultados muestran que la calidad de la predicción decae a medida que aumenta la distanca a la pared. Esto se debe a que la huella de las estructuras del flujo en la pared está relacionada de forma lineal para todas las escalas que tienen un tamaño del orden de la distancia a la pared. Por el contrario, las escalas más pequeñas no imponen su huella directamente en la pared y la relación no lineal entre las dos se atenúa con la distancia. Para estos resultados, también se ha evaluado el espectro de energía del flujo, mostrando que la pérdida de precisión de las predicciones comienza con las escalas más pequeñas. Sin embargo, las predicciones son capaces de recuperar las estructuras más grandes del flujo, que son las que transportan la mayor parte de la energía. Además, esta red se compara con una red que no aprovecha las ventajas de compresibilidad de la POD. Los resultados son en general similares, aunque la red basada en la POD parece ligeramente más robusta cuando aumenta la distancia desde la pared.

Si bien las redes anteriores han utilizado información de pared completamente resuelta proveniente de la simulación de DNS, en aplicaciones reales es difícil incrustar una malla tan fina de sensores en la pared. Entre los distintos tipos de arquitectura que componen las redes neuronales, las redes de generación por confrontación (Generative Adversarial Networks, GAN) se han destacado por su capacidad para recuperar detalles a pequeña escala en imágenes de baja resolución. De hecho, ya se han propuesto como herramienta para recuperar la resolución en medidas de flujo. En esta tesis se presenta una alternativa basada en GANs para recuperar información de pared de alta resolución. Se han evaluado las prestaciones de las GANs para diferentes ratios de pérdida de información, demostrando que incluso en casos extremos las GANs son capaces de recuperar las escalas más energéticas de la huella de flujo en la pared.

Finalmente, dado que el uso de las GANs ha demostrado ser exitoso tanto para los campos de pared como para los de flujo, se propone usarlos para hacer una estimación directa del flujo a partir de las cantidades de pared. Los resultados ofrecen una mejora significativa en la reconstrucción del flujo en comparación con los resultados anteriores. Esta mejora se puede atribuir a la complejidad de las arquitecturas GAN, que permite optimizar mejor la cantidad

de información introducida en la red sin sufrir problemas de desvanecimiento de gradiente. Las prestaciones de reconstrucción también se han probado con datos de entrada de baja resolución. Aunque en este caso las predicciones también empeoran al alejarse del muro, la tasa de disminución es menor que en los casos anteriores. Esto permite tener reconstrucciones que capturan las estructuras coherentes lejos de la pared.

Palabras clave: turbulencia de pared, estructuras coherentes, aprendizaje profundo.

Preface

This thesis deals with the reconstruction of coherent structures in wall-bounded turbulent flows by means of deep neural networks. A brief introduction of the basic concepts and methods is presented in the first part. The second part contains three articles. The papers are adjusted to comply with the present thesis format for consistency, but their contents have not been altered as compared with their original counterparts.

Paper 1. Güemes, A., Discetti, S. & Ianiro, A., 2019. *Sensing the turbulent large-scale motions with their wall signature.* Phys. Fluids **31** (12), 125112.

Paper 2. Guastoni, L., Güemes, A., Ianiro, A., Discetti, S., Schlatter, P., Azizpour, H. & Vinuesa, R., 2021. *Convolutional-network models to predict wall-bounded turbulence from wall quantities.* J. Fluid Mech., **Under review**.

Paper 3. Güemes, A., Discetti, S., Ianiro, A., Sirmacek, B., Azizpour, H. & Vinuesa, R., 2021. *From sparse wall measurements to turbulent velocity fields through deep learning.* Phys. Fluids, **Under review**.

July 2021, Leganés
Alejandro Güemes Jiménez

Division of work between authors

The advisors for the project are Professor Andrea Ianiro (AI) and Professor Stefano Discetti (SD).

Paper 1. The data pre-processing and the application of reconstruction algorithms have been carried out by Alejandro Güemes (AG). The paper has been written by AG and reviewed by SD and AI.

Paper 2. The simulations have been generated by Luca Guastoni (LG). The analysis of linear and nonlinear predictive models has been carried out by AG and LG. The paper has been written by LG and AG with input from SD, AI, Philipp Schlatter (PS), Hossein Azizpour (HA) and Ricardo Vinuesa (RV).

Paper 3. The application of predictive models to reconstruct flow fields from coarse wall measurements has been carried out by AG. The paper has been written by AG and reviewed by SD, AI, Beril Sirmacek (BS), HA and RV.

Other publications

The following papers, although related, are not included in this thesis.

Güemes, A., Sanmiguel Vila, C., Örlü, R., Vinuesa, R., Schlatter, P., Ianiro, A. & Discetti, S., 2019. *Flow organization in the wake of a rib in a turbulent boundary layer with pressure gradient.* Exp. Therm. Fluid Sci. **108**, pp. 115–124.

Güemes, A., Fajardo, P. & Raiola, M., 2020. *Experimental assessment of RANS models for wind load estimation over solar-panel arrays.* Appl. Sci. **11** (6), 2496.

Castellanos, R., Sanmiguel Vila, C., Güemes, A. & Discetti, S., 2020. *On the uncertainty of boundary-layer parameters from Ensemble PTV data.* Meas. Sci. Technol. **32**, 084006.

Conferences

Part of the work in this thesis has been presented at the following international conferences. The presenting author is underlined.

Güemes, A., Vaquero, A., Flores, O., Discetti, S. & Ianiro, A., *Identifying the wall signature of large-scale motions with Extended POD.* 8th iTi Conference in Turbulence. Italy, 2018.

Sanmiguel Vila, C., Güemes, A., Örlü, R., Vinuesa, R., Schlatter, P., Ianiro, A. & Discetti, S., *Pressure-gradient effects on turbulent boundary layers with square ribs.* 15th International Conference on Fluid Control, Measurements and Visualization. Italy, 2019.

Güemes, A., Ianiro, A. & Discetti, S., *Experimental assessment of large-scale motions in turbulent boundary layers.* 13th International Symposium on Particle Image Velocimetry. Germany, 2019.

Güemes, A., Discetti, S. & Ianiro, A., *Low-order flow reconstruction with convolutional neural networks.* 1st VKI Lecture Series on Machine Learning for Fluid Mechanics: Analysis, Modeling, Control and Closures. Belgium, 2020.

Foroozan, F., Güemes, A., Guerrero, V., Ianiro, A. & Discetti, S., *Unsupervised discovery of boundary layer transition theory.* 1st VKI Lecture Series on Machine Learning for Fluid Mechanics: Analysis, Modeling, Control and Closures. Belgium, 2020.

Guastoni, L., Güemes, A., Ianiro, A., Discetti, S., Schlatter, P., Azizpour, H. & Vinuesa, R., *Predictions in wall-bounded turbulence through convolutional-network models using wall quantities.* APS Division of Fluid Dynamics. USA, 2020.

Arivazhagan, G.B., Guastoni, L., Güemes, A. Ianiro, A., Discetti, S., Schlatter, P., Azizpour, H. & Vinuesa, R., *Predicting the near-wall region of turbulence through convolutional neural networks.* 13th International ERCOFTAC Symposium on Engineering, Turbulence, Modeling and Measurements. Greece, 2021.

Güemes, A., Tober, H., Discetti, S., Ianiro, A., Sirmacek, B., Azizpour, H. & Vinuesa, R., *Reconstructing flow fields from coarse wall measurements with deep learning.* 19th International Symposium on Flow Visualization. China, 2021.

Güemes, A., Sanmiguel Vila, C. & Discetti, S., *High-resolution PIV with semi-supervised GANs.* 19th International Symposium on Flow Visualization. China, 2021.

Foroozan, F., Güemes, A., Castellanos, R., Raiola, M., Ianiro, A., & Discetti, S., *Correlation between wall and flow fields in turbulent boundary layer: a synchronized measurement.* 19th International Symposium on Flow Visualization. China, 2021.

Güemes, A., Sanmiguel Vila, C. & Discetti, S., *GANs-based PIV resolution enhancement without the need of high-resolution input.* 14th International Symposium on Particle Image Velocimetry. USA, 2021.

Cuellar, A., Güemes, A., Discetti, S. & Ianiro, A., *Convolutional neural networks for flow sensing in wall turbulence.* 8th International Congress of Serbian Society of Mechanics. Serbia, 2021.

Contents

Acknowledgements v

Abstract vii

Resumen ix

Preface xii

Part I - Overview and summary

Chapter 1. Introduction 1
- 1.1. Objectives 3
- 1.2. Thesis structure 4

Chapter 2. Wall turbulence and coherent structures 6
- 2.1. The structure of wall-bounded turbulence 6
- 2.2. On the presence of coherent structures 9
- 2.3. Coherent structures with proper orthogonal decomposition 12
- 2.4. The imprint of coherent structures at the wall 16

Chapter 3. Deep neural networks as a tool for turbulence 20
- 3.1. Introduction to deep neural networks 20
- 3.2. Application of deep neural networks in fluid mechanics 22
- 3.3. POD-based convolutional networks for flow reconstruction 27
- 3.4. Super-resolution GANs for wall and flow measurements 29

Chapter 4. Main contributions and conclusions 34
- 4.1. Future work 35
- 4.2. Paper highlights 36

Bibliography 38

Part II - Papers

Summary of the papers 51

Paper 1. Sensing the turbulent large-scale motions with their wall signature 53

Paper 2. Wall-bounded turbulence from wall quantities 79

Paper 3. From sparse wall measurements to turbulent velocity fields 125

Part I

Overview and summary

Chapter 1

Introduction

Any body in relative motion with respect to its surrounding fluid faces drag forces. Thus, it will expend energy to conserve its total momentum, which will otherwise be reduced due to the drag forces acting on the body. Among the different drag sources, one of the most important is the friction drag exerted by the surrounding fluid on the body's wall or surface. Jiménez (2012) reported that *"roughly half the energy spent in transporting fluids through pipes and canals, or vehicles through air and water, is dissipated by turbulence in the immediate vicinity of walls"*. In the case of aviation industry, friction drag accounts for more than half of the total aircraft drag (Schrauf 2005). For instance, greenhouse gas emissions emitted by the aviation industry in the European Union were 3.6% of the total in 2016, and 13.4% of transport emissions (European Environment Agency *et al.* 2019). Consequently, it is of utmost importance to take into account this effect during the design process of any engineering device that involves a fluid in motion, such as aircraft, ships, windmills, or turbomachinery, to name a few examples. Based on this premise, world authorities have proposed different goals to mitigate any unnecessary energy consumption as well as its effect on the environment. The European Commission, e.g., has set the goal of reducing greenhouse emissions until Europe becomes a carbon neutral region by 2050 (European Commission 2018). In the same way, the Paris agreement sets the objective of keeping the global temperature increase below 2°C with respect to pre-industrial levels and pursue efforts to keep it to 1.5°C (UNFCCC 2015). Within this framework, a strong investment is taking place to better understand the forces arising from the interaction between a solid object in motion with respect to its surrounding fluid. This knowledge may allow the development of control techniques that improve the efficiency of these devices, therefore reducing their energy consumption.

The mechanism behind friction forces in fluids is strongly dependent on the flow regime, either laminar or turbulent. The latter is highly predominant in natural flows and real-life applications. To provide a rigorous and unique definition of turbulence is a challenging task; for instance, Tsinober (2009) reports more than 10 different definitions by leading scientists in the field of turbulence. Here the definition by Bradshaw (2013) is cited, which states that *"turbulence is a three-dimensional time-dependent motion in which vortex*

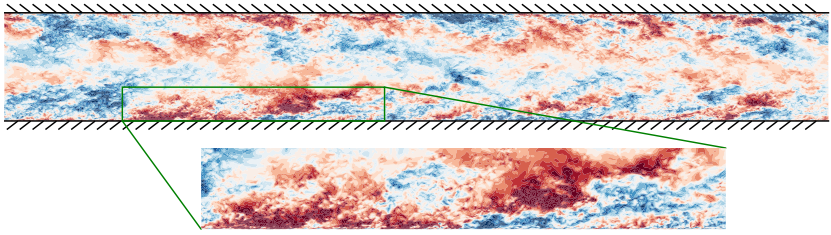

Figure 1.1: Instantaneous velocity fluctuations of a turbulent channel flow. Red colour refers to positive fluctuations, while blue refers to negative ones. Data has been acquired from a direct numerical simulation of a turbulent channel flow available at the Johns Hopkins Turbulence Database.

stretching causes velocity fluctuations to spread to all wavelengths between a minimum determined by viscous forces and a maximum determined by the boundary conditions of the flow. It is the usual state of fluid motion except at low Reynolds numbers". In other words, turbulence is a chaotic three-dimensional phenomenon, with nonlinear dynamics that involve a wide range of scales, both in space and time, which is often modelled as a series of eddies or whirls of different sizes. There are several theories that explain how these different-size eddies relate to each other, but in general they are based on the energy cascade view proposed by Richardson (1920) and extended by Kolmogorov (1941). This view states that energy production occurs at the largest scales, and then energy flows to the smallest eddies where it is dissipated into heat.

Despite being chaotic, among the multiple scales of turbulence it is possible to identify patterns, as intuitive from the inspection of figure 1.1. It provides a clear indication that, within the chaos of turbulence, patterns or structures that maintain their spatial coherence for a time can be identified. Jiménez (2018*a*) defined these flow regions, commonly known as *coherent structures*, as *"eddies with enough internal dynamics to behave relatively autonomously from any remaining incoherent part of the flow"*. According to the theory of Richardson and Kolmogorov, they are responsible for carrying the bulk of energy that will eventually be dissipated at the smallest scales. In wall-bounded flows, this process occurs close to the wall. Unfortunately, identifying these coherent structures and differentiating them from the rest of the range of scales in the fluid are a complex task, which has received the attention of turbulence researchers in the last 40 years (Hussain 1983; Jeong *et al.* 1997; Jiménez 2012). If we zoom into the big vortices or eddies seen in figure 1.1, it can be seen that they are composed of smaller eddies. This process can be repeated until we reach the dissipative scales, where energy dissipation is taking place. Coherent structures and their dynamics can be controlled by means of actuation techniques with the aim of modifying the flow, reaching a new state that implies a gain (Hof *et al.*

2010). While controlling coherent structures focuses on the large scales that carry the bulk of energy, control could be also applied to the near-wall viscous region, where the energy is dissipating, as shown by Jiménez (1994). It is at this point that the question of which methods are more efficient to flow control arises, whether they are those that focus only on the viscous region close to the wall, where the energy is dissipating, or those that focus on controlling the largest scales that carry the bulk of energy. Thanks to the state-of-art high-resolution numerical and experimental data that have been produced in the last years, it has been possible to better understand the importance of coherent structures, as shown by Jiménez (2018*a*). The knowledge acquired from direct numerical simulations (DNSs) and experiments of simple-geometry flows can be translated to more complex applications, such as the flow over commercial aircraft wings (Vinuesa *et al.* 2018), turbines (Ouro *et al.* 2017), insects (Yao & Yeo 2020) or inside biological systems (García-Villalba *et al.* 2021).

Turbulent flows can be modelled through the Navier-Stokes equations as dynamical systems with high dimensionality. The existence of coherent structures represents an opportunity to approximate turbulence as a system with a much lower dimensionality, which can be tractable, for instance, for real-life control applications. Acquiring an approximate representation of coherent structures for simple-topology flows is relatively accessible through numerical simulations, where the entire flow state is known, or experiments, where different measurement techniques may be used to acquire complementary information of the flow. However, when dealing with real-life applications this is not practical. Carrying out a real-time numerical simulation of a turbulent flow affecting an industrial application is not feasible, both due to the current-available computational capabilities and to the uncertainty in the initial and boundary conditions of the system. In the case of experiments, available techniques are usually intrusive, as flow visualization or hot-wire anemometry, or require an experimental setup that could be too complex to install in real-life application, such as particle image velocimetry (PIV, see Adrian (1991)). For instance, figure 1.2 shows the experimental apparatus required by PIV to acquire instantaneous visualizations of the flow around a test model, which evidently cannot be installed on an aircraft in operation. To overcome this problem, measurement techniques that do not interfere with the flow and can be embedded in the wall are of interest, such as hot films or pressure taps. These wall-sensing techniques allow to acquire wall quantities such as pressure, shear stress or heat flux, which can later be used to infer the flow state far from the wall. In any case, these techniques also pose a technological challenge, since their spatial and temporal resolution must be sufficient to capture the widest possible range of scales.

1.1. Objectives

The main objective of this dissertation is to develop a methodology to know the instantaneous state of a turbulent flow, focusing on its coherent structures, from wall signals employing sensors embedded in the wall. For that purpose, this thesis

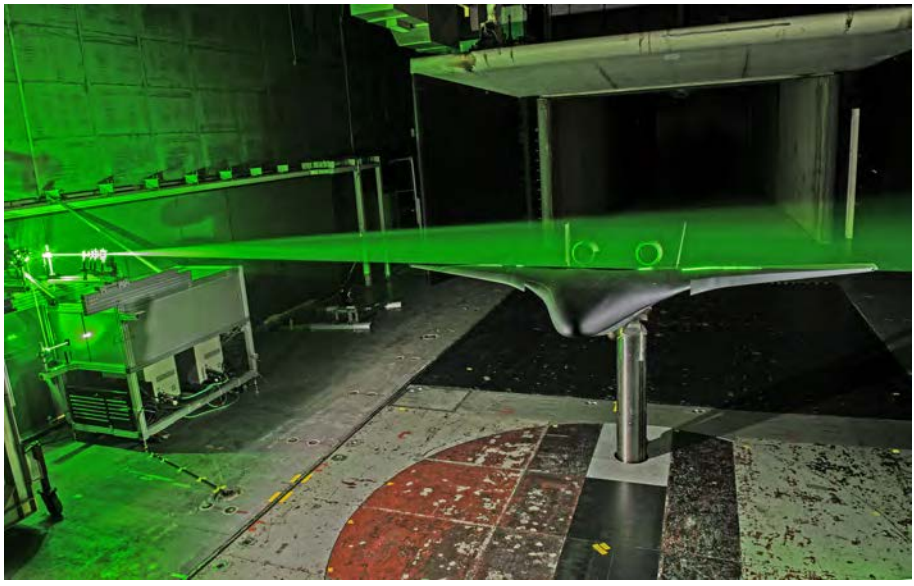

Figure 1.2: Experimental setup at NASA Langley Research Center to investigate the flow around a hybrid-wing-body model. Courtesy of NASA.

explores the relation between coherent structures and shear-stress and pressure fields at the wall, for which both linear and nonlinear reconstruction techniques have been considered. While linear techniques have already been assessed previously for this type of problem, the exhaustive application of nonlinear methods, and specifically of neural networks, has not begun to be explored until the last decade. This study is very timely, thanks to the availability of advanced neural network architectures and new computational resources. Taking this into account, the objective is to develop a set of techniques based on deep learning to address the nonlinear relationship between the quantities measured at the wall and the characteristics of the flow.

1.2. Thesis structure

The present thesis is structured in two parts. Part I starts with an introduction (Chapter 1) and continues with a review of the state of the art of wall turbulence (Chapter 2), focusing on the relationship between the structures that populate turbulent flows and their footprint on the wall. Chapter 2 also introduces the proper orthogonal decomposition, a statistical tool commonly used to represent coherent structures. Chapter 3 presents a review of the main characteristics of neural networks, as well as their applicability as a tool for the study and control of turbulence. The neural networks that have been developed in this thesis for the reconstruction of turbulent flows from wall probes are also presented.

Finally, Chapter 4 summarizes the main contributions and conclusions of this thesis.

The original contributions that have been produced during the development of this thesis are presented in Part II. The first article presents the original version of the neural network developed during this thesis. The second article improves this architecture, making it capable of predicting larger fields and capturing a greater range of energy scales. Finally, the last contribution evaluates the capacity of generative adversarial networks to improve these predictions when the wall data does not contain all the turbulent scales. The aforementioned papers are:

Paper 1: *Sensing the turbulent large-scale motions with their wall signature*

Paper 2: *Convolutional-network models to predict wall-bounded turbulence from wall quantities*

Paper 3: *From coarse wall measurements to turbulent velocity fields through deep learning*

Chapter 2

Wall turbulence and coherent structures

Following the concepts introduced in Chapter 1, this chapter focuses on wall-bounded flows. Wall-bounded turbulent flows are a type of shear flow where the turbulent cascade is fed with the energy extracted from the instabilities originated in the mean-velocity gradients along the inhomogeneous directions.

The content of this chapter is divided into three main sections. Section 2.1 introduces the main characteristics of wall turbulence, establishing the nomenclature and the dimensionless parameters that will be used in the rest of the present thesis to compare the different types of flows. Next, § 2.2 focuses on the dynamics of coherent structures that populate wall-bounded flows. Finally, § 2.3 shows the effect that these coherent structures have on the wall quantities, as well as the available methodologies in the literature to infer the flow state from those wall quantities. In order to maintain a constant nomenclature throughout this thesis, the main dimensions of space and velocity that are used in both Part I and Part II are defined here. Lower-case letters x, y, and z refer to the streamwise, wall-normal and spanwise directions respectively, with their corresponding instantaneous fluctuating velocities indicated as u, v and w. Pressure fluctuations are referred to as p, while the streamwise and spanwise wall-shear stresses are referred to as τ_{w_x} and τ_{w_z} respectively, using the subindex w as indication of wall quantities. Time-averaged quantities are referred with upper case letters, while root-mean-squared quantities are indicated with $'$.

2.1. The structure of wall-bounded turbulence

The interaction between the surface of a body and the flow that surrounds it gives the rise to wall turbulence. The geometry of these bodies can be overly complex. However, to better understand the involved flow physics, it is worth simplifying the geometrical parameters. Therefore, the study of wall turbulence is usually focused on three canonical problems, namely channel, pipe and boundary-layer flows. Channel flows are used to describe the flow between two parallel infinite plates, which is driven by a favourable pressure gradient parallel to the plates, and whose statistics are independent in the wall-parallel directions. Pipe flows are similar, since they are also driven by a favourable pressure gradient. However, the flow is confined by a circular boundary in the plane normal to the pressure gradient, thus forming the pipe geometry. Last

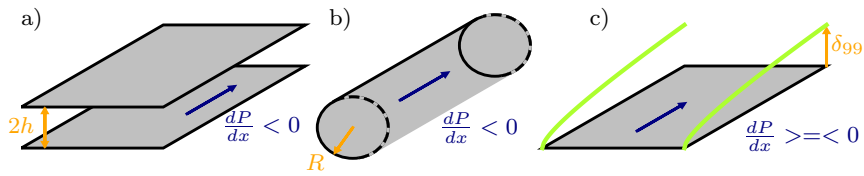

Figure 2.1: Schematic view of the three canonical flows in wall turbulence, which are a) channel, b) pipe and c) boundary-layer flows.

but not least, turbulent boundary layers refer to the region formed when a non-turbulent flow passing parallel to a smooth flat plate decreases its velocity until matching the wall velocity. Unlike channel and pipe flows, boundary layers develop along the direction of the pressure gradient, which for this type of flow could be favourable, but also zero or adverse. These flows are often characterized in terms of their length scale, which indicates the size in the wall-normal direction of the largest structures present in these flows. These length scales are the half-height h of the channel, the radius R of the pipe, and the boundary-layer thickness δ_{99}, but can be defined overall as δ. Figure 2.1 provides a schematic representation of each of these canonical flows. Apart from δ, it is also important to define another characteristic that divides these canonical flows, which is the boundary condition far from the wall. With this in mind, canonical flows can be divided into internal (pipe and channel) and external (boundary layer) flows.

The Reynolds number provides the ratio between the inertia and viscous forces acting on the flow, which is intimately related to how turbulent a flow is. There are many definitions for this parameter and not all of them are commonly used in the study of wall turbulence. The definition that is most used in wall-turbulence research is the friction-based Reynolds number, or Kárman number, which is defined as:

$$Re_\tau = \frac{\delta u_\tau}{\nu}, \tag{2.1}$$

where ν is the dynamic viscosity of the fluid and u_τ is the wall-friction velocity, defined as:

$$u_\tau = \sqrt{\tau_w/\rho}, \tag{2.2}$$

where ρ is the working fluid density. This velocity is fictitious, and it is based on the wall-shear stress τ_w. The viscous length scale, i.e., the smallest eddy scale in the flow, can be defined as:

$$\ell^* = \frac{\nu}{u_\tau}, \tag{2.3}$$

thus equation 2.1 can be redefined as the ratio between the largest and smallest scales present in the flow:

$$Re_\tau = \frac{\delta}{\ell^*}. \tag{2.4}$$

Following these definitions, the superscript + indicates that a quantity has been scaled with viscous quantities, using u_τ for velocities and ℓ^* for distances.

For turbulent boundary layers with freestream velocity U_∞, it is important to define the displacement thickness:

$$\delta^*(x) = \int_0^\infty \left(1 - \frac{U(x,y)}{U_\infty}\right) dy, \tag{2.5}$$

and momentum thickness:

$$\theta(x) = \int_0^\infty \frac{U(x,y)}{U_\infty}\left(1 - \frac{U(x,y)}{U_\infty}\right) dy, \tag{2.6}$$

as integral quantities of the flow. The former provides a value for the amount of fluid mass displaced by the wall, while the latter does the same for the momentum, and they can used to the defined the displacement-thickness-based:

$$Re_{\delta^*} = \frac{\delta^* U_\infty}{\nu}, \tag{2.7}$$

and the momentum-thickness-based:

$$Re_\theta = \frac{\theta U_\infty}{\nu}, \tag{2.8}$$

Reynolds numbers. Since the boundary layer is not homogeneous along the streamwise direction, these quantities evolve in this direction, and they are useful to quantify whether a turbulent boundary layer is well-behaved i.e., it is not affected by the transition-to-turbulence or non-equilibrium effects, or not.

The presence of the wall imposes a hierarchical organization of different-size eddies, since the centre of an eddy with characteristic size ℓ cannot be closer to the wall than $\ell/2$. Consequently, the wall-normal distance to the wall can be used to provide a structured view of wall-bounded turbulence. If the viscous scaling previously defined is used, Navier-Stokes equations can be operated to define a function describing the mean-velocity profile as:

$$U^+ = f(y^+), \tag{2.9}$$

which is only dependent of the wall distance. This definition, proposed by Prandtl (1925) and commonly known as *law of the wall*, only stands when $y/\delta \ll 1$, thus defining the first region of the flow, the *inner layer*. There is not a specific distance where this region starts, but it is widely accepted as $y/\delta < 0.1$. At closer distances to the wall, when $y^+ \sim \mathcal{O}(1)$, equation 2.9 can be redefined as:

$$U^+ = y^+, \tag{2.10}$$

which implies that Reynolds shear stresses are negligible compared to viscous shear stresses. This region is the *viscous sublayer*, and stands for $y^+ \lesssim 5$

(Reichardt 1951). Just above this layer is the *buffer layer*, where the viscous effects are still important but the turbulent shear stresses are not negligible. The limits of this region are not exact, but it is commonly accepted to be $10 \lesssim y^+ \lesssim 100$.

Using outer scaling and dimensional analysis, Von Kármán (1930) proposed a function for the mean-velocity profile of the form:

$$\frac{U_o - U(y)}{u_\tau} = F\left(\frac{y}{\delta}\right), \tag{2.11}$$

where the subindex o stands for outer. Equation 2.11, known as the *velocity defect law*, is applicable in the *outer layer*, where the viscous effects are negligible. Since this is true when $y/\delta \sim \mathcal{O}(1)$, the outer layer lies in wall distances where $y^+ \gtrsim 100$.

For a sufficiently large Reynolds number $Re_\tau \gg 1$, there is a region of the flow where $y/\delta < 0.1$ and $y^+ \gtrsim 100$, which imposes that solutions for equations 2.9 and 2.11 must coexist. This region is known as the *logarithmic layer*, since the mean-velocity profile of the flow can be described by a logarithmic function as:

$$U^+ = \frac{1}{\kappa}\log(y^+) + A, \tag{2.12}$$

where constants κ and A are referred as universal, although scattered values can be found in the literature (George 2007; Marusic *et al.* 2010*b*; Bailey *et al.* 2014). Figure 2.2 provides a schematic representation of the different regions.

With the structured view of wall-bounded flows came the necessity of defining a model describing its dynamic mechanism. Since the first attempt of Theodorsen (1952), several models have been proposed. One of the most accepted models is the *attached-eddy hypothesis* proposed by Townsend (1976). This model defines the logarithmic layer as a region where the energy production and dissipation rates are in balance, describing it as the superposition of several wall-attached vortices of different size aligned in the streamwise direction. Furthermore, Perry & Chong (1982) contributed to this view by proving mathematically that these structures should grow linearly from their birth until death. Another model is the *hairpin forest* proposed by Adrian *et al.* (2000), which describes wall turbulence as a series of horseshoe-vortex packets. Although instantaneous realizations have been provided for transitional boundary layers (Wu & Moin 2010), there is no evidence of their instantaneous existence for high enough Reynolds numbers (Schlatter & Örlü 2010; Sillero *et al.* 2013).

2.2. On the presence of coherent structures

This section is intended to provide a panoramic view of the current knowledge about the coherent structures populating wall-bounded turbulent flows, but for a deeper and exhaustive review the reader is referred to the work of Jiménez (2018*a*). The first quantitative evidences of the existence of coherent structures in the viscous and buffer layers were provided by experimental studies.

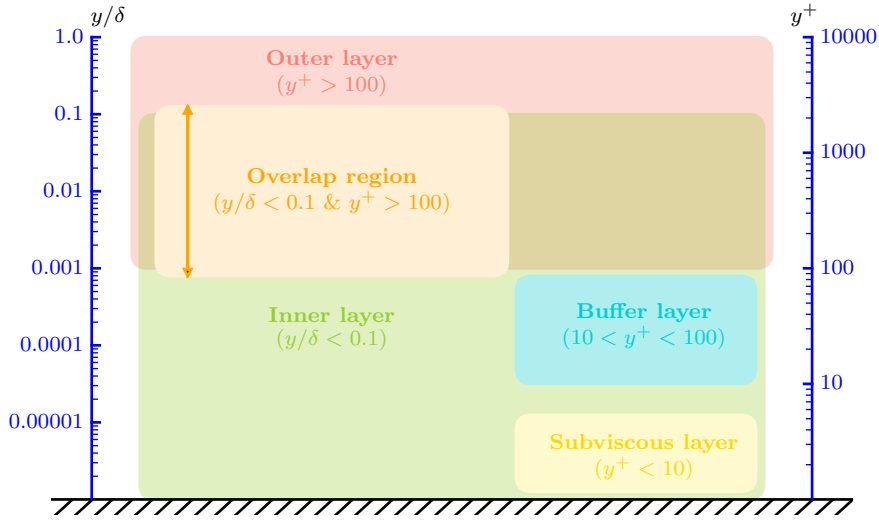

Figure 2.2: Schematic overview of the structure of a wall-bounded turbulent flow with $Re_\tau \approx 10000$.

These structures, known as *u-streaks*, were found with two-point correlations of hot-wire measurements, and were characterized as regions elongated in the streamwise direction with high and low momentum (Kline *et al.* 1967; Favre *et al.* 1967; Blackwelder & Kovasznay 1972). These near-wall structures are typically separated in the spanwise direction by $z^+ = 100$, and have quasi-streamwise vortices staggered between them (Robinson 1991; Jiménez & Moin 1991). Furthermore, Jiménez & Moin (1991) showed that these streaks and vortices sustain an autonomous nonlinear cycle by themselves, without the assistance of energy coming from the outer region, thus proposing a minimal-flow unit to sustain turbulence. This autonomous cycle has been extended in several works (Waleffe 1995; Jiménez & Pinelli 1999; Moehlis *et al.* 2004), but it can be resumed as follows: The quasi-streamwise vortices, with streamwise length scales of the order of 100 ℓ_x^+, generate streaks by stirring energy from the mean-velocity profile, which has a steep gradient at those wall distances. The generated positive- and negative-momentum streaks, with streamwise length scales of the order of $\mathcal{O}(10^3)$ ℓ_x^+, become unstable at some point, creating new vortices.

While Blackwelder & Kovasznay (1972) also provided statistical evidences of the presence of coherent structures in the logarithmic layer, its instantaneous structure could not be verified until the advent of PIV allowed to acquire instantaneous velocity fields (Ganapathisubramani *et al.* 2003; Tomkins & Adrian 2003). These new insights were complemented by computational advances in

numerical simulations that allowed to study coherent structures from a structural point of view (Hoyas & Jiménez 2006). One of the most relevant findings is the imprint that coherent structures populating the logarithmic layer have in the near-wall region (Hoyas & Jiménez 2006, 2008; Hutchins & Marusic 2007; Mathis *et al.* 2009; Dogan *et al.* 2019). This imprint can be ascribed to the linear superposition of scales in the wall-normal direction, but also to the amplitude modulation of the small-scale velocity fluctuations originated from the large coherent structures. The view of minimal-flow units was also extended to the logarithmic region by Flores & Jiménez (2010), showing the presence of the same elongated u-streak and quasi-periodic vortices but with a turbulent behaviour. The size of the minimum-flow unit restricts the range of scales that can be accommodated within, where the minimum size to prevent turbulence from decaying is that proposed by Jiménez & Moin (1991). This gives an insight that coherent structures do not have a single size, but rather overlap over a range of scales, affecting each other, which agrees with the aforementioned attached-eddy hypothesis proposed by Townsend (1976). This theory is also endorsed by the structural analysis of regions with high tangential Reynolds shear stress carried out by Lozano-Durán *et al.* (2012). When moving towards or away from the wall, these structures are referred as *sweeps* and *ejections* respectively, and they are located in the side wall of the long u-streak structures present in the flow. In some sense, they can be interpreted as manifestations of the quasi-streamwise vortices found in the minimal boxes of Jiménez & Moin (1991) and Flores & Jiménez (2010). Moreover, Ganapathisubramani *et al.* (2003) reported in the past that sweeps and ejections contribute significantly to the total amount of Reynolds stress, thus becoming significantly relevant for any wall-bounded flow. In the same manner, Del Alamo *et al.* (2006) showed that structures defined with the discriminant of the velocity gradient, usually known as *vortex clusters*, can also be related to the long u-streak structures present in the flow. Figure 2.3 shows examples of these types of coherent structures. It is important to remark that these structures could be found detached from or attached to the wall. While the former are more numerous, the latter contribute almost the entire volume and are responsible for most of the total momentum transfer (Lozano-Durán & Jiménez 2014).

Regarding their size, coherent structures are of the order of several δ in the streamwise direction, as reported for internal flows in several works (Kim & Adrian 1999; Guala *et al.* 2006; Monty *et al.* 2007). In the case of turbulent boundary layers, Hutchins & Marusic (2007*a*) showed that coherent structures in the logarithmic layer can reach similar lengths, but they present a meandering behaviour in the spanwise direction. One effect of this meandering behaviour can be observed in the energy spectra reported by Monty *et al.* (2009), where the outer energy peak does not appear for external flows. Subsequently, the differences in the outer region between internal and external flows extend up to the buffer and viscous layers. Furthermore, Hutchins & Marusic (2007) provided experimental evidences that coherent structures modify the near-wall scales by

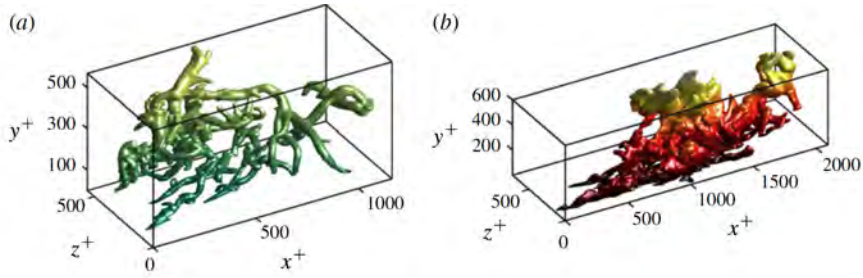

Figure 2.3: a) Vortex cluster and b) sweep wall-attached structures. Adapted from Lozano-Durán & Jiménez (2014).

means of linear superposition and amplitude modulation. This effect was later quantified by Mathis *et al.* (2009).

The new computational and storage resources that have become available in the last decade have made possible to study coherent structures from a dynamical point of view i.e., considering their evolution over time. This is the case of the work reported by Lozano-Durán & Jiménez (2014), where sweeps and ejections were tracked along their *lives*. The results show that large coherent structures attached to the wall have lifetimes proportional to their sizes. This is a significant property for any flow control technique, since coherent structures responsible for the majority of the momentum transfer are those that last more and thus could be easier to control with an actuation. Temporal analysis has also been taken into account in the work of Encinar & Jiménez (2020), which introduces a band-pass filtering method to extend the analysis of sustained turbulence proposed by Jiménez & Moin (1991) to structures of arbitrary size. The identified filtered structures are consistent with the sweeps and ejections of Lozano-Durán *et al.* (2012) and Lozano-Durán & Jiménez (2014), and their length and time scales exhibit a self-similar behaviour with respect to the proposed filter.

Among the different techniques that have been used to study turbulence, the use of data-driven methods stands out. § 2.3 presents the concept of proper orthogonal decomposition (POD), one of those data-driven methods that have been previously applied to turbulence.

2.3. Coherent structures with proper orthogonal decomposition

Principal component analysis is a statistical tool developed in the works of Pearson (1901) and Hotelling (1933) that decomposes any signal as the linear combination of a set of orthonormal basis functions, allowing to capture hidden patterns in large datasets. Introduced in the area of fluid mechanics by Lumley (1967) under the name of POD, it allows to decompose a fluctuating velocity

field $\boldsymbol{u}(\boldsymbol{x},t)$ as a linear combination of orthonormal basis functions, composed of spatially orthonormal functions $\boldsymbol{\phi}_i(\boldsymbol{x})$ times a temporal orthonormal basis, such that:

$$\boldsymbol{u}(\boldsymbol{x},t) \approx \sum_{i=1}^{N_m} \psi_i(t)\sigma_i\boldsymbol{\phi}_i(\boldsymbol{x}), \tag{2.13}$$

with N_m being the number of POD modes, σ_i weighting the contribution of each mode to each instantaneous field and $\psi_i(t)$ referring to the time coefficients or temporal modes. Note that when N_m tends to infinity, equation 2.13 becomes an equality.

In the original work of Lumley (1967), POD aims to maximize the projection of the velocity fluctuations $\boldsymbol{u}(\boldsymbol{x},t)$ on their spatial basis $\boldsymbol{\phi}_i(\boldsymbol{x})$ that belong to a general Hilbert space $\mathcal{H}$:

$$\max_{\boldsymbol{\phi}_i(\boldsymbol{x})\epsilon\mathcal{H}} \left\langle \int_\Omega \big(\boldsymbol{u}(\boldsymbol{x},t), \boldsymbol{\phi}_i(\boldsymbol{x})\big) d\boldsymbol{x} \right\rangle, \tag{2.14}$$

where $\big(\cdot,\cdot\big)$ indicates a scalar product, Ω refers to the entire observation domain and $\langle\cdot,\cdot\rangle$ indicates an ensemble average. The choice of the basis functions $\boldsymbol{\phi}_i(\boldsymbol{x})$ affects the solution of equation 2.14. In the case of POD, this set of basis functions is chosen to be orthonormal i.e., fulfilling the orthogonality constraint and having unitary norm:

$$\left\langle \int_\Omega \big(\boldsymbol{\phi}_i(\boldsymbol{x}), \boldsymbol{\phi}_j(\boldsymbol{x})\big) d\boldsymbol{x} \right\rangle = \delta_{ij}, \tag{2.15}$$

where δ_{ij} is the Kronecker delta. Holmes *et al.* (2012) shows that the problem defined by equations 2.14-2.15 is equivalent to finding the solution of the Fredholm equation:

$$\left\langle \int_\Omega \big(\mathcal{R}(\boldsymbol{x},\boldsymbol{x}'), \boldsymbol{\phi}_i(\boldsymbol{x}')\big) d\boldsymbol{x}' \right\rangle = \sigma_i^2\boldsymbol{\phi}_i(\boldsymbol{x}), \tag{2.16}$$

where $\mathcal{R}(\boldsymbol{x},\boldsymbol{x}')$ is the two-point spatial correlation. Since any combination of the first r POD modes composes an orthonormal basis that minimizes the Frobenius norm of the flow field, POD provides flow field modes ordered by their energy contribution. Furthermore, Holmes *et al.* (2012) showed that eigenfunctions of equation 2.16 are Fourier modes if the two-point correlation is *homogeneous*, i.e. it is *translation invariant*.

An alternative approach to compute the POD modes of a set of instantaneous turbulent velocity fields is the *snapshot method* suggested by Sirovich (1987). This method proposes to project $\boldsymbol{u}(\boldsymbol{x},t)$ on their temporal basis $\psi_i(t)$, such that:

$$\max_{\psi_i(t)\epsilon\mathcal{H}} \left\langle \int_\mathcal{T} \big(\boldsymbol{u}(\boldsymbol{x},t), \psi_i(t)\big) dt \right\rangle, \tag{2.17}$$

where $\mathcal{T}$ refers to the entire temporal domain. Equation 2.16, together with the orthogonality condition for $\psi_i(t)$, defined as:

$$\left\langle \int_{\mathcal{T}} \big(\psi_i(t), \psi_j(t)\big) dt \right\rangle = \delta_{ij}, \tag{2.18}$$

compose a problem defined by the Fredholm equation such that:

$$\left\langle \int_{\Omega} \big(\mathcal{R}_t(t,t'), \psi_i(t')\big) dt' \right\rangle = \sigma_i^2 \psi_i(t), \tag{2.19}$$

where $\mathcal{R}_t(t,t')$ is the two-point temporal correlation. This approach is useful when the available temporal information is significantly smaller than the spatial one, which is often the case for experimental and numerical flow fields in which the available number of spatial points is larger than the number of available snapshots.

One of the first applications of POD to extract coherent structures in wall turbulence was shown by Moin & Moser (1989). Their results showed that coherent structures detected with POD contribute to a significant part of the turbulent kinetic energy. A similar approach was followed by Gurka *et al.* (2006) to study coherent structures based on vorticity measurements. Sanmiguel Vila *et al.* (2017) used POD to study the flow organization of turbulent boundary layers subjected to an adverse pressure gradient. More recently, Discetti *et al.* (2019) applied this approach to study experimentally the coherent structures present in a turbulent pipe flow. The advantage of studying coherent structures with this method resides in that the most energetic POD modes are ascribed to the largest structures present in turbulent flows (Wu 2014). This statement is based on the fact that the first POD mode contains most of the energy of coherent structures with characteristic lengths larger than the temporal or spatial domain in which POD is applied, thus becoming a useful tool to identify and represent coherent structures with a data-driven approach.

Apart from fluid mechanics, POD has been extensively applied in many research areas, such as data compression (Andrews *et al.* 1967), signal analysis (Algazi & Sakrison 1969), oceanography (Preisendorfer & Mobley 1988) or astrophysics (Soummer *et al.* 2012), to name a few. In the case of fluid mechanics, POD has not been only used to study coherent structures. For example, Mendez *et al.* (2017) proposed a method to remove the background from PIV images, which is especially useful for eliminating static reflections that occur when the PIV laser hits an object. Nonetheless, the classical application of the method is the data decomposition of flow fields, for which several examples can be found. POD has been used to apply a decomposition of the data in several flow topologies, such as the cylinder wake (Deane *et al.* 1991; Feng *et al.* 2011; Raiola *et al.* 2016), turbulent jets (Glauser *et al.* 1991; Berry *et al.* 2017; Mendez *et al.* 2019), flapping wings (Liang & Dong 2015; Troshin & Seifert 2018; Raiola *et al.* 2021), wall-mounted obstacles (Mallor *et al.* 2018, 2019; Güemes *et al.* 2019), or mixing layers (Delville *et al.* 1999; Citriniti & George 2000; Kaiser *et al.* 2014) among many others (Rowley & Dawson 2017). It is worth

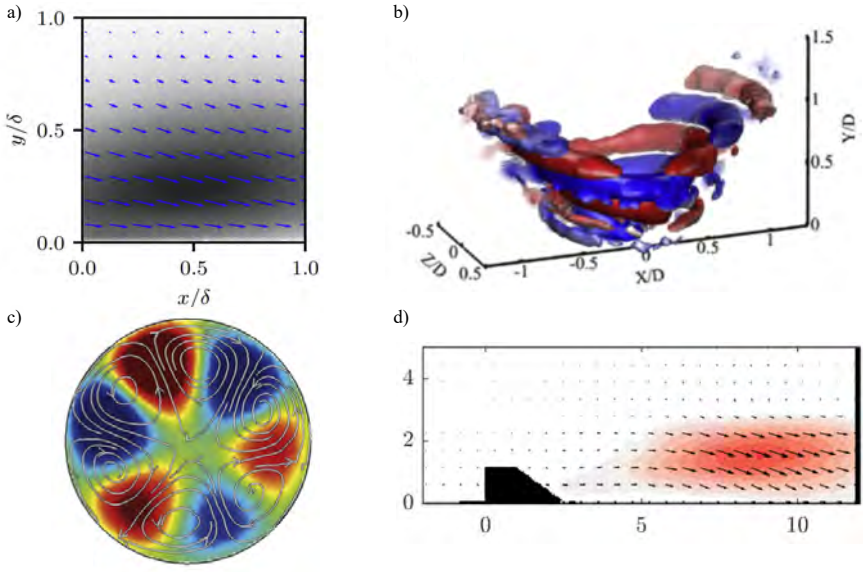

Figure 2.4: First POD mode $\boldsymbol{\phi}_1(\boldsymbol{x})$ for a) experimental turbulent boundary layer, b) experimental swirling jet, c) numerical turbulent pipe flow and d) obstacle wake immersed in a turbulent boundary layers. Data for panel a) comes from an experimental campaign carried out during the development of the present thesis in the wind-tunnel facility at Universidad Carlos III de Madrid. Panels b), c) and d) have been adapted from Ceglia *et al.* (2014), Hellström & Smits (2014), and Güemes *et al.* (2019) respectively.

noting that many of these works have also used POD to study the coherent structures present in turbulent shear flows. Figure 2.4 shows a collection of coherent structures identified for free- and wall-shear turbulent flows.

2.3.1. *Discrete formulation of proper orthogonal decomposition*

Commonly, the data acquired from experiments or numerical simulations come on discrete form, i.e., N_t samples are acquired on a spatial grid composed of N_p grid points, therefore requiring a discrete formulation. The three fluctuating velocity components of each instantaneous field are rearranged into a snapshot matrix U such as:

$$\mathbf{U} = \begin{bmatrix} u_{x_1}^{t_1} & \dots & u_{x_{N_p}}^{t_1} & v_{x_1}^{t_1} & \dots & v_{x_{N_p}}^{t_1} & w_{x_1}^{t_1} & \dots & w_{x_{N_p}}^{t_1} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ u_{x_1}^{t_{N_t}} & \dots & u_{x_{N_p}}^{t_{N_t}} & v_{x_1}^{t_{N_t}} & \dots & v_{x_{N_p}}^{t_{N_t}} & w_{x_1}^{t_{N_t}} & \dots & w_{x_{N_p}}^{t_{N_t}} \end{bmatrix}. \tag{2.20}$$

The ratio N_t/N_p conditions the choice between the original method proposed by Lumley (1967) and snapshot method of Sirovich (1987). When $N_t/N_p > 1$, POD temporal modes can be evaluated solving the eigenvalue problem of the temporal correlation matrix T as follows:

$$\mathbf{T} = \mathbf{U}\mathbf{U}^T = \mathbf{\Psi}\mathbf{\Lambda}\mathbf{\Psi}^T, \tag{2.21}$$

where $\mathbf{\Psi}$ is a matrix whose rows contain the temporal POD modes, while $\mathbf{\Lambda}$ is a diagonal matrix with elements $\lambda_i = \sigma_i^2$, which represent the variance content of each mode. Note that the singular value decomposition of the temporal correlation matrix $\mathbf{T}$ shown in equation 2.21 is the discrete version of the Fredholm equation presented in equation 2.19. Considering that equation 2.13 can be expressed in matrix form as:

$$\boldsymbol{U} = \mathbf{\Psi}\mathbf{\Sigma}\mathbf{\Phi}^T, \tag{2.22}$$

where Σ is a diagonal matrix containing the square root of the mode variances σ_i, the spatial POD modes $\boldsymbol{\phi}_i(\boldsymbol{x})$ are obtained by projecting the flow fields on the temporal POD modes computed with equation 2.21 as follows:

$$\mathbf{\Phi}^T = \mathbf{\Sigma^{-1}}\mathbf{\Psi}^T\mathbf{T}. \tag{2.23}$$

In the cases where $N_t/N_p < 1$, the discrete approach can be formulated to solve the spatial correlation matrix:

$$\mathbf{S} = \mathbf{U}^T\mathbf{U} = \mathbf{\Phi}\mathbf{\Lambda}\mathbf{\Phi}^T. \tag{2.24}$$

2.4. The imprint of coherent structures at the wall

The aforementioned effect that coherent structures have on the near-wall region can be extended to the wall if the linear behaviour of the mean-velocity profile in the viscous sublayer is taken into account. Early studies already showed the connection between coherent structures and wall-shear stresses (Brown & Thomas 1977), and more evidences were provided by Örlü & Schlatter (2011), which showed such relation in terms of energy spectra at the wall. Additionally, Sanmiguel Vila & Flores (2018) gave an insight of this wall imprint, showing the conditional average of the wall-normal fluctuating velocity with the pressure fluctuations at the flow for a turbulent channel flow.

Considering the imprint that coherent structures have on the wall, several methods that try to reconstruct wall-bounded flows from wall information can be found in the literature. One of those methods is based on the linear stochastic estimation (LSE), originally proposed by Adrian (1979), and exploits the self-similar linear scaling of the wall-attached eddies along the wall-normal direction. Using statistical information, LSE provides a linear approximation of an unknown quantity by solving the following minimization problem:

$$\left\langle \left| \boldsymbol{u}(\boldsymbol{x},t) - \boldsymbol{u}^{\dagger}(\boldsymbol{x},t) \right|^2 \right\rangle = \left\langle \left| \boldsymbol{u}(\boldsymbol{x},t) - \mathcal{A}(\boldsymbol{x},\boldsymbol{x'})\boldsymbol{u}(\boldsymbol{x'},t) \right|^2 \right\rangle, \tag{2.25}$$

where superindex † refers to predicted quantities. The operator $\mathcal{A}(\boldsymbol{x},\boldsymbol{x'})$ can be obtained by solving the following system of linear equations:

$$\big\langle \boldsymbol{u}(\boldsymbol{x},t)\boldsymbol{u}(\boldsymbol{x'},t)\big\rangle = \mathcal{A}(\boldsymbol{x},\boldsymbol{x'})\big\langle \boldsymbol{u}(\boldsymbol{x'},t)\boldsymbol{u}(\boldsymbol{x'},t)\big\rangle, \tag{2.26}$$

where $\big\langle \boldsymbol{u}(\boldsymbol{x},t)\boldsymbol{u}(\boldsymbol{x'},t)\big\rangle$ contains the cross-correlation information between the input and output variables, while $\big\langle \boldsymbol{u}(\boldsymbol{x'},t)\boldsymbol{u}(\boldsymbol{x'},t)\big\rangle$ is the autocorrelation information of the input variables. The linear operator $\mathcal{A}(\boldsymbol{x},\boldsymbol{x'})$ obtained from LSE can be used to reconstruct an instantaneous velocity as:

$$\boldsymbol{u}^{\dagger}(\boldsymbol{x},t) = \mathcal{A}(\boldsymbol{x},\boldsymbol{x'})\boldsymbol{u}(\boldsymbol{x'},t), \tag{2.27}$$

where $\boldsymbol{u}(\boldsymbol{x'},t)$ does not need to be necessarily a velocity quantity. In the beginning LSE was used to obtain conditional averages based on events in isotropic flows and was later applied to a variety of cases. The work of Encinar *et al.* (2018) and Encinar & Jiménez (2019) exploited the reconstructing capabilities of LSE by showing that wall information can be used to reconstruct turbulent flow fields at large Reynolds numbers.

Based on the findings of linear superposition and amplitude modulation (Hutchins & Marusic 2007; Mathis *et al.* 2009), Marusic and coworkers presented a method to estimate near-wall velocity fluctuations from the logarithmic layer (Marusic *et al.* 2010; Mathis *et al.* 2011). This model is based on the fact that current measuring technologies do not have the spatial and temporal resolution to capture the velocity fluctuations in the near-wall region at the Reynolds numbers of typical industrial applications. Assuming the linearity of the mean-velocity profile at the viscous sublayer, this model was extended by Mathis *et al.* (2013) to predict the shear-stress fluctuations at the wall. Unfortunately, the mathematical description of the model is not invertible i.e., the coherent structures at the logarithmic layer cannot be predicted from the wall. A refined version of the predictive model of Marusic *et al.* (2010) was provided Baars *et al.* (2016), by adding to the model a spectral LSE. This technique, originally proposed by Tinney *et al.* (2006), translates LSE to the frequency space to remove noise from coherent sources. Another option is to use the linearized version of the Navier-Stokes equations as a predictive model. This technique was applied to moderate-Reynolds-number flows by Illingworth *et al.* (2018), with successful results when reconstructing the large-scale coherent structures present in the flow.

Another linear method commonly used to reconstruct turbulent flow fields from a reduced set of measurements is the extended proper orthogonal decomposition (EPOD), originally proposed by Maurel *et al.* (2001). EPOD is a statistical tool that allows to draw the correlation between the flow features extracted from two different datasets. For instance, given a signal $\boldsymbol{w}(\boldsymbol{x},t)$ containing the pressure and shear-stress fluctuations at the wall, it is possible to calculate its extended POD modes as its projection onto the temporal basis defined for $\boldsymbol{u}(\boldsymbol{x},t)$ as:

$$\sigma_i^e \boldsymbol{\phi}_i^e(\boldsymbol{x}) = \left\langle \int_{\Omega} \big(\psi_i(t), \boldsymbol{w}(\boldsymbol{x},t)\big) d\boldsymbol{x} \right\rangle, \tag{2.28}$$

where superindex e stands for *extended*. The strength of equation 2.28 is that allows to extract the correlation between synchronized measurements through the extended POD modes, being possible to quantify the relationship between the coherent structures populating each one of them.

Rearranging the wall information contained in $\boldsymbol{w}(\boldsymbol{x}, t)$ in a matrix form:

$$\mathbf{W} = \begin{bmatrix} \tau_{w_{x_{x_1}}}^{t_1} & \dots & \tau_{w_{x_{x_{N_p}}}}^{t_1} & \tau_{w_{z_{x_1}}}^{t_1} & \dots & \tau_{w_{z_{x_{N_p}}}}^{t_1} & p_{w_{x_1}}^{t_1} & \dots & p_{w_{N_p}}^{t_1} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \tau_{w_{x_{x_1}}}^{t_{N_t}} & \dots & \tau_{w_{x_{x_{N_p}}}}^{t_{N_t}} & \tau_{w_{z_{x_1}}}^{t_{N_t}} & \dots & \tau_{w_{z_{x_{N_p}}}}^{t_{N_t}} & p_{w_{x_1}}^{t_{N_t}} & \dots & p_{w_{N_p}}^{t_{N_t}} \end{bmatrix}, \quad (2.29)$$

and considering that $\mathbf{W}$ can be decomposed as:

$$\boldsymbol{W} = \boldsymbol{\Psi}_{\boldsymbol{w}} \boldsymbol{\Sigma}_{\boldsymbol{w}} \boldsymbol{\Phi}_{\boldsymbol{w}}{}^{T}, \quad (2.30)$$

it is possible to express equation 2.28 in discrete form analysis as:

$$\boldsymbol{\Sigma}^{e} \boldsymbol{\Phi}^{eT} = \boldsymbol{\Psi} \boldsymbol{W} = \boldsymbol{\Psi} \boldsymbol{\Psi}_{\boldsymbol{w}} \boldsymbol{\Sigma}_{\boldsymbol{w}} \boldsymbol{\Phi}_{\boldsymbol{w}}{}^{T} = \boldsymbol{\Xi} \boldsymbol{\Sigma}_{\boldsymbol{w}} \boldsymbol{\Phi}_{\boldsymbol{w}}{}^{T}, \quad (2.31)$$

where the matrix $\boldsymbol{\Xi}$ contains the correlation between the temporal modes.

Assuming that the POD for both quantities has reached statistical convergence, it is possible to estimate any of the quantities through a linear projection on the extended POD modes as:

$$\boldsymbol{U}^{\dagger} = \boldsymbol{\Xi} \boldsymbol{\Psi}_{\boldsymbol{w}} \boldsymbol{\Sigma} \boldsymbol{\Phi}^{T}. \quad (2.32)$$

Discetti *et al.* (2018) showed that the energy-optimality of POD modes mentioned in § 2.3 can then be retained by employing statistical filters on the temporal correlation matrix.

EPOD has been shown to provide accurate predictions in several types of flows, such as the far-field predictions of turbulent jets carried out by Tinney *et al.* (2008). In the case of wall turbulence, there are multiple examples. For instance, Discetti *et al.* (2018) reconstructed wall-normal velocity fields in a channel flow by using a limited amount of velocity probes, equally distributed along the wall-normal direction. Their results showed that EPOD could be used to reconstruct large-domain fields in a turbulent channel flow. Similar results were obtained for the case of a turbulent pipe flow by Discetti *et al.* (2019).

One of the most important properties of EPOD is its connection to LSE. Borée (2003) proved mathematically that LSE and EPOD are equivalent methods. Following Holmes *et al.* (2012), this relation can be resumed in two main properties. First, suitable averages of LSE events produce POD modes and second, all LSE events are linear combinations of POD modes.

Other examples of predictive models can be found in the literature, like the sequential approach proposed by Suzuki & Hasegawa (2017) or the transfer functions applied by Sasaki *et al.* (2019). Linear methods exploit the self-similarity of the wall-attached coherent structures when doing the predictions. However, they have shown problems when predicting velocity fluctuations at a specific wall distance y^+ whose characteristic length scales are smaller than this

distance, since these small fluctuations do not have an imprint on the wall. It is to be expected that nonlinear methods can significantly outperform linear ones by also modelling the nonlinear part of the relation between coherent structures and their imprint on the wall.

Chapter 3

Deep neural networks as a tool for turbulence

Chapter 2 has shown the capability of POD to provide a compact spatial representation of the coherent structures populating turbulent flows. The capabilities of state-of-art methods can cover only the linear part of the relation between coherent structures and their imprint at the wall. As it will be shown later in this chapter, neural networks are chosen as estimation methods due to their capability to establish approximations for nonlinear relations, thus paving the way toward the estimation of the nonlinear part of the relation between wall quantities and velocity fields far from the wall.

This chapter presents a simple introduction to deep neural networks in § 3.1, while § 3.2 provides a review of their most recent applications in fluid mechanics. Section 3.3 presents the neural-network architecture developed in the present thesis to reconstruct turbulent flow fields from wall measurements, exploiting the nonlinear capabilities inherent to neural networks and the compact representations of turbulent flow fields provided by POD. Finally, § 3.4 shows how neural networks can be applied to recover wall information, which is especially useful when the amount of wall data is limited, as it is often the case in real-life applications.

3.1. Introduction to deep neural networks

Although the terms artificial intelligence (AI), machine learning and deep learning are used interchangeably by the general public, they do not refer to the same concepts. The term AI refers to a broad set of techniques that are used for computers to mimic the human behaviour. Machine learning is a set of techniques used within AI that apply algorithms to help computers learn models from sets of data, allowing them to solve different regression or classification problems. Other research areas, such as robotics or reinforcement learning, are also under the umbrella of AI, at the same level of machine learning. Finally, deep learning is a specific branch of machine learning based on the use of neural networks. Decision trees, support vector machines, Bayesian networks or genetic algorithms are other types of machine learning techniques at the same level of deep learning.

The *perceptron*, introduced by Rosenblatt (1958), is the basic unit composing a neural network. Devised with the aim of mimicking the behaviour of human

neurons, it can be formulated as a mapping function f that transforms a real-valued input vector $\boldsymbol{x}$ to a single output y. This transformation is linear, and it can be expressed mathematically as:

$$y = f(\boldsymbol{x}) = \sum_{i=1}^{N_i} w_i x_i + b, \tag{3.1}$$

where w_i are the coefficients of the mapping function, the so-called weights, b is a bias term in the mapping function independent of the inputs, and N_i is the number of inputs contained by $\boldsymbol{x}$. Nonlinear effects are incorporated by adding a second mapping function g, commonly known as *activation function*, that applies a threshold to y in order to obtain the final output z, such as:

$$z = \begin{cases} 1, & \text{if } y \geq \theta \\ 0, & \text{otherwise} \end{cases}, \tag{3.2}$$

where θ is the threshold value. The definition of this activation function is a research area in itself and numerous examples can be found in the literature, such as tangent hyperbolic, sigmoid, or rectified linear unit activation functions among others. The activation function used in equation 3.2 is known as Heaviside step function.

Perceptrons can be combined within a layer to produce an output vector $\boldsymbol{z}$ of j components such that:

$$z_j = \begin{cases} 1, & \text{if } \sum_{i=1}^{N_i} w_{ji} x_i + b_j \geq \theta \\ 0, & \text{otherwise} \end{cases}, \tag{3.3}$$

and several layers can be stacked together, thus using hidden layers, i.e., layers that see neither the input nor the output. Both single-layer and multi-layer perceptrons are fully-connected neural networks, and they are classified as feedforward networks since their weights do not have cyclical connections. The use of the term *deep* is not associated with a specific number of layers or neurons per layer, but the more there are, the deeper the network will be.

Despite having been introduced more than 60 years ago, the use of neural networks has not been practical until the last decade, when the advent of powerful computational resources in combination with the backpropagation algorithm for training the weights made them computationally affordable. The current concept of backpropagation was introduced by Rumelhart *et al.* (1986). Given a loss function $\mathcal{L}$ to evaluate the error of the predicted output with respect to its target value, the weights of the neural network are updated using the gradient of the error with respect to the weights as:

$$\Delta w = -\eta \frac{\partial \mathcal{L}}{\partial w_{ji}}, \tag{3.4}$$

where η is the learning-rate parameter.

Fully-connected layers represent just one type of feedforward neural networks. For instance, convolutional neural networks (CNNs) were developed

mainly during the 80s and the 90s (Fukushima 1980, 1988; LeCun *et al.* 1989, 1998) and have been extensively used for image-related problems. They use a convolution operator to extract feature maps $F(i,j)$ from an input $I(i,j)$ as:

$$F(i,j) = \sum_{m}^{M} \sum_{n}^{N} I(i-m, j-n) K(m,n), \tag{3.5}$$

where K is the kernel operator of size $M \times N$. Note that the input can be three dimensional, such as RGB images, and that the kernel may also have a third dimension to extract more than one feature map. Their main advantages are the sparse interactions, parameter sharing and equivariant representations (Goodfellow *et al.* 2016). While in fully-connected networks every neuron in one layer receives information from all neurons in the previous one, CNNs use kernels much smaller than the inputs, connecting only those inputs that are spatially close, thus leading to sparse interactions. As the same kernel is applied to the entire input, the number of weights is reduced substantially with respect to fully-connected networks, which is known as parameter sharing. Finally, the equivariance means that, as the same kernel is applied to the entire input, if the input changes the feature maps change in the same manner.

Feedforward is not the only type of neural networks. For example, recurrent neural networks are one of the most famous types, which are characterized by using a sequential processing of the input data more similar to the way of working of the human brain. However, as neither recurrent neural network nor other types of neural networks have been used in the development of this thesis, no further details will be provided. The interested reader is referred to Goodfellow *et al.* (2016) for a comprehensive review of these concepts.

3.2. Application of deep neural networks in fluid mechanics

Thanks to new computational resources such as graphic processing units (GPUs), in recent years numerous, studies have flourished that apply deep neural networks in different fields of research, such as socioeconocmics (Jean *et al.* 2016), medicine (De Fauw *et al.* 2018; Zeleznik *et al.* 2021), wildlife ecology (Norouzzadeh *et al.* 2018), climate (Ham *et al.* 2019; Waldmann & Griffith 2019), sustainability (Vinuesa *et al.* 2020), or physics (Udrescu & Tegmark 2020; Kwon *et al.* 2020). As might be expected, fluid mechanics has not been unrelated to this renewed interest, as reviewed in the works of Kutz (2017) and Brunton *et al.* (2020). In the rest of this section, a review will be made of the main areas of fluid mechanics where deep neural networks have been applied, as well as the most notable works.

The first area is the application of neural networks for turbulence modelling, which can be grouped into Reynolds-averaged Navier-Stokes (RANS) and large-eddy simulation (LES) models. In the case of RANS models, the work of Ling *et al.* (2016) is one of the most relevant. They used fully-connected neural networks to compute the Reynolds stress anisotropy tensor from the mean strain rate and rotation rate tensors. Moreover, they incorporated a custom

layer based on representation theory to ensure the rotational invariance of the predicted anisotropy tensor (Pope 1975). Evaluated on duct and wavy flows, results showed that their neural network was able to produce better predictions than linear and nonlinear eddy viscosity models, although DNS predictions were not perfectly reproduced. Another possible approach is the one proposed by Thuerey *et al.* (2020), whose model used convolutional networks to predict the velocity and pressure fields around an airfoil from freestream velocity. In this case, RANS simulations were used as variables to be predicted, so its neural network replaced the numerical simulation.

LES models produce more accurate results than RANS, but their modelling process is more complex and expensive. They use different methods to model the smallest length scales in the flow that, unlike in DNS simulations, are not resolved. Deep neural networks can be used to substitute or complement these models. For example, Maulik & San (2017) and Maulik *et al.* (2019) used a fully-connected neural network to predict the local closure term in the vorticity equation. They tested their model for a two-dimensional decaying turbulent flow, with successful results when modelling the dissipation of turbulent kinetic energy at length scales smaller than the LES grid. A different example was provided by Beck *et al.* (2019), where the residuals of the continuity, momentum and energy equations were predicted with convolutional networks for three-dimensional isotropic turbulence.

More recently, Park & Choi (2021) have used a fully-connected neural network to predict the Reynolds shear stress in a turbulent channel flow of $Re_\tau \approx 180$. They have analysed the benefits and drawbacks of different inputs, concluding that the best overall results are obtained using the velocity gradient rate. Additionally, this model has been proved to still provide accurate results when applied to a larger Re_τ flow, not seen during the training process. This work is similar to the one of Gamahara & Hattori (2017), where the prediction of Reynolds stresses was investigated by using different inputs, although Park & Choi (2021) have better results when using the model for advancing the simulation.

LES modelling with neural networks can be applied to more complex flows. For instance, Lapeyre *et al.* (2019) recovered the flame surface density of a premixed turbulent flame from temperature measurements. Moreover, they used convolutional networks instead of fully-connected ones, thus exploiting the spatial information of the flow. In summary, neural networks can be used for turbulence modelling in a wide range of ways, where the choice of input and output variables is more influenced by the user than by the equations themselves. Despite the successful results collected in these studies, there is still a long way to go for neural networks to be able to improve these types of simulations to the point of reaching DNS precision. Furthermore, it must be considered that all these models require high-resolution data obtained by DNS simulations. Because of this, in the coming years, it will be of utmost

importance to train networks able to provide accurate predictions with data not seen during training.

The second major area where neural networks are employed in fluid mechanics is the resolution enhancement of flow fields, being the convolutional networks the dominant ones for this type of application. One of the first examples found in the literature is the work of Fukami *et al.* (2019), where convolutional networks were used to enhance the resolution of two-dimensional decaying isotropic turbulence. A similar approach was followed by Liu *et al.* (2020), who targeted the resolution enhancement of turbulent two-dimensional decaying isotropic turbulence and channel flows. Moreover, they presented a version of their architecture where temporal information was taken into account. Specifically, the previous and following low-resolution fields were used as input together with the one corresponding to the high-resolution field to be predicted. They reported better results for the network including temporal information, being able to recover more small-scale details. This type of networks is also able to recover temporal information during the resolution enhancement, as shown by Fukami *et al.* (2021). Using two low-resolution fields, located at the beginning and end of a 9-snapshot well-resolved temporal sequence, they proposed a methodology to recover high-resolution sequences of flow fields. They evaluated their methodology at homogeneous isotropic and wall-bounded flows, obtaining accurate resolution enhancement for both cases.

However, CNNs have not been the only ones used for this task, nor the most successful. Generative adversarial networks (GANs) were introduced by Goodfellow *et al.* (2014) and have been very successful at this task. The main concept behind this type of networks relies on game theory, since there are actually two networks working collaboratively to produce better results. The first network is known as the generator, and it is in charge of producing an artificial output that mimics the reality. The second network, known as the discriminator, has the task of distinguishing between reality and artificial outputs. This type of networks has been successfully applied to resolution enhancement of pictures, as shown in the works of Ledig *et al.* (2017) and Wang *et al.* (2018). This ability to retrieve small-scale details has not been overlooked in fluid mechanics. For instance, Deng *et al.* (2019) used GANs to recover high-resolution flow fields in the wake behind one or two side-by-side cylinders. Another example was provided by Werhahn *et al.* (2019), who employed them to recover high-resolution volumetric density of smoke simulations. The work by Xie *et al.* (2018) is also interesting, since they embedded temporal information into a second discriminator's loss function with the purpose of producing predictions with higher temporal coherence. Special mention deserves the study carried out by Kim *et al.* (2021). They employed an unsupervised version of GANs, known as cycle-GAN (Zhu *et al.* 2017), to enhance the resolution of homogeneous turbulence and channel flows when paired high- and low-resolution data is not available.

Generally, all these previous works train their neural networks by minimizing a loss function between the target and the predicted variables. These functions are usually simple, such as the mean-squared error, but they can become more complex, allowing to embed physical information into the model. This was the case of Raissi *et al.* (2019) work, which included the Navier-Stokes equations in the loss function. They tested their model in the cylinder wake, being able to learn two constants affecting the convective and diffusive terms in the Navier-Stokes equations. This work was extended to target more complex flows, such as the blood flow in the carotid artery with an aneurysm (Raissi *et al.* 2020). Nonetheless, calculating the residuals of the Navier-Stokes equations requires time-resolved data, which, depending on the application, can be challenging to acquire.

Another of the most notable applications of neural networks is the dimensionality reduction. As stated in Chapter 2, POD has been one of the most successful techniques for this task since its introduction in fluid mechanics (Lumley 1967). The relationship between POD and neural networks is very close, since it was shown by Baldi & Hornik (1989) that fully-connected neural networks with a single hidden layer and linear activation function are equivalent to POD. Moreover, Hornik *et al.* (1989) proved that multi-layer neural networks are universal approximators for any continuous function. Milano & Koumoutsakos (2002) were pioneers in the application of neural networks for low-order modelling. They compared the performance of POD and neural networks when reconstructing the near-wall region of a turbulent channel flow. Their results proved that neural networks are able to capture the nonlinearities present in turbulent flows, thus providing better dimensionality reduction capabilities than POD.

The emergence of autoencoders (DeMers & Cottrell 1993) brought a renewed interest in the application of neural networks for dimensionality reduction. An autoencoder is composed of two mapping functions, the encoder and the decoder. The encoder maps the input data to a latent representation whose dimensionality is much lower than the input. Then the decoder maps back that latent representation into the original input. The objective of these networks is to reduce the dimension of the data as much as possible without losing information. Deep autoencoders were employed by Hinton & Salakhutdinov (2006) to map high-dimensional data into a low-order representation, retaining more information than POD. Despite having better encoding, autoencoders were not able to provide a modal decomposition where the modes could be inspected as in POD. This issue was tackled by Murata *et al.* (2020) by proposing an autoencoder with several branches in the decoder. Each of the branches was assigned to a single mode, thus providing a visual representation. Later, the output was reconstructed as a direct combination of these modes. They evaluated their model on the cylinder's wake, whose first two POD modes contain 99% of the flow energy, obtaining accurate results. The other difference between POD and the autoencoder proposed by Murata *et al.* (2020) is that the latter did not

impose a hierarchical organization of the modes based on their energy content. A convolutional-based hierarchical autoencoder was proposed by Fukami *et al.* (2020), showing that their methodology was able to capture in its first mode the high-energy dominant structures represented in the first POD mode. Due to its undeniable potential, several examples of the use of autoencoders can be found in the literature. For instance, Fukami *et al.* (2019*b*) proposed to employ autoencoders as a synthetic inflow turbulence generator for turbulent channel flow of $Re_\tau = 180$. Generating physical inflow is of utmost importance, because otherwise the fluctuations dissipate rapidly. Another application was shown by Glaws *et al.* (2020), who used convolutional autoencoders to compress DNS data after its generation. This allows to reduce the computational resources required to store the large amount of data produced by DNS.

Last but not least, we must review the application of neural networks to predict the instantaneous state of a flow from certain variables that can be acquired with wall sensors, whether they are pressure transducers, IR thermography or hot-film sensors. For example, Erichson *et al.* (2020) applied shallow fully-connected neural networks to predict the flow fields in a cylinder wake from measurements at cylinder surface. In the work of Kim & Lee (2020) CNNs were used to predict the heat transfer at the wall using the two components of the wall-shear-stress and the pressure fluctuations. They tested their methodology at a turbulent channel flow of $Re_\tau = 180$ with successful results, and also validated its application for larger Re_τ flows, thus claiming that the relationship between the heat flux and the inputs is almost independent of the frictional Reynolds number within the tested range. Finally, Guastoni *et al.* (2020*a*) presented preliminary results for the use of CNNs to predict wall-parallel velocity fields from wall-shear measurements. This work is intimately related to the main research line of this thesis, so more details will be given in the following section.

These are the main areas in which neural networks have been applied in fluid mechanics. However, there are more works of great relevance that are not grouped in these areas. For instance, the lift coefficient and the displacement of a freely vibrating cylinder was estimated in Raissi *et al.* (2019) by using a fully-connected network fed only with flow measurements from the surrounding. Moreover, they also showed the capability of neural networks to estimate the concentration of a passive scalar in the flow. Another example can be found in Hack & Zaki (2016), where neural networks were used to classify streaks in a pressure-gradient boundary layer between those that were going to transition into turbulence and those that returned to a laminar state. Despite not using neural networks, it is essential to highlight the methodology presented by the series of works by Prof. J. Jiménez (Jiménez 2018, 2020*a*,*b*,*c*). In these works, machine learning is used to extract a physics-based theory for describing decaying two-dimensional isotropic turbulence, finding than dipoles are as important as classical individual vortex cores for this type of flows. Neural networks could be used with a similar approach, as demonstrated by Iten *et al.* (2020).

These authors used neural networks to extract physical laws in agreement with classical theories, such as the conservation of angular momentum, the damped pendulum or the model for a heliocentric solar system.

As a final note, it is important to remark that most of the studies mentioned before are very recent, many of them having been published in the last 3 years. This shows that the application of neural networks in fluid mechanics has aroused great interest in the scientific community due to its potential.

3.3. POD-based convolutional networks for flow reconstruction

This section presents the POD-based CNNs developed in the present thesis for the prediction of turbulent flow fields from wall measurements. As shown below, there is an important conceptual difference between the networks used in Paper 1 and the network used in the rest of the papers.

Figure 3.1a) shows a schematic representation of the POD-based neural employed in Paper 1. This type of network is tailored for high-Re_τ data where only the coherent structures with high-energy content are of interest. For that purpose, the network presented in figure 3.1a) is used to generate the predictions. The network is fed with two-dimensional fields of wall quantities. In the case of Paper 1, these quantities were streamwise and spanwise wall-shear stresses. The output of the network is a single ψ_i coefficient. For Paper 1, ten different networks were trained to predict the first 10 POD coefficients. The instantaneous coherent structures populating the flow field are obtained by projecting these predicted coefficients into their POD spatial basis functions. The architecture of the network stacks together 5 blocks composed of a convolutional layer of 64 filters with size 4×4, followed by a rectified linear unit (ReLU, see Nair & Hinton 2010) activation function and a max-pooling layer that reduces the data dimension by a factor of 2. Finally, two fully-connected layers of output size 64 and 1 are added to the model, using a ReLU activation function for the first one. Using the Adam algorithm (Kingma & Ba 2014), the weights of the network are optimized by minimizing the loss function:

$$\mathcal{L} = \frac{1}{N_r} \sum_{i=1}^{N_r} |\psi_i^\dagger - \psi_i|^2, \tag{3.6}$$

where N_r refers to the number of POD modes taken into account for the reconstruction. Acronym CNN-POD is used to refer to this network.

In the case of Papers 2 and 3, the fully-connected layers at the end of the network are removed, thus leading to a fully-convolutional network, which is commonly used in tasks where the input and output data share structural similarities (Long *et al.* 2015). This method divides the turbulent flow fields into N_s two-dimensional subdomains of $N_p \times N_p$ grid points, which are decomposed into POD modes. The number of subdomains is chosen based on Re_τ characterizing the flow, with the purpose of ensuring that around 90% of the flow kinetic energy is contained within $\mathcal{O}(10^2)$ POD modes that can be

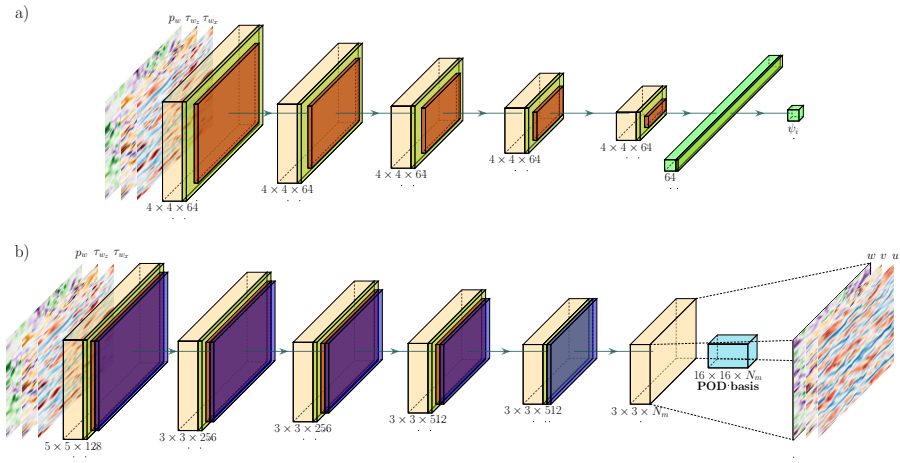

Figure 3.1: Schematic view of the FCN-POD architectures for a) fully-connected layer, and b) fully-convolutional networks. The colour coding for each layer is: 2D-convolution ( ), ReLU-activation ( ), fully-connected ( ), batch-normalization ( ), and max-pooling ( ) layers. The kernel size and the number of filters are shown at the bottom of the convolution layers.

translated to convolutional filters. Figure 3.1 shows the architecture in charge of reconstructing a three-dimensional tensor of size $N_x \times N_z \times N_r$ containing the truncated number of POD coefficients N_r at each subdomain, where N_x and N_z refers to the number of subdomains in which the turbulent flow fields are decomposed in the streamwise and spanwise directions respectively. Note that the total number of subdomains $N_x \times N_z$ is also referred previously as N_s. This three-dimensional tensor is converted later into the flow field by projecting each POD coefficient in its corresponding basis. With respect to the input information, Paper 2 and 3 also use the instantaneous pressure fluctuations together with the wall-shear-stress components. Similar to the first version of the architecture, the network weights are optimized with Adam algorithm (Kingma & Ba 2014). For this network, the loss function is defined as:

$$\mathcal{L} = \frac{1}{N_r} \sum_{i=1}^{N_r} \sum_{s_i=1}^{N_x} \sum_{s_j=1}^{N_z} |\psi^{\dagger}_{i_{s_i,s_j}} - \psi_{i_{s_i,s_j}}|^2, \tag{3.7}$$

where subindeces s_i and s_j refer to the spatial location in the streamwise and spanwise directions of each subdomain. This network is referred with FCN-POD acronym.

Figure 3.2 shows an example of a predicted turbulent flow field at different wall distances using the FCN-POD approach. As it can be observed, when moving away from the wall the network starts to lose the small-scale details,

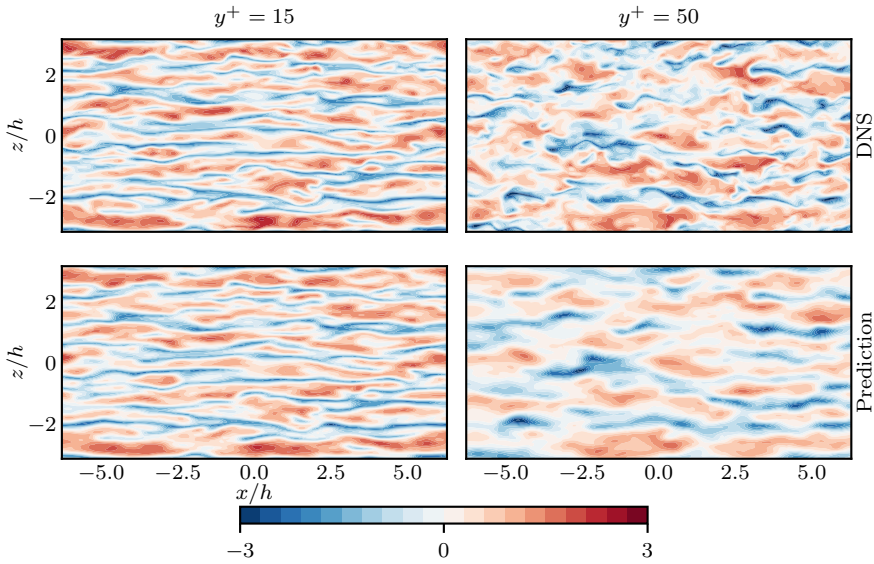

Figure 3.2: Countour map for the streamwise velocity fluctuation fields scaled with their corresponding standard deviation. Top panels denote DNS, while bottom ones refer to FCN-POD predictions. Left columns refers to $y^+ = 15$, while rigth columns stand for $y^+ = 50$.

capturing only the large-scale fluctuations of the velocity fields. In any case, this is still a useful reconstruction, since their large-scale data is ascribed to the elongated u-streaks mentioned in Chapter 2. Moreover, Paper 2 compared the performance of FCN-POD with the fully-convolutional-network (FCN) implementation of Guastoni *et al.* (2020*a*). This network applies several convolutional layers to the same inputs in order to obtain directly the velocity fields, i.e., without POD compression. It was shown that FCN-POD is able to provide better predictions further from the wall than FCN approach thanks to partial encoding of physical information in the POD spatial basis functions. While Paper 2 and Paper 3 have used a database without an overlap region, the work of Paper 1 has been carried out with a turbulent channel flow much closer to this condition, showing that this approach is still valid when targeting the large coherent structures present in the flow. This work also evaluated the effect of wall resolution, showing that in the tested range there was a slight performance decrease when reducing the input resolution.

3.4. Super-resolution GANs for wall and flow measurements

While the reconstruction of turbulent flow fields from low-resolution wall measurements can be achieved by the neural network presented in § 3.3, Paper

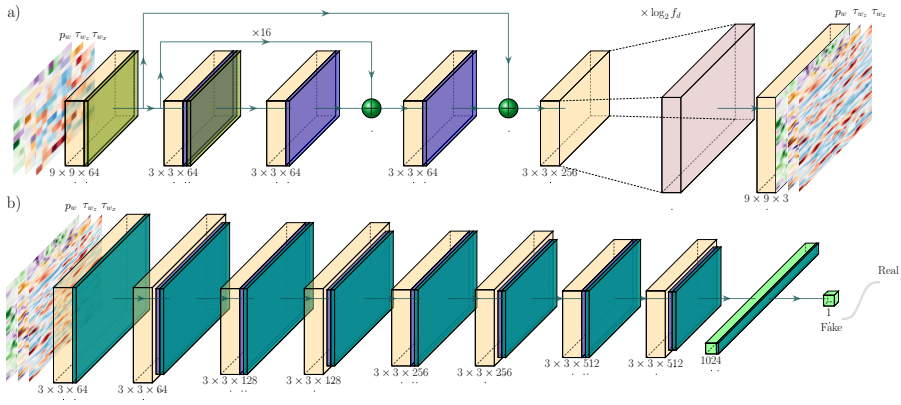

Figure 3.3: Schematic view of the SRGAN architectures for a) generator G, and b) discriminator D networks. The colour coding for each layer is: 2D-convolution ( ), parametric-ReLU-activation ( ), batch-normalization ( ), sub-pix-convolution ( ), Leaky-ReLU-activation ( ), and fully-connected ( ) layers. The kernel size and the number of filters are shown at the bottom of the convolution layers.

1 showed a performance decrease when reducing the wall resolution. This section presents a GAN-based methodology, introduced in Paper 3, to deal with low-resolution wall measurements, both to increase the resolution of these measurements and to reconstruct turbulent flow fields. Similarly to POD-based networks, this architecture is fed with fluctuating streamwise and spanwise wall-shear-stress and pressure fields. However, the output is not a POD temporal projection, but the actual fields in their physical dimensions, for both wall and flow reconstructions.

The choice of GANs is based on their proven effectiveness for enhancing low-resolution images. This type of GANs, known as super-resolution GANs (SRGANs), consists of two networks: a generator (G) and a discriminator (D). Applying this methodology to wall measurements, G is in charge of producing an artificial high-resolution wall or flow field $\tilde{H}_R$ from a conditioned input, i.e., the low-resolution wall field L_R, whereas the purpose of D is to distinguish between the target high-resolution fields H_R and their artificial counterparts. As it has been seen in § 3.2, SRGANs have been already used in different fluid-mechanics-related problems. The SRGAN architecture presented in figure 3.3 is based on the Ledig *et al.* (2017) work. Figure 3.3a) presents the architecture of G, which is characterized by two main features: the use of 16 residual networks (He *et al.* 2016) in the core and the use of sub-pixel convolutional layers (Shi *et al.* 2016) to increase the resolution at the end of the network. Note that the number of sub-pixel convolutional layers required to match high-resolution fields

is equal to $\log_2(f_d)$, being f_d the downsampling ratio between the high- and low-resolution wall fields. The residual blocks are composed of two convolutional layers of 64 filters of size 3×3, each of them followed by a batch-normalization layer (Ioffe & Szegedy 2015) and separated by a parametric ReLU (He *et al.* 2015) activation function.

For the D architecture shown in figure 3.3b), the guidelines of Radford *et al.* (2015) have been followed. Convolutional layers are used together with batch-normalization layers and Leaky ReLU (Maas *et al.* 2013) activation functions to extract features from the wall fields. The number of filters in the convolutional layer is doubled every two layers, while the image size is reduced by a factor of two by striding the kernel. Finally, two fully-connected layers with output size of 1024 and 1 are used, applying a sigmoid activation function to the architecture output in order to quantify the probability of the input high-resolution wall field of being real (0) or fake (1).

The discriminator loss function $\mathcal{L}_D$ is computed based on the binary cross-entropy for both real and fake images as:

$$\mathcal{L}_D = -\mathbb{E}[\log D(H_R)] - \mathbb{E}[\log(1 - D(G(L_R)))]. \tag{3.8}$$

In the case of G, its loss function $\mathcal{L}_G$ has been computed following the perceptual loss defined by Ledig *et al.* (2017), where the pixel-based mean squared error is combined with an adversarial contribution coming from $\mathcal{L}_D$. Mathematically, the error is formulated as:

$$\mathcal{L}_G = \frac{1}{N_x N_z} \sum_{i=1}^{N_x} \sum_{j=1}^{N_z} |G(L_R)_{i,j} - H_{Ri,j}|^2 - \lambda \mathcal{L}_D, \tag{3.9}$$

where λ is a scalar factor weighting the contribution of the adversarial loss. Figure 3.4 shows the resolution enhancement of an instantaneous streamwise wall-shear-stress fields of a turbulent channel flow. While figure 3.2 has shown that moving away from the wall reduce the efficiency of predicting turbulent flow fields, these results show that reducing the amount of information provided to the network as input also affects negatively to the resolution enhancement. Nonetheless, the results show the ability of SRGAN to recover the largest traces of coherent structures on the wall, even in cases where the amount of information is significantly reduced.

Figure 3.5 shows the reconstruction of instantaneous streamwise velocity fields of a turbulent channel flow from wall measurements with different resolution. Similar to FCN-POD results of Paper 2, when moving away from the wall the quality of flow predictions is decreased. However, SRGAN network is able to recover more small-scale details than FCN-POD and FCN for the same wall distance. Regarding the effect of wall resolution in the flow reconstruction, Paper 3 tackled more challenging cases than Paper 1, showing that for extreme cases there is a significant accuracy decrease. Nonetheless, Paper 3 showed that even in the most challenging cases of wall distance, wall resolution, or

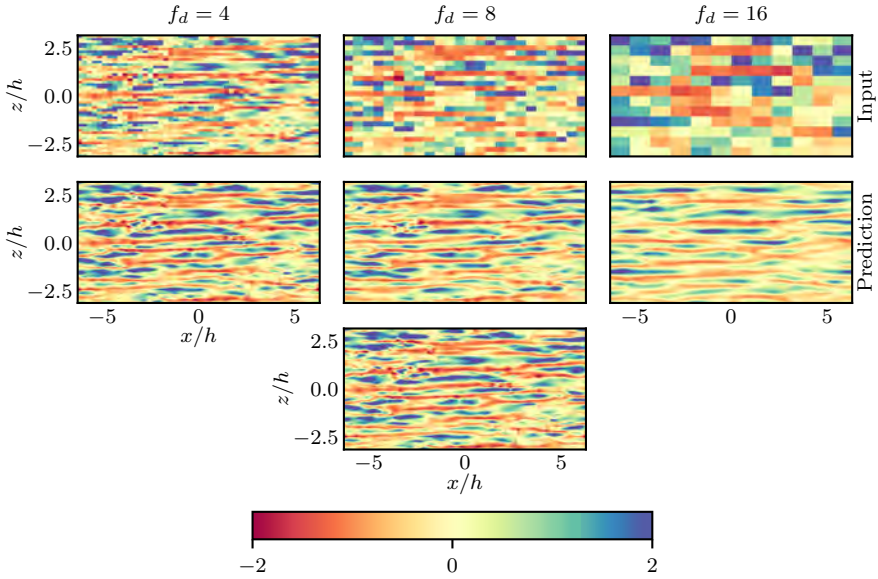

Figure 3.4: Comparison of the streamwise wall-shear-stress fluctuating fields at $Re_\tau = 180$, scaled with their corresponding standard deviation. Reference DNS is reported at bottom panel, while top- and middle-row panels refer to low-resolution inputs and SRGAN predictions respectively. Each column contains a different case in terms of ratio f_d between low- and high-resolution data, covering $f_d = 4$ (left), $f_d = 8$ (center), and $f_d = 16$ (right).

a combination of both, SRGAN is able to capture the large-scale structures present in the flow.

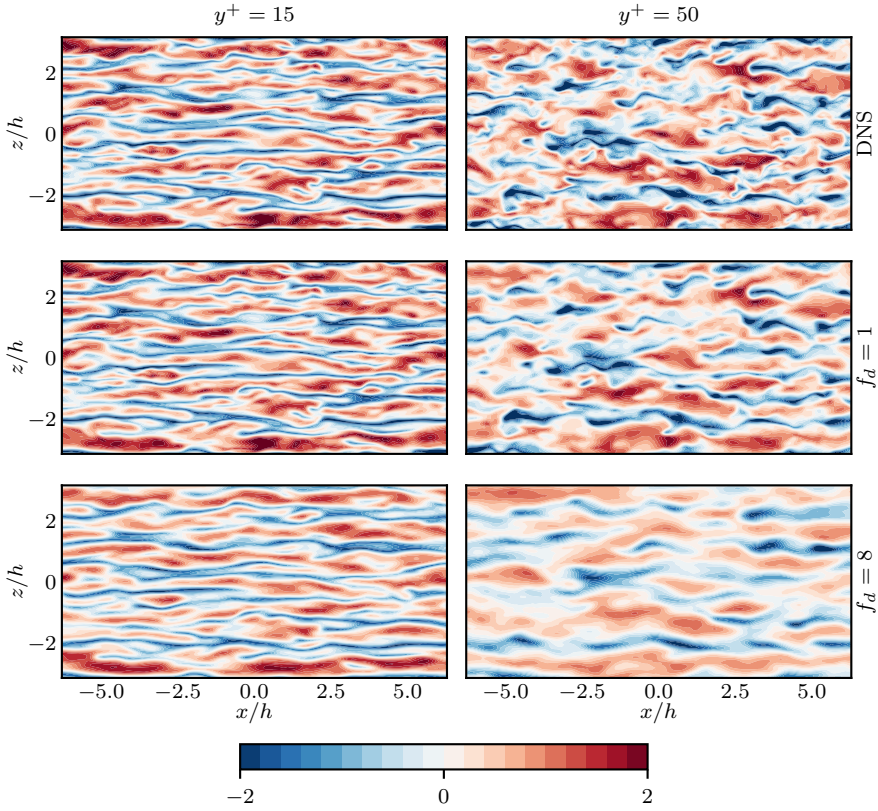

Figure 3.5: Countour map for the streamwise velocity fluctuation fields scaled with their corresponding standard deviation. Top panels denote DNS, while middle and bottom ones refer to SRGAN predictions with f_d equal to 1 and 8 respectively. Left columns refer to $y^+ = 15$, while rigth columns stand for $y^+ = 50$.

Chapter 4

Main contributions and conclusions

The present thesis has dealt with the estimation of coherent structures populating wall-bounded turbulent flows from wall information. Starting from existing linear methods in the literature, such as LSE or EPOD, it has been proposed to analyse the suitability of nonlinear methods based on neural networks. This choice has been based on the fact that turbulent flows are a nonlinear phenomenon both in spatial and temporal dimensions. As it has been shown in Chapter 2 of this thesis, among the different techniques that exist to study coherent structures, POD stands out, since it allows a dimensional reduction of the data, while providing a spatial representation of coherent structures, sorting them according to their energy content. For this reason, it has been combined with neural networks to produce an estimation of the instantaneous flow field. This method has been compared with EPOD, showing that the capture of nonlinearities by the proposed neural network allows the reconstruction of a greater number of scales present in the flow.

Furthermore, the effect of data quality on these neural networks has been investigated. Specifically, it has been analysed how the resolution reduction influences the predictions. Reducing the resolution of the wall data eliminates the smaller scale ranges in the wall, which tends to reduce the ability of the proposed neural network to estimate those scale ranges in the velocity field. To compensate for this reduction in efficiency when predicting velocity fields, a second neural network based on GANs has been proposed to improve the resolution of the wall data. Furthermore, this method, which does not use POD, has been extended to reconstruct directly instantaneous velocity fields from wall measurements with coarse resolution. This method is interesting when combined with transfer learning, as it would allow networks trained with synthetic data to be transferred to real-life applications where acquiring high-resolution measurements can be difficult.

Figure 4.1 summarizes the main achievements of this thesis. Starting from coarse data acquired from the wall of a body immersed in a turbulent flow, the proposed neural networks allow to reconstruct instantaneous velocity fields at different distances from the wall. While Papers 2 and 3 have focused on reconstructing as many scales as possible, Paper 1 had the objective of reconstructing the coherent structures present in the flow. Their results demonstrate that

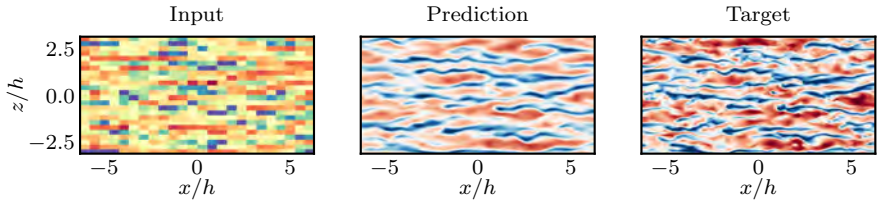

Figure 4.1: Countour map for the SRGAN wall-information input with $f_d = 8$ (left panel), SRGAN-prediction (center panel) and DNS-reference (right panel) streamwise velocity fluctuation fields at $y^+ = 30$. All quantities are scaled with their corresponding standard deviation.

nonlinear flow reconstruction methods based on neural networks can be used to estimate the presence of coherent structures in the logarithmic layer. This is of utmost importance for any application that intends to carry out an active control of the flow from wall measurements, since Lozano-Durán & Jiménez (2014), among others, have shown that wall-attached coherent structures present in this area are responsible for most of the momentum transfer.

4.1. Future work

The promising results of this thesis offer several interesting pathways. Some of the possible future lines of research are listed here:

- While flow reconstructions presented here refer to wall-parallel velocity fields, turbulence and coherent structures are mainly three-dimensional phenomena. Therefore, it is necessary to extend the network to allow predicting volumes.
- The data used in this thesis have been generated from numerical simulations. Therefore, it is necessary to carry out an experimental campaign to demonstrate that neural networks trained with these data can be used to provide predictions in experimental environments. In parallel to this thesis, work has been done on the design and assembly of an experimental setup that will allow evaluating these techniques in an experimental turbulent-boundary-layer flow.
- Although an experimental demonstration is required, it may not always be possible to train a neural network with experimental data. This thesis has already analysed the effect of reducing the amount of information available, but it is necessary to evaluate how the noise present in experimental measurements can affect predictions.

Finally, the significant pace of development that deep learning has experienced in recent years in terms of hardware and algorithms does not seem to be slowing down. Therefore, it is important to transfer the knowledge that can be developed in other fields of research to the study of turbulence.

4.2. Paper highlights

This section highlights the contributions of each of the papers included in the present thesis.

Paper 1

Sensing the turbulent large-scale motions with their wall signature

- A novel convolutional neural network is presented with the ability to predict wall-parallel velocity fields from wall measurements. The network exploits the compression capabilities of proper orthogonal decomposition.
- The method is tested on a DNS database of a turbulent channel flow with friction-based Reynold number $Re_\tau = 1000$.
- Predicting the large coherent structures present in flow fields of size $h \times h$, the new architecture shows superior results when compared to those obtained by a commonly used linear method such as the extended proper orthogonal decomposition.

Paper 2

Convolutional-network models to predict wall-bounded turbulence from wall quantities

- Extension of the POD-based network to predict larger domains with more energy content.
- Prediction of wall-parallel velocity fields from wall measurements in turbulent open-channel flows in the range of friction-based Reynolds number $Re_\tau = [180 - 550]$.
- As the distance between the flow and wall increases, it is found that the FCN-POD architecture is only able to recover the large coherent structures which have an imprint at the wall, although not limited to the linearly-correlated part of the flow.

Paper 3

From coarse wall measurements to turbulent velocity fields with deep learning

- Application of GAN-based networks for resolution enhancement of coarse wall measurements and reconstruction of wall-parallel velocity fields in turbulent open-channel flows with friction-based Reynolds number $Re_\tau = 180$. The effect of using coarse wall measurements is evaluated using three downsampling factors $f_d = [4, 8, 16]$ from the original DNS resolution.
- Results show that flow reconstruction with GANs provides better results than previous architectures.
- The downsampling ratio between the coarse and fine wall measurements is modified to take into account the viscous length scale, which may vary from one work to another.

- Although increasing the downsampling ratio or the distance between the wall and flow fields implies a decrease of accuracy, it is shown that GAN-based networks are able to capture the large-scale coherent structures present in the flow.

Bibliography

Adrian, R. J. 1979 Conditional eddies in isotropic turbulence. *Phys. Fluids* **22** (11), 2065–2070.

Adrian, R. J. 1991 Particle-imaging techniques for experimental fluid mechanics. *Annu. Rev. Fluid Mech.* **23** (1), 261–304.

Adrian, R. J., Meinhart, C. D. & Tomkins, C. D. 2000 Vortex organization in the outer region of the turbulent boundary layer. *J. Fluid Mech.* **422**, 1–54.

Algazi, V. & Sakrison, D. 1969 On the optimality of the Karhunen-Loève expansion. *IEEE Trans. Inform. Theory* **15** (2), 319–321.

Andrews, C., Davies, J. & Schwarz, G. 1967 Adaptive data compression. *Proc. IEEE* **55** (3), 267–277.

Baars, W. J., Hutchins, N. & Marusic, I. 2016 Spectral stochastic estimation of high-Reynolds-number wall-bounded turbulence for a refined inner-outer interaction model. *Phys. Rev. Fluids* **1** (5), 054406.

Bailey, S. C., Vallikivi, M., Hultmark, M. & Smits, A. 2014 Estimating the value of von Kármán's constant in turbulent pipe flow. *J. Fluid Mech.* **749**, 79.

Baldi, P. & Hornik, K. 1989 Neural networks and principal component analysis: Learning from examples without local minima. *Neural Networks* **2** (1), 53–58.

Beck, A., Flad, D. & Munz, C.-D. 2019 Deep neural networks for data-driven LES closure models. *J. Comput. Phys.* **398**, 108910.

Berry, M., Magstadt, A. & Glauser, M. N. 2017 Application of POD on time-resolved schlieren in supersonic multi-stream rectangular jets. *Phys. Fluids* **29** (2), 020706.

Blackwelder, R. F. & Kovasznay, L. S. 1972 Time scales and correlations in a turbulent boundary layer. *Phys. Fluids* **15** (9), 1545–1554.

Borée, J. 2003 Extended proper orthogonal decomposition: a tool to analyse correlated events in turbulent flows. *Exp. Fluids* **35** (2), 188–192.

Bradshaw, P. 2013 *An introduction to turbulence and its measurement: thermodynamics and fluid mechanics series*. Elsevier.

Brown, G. L. & Thomas, A. S. 1977 Large structure in a turbulent boundary layer. *Phys. Fluids* **20** (10), S243–S252.

Brunton, S. L., Noack, B. R. & Koumoutsakos, P. 2020 Machine learning for fluid mechanics. *Annu. Rev. Fluid Mech.* **52**, 477–508.

Ceglia, G., Discetti, S., Ianiro, A., Michaelis, D., Astarita, T. & Cardone, G. 2014 Three-dimensional organization of the flow structure in a non-reactive model aero engine lean burn injection system. *Exp. Therm. Fluid Sci.* **52**, 164–173.

Citriniti, J. H. & George, W. K. 2000 Reconstruction of the global velocity field in the axisymmetric mixing layer utilizing the proper orthogonal decomposition. *J. Fluid Mech.* **418**, 137–166.

De Fauw, J., Ledsam, J. R., Romera-Paredes, B., Nikolov, S., Tomasev, N., Blackwell, S., Askham, H., Glorot, X., O'Donoghue, B., Visentin, D. *et al.* 2018 Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nat. Med.* **24** (9), 1342–1350.

Deane, A., Kevrekidis, I., Karniadakis, G. E. & Orszag, S. 1991 Low-dimensional models for complex geometry flows: application to grooved channels and circular cylinders. *Phys. Fluids* **3** (10), 2337–2354.

Del Alamo, J. C., Jimenez, J., Zandonade, P. & Moser, R. D. 2006 Self-similar vortex clusters in the turbulent logarithmic region. *J. Fluid Mech.* **561**, 329–358.

Delville, J., Ukeiley, L., Cordier, L., Bonnet, J.-P. & Glauser, M. 1999 Examination of large-scale structures in a turbulent plane mixing layer. Part 1. Proper orthogonal decomposition. *J. Fluid Mech.* **391**, 91–122.

DeMers, D. & Cottrell, G. W. 1993 Non-linear dimensionality reduction. In *Adv. Neur. In.*, pp. 580–587.

Deng, Z., He, C., Liu, Y. & Kim, K. C. 2019 Super-resolution reconstruction of turbulent velocity fields using a generative adversarial network-based artificial intelligence framework. *Phys. Fluids* **31** (12), 125111.

Discetti, S., Bellani, G., Örlü, R., Serpieri, J., Sanmiguel Vila, C., Raiola, M., Zheng, X., Mascotelli, L., Talamelli, A. & Ianiro, A. 2019 Characterization of very-large-scale motions in high-Re pipe flows. *Exp. Therm. Fluid Sci.* **104**, 1–8.

Discetti, S., Raiola, M. & Ianiro, A. 2018 Estimation of time-resolved turbulent fields through correlation of non-time-resolved field measurements and time-resolved point measurements. *Exp. Therm. Fluid Sci.* **93**, 119–130.

Dogan, E., Örlü, R., Gatti, D., Vinuesa, R. & Schlatter, P. 2019 Quantification of amplitude modulation in wall-bounded turbulence. *Fluid Dyn. Res.* **51**, 011408.

Encinar, M. P. & Jiménez, J. 2019 Logarithmic-layer turbulence: A view from the wall. *Phys. Rev. Fluids* **4** (11), 114603.

Encinar, M. P. & Jiménez, J. 2020 Momentum transfer by linearised eddies in turbulent channel flows. *J. Fluid Mech.* **895**.

Encinar, M. P., Lozano-Durán, A. & Jiménez, J. 2018 Reconstructing channel turbulence from wall observations. In *Procs. CTR Summer School*.

Erichson, N. B., Mathelin, L., Yao, Z., Brunton, S. L., Mahoney, M. W. & Kutz, J. N. 2020 Shallow neural networks for fluid flow reconstruction with limited sensors. *P. Roy. Soc. A - Math. Phys.* **476** (2238), 20200097.

European Commission 2018 A clean planet for all. A European strategic long-term vision for a prosperous, modern, competitive and climate neutral economy. COM 773.

European Environment Agency, European Aviation Safety Agency & Eurocontrol 2019 European aviation environmental report.

Favre, A., Gaviglio, J. & Dumas, R. 1967 Structure of velocity space-time correlations in a boundary layer. *Phys. Fluids* **10** (9), S138–S145.

Feng, L.-H., Wang, J.-J. & Pan, C. 2011 Proper orthogonal decomposition analysis

of vortex dynamics of a circular cylinder under synthetic jet control. *Phys. Fluids* **23** (1), 014106.

Flores, O. & Jiménez, J. 2010 Hierarchy of minimal flow units in the logarithmic layer. *Phys. Fluids* **22** (7), 071704.

Fukami, K., Fukagata, K. & Taira, K. 2019*a* Super-resolution reconstruction of turbulent flows with machine learning. *J. Fluid Mech.* **870**, 106–120.

Fukami, K., Fukagata, K. & Taira, K. 2021 Machine-learning-based spatio-temporal super resolution reconstruction of turbulent flows. *J. Fluid Mech.* **909**.

Fukami, K., Nabae, Y., Kawai, K. & Fukagata, K. 2019*b* Synthetic turbulent inflow generator using machine learning. *Phys. Rev. Fluids* **4** (6), 064603.

Fukami, K., Nakamura, T. & Fukagata, K. 2020 Convolutional neural network based hierarchical autoencoder for nonlinear mode decomposition of fluid field data. *Phys. Fluids* **32** (9), 095110.

Fukushima, K. 1980 Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biol. Cybern.* **36** (4), 193–202.

Fukushima, K. 1988 Neocognitron: A hierarchical neural network capable of visual pattern recognition. *Neural Networks* **1** (2), 119–130.

Gamahara, M. & Hattori, Y. 2017 Searching for turbulence models by artificial neural network. *Phys. Rev. Fluids* **2** (5), 054604.

Ganapathisubramani, B., Longmire, E. K. & Marusic, I. 2003 Characteristics of vortex packets in turbulent boundary layers. *J. Fluid Mech.* (478), 35–46.

García-Villalba, M., Rossini, L., Gonzalo, A., Vigneault, D., Martinez-Legazpi, P., Durán, E., Flores, O., Bermejo, J., McVeigh, E., Kahn, A. M. *et al.* 2021 Demonstration of patient-specific simulations to assess left atrial appendage thrombogenesis risk. *Front. Physiol.* **12**.

George, W. K. 2007 Is there a universal log law for turbulent wall-bounded flows? *Philos. T. R. Soc. A* **365** (1852), 789–806.

Glauser, M., Zheng, X. & Doering, C. R. 1991 The dynamics of organized structures in the axisymmetric jet mixing layer. In *Fluid Mech. A*, pp. 253–265. Springer.

Glaws, A., King, R. & Sprague, M. 2020 Deep learning for in situ data compression of large turbulent flow simulations. *Phys. Rev. Fluids* **5** (11), 114602.

Goodfellow, I., Bengio, Y., Courville, A. & Bengio, Y. 2016 *Deep learning*. MIT Press Cambridge.

Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. & Bengio, Y. 2014 Generative adversarial networks. arXiv preprint arXiv:1406.2661.

Guala, M., Hommema, S. & Adrian, R. 2006 Large-scale and very-large-scale motions in turbulent pipe flow. *J. Fluid Mech.* **554**, 521–542.

Guastoni, L., Encinar, M. P., Schlatter, P., Azizpour, H. & Vinuesa, R. 2020 Prediction of wall-bounded turbulence from wall quantities using convolutional neural networks. *J. Phys. Conf. Ser.* **1522** (1), 012022.

Güemes, A., Vila, C. S., Örlü, R., Vinuesa, R., Schlatter, P., Ianiro, A. & Discetti, S. 2019 Flow organization in the wake of a rib in a turbulent boundary layer with pressure gradient. *Exp. Ther. Fluid Sci.* **108**, 115–124.

Gurka, R., Liberzon, A. & Hetsroni, G. 2006 POD of vorticity fields: A method

for spatial characterization of coherent structures. *Int. J. Heat Fluid Fl.* **27** (3), 416–423.

Hack, M. P. & Zaki, T. A. 2016 Data-enabled prediction of streak breakdown in pressure-gradient boundary layers. *J. Fluid Mech.* **801**, 43–64.

Ham, Y.-G., Kim, J.-H. & Luo, J.-J. 2019 Deep learning for multi-year ENSO forecasts. *Nature* **573** (7775), 568–572.

He, K., Zhang, X., Ren, S. & Sun, J. 2015 Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proc. ICCV IEEE*, pp. 1026–1034.

He, K., Zhang, X., Ren, S. & Sun, J. 2016 Deep residual learning for image recognition. In *Proc. CVPR IEEE*, pp. 770–778.

Hellström, L. H. & Smits, A. J. 2014 The energetic motions in turbulent pipe flow. *Phys. Fluids* **26** (12), 125102.

Hinton, G. E. & Salakhutdinov, R. R. 2006 Reducing the dimensionality of data with neural networks. *Science* **313** (5786), 504–507.

Hof, B., De Lozar, A., Avila, M., Tu, X. & Schneider, T. M. 2010 Eliminating turbulence in spatially intermittent flows. *Science* **327** (5972), 1491–1494.

Holmes, P., Lumley, J. L., Berkooz, G. & Rowley, C. W. 2012 *Turbulence, coherent structures, dynamical systems and symmetry*. Cambridge University Press.

Hornik, K., Stinchcombe, M. & White, H. 1989 Multilayer feedforward networks are universal approximators. *Neural Networks* **2** (5), 359–366.

Hotelling, H. 1933 Analysis of a complex of statistical variables into principal components. *J. Educ. Psychol.* **24** (6), 417.

Hoyas, S. & Jiménez, J. 2006 Scaling of the velocity fluctuations in turbulent channels up to $Re_\tau = 2003$. *Phys. Fluids* **18** (1), 011702.

Hoyas, S. & Jiménez, J. 2008 Reynolds number effects on the Reynolds-stress budgets in turbulent channels. *Phys. Fluids* **20** (10), 101511.

Hussain, A. F. 1983 Coherent structures—reality and myth. *Phys. Fluids* **26** (10), 2816–2850.

Hutchins, N. & Marusic, I. 2007*a* Evidence of very long meandering features in the logarithmic region of turbulent boundary layers. *J. Fluid Mech.* **579**, 1–28.

Hutchins, N. & Marusic, I. 2007*b* Large-scale influences in near-wall turbulence. *Philos. T. R. Soc. A* **365** (1852), 647–664.

Illingworth, S. J., Monty, J. P. & Marusic, I. 2018 Estimating large-scale structures in wall turbulence using linear models. *J. Fluid Mech.* **842**, 146–162.

Ioffe, S. & Szegedy, C. 2015 Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proc. ICML*, pp. 448–456.

Iten, R., Metger, T., Wilming, H., Del Rio, L. & Renner, R. 2020 Discovering physical concepts with neural networks. *Phys. Rev. Letters* **124** (1), 010508.

Jean, N., Burke, M., Xie, M., Davis, W. M., Lobell, D. B. & Ermon, S. 2016 Combining satellite imagery and machine learning to predict poverty. *Science* **353** (6301), 790–794.

Jeong, J., Hussain, F., Schoppa, W. & Kim, J. 1997 Coherent structures near the wall in a turbulent channel flow. *J. Fluid Mech.* **332** (185-214), 188.

Jiménez, J. 1994 On the structure and control of near wall turbulence. *Phys. Fluids* **6** (2), 944–953.

Jiménez, J. 2012 Cascades in wall-bounded turbulence. *Annu. Rev. Fluid Mech.* **44**, 27–45.

Jiménez, J. 2018*a* Coherent structures in wall-bounded turbulence. *J. Fluid Mech.* **842**.

Jiménez, J. 2018*b* Machine-aided turbulence theory. *J. Fluid Mech.* **854**.

Jiménez, J. 2020*a* Computers and turbulence. *Eur. J. Mech. B-Fluid.* **79**, 1–11.

Jiménez, J. 2020*b* Dipoles and streams in two-dimensional turbulence. *J. Fluid Mech.* **904**.

Jiménez, J. 2020*c* Monte Carlo science. *J. Turbul.* **21** (9-10), 544–566.

Jiménez, J. & Moin, P. 1991 The minimal flow unit in near-wall turbulence. *J. Fluid Mech.* **225**, 213–240.

Jiménez, J. & Pinelli, A. 1999 The autonomous cycle of near-wall turbulence. *J. Fluid Mech.* **389**, 335–359.

Kaiser, E., Noack, B. R., Cordier, L., Spohn, A., Segond, M., Abel, M., Daviller, G., Östh, J., Krajnović, S. & Niven, R. K. 2014 Cluster-based reduced-order modelling of a mixing layer. *J. Fluid Mech.* **754**, 365–414.

Kim, H., Kim, J., Won, S. & Lee, C. 2021 Unsupervised deep learning for super-resolution reconstruction of turbulence. *J. Fluid Mech.* **910**, A29.

Kim, J. & Lee, C. 2020 Prediction of turbulent heat transfer using convolutional neural networks. *J. Fluid Mech.* **882**.

Kim, K. & Adrian, R. 1999 Very large-scale motion in the outer layer. *Phys. Fluids* **11** (2), 417–422.

Kingma, D. P. & Ba, J. 2014 Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.

Kline, S. J., Reynolds, W. C., Schraub, F. & Runstadler, P. 1967 The structure of turbulent boundary layers. *J. Fluid Mech.* **30** (4), 741–773.

Kolmogorov, A. N. 1941 Energy dissipation in locally isotropic turbulence. In *Dokl. Akad. Nauk. SSSR*, , vol. 32, pp. 19–21.

Kutz, J. N. 2017 Deep learning in fluid dynamics. *J. Fluid Mech.* **814**, 1–4.

Kwon, H., Yoon, H., Lee, C., Chen, G., Liu, K., Schmid, A., Wu, Y., Choi, J. & Won, C. 2020 Magnetic Hamiltonian parameter estimation using deep learning techniques. *Science Advances* **6** (39), eabb0872.

Lapeyre, C. J., Misdariis, A., Cazard, N., Veynante, D. & Poinsot, T. 2019 Training convolutional neural networks to estimate turbulent sub-grid scale reaction rates. *Combust. Flame* **203**, 255–264.

LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W. & Jackel, L. D. 1989 Backpropagation applied to handwritten zip code recognition. *Neural Comput.* **1** (4), 541–551.

LeCun, Y., Bottou, L., Bengio, Y. & Haffner, P. 1998 Gradient-based learning applied to document recognition. *Proc. IEEE* **86** (11), 2278–2324.

Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z. *et al.* 2017 Photo-realistic single image super-resolution using a generative adversarial network. In *Proc. CVPR IEEE*, pp. 4681–4690.

Liang, Z. & Dong, H. 2015 On the symmetry of proper orthogonal decomposition modes of a low-aspect-ratio plate. *Phys. Fluids* **27** (6), 063601.

Ling, J., Kurzawski, A. & Templeton, J. 2016 Reynolds averaged turbulence

modelling using deep neural networks with embedded invariance. *J. Fluid Mech.* **807**, 155–166.

Liu, B., Tang, J., Huang, H. & Lu, X.-Y. 2020 Deep learning methods for super-resolution reconstruction of turbulent flows. *Phys. Fluids* **32** (2), 025105.

Long, J., Shelhamer, E. & Darrell, T. 2015 Fully convolutional networks for semantic segmentation. In *Proc. CVPR IEEE*, pp. 3431–3440.

Lozano-Durán, A., Flores, O. & Jiménez, J. 2012 The three-dimensional structure of momentum transfer in turbulent channels. *J. Fluid Mech.* **694**, 100.

Lozano-Durán, A. & Jiménez, J. 2014 Time-resolved evolution of coherent structures in turbulent channels: characterization of eddies and cascades. *J. Fluid Mech.* **759**, 432.

Lumley, J. L. 1967 The structure of inhomogeneous turbulent flows. In *Atmospheric turbulence and radio wave propagation*, pp. 166–178.

Maas, A. L., Hannun, A. Y. & Ng, A. Y. 2013 Rectifier nonlinearities improve neural network acoustic models. In *Proc. ICML*, p. 3.

Mallor, F., Raiola, M., Sanmiguel Vila, C., Örlü, R., Discetti, S. & Ianiro, A. 2019 Modal decomposition of flow fields and convective heat transfer maps: An application to wall-proximity square ribs. *Exp. Therm. Fluid Sci.* **102**, 517–527.

Mallor, F., Sanmiguel Vila, C., Ianiro, A. & Discetti, S. 2018 Wall-mounted perforated cubes in a boundary layer: Local heat transfer enhancement and control. *Int. J. Heat Mass Tran.* **117**, 498–507.

Marusic, I., Mathis, R. & Hutchins, N. 2010*a* Predictive model for wall-bounded turbulent flow. *Science* **329** (5988), 193–196.

Marusic, I., McKeon, B. J., Monkewitz, P. A., Nagib, H. M., Smits, A. J. & Sreenivasan, K. R. 2010*b* Wall-bounded turbulent flows at high Reynolds numbers: recent advances and key issues. *Phys. Fluids* **22** (6), 065103.

Mathis, R., Hutchins, N. & Marusic, I. 2009 Large-scale amplitude modulation of the small-scale structures in turbulent boundary layers. *J. Fluid Mech.* **628**, 311–337.

Mathis, R., Hutchins, N. & Marusic, I. 2011 A predictive inner–outer model for streamwise turbulence statistics in wall-bounded flows. *J. Fluid Mech.* **681**, 537–566.

Mathis, R., Marusic, I., Chernyshenko, S. I. & Hutchins, N. 2013 Estimating wall-shear-stress fluctuations given an outer region input. *J. Fluid Mech.* **715**, 163–180.

Maulik, R. & San, O. 2017 A neural network approach for the blind deconvolution of turbulent flows. *J. Fluid Mech.* **831**, 151–181.

Maulik, R., San, O., Rasheed, A. & Vedula, P. 2019 Subgrid modelling for two-dimensional turbulence using neural networks. *J. Fluid Mech.* **858**, 122–144.

Maurel, S., Borée, J. & Lumley, J. 2001 Extended proper orthogonal decomposition: Application to jet/vortex interaction. *Flow Turbul. Combust.* **67** (2), 125–136.

Mendez, M., Balabane, M. & Buchlin, J.-M. 2019 Multi-scale proper orthogonal decomposition of complex fluid flows. *J. Fluid Mech.* **870**, 988–1036.

Mendez, M., Raiola, M., Masullo, A., Discetti, S., Ianiro, A., Theunissen, R. & Buchlin, J.-M. 2017 POD-based background removal for particle image velocimetry. *Exp. Therm. Fluid Sci.* **80**, 181–192.

Milano, M. & Koumoutsakos, P. 2002 Neural network modeling for near wall turbulent flow. *J. Comput. Phys.* **182** (1), 1–26.

Moehlis, J., Faisst, H. & Eckhardt, B. 2004 A low-dimensional model for turbulent shear flows. *New J. Phys.* **6** (1), 56.

Moin, P. & Moser, R. D. 1989 Characteristic-eddy decomposition of turbulence in a channel. *J. Fluid Mech.* **200**, 471–509.

Monty, J., Hutchins, N., Ng, H., Marusic, I. & Chong, M. 2009 A comparison of turbulent pipe, channel and boundary layer flows. *J. Fluid Mech.* **632**, 431.

Monty, J., Stewart, J., Williams, R. & Chong, M. 2007 Large-scale features in turbulent pipe and channel flows. *J. Fluid Mech.* **589**, 147–156.

Murata, T., Fukami, K. & Fukagata, K. 2020 Nonlinear mode decomposition with convolutional neural networks for fluid dynamics. *J. Fluid Mech.* **882**.

Nair, V. & Hinton, G. E. 2010 Rectified linear units improve restricted boltzmann machines. In *Proc. ICML*, pp. 807–814.

Norouzzadeh, M. S., Nguyen, A., Kosmala, M., Swanson, A., Palmer, M. S., Packer, C. & Clune, J. 2018 Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning. *Proc. Natl. Acad. Sci.* **115** (25), E5716–E5725.

Örlü, R. & Schlatter, P. 2011 On the fluctuating wall-shear stress in zero pressure-gradient turbulent boundary layer flows. *Phys. Fluids* **23** (2), 021704.

Ouro, P., Harrold, M., Stoesser, T. & Bromley, P. 2017 Hydrodynamic loadings on a horizontal axis tidal turbine prototype. *J. Fluids Struct.* **71**, 78–95.

Park, J. & Choi, H. 2021 Toward neural-network-based large eddy simulation: application to turbulent channel flow. *J. Fluid Mech.* **914**.

Pearson, K. 1901 On lines and planes of closest fit to systems of points in space. *Philos. Mag.* **2** (11), 559–572.

Perry, A. & Chong, M. 1982 On the mechanism of wall turbulence. *J. Fluid Mech.* **119**, 173–217.

Pope, S. 1975 A more general effective-viscosity hypothesis. *J. Fluid Mech.* **72** (2), 331–340.

Prandtl, L. 1925 Bericht über die entstehung der turbulenz. *Z. Angew. Math. Mech.* **5**, 136–139.

Preisendorfer, R. W. & Mobley, C. D. 1988 Principal component analysis in meteorology and oceanography. *Dev. Atmosph.* **17**.

Radford, A., Metz, L. & Chintala, S. 2015 Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint arXiv:1511.06434.

Raiola, M., Discetti, S. & Ianiro, A. 2021 Data-driven identification of unsteady-aerodynamics phenomena in flapping airfoils. *Exp. Therm. Fluid Sci.* **124**, 110234.

Raiola, M., Ianiro, A. & Discetti, S. 2016 Wake of tandem cylinders near a wall. *Exp. Therm. Fluid Sci.* **78**, 354–369.

Raissi, M., Perdikaris, P. & Karniadakis, G. E. 2019*a* Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707.

Raissi, M., Wang, Z., Triantafyllou, M. S. & Karniadakis, G. E. 2019*b* Deep learning of vortex-induced vibrations. *J. Fluid Mech.* **861**, 119–137.

Raissi, M., Yazdani, A. & Karniadakis, G. E. 2020 Hidden fluid mechanics:

Learning velocity and pressure fields from flow visualizations. *Science* **367** (6481), 1026–1030.

Reichardt, H. 1951 Vollständige darstellung der turbulenten geschwindigkeitsverteilung in glatten leitungen. *Z. Angew. Math. Mech.* **31** (7), 208–219.

Richardson, L. F. 1920 The supply of energy from and to atmospheric eddies. *Proc. R. Soc. Lond. A* **97** (686), 354–373.

Robinson, S. K. 1991 Coherent motions in the turbulent boundary layer. *Annu. Rev. Fluid Mech.* **23** (1), 601–639.

Rosenblatt, F. 1958 The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol. Rev.* **65** (6), 386.

Rowley, C. W. & Dawson, S. T. 2017 Model reduction for flow analysis and control. *Annu. Rev. Fluid Mech.* **49**, 387–417.

Rumelhart, D. E., Hinton, G. E. & Williams, R. J. 1986 Learning representations by back-propagating errors. *Nature* **323** (6088), 533–536.

Sanmiguel Vila, C. & Flores, O. 2018 Wall-based identification of coherent structures in wall-bounded turbulence. *J. Phys. Conf. Ser.* **1001** (1), 012007.

Sanmiguel Vila, C., Örlü, R., Vinuesa, R., Schlatter, P., Ianiro, A. & Discetti, S. 2017 Adverse-pressure-gradient effects on turbulent boundary layers: statistics and flow-field organization. *Flow Turbul. Combust.* **99** (3), 589–612.

Sasaki, K., Vinuesa, R., Cavalieri, A. V., Schlatter, P. & Henningson, D. S. 2019 Transfer functions for flow predictions in wall-bounded turbulence. *J. Fluid Mech.* **864**, 708–745.

Schlatter, P. & Örlü, R. 2010 Assessment of direct numerical simulation data of turbulent boundary layers. *J. Fluid Mech.* **659**, 116.

Schrauf, G. 2005 Status and perspectives of laminar flow. *The Aeronautical Journal* **109** (1102), 639–644.

Shi, W., Caballero, J., Huszár, F., Totz, J., Aitken, A. P., Bishop, R., Rueckert, D. & Wang, Z. 2016 Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proc. CVPR IEEE*, pp. 1874–1883.

Sillero, J. A., Jiménez, J. & Moser, R. D. 2013 One-point statistics for turbulent wall-bounded flows at Reynolds numbers up to $\delta^+ \approx 2000$. *Phys. Fluids* **25** (10), 105102.

Sirovich, L. 1987 Turbulence and the dynamics of coherent structures. I. Coherent structures. *Q. Appl. Math.* **45** (3), 561–571.

Soummer, R., Pueyo, L. & Larkin, J. 2012 Detection and characterization of exoplanets and disks using projections on Karhunen-Loève eigenimages. *Astrophys. J. Lett.* **755** (2), L28.

Suzuki, T. & Hasegawa, Y. 2017 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.

Theodorsen, T. 1952 Mechanisms of turbulence. In *Proc. 2nd Midwestern Conf. of Fluid Mechanics*.

Thuerey, N., Weissenow, K., Prantl, L. & Hu, X. 2020 Deep learning methods for Reynolds-averaged Navier–Stokes simulations of airfoil flows. *AIAA J.* **58** (1), 25–36.

Tinney, C., Coiffet, F., Delville, J., Hall, A., Jordan, P. & Glauser, M. N. 2006 On spectral linear stochastic estimation. *Exp. Fluids* **41** (5), 763–775.

Tinney, C. E., Ukeiley, L. & Glauser, M. N. 2008 Low-dimensional characteristics of a transonic jet. Part 2. Estimate and far-field prediction. *J. Fluid Mech.* **615**, 53–92.

Tomkins, C. D. & Adrian, R. J. 2003 Spanwise structure and scale growth in turbulent boundary layers. *J. Fluid Mech.* **490**, 37–74.

Townsend, A. 1976 *The structure of turbulent shear flow*. Cambridge University Press.

Troshin, V. & Seifert, A. 2018 Modeling of a pitching and plunging airfoil using experimental flow field and load measurements. *Exp. Fluids* **59** (1), 1–14.

Tsinober, A. 2009 *An informal conceptual introduction to turbulence*. Springer.

Udrescu, S.-M. & Tegmark, M. 2020 AI Feynman: A physics-inspired method for symbolic regression. *Science Advances* **6** (16), eaay2631.

UNFCCC, V. 2015 Adoption of the Paris agreement. Proposal by the President.

Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., Felländer, A., Langhans, S. D., Tegmark, M. & Nerini, F. F. 2020 The role of artificial intelligence in achieving the sustainable development goals. *Nat. Commun.* **11** (1), 1–10.

Vinuesa, R., Negi, P. S., Atzori, M., Hanifi, A., Henningson, D. S. & Schlatter, P. 2018 Turbulent boundary layers around wing sections up to Re_c= 1,000,000. *Int. J. Heat Fluid Fl.* **72**, 86–99.

Von Kármán, T. 1930 Mechanische ahnlichkeit und turbulenz. *Math. Phys. Klasse* .

Waldmann, I. P. & Griffith, C. A. 2019 Mapping Saturn using deep learning. *Nat. Astron.* **3** (7), 620–625.

Waleffe, F. 1995 Hydrodynamic stability and turbulence: Beyond transients to a self-sustaining process. *Stud. Appl. Math.* **95** (3), 319–343.

Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y. & Change Loy, C. 2018 Esrgan: Enhanced super-resolution generative adversarial networks. In *Proc. ECCV*.

Werhahn, M., Xie, Y., Chu, M. & Thuerey, N. 2019 A multi-pass GAN for fluid flow super-resolution. *Proc. ACM on Computer Graphics and Interactive Techniques* **2** (2), 1–21.

Wu, X. & Moin, P. 2010 Transitional and turbulent boundary layer with heat transfer. *Phys. Fluids* **22** (8), 085105.

Wu, Y. 2014 A study of energetic large-scale structures in turbulent boundary layer. *Phys. Fluids* **26** (4), 045113.

Xie, Y., Franz, E., Chu, M. & Thuerey, N. 2018 tempoGAN: A temporally coherent, volumetric GAN for super-resolution fluid flow. *ACM T. Graphic.* **37** (4), 1–15.

Yao, J. & Yeo, K. 2020 Forward flight and sideslip manoeuvre of a model hawkmoth. *J. Fluid Mech.* **896**.

Zeleznik, R., Foldyna, B., Eslami, P., Weiss, J., Alexander, I., Taron, J., Parmar, C., Alvi, R. M., Banerji, D., Uno, M. *et al.* 2021 Deep convolutional neural networks to predict cardiovascular risk from computed tomography. *Nat. Commun.* **12** (1), 1–9.

Zhu, J.-Y., Park, T., Isola, P. & Efros, A. A. 2017 Unpaired image-to-image

translation using cycle-consistent adversarial networks. In *Proc. ICCV IEEE*, pp. 2223–2232.

Part II

Papers

Summary of the papers

Paper 1

Sensing the turbulent large-scale motions with their wall signature

This study assesses the capability of extended proper orthogonal decomposition (EPOD) and convolutional neural networks (CNNs) to reconstruct large-scale and very-large-scale motions (LSMs and VLSMs respectively) employing wall-shear-stress measurements in wall-bounded turbulent flows. Both techniques are used to reconstruct the instantaneous LSM evolution in the flow field as a combination of proper orthogonal decomposition (POD) modes, employing a limited set of instantaneous wall-shear-stress measurements. Due to the dominance of nonlinear effects, only CNNs provide satisfying results. Being able to account for nonlinearities in the flow, CNNs are shown to perform significantly better than EPOD in terms of both instantaneous flow-field estimation and turbulent-statistics reconstruction. CNNs are able to provide a more effective reconstruction performance employing more POD modes at larger distances from the wall and employing lower wall-measurement resolutions. Furthermore, the capability of tackling nonlinear features of CNNs results in estimation capabilities that are weakly dependent on the distance from the wall.

Paper 2

Convolutional-network models to predict wall-bounded turbulence from wall quantities

Two models based on convolutional neural networks are trained to predict the two-dimensional instantaneous velocity-fluctuation fields at different wall-normal locations in a turbulent open-channel flow, using the wall-shear-stress components and the wall pressure as inputs. The first model is a fully-convolutional neural network (FCN) which directly predicts the fluctuations, while the second one reconstructs the flow fields using a linear combination of orthonormal basis functions, obtained through proper orthogonal decomposition (POD), hence named FCN-POD. Both models are trained using data from direct numerical simulations (DNS) at friction Reynolds numbers $Re_\tau = 180$ and 550. Being able to predict the nonlinear interactions in the flow, both models show better predictions than the extended proper orthogonal decomposition (EPOD), which

establishes a linear relation between input and output fields. The performance of the models is compared based on predictions of the instantaneous fluctuation fields, turbulence statistics and power-spectral densities. FCN exhibits the best predictions closer to the wall, whereas FCN-POD provides better predictions at larger wall-normal distances. We also assessed the feasibility of transfer learning for the FCN model, using the model parameters learned from $Re_\tau = 180$ dataset to initialize those of the model that is trained on the $Re_\tau = 550$ dataset. After training the initialized model at the new Re_τ, our results indicate the possibility to match the reference-model performance up to $y^+ = 50$, with 50% and 25% of the original training data. We expect that these non-intrusive sensing models will play an important role in applications related to closed-loop control of wall-bounded turbulence.

Paper 3

From sparse wall measurements to turbulent velocity fields through deep learning

This work evaluates the applicability of super-resolution generative adversarial networks (SRGANs) as a methodology for the reconstruction of turbulent-flow quantities from coarse wall measurements. The method is applied both for the resolution enhancement of wall fields and the estimation of wall-parallel velocity fields from coarse wall measurements of shear stress and pressure. The analysis has been carried out with a database of a turbulent open-channel flow with friction Reynolds number $Re_\tau = 180$ generated through direct numerical simulation. Coarse wall measurements have been generated with three different downsampling factors $f_d = [4, 8, 16]$ from the high-resolution fields, and wall-parallel velocity fields have been reconstructed at four inner-scaled wall-normal distances $y^+ = [15, 30, 50, 100]$. We first show that SRGAN can be used to enhance the resolution of coarse wall measurements. If compared with direct reconstruction from the sole coarse wall measurements, SRGAN provides better instantaneous reconstructions, both in terms of mean-squared error and spectral-fractional error. Even though lower resolutions in the input wall data make it more challenging to achieve highly accurate predictions, the proposed SRGAN-based network yields very good reconstruction results. Furthermore, it is shown that even for the most challenging cases the SRGAN is capable of capturing the large-scale structures that populate the flow. The proposed novel methodology has great potential for closed-loop control applications relying on non-intrusive sensing.

1

Paper 1

Sensing the turbulent large-scale motions with their wall signature

Alejandro Güemes[1], Stefano Discetti[1] and Andrea Ianiro[1]

[1] Aerospace Research Group, Universidad Carlos III de Madrid, Avenida de la Universidad 30, Leganés 28911, Spain

Physics of Fluids (2019), vol. **31** (12), 125112

This study assesses the capability of extended proper orthogonal decomposition (EPOD) and convolutional neural networks (CNNs) to reconstruct large-scale and very-large-scale motions (LSMs and VLSMs respectively) employing wall-shear-stress measurements in wall-bounded turbulent flows. Both techniques are used to reconstruct the instantaneous LSM evolution in the flow field as a combination of proper orthogonal decomposition (POD) modes, employing a limited set of instantaneous wall-shear-stress measurements. Due to the dominance of nonlinear effects, only CNNs provide satisfying results. Being able to account for nonlinearities in the flow, CNNs are shown to perform significantly better than EPOD in terms of both instantaneous flow-field estimation and turbulent-statistics reconstruction. CNNs are able to provide a more effective reconstruction performance employing more POD modes at larger distances from the wall and employing lower wall-measurement resolutions. Furthermore, the capability of tackling nonlinear features of CNNs results in estimation capabilities that are weakly dependent on the distance from the wall.

1. Introduction

Wall-bounded turbulent flows are present in many industrial and aeronautical applications, including all kind of vehicles immersed in a fluid or turbomachinery flows. Understanding the dynamics of these flows is of utmost importance to enabling sensing and control to improve the performances in terms of skin friction, heat transfer, etc.

Among the spectral pipeline of scales in a turbulent boundary layer, large-scale and very-large-scale motions (LSMs and VLSMs respectively, shown to have a wavelength of up to more than 10 flow characteristics lengths) are one of the main sources of momentum and energy transport. LSMs and VLSMs are regions of the flow that present spatial and temporal correlation, whose evidences are available in several research studies (Kline *et al.* 1967; Favre *et al.* 1967; Blackwelder & Kovasznay 1972; Wark & Nagib 1991; Kim & Adrian 1999; Hutchins & Marusic 2007*a*). LSMs have been documented as long features of uniform high or low momentum in the logarithmic region of a turbulent

boundary layer. Hutchins & Marusic (2007*a*) proved using a spanwise rake of hot-wires that these elongated regions extended up to a length of the order of 20 boundary layer thicknesses δ_{99}. These features were also found experimentally in turbulent pipe flows and channel flows, as reported by Monty *et al.* (2007) and Discetti *et al.* (2019).

LSMs have been proved both experimentally (Hutchins & Marusic 2007; Mathis *et al.* 2009) and numerically (Hoyas & Jiménez 2006, 2008) to have a direct impact on the inner-scales fluctuations close to the wall. It has been shown that the LSMs present in the logarithmic region have an amplitude and frequency modulation effect on the streamwise fluctuations close to the wall (Marusic *et al.* 2010; Mathis *et al.* 2011), which has also been extended to the wall-normal and spanwise fluctuations (Ganapathisubramani *et al.* 2012). Mathis *et al.* (2013) showed that it is possible to estimate the wall-shear-stress fluctuations from the one-point measurements of LSMs in the logarithmic region. These works (Marusic *et al.* 2010; Mathis *et al.* 2011; Ganapathisubramani *et al.* 2012; Mathis *et al.* 2013) proposed a model that assumes that the near-wall small-scale structures are universal(Jiménez & Pinelli 1999) and that the effect of Reynolds number can be modelled as an amplitude and frequency modulation effect of the LSMs on the small scales. Furthermore, they showed that the spatial resolution of the measurements was a driving parameter to obtain accurate predictions.

Due to the relevance that LSMs have on the flow, it has become of utmost importance to define techniques that allow to detect and control them. Abbassi *et al.* (2017) proposed the use of jets to reduce the skin-friction drag in turbulent boundary layers. Piezo-ceramic actuators to generate a transverse travelling wave along the wall surface (Bai *et al.* 2014) or a spanwise-oscillating surface (Gouder *et al.* 2013) have also been explored for the same purpose. Suction and blowing wall jets that weaken the near-wall streaks (Lee *et al.* 1997) have also been proven as an effective solution.

Closed-loop control strategies require proper sensorization strategies for the flow; the clear role played by LSMs thus set the requirements of an effective real-time LSMs-detection solution with minimum intrusiveness. Pointwise probes have been often explored as a good compromise between the needs of low-processing time and minimization of flow disturbance. The flow state is then estimated using only the probe data, thus requiring robust algorithms leveraging statistical evidence of correlations between flow features and probe data. Linear stochastic estimation (Adrian 1996) has been successfully used to estimate turbulent fields (Tinney *et al.* 2006, 2008; Baars *et al.* 2016; Suzuki & Hasegawa 2017; Illingworth *et al.* 2018; Encinar *et al.* 2018) based on a limited set of probes. One of the most promising approaches that have been proposed is the extended proper orthogonal decomposition (EPOD, see e.g. Borée (2003)). Using the projection of the snapshot matrix of a given quantity on the temporal basis corresponding to another one, it has been shown that it is possible to extract the correlation between synchronized measurements

through the extended POD modes (Borée 2003). The compactness of the POD basis allows data-driven filtering, thus providing more robust estimations than the simple application of linear stochastic estimation. Using this correlation it has been shown that accurate far-field predictions of a turbulent jet can be obtained, (Tinney *et al.* 2008) as a combination of the flow-field POD modes through time coefficients linearly-estimated based on instantaneous point-wise measurements. The EPOD has been used previously to reconstruct turbulent flow fields using velocity probes in a turbulent channel (Discetti *et al.* 2018) and in a high-Reynolds-number pipe flow (Discetti *et al.* 2019).

The EPOD approach in Discetti *et al.* (2018), however, does not allow to take into account nonlinearities in the relation between probe data and flow features. More specifically, the amplitude modulation effects that LSMs have on the near-wall scales (Hutchins & Marusic 2007; Mathis *et al.* 2009) could not be identified using EPOD. One promising path to embed these non-linear effects is the use of artificial neural networks (ANNs), firstly introduced with the *perceptron* concept (Rosenblatt 1958). ANNs are a set of layers that transform a fixed-size input (for example, a flow velocity field) into a fixed-sized output (for example, a coefficient) by computing a weighted sum of their inputs from the previous layer and applying a non-linear function to the result (LeCun *et al.* 2015). It has been shown that an ANN with a single hidden layer is equivalent to perform a POD on the data (Baldi & Hornik 1989) while non-linear effects can be introduced making use of several hidden layers.

ANNs have been applied to a large variety of research fields, including fluid mechanics. The main applications include the solution of the Reynolds-Averaged Navier-Stokes (RANS) and of the large-eddy simulation (LES) problems, improving turbulence modelling (Ling *et al.* 2016; Wu *et al.* 2017; Maulik *et al.* 2019; Nikolaou *et al.* 2019). Raissi *et al.* (2019) used a multi-layer perceptron (MLP) to estimate the lift coefficient and the displacement of a freely vibrating cylinder from flow measurements in the surroundings. Additionally, they were able to estimate the flow fields from the concentration of a passive scalar in the flow. Another possible application has been shown in Hack & Zaki (2016), where streaks in pressure-gradient boundary layers are classified into those that will break into turbulence and those which will not.

As reviewed by Brunton *et al.* (2020), several authors employed machine learning for modelling fluid flows. A promising application of ANNs in turbulence is the direct reconstruction of turbulent flows from external measurements. In particular, Milano & Koumoutsakos (2002) carried out a reconstruction of the near-wall flow using a neural network to approximate the higher-order terms of the second-order model as:

$$\underline{u}(\underline{x},t) = \underline{\omega}\, y + \frac{Re}{2}\frac{d\underline{P_w}}{\underline{x}} y^2 + ANN(\underline{P_w}, \underline{\tau_w}), \tag{1}$$

where $\underline{\omega}$ is the wall tangential vorticity, $\underline{P_w}$ is the wall pressure field,and $\underline{\tau_w}$ the wall-shear-stress field. The model provided a good reconstruction up to $y^+ \approx 40$ using wall-shear and pressure measurements. Superscript + refers to

inner-scaling, based on the friction velocity $u_\tau = \sqrt{\tau_w/\rho}$ and viscous length $l_* = \nu/u_\tau$ (being τ_w, ρ, and ν the mean wall-shear stress, density and viscosity respectively). An ANN model was proposed in Maulik & San (2017) for the deconvolution and regularization of low-pass spatially filtered turbulent fields. In Fukami *et al.* (2019) two different ANN models were used to increase the resolution of turbulent fields; they use DNS fields to train the model to upgrade low-resolution fields into fully-resolved ones. Lusch *et al.* (2018) has shown that ANNs can be used to predict the nonlinear dynamics of the flow around a cylinder. In Srinivasan *et al.* (2019) both MLP and long short-term memory (LSTM) ANNs have been used to predict the temporal evolution of a low-order model of a turbulent flow. Due to the sequential nature of LSTM, their results outperformed those obtained by the MLP, providing very accurate turbulence statistics. Based on the results of these studies, the ANNs appear to be a perfect candidate to estimate the dynamics of POD modes from a limited set of time-resolved wall measurements.

Among ANNs, convolutional neural networks (CNNs) take advantage of their input hierarchical shape (Fukushima 1980; LeCun 1985; Fukushima 1988; LeCun *et al.* 1989). CNNs have been applied to a large variety of research fields (Liu *et al.* 2017; Ou & Li 2019), and are recently being exploited for turbulent-flow estimation. For instance, CNNs were used in Fukami *et al.* (2019) to increase the resolution of the flow field, employing as input the flow field itself but with lower resolution. Sasaki *et al.* (2019) reports a remarkable application in which convolutional kernels with different inputs are used to estimate the streamwise velocity fluctuations of a zero-pressure-gradient turbulent boundary layer. In Guastoni *et al.* (2020*a*) CNNs have been used to reconstruct instantaneous velocity fields parallel to the wall up to distances of $y^+ = 50$. Their results show that first- and second-order statistics can be computed accurately from the reconstructed flow fields, although the reconstruction error increases with the wall distance.

Aiming at the identification of a suitable sensing device for the detection of large- and very-large-scale motions in a turbulent wall-bounded flow, we explore the performances of CNNs for the identification of the large-scale instantaneous flow fields from a limited set of shear-wall probes. A comparison with the EPOD is carried out to assess whether the non-linearity provided by CNNs can improve the quality of the flow field reconstruction.

This work is based on a computational dataset which is used both for the training of the estimation methodologies and for the assessment of the reconstruction performances. Wall-shear-stress values at a limited number of locations are employed to reconstruct the flow field in a turbulent channel flow at different wall-parallel planes ranging from $y^+ = 50$ to $y^+ = 200$, being the friction-based Reynolds number equal to $Re_\tau = 1000$. The paper is organised as follows: Sec. 2 describes the different methodologies employed to reconstruct the LSMs in the flow field and the database used to assess their performance. In

Sec. 3 the obtained results are presented and discussed, prior to the conclusions in Sec. 4.

2. Methodology

The reference frame $\underline{x} = [x, y, z]$ adopted in this work is aligned with the streamwise x, wall-normal y and spanwise z directions, being their corresponding fluctuating velocity components denoted as $\underline{u} = [u, v, w]$. Velocity fields are decomposed according to the Reynolds decomposition; the fluctuating velocity components are then expressed hereafter as a linear combination of a spatial basis, composed of spatially orthonormal functions $\underline{\phi}_i(\underline{x})$, times a temporal basis, made of temporally orthonormal functions $\psi_i(t)$, by means of the proper orthogonal decomposition (POD) (Lumley 1967):

$$\underline{u}(\underline{x}, t) \approx \sum_{i=1}^{Nm} \psi_i(t)\sigma_i\underline{\phi}_i(\underline{x}), \tag{2}$$

with N_m being the number of POD modes and σ_i weighs the contribution of each mode to the snapshots. POD modes are optimal in the least square sense: the first r POD modes form an orthogonal basis which minimizes the Frobenius norm of the flow field reconstruction error, i.e. provides the best rank-r approximation of the snapshots. This implies that POD provides flow field modes ordered by their energy content. Owing to the prevalent role in momentum and energy transport of LSMs, it has indeed been shown that the first POD modes can be thought of as representations of the LSMs present in the flow (Wu 2014).

The practical implementation with the dataset described in Sec. 2.3 follows the snapshot POD method (Sirovich 1987). The velocity fields are reshaped in a vector form and rearranged in a matrix $\underline{\underline{U}}$ of size $n_t \times n_p$, where n_t is the number of available flow fields and n_p is the number of grid points. The corresponding SVD is:

$$\underline{\underline{U}} = \underline{\underline{\Psi}}\,\underline{\underline{\Sigma}}\,\underline{\underline{\Phi}}. \tag{3}$$

Assuming statistical convergence, flow fields at a generic time instant (inside and outside of the training dataset) can be quite accurately be described as a linear combination of the modes contained in the $\underline{\underline{\Sigma}}\,\underline{\underline{\Phi}}$ matrix by the corresponding time coefficients, if available. One strategy to obtain these time coefficients for time instants outside of the dataset is the synchronized use of pointwise probes and flow field data, with the probes being trained to *reconstruct* the velocity fields. The velocity field at a given instant can then be expressed as:

$$\underline{u}^* = \underline{\psi}^*\,\underline{\underline{\Sigma}}\,\underline{\underline{\Phi}}, \tag{4}$$

where estimated quantities are denoted with the superscript $*$. In the following we describe two methodologies employed for the estimation of the time coefficients $\underline{\psi}^*$.

2.1. *Extended proper orthogonal decomposition*

The snapshot POD can be applied both to the flow field and to any other quantity measured by synchronized probes immersed in the flow or at the wall, e.g. corresponding to a set of probes $\underline{\underline{U}}_{pr}$:

$$\underline{\underline{U}}_{pr} = \underline{\underline{\Psi}}_{pr} \underline{\underline{\Sigma}}_{pr} \underline{\underline{\Phi}}_{pr}. \tag{5}$$

The extended POD modes (Borée 2003), corresponding to the projection of the probe measurements, on the flow-field temporal basis can be estimated as:

$$\underline{\underline{\Psi}}_{pr}^T \underline{\underline{U}} = \underline{\underline{\Psi}}_{pr}^T \underline{\underline{\Psi}} \underline{\underline{\Sigma}} \underline{\underline{\Phi}} = \underline{\underline{\Xi}} \underline{\underline{\Sigma}} \underline{\underline{\Phi}} = \underline{\underline{\Sigma}}_e \underline{\underline{\Phi}}_e, \tag{6}$$

where the subscript e refers to extended POD modes and the matrix $\underline{\underline{\Xi}} = \underline{\underline{\Psi}}_{pr}^T \underline{\underline{\Psi}}$ is a matrix containing information about the temporal correlation between flow and probe modes. If the dataset is sufficiently large to reach statistical convergence, it is possible to assume that the matrix $\underline{\underline{\Xi}}$ represents a good estimate of the relation between the POD time coefficients at a certain probe snapshot and those of an out-of-sample estimate of the flow field. Taking this and equation 6 into account, it is possible to estimate the temporal POD coefficients of the flow field using the wall-field POD time coefficient and the temporal correlation matrix:

$$\underline{\underline{\Psi}}^* = \underline{\underline{\Psi}}_{pr} \underline{\underline{\Xi}}. \tag{7}$$

A flow field can thus be reconstructed by means of a set of probes once the temporal correlation matrix is known. In this work we will explore the use of wall sensors, since it is of more practical application if compared to probes immersed in the flow.

The EPOD reconstruction is equivalent to a linear stochastic estimation and as shown in Discetti *et al.* (2018) it is an energy preserving process. The energy-optimality of POD modes can then be enforced by employing statistical filters (Discetti *et al.* 2018) and improving the reconstruction with respect to LSE by cleaning the data from uncorrelated features. This approach has been shown to be successful in several application, but it is not able to take into account non-linear effects.

$$\underline{\underline{U}}^* = \underline{\underline{\Psi}}^* \underline{\underline{\Sigma}} \underline{\underline{\Phi}} \tag{8}$$

2.2. *Artificial neural network*

The alternative approach pursued here to estimate the POD time coefficient is through the use of ANNs. Probe data can be rearranged in a tensor form $\underline{\underline{\underline{W}}}$ of size $n_t \times n_w \times n_h$, where n_w is the number of grid points in the width direction and n_h is the number of grid points in the height direction ($n_p = n_w \times n_h$).

The tensor $\underline{\underline{\underline{W}}}$ is passed through a set of convolutional and max-pooling layers that halve sequentially the field size. The number of convolutional and max-pooling layers depends on the initial wall-shear stress field resolution, thus three different CNN architectures have been set depending on this feature.

When the tensor has been reduced to the minimum possible size, it is reshaped in a vector form $\underline{v}$ by means of a flatten layer, and it is passed through two fully-connected layers with weights $\underline{\underline{w}}$. The last fully-connected layer outputs a single value that corresponds to the temporal POD mode of the flow field associated to the wall-shear-stress field.

$$\underline{\psi}_i^* = \underline{\underline{w}}_2 \times \underline{\underline{w}}_1 \times \underline{v} \tag{9}$$

Finally, the temporal POD-mode estimation obtained through CNNs can be used to reconstruct the flow fields from a set of wall probes applying Eq. 8. Different CNNs have been trained for different probe spacings, wall-distance positions of the flow fields and the first 10 temporal POD modes, resulting in a total of 120 different CNNs. A simplified sketch of the CNN architectures can be seen in Fig. 1.

A rectified linear unit ($ReLU$) activation function (Nair & Hinton 2010) has been used in the fully connected layers of the CNNs architecture to take into account the non-linearities of the flow, while an Adam optimizer (Kingma & Ba 2014) with a constant learning rate of 0.001 has been selected to update the weights of the different neuron layers.

$$ReLU(x) = max\{0, x\} \tag{10}$$

2.3. *Dataset description*

The different methods have been tested employing a Direct Numerical Simulation (DNS) dataset that includes full flow-field information from a numerical simulation in a channel flow at friction Reynolds number $Re_\tau \approx 1000$. The channel-flow database is available at the Johns Hopkins Turbulence Database (http://turbulence.pha.jhu.edu/). The channel has a dimension of 2 half-channel-heights h from wall to wall, $3\pi h$ in the span-wise direction and $8\pi h$ in the stream-wise direction. The DNS database covers one channel-flow through-time $8\pi h/U_b$ (where U_b is the channel bulk velocity) with a DNS time step of

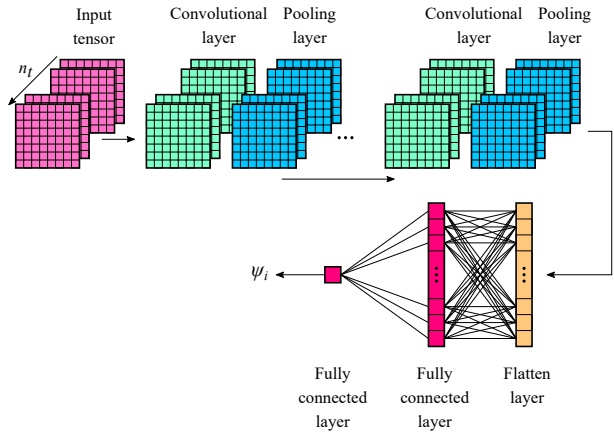

Figure 1: Schematic view of the CNN architecture.

$\delta_t = 0.0014h/U_b$. More details on the numerical simulations can be found by Li *et al.* (2008); Yu *et al.* (2012).

To train the proposed methodologies and to assess their capability of reconstructing the LSMs, we collected both wall-shear-stress fields and $[u, v]$ fields in streamwise-spanwise planes located at four distances from the wall. These fields were representative of square $h \times h$ domains parallel to the wall located at $y^+ = [50, 100, 150, 200]$. The fields resolution is equal to 128×128 points. The wall-friction fields were extracted at the same temporal instants in the same $[x, z]$ region with three different spatial resolutions, $[128 \times 128, 64 \times 64, 32 \times 32]$, covering in all cases the same $h \times h$ area.

Only 24 flow and wall-shear fields, separated by $t^+ = 1000$ and therefore uncorrelated, can be extracted in the same position due to the limited timespan of the simulation. To acquire a well-converged dataset, the flow and wall-shear fields have been extracted at different positions in the streamwise and spanwise directions of the bottom part of the channel domain, taking advantage of the simulation homogeneity in those directions. Specifically, flow fields have been obtained in 26 streamwise locations and in 10 spanwise locations. Fields have been extracted each 0.25 convective times, leading to 100 different temporal fields at each location. The total amount of available flow and wall fields is 26000. Taking into account that this quantity is more than enough to reach statistical convergence and that the dataset is aimed to be used for training, validation and testing of different LSMs estimation methods, three field subsets have been extracted: 7000 fields for training, 500 fields for validation (only used

Case	y^+	Resolution [pt]	EPOD symbol	CNN symbol
H050	50	128×128	(–□–)	(–■–)
H100	100			
H150	150			
H200	200			
M050	50	64×64	(–○–)	(–●–)
M100	100			
M150	150			
M200	200			
L050	50	32×32	(–△–)	(–▲–)
L100	100			
L150	150			
L200	200			

Table 1: Dataset description. For each case, 7000 fields have been extracted for training, 500 for validation, and 7500 for testing.

in the CNN) and 7500 fields for testing. It is important to mention that the field subsets have been separated in the time domain, i.e. belong to snapshots sets extracted respectively at the beginning and at the end of the sequence. A sketch of the fields extracted from the DNS domain is given in figure 2, while a description of each dataset is provided in table 1.

The POD modes have been computed using the fluctuating velocity fields. Figure 3 reports both the modes energy distribution (i.e the eigenvalue distribution $\lambda_i = \sigma_i^2$) and the cumulative energy content for the flow-field snapshots. It is well known that the first POD modes correspond to the large-scale motions present in the flow (Liu *et al.* 2001; Wu 2014) while following modes are associated with smaller and less-energetic structures. According to previous literature, LSMs and VLSMs account for approximately 40% of the turbulent kinetic energy (Discetti *et al.* 2019; Guala *et al.* 2006; Sanmiguel Vila *et al.* 2017) which approximately correspond here to the first 10 POD modes. In the following, we thus focus on the estimation of the time coefficients for these modes only, resulting in a low-order reconstruction of the LSMs and VLSMs. Turbulent flows are characterized by having a rich range of spatial scales, which in this case makes necessary a large number of modes (between 400 and 900) to reach 90% of the cumulative energy spectrum. When analyzing the energy contribution of each mode at different wall-normal positions, it can be seen that, the closer to the wall, the smaller is the contribution of the LSMs, which is not surprising since close to the wall the smallest scales energy contribution is more significant (Ganapathisubramani *et al.* 2012).

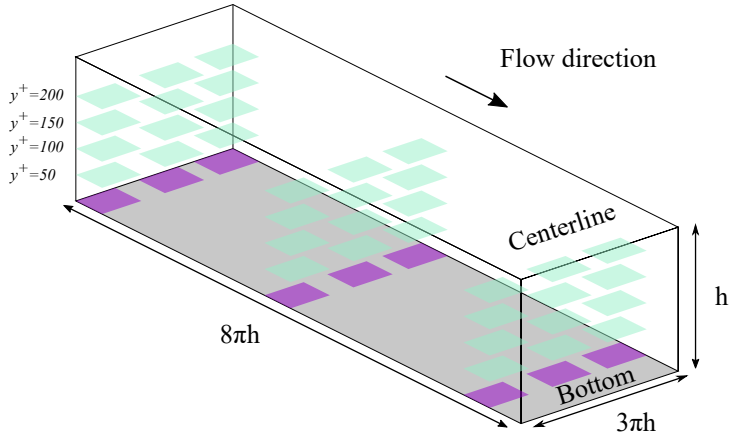

Figure 2: Schematic representation of the flow and wall fields extracted from the Channel Flow DNS spatial domain. Axes are not equally scaled.

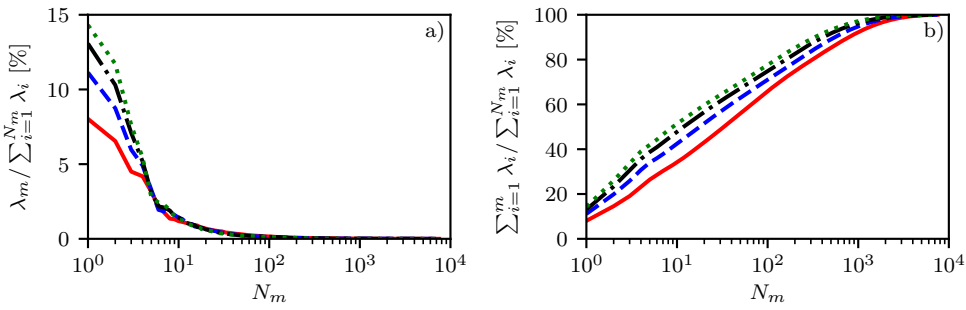

Figure 3: a) Energy spectrum and b) cumulative energy contribution of the flow field modes of the channel simulation. Colours refer to the wall distance position of the flow fields: $y^+ = 50$ (——), $y^+ = 100$ (– –), $y^+ = 150$ (–·–·), $y^+ = 200$ (·····).

3. Results

This section presents the LSM reconstruction with 10 POD modes using EPOD and CNN techniques. The main implementation difference is that EPOD is able to estimate simultaneously the time coefficients of all the POD modes, while a

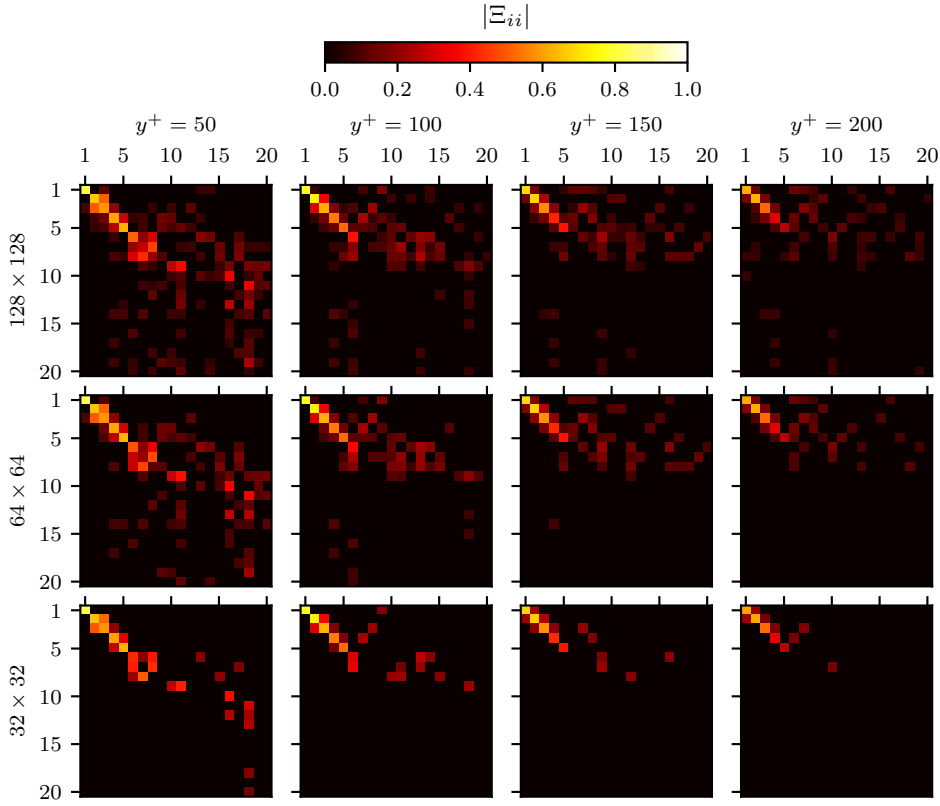

Figure 4: Time correlation coefficient matrix $\underline{\underline{\Xi}}$ for EPOD estimation. Columns report data at fixed wall distance y^+. Rows report data at fixed wall-shear-stress resolution.

CNN needs a separate training for each mode. From the practical viewpoint, EPOD and CNN work with the same input. As explained in Sec. 2 the input data are the streamwise wall-shear-stress fields, while the predicted output is the first 10 POD time coefficients. As argued in the previous section, this is normally sufficient to detect the large-scale motions. Having available the target fields to be reconstructed, the actual temporal POD modes of each field can be computed and it is possible to employ the low-order reconstruction of the fields obtained with the first 10 modes as reference field.

3.1. *A-priori assessment*

Both methods (EPOD and CNN) allow an *a priori* assessment of the reconstruction ability, without having to rely on a direct comparison with a ground

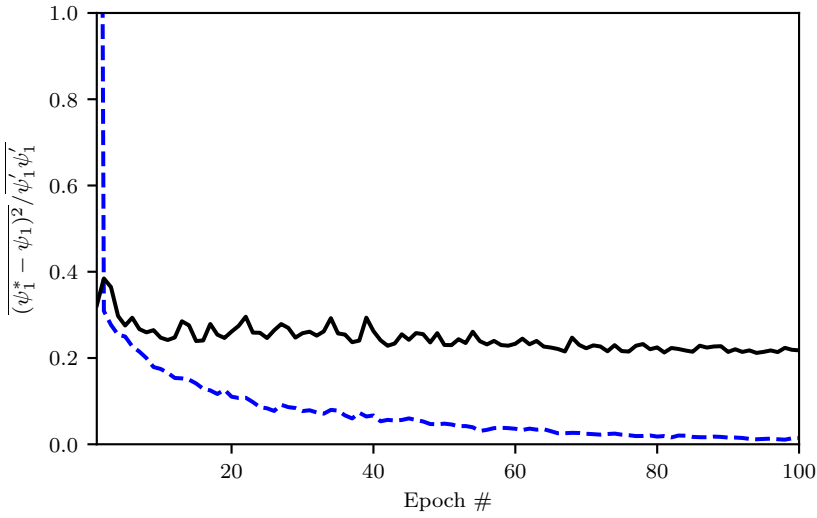

Figure 5: Mean-squared-error evolution for the training (– –) and validation (——) datasets for the first temporal POD-mode coefficient in the case H050 (see table 1).

truth, which is of course not available in experiments. The estimation performances of the EPOD can be deduced for instance from the temporal correlation coefficient matrix $\underline{\underline{\Xi}}$. Fig. 4 shows the $\Xi_{i,i}$ values for the first 20 temporal POD modes for all the tested wall-distance positions and wall-shear-stress-field resolutions. It can be seen that the matrices exhibit a diagonal dominance in the north-west corner, which corresponds to a direct relation between the most energetic modes of flow fields and wall-shear stress. While this dominance is not a necessary condition for high-quality reconstruction, it is an indicator of a "healthy" correspondence of the most energetic features in flow fields and probe data. This is needed to achieve a compact and efficient reconstruction, since the energy of each mode is not spread over a range of modes. However, this dominance disappears starting from the 5th-6th mode. Moreover, this correlation between low-frequency features in the flow and low-frequency features at the wall impedes the EPOD to capture the amplitude modulation effects that the LSMs produce over high-frequency near-wall scales (Mathis *et al.* 2009). The wall-distance is the most relevant parameter; as expected, the peak of the correlation coefficient of the first modes is larger for the flow fields located closer to the wall. On the other hand, the resolution level of the wall measurements seems to have a minor effect, which could indicate that coarser meshes could be employed without significantly affecting the estimation quality. This is clearly an important indication to simplify the practical implementation of this method in experiments.

For CNNs, the validation error is a directly interpretable indicator of the quality of the reconstruction. Figure 5 shows the evolution of both training and validation datasets for the first mode of the H050 case (see table 1). During training, a callback function was used initially to select the number of iterations, also called epochs, based on the mean-square error evolution of the validation dataset. However, after the first training it was noticed that the validation mean-square error had a significant scatter and a low rate of decrease, while the training error at the stopping epoch was still clearly descending for increasing number of epochs. Because of this, it was decided to fix the number of epochs equal to 100.

An important metric for comparison is the training time of EPOD and CNN. Both data-driven techniques have been trained in a standard workstation (Intel® Core™ i9-8950K CPU @ 2.9 GHz with NVIDIA GeForce® GTX 1050 Ti GPU with Max-Q design). EPOD has required 7 minutes to train case M050, while CNNs have required 44 minutes for the training of the same case. It has to be reminded that 10 different CNNs have been trained, one for each of the first 10 first temporal POD modes. Considering the required time to train the CNNs for a single temporal POD mode, it is reduced to 4 minutes. The larger computational cost required by CNNs is counterbalanced by the possibility of tackling non-linear effects related with the LSMs present in the flow. The estimation time for each field is of the order of $0.01s$ per snapshot for CNN, while it is practically real-time for EPOD. This is important to allow the possibility of active closed-loop flow control based on flow field estimation. However, it has to be remarked that the selected architecture is not conceived for direct implementation in active flow control, and significant improvements could be obtained employing dedicated high-performance hardware.

3.2. *LSMs-reconstruction performance*

For the *a posteriori* assessments of the LSM reconstruction techniques, several metrics have been used. The first one is the squared correlation coefficient R^2 of the temporal mode coefficients, which quantifies the quality of the reconstruction of ψ_i^* before reconstructing the flow fields. It is defined as:

$$R^2_{\psi_i} = \frac{\overline{\psi_i^* \psi_i}^2}{\overline{\psi_i^{*2}}\,\overline{\psi_i^2}}. \tag{11}$$

The squared correlation coefficient is reported in figure 6 for both EPOD and CNN techniques. Here and in the following, ensemble-averaged values are indicated as $\bar{\cdot}$. For colour code, the reader is referred to table 1. CNN abundantly surpasses EPOD in reconstruction quality for all the analyzed cases, independently of wall-distance and wall-shear-stress resolution. The squared correlation coefficient achieved by EPOD for the case H50 (the one closer to the wall and with higher resolution, thus the most favourable case) is lower than the value reached by CNN for case L200 (the most challenging case, further

from the wall and with lower resolution) for all the first 10 modes. Additionally, EPOD shows a steep drop in performance beyond the third POD mode in most cases, while CNNs are significantly more robust. While the analysis here is limited to the first 10 POD modes due to the available computational resources, it can be expected that CNNs can pave the way to successful reconstruction of also higher-order modes.

Overall, the correlation coefficient spotlights the limits of EPOD for LSM detection. This is most likely to be addressed to non-linear effects of LSMs on the wall signature, which are not modelled by EPOD and become more intense when moving away from the wall. CNNs exhibit a negligible decrease of the squared correlation coefficient when going away from the wall, thus extending the capability of flow-field estimation from the wall signature. However, this effect could be an artifact of the number of selected modes. The energy content extracted with the first 10 POD modes is different at each distance from the wall, as shown in figure 3. In the regions closer to the wall ($y^+ = 50$) the first modes do not represent the LSMs but the near-wall scales affected by them. Thus, the reconstruction of the full instantaneous fields would feel a significant larger effect of the wall distance on the reconstruction performance. In terms of wall-shear-stress resolution, CNNs show a weak sensitivity, though stronger than EPOD. This can be probably ascribed to the fact that the non-linear effects require larger datasets in order to be adequately characterized. This trend can be explained considering that reducing the resolution would actually tend to the asymptotic result corresponding to a 2×2 resolution which would correspond to a simple MLP. Baldi & Hornik (1989) showed that this neural-network architecture produces the same output as POD. Consequently low-resolution configurations were not explored also because, with such low-resolution inputs, it would be much more interesting to employ temporal signals as input in both EPOD (as shown by Discetti *et al.* (2019) and Discetti *et al.* (2018)) and ANNs (using long short-term memory for example (Deng *et al.* 2019)).

Figure 7 reports the reconstruction of the streamwise velocity fluctuation fields at $y^+ = 150$ using the largest wall-shear-stress resolution for the estimation. This wall distance was chosen since it has been proven that the geometrical center of the logarithmic region (defined as $y^+ = 3.9\sqrt{Re_\tau}$) is the point where high-momentum LSMs start providing a dominating effect over the near-wall small-scale fluctuations (Mathis *et al.* 2009) (while for smaller wall distances low-momentum events are dominant). The streamwise velocity fluctuations at this point have been used in previous studies to predict near-wall (Marusic *et al.* 2010; Mathis *et al.* 2011) and wall-shear-stress (Mathis *et al.* 2013) turbulent statistics. In figure 7 an instantaneous field is compared with the low-order reconstruction performed with EPOD and with that performed with CNNs. The exact reconstruction using the first 10 modes is included as ground truth for reference. Visual inspection shows that the CNNs are able to capture the high- and low-momentum regions of the turbulent field with high accuracy. On

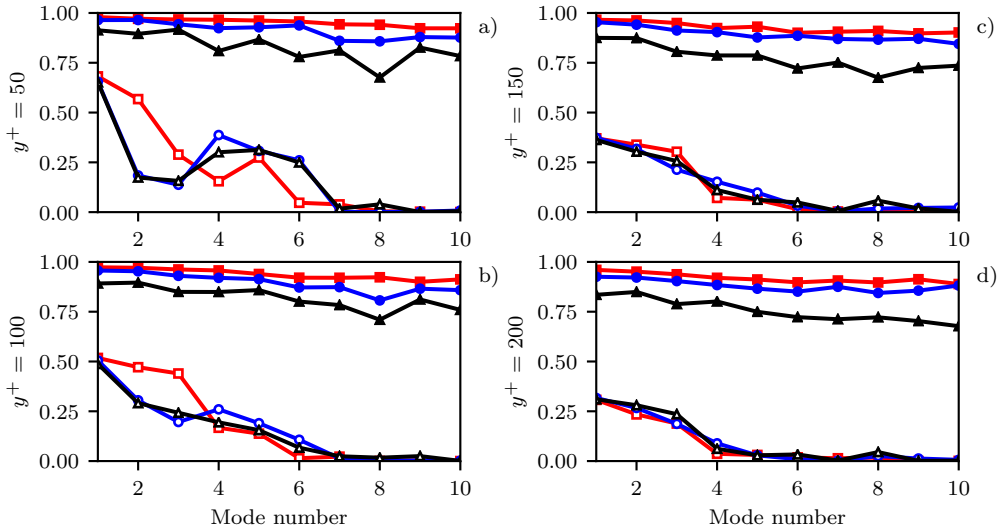

Figure 6: Squared correlation coefficient $R^2_{\psi_i}$ of the reconstructed time coefficients. From top to bottom, wall-distance positions correspond to $y^+ = [50, 100, 150, 200]$. For colour code, see table 1.

the other hand, as expectable from the low R^2 values, EPOD reconstruction does not capture any significant flow feature.

A quantitative analysis is carried out by defining the error in the reconstruction of the instantaneous velocity fields as the root mean squared error of the fields scaled with the streamwise Reynolds stress:

$$error = \frac{\sqrt{\overline{(u'^* - u')^2}}}{\sigma_{u'}}. \tag{12}$$

The errors reported in figure 8 show that the reconstruction quality of the instantaneous fields with EPOD is rather poor, reaching levels around 100% of streamwise Reynolds stress deviation. The error does not exhibit any significant dependence with the wall-shear resolution. Regarding CNN, it can be seen that wall-shear-stress resolution has a significant larger impact than the wall-distance position, reaching a 40% error in the L050 case. This result is in agreement with the behaviour of squared correlation coefficient shown in figure 6.

The inner-scaled streamwise Reynolds stresses estimated in each of the four planes, accounting for the reconstructed LSM velocity fields, are reported in figure 8. It can be observed that CNNs provided a reasonably good estimation

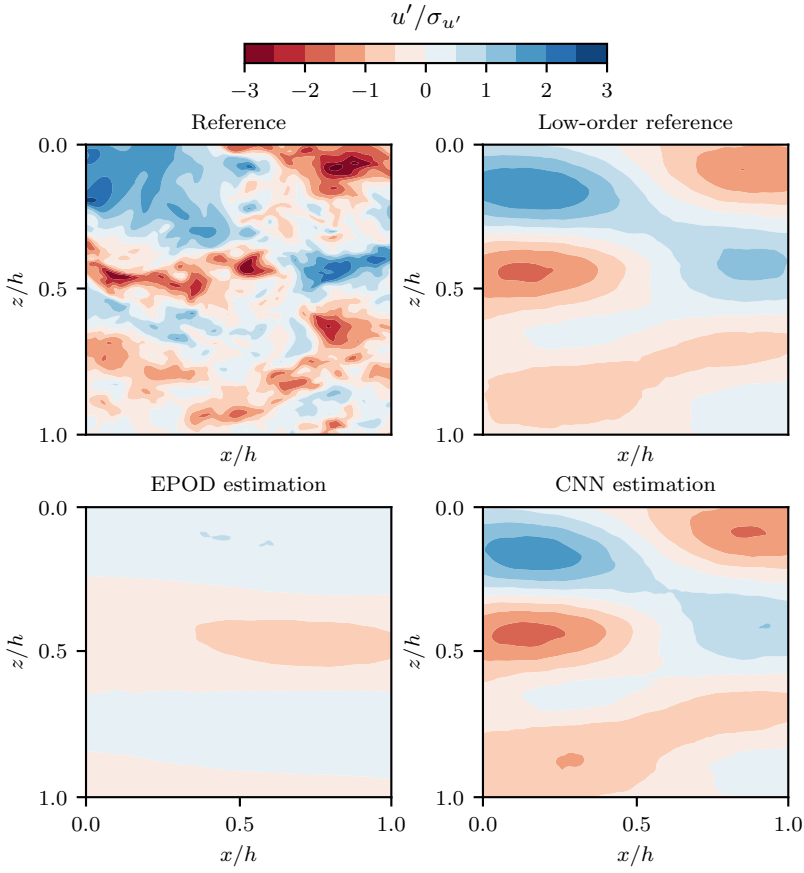

Figure 7: Instantaneous streamwise velocity fluctuations fields reconstructed with 10 POD modes for case H150. The first row shows the instantaneous field and the reference low-order reconstruction with 10 modes. The second row shows, from left to right, EPOD reconstruction, and CNN reconstruction. All fields have been scaled with the streamwise velocity fluctuation standard deviation $\sigma_{u'}$ of the nominal case.

independently of the wall-shear-stress resolution. Similarly, the turbulent statistics computed with EPOD are not affected by the wall-shear-stress resolution, which is in line with the results of figure 6. EPOD is significantly underestimating the streamwise Reynolds stress values, with this effect being stronger at larger distances from the wall. This result is not surprising, since EPOD is a linear operator which is expected to be severely challenged when the relation between probe and field fluctuations become strongly non-linear.

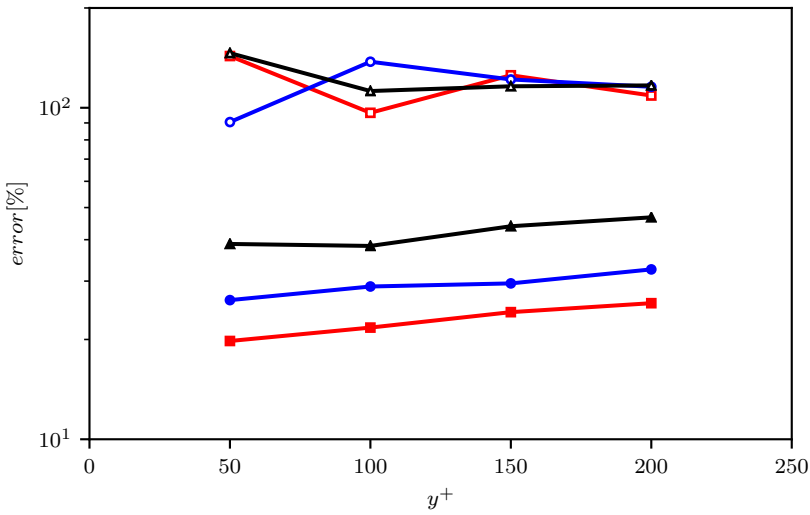

Figure 8: Mean error (according to Eq. 12) of the instantaneous streamwise velocity fluctuation fields reconstructed with the first 10 POD modes versus wall-normal positions. For colour code, see table 1.

4. Conclusions

The ability of two different data-driven techniques to reconstruct the LSMs present in a turbulent channel flow using wall-shear stress measurements has been assessed. Both techniques have been used to perform a low-order estimation within the POD basis described by the flow field measurements. The effect of the wall-shear-stress field resolution and of the wall-distance of the plane to be reconstructed have been addressed.

The extended proper orthogonal decomposition (EPOD) has shown a limited capacity to reconstruct the streamwise Reynolds stress of the LSMs. The reconstruction performance is severely affected by the wall distance, a limit that can be addressed to impossibility to model non-linearity with EPOD. The reconstruction error reaches levels of 100% streamwise Reynolds stress deviation, with a corresponding significant underestimation of the streamwise velocity fluctuations of LSMs.

Convolutional neural networks (CNNs) have shown an outstanding capability to reconstruct LSMs in terms of turbulent statistics and instantaneous flow fields. The reconstruction performance is only weakly affected by the wall distance and the wall-shear-stress resolution. Nevertheless, this effect could be addressed to the number of selected POD modes, as long as their energy content varies with wall distance. In all tested cases the reconstruction error is significantly lower than the error achieved by EPOD. It is impressive that the

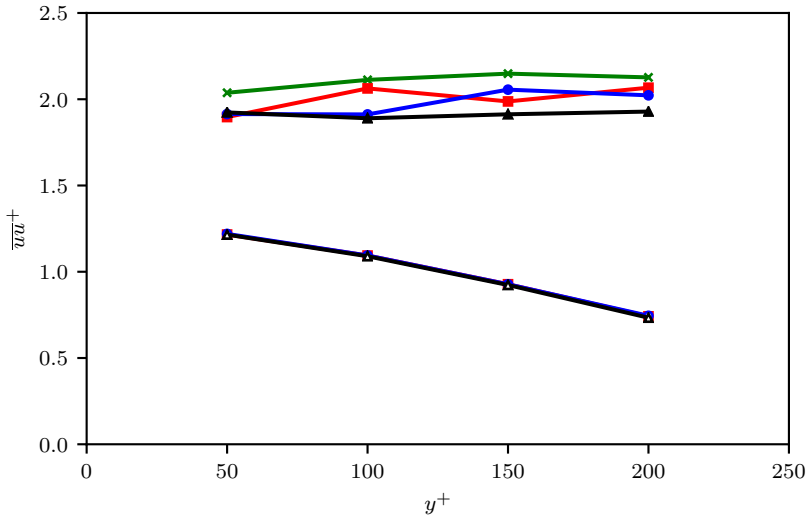

Figure 9: Reconstruction with 10 POD modes of inner-scaled streamwise velocity fluctuations versus the wall-normal positions. For colours, see table 1. Nominal POD reconstruction is shown as (-×-).

reconstruction error of CNN of the most challenging case (minimum wall-shear resolution and maximum wall distance) is about one half of the error of the simplest case for EPOD (i.e. maximum wall-shear resolution and minimum wall distance). This opens the path to more robust estimation of flow fields from wall-data.

It would be interesting, in future investigations, to analyze the capabilities of CNNs to reconstruct the LSMs in a flow with a different Reynolds number to that of the training dataset. This would enable flow-field reconstructions based on a single computational or experimental training. One key enabler is that POD spatial modes have been shown to be independent of the Reynolds number (Liu *et al.* 1994). In the same way, it is worth analyzing whether the reconstruction algorithms trained in a turbulent channel flow would be able to obtain similar results in other canonical flows, i.e. turbulent-boundary-layer or pipe flows, suggesting a universal mechanism of non-linear interactions in wall-bounded flows. Both purposes could be tackled using the transfer-learning technique, that allows to transfer knowledge acquired from a solved task, like channel-flow reconstruction, to a new one, such as the reconstruction of a pipe flow. For example, Guastoni *et al.* (2020*a*) has shown that transfer learning can be used to reduce the computational cost of training a CNN for the reconstruction at a new wall distance by using a previously-trained CNN at a different wall distance.

With respect to a future experimental set-up to test the capabilities of CNNs to reconstruct LSMs, recent investigations have shown that it is possible to obtain time-resolved heat-transfer maps with resolutions of the order of 100×100 (Nakamura & Yamada 2013; Raiola *et al.* 2017; Mallor *et al.* 2019). Whereas the relationship between heat-transfer maps and wall-shear-stress measurements have been shown by Mayer *et al.* (1998), it might be expected that this input resolution is enough for CNNs to provide accurate LSM reconstructions.

Acknowledgements

This work has been partially supported by the Grant DPI2016-79401-R funded by the Spanish State Research Agency (SRA) and European Regional Development Fund (ERDF). AG acknowledges Dr. A. Sánchez for insightful discussions about CNN architecture. The authors acknowledge Dr. R. Vinuesa for insightful comments and discussions.

References

Abbassi, M., Baars, W., Hutchins, N. & Marusic, I. 2017 Skin-friction drag reduction in a high-Reynolds-number turbulent boundary layer via real-time control of large-scale structures. *Int. J. Heat Fluid Fl.* **67**, 30–41.

Adrian, R. J. 1996 Stochastic estimation of the structure of turbulent fields. In *Eddy structure identification*, pp. 145–195. Springer.

Baars, W. J., Hutchins, N. & Marusic, I. 2016 Spectral stochastic estimation of high-Reynolds-number wall-bounded turbulence for a refined inner-outer interaction model. *Phys. Rev. Fluids* **1** (5), 054406.

Bai, H., Zhou, Y., Zhang, W., Xu, S., Wang, Y. & Antonia, R. 2014 Active control of a turbulent boundary layer based on local surface perturbation. *J. Fluid Mech.* **750**, 316–354.

Baldi, P. & Hornik, K. 1989 Neural networks and principal component analysis: Learning from examples without local minima. *Neural Networks* **2** (1), 53–58.

Blackwelder, R. F. & Kovasznay, L. S. 1972 Time scales and correlations in a turbulent boundary layer. *Phys. Fluids* **15** (9), 1545–1554.

Borée, J. 2003 Extended proper orthogonal decomposition: a tool to analyse correlated events in turbulent flows. *Exp. Fluids* **35** (2), 188–192.

Brunton, S. L., Noack, B. R. & Koumoutsakos, P. 2020 Machine learning for fluid mechanics. *Annu. Rev. Fluid Mech.* **52**, 477–508.

Deng, Z., Chen, Y., Liu, Y. & Kim, K. C. 2019 Time-resolved turbulent velocity field reconstruction using a long short-term memory (LSTM)-based artificial intelligence framework. *Phys. Fluids* **31** (7), 075108.

Discetti, S., Bellani, G., Örlü, R., Serpieri, J., Sanmiguel Vila, C., Raiola, M., Zheng, X., Mascotelli, L., Talamelli, A. & Ianiro, A. 2019 Characterization of very-large-scale motions in high-Re pipe flows. *Exp. Therm. Fluid Sci.* **104**, 1–8.

Discetti, S., Raiola, M. & Ianiro, A. 2018 Estimation of time-resolved turbulent fields through correlation of non-time-resolved field measurements and time-resolved point measurements. *Exp. Therm. Fluid Sci.* **93**, 119–130.

Encinar, M. P., Lozano-Durán, A. & Jiménez, J. 2018 Reconstructing channel turbulence from wall observations. In *Procs. CTR Summer School*.

Favre, A., Gaviglio, J. & Dumas, R. 1967 Structure of velocity space-time correlations in a boundary layer. *Phys. Fluids* **10** (9), S138–S145.

Fukami, K., Fukagata, K. & Taira, K. 2019 Super-resolution reconstruction of turbulent flows with machine learning. *J. Fluid Mech.* **870**, 106–120.

Fukushima, K. 1980 Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biol. Cybern.* **36** (4), 193–202.

Fukushima, K. 1988 Neocognitron: A hierarchical neural network capable of visual pattern recognition. *Neural Networks* **1** (2), 119–130.

Ganapathisubramani, B., Hutchins, N., Monty, J., Chung, D. & Marusic, I. 2012 Amplitude and frequency modulation in wall turbulence. *J. Fluid Mech.* **712**, 61–91.

Gouder, K., Potter, M. & Morrison, J. F. 2013 Turbulent friction drag reduction using electroactive polymer and electromagnetically driven surfaces. *Exp. Fluids* **54** (1), 1441.

Guala, M., Hommema, S. & Adrian, R. 2006 Large-scale and very-large-scale motions in turbulent pipe flow. *J. Fluid Mech.* **554**, 521–542.

Guastoni, L., Encinar, M. P., Schlatter, P., Azizpour, H. & Vinuesa, R. 2020 Prediction of wall-bounded turbulence from wall quantities using convolutional neural networks. *J. Phys. Conf. Ser.* **1522** (1), 012022.

Hack, M. P. & Zaki, T. A. 2016 Data-enabled prediction of streak breakdown in pressure-gradient boundary layers. *J. Fluid Mech.* **801**, 43–64.

Hoyas, S. & Jiménez, J. 2006 Scaling of the velocity fluctuations in turbulent channels up to $Re_\tau = 2003$. *Phys. Fluids* **18** (1), 011702.

Hoyas, S. & Jiménez, J. 2008 Reynolds number effects on the Reynolds-stress budgets in turbulent channels. *Phys. Fluids* **20** (10), 101511.

Hutchins, N. & Marusic, I. 2007*a* Evidence of very long meandering features in the logarithmic region of turbulent boundary layers. *J. Fluid Mech.* **579**, 1–28.

Hutchins, N. & Marusic, I. 2007*b* Large-scale influences in near-wall turbulence. *Philos. T. R. Soc. A* **365** (1852), 647–664.

Illingworth, S. J., Monty, J. P. & Marusic, I. 2018 Estimating large-scale structures in wall turbulence using linear models. *J. Fluid Mech.* **842**, 146–162.

Jiménez, J. & Pinelli, A. 1999 The autonomous cycle of near-wall turbulence. *J. Fluid Mech.* **389**, 335–359.

Kim, K. & Adrian, R. 1999 Very large-scale motion in the outer layer. *Phys. Fluids* **11** (2), 417–422.

Kingma, D. P. & Ba, J. 2014 Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.

Kline, S. J., Reynolds, W. C., Schraub, F. & Runstadler, P. 1967 The structure of turbulent boundary layers. *J. Fluid Mech.* **30** (4), 741–773.

LeCun, Y. 1985 A learning scheme for asymmetric threshold networks. In *Proc. Cognitiva*, , vol. 85, pp. 599–604.

LeCun, Y., Bengio, Y. & Hinton, G. 2015 Deep learning. *Nature* **521** (7553), 436.

LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard,

W. & Jackel, L. D. 1989 Backpropagation applied to handwritten zip code recognition. *Neural Comput.* **1** (4), 541–551.

Lee, C., Kim, J., Babcock, D. & Goodman, R. 1997 Application of neural networks to turbulence control for drag reduction. *Phys. Fluids* **9** (6), 1740–1747.

Li, Y., Perlman, E., Wan, M., Yang, Y., Meneveau, C., Burns, R., Chen, S., Szalay, A. & Eyink, G. 2008 A public turbulence database cluster and applications to study lagrangian evolution of velocity increments in turbulence. *J. Turbul.* **9**, N31.

Ling, J., Kurzawski, A. & Templeton, J. 2016 Reynolds averaged turbulence modelling using deep neural networks with embedded invariance. *J. Fluid Mech.* **807**, 155–166.

Liu, W., Wang, Z., Liu, X., Zeng, N., Liu, Y. & Alsaadi, F. E. 2017 A survey of deep neural network architectures and their applications. *Neurocomputing* **234**, 11–26.

Liu, Z., Adrian, R. & Hanratty, T. 2001 Large-scale modes of turbulent channel flow: transport and structure. *J. Fluid Mech.* **448**, 53–80.

Liu, Z.-C., Adrian, R. J. & Hanratty, T. J. 1994 Reynolds number similarity of orthogonal decomposition of the outer layer of turbulent wall flow. *Phys. Fluids* **6** (8), 2815–2819.

Lumley, J. L. 1967 The structure of inhomogeneous turbulent flows. In *Atmospheric turbulence and radio wave propagation*, pp. 166–178.

Lusch, B., Kutz, J. N. & Brunton, S. L. 2018 Deep learning for universal linear embeddings of nonlinear dynamics. *Nat. Commun.* **9** (1), 4950.

Mallor, F., Raiola, M., Sanmiguel Vila, C., Örlü, R., Discetti, S. & Ianiro, A. 2019 Modal decomposition of flow fields and convective heat transfer maps: An application to wall-proximity square ribs. *Exp. Therm. Fluid Sci.* **102**, 517–527.

Marusic, I., Mathis, R. & Hutchins, N. 2010 Predictive model for wall-bounded turbulent flow. *Science* **329** (5988), 193–196.

Mathis, R., Hutchins, N. & Marusic, I. 2009 Large-scale amplitude modulation of the small-scale structures in turbulent boundary layers. *J. Fluid Mech.* **628**, 311–337.

Mathis, R., Hutchins, N. & Marusic, I. 2011 A predictive inner–outer model for streamwise turbulence statistics in wall-bounded flows. *J. Fluid Mech.* **681**, 537–566.

Mathis, R., Marusic, I., Chernyshenko, S. I. & Hutchins, N. 2013 Estimating wall-shear-stress fluctuations given an outer region input. *J. Fluid Mech.* **715**, 163–180.

Maulik, R. & San, O. 2017 A neural network approach for the blind deconvolution of turbulent flows. *J. Fluid Mech.* **831**, 151–181.

Maulik, R., San, O., Rasheed, A. & Vedula, P. 2019 Subgrid modelling for two-dimensional turbulence using neural networks. *J. Fluid Mech.* **858**, 122–144.

Mayer, R., Henkes, R. A. W. M. & Van Ingen, J. L. 1998 Quantitative infrared-thermography for wall-shear stress measurement in laminar flow. *Int. J. Heat Mass Tran.* **41** (15), 2347–2356.

Milano, M. & Koumoutsakos, P. 2002 Neural network modeling for near wall turbulent flow. *J. Comput. Phys.* **182** (1), 1–26.

Monty, J., Stewart, J., Williams, R. & Chong, M. 2007 Large-scale features in turbulent pipe and channel flows. *J. Fluid Mech.* **589**, 147–156.

Nair, V. & Hinton, G. E. 2010 Rectified linear units improve restricted boltzmann machines. In *Proc. ICML*, pp. 807–814.

Nakamura, H. & Yamada, S. 2013 Quantitative evaluation of spatio-temporal heat transfer to a turbulent air flow using a heated thin-foil. *Int. J. Heat Mass Tran.* **64**, 892–902.

Nikolaou, Z., Chrysostomou, C., Vervisch, L. & Cant, S. 2019 Progress variable variance and filtered rate modelling using convolutional neural networks and flamelet methods. *Flow Turbul. Combust.* **103**, 1–17.

Ou, J. & Li, Y. 2019 Vector-kernel convolutional neural networks. *Neurocomputing* **330**, 253–258.

Raiola, M., Greco, C. S., Contino, M., Discetti, S. & Ianiro, A. 2017 Towards enabling time-resolved measurements of turbulent convective heat transfer maps with ir thermography and a heated thin foil. *Int. J. Heat Mass Tran.* **108**, 199–209.

Raissi, M., Wang, Z., Triantafyllou, M. S. & Karniadakis, G. E. 2019 Deep learning of vortex-induced vibrations. *J. Fluid Mech.* **861**, 119–137.

Rosenblatt, F. 1958 The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol. Rev.* **65** (6), 386.

Sanmiguel Vila, C., Örlü, R., Vinuesa, R., Schlatter, P., Ianiro, A. & Discetti, S. 2017 Adverse-pressure-gradient effects on turbulent boundary layers: statistics and flow-field organization. *Flow Turbul. Combust.* **99** (3-4), 589–612.

Sasaki, K., Vinuesa, R., Cavalieri, A. V., Schlatter, P. & Henningson, D. S. 2019 Transfer functions for flow predictions in wall-bounded turbulence. *J. Fluid Mech.* **864**, 708–745.

Sirovich, L. 1987 Turbulence and the dynamics of coherent structures. I. Coherent structures. *Q. Appl. Math.* **45** (3), 561–571.

Srinivasan, P., Guastoni, L., Azizpour, H., Schlatter, P. & Vinuesa, R. 2019 Predictions of turbulent shear flows using deep neural networks. *Phys. Rev. Fluids* **4** (5), 054603.

Suzuki, T. & Hasegawa, Y. 2017 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.

Tinney, C., Coiffet, F., Delville, J., Hall, A., Jordan, P. & Glauser, M. N. 2006 On spectral linear stochastic estimation. *Exp. Fluids* **41** (5), 763–775.

Tinney, C. E., Ukeiley, L. & Glauser, M. N. 2008 Low-dimensional characteristics of a transonic jet. Part 2. Estimate and far-field prediction. *J. Fluid Mech.* **615**, 53–92.

Wark, C. E. & Nagib, H. 1991 Experimental investigation of coherent structures in turbulent boundary layers. *J. Fluid Mech.* **230**, 183–208.

Wu, J.-L., Wang, J.-X., Xiao, H. & Ling, J. 2017 A priori assessment of prediction confidence for data-driven turbulence modeling. *Flow Turbul. Combust.* **99** (1), 25–46.

Wu, Y. 2014 A study of energetic large-scale structures in turbulent boundary layer. *Phys. Fluids* **26** (4), 045113.

Yu, H., Kanov, K., Perlman, E., Graham, J., Frederix, E., Burns, R., Szalay, A., Eyink, G. & Meneveau, C. 2012 Studying lagrangian dynamics of turbulence using on-demand fluid particle tracking in a public turbulence database. *J. Turbul.* **13**, N12.

Paper 2

2

Convolutional-network models to predict wall-bounded turbulence from wall quantities

Luca Guastoni[1,2], Alejandro Güemes[3], Andrea Ianiro[3], Stefano Discetti[3], Phillip Schlatter[1,2], Hossein Azizpour[4,2] and Ricardo Vinuesa[1,2]

[1] Linné FLOW Centre, KTH Mechanics, S-100 44 Stockholm, Sweden
[2] Swedish e-Science Research Centre (SeRC), Stockholm, Sweden
[3] Aerospace Research Group, Universidad Carlos III de Madrid, Leganés, Spain
[4] Division of Robotics, Perception and Learning, KTH Royal Institute of Technology, Stockholm, Sweeden

Journal of Fluid Mechanics, ***Under review*** (2021)

Two models based on convolutional neural networks are trained to predict the two-dimensional velocity-fluctuation fields at different wall-normal locations in a turbulent-open channel flow, using the wall-shear-stress components and the wall pressure as inputs. The first model is a fully-convolutional neural network (FCN) which directly predicts the fluctuations, while the second one reconstructs the flow fields using a linear combination of orthonormal basis functions, obtained through proper orthogonal decomposition(POD), hence named FCN-POD. Both models are trained using data from two direct numerical simulations (DNS) at friction Reynolds numbers $Re_\tau = 180$ and 550. Thanks to their ability to predict the nonlinear interactions in the flow, both models show a better prediction performance than the extended proper orthogonal decomposition (EPOD), which establishes a linear relation between input and output fields. The performance of the various models is compared based on predictions of the instantaneous fluctuation fields, turbulence statistics and power-spectral densities. The FCN exhibits the best predictions closer to the wall, whereas the FCN-POD model provides better predictions at larger wall-normal distances. We also assessed the feasibility of performing transfer learning for the FCN model, using the weights from $Re_\tau = 180$ to initialize those of the $Re_\tau = 550$ case. Our results indicate that it is possible to obtain a performance similar to that of the reference model up to $y^+ = 50$, with 50% and 25% of the original training data. These non-intrusive sensing models will play an important role in applications related to closed-loop control of wall-bounded turbulence

Key words: turbulent flows, turbulence simulation

1. Introduction

In this work, we assess the potential of deep neural networks (DNNs, see LeCun *et al.* 2015) to perform *non-intrusive sensing*, that consists in using measurable quantities in a fluid flow to reconstruct its behaviour in another location in the domain or to predict its dynamics in the future, without using probes that affect the flow itself. For instance, it is possible to accurately measure time-resolved quantities at the wall, such as the wall-shear stress or the pressure, and then correlate these measurements with the flow farther away. A reliable flow estimation is typically a prerequisite for closed-loop control applications, where the actuation is applied with the aim of suppressing the effect of certain structures in the flow (Choi *et al.* 1994). In order to effectively perform closed-loop control it is necessary to monitor the instantaneous state of the flow so as to devise the best way to affect it, however characterizing the flow state can be extremely challenging, particularly at very high Reynolds numbers where the near-wall structures become progressively smaller.

Before the appearance of DNNs, flow-field predictions were performed mainly through linear methods. Among them, the linear stochastic estimation (LSE) introduced by Adrian (1988) stands out. Recently Suzuki & Hasegawa (2017) and Encinar & Jiménez (2019) have used LSE to reconstruct the velocity field on a wall-parallel plane in a turbulent channel flow employing wall measurements. The latter study showed that LSE can only reconstruct the large wall-attached eddies in the outer part of the logarithmic region. An extension of the LSE method in the spectral domain (Tinney *et al.* 2006) was shown to be more suitable for noisy predictions in turbulent flows. More recently Baars & Tinney (2014) proposed a method based on proper orthogonal decomposition (POD) to improve the spectral-LSE approach. Borée (2003) reported the possibility of projecting a synchronized field on the POD temporal modes of another quantity; this method is known as extended POD (EPOD). The correlation matrix between the temporal POD coefficients of two given quantities can be used to predict one based on the other. The work of Borée (2003) proved EPOD to be equivalent to LSE when all modes are retained in the reconstruction. Using remote probes, EPOD has been used to provide predictions in turbulent jets (Tinney *et al.* 2008), channel flows (Discetti *et al.* 2018), pipe flows (Discetti *et al.* 2019) and wall-mounted obstacles (Hosseini *et al.* 2016; Bourgeois *et al.* 2013). Note however that in the latter work quadratic terms are included in the model of the POD coefficient dynamics, hinting at the need of non-linear estimation even for a relatively-simple, predominantly-oscillatory flow. In this regard, Sasaki *et al.* (2019) recently assessed the capabilities of both linear and non-linear transfer functions with single and multiple inputs to provide turbulent-flow predictions. They documented a significant improvement in the predictions when the transfer functions were designed to account for nonlinear interactions between the inputs and the flow field. The improved prediction capabilities of nonlinear methods over linear ones were also reported by Mokhasi *et al.* (2009) and Nair & Goza (2020).

DNNs are non-linear models that have found application in many research areas (Jean *et al.* 2016; De Fauw *et al.* 2018; Norouzzadeh *et al.* 2018; Ham *et al.* 2019; Udrescu & Tegmark 2020; Vinuesa *et al.* 2020). Due to their potential applications in flow modelling, identification of turbulence features and flow control, DNNs have recently received extensive attention in the fluid-mechanics research community (Kutz 2017; Jiménez 2018; Duraisamy *et al.* 2019; Brunton *et al.* 2020). Here we provide a brief overview of the recent neural network applications in fluid mechanics, before coming back to the flow estimation problem we investigated in this work. In the case of turbulence modelling, DNNs have been reported to improve the results of Reynolds-averaged Navier–Stokes (RANS, Ling *et al.* 2016; Wu *et al.* 2017) models and large-eddy simulations (LES, Maulik *et al.* 2019; Lapeyre *et al.* 2019; Beck *et al.* 2019). There are also a number of on-going efforts towards including the constraints from the Navier–Stokes equations into prediction models through the so-called physics-informed neural networks (Wang *et al.* 2017; Raissi *et al.* 2019). Furthermore, several artificial-intelligence-based solutions have been proposed to perform optimal control on different types of flows, such as the wake behind one or more cylinders (Rabault *et al.* 2019; Raibaudo *et al.* 2020). Other promising applications of machine learning to fluid mechanics include generation of inflow conditions (Fukami *et al.* 2019*b*) and extraction of flow patterns (Raissi *et al.* 2020).

DNNs have also been used in temporal prediction of dynamical systems. As an example, Srinivasan *et al.* (2019) compared the capabilities of the multi-layer perceptron (MLP, also known as fully-connected-layer neural network) and several long-short-term memory (LSTM) networks to predict the coefficients of a low-order model for near-wall turbulence (Moehlis *et al.* 2004). While the most relevant flow features are captured by both architectures, the LSTM network outperformed the MLP in terms of ability to predict turbulence statistics and the dynamics of the flow. This work has been extended by Eivazi *et al.* (2021), where the LSTM network has been compared with a Koopman-based framework which accounts for non-linearities through external forcing. Although both approaches provide accurate predictions of the dynamical evolution of the system, the latter outperforms the LSTM in terms of time and data required for training. Similar temporal predictions of the near-wall model (Moehlis *et al.* 2004) were conducted by Pandey *et al.* (2020) using echo state networks (ESN). Moreover, nonlinear autoregressive exogeneous networks (NARXs) have been used by Lozano-Durán *et al.* (2020) to exploit the relation between the temporal dynamics of the Fourier coefficients of a minimal turbulent channel flow. Their results showed accurate predictions of the bursting events in the logarithmic layer from buffer-region data. Other related work, in the context of control of the Kuramoto–Sivashinsky (KS) chaotic system, was recently conducted by Bucci *et al.* (2019). Note however that the use of temporal sequences implies a high computational cost to generate well-resolved temporal data. Furthermore, longer sequences require higher memory requirements in order to predict the future behaviour. For these reasons, several neural-network-based models that

learn spatial relations have been proposed in the literature. Convolutional neural networks (CNNs) have become increasingly popular during the last years due to the hierarchical structure of their input (Fukushima 1980, 1988; LeCun *et al.* 1989, 1998). For instance, Fukami *et al.* (2019, 2021) have shown that flow fields from the laminar wake of a cylinder and a turbulent channel can be reconstructed from extremely coarse data with remarkable success. CNNs have also been used to investigate the dynamical features of the flow without *a-priori* knowledge, as shown by Jagodinski *et al.* (2020).

Neural networks are mathematical models which exhibit very appealing properties, such as being *universal approximators* (Hornik *et al.* 1989). This means that they can potentially represent any continuous function with the adequate model parametrization, even though there is no guarantee that it is possible to infer the value of the parameters from data sampled from the original function. Nonetheless, neural-network parameters are typically tuned through data-driven training, and as such they have been compared and used together with other data-driven methods. For instance, the relationship between proper orthogonal decomposition (POD, see Lumley 1967) and the MLP is well documented in the literature (Bourlard & Kamp 1988; Baldi & Hornik 1989). These works showed that a MLP with a single hidden layer is equivalent to POD if a linear activation function is used. One of the early applications on a fluid-dynamics dataset was proposed by Milano & Koumoutsakos (2002), who compared the results of POD-based neural networks with linear and non-linear functions for the prediction of near-wall velocity fields, showing that nonlinear POD has significantly better predictive capabilities. More recently, the emergence of autoencoder architectures has motivated a renewed interest in the application of neural networks for dimensionality reduction. Hinton & Salakhutdinov (2006) proposed the use of deep autoencoders to obtain a low-order representation of high-dimensional data, showing that this approach is able to retain more information than POD. It is interesting to note that this work avoids the inherent difficulty of optimizing weights in deep autoencoders by training each layer with a Restricted Boltzmann Machine. Murata *et al.* (2020) used an autoencoder with convolutional layers to obtain a low-order representation of the flow around a cylinder. Their results suggest that CNN autoencoders with linear activation functions reproduce the same dimensionality reduction as POD, while the use of nonlinear activation functions improves the reconstruction performance. On a related note, flow reconstruction based on shallow neural networks was studied by Erichson *et al.* (2020) in several fluid-mechanics examples.

This work is not the first in which neural networks are used to perform non-intrusive sensing in wall-bounded flows: in a seminal study over 20 years ago, Lee *et al.* (1997) tested a two-layer neural network (with a single non-linear activation function) to predict the wall actuation, based on the wall-shear-stress components, in order to reduce the drag at the wall. Guzmán Iñigo *et al.* (2014) predicted the velocity field in a transitional boundary-layer flow using wall-shear

measurements, by projecting the velocity fields on a POD basis and using a dynamic observer to predict the dynamics of the flow. More recently, Kim & Lee (2020) used the two wall-shear-stress rcomponents to predict the instantaneous wall-normal heat flux with satisfactory results. The same wall information was used by Guastoni *et al.* (2020*a*) to predict the instantaneous streamwise flow fields at several wall-normal positions using convolutional networks. Their results show that these neural networks provide better predictions than linear methods (see below) in terms of instantaneous predictions and second-order statistics. The predictions were limited to one velocity component and to a low Reynolds number. In this work, all the velocity components are predicted and a higher Re is investigated as well. Furthermore, in the work by Güemes *et al.* (2019) the information of the most-energetic scales was encoded into a POD basis, and a CNN was used to predict that information at different wall-normal locations from streamwise wall-shear-stress measurements. Their results demonstrated that CNNs can significantly outperform linear methods in the prediction of POD time coefficients for low-order reconstruction of the velocity fields. Their convolutional network would predict only one POD temporal coefficient at a time, in this work all the coefficients are computed at the same time thanks to the implementation of an improved network architecture.

DNNs perform best when training and test data are taken from the same distribution, *i.e.* for the same flow and at the same Reynolds number in our case. However, in a real-world application the flow conditions will be continuously varying and/or it might be unfeasible to perform a full training at exactly the same conditions. If the flow behaviour at the Reynolds number of interest is roughly the same as the training Reynolds number and if the neural network can successfully approximate it, then the model should be able to perform consistently across the different Reynolds numbers, as shown in active flow control applications by Tang *et al.* (2020). In order to improve the performance at a different Reynolds number, it would then be possible to exploit initial training at a certain flow condition and transfer this knowledge to another one. Such knowledge transfer could reduce significantly the amount of data needed for training and improve the network applicability for industrial applications. Transfer learning (Pan & Yang 2009) is the suitable learning framework to address this issue, and it is discussed in detail below.

The methods proposed by Guastoni *et al.* (2020*a*) and Güemes *et al.* (2019), henceforth referred to as fully-convolutional network (FCN) and FCN-POD respectively, are extended in the present study. Both models are able to provide a nonlinear characterization of the relation between wall features and the flow on wall-parallel planes. The purpose of this work is to provide a detailed comparison of the two aforementioned nonlinear methods regarding their capabilities to predict turbulent flow fields from wall information. Their improvement over linear methods is measured using EPOD as a reference. Furthermore, transfer learning was applied to the FCN approach with the purpose of evaluating to

what extent a network trained at one Reynolds number can be used at a different one.

The remainder of this article is organised as follows. Section 2.1 describes the numerical databases used for training and testing the neural networks and §2.2 provides a brief comparison between the considered models. The FCN and FCN-POD methods are further detailed in §2.4 and §2.5 respectively, while EPOD is discussed in §2.3. Results from the considered prediction methods are presented and compared in §3, including instantaneous fields in §3.1, second-order statistics in §3.2, and spectra in §3.3. Furthermore, an assessment of transfer learning between different Reynolds numbers is presented in §4. Finally, the main conclusions of the work are presented in §5. Two Appendices are provided containing additional information about the training of the neural-network-based models and regarding the predicted instantaneous flow fields.

2. Methodology

2.1. *Datasets*

All the DNN variants in this study have been trained using the data generated from direct numerical simulations (DNS) of a turbulent open-channel flow. Periodic boundary conditions are imposed in the x- and z-directions (which are the streamwise and spanwise coordinates, respectively), and a no-slip condition is applied at the lower boundary ($y = 0$, where y is the wall-normal coordinate). Differently from a standard channel-flow simulation, a symmetry condition is imposed at the upper boundary. In standard channel flows, the wall-attached coherent structures may extend beyond the channel centerline, thus affecting the other wall (Lozano-Durán *et al.* 2012). On the other hand, in open-channel flows there is no upper wall. This makes the simulation more suitable to investigate to which extent the neural networks are able to learn the dynamics of near-wall turbulence, since the interaction of the large scales with both walls is not present.

The simulation is performed using the pseudo-spectral code SIMSON (Chevalier *et al.* 2007) with constant mass flow rate, in a domain $\Omega = L_x \times L_y \times L_z = 4\pi h \times h \times 2\pi h$ (where h is the channel height), as shown in figure 1. Two friction Reynolds numbers Re_τ (based on h and the friction velocity $u_\tau = \sqrt{\tau_w/\rho}$, where τ_w is the wall-shear stress and ρ is the fluid density) are considered, as summarized in table 1. The flow field is represented with N_y Chebyshev modes in the wall-normal direction and with N_x and N_z Fourier modes in the streamwise and spanwise directions, respectively. The instantaneous fields are obtained at constant time intervals from the time-advancing scheme, which is a second-order Crank–Nicholson scheme for the linear terms and a third-order Runge–Kutta method for the nonlinear terms. Dealiasing using a standard 3/2 rule is employed in the wall-parallel Fourier directions. The velocity fields to be used as ground truth for training and testing are sampled at the following inner-scaled wall-normal coordinates: $y^+ = 15,\ 30,\ 50$ and 100. Note that '+' denotes viscous scaling, *i.e.* in terms of the friction velocity u_τ or the

viscous length $\ell^* = \nu/u_\tau$ (where ν is the fluid kinematic viscosity). A dataset is defined as a collection of samples, each consisting of the shear-stress and pressure fields at the wall as inputs, along with the corresponding velocity fields at the target wall-normal locations as outputs. The training/validation dataset at $Re_\tau = 180$ consists of 50,400 instantaneous fields, with a sampling period of $\Delta t_s^+ = 5.08$. The sampling period at $Re_\tau = 550$ is set to $\Delta t_s^+ = 1.49$ and the training/validation dataset includes 19,920 fields. The resolution in viscous units of the individual fields is the same in both Re_τ datasets. This is necessary to represent all the flow features; however this also implies that the number of Fourier modes in the wall-parallel directions is higher at $Re_\tau = 550$, as shown in table 1, even if the domain Ω is the same. A higher number of modes corresponds to a higher number of spatial locations in which the different quantities are sampled from the DNS and this partially compensate the lower number of fields since the two proposed methods act locally on the input data, as it will be shown in §2.2.

In both Re_τ cases, the dataset is split into training and validation sets, with a ratio of 4:1. The training and validation sets are obtained from a group of randomly-initialized simulations. With the number of samples chosen for our investigation at either Re_τ, there is a simulation whose samples are used both for training or for validation. Note that there is no temporal overlapping between the two datasets, as the first samples from the "shared" simulation are sequentially added to the training set until the required number of training samples is reached. The remaining samples are then added to the validation set. This is done to reduce the correlation between the two datasets. While using the separate groups of simulations or adding an additional time separation between the last sample of the training set and the first sample of the validation would have further reduced the correlation, this was not enforced, since the validation error is only used during training to tune the network hyperparameters and avoid over-fitting the training dataset. If the training provides satisfactory results, it is common practice to use both training and validation sets for the training of the final models. This was not implemented in our case, but it allows the model to make use of all the available data to improve its performance. The validation error may be significantly different from the one computed during testing, hence all the comparisons in this work are based on the error computed on the test dataset. The predictions used to assess the performance of the trained models were obtained from additional simulations. This is done both at $Re_\tau = 180$ and $Re_\tau = 550$, and it is necessary to ensure that the training and test datasets are completely uncorrelated, both in space and time. Test samples were taken from simulations initialized with a random seed, different from that of the training-data simulation. The correlation between train simulations and between train and test simulations was checked with the cross-correlations $\rho_{ij}(h)$, defined as (Makarashvili *et al.* 2017):

$$\rho_{ij}(h) = \frac{\gamma_{ij}(h)}{\sqrt{\gamma_i(0)\gamma_j(0)}}, \tag{1}$$

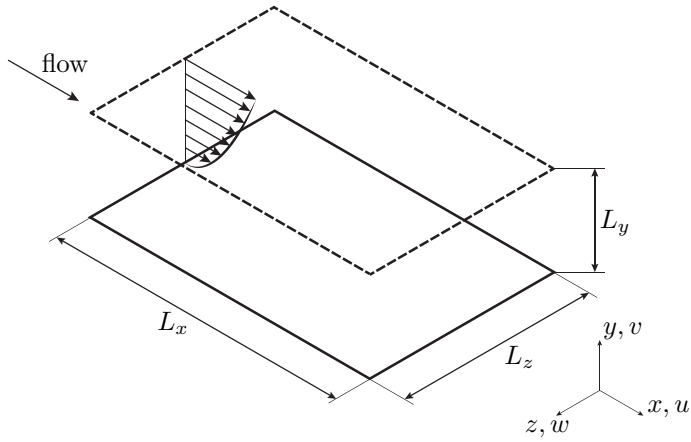

Figure 1: Computational domain and frame of reference for the DNS of the turbulent open channel considered in this study.

Re_τ	$N_x \times N_z \times N_y$	#Train.+Val. fields	$\Delta t^+_{s,\text{train}}$	T^+_{train}	#Test fields	$\Delta t^+_{s,\text{test}}$	T^+_{test}
180	$192 \times 192 \times 65$	50,400	5.08	$\approx 2.56 \times 10^5$	3,125	1.69	$\approx 5.3 \times 10^3$
550	$512 \times 512 \times 193$	19,920	1.49	$\approx 2.97 \times 10^4$	3,320	1.49	$\approx 5 \times 10^3$

Table 1: Description of the DNS datasets used for computing the EPOD and training/testing the CNN-based models.

where h is the time lag and $\gamma_{ij}(h) = \mathbb{E}[(x_{i,t+h} - \mu_i)(x_{j,t} - \mu_j)]$ are the cross-covariances. i and j refer to the two simulations and x is the considered quantity, in this case one of the velocity components at a given wall-normal location. By computing $\rho_{ij}(0)$, it was possible to verify that the simulations had run sufficiently long to ensure that they are overall stastistically uncorrelated, before starting the sampling to build the datasets.

The size of the test dataset (more than 3,000 fields for both Re_τ) is sufficient to achieve convergence of the turbulence statistics from the predicted flow, and then these are compared with the reference values from the DNS.

2.2. *Summary of the methods*

In this study we consider three different data-driven methods to estimate the instantaneous two-dimensional fields of the velocity fluctuations, at a given wall-normal distance. Two models (FCN and FCN-POD) are introduced and compared with the Extended POD method, which is used as a baseline model, highlighting the capabilities (and limitations) of linear estimation. Despite using the same input data and predicting the same output quantities, the three models leverage on different tools in order to extract the information from

the inputs and estimate the outputs. Here we briefly review the differences and similarities between the models, each method is described in detail in the dedicated sections.

EPOD and FCN-POD use proper orthogonal decomposition (POD, see Lumley 1967): this approach is advantageous as it allows to filter out the noise content represented by small and uncorrelated scales, thus taking advantage of the energy optimality of POD modes. Additionally, spatial and temporal dynamics are separated. EPOD decomposes both the input and the output: the flow field at a certain wall-normal distance is reconstructed as a linear combination of orthogonal modes $\boldsymbol{\phi}_i(\boldsymbol{x})$:

$$\boldsymbol{u}(\boldsymbol{x}, t) \approx \sum_{i=1}^{N_m} \psi_i(t)\sigma_i\boldsymbol{\phi}_i(\boldsymbol{x}), \tag{2}$$

where N_m is the total number of POD modes, $\psi_i(t)$ is the temporal POD coefficient corresponding to mode i, and σ_i is its corresponding root-squared energy contribution. While the orthogonal modes are computed from a POD of the training dataset, the temporal coefficients are estimated by decomposing input wall quantities using POD and assuming a linear correlation between the known temporal modes of the wall quantities $\psi_{w,i}(t)$ and the temporal modes $\psi_i(t)$ of the output flow field of the test dataset. F urther details are provided in §2.3. On the other hand, FCN-POD applies POD only on the output flow fields, using the spatial modes from the training set and predicting the corresponding temporal coefficients using a neural network. This is a development of the model used by Güemes *et al.* (2019), note however that in that study the domain employed and reconstructed provided a compact POD eigenspectrum, here the availability of a larger domain in the streamwise and spanwise directions spreads the energy content over a wider set of POD modes. This makes the predictions of temporal coefficients more difficult, especially those associated with the least energetic modes. To address this issue, the large instantaneous flow fields were subdivided into $N_{s_x} \times N_{s_z}$ smaller regions (henceforth referred to as subdomains). The neural network predicts the temporal coefficients for all subdomains at the same time.

The FCN and FCN-POD models consider the instantaneous two-dimensional fields of the streamwise and spanwise wall-shear-stress components and of the wall pressure as inputs. In the physical-coordinates representation of these fields, the presence of coherent features motivates the use of convolutional layers in our neural-network models to process the information. In these layers, the inputs are processed at the same time, hence a convolution in three dimensions is performed and it is defined as:

$$F_{i,j} = \sum_l \sum_m \sum_n I_{i-m,j-n,l} K_{m,n,l}, \tag{3}$$

following Goodfellow *et al.* (2016), where $\mathbf{I} \in \mathbb{R}^{d_1 \times d_2 \times d_3}$ is the input, $\mathbf{K} \in \mathbb{R}^{k_1 \times k_2 \times k_3}$ is the so-called *kernel* (or *filter*) containing the learnable parameters

of the layer, and the transformed output $\mathbf{F}$ is the *feature map*. Note that we consider $d_3 = k_3$, meaning that the resulting feature map is two-dimensional. Multiple feature maps can be stacked and sequentially fed into another convolutional layer as input. This allows the next layer to combine the features individually identified in each feature map, enabling the prediction of larger and more complex features for progressively deeper convolutional networks. Note that a convolutional layer can be rewritten as a matrix multiplication, hence it is mathematically equivalent to a fully-connected layer Ma & Lu (2017). If we assume that the input features are spatially localized, using a convolutional layer allows translational equivariance to be enforced in both periodic directions. Furthermore, since $k_i \ll d_i \ \forall i \neq 3$, the use of kernels greatly reduces the number of parameters that need to be learned during training (in comparison with fully-connected MLP networks).

Both neural-network-based models use fully-convolutional neural network (hence FCN in the names). This architecture is similar, but conceptually different from CNNs, which typically have several convolutional layers followed by one or more fully-connected layers (which are the building blocks of MLP networks). In CNNs, the localized information processed by the individual convolutional kernels is combined to obtain a global prediction, whereas in FCNs only convolutional layers are present and the network architecture is based on the assumption that the relation between input and output variables is spatially localized. The input region from which a single point of the output can draw information is called *receptive field* and it can be computed based on the network architecture, as described by Dumoulin & Visin (2016). In the FCN model, the instantaneous two-dimensional velocity fluctuations are directly predicted from the input fields by using the fully-convolutional neural network. Additional details are provided in §2.4. On the other hand, the neural network in the FCN-POD model predicts the temporal coefficients for each of the $N_{s_x} \times N_{s_z}$ subdomains, as described in §2.5.

2.3. *Extended POD*

EPOD provides a linear relation between input and output and it is the reference method for all the following comparisons. Doing so, it is possible to assess the prediction improvement with nonlinear, neural-network-based methods in the context of wall-bounded turbulence. Wall quantities in each field can be rearranged as a row vector and used to assemble a snapshot matrix $\mathbf{W}$. Using the method proposed by Sirovich (1987), it is possible to decompose this matrix into POD modes as:

$$\mathbf{W} = \mathbf{\Psi}_w \mathbf{\Sigma}_w \mathbf{\Phi}_w, \tag{4}$$

with $\mathbf{\Psi}_w$ and $\mathbf{\Phi}_w$ being the temporal and spatial mode matrices respectively, and $\mathbf{\Sigma}_w$ being a diagonal matrix containing the singular values. The extended POD modes (Borée 2003), corresponding to the projection of the wall quantities on the flow-field temporal basis, are defined as

$$\mathbf{L} = \mathbf{\Psi}_w^T \mathbf{U}, \tag{5}$$

where $\mathbf{U}$ is the snapshot matrix in which the instantaneous velocity fluctuations in the three directions are rearranged:

$$\mathbf{U} = \begin{bmatrix} u_{x_1}^{t_1} & \dots & u_{x_{N_p}}^{t_1} & v_{x_1}^{t_1} & \dots & v_{x_{N_p}}^{t_1} & w_{x_1}^{t_1} & \dots & w_{x_{N_p}}^{t_1} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ u_{x_1}^{t_{N_t}} & \dots & u_{x_{N_p}}^{t_{N_t}} & v_{x_1}^{t_{N_t}} & \dots & v_{x_{N_p}}^{t_{N_t}} & w_{x_1}^{t_{N_t}} & \dots & w_{x_{N_p}}^{t_{N_t}} \end{bmatrix}. \tag{6}$$

Here N_t refers to the total number of snapshots (equal to the number of instantaneous flow fields N_f for the EPOD) and N_p refers to the total number of grid points in one field.

If the dataset is sufficiently large to reach statistical convergence, the matrix $\mathbf{L}$ describes the relationship between the temporal POD coefficients of a certain distribution of wall features, and those of the corresponding flow field. Once the temporal correlation matrix is known, an out-of-sample flow field $\mathbf{u}$ can be reconstructed using $\mathbf{L}$ and the instantaneous realization of wall features as follows:

$$\mathbf{u} = \boldsymbol{\psi}_w \mathbf{L}, \tag{7}$$

where $\boldsymbol{\psi}_w$ is the vector containing the temporal coefficients of the wall fields used for prediction. Note that $\boldsymbol{\psi}_w$ is retrieved by projecting the out-of-sample wall field $\boldsymbol{w}$ on the POD basis: $\boldsymbol{\psi}_w = \boldsymbol{w}\boldsymbol{\Phi}_w^T\boldsymbol{\Sigma}_w^{-1}$, where $\boldsymbol{\Phi}_w^T\boldsymbol{\Sigma}_w^{-1}$ is readily available from the training dataset.

An important remark is that, due to the homogeneity in the streamwise and spanwise directions and the large number of fields for the proposed Reynolds numbers, the wall-based matrix $\boldsymbol{\Sigma}_w$ can be ill-conditioned. In fact due to the correlation between subsequent time-resolved snapshots, the rank of $\boldsymbol{\Sigma}_w$ is smaller than N_f, which is the number of snapshots (here smaller than the number of points). A well- or ill-conditioned matrix is based on the condition number: if the condition number is large the matrix is ill-conditioned, while if it is small the matrix is well-conditioned. Defining the condition number as $\kappa(\Sigma_w) = \frac{max(\Sigma_w)}{min(\Sigma_w)}$, the large difference between the first and the last POD modes leads to an ill-conditioned matrix, which is difficult to be inverted numerically. To address this issue a reduced-order representation of the matrix $\boldsymbol{\Sigma}_w$ is employed, using a number of modes equal to the rank of the matrix. The energy contribution of the excluded modes is zero, and thus the predictions are numerically equivalent to a full-rank estimation. Even if $\boldsymbol{\Sigma}_w$ would have rank equal to N_f, it might be adequate to truncate the matrix $\boldsymbol{L}$ (Discetti *et al.* 2018). Decomposing the flow quantities as: $\mathbf{U} = \boldsymbol{\Psi}\boldsymbol{\Sigma}\boldsymbol{\Phi}$, similarly to what is done for the wall quantities in equation (4), it can be observed that the product of the two matrices $\boldsymbol{\Psi}\boldsymbol{\Psi}_w^T$ in equation (5) returns a unitary-norm matrix with rank equal to those of $\boldsymbol{\Psi}$ and $\boldsymbol{\Psi}_w^T$, which are bases in the $\mathbb{R}^{N_f}$ vector space. As a consequence, a certain j^{th} wall mode, uncorrelated with any field mode, would not result in a corresponding null row or column. To ensure removing the uncorrelated content from the matrix $\boldsymbol{\Psi}\boldsymbol{\Psi}_w^T$, Discetti *et al.* (2018) proposed to set to zero all the entries of the matrix with absolute values

smaller than a threshold proportional to the matrix standard deviation. In the present work we have found an error drop of approximately 10 percentage points with respect to the standard EPOD procedure. However, since the EPOD is used as a benchmark for the performance of linear methods with respect to the FCN-based approaches proposed herein, the filtered EPOD by Discetti *et al.* (2018) is not included in this comparison for brevity.

2.4. *Fully-convolutional neural-network predictions*

Fully-convolutional networks are commonly used in applications where the input and output domains have structural similarities. *Image segmentation* (Long *et al.* 2015) is one such case, since the output has the same spatial dimension as the input, as in our predictions of two-dimensional flow fields. The inputs of the network are the wall-shear-stress components in the streamwise and spanwise directions, as well as the pressure at the wall. Each of the inputs of the network is normalized using the respective mean and standard deviation, as computed from the training/validation set. The predictions are performed using the same mean and standard deviation values on the test dataset inputs. The outputs are the instantaneous velocity fluctuations, denoted as u, v and w (corresponding to the streamwise, wall-normal and spanwise velocities, respectively), at a given distance from the wall. Note that the predictions are carried out at the same time as the input fields. In our previous work (Guastoni *et al.* 2020*a*), a similar FCN was used to predict the streamwise component of instantaneous flow fields. In the present study the predictions are extended to the wall-normal and spanwise components, however the back-propagation algorithm that is used to train the networks works best when the error in the prediction of three outputs (*i.e.* the three velocity components) has a similar magnitude for all of them. Thus, the fluctuations are scaled as follows:

$$\widehat{u} = u, \quad \widehat{v} = v\frac{u_{\mathrm{RMS}}}{v_{\mathrm{RMS}}}, \quad \widehat{w} = w\frac{u_{\mathrm{RMS}}}{w_{\mathrm{RMS}}}, \tag{8}$$

where RMS refers to root-mean-squared quantities. During inference (*i.e.* when the predictions are computed from the inputs in the test dataset), the outputs of the network are re-scaled back to their original magnitude. The network is trained to minimize the following loss function:

$$\mathcal{L}_{\mathrm{FCN}}(\widehat{\boldsymbol{u}}_{\mathrm{FCN}}; \widehat{\boldsymbol{u}}_{\mathrm{DNS}}) = \frac{1}{N_x N_z}\sum_{i=1}^{N_x}\sum_{j=1}^{N_z}\left|\widehat{\boldsymbol{u}}_{\mathrm{FCN}}(i,j) - \widehat{\boldsymbol{u}}_{\mathrm{DNS}}(i,j)\right|^2, \tag{9}$$

which is the mean-squared error (MSE) between the instantaneous prediction $\widehat{\boldsymbol{u}}_{\mathrm{FCN}}$ and the true velocity fluctuations $\widehat{\boldsymbol{u}}_{\mathrm{DNS}}$, as computed by the DNS and scaled in the same way as the network outputs. Training details are further discussed in appendix Appendix A. The FCN architecture is shown in figure 2. Each convolution operation (except for the last one) is followed by batch normalization (Ioffe & Szegedy 2015) and a rectified-linear-unit (ReLU, see Nair & Hinton 2010) activation function.

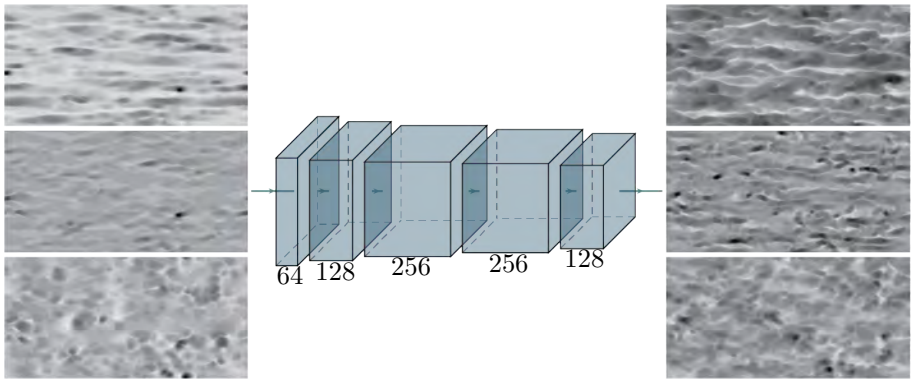

Figure 2: Schematic representation of the considered FCN architecture. The input fields are on the left (from top to bottom: streamwise wall-shear stress, spanwise wall-shear stress and wall pressure) and the outputs are on the right (from top to bottom: streamwise, wall-normal and spanwise fluctuations). The numbers indicate the number of kernels applied to each of the layers. The kernels (not represented in the figure) have size (5×5) in the first convolutional layer, and (3×3) in the subsequent layers. A darker colour corresponds to a higher value of the represented quantity.

The chosen inputs and outputs allow the FCN to learn only the spatial relation between the quantities at the wall and the fluctuations farther away from it. Note that it would also be possible to consider predictions in time, and in that case convolutional neural networks could be used (van den Oord *et al.* 2016) treating time as another spatial coordinate, or it would be possible to use recurrent networks, specifically designed to learn temporal sequences as we have recently shown with long-short term memory (LSTM) networks (Srinivasan *et al.* 2019; Guastoni *et al.* 2019). In both cases, the need of multiple samples in time makes the model less flexible than one that relies only on spatial correlations, both during training and testing. These models usually assume a constant sampling time for the data sequence, which might be difficult to enforce, for example if the fields are taken from a numerical simulation with adaptive time step. During inference, models that work with sequences would require input fields at different times to perform the prediction. On the other hand, a single input sample is sufficient for the FCN to predict the output. Input and output fields of the FCN model are obtained from a simulation with periodic boundary conditions in the streamwise and spanwise directions. Such constraints could be added to the loss function, however this would imply that periodicity would only be satisfied in a least-square sense and for this reason they are not used in the present study. Instead, in our implementation we are able to strictly enforce periodicity in both wall-parallel directions by leveraging the fact that the convolutional-network output is deterministic and influenced only by the

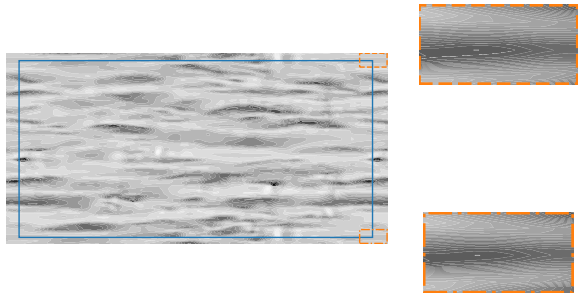

Figure 3: Example of padding of the input streamwise wall-shear stress field. The solid blue line indicates the size of the original field. The orange boxes show the information available to the FCN for the reconstruction of a single point of the output fields, at both edges of the domain. A darker colour corresponds to a higher value of the represented quantity.

local information in the receptive field. In other words, if the network receives a certain local input, the local output value will always be the same, regardless of the local position within the input field. In order to have the same values on both edges of the domain, the inputs fields are padded in the periodic directions, *i.e.*, they are extended on both ends, using the values from the other side of the fields. The padding procedure is exemplified in figure 3, in which it is highlighted how the information in the receptive field at the two ends of the domain is indeed the same. The receptive field for this architecture is 15×15 points, hence 16 points are added to each field in both streamwise and spanwise directions. Note that this padding would determine a network output size that is slightly larger than the size of the velocity fields from the DNS, and therefore the network output is cropped to match the size of the reference flow fields. The padding involves a computational overhead due to the increased size of the input fields, however it is important to highlight that the padding is architecture-dependent and not input-dependent, meaning that the input is $\approx 17\%$ bigger with a 192×192 field resolution (at $Re_\tau = 180$), but only $\approx 6\%$ bigger when the fields have a size of 512×512 (at $Re_\tau = 550$).

2.5. *POD-based predictions with convolutional neural networks*

The FCN-POD model is built upon the previous work by Güemes *et al.* (2019). In the present work a different neural network architecture, a fully-convolutional one, is used instead of a CNN. The network takes as inputs the wall-shear-stress and wall-pressure fields to predict the temporal coefficients of the POD modes in each of the $N_{s_x} \times N_{s_z}$ subdomains, in which the output velocity-fluctuations field is divided. The size of the subdomains in the streamwise and spanwise directions is roughly $(h \times 0.5h)$ for the $Re_\tau = 180$ case, and $(0.4h \times 0.2h)$ at

$Re_\tau = 550$. This is smaller than the domain considered in the previous cited work, where the domain had size of $h \times h$. The advantage of this approach compared to directly decomposing the full field lies in the fact that, in these subdomains, the first POD modes contain a very large fraction of the total energy content. This is a direct consequence of including the energy of the structures larger than the domain into the first POD mode (Liu *et al.* 2001; Wu 2014) As a result, a lower number of modes is needed in each subdomain to reconstruct a large fraction of the total energy, as shown in figures 4a) and c). The same amount of energy would have been spread over a larger number of modes if the entire field had been considered at once.

The choice of the size of the subdomains is the result of a compromise between network capability of reconstructing the majority of the energy content of the flow and compression of the information. For $Re_\tau = 180$ the flow fields were divided into 12×12 subdomains, while for $Re_\tau = 550$ a discretization into 32×32 subdomains is chosen. Note that the number of subdomains was selected to ensure that the first POD mode contains a similar level of energy in both cases.

Also in this case the snapshot matrix $\mathbf{U}$ is assembled, similarly to what was done for EPOD in section 2.3. The most important difference is that in this case the total number of snapshots N_t is equal to the number of instantaneous flow fields N_f times the number of subdomains per each flow field ($N_{s_x} \times N_{s_z} \times N_f$), and N_p refers to the total number of grid points in one subdomain. This organization of the snapshots implies that the spatial modes are assumed to be the same for each individual subdomain. These POD spatial modes can be evaluated solving the eigenvalue problem of the spatial correlation matrix $\mathbf{C}$ as follows:

$$\mathbf{C} = \mathbf{U}^T\mathbf{U} = \mathbf{\Phi}^T\mathbf{\Lambda}\mathbf{\Phi}, \tag{10}$$

where $\mathbf{\Phi}$ is a matrix the rows of which contain the spatial POD modes, while $\mathbf{\Lambda}$ is a diagonal matrix with elements $\lambda_i = \sigma_i^2$, which represent the variance content of each mode. The POD coefficients $\psi_i(t)$ are obtained by projecting the flow fields on the spatial POD modes computed with equation (10). Note that this economy-size decomposition returns a number of POD modes N_m equal to $3N_p$ and that for such a discrete dataset the flow field in the subdomain is represented with N_m modes without approximation. It has to be recalled that this formulation applies for regular grids; for non-regular grids, the snapshot matrix $\mathbf{U}$ should be adapted to take into account the area of each grid point.

The temporal POD coefficients of each instantaneous flow field were rearranged in a tensor of size $N_{s_x} \times N_{s_z} \times N_r$ to train a FCN, with $N_{s_x} \times N_{s_z}$ being 12×12 and 32×32 for $Re_\tau = 180$ and 550, respectively and N_r the number POD modes to be predicted, with $N_r < N_m$. As shown in figures 4b) and d), the first 64 POD modes account for 90% of the total energy at $Re_\tau = 180$, while 128 modes are needed at $Re_\tau = 550$ to retain a similar amount of energy. Therefore our predictions are truncated at $N_r = 64$ and 128 for $Re_\tau = 180$ and 550 respectively. As illustrated in figure 5, each filter corresponds to the

Re_τ	N_f	$N_{s_x} \times N_{s_z}$	N_t	N_p	N_m	N_r
180	40 320	12×12	5 806 080	256	768	64
550	15 936	32×32	16 318 464	256	768	128

Table 2: Description of the POD settings for the subdomain decomposition. Note that only the training set has been used to compute the POD modes.

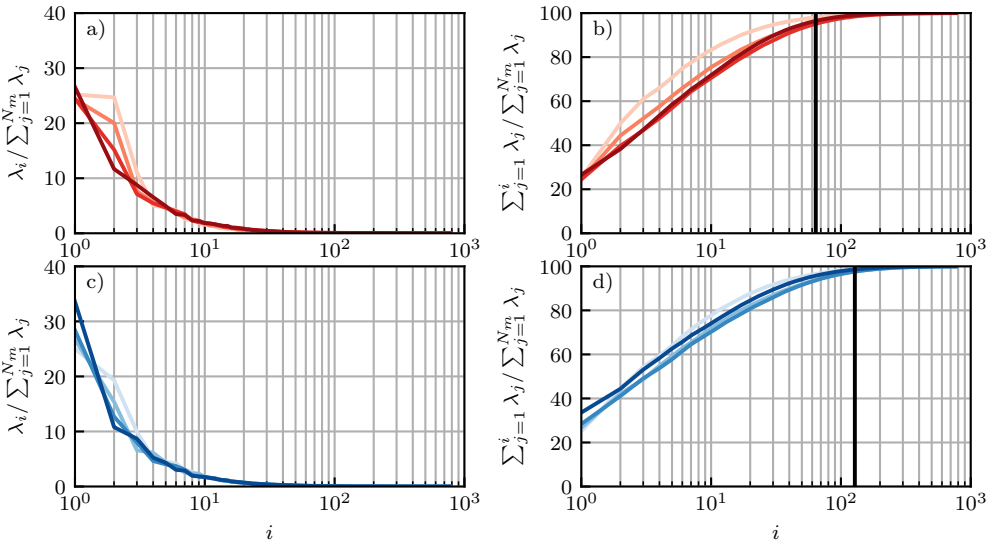

Figure 4: Distribution of a), c) POD eigenvalues $\lambda_i = \sigma_i^2$ (where i denotes mode number) and b), d) cumulative eigenspectrum $\sum_{j=1}^{i} \lambda_j$ normalized with the cumulative sum of the eigenvalues $\sum_{j=1}^{N_m} \lambda_j$ of each case. The colour refers to the wall-normal locations described in §2.1, where darker colors indicate larger distance from the wall. Note that a) and b) correspond to $Re_\tau = 180$, while c) and d) correspond to $Re_\tau = 550$. The solid black vertical lines in b) and d) refer to the number of modes selected for prediction in this study.

$N_{s_x} \times N_{s_z}$ POD coefficients of a given mode number. In general, the energy distribution reported for both Re_τ cases is very similar. The main significant difference is that the energy distribution becomes more compact at $y^+ = 15$ for the low-Re_τ case (see figure 4). Table 2 shows the actual values of the POD settings for each Re_τ number.

In order to reconstruct the instantaneous fluctuation fields, the time coefficients were predicted using a FCN, with the wall-shear-stress components and

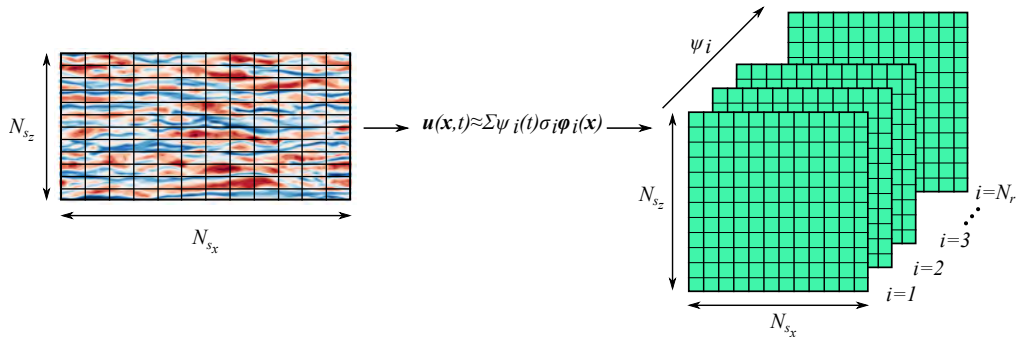

Figure 5: Schematic representation of the encoding of turbulent flow fields into target tensors of the network, containing the temporal POD coefficients. When $Re_\tau = 180$, the flow field is divided into 12×12 subdomains in the streamwise and spanwise directions respectively; a 3D tensor is built, where the first two dimensions correspond to the streamwise and spanwise position of the subdomain, while the third one accounts for the number of POD modes to be predicted.

the pressure at the entire wall as inputs. It is important to note that the temporal coefficients for the subdomains are predicted together, and each subdomain is receiving information from the surrounding ones through the convolutional operations. The predicted $N_{s_x} \times N_{s_z} \times N_r$ POD temporal coefficients belonging to the subdomains are used to reconstruct their respective fluctuation fields, with the orthonormal basis functions are retrieved from the training data. The underlying assumption is *ergodicity*, *i.e.* both the training and test datasets share the same statistical features and, consequently, the same spatial modes. This requires a sufficiently large training dataset to ensure convergence of the spatial modes, which is generally ascribed to the convergence of second-order statistics. Once the velocity-fluctuation fields are reconstructed within each subdomain, they are assembled together to provide the full-domain prediction. Note that the *tiling* of the subdomains does not provide any guarantee of smoothness across the edges of the subdomains because of the finite number of modes that are used to reconstruct the flow and because of the prediction error in the temporal coefficients. Also, the predictions are performed at the same instant as that of the input fields. The implemented neural network does not require the knowledge of the input at previous timesteps, thus avoiding the limitations of availability of time sequences, as discussed above. The network is

trained to minimize the loss function:

$$\mathcal{L}_{\mathrm{FCN-POD}}(\psi_{\mathrm{POD}};\psi_{\mathrm{DNS}}) = \frac{1}{N_{s_x}N_{s_z}N_r}\sum_{i=1}^{N_{s_x}}\sum_{j=1}^{N_{s_z}}\sum_{k=1}^{N_r}\left|\psi_{\mathrm{POD}}(i,j,k)-\psi_{\mathrm{DNS}}(i,j,k)\right|^2, \tag{11}$$

which is the MSE between the predicted and the actual POD temporal coefficients of the DNS data. Parameter settings for training are detailed in appendix Appendix A.

The neural-network architecture considered here blends the FCN shown in figure 2 and the network used by Güemes *et al.* (2019) (see figure 1 in that work). As in the FCN approach, each convolution operation (except for the last one) is followed by batch normalization (Ioffe & Szegedy 2015) and a ReLU (Nair & Hinton 2010) activation function. After each activation function a max pooling layer is added. Differently from what was done in the FCN approach, here the velocity components were not scaled before the decomposition, in order to keep the physical encoding based on the turbulent kinetic energy (TKE) of the flow. Note that by modifying the relative contribution of the velocity components to the energy norm, the modes would have been sorted based on a norm different from the TKE. The main difference with respect to the network used by Güemes *et al.* (2019) is the fact that here a single network is used to predict the full set of POD coefficients, instead of using different networks to predict each mode. Additionally, the work by Güemes *et al.* (2019) focused on a smaller region of the flow field, and therefore no subdomains were required for the region of interest. Lastly, the final fully-connected layer in Güemes *et al.* (2019) was not considered here, in order to have an architecture more directly comparable with the FCN.

3. Results

The predictions of the trained models are compared with the data obtained from the DNS at $Re_\tau = 180$ and 550. The performance assessment is carried out first from a qualitative point of view and subsequently from a quantitative perspective, based on predictions of instantaneous fields, turbulence statistics and spectra.

3.1. *Instantaneous predictions*

The predicted fluctuation fields are first qualitatively inspected. Note that the fluctuation flow fields are the direct output of the FCN models, while in the FCN-POD models the temporal coefficients need to be processed to reconstruct the fluctuations, as outlined above. In this work the sampling period in the simulation is fixed, however we showed in our previous work (Guastoni *et al.* 2020*a*) that using less correlated samples during training (*i.e.* larger sampling period, that corresponds to a longer sampled time for the same number of samples gathered) can effectively improve the quality of the instantaneous

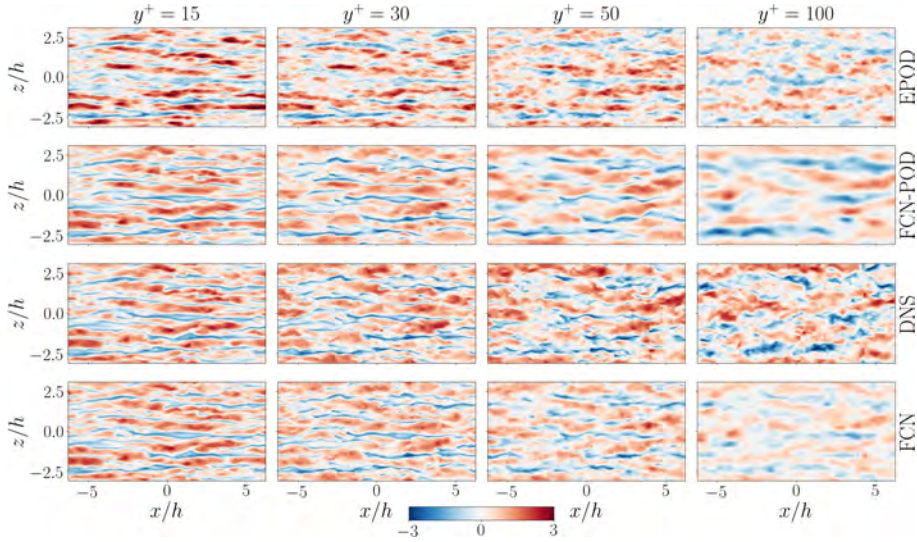

Figure 6: Comparison of the streamwise fluctuation fields at $Re_\tau = 180$, scaled with the corresponding u_{RMS}, from EPOD (1st row), FCN-POD (2nd row), reference DNS (3rd row) and FCN (4th row). Results at $y^+ = 15$ (1st column), $y^+ = 30$ (2nd column), $y^+ = 50$ (3rd column) and $y^+ = 100$ (4th column).

predictions of the FCN method, provided that the neural-network capacity is sufficient to generalize over the training dataset.

In figure 6, the predictions of an instantaneous field of streamwise velocity fluctuations based on the various methods (namely FCN, FCN-POD and EPOD) are compared with the reference DNS. The predictions of the wall-normal and spanwise fluctuations at the same instant are shown in appendix Appendix B. At $y^+ = 15$ all the methods provide accurate results, although the EPOD overestimates the fluctuations from the high-speed streaks. At $y^+ = 30$ such overestimation is reduced, although EPOD does not seem to reach a level of accuracy comparable to the other two models. The CNN-based methods start to exhibit some deviations with respect to the reference at $y^+ = 50$, where the FCN-POD field is smoother and the FCN is slightly noisier than the DNS. Farther from the wall, the footprint of the large scales at the wall (through linear superposition, see Dogan *et al.* 2019) is less pronounced, and therefore the ability of EPOD (which is a linear method) to predict the flow in this region is significantly reduced. In fact, the fields predicted through EPOD at $y^+ > 15$ are qualitatively very similar to the DNS, although the fluctuations become increasingly attenuated at larger y^+. Furthermore, the FCN-POD method tends to merge neighbouring regions with high- or low-velocity fluctuations, predicting more elongated streak-like patterns than in the reference field. This is more evident at $y^+ = 100$, leading to an overestimation of the amplitude

of the regions in the flow where the velocity fluctuations are higher. At this location, the FCN is not able to provide a reliable prediction of the flow field, capturing only the regions in which the magnitude of the fluctuations is higher. The corresponding structures probably have a distinct footprint at the wall, which allows the FCN to identify them.

As discussed in §2.5, the FCN-POD approach does not guarantee flow smoothness across the edges of the subdomains. Close inspection of the predictions from the FCN-POD method reveals the edges of the subdomains at all y^+, and the tiling is more evident in the streamwise direction because of the discontinuities located at the same spanwise location, orthogonally to the main flow structures. Despite these limitations we can observe that the velocity fields are generally smooth, without steep discontinuities at the edges of the subdomains: the variations of the velocity magnitude at the edges are of the same order as the fluctuations at the corresponding wall-normal distance.

The qualitative observations discussed above are complemented with a quantitative assessment of the instantaneous prediction performance, by analyzing the MSE $\mathcal{L}$ between the instantaneous predictions (denoted by 'Pred') and the reference (defined for each of the fluctuations independently), as shown in figure 7. The error in the streamwise velocity fluctuations is defined as:

$$\mathcal{L}(u_{\mathrm{Pred}}; u_{\mathrm{DNS}}) = \frac{1}{N_x N_z} \sum_{i=1}^{N_x} \sum_{j=1}^{N_z} \left| u_{\mathrm{Pred}}(i,j) - u_{\mathrm{DNS}}(i,j) \right|^2 , \tag{12}$$

and similarly for the wall-normal and spanwise fluctuations. Note that the velocity-fluctuation fields need to be reconstructed from the temporal coefficients in POD-based models (FCN-POD and EPOD). Both neural-network models (FCN and FCN-POD) are trained using a stochastic algorithm and, in order to show the robustness of the optimal configuration, the statistics at each y^+ are averaged over 3 different models, with different initial random weight initializations. Since the EPOD algorithm is completely deterministic, one single prediction is needed. The neural-network-based models consistently provide a lower error than the EPOD, with the FCN yielding a slightly better instantaneous performance closer to the wall than the FCN-POD approach. The gap between the two is reduced when moving away from the wall, where the prediction error of the streamwise fluctuations is approximately the same for both models at $y^+ = 50$, and it is slightly higher for the FCN at $y^+ = 100$.

The FCN architecture reported by Guastoni *et al.* (2020*a*) would only predict the streamwise velocity component of the velocity field at the target y^+. The addition of the two other components implies that the FCN has multiple outputs that need to be optimized at the same time. We note that adding the two additional fluctuating components as outputs leads to slightly less accurate predictions with respect to those reported by Guastoni *et al.* (2020*a*) for one single output. This is not surprising, since the capacity of the network remained unchanged, however we tested a variation of the model architecture based on this observation, in order to have more layers dedicated

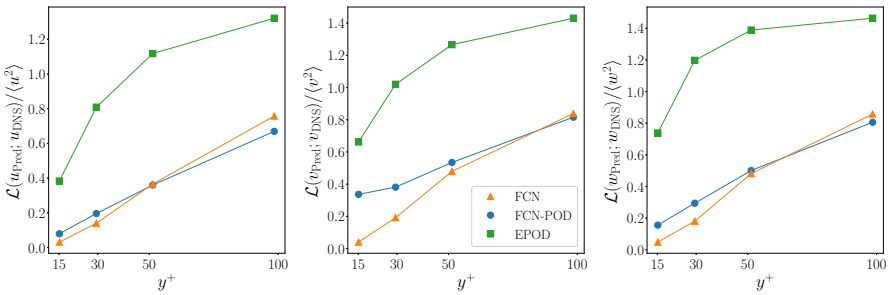

Figure 7: Mean-squared-error in the instantaneous fields (scaled with the corresponding RMS components) predicted by the three models at $Re_\tau = 180$, for the streamwise (left), wall-normal (middle) and spanwise (right) velocity fluctuations.

to the prediction of each individual component. This network variation has a common part, identical to the original FCN up to the 4th convolutional layer, in which the weights are optimized using the information from the error gradients computed for all the outputs. The last two convolution operations are replicated for each velocity component and the weights of these layers are updated only with the error associated to the respective output. Such a network, despite its higher capacity, provided worse predictions. A strong causal relation between the different components of the velocity (Lozano-Durán *et al.* 2020) can be a possible explanation for this result, which shows that updating all the weights with information from the three components at the same time can be beneficial for the quality of the predictions. Note that it is not trivial to design an architecture able to provide the best trade-off between single-component predictions and usage of the information from all the components, and obtaining such an architecture would require further investigation. For the FCN-POD model the multiple-component predictions were obtained as discussed by Güemes *et al.* (2019). The temporal POD coefficients can be projected on spatial POD modes involving the three velocity components, thus requiring only one output to predict the three fluctuations. The network architecture is different than the one used by Güemes *et al.* (2019), since it predicts directly all the needed time coefficients for each snapshot. While the final fully-connected layer included in the network architecture by Güemes *et al.* (2019) improves the robustness of the prediction, the FCN-POD implementation used here has a much smaller number of weights, thus significantly reducing the computational cost, and retains a larger number of POD modes (and thus more energy).

Predictions of the streamwise fluctuation fields from the various methods at $Re_\tau = 550$ are shown in figure 8, while the results for the wall-normal and spanwise components are presented in Appendix Appendix B. Despite the higher friction Reynolds number, the FCN maintains a performance similar to the one

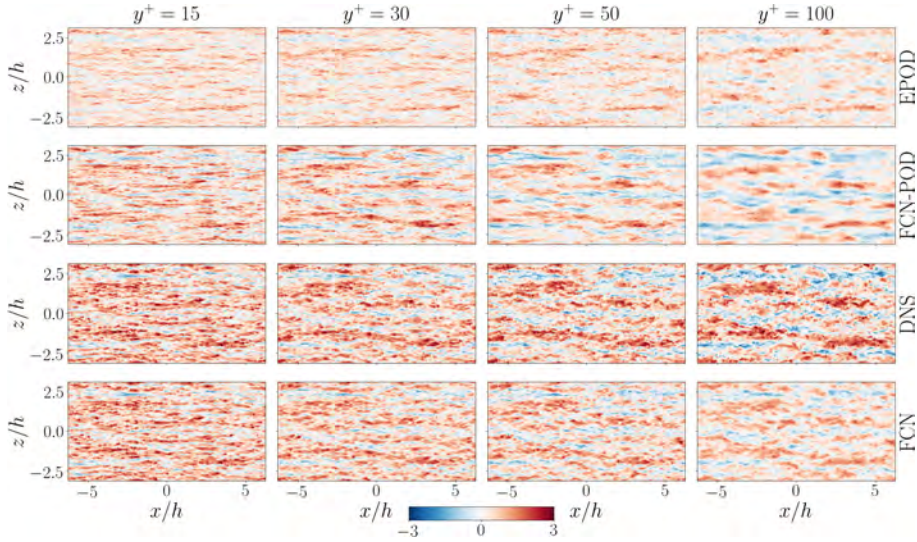

Figure 8: Comparison of the streamwise fluctuation fields at $Re_\tau = 550$, scaled with the corresponding u_{RMS}, from EPOD (1st row), FCN-POD (2nd row), reference DNS (3rd row) and FCN (4th row). Results at $y^+ = 15$ (1st column), $y^+ = 30$ (2nd column), $y^+ = 50$ (3rd column) and $y^+ = 100$ (4th column).

achieved for $Re_\tau = 180$, at all wall-normal locations. Note that the FCN has the same architecture as the lower-Reynolds-number case, *i.e.* it has the same number of trainable parameters, while in the case of the FCN-POD approach, the network was modified to reconstruct approximately the same amount of energy as at $Re_\tau = 180$. Despite the higher number of employed subdomains, the tiling is more apparent at $Re_\tau = 550$. The prediction performance of the FCN-POD model degrades less quickly than the FCN when moving away from the wall, however the latter still performs better at $y^+ = 50$, as shown in figure 9. On the other hand, the EPOD also exhibits similar error levels as those reported for $Re_\tau = 180$, except at $y^+ = 15$, where the reconstruction of the streamwise-fluctuation field is significantly worse.

3.1.1. *Inclination of coherent structures*

The coherent structures in wall-bounded turbulence are inclined (Marusic & Heuer 2007), with a slope that can be computed by finding the maximum spatial correlation $\mathcal{R}_{ij}(\delta x)$ between the inputs at the wall (index i) and the outputs (index j), with δx representing the distance in the streamwise direction at which the correlation is computed. By including a streamwise *shift* in the output fields, it is possible to obtain the maximum correlation at $\delta x = 0$, ensuring that the footprint of the coherent structure is included in the receptive field of the output. The use of such a shift was also discussed by Sasaki *et al.* (2019)

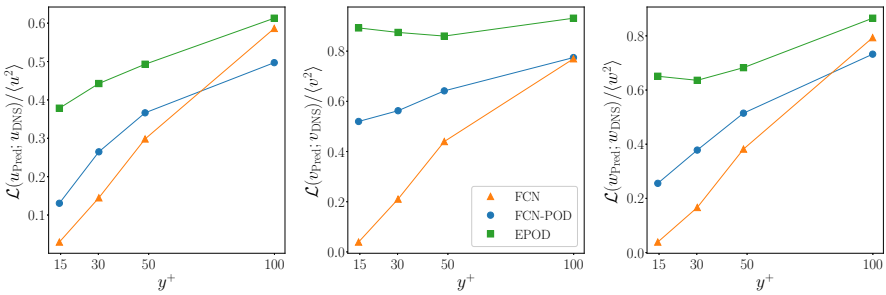

Figure 9: Mean-squared-error in the instantaneous fields (scaled with the corresponding RMS components) predicted by the three models at $Re_\tau = 550$, for the streamwise (left), wall-normal (middle) and spanwise (right) velocity fluctuations.

in a similar context. By considering the maximum correlation between the wall-shear stress in the streamwise direction and the streamwise velocity at a certain y^+, we obtain an angle:

$$\theta_{\mathrm{shift}} = \tan^{-1}\left(\frac{\arg\max_{\delta x} R_{11}(\delta x)}{\delta y}\right) \approx 11^\circ \tag{13}$$

at $Re_\tau = 180$, where δy indicates the considered distance from the wall. The obtained value is slightly lower than previous observations (Marusic & Heuer 2007; Sasaki *et al.* 2019), possibly because of the low Reynolds number. The inclusion of the shift information has been tested using two alternative implementations: first, by modifying the target output field, *i.e.* considering a field that has been sampled later in the simulation, although the accuracy of the introduced shift is limited by the value of the sampling period. The second approach makes use of the periodicity of the output fields, which are translated in the streamwise direction until the maximum correlation is obtained at $\delta x = 0$. This approach allows to more accurately introduce the shift, however in this case the underlying hypothesis is that the shift is sufficiently small so that temporal dynamics modify the flow in a negligible manner. None of the two shift implementations provided the expected improvement, and we observed a significant degradation of the prediction performance. These results could possibly be explained by the fact that coherent structures of different size have different inclinations, and imposing a single value is detrimental for the overall network performance, despite having chosen the angle that provides the maximum spatial correlation. Furthermore, the quality of the predictions is measured using the MSE between the prediction and the reference: this error indicator considers all wavelengths at the same time, without considering how the different wavelengths are affected by the shift. Further investigation of this aspect, also at higher Re_τ, will be conducted in future work.

3.2. *Predictions of turbulence statistics*

By averaging over the fields obtained from the neural-network models and EPOD, it is possible to evaluate the turbulence statistics of the predicted flow. First we consider the dataset at $Re_\tau = 180$: the predicted RMS fluctuations of the three components are shown in figure 10, together with the reference DNS profiles. The error in these statistical quantities is defined as:

$$E^+_{\mathrm{RMS}}(u) = \frac{\left| u^+_{\mathrm{RMS,Pred}} - u^+_{\mathrm{RMS,DNS}} \right|}{u^+_{\mathrm{RMS,DNS}}}, \tag{14}$$

for the streamwise component, and similarly for the other two components. As above, the subscripts 'DNS' and 'Pred' refer to the reference and predicted profiles, respectively. An important premise is that neither of the neural-network-based models is explicitly optimized to reproduce the statistics of the original simulation. This prevents the neural networks from learning only the average behaviour of the flow, however the predictions may be less statistically accurate, with the aim of maximizing the instantaneous performance. Note that here we favor instantaneous performance because our motivation is to use non-intrusive sensing for closed-loop flow control.

The prediction errors in the various RMS profiles are summarized in table 3, and they are averaged over the different training runs for the FCN and FCN-POD models. Note that the average is performed over the fluctuation-intensity values and not on the predictions, because that would alter the statistical properties of the predicted flow fields. The comparison of the errors from the different models shows that the statistical performance mimics the one observed for the instantaneous predictions at $y^+ = 15$ and $y^+ = 30$, with the FCN performing better than the FCN-POD and EPOD models. Furthermore, the FCN model provides a similar performance for the fluctuations of all three velocity components, while POD-based methods are more accurate in the predictions of u^+_{RMS}. This is related to the choice of not scaling the different velocity components in the FCN-POD and EPOD approaches, and the fact that near the wall the most energetic dynamics of the flow are in the streamwise direction. Taking into account the standard deviation in the results of the neural-network-based methods, the three models provide similar error levels at $y^+ = 50$. At $y^+ = 100$ the scenario is opposite to what we observed close to the wall: the FCN exhibits the highest errors, while the EPOD provides the best results. The error in the prediction of u^+_{RMS} from the FCN-POD model is between those of the two other models, while the wall-normal and spanwise intensities are closer to the errors from the FCN, due the reasons outlined above.

The statistical analysis is repeated also for the models trained at $Re_\tau = 550$, with the predicted RMS fluctuations shown in figure 11 and the relative error with respect to the reference simulation in table 4. FCN-based models do not show a significant variation in the prediction of the streamwise fluctuations with respect to the results at $Re_\tau = 180$, whereas the EPOD exhibits higher errors at this Reynolds number (also in the other two fluctuating components). The FCN

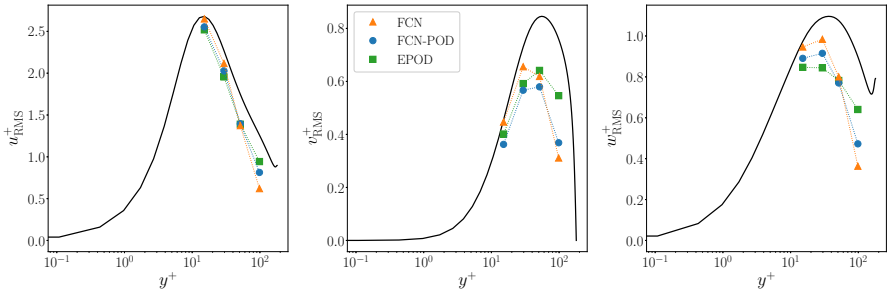

Figure 10: Comparison between the DNS (——) and the predictions of streamwise (left), wall-normal (middle) and spanwise (right) velocity fluctuations at $Re_\tau = 180$.

$E^+_{\mathrm{RMS}}(\cdot)$ [%]		$y^+ = 15$	$y^+ = 30$	$y^+ = 50$	$y^+ = 100$
	EPOD	6.03	13.87	20.50	25.15
u	FCN	1.16 ($\pm$0.74)	6.79 ($\pm$1.31)	21.47 ($\pm$1.97)	50.82 ($\pm$2.19)
	FCN-POD	4.70 ($\pm$0.02)	10.70 ($\pm$0.02)	20.15 ($\pm$0.03)	35.46 ($\pm$0.02)
	EPOD	11.68	18.89	23.97	28.10
v	FCN	1.74 ($\pm$1.00)	10.18 ($\pm$1.67)	26.66 ($\pm$0.76)	59.05 ($\pm$1.61)
	FCN-POD	20.29 ($\pm$0.02)	22.32 ($\pm$0.02)	31.32 ($\pm$0.01)	51.48 ($\pm$0.04)
	EPOD	13.01	22.48	27.27	28.72
w	FCN	2.79 ($\pm$0.36)	9.65 ($\pm$1.07)	25.60 ($\pm$1.214)	59.59 ($\pm$1.310)
	FCN-POD	8.50 ($\pm$0.04)	15.95 ($\pm$0.06)	28.38 ($\pm$0.004)	47.42 ($\pm$0.001)

Table 3: Percentage error in the prediction of the various RMS fluctations at the different wall-normal locations. Results at $Re_\tau = 180$.

has a consistent behaviour also for the fluctuations in the y- and z-directions, however the FCN-POD performs slightly worse than before, following the same trend but with higher error levels. The FCN-POD method outperforms the FCN approach only at $y^+ = 100$, confirming the results of the instantaneous performance at $Re_\tau = 550$.

3.3. *Predictions of power-spectral density*

The energetic scales present in the predicted fields, as well as their associated energy, are compared with those in the reference DNS data through spectral analysis. In figure 12 we show the pre-multiplied two-dimensional power-spectral density of the streamwise, wall-normal and spanwise fluctuations (denoted by ϕ_{uu}, ϕ_{vv} and ϕ_{ww} respectively) at $Re_\tau = 180$, where λ_x and λ_z denote the streamwise and spanwise wavelengths, whereas k_x and k_z are the corresponding wavenumbers. These results confirm the observations made in §3.1: at $y^+ = 15$,

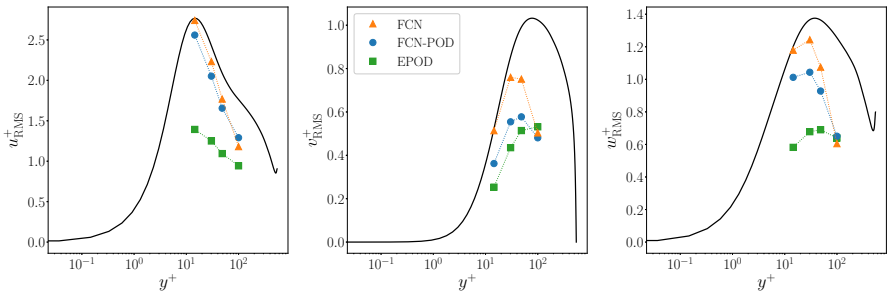

Figure 11: Comparison between the DNS (——) and the predictions of streamwise (left), wall-normal (middle) and spanwise (right) velocity fluctuations at $Re_\tau = 550$.

$E^+_{\mathrm{RMS}}(\cdot)$ [%]		$y^+ = 15$	$y^+ = 30$	$y^+ = 50$	$y^+ = 100$
	EPOD	49.69	48.53	47.62	46.53
u	FCN	0.98 (±0.66)	8.10 (±0.62)	15.33 (±0.22)	33.06 (±0.30)
	FCN-POD	7.53 (±0.27)	15.53 (±0.39)	20.75 (±0.85)	26.79 (±0.36)
	EPOD	51.57	49.13	48.13	48.07
v	FCN	1.74 (±0.11)	11.21 (±1.41)	24.20 (±1.38)	50.82 (±0.26)
	FCN-POD	30.73 (±0.06)	35.20 (±0.09)	41.75 (±0.10)	53.04 (±0.18)
	EPOD	51.56	50.35	49.42	49.08
w	FCN	1.86 (±0.60)	9.03 (±0.31)	21.21 (±1.27)	51.83 (±0.38)
	FCN-POD	15.74 (±0.05)	23.63 (±0.07)	31.97 (±0.10)	48.20 (±0.24)

Table 4: Percentage error in the prediction of the various RMS fluctations at the different wall-normal locations. Results at $Re_\tau = 550$.

all the considered models are able to correctly predict the energy content of the flow at all wavelengths, with the FCN slightly outperforming the two POD-based approaches. Note that the FCN-POD model is able to reconstruct the energy content of the flow at wavelengths that are longer than the size of the subdomains, proving that this is not a limiting factor for the model. However, a small jump, probably due to a lack of smoothness at the edges of the subdomains, can be observed in the streamwise wavelength for the 10%-energy level in the wall-normal and spanwise components. These jumps are found at a wavelength $\lambda_x^+ \approx 180$, corresponding to the subdomain size. At $y^+ = 30$ there is a slight energy attenuation which becomes increasingly more noticeable when the predicted flow is farther away from the wall. At $y^+ = 100$ the POD-based methods perform better than the FCN model, a fact that can be explained by considering two concurring aspects: the first is that POD methods only predict the temporal dynamics of the system, thus the overall energy-scale distribution

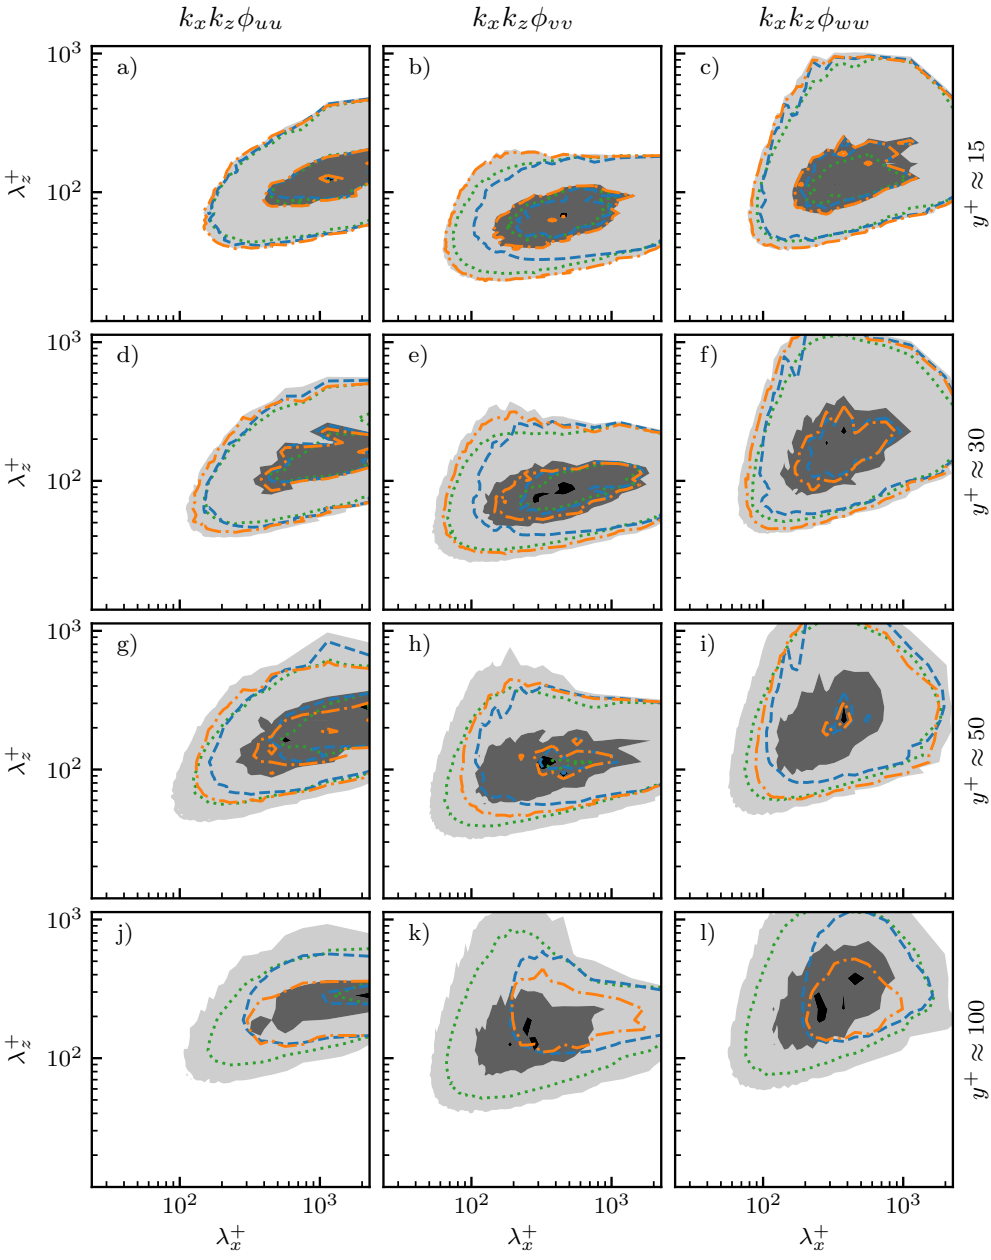

Figure 12: Pre-multiplied two-dimensional power-spectral densities for $Re_\tau = 180$. The contour levels contain 10%, 50% and 90% of the maximum DNS power-spectral density. Shaded contours refer to the reference data, while contour lines refer to (—·—) FCN, (– –) FCN-POD and (·····) EPOD predictions, respectively.

stored in the POD spatial modes does not need to be predicted. This allows to reconstruct more than 50% of the flow fields, at least in the streamwise component. The second aspect is the fact that the receptive field from the FCN method, while sufficient for planes closer to the wall, is not large enough to reproduce the large scales present at larger y^+.

Figure 13 reports the reference and predicted power-spectral densities at $Re_\tau = 550$. As opposed to what was observed for the instantaneous predictions and the turbulence statistics, the spectra highlight the differences and similarities between the models used for the two Reynolds numbers. As noted above, the FCN architecture is the same for both Re_τ, however this implies that the receptive field is smaller at higher Reynolds number when it is measured in outer units. This can potentially help in the prediction of the small scales, but it can also be detrimental for the larger scales. Note that the different size of the receptive field does not seem to affect the predicted energy content of the FCN, which shows the same trends as for the low-Re_τ case. On the other hand, the spectra of the FCN-POD with 12×12 subdomains (the same amount as in the low-Re_τ case, not shown) exhibits spurious periodic peaks due to the tiling. This observation motivated the increased number of subdomains considered at $Re_\tau = 550$. As in the low-Re_τ case, the FCN-POD approach is able to reconstruct the scales larger than the subdomain size, but the countour line exhibits a small jump in the streamwise wavelength for the 10%-energy level. This jump appears at a wavelength $\lambda_x^+ \approx 200$, corresponding to the subdomain size employed in the high-Re_τ case. This jump is also appreciated in the streamwise component at $y^+ = 15$. The POD-based approaches are outperformed by the FCN in the range of $y^+ = 15 - 50$. Farther from the wall, the accuracy of the FCN is matched by the FCN-POD method. It is interesting to note that the EPOD does not follow the same attenuation process as the FCN-based methods in the wall-normal and spanwise components. As one moves farther from the wall, the FCN-based methods fail to reproduce a wider range of small scales, whereas the EPOD exhibits more difficulties predicting the large scales. When it comes to the spectral peaks far from the wall, while the FCN-based methods produce noisier predictions than in the low-Re_τ case, the EPOD is not able to reproduce that part of the spectra.

4. Transfer learning

A number of more advanced techniques have also started to be adopted from the specialized machine-learning literature and applied to fluid-dynamics research. One notable example is *transfer learning* (Pan & Yang 2009), a method that allows to transfer knowledge from one neural-network model to another one, thus reducing the amount data and time required for training. Guastoni *et al.* (2020*a*) showed that the training time at a given wall-normal location may be significantly reduced if the network parameters are initialized using the optimized parameters of a previously-trained network at another wall-normal location. Similarly, Kim & Lee (2020) used the convolutional network trained at a low Reynolds number to predict the flow at a higher Reynolds number.

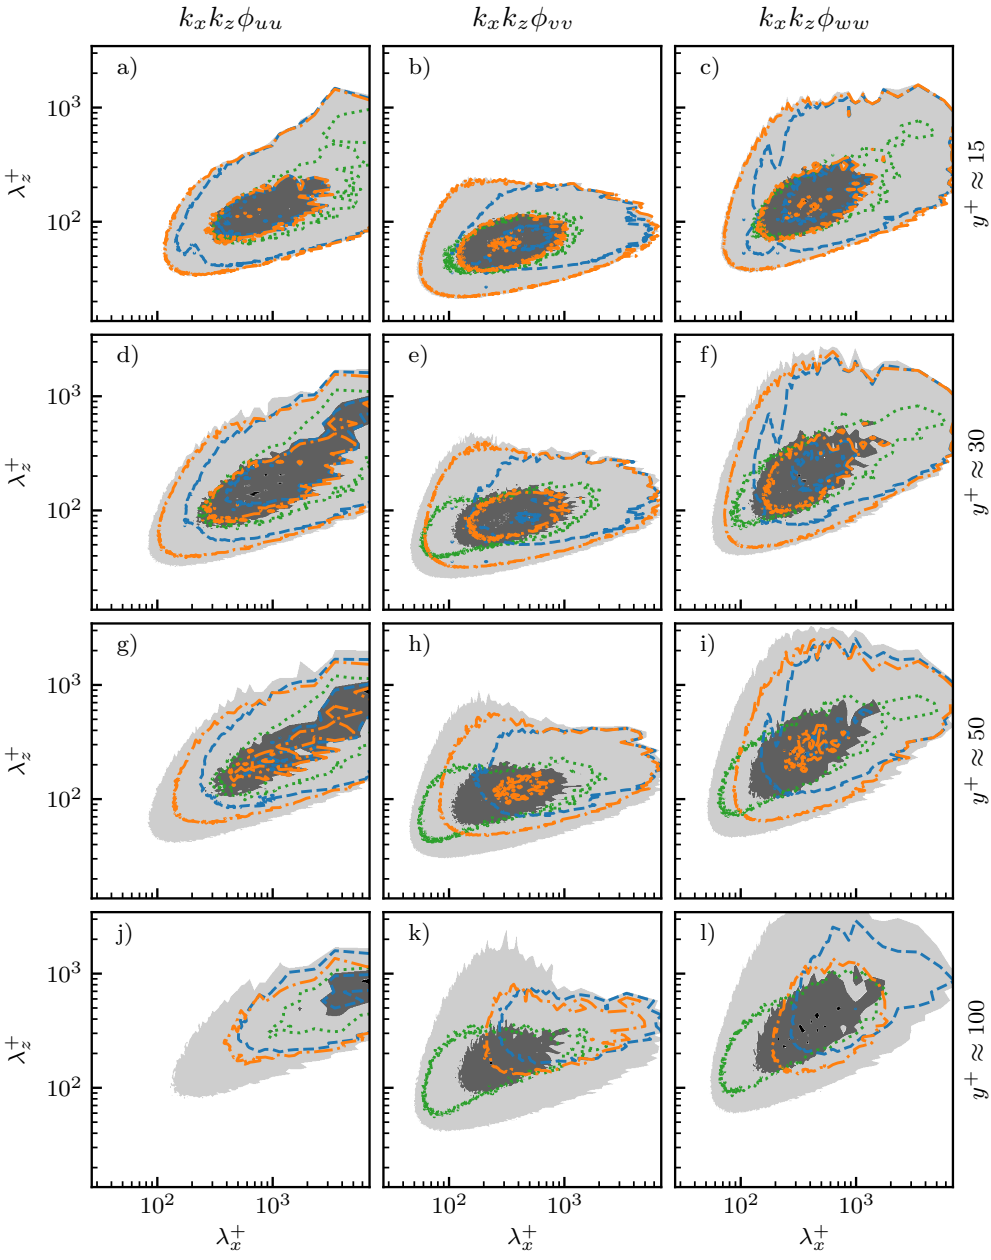

Figure 13: Pre-multiplied two-dimensional power-spectral densities for $Re_\tau = 550$. The contour levels contain 10%, 50% and 90% of the maximum DNS power-spectral density. Shaded contours refer to the reference data, while contour lines refer to (—·—) FCN, (— —) FCN-POD and (·····) EPOD predictions, respectively.

Transfer learning represents an appealing solution for the main drawback of neural networks, which is the need to train them with a sufficient amount of data. Training typically requires specialized hardware and in our specific application the computational cost of generating the training and test datasets is not negligible. Furthermore, this cost grows as Re_τ increases, making the generation of training data through DNS unfeasible at the Reynolds numbers that are relevant for engineering applications. In this regard, it is important to make an efficient use of the data and the trained models at our disposal. In this work, the possibility of transferring knowledge between models trained at different friction Reynolds numbers is investigated. At a fixed wall-normal distance, the weights of the FCN model trained on the dataset at $Re_\tau = 180$ are loaded before training the network with the higher-Re_τ dataset. This is possible because the network has the same amount of trainable parameters in both cases, as noted above. The learning rate is the only parameter that needs to be modified: a lower value has to be set, in order to prevent the optimizer from diverging too quickly from the weight configuration used for initialization. While in Guastoni *et al.* (2020*a*) we froze the first layers of the initialized network because the input was the same at the different wall-normal locations, in this case all the layers are trainable because the input distribution changes from one Re_τ to the other one.

The transfer-learning numerical experiments that follow are designed with two objectives in mind. First, we investigate the effect of using the parameters of previously-trained model to initialize a model that will be trained at a different target Re_τ than the first model. Second, we verify whether the non-random initialization that we described can potentially improve the training process, reducing the amount of data at higher Re_τ that is needed to achieve the same performance of a model trained from scratch, using the entire training/validation dataset, as we did in section 3. To this end, we initialized a FCN model with the parameters learnt on the $Re_\tau = 180$ dataset and we trained such model on the full training/validation datasets at $Re_\tau = 550$. Subsequently, the same initialization is used for models that are trained on a reduced dataset at $Re_\tau = 550$, namely 25% and 50% of the full dataset. Differently from the previous sections, only one training run was performed for each case. In order to compare models trained with datasets of different sizes, we considered the number of weight updates through the optimization algorithm during training. In figure 14 the validation and test losses are compared for the models trained with the full dataset and a random initialization. When the initialized model is trained on the full dataset, the performance is consistently better than that of the random initialization, both in terms of validation and test loss. The improvement is more evident close to the wall, whereas at $y^+ = 100$ the two models provide approximately the same results after the first 150,000 updates. Transferring knowledge between different Reynolds numbers is then not only feasible, but also advantageous in terms of performance when the same amount of data is considered. Note that we kept the same ratio (4:1) when dividing training and validation sets even when the size of the dataset was reduced.

$E^+_{\mathrm{RMS}}(\cdot)$ [%]		$y^+ = 15$	$y^+ = 30$	$y^+ = 50$	$y^+ = 100$
u	Ref.	0.98 ($\pm$0.66)	8.10 ($\pm$0.62)	15.33 ($\pm$0.22)	33.06 ($\pm$0.30)
	100%	1.22	7.15	16.09	32.75
	50%	2.94	7.11	16.33	34.11
	25%	1.15	7.74	14.78	33.78
v	Ref.	1.74 ($\pm$0.11)	11.21 ($\pm$1.41)	24.20 ($\pm$1.38)	50.82 ($\pm$0.26)
	100%	1.86	9.40	24.40	51.59
	50%	2.40	9.46	25.96	50.90
	25%	1.71	11.33	23.15	50.43
w	Ref.	1.86 ($\pm$0.60)	9.03 ($\pm$0.31)	21.21 ($\pm$1.27)	51.83 ($\pm$0.38)
	100%	1.75	8.65	21.05	53.04
	50%	2.61	7.70	21.34	52.60
	25%	1.22	9.67	20.35	49.75

Table 5: Percentage error in the prediction of the various RMS fluctations at the different wall-normal locations from models initialized with parameters from the $Re_\tau = 180$ FCN. The statistics of the different initialized models are computed after 250,000 updates, and they are shown together the models with full dataset and random initialization, which is included as a reference. Results at $Re_\tau = 550$.

In particular, when 50% or 25% of the samples are used, the validation set becomes too small to provide a reliable estimate of the error. On the other hand, the size of the test dataset is the same as the previous experiments. Up to $y^+ = 50$, the initialized networks are able to provide a performance that is very similar to that of the reference model with the same number of updates, with significant savings in terms of amount of data needed to train the network. On the other hand, at $y^+ = 100$ a sufficient number of samples becomes a necessary condition to ensure the convergence to an optimal configuration: the loss of the network trained with 50% of the training dataset does not improve after the first 100,000 updates, while with 25% of the original dataset the network exhibits overfitting. The initialized models are able to provide a comparable accuracy also from the statistical point of view, as reported in table 5. We stress once again that the networks are not explicitly optimized to reduce the error in the statistics and that small variations in these error figures can be ascribed to the stochastic nature of the optimization algorithm. Overall, these results demonstrate the feasibility of knowledge transfer from models at different Re_τ: with careful tuning of the hyperparameters it should be possible to substantially reduce the training time, as well as the amount of data needed for training. Although not tested, the transfer between different wall-normal locations described in Guastoni *et al.* (2020*a*) is still applicable in this case, thus enabling a more efficient prediction of the flow at different wall-normal locations.

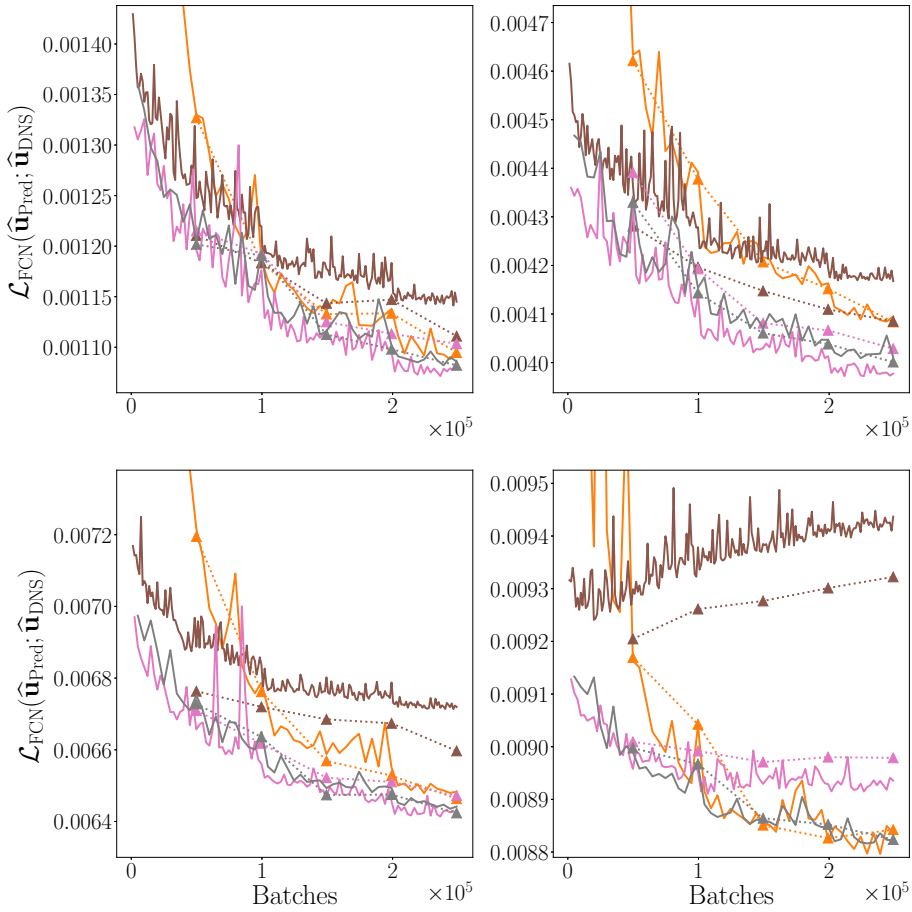

Figure 14: Validation (——) and test (- - -) loss in the FCN prediction at (from left to right, top to bottom): $y^+ = 15$, 30, 50 and 100. Orange represents the models trained with the full dataset and random initialization, grey the models trained with the full dataset and initialized with previously-trained networks, pink and brown represent models initialized with the parameters from the $Re_\tau = 180$ network, trained with 50% and 25% of the original dataset, respectively.

5. Conclusions

In this work, we introduced and compared two different models based on fully-convolutional neural networks, for prediction of the velocity fluctuations at a given wall-normal distance, using quantities measured at the wall as inputs. The

FCN and FCN-POD models are improved versions of previous architectures, used by Guastoni *et al.* (2020*a*) and Güemes *et al.* (2019), respectively. Both of them are able to provide predictions in very good agreement with the reference data, simulated by means of the pseudo-spectral DNS code SIMSON (Chevalier *et al.* 2007), up to $y^+ = 50$. Such an agreement is verified by comparing the error in instantaneous predictions, turbulence statistics (namely RMS fluctuations) and the energy content at the different wavelengths (*i.e.* spectral analysis). Both models show better prediction capabilities than EPOD (which is a linear method) in almost all the wall-normal locations and investigated features, thanks to their ability to predict nonlinear scale interactions. Furthermore, we showed that these architectures can be used at two different friction Reynolds numbers ($Re_\tau = 180$ and 550) with minimal modifications, providing satisfactory results on both datasets.

The two models are designed under the assumption that local information at the wall is sufficient to predict the flow farther away, however the FCN-POD model partially encodes further physical information of the system by using the spatial modes obtained through POD of the training dataset. On the other hand, features like the periodicity of the flow are enforced in the FCN by exploiting the mathematical characteristics of the model. These architectural differences are associated with performance discrepancies at the tested wall-normal locations: the FCN provides higher accuracy than the FCN-POD model closer to the wall, *i.e.* up to $y^+ = 30$ at $Re_\tau = 180$ and up to $y^+ = 50$ at $Re_\tau = 550$. Farther from the wall, the FCN-POD method produces the most accurate predictions. The choice between these two models is motivated by the application into which the prediction model is integrated.

Despite the encouraging results discussed here, both models can be improved in terms of network architecture and training. An attempt to embed further physical information into the FCN did not result in improved predictions, as reported in §3.1.1. The correct way of incorporating this information to enhance the predictions is an active area of research. As another example, the high-frequency noise in the FCN predictions could be reduced with appropriate filtering, possibly adding a trainable layer to the network to perform this operation. The FCN-POD model has a higher number of hyperparameters to be set, such as the number of predicted temporal modes or the size of the subdomains. A more thorough inspection of the hyperparameter space may provide a significant improvement in the prediction performance. Differently from the FCN, in the FCN-POD model the velocity components are not scaled to have the same magnitude: such a modification could help to predict the wall-normal and spanwise components of the velocity more accurately, even though it would also modify the POD mode sorting because of the different energy norm. Furthermore, the FCN-POD results exhibit lack of smoothness at the subdomain edges in the flow predictions. Finally, both models are trained to minimize a loss function based on the instantaneous error. Such a function

could be modified to improve other physical characteristics of the predicted flow, for example the turbulence statistics and the spectral energy content.

To reduce the training time in view of industrial applications, the implementation of transfer learning was tested for the FCN model. Transfer learning can exploit a network trained at a lower Reynolds number to provide the weight initialization for training at a higher Reynolds number, thus reducing the requirements in terms of training time and data. The results are very encouraging, showing that it is possible to train the network with 50% and even 25% of the original training dataset, obtaining a performance similar to that of the reference model up to $y^+ = 50$.

Once the neural networks are trained, they are computationally cheap to evaluate, and they can become even cheaper by *pruning* the parts that have a negligible contribution to the final result off the network. Such an operation is not possible *a priori*, since the training determines how the inputs have to be processed to obtain the output. By reducing the computational cost of the evaluation it is possible to deploy the model using low-powered hardware and/or potentially run it in real-time. Thus, the proposed FCN-based methods could be used for non-intrusive sensing of the flow, which is needed for closed-loop control applications. Furthermore, since the FCN models are able to reproduce non-linear interactions in wall-bounded turbulence, new promising avenues in turbulence research could be opened by the network interpretation (Fan *et al.* 2021), as shown by Iten *et al.* (2020), who demonstrated that neural networks can provide relevant physical insights.

Acknowledgements

All the codes employed in this work will be released as open-source on GitHub upon publication in the peer-reviewed literature. The authors acknowledge the funding provided by the Swedish e-Science Research Centre (SeRC), the Göran Gustafsson Foundation and the Knut and Alice Wallenberg (KAW) Foundation. The numerical simulations were carried out on resources provided by the Swedish National Infrastructure for Computing (SNIC) at PDC and HPC2N. Part of this study was conducted in the context of the 4th Madrid Turbulence Workshop, in the context of the COTURB project (Coherent Structures in Wall-bounded Turbulence), funded by the European Research Council (ERC), under grant ERC-2014.AdG-669505. S D and AI were partially supported by the Grant DPI2016-79401-R funded by the Spanish State Research Agency (SRA) and European Regional Development Fund (ERDF).

Declaration of Interests

The authors report no conflict of interest.

Appendix A.

This Appendix contains the training details of the neural network models in the FCN and FCN-POD method. Both models introduced in section 2 need

to be tuned to perform optimally on the chosen dataset. The networks were trained using the Adam (Kingma & Ba 2014) optimization algorithm for 50 epochs, with a scheduled exponential learning-rate decay. We used the optimizer hyperparameters suggested in the original paper, except the $\hat{\epsilon}$ parameter that was set to 0.1 following TensorFlow recommendations (Abadi *et al.* 2016). The total number of trainable parameters in the FCN is 1,264,131 and it does not depend on the Re_τ of the dataset. On the other hand, this number is Re_τ-dependent for the FCN-POD model, as the number of subdomains and reconstructed POD modes changes. For the $Re_\tau = 180$ case the number of trainable parameters is 4,733,248, while for the $Re_\tau = 550$ case it is 5,028,224.

For the FCN model the $Re_\tau = 180$ dataset, composed by input and output fields, has a size of about 70 GB per wall-normal distance, while the $Re_\tau = 550$ dataset occupies 120 GB. Depending on the amount of available RAM it might not be possible to load the entire dataset in memory, however our implementation does not require to load it at once, allowing training on less-performing computer as well. The training was performed on two different GPUs, namely NVIDIA K80 and NVIDIA RTX2080Ti. The former has 24 GB of GPU memory, whereas the latter only 12 GB. The turbulent field snapshots have a comparatively high resolution with respect to the samples used in image processing and the GPU memory can also be a limiting parameter. For a given dataset, the amount of memory that is allocated on the GPU depend on the batch size, that is the number of samples used to estimate the loss function gradient at each update of the learnable parameters of the network. In the $Re_\tau = 550$, the batch size is limited to 2 samples when the faster RTX2080Ti is used. With this setting the gradient approximation might be too noisy and the training performance might be degraded, therefore we suggest to use at least a batch size of 4, using two RTX2080Ti GPUs.

Due to the reduced-order nature of FCN-POD method, the size of the training/validation datasets is lower: 40 GB for $Re_\tau = 180$ dataset and 70 GB for $Re_\tau = 550$ dataset. Thanks to the max pooling operations that are not present in the FCN model, the size of the tensors allocated on GPU is smaller and this allows to train the model with a batch size of 8 for both datasets.

Appendix B.

This Appendix contains the wall-normal and spanwise fluctuations corresponding to the fields shown in figures 6 and 8 for $Re_\tau = 180$ and 550, respectively. If we observe the wall-normal fluctuations at $Re_\tau = 180$ in figure 15, the FCN model provides better predictions than POD-based methods at $y^+ = 15$. The tiling of the subdomains is more apparent in the FCN-POD predictions, while the EPOD predictions contain small features that are not present in the DNS reference field. Similar observations can be made at $y^+ = 30$. As we move farther away from the wall, at $y^+ = 50$, the accuracy of the neural-network-based models is similar. For the FCN-POD the flow features in the streamwise direction appear smoothed out with respect to the DNS fields; on the other hand the

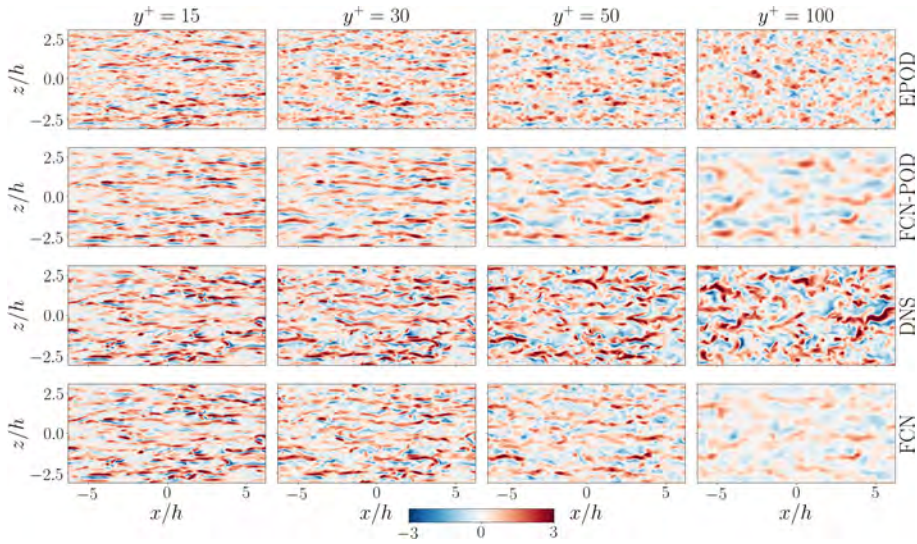

Figure 15: Comparison of the wall-normal fluctuation fields at $Re_\tau = 180$, scaled with the corresponding v_{RMS}, from EPOD (1st row), FCN-POD (2nd row), reference DNS (3rd row) and FCN (4th row). Results at $y^+ = 15$ (1st column), $y^+ = 30$ (2nd column), $y^+ = 50$ (3rd column) and $y^+ = 100$ (4th column).

same features have more jagged edges in the FCN predictions, possibly as a consequence of the localized nature of the predictions. As observed for the streamwise velocity fluctuations, the performance degrades as we move farther away from the wall. Note that some features are completely missing in the EPOD predictions, due to the fact that the three velocity components are predicted at the same time. From the observation of the spanwise fluctuations predictions, the behaviour of the three methods is similar to the one in the wall-normal components. Note however that, closer to the wall, the FCN-POD provides more accurate predictions for the spanwise fluctuation component than for the wall-normal component, as it can be observed in figure 7.

When a higher Reynolds number is considered, the error trends of the neural-network-based models in figure 9 are comparable to the one observed at $Re_\tau = 180$, for the wall-normal and spanwise fluctuations. On the other hand, the EPOD predictions improve when moving from $y^+ = 15$ to $y^+ = 30$, before the performance starts degrading as we move farther from the wall. Additionally, figures 17 and 18 show that, close to the wall, the intensity of the fluctuations in the wall-normal and spanwise direction is significantly underestimated by the POD-based methods, when compared to FCN. As the wall-normal distance increases, the fluctuations range predicted by the FCN reduces as well and at $y^+ = 100$ the range is comparable for all three methods.

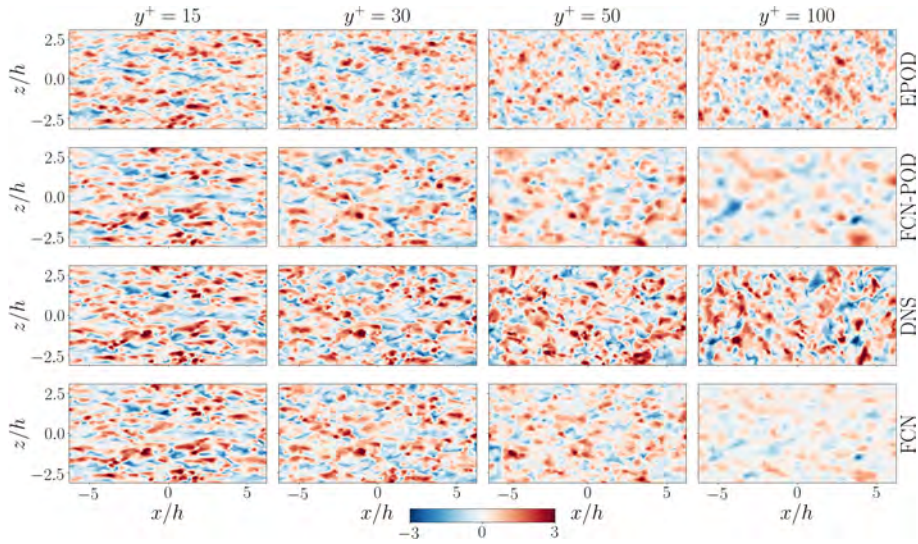

Figure 16: Comparison of the spanwise fluctuation fields at $Re_\tau = 180$, scaled with the corresponding w_{RMS}, from EPOD (1st row), FCN-POD (2nd row), reference DNS (3rd row) and FCN (4th row). Results at $y^+ = 15$ (1st column), $y^+ = 30$ (2nd column), $y^+ = 50$ (3rd column) and $y^+ = 100$ (4th column).

References

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M. *et al.* 2016 Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pp. 265–283.

Adrian, R. 1988 Stochastic estimation of organized turbulent structure: Homogeneous shear flow. *J. Fluid Mech.* **190**, 531–559.

Baars, W. J. & Tinney, C. E. 2014 Proper orthogonal decomposition-based spectral higher-order stochastic estimation. *Phys. Fluids* **26** (5), 055112.

Baldi, P. & Hornik, K. 1989 Neural networks and principal component analysis: Learning from examples without local minima. *Neural Networks* **2** (1), 53–58.

Beck, A., Flad, D. & Munz, C.-D. 2019 Deep neural networks for data-driven LES closure models. *J. Comput. Phys.* **398**, 108910.

Borée, J. 2003 Extended proper orthogonal decomposition: a tool to analyse correlated events in turbulent flows. *Exp. Fluids* **35** (2), 188–192.

Bourgeois, J., Noack, B. & Martinuzzi, R. 2013 Generalized phase average with applications to sensor-based flow estimation of the wall-mounted square cylinder wake. *J. Fluid Mech.* **736**, 316–350.

Bourlard, H. & Kamp, Y. 1988 Auto-association by multilayer perceptrons and singular value decomposition. *Biol. Cybern.* **59** (4-5), 291–294.

Brunton, S. L., Noack, B. R. & Koumoutsakos, P. 2020 Machine learning for fluid mechanics. *Annu. Rev. Fluid Mech.* **52**, 477–508.

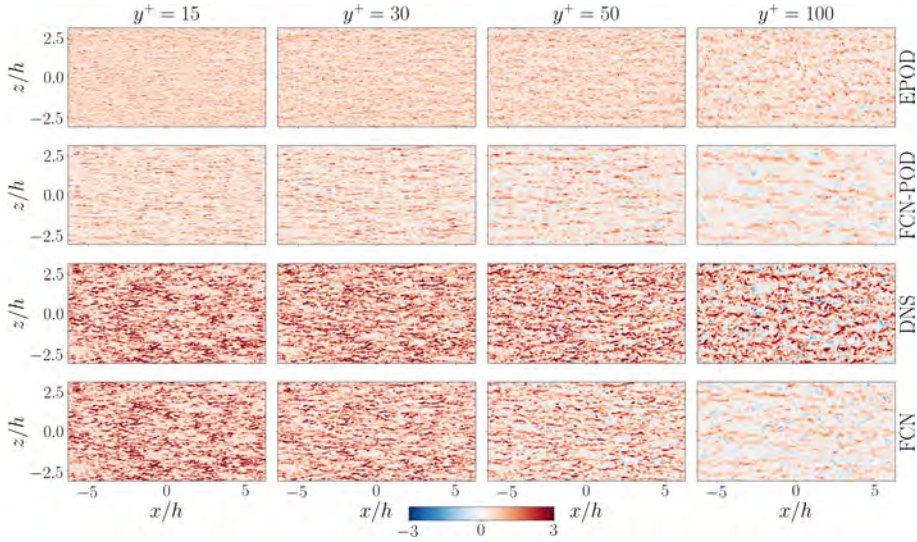

Figure 17: Comparison of the wall-normal fluctuation fields at $Re_\tau = 550$, scaled with the corresponding v_{RMS}, from EPOD (1st row), FCN-POD (2nd row), reference DNS (3rd row) and FCN (4th row). Results at $y^+ = 15$ (1st column), $y^+ = 30$ (2nd column), $y^+ = 50$ (3rd column) and $y^+ = 100$ (4th column).

Bucci, M. A., Semeraro, O., Allauzen, A., Wisniewski, G., Cordier, L. & Mathelin, L. 2019 Control of chaotic systems by deep reinforcement learning. *Proc. R. Soc. A* **475**, 20190351.

Chevalier, M., Schlatter, P., Lundbladh, A. & Henningson, D. S. 2007 A pseudospectral solver for incompressible boundary layer flows. *Tech. Rep.*. TRITA-MEK 2007:07. KTH Mechanics, Stockholm, Sweden.

Choi, H., Moin, P. & Kim, J. 1994 Active turbulence control for drag reduction in wall-bounded flows. *J. Fluid Mech.* **262**, 75–75.

De Fauw, J., Ledsam, J. R., Romera-Paredes, B., Nikolov, S., Tomasev, N., Blackwell, S., Askham, H., Glorot, X., O'Donoghue, B., Visentin, D. *et al.* 2018 Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nat. Med.* **24** (9), 1342–1350.

Discetti, S., Bellani, G., Örlü, R., Serpieri, J., Sanmiguel Vila, C., Raiola, M., Zheng, X., Mascotelli, L., Talamelli, A. & Ianiro, A. 2019 Characterization of very-large-scale motions in high-Re pipe flows. *Exp. Therm. Fluid Sci.* **104**, 1–8.

Discetti, S., Raiola, M. & Ianiro, A. 2018 Estimation of time-resolved turbulent fields through correlation of non-time-resolved field measurements and time-resolved point measurements. *Exp. Therm. Fluid Sci.* **93**, 119–130.

Dogan, E., Örlü, R., Gatti, D., Vinuesa, R. & Schlatter, P. 2019 Quantification of amplitude modulation in wall-bounded turbulence. *Fluid Dyn. Res.* **51**, 011408.

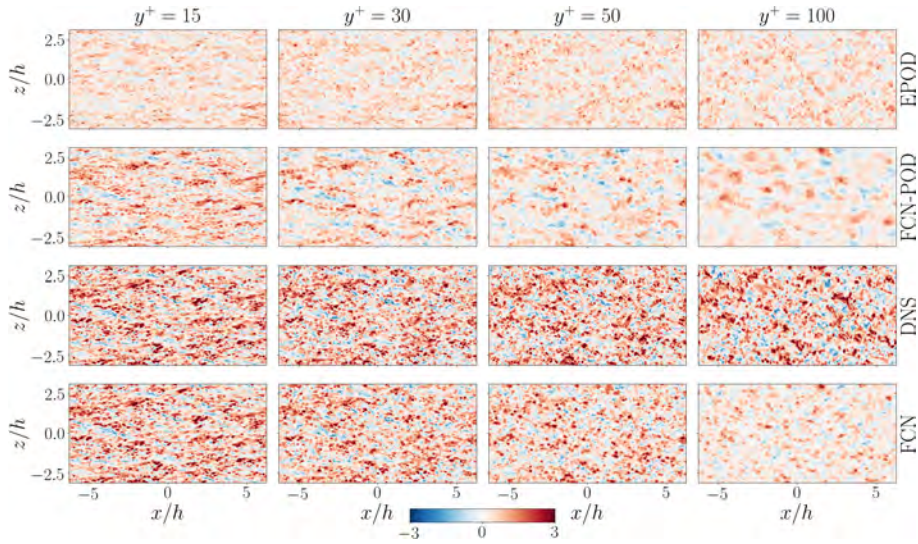

Figure 18: Comparison of the spanwise fluctuation fields at $Re_\tau = 550$, scaled with the corresponding w_{RMS}, from EPOD (1st row), FCN-POD (2nd row), reference DNS (3rd row) and FCN (4th row). Results at $y^+ = 15$ (1st column), $y^+ = 30$ (2nd column), $y^+ = 50$ (3rd column) and $y^+ = 100$ (4th column).

Dumoulin, V. & Visin, F. 2016 A guide to convolution arithmetic for deep learning. arXiv preprint arXiv:1603.07285.

Duraisamy, K., Iaccarino, G. & Xiao, H. 2019 Turbulence modeling in the age of data. *Annu. Rev. Fluid Mech.* **51**, 357–377.

Eivazi, H., Guastoni, L., Schlatter, P., Azizpour, H. & Vinuesa, R. 2021 Recurrent neural networks and koopman-based frameworks for temporal predictions in a low-order model of turbulence. *Int. J. Heat Fluid Fl.* **90**, 108816.

Encinar, M. P. & Jiménez, J. 2019 Logarithmic-layer turbulence: A view from the wall. *Phys. Rev. Fluids* **4** (11), 114603.

Erichson, N. B., Mathelin, L., Yao, Z., Brunton, S. L., Mahoney, M. W. & Kutz, J. N. 2020 Shallow neural networks for fluid flow reconstruction with limited sensors. *P. Roy. Soc. A - Math. Phys.* **476** (2238), 20200097.

Fan, F., Xiong, J., Li, M. & Wang, G. 2021 On interpretability of artificial neural networks: A survey. *IEEE T. Radiat. Plasma Med. Sci.* .

Fukami, K., Fukagata, K. & Taira, K. 2019*a* Super-resolution reconstruction of turbulent flows with machine learning. *J. Fluid Mech.* **870**, 106–120.

Fukami, K., Fukagata, K. & Taira, K. 2021 Machine-learning-based spatio-temporal super resolution reconstruction of turbulent flows. *J. Fluid Mech.* **909**.

Fukami, K., Nabae, Y., Kawai, K. & Fukagata, K. 2019*b* Synthetic turbulent inflow generator using machine learning. *Phys. Rev. Fluids* **4** (6), 064603.

Fukushima, K. 1980 Neocognitron: A self-organizing neural network model for a

mechanism of pattern recognition unaffected by shift in position. *Biol. Cybern.* **36** (4), 193–202.

Fukushima, K. 1988 Neocognitron: A hierarchical neural network capable of visual pattern recognition. *Neural Networks* **1** (2), 119–130.

Goodfellow, I., Bengio, Y., Courville, A. & Bengio, Y. 2016 *Deep learning*. MIT Press Cambridge.

Guastoni, L., Encinar, M. P., Schlatter, P., Azizpour, H. & Vinuesa, R. 2020 Prediction of wall-bounded turbulence from wall quantities using convolutional neural networks. *J. Phys. Conf. Ser.* **1522** (1), 012022.

Guastoni, L., Srinivasan, P. A., Azizpour, H., Schlatter, P. & Vinuesa, R. 2019 On the use of recurrent neural networks for predictions of turbulent flows. In *Proc. of the 11th Int. Symp. on Turbulence and Shear Flow Phenomena (TSFP11), Southampton, UK, July 30 - August 2*.

Güemes, A., Discetti, S. & Ianiro, A. 2019 Sensing the turbulent large-scale motions with their wall signature. *Phys. Fluids* **31** (12), 125112.

Guzmán Iñigo, J., Sipp, D. & Schmid, P. J. 2014 A dynamic observer to capture and control perturbation energy in noise amplifiers. *J. Fluid Mech.* **758**, 728–753.

Ham, Y.-G., Kim, J.-H. & Luo, J.-J. 2019 Deep learning for multi-year ENSO forecasts. *Nature* **573** (7775), 568–572.

Hinton, G. E. & Salakhutdinov, R. R. 2006 Reducing the dimensionality of data with neural networks. *Science* **313** (5786), 504–507.

Hornik, K., Stinchcombe, M. & White, H. 1989 Multilayer feedforward networks are universal approximators. *Neural Networks* **2** (5), 359–366.

Hosseini, Z., Martinuzzi, R. J. & Noack, B. R. 2016 Modal energy flow analysis of a highly modulated wake behind a wall-mounted pyramid. *J. Fluid Mech.* **798**, 717–750.

Ioffe, S. & Szegedy, C. 2015 Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proc. ICML*, pp. 448–456.

Iten, R., Metger, T., Wilming, H., Del Rio, L. & Renner, R. 2020 Discovering physical concepts with neural networks. *Phys. Rev. Letters* **124** (1), 010508.

Jagodinski, E., Zhu, X. & Verma, S. 2020 Uncovering dynamically critical regions in near-wall turbulence using 3D convolutional neural networks. arXiv preprint arXiv:2004.06187.

Jean, N., Burke, M., Xie, M., Davis, W. M., Lobell, D. B. & Ermon, S. 2016 Combining satellite imagery and machine learning to predict poverty. *Science* **353** (6301), 790–794.

Jiménez, J. 2018 Machine-aided turbulence theory. *J. Fluid Mech.* **854**.

Kim, J. & Lee, C. 2020 Prediction of turbulent heat transfer using convolutional neural networks. *J. Fluid Mech.* **882**.

Kingma, D. P. & Ba, J. 2014 Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.

Kutz, J. N. 2017 Deep learning in fluid dynamics. *J. Fluid Mech.* **814**, 1–4.

Lapeyre, C. J., Misdariis, A., Cazard, N., Veynante, D. & Poinsot, T. 2019 Training convolutional neural networks to estimate turbulent sub-grid scale reaction rates. *Combust. Flame* **203**, 255–264.

LeCun, Y., Bengio, Y. & Hinton, G. 2015 Deep learning. *Nature* **521** (7553), 436.

LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard,

W. & Jackel, L. D. 1989 Backpropagation applied to handwritten zip code recognition. *Neural Comput.* **1** (4), 541–551.

LeCun, Y., Bottou, L., Bengio, Y. & Haffner, P. 1998 Gradient-based learning applied to document recognition. *Proc. IEEE* **86** (11), 2278–2324.

Lee, C., Kim, J., Babcock, D. & Goodman, R. 1997 Application of neural networks to turbulence control for drag reduction. *Phys. Fluids* **9** (6), 1740–1747.

Ling, J., Kurzawski, A. & Templeton, J. 2016 Reynolds averaged turbulence modelling using deep neural networks with embedded invariance. *J. Fluid Mech.* **807**, 155–166.

Liu, Z., Adrian, R. & Hanratty, T. 2001 Large-scale modes of turbulent channel flow: transport and structure. *J. Fluid Mech.* **448**, 53–80.

Long, J., Shelhamer, E. & Darrell, T. 2015 Fully convolutional networks for semantic segmentation. In *Proc. CVPR IEEE*, pp. 3431–3440.

Lozano-Durán, A., Bae, H. J. & Encinar, M. P. 2020 Causality of energy-containing eddies in wall turbulence. *J. Fluid Mech.* **882**.

Lozano-Durán, A., Flores, O. & Jiménez, J. 2012 The three-dimensional structure of momentum transfer in turbulent channels. *J. Fluid Mech.* **694**, 100.

Lumley, J. L. 1967 The structure of inhomogeneous turbulent flows. In *Atmospheric turbulence and radio wave propagation*, pp. 166–178.

Ma, W. & Lu, J. 2017 An equivalence of fully connected layer and convolutional layer. arXiv preprint arXiv: 1712.01252.

Makarashvili, V., Merzari, E., Obabko, A., Siegel, A. & Fischer, P. 2017 A performance analysis of ensemble averaging for high fidelity turbulence simulations at the strong scaling limit. *Computer Physics Communications* **219**, 236 – 245.

Marusic, I. & Heuer, W. D. C. 2007 Reynolds number invariance of the structure inclination angle in wall turbulence. *Phys. Rev. Lett.* **99**, 114504.

Maulik, R., San, O., Rasheed, A. & Vedula, P. 2019 Subgrid modelling for two-dimensional turbulence using neural networks. *J. Fluid Mech.* **858**, 122–144.

Milano, M. & Koumoutsakos, P. 2002 Neural network modeling for near wall turbulent flow. *J. Comput. Phys.* **182** (1), 1–26.

Moehlis, J., Faisst, H. & Eckhardt, B. 2004 A low-dimensional model for turbulent shear flows. *New J. Phys.* **6** (1), 56.

Mokhasi, P., Rempfer, D. & Kandala, S. 2009 Predictive flow-field estimation. *Physica D* **238**, 290–308.

Murata, T., Fukami, K. & Fukagata, K. 2020 Nonlinear mode decomposition with convolutional neural networks for fluid dynamics. *J. Fluid Mech.* **882**.

Nair, N. J. & Goza, A. 2020 Leveraging reduced-order models for state estimation using deep learning. *J. Fluid Mech.* **897**, R1.

Nair, V. & Hinton, G. E. 2010 Rectified linear units improve restricted boltzmann machines. In *Proc. ICML*, pp. 807–814.

Norouzzadeh, M. S., Nguyen, A., Kosmala, M., Swanson, A., Palmer, M. S., Packer, C. & Clune, J. 2018 Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning. *Proc. Natl. Acad. Sci.* **115** (25), E5716–E5725.

van den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A. & Kavukcuoglu, K. 2016 Wavenet: A generative model for raw audio. arXiv preprint arXiv:1609.03499.

Pan, S. J. & Yang, Q. 2009 A survey on transfer learning. *IEEE Transactions on knowledge and data engineering* **22** (10), 1345–1359.

Pandey, S., Schumacher, J. & Sreenivasan, K. R. 2020 A perspective on machine learning in turbulent flows. *J. Turbul.* pp. 1–18.

Rabault, J., Kuchta, M., Jensen, A., Réglade, U. & Cerardi, N. 2019 Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control. *J. Fluid Mech.* **865**, 281–302.

Raibaudo, C., Zhong, P., Noack, B. R. & Martinuzzi, R. J. 2020 Machine learning strategies applied to the control of a fluidic pinball. *Phys. Fluids* **32** (1), 015108.

Raissi, M., Perdikaris, P. & Karniadakis, G. E. 2019 Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707.

Raissi, M., Yazdani, A. & Karniadakis, G. E. 2020 Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations. *Science* **367** (6481), 1026–1030.

Sasaki, K., Vinuesa, R., Cavalieri, A. V., Schlatter, P. & Henningson, D. S. 2019 Transfer functions for flow predictions in wall-bounded turbulence. *J. Fluid Mech.* **864**, 708–745.

Sirovich, L. 1987 Turbulence and the dynamics of coherent structures. I. Coherent structures. *Q. Appl. Math.* **45** (3), 561–571.

Srinivasan, P., Guastoni, L., Azizpour, H., Schlatter, P. & Vinuesa, R. 2019 Predictions of turbulent shear flows using deep neural networks. *Phys. Rev. Fluids* **4** (5), 054603.

Suzuki, T. & Hasegawa, Y. 2017 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.

Tang, H., Rabault, J., Kuhnle, A., Wang, Y. & Wang, T. 2020 Robust active flow control over a range of reynolds numbers using an artificial neural network trained through deep reinforcement learning. *Phys. Fluids* **32** (5), 053605.

Tinney, C., Coiffet, F., Delville, J., Hall, A., Jordan, P. & Glauser, M. N. 2006 On spectral linear stochastic estimation. *Exp. Fluids* **41** (5), 763–775.

Tinney, C. E., Ukeiley, L. & Glauser, M. N. 2008 Low-dimensional characteristics of a transonic jet. Part 2. Estimate and far-field prediction. *J. Fluid Mech.* **615**, 53–92.

Udrescu, S.-M. & Tegmark, M. 2020 AI Feynman: A physics-inspired method for symbolic regression. *Science Advances* **6** (16), eaay2631.

Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., Felländer, A., Langhans, S. D., Tegmark, M. & Nerini, F. F. 2020 The role of artificial intelligence in achieving the sustainable development goals. *Nat. Commun.* **11** (1), 1–10.

Wang, J. X., Wu, J. L. & Xiao, H. 2017 Physics-informed machine learning approach for reconstructing Reynolds stress modeling discrepancies based on DNS data. *Phys. Rev. Fluids* **2** (3), 034603.

Wu, J.-L., Wang, J.-X., Xiao, H. & Ling, J. 2017 A priori assessment of prediction confidence for data-driven turbulence modeling. *Flow Turbul. Combust.* **99** (1), 25–46.

Wu, Y. 2014 A study of energetic large-scale structures in turbulent boundary layer. *Phys. Fluids* **26** (4), 045113.

Paper 3

3

Güemes, A., Discetti, S., Ianiro, A., Sirmacek, B., Azizpour, H. & Vinuesa, R., 2021. From sparse wall measurements to turbulent velocity elds through deep learning. *Phys. Fluids*. **Under review**. https://arxiv.org/abs/2103.07387v2