

AI-BASED SENSING OF TURBULENT WALL-BOUNDED FLOWS

by

ANTONIO CUÉLLAR MARTÍN

A dissertation submitted in partial fulfillment of the requirements for the
degree of Doctor of Philosophy in

Aerospace Engineering

Universidad Carlos III de Madrid

Advisor(s):

ANDREA IANIRO
STEFANO DISCETTI

Tutor:

STEFANO DISCETTI

December, 2024

This thesis is distributed under license “Creative Commons **Attribution - Non Commercial - Non Derivatives**”.



*A mi madre,
por creer siempre en mí
y caminar a mi lado en cada paso*

Preface

Agradecimientos

Quiero comenzar estas palabras expresando mi más profundo agradecimiento a mis tutores de tesis, Stefano y Andrea, quienes desde el principio me apoyasteis y me ofrecisteis la oportunidad de embarcarme en esta emocionante aventura científica. Fue durante mi trabajo de fin de máster cuando me propusisteis comenzar con este trabajo, y desde entonces no habéis dejado de guiarme con vuestra experiencia y conocimiento. Me habéis transmitido no solo vuestra pasión por la aerodinámica experimental, sino también la importancia del rigor y la precisión en la ciencia, valores que siempre llevaré conmigo en mi carrera. Siempre habéis sido pacientes conmigo, encontrando tiempo para apoyarme en cada etapa de mi investigación más allá del tiempo material del que disponemos, y especialmente en estos últimos meses. Me habéis ofrecido grandes consejos cuando más lo he necesitado y manteniendo siempre un compromiso firme con mi formación. Además, agradezco enormemente el haberme involucrado en diferentes proyectos y actividades que han enriquecido este trabajo y me han permitido crecer como investigador. Sin vuestro apoyo y constante dedicación, este trabajo no habría sido posible.

A todo el grupo de investigación quiero agradecer vuestra presencia y apoyo constante. Directa o indirectamente siempre aprendemos los unos de los otros, tanto dando clase como en el laboratorio, y me llevo algo muy valioso de cada uno de vosotros que considero fundamental en este proceso y lo habéis hecho más enriquecedor y llevadero. Especialmente quiero agradecer, Marco, por haberme ayudado siempre que te he pedido ayuda y haberme dedicado mucho más tiempo del que estaba esperando. Y a Qihong, por estar siempre para escuchar y animarme.

He tenido la oportunidad de trabajar con grandes personas como Ricardo, de KTH, y Óscar, he aprendido mucho de vosotros y os agradezco gratamente vuestra predisposición constante a colaborar con nosotros y vuestra contribución a esta tesis. Muy especialmente, quiero agradecer, Álex, la gran ayuda que diste en los primeros momentos de esta tesis, que después me sirvieron para desarrollar todo el trabajo de los siguientes años. Espero que todo esto solo sea el comienzo de futuras colaboraciones en las que volveremos a trabajar juntos. En estos años he tenido la oportunidad asistir a diversos congresos científicos donde he podido conocer a gente de quien he aprendido, que me ha dado ideas, ha valorado mi trabajo, y de una u otra forma también han repercutido en mejoras de mi trabajo científico, dándole un valor añadido que también es de agradecer. And a great example of this is Woutijn from TU Delft, with whom I have had the pleasure of meeting on several occasions and whom I thank for his willingness to collaborate in the work on the last part of this and future theses. E non posso dimenticare

di ringraziarvi per il supporto, il lavoro, la gentilezza e la predisposizione che tutto il team Polito mi ha dimostrato in ogni momento. Enrico, Sara, Jacopo, Giocchino, grazie di cuore per tutto questo, per avermi aiutato a rendere possibile una parte fondamentale di questa tesi e per avermi fornito un grande contributo. Porto con me anche una grande amicizia e stima nei vostri confronti, cosa a cui tengo molto. Ho solo belle parole per tutti voi.

A todos los amigos con quienes he podido desconectar del trabajo en estos años, mil gracias porque con vosotros se ha hecho mucho más llevadero. De estos años me llevo a grandes personas como Gala, Eric y Paula. Y especialmente a Vanesa, te agradezco estar siempre ahí después de tanto recorrido juntos, y que a pesar de la distancia y el tiempo siempre encontramos algún momento para vernos, escucharnos y ponernos al día.

Y por último a mi familia, quiero agradecer todo el apoyo a lo largo de este camino. A pesar de las largas horas de estudio y las ausencias, siempre habéis creído en mí y en este proyecto, y vuestra confianza ha sido mi mayor motivación. A María, que siempre me has acompañado. A mis tíos Antonio y Jesús que me apoyáis cada día y me habéis dado todas las facilidades del mundo para hacer esta tesis. A Edu por aguantarme y acompañarme en todo este camino como nadie. Y especialmente a mi madre, por superarte cada día, por haberte esforzado tanto para que tuviésemos la mejor formación, una carrera próspera, por acompañarme siempre, por quererme como nadie, creer en mi y en cada proyecto que inicio, por los valores que nos has inculcado, por habernos llevado a estudiar un poco más lejos de casa a pesar del déficit de horas de sueño, por haberme lanzado a aprender idiomas, a conocer el mundo, a conocer gente nueva, a esforzarme por conseguir aquello que siempre he querido y por haberme convencido de que hiciese ingeniería aeroespacial. Este logro también es gracias a todos vosotros.

Funding

This thesis has been funded by the Ministry of Universities under the programme “Formación de Profesorado Universitario (FPU) 2020”, with the reference number FPU20/02551. The research has received funding from the European Research Council (ERC) under the European Union’s Horizon H2020 research and innovation programme (grant agreement No 949085, “NEXTFLOW”), by MCIN/AEI/10.13039/501100011033 and by the “European Union NextGenerationEU/PRTR” within the project “ACCREDITATION” (Grant No TED2021-131453B-I00), and by MCIU/AEI/10.13039/501100011033 and by “ERDF A way of making Europe” within the project “EXCALIBUR” (Grant No PID2022-138314NB-I00).

Resumen

Detección de flujos turbulentos confinados por paredes basada en IA

Antonio Cuéllar Martín, Departamento de Ingeniería Aeroespacial, Universidad Carlos III de Madrid

Esta tesis investiga la estimación del campo de velocidad en un canal mediante sensores no intrusivos en la pared. Esto es crucial para el desarrollo de estrategias de control activo en flujos turbulentos limitados por paredes. La estimación del flujo a partir de magnitudes medidas en la pared es un reto pendiente desde hace tiempo. El principal objetivo de esta investigación es desarrollar modelos basados en datos para predecir el comportamiento tridimensional del flujo de fluidos.

Se emplearon datos de simulaciones numéricas directas para entrenar modelos de aprendizaje profundo. Se demostró que las redes neuronales convolucionales tridimensionales con entrenamiento adversativo predicen con precisión los campos de flujo desde mediciones de la pared, con una reducción significativa del coste computacional en base a los estimadores planares. En particular, la red propuesta es capaz de estimar estructuras coherentes adheridas a la pared debido a que su huella se detecta en la pared. La arquitectura de red neuronal propuesta demostró un excelente rendimiento incluso en presencia de ruido. Además, se ha explorado el efecto de reducir la cantidad de información disponible en la pared. Las mediciones de presión proporcionarían mejores reconstrucciones del flujo con suficientes sensores como para muestrear adecuadamente las escalas de flujo, mientras que la tensión de cizalladura de la pared en sentido de la corriente debería preferirse para la medición de las fluctuaciones de velocidad en sentido de la corriente y si el número de sensores es limitado.

Para allanar el camino hacia una demostración experimental, se ha diseñado una campaña experimental para la adquisición de medidas sincronizadas de velocidades en la región próxima a la pared y de transferencia de calor en la pared. El principal reto, la medición de los campos instantáneos de la pared tiempo resueltos, se ha abordado con un sensor de lámina delgada

calentada y termografía IR. La relación señal-ruido de este tipo de mediciones en el aire suele ser baja. Se propone una receta de filtrado para extraer datos fiables en este escenario.

Esta investigación proporciona un nuevo marco para utilizar métodos basados en datos en flujos confinados por paredes, con aplicaciones potenciales en el control activo del flujo para sistemas de ingeniería. La integración del aprendizaje automático abre nuevas vías hacia la predicción y el control del flujo en tiempo real.

Abstract

AI-based sensing of turbulent wall-bounded flows

Antonio Cuéllar Martín, Department of Aerospace Engineering, Universidad Carlos III de Madrid

This thesis investigates the estimation of the velocity field in a channel flow from non-intrusive wall-embedded sensors. This is crucial for the development of active control strategies in wall-bounded turbulent flows. Flow estimation based on wall-measured quantities is a long-standing challenge. The main objective of this research is to develop data-driven models to predict the three-dimensional fluid flow behaviour.

Datasets from direct numerical simulations were employed to train deep-learning models. Three-dimensional convolutional neural networks with adversarial training were shown to accurately predict flow fields from wall measurements, with a significant reduction in computational cost with respect to planar estimators. In particular, the proposed network is capable of estimating wall-attached coherent structures due to their footprint being sensed at the wall. The proposed neural network architecture demonstrated excellent performance even in the presence of noise. Furthermore, the effect of reducing the amount of information available at the wall has been explored. Pressure measurements would provide better flow reconstructions if the number of sensors is large enough to sample flow scales properly, while streamwise wall shear stress should be preferred if the target is the measurement of streamwise velocity fluctuations and if the number of sensors is limited.

To pave the way towards an experimental demonstration, an experimental campaign for the acquisition of synchronised measurements of velocities in the near-wall region and heat transfer at the wall has been designed. The main challenge, measuring time-resolved instantaneous wall fields, has been addressed with a heated-thin-foil sensor coupled with IR thermography. The signal-to-noise ratio of these types of measurements in air is generally low. A recipe to extract reliable data in this scenario is proposed.

This research provides a new framework for using data-driven methods in wall-bounded flows, with potential applications in active flow control for engineering systems. The integration of machine learning opens new avenues towards real-time flow prediction and control.

Published and submitted content

This PhD thesis presents a methodology to conduct flow field estimations in wall-bounded flows from wall measurements employing deep neural networks. Simulated datasets have been employed to develop the methodology and to assess the impact of several effects related to its implementation in experiments. Furthermore, an experimental campaign to obtain a dataset is described. The first part of this thesis contains an introduction, description of the problem, motivation, theoretical background, methodological description of the techniques employed and main conclusions of this thesis. The second part contains the scientific articles that develop this thesis in detail. These articles included here have been modified to have the same format as this thesis, but their contents are the same of the published version. Whenever material from these sources is included in this thesis, it is singled out with typographic means and references.

Paper 1. A. CUÉLLAR, A. GÜEMES, A. IANIRO, Ó. FLORES, R. VINUESA & S. DISCETTI 2024 *Three-dimensional generative adversarial networks for turbulent flow estimation from wall measurements* *Journal of Fluid Mechanics*, **991** A1, [doi:10.1017/jfm.2024.432](https://doi.org/10.1017/jfm.2024.432)

Paper 2. A. CUÉLLAR, A. IANIRO & S. DISCETTI 2024 *Some effects of limited wall-sensor availability on flow estimation with 3D-GANs* *Theoretical and Computational Fluid Dynamics*, [doi:10.1007/s00162-024-00718-w](https://doi.org/10.1007/s00162-024-00718-w)

Paper 3. A. CUÉLLAR, E. AMICO, J. SERPIERI, G. CAFIERO, W. BAARS, S. DISCETTI & A. IANIRO 2024 *Measuring time-resolved heat transfer fluctuations on a heated-thin foil in a turbulent channel airflow* Under review, [Preprint arXiv: 2410.12778](https://arxiv.org/abs/2410.12778) [[physics.flu-dyn](https://arxiv.org/abs/2410.12778)]

Author contributions

The categories of the different author contributions are described in Brand, A., et al. Beyond authorship: Attribution, contribution, collaboration, and credit. Learned Publishing, 2015, vol. 28, no 2.

Paper 1 All authors contributed to the study’s conception and design. Conceptualisation: SD; Methodology: AC, AG, AI, OF, RV, SD; Software: AC, AG; Formal analysis: AC, OF; Investigation: AC, AG, AI, OF, RV, SD Resources: AI, SD; Data curation: AC; Writing - original draft preparation: AC, AI, OF, SD; Writing - review and editing: AC, AG, AI, OF, RV, SD; Visualization: AC; Supervision: AI, SD; Project Administration: AC, AI, SD; Funding acquisition: AC, SD. All authors read and approved the final manuscript. This article is wholly included in this thesis, reproduced as paper 1 in part [III](#).

Paper 2 All authors contributed to the study’s conception and design. Conceptualisation: AC, AI, SD; Methodology: AC, AI, SD; Software: AC; Formal analysis: AC; Investigation: AC; Resources: AI, SD; Data curation: AC; Writing - original draft preparation: AC; Writing - review and editing: AC, AI, SD; Visualization: AC; Supervision: AI, SD; Project Administration: AC, AI, SD; Funding acquisition: AC, SD. All authors read and approved the final manuscript. This article is wholly included in this thesis, reproduced as paper 2 in part [III](#).

Paper 3 All authors contributed to the study’s conception and design. Conceptualisation: AI, SD; Methodology: AC, EA, JS, GC, WB, AI, SD; Software: AC, JS, GC, AI; Formal analysis: AC; Investigation: AC, EA, JS, GC; Resources: JS, GC, AI, SD, WB; Data curation: AC; Writing - original draft preparation: AC; Writing - review and editing: AC, EA, GC, JS, WB, AI, SD; Visualization: AC; Supervision: AI, SD; Project Administration: AC, AI, SD; Funding acquisition: AC, AI, SD. All authors read and approved the final manuscript. The preprint of this article, currently under review, is wholly included in this thesis, reproduced as paper 3 in part [III](#).

Acronyms correspond to:

AC	Antonio Cuéllar, Universidad Carlos III de Madrid, PhD Candidate
SD	Stefano Discetti, Universidad Carlos III de Madrid, Advisor
AI	Andrea Ianiro, Universidad Carlos III de Madrid, Advisor
AG	Alejandro Güemes, Universidad Carlos III de Madrid
OF	Óscar Flores, Universidad Carlos III de Madrid
RV	Ricardo Vinuesa, Kungliga Tekniska Högskolan (Sweden)
EA	Enrico Amico, Politecnico di Torino (Italy)
GC	Gioacchino Cafiero, Politecnico di Torino (Italy)
JS	Jacopo Serpieri, Politecnico di Torino (Italy)
WB	Woutijn J. Baars, Technische Universiteit Delft (The Netherlands)

Abstracts

Paper 1

Different types of neural networks have been used to solve the flow sensing problem in turbulent flows, namely to estimate velocity in wall-parallel planes from wall measurements. Generative adversarial networks (GANs) are among the most promising methodologies, due to their more accurate estimations and better perceptual quality. This work tackles this flow sensing problem in the vicinity of the wall, addressing for the first time the reconstruction of the entire three-dimensional (3-D) field with a single network, i.e. a 3-D GAN. With this methodology, a single training and prediction process overcomes the limitation presented by the former approaches based on the independent estimation of wall-parallel planes. The network is capable of estimating the 3-D flow field with a level of error at each wall-normal distance comparable to that reported from wall-parallel plane estimations and at a lower training cost in terms of computational resources. The direct full 3-D reconstruction also unveils a direct interpretation in terms of coherent structures. It is shown that the accuracy of the network depends directly on the wall footprint of each individual turbulent structure. It is observed that wall-attached structures are predicted more accurately than wall-detached ones, especially at larger distances from the wall. Among wall-attached structures, smaller sweeps are reconstructed better than small ejections, while large ejections are reconstructed better than large sweeps as a consequence of their more intense footprint.

Paper 2

In this work we assess the impact of the limited availability of wall-embedded sensors on the full 3D estimation of the flow field in a turbulent channel with $Re_\tau = 200$. The estimation technique is based on a 3D generative adversarial network (3D-GAN). We recently demonstrated that 3D-GANs are capable of estimating fields with good accuracy by employing fully-resolved wall quantities (pressure and streamwise/spanwise wall shear stress on a grid with DNS resolution). However, the practical implementation in an experimental setting is challenging due to the large number of sensors required. In this work, we aim to estimate the flow fields with substantially fewer sensors. The impact of the reduction of the number of sensors on the quality of the flow reconstruction is assessed in terms of accuracy degradation and spectral length-scales involved. It is found that the accuracy degradation is mainly due to the spatial undersampling of scales, rather than the reduction of the number of sensors per se. We explore the performance of the estimator in case only one wall quantity is available. When a large number of sensors is

available, pressure measurements provide more accurate flow field estimations. Conversely, the elongated patterns of the streamwise wall shear stress make this quantity the most suitable when only few sensors are available. As a further step towards a real application, the effect of sensor noise is also quantified. It is shown that configurations with fewer sensors are less sensitive to measurement noise.

Paper 3

We present an experimental setup to perform time-resolved convective heat transfer measurements in a turbulent channel flow with air as the working fluid. We employ a heated thin foil coupled with high-speed infrared thermography. The measurement technique is challenged by the thermal inertia of the foil, the high frequency of turbulent fluctuations, and the measurement noise of the infrared camera. We discuss in detail the advantages and drawbacks of all the design choices that were made, thereby providing a successful implementation strategy to obtain high-quality data. This experimental approach could be useful for experimental studies employing wall-based measurements of turbulence, such as flow control applications in wall-bounded turbulence.

Other research merits

Conference - Oral contributions

A. CUÉLLAR, A. GÜEMES, A. IANIRO, Ó. FLORES, R. VINUESA & S. DISCETTI *From wall measurements to three-dimensional turbulent-flow fields via GANs* 1st Spanish Fluid Mechanics Conference 2022 (SFMC22), Cádiz, Spain, 2022

A. CUÉLLAR, A. GÜEMES, A. IANIRO, Ó. FLORES, R. VINUESA & S. DISCETTI *Generative Adversarial Networks for the Estimation of 3D Turbulent Fields with Wall-Measurements* 1st International Conference on Mathematical Modelling in Mechanics and Engineering (ICME 2022), Belgrade, Serbia, 2022

A. CUÉLLAR, L. FRANCESCHELLI, C. MÁRQUEZ, S. DISCETTI & A. IANIRO *Experimental investigation of turbulent swirling jets* 2nd Spanish Fluid Mechanics Conference 2023 (SFMC23), Barcelona, Spain, 2023

A. CUÉLLAR, A. GÜEMES, A. IANIRO, Ó. FLORES, R. VINUESA & S. DISCETTI *From wall measurements to three-dimensional turbulent-flow fields* 18th European Turbulence Conference (ETC18), Valencia, Spain, 2023

A. CUÉLLAR, A. IANIRO & S. DISCETTI *Effect of the wall-sensor availability on the quality of the flow field estimation in turbulent wall-bounded flows* Euromech Colloquium 631 ‘Control of skin friction and convective heat transfer in wall-bounded flows’, Madrid, Spain, 2024

A. CUÉLLAR, A. IANIRO, & S. DISCETTI *Effect of the wall-sensor availability on the flow field estimation with 3D-GANs in turbulent wall-bounded flows* European Drag Reduction and Flow Control Meeting (EDRFCM 2024), Turin, Italy, 2024

Conference - Poster contribution

A. CUÉLLAR, A. GÜEMES, A. IANIRO, Ó. FLORES, R. VINUESA & S. DISCETTI *From wall measurements to three-dimensional turbulent-flow fields* iTi X 2023 Conference on Turbulence, Bertinoro, Italy, 2023

Research Stay

3-month research stay at Politecnico di Torino (May - July 2024 and September 2024) under the invitation of Professor Gioacchino Cafiero.

During this research stay abroad, an experimental campaign was carried out at the host research centre. Particle image velocimetry and infrared thermography have been used to perform synchronized measurements of velocity and wall convective heat transfer in a channel flow. One of the main objectives of this campaign is to obtain an experimental database to test the flow-sensing techniques developed in this thesis.

Dissemination activities

Semana de la Ciencia y la Innovación de Madrid 2022, *¿Cómo funciona un turborreactor?*

Semana de la Ciencia y la Innovación de Madrid 2023, *Aprende a innovar en el modelado aeroespacial*

Noche Europea de los Investigadores y de las Investigadoras de Madrid 2024, *Experimenta con cámaras térmicas utilizadas en investigación aeronáutica: ¡ven a una escuela de calor!*

Semana de la Ciencia y la Innovación de Madrid 2024, *Aprende con una cámara térmica*

Contents

Preface	vi
Agradecimientos	vi
Resumen	ix
Abstract	xi
Published and submitted content	xiii
Other research merits	xvii
Conference - Oral contributions	xvii
Conference - Poster contribution	xviii
Research Stay	xviii
Dissemination activities	xviii
Contents	xix
I Overview and summary	1
1 Introduction	3
1.1 Brief history of turbulence in science	4
1.2 The turbulence control challenge	7
2 Wall-bounded turbulence	11
2.1 Wall-bounded flows	12
2.2 Turbulent coherent structures	16
2.3 Strategies to discover patterns in wall-bounded flows	19

2.4 Estimation in wall-bounded flows	20
3 Deep Learning for fluid flows	23
3.1 Overview on ML approaches	24
3.2 Machine learning and its impact on fluid dynamics	25
3.3 Introduction to Deep Neural Networks	26
3.4 Generative Adversarial Networks	33
3.5 3D-GANs for flow estimation	35
4 A technique to measure instantaneous convective heat transfer in wall-bounded flows	39
4.1 The link between wall-shear stress and convective heat transfer coefficient	40
4.2 Experimental background on IR thermography for wall-bounded flows	41
4.3 Sensor for IR: implementation and thermal model	42
5 Main contributions and conclusions	47
5.1 Main highlights	48
5.2 Future work	51
5.2.1 Development of more efficient estimation techniques	51
5.2.2 Addressing effects towards experimental flow sensing implementation	52
5.2.3 Experimental validation with synchronized PIV and IR measurements	53
5.2.4 Going beyond the training dataset	56
II Papers	59
Paper 1	61
Paper 2	99
Paper 3	125
Bibliography	141

Part I

Overview and summary

Chapter 1

Introduction

Turbulence is one of the most important and studied phenomena in fluid dynamics. It is present in numerous natural mechanisms and configurations such as atmospheric and oceanic currents or within our body's circulatory system [1]. Turbulence is observed in daily situations, such as in a car exhaust or in the rapids of a river. Some examples are illustrated in figure 1.1. It affects bodies in motion in fluids, such as automobiles, aeroplanes, ships and wind farms. Additionally, turbulence impacts natural flyers such as birds, influencing their flight patterns and energy expenditure as they navigate through turbulent air currents. The complexity of this phenomenon makes it difficult to predict its behaviour, notwithstanding the efforts motivated by its ubiquitous presence in our daily lives.



Figure 1.1: (a) Smoke at a volcano outlet, (b) Urban flow, reproduced from [2], (c) Ship wake, © Lancaster University, (d) off-shore wind farm

This chapter starts with a brief historical context about the study of turbulence, referring to some important milestones, and introducing today's main challenges. In particular, the current interest behind flow control is discussed, which pushes the development of smart flow sensing systems, hence setting the motivation of this thesis.

1.1 Brief history of turbulence in science

Leonardo da Vinci (1452–1519) was a pioneer in the observation of turbulent flows. Indeed, he coined the word *turbulence* more than 500 years ago in the ‘Atlantic Codex’, where we can read:

*doue la turbolenza dell’acqua si genera
doue la turbolenza dell’acqua si mantiene plugho
doue la turbolenza dell’acqua si posa.*

He systematically analysed numerous cases of turbulent flows. His contributions are even more remarkable if we account that the theoretical tools to conduct a formal analysis were not yet developed. In particular, the works by Navier and Stokes building the mathematical background of fluid dynamics came about 300 years after da Vinci’s scientific contribution. Nevertheless, his studies led to great technological advances at the time. He was able to develop experiments and get explanations following logical reasoning from his observations, approaching the procedure of the scientific method established after him by Galileo Galilei [3]. His drawings sketch different setups, even with velocity streamlines and vortices, such as those seen in figure 1.2. To mention some examples, he studied the water jet from a channel falling on a tank [4], the courses of rivers used for transportation or as energy sources [3], wave dynamics or the cardiovascular vortexes in the aorta artery [5].



Figure 1.2: A Deluge, by Leonardo da Vinci (No. 192), Windsor Castle, Print Room (Royal Collection Trust © Her Majesty Queen Elizabeth II, 2020), reproduced with permission

Da Vinci’s contribution provided notions of what years later would become the fundamental underpinnings in this field of knowledge. He was aware of the 3D multi-scale nature eddies as an

essential part of turbulent flows. This notion would materialise centuries later in Kolmogorov's 'cascade model' in 1941 [6]. He also described the water velocity profile in a pipe, noticing its dependence on the distance to the wall, surface roughness and wall friction. His drawings also report currents with reduced speed near a fixed solid body. Nowadays, we can talk about the 'no-slip wall boundary condition' and 'boundary layers' [5]. It was relatively recently when Froude's experiments determined that the flow speed increases with the distance from the wall in 1874 [7] and when Prandtl first introduced the 'boundary layer' concept and mathematical formulation in 1904 [8].

Different definitions of turbulence can be found in the literature from varied points of view and covering different aspects. Those are some of them:

- *'Turbulence is the name given to the imperfectly-understood class of chaotic solutions to the Navier–Stokes equations in which many degrees of freedom are excited'* [9].
- *'The distinguishing feature of turbulent flow is that its velocity field appears to be random and varies unpredictably. The flow, however, does satisfy a set of differential equations, the Navier–Stokes equations, which are not random. This contrast is the source of much of what is interesting in turbulence theory'* [10].
- *'Creation of small-scale activity and dissipation is the principle of turbulence. Classical fluid-dynamical instabilities play a role of the fuel, vortex stretching is the engine, and viscous dissipation is the breaks'* [11].
- *'We have therefore defined turbulence as random fluctuations of the thermodynamic characteristics of vortex flows, thereby distinguishing it at the outset from any kind of whatever random irrotational, i.e., potential flows, ...'* [12].

The advancements mentioned above, conducted just over a hundred years ago, allowed for the subsequent turbulence research we continue developing today. Various physicists such as Sommerfeld, Einstein or Feynman have stated that *'turbulence is the last great unsolved problem of classical physics'* [1], and it remains one of the most challenging research areas of fluid dynamics to the scientific community. Many fundamental questions are still unresolved. This list with some of the most important pending challenges might provide a view of the complexity of turbulent phenomena and the special interest it arouses:

- Complete statistical formulation for turbulence [13], [14]
- Link between the Navier-Stokes equations with chaos and randomness in turbulence [15]

- Wall-bounded turbulence and complex geometries [16]–[18]
- Active and passive control for efficiency improvement [19], [20]
- Small-scale modelling impact on the larger scales for large eddy simulations (LES) [21], [22]
- Reynolds stress modelling to close the Navier-Stokes equations [23], [24]
- Effects behind the laminar-to-turbulent flow transition [25], [26]
- The mechanisms behind turbulent flows [26], [27]
- Distribution of energy in a multi-scale flow [28]
- Transport and mixing in turbulent flows [29]–[32]
- Transport and interaction in turbulent multi-phase flows [33], [34]

The significant advancements reported in the last decades have been possible, among other reasons, thanks to the technological leap forward of these times, both in terms of experiments and simulations. For instance, today’s digital cameras, lasers, and computers allow us to obtain much more complete flow characterisations with particle image velocimetry (PIV) [35]–[37] than what was possible with only hot-wire anemometers [38], [39]. Besides, state-of-the-art facilities such as (among many others) the long pipe ‘CICLoPE’ in Predappio (Italy) [40] or the largest setup for highly turbulent Rayleigh-Bénard convection ‘Barrel of Ilmenau’ in Ilmenau (Germany) [41], open up the possibility of performing experiments with the potential to validate theories and discover new physical mechanisms in controlled conditions.

In parallel, progress in computing capabilities, with an exponential growth trend in the time of floating point operations per second at which they can be solved, has enabled simulations that were not possible some time ago [42]–[44], bringing high-fidelity simulations progressively closer to the real applications. Recently, this growth has been largely driven by the development of GPUs, which have revolutionised computational performance through their ability to handle massively parallel processing efficiently. This raised the question of whether simulations could replace experiments. However, the computational cost of a direct numerical simulation (DNS) of the Navier-Stokes equations increases with the third power of the Reynolds number [45] for one single time step. Cheaper alternatives, such as Large-Eddy Simulations (LES) and Reynolds-averaged Navier-Stokes equations (RANS) for wall-bounded turbulence, may need the previous development of a physical theory for the structure of the boundary layer [46].

1.2 The turbulence control challenge

One of the main current societal challenges is to shift towards sustainable development models. This implies remarkable changes in our industrial model, consumption habits and other socio-cultural aspects. Concerning the 2030 Agenda for Sustainable Development [47], more measures must be implemented to achieve emission reduction goals, including the net zero emission goal for 2050. Recently, the Centre for Research on Energy and Clean Air (CREA) pointed out an 8% drop in 2023 in CO₂ emissions in Europe [48], the second most significant reduction ever registered after the year 2020 —with a clear influence from the COVID-19 pandemic. The current level of CO₂ emissions is the lowest since the 1960s. Most of this decrease is due to the use of alternative and clean ways of producing electricity [48], but demand reduction has also contributed significantly. This can be intensified by seeking higher levels of efficiency in each device that requires a certain amount of energy consumption. Some applications can make use of both green energy production and highly efficient designs to reduce the emissions of CO₂ and other pollutants. Nevertheless, the restricted access and availability of green energy, and more importantly green power, in environments such as aviation, makes this issue more challenging. As long as access to high levels of power obtained from sustainable sources is not possible, the development of highly efficient devices becomes even more relevant in the aviation field.

Turbulent flows are present in various industrial applications [1] and have a great influence on their performances, efficiency and energetic consumption. In particular, wall-bounded flows are very relevant as they are directly in contact with the devices and turbulent transport of momentum is related to skin friction, drag, and pressure losses. A good knowledge and understanding of these flows might open up new possibilities for turbulence control that could mitigate or exacerbate the effects of turbulence when they are undesired or wanted respectively, making new devices more efficient. Some prominent examples include noise reduction, mixing maximisation, lift enhancement or skin friction reduction [49].

The viscous drag from the wall-bounded turbulence around vehicles and pipelines has been estimated to account for about 50% of the total human-made CO₂ emissions [50]. The control of turbulence is thus key to the reduction of CO₂ emissions. Among all possible applications, the aerospace industry has a clear dependence on this fluid dynamics problem, influenced by effects such as the transfer of energy and momentum through the boundary layer, the development of coherent structures, their propagation in the vicinity of a wall or the friction on the skin of an immersed surface. Emissions in aviation are responsible for 12% of CO₂ emissions of all

transportation means, according to [Air Transport Action Group \(ATAG\)](#). As seen in figure [1.3](#), in a typical aircraft at cruise conditions, skin friction can represent nearly 50% of its total drag amount [\[51\]](#).

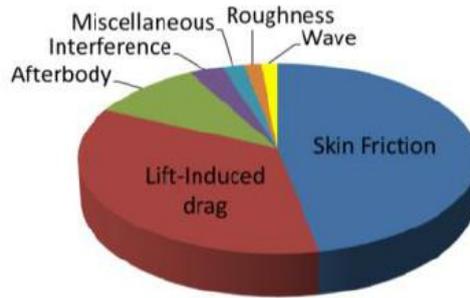


Figure 1.3: Transport aircraft drag breakdown. Figure 1 from [\[51\]](#), reproduced with permission

Turbulent flow control is generally classified as passive or active. Passive control strategies rely on fixed and pre-designed features of the system that do not require external energy or dynamic adjustments, while active control ones employ some external forcing to manipulate the system behaviour. These techniques have demonstrated great potential in altering wall-bounded flow behaviour to reduce skin friction. Passive control of skin friction has been shown with surface modifications such as large eddy break-up devices (LEBUs) and riblets [\[52\]](#), [\[53\]](#). Alternatively, active control could allow us to promote or suppress certain flow features. The actuation can be typically tuned depending on the operating conditions [\[54\]](#), [\[55\]](#). For example, whenever it is detected that a turbulent event with a negative impact on the efficiency of the system is present or is being developed, an actuator could prevent it or reduce its effect. Active control devices, however, need a power source and should be carefully designed to ensure that their operation results in a net energy saving [\[49\]](#). Considering a turbulent flow as a superposition of different flow features, each with a certain role in the flow dynamics, we might think of controlling specific flow features through small-amplitude targeted forcing systems, which consume less power and might be smaller and lighter [\[56\]](#). For example, Ref. [\[57\]](#) demonstrated a drag reduction of 13% at $Re_\tau = 12800$ with an actuation targeting the large-scale eddies of a turbulent boundary layer; this reduction should increase with increasing Re_τ with a significantly lower power actuation cost.

The implementation of active control systems for these purposes may require continuous monitoring of the flow and its state [\[58\]](#). Intrusive sensing techniques, such as hot-wire anemometry, alter the natural state of turbulent structures near the wall and might bring unwanted

effects. This motivates the development of non-intrusive sensing techniques that do not disturb the flow. For active control systems for wall-bounded turbulence, wall-embedded sensors seem a reasonable approach, based on the evidenced correlations in wall-bounded flows. A detailed section covering the literature on flow estimation methodologies from wall measurements is included in chapter [2](#).

In this context, **the fundamental objective of this thesis is the development of a flow-sensing algorithm capable of making 3D velocity estimations from instantaneous mappings of wall measurements.** A review of state-of-the-art estimation approaches may allow us to identify some of the most promising ones. An extension of the methodology to 3D estimation would overcome certain limitations from 2D estimators.

The recent disruption of artificial intelligence (AI) provides new opportunities for flow control. Modern AI algorithms are capable of learning by extracting patterns and information from data, even when modelling based on first principles is particularly challenging. In the context of turbulence, AI algorithms bring forth the advantage of being capable of dealing with data representative of physical scenarios even in the presence of strong non-linear behaviours. This is a particularly intriguing skill in view of tackling different problems in the field of fluid dynamics. In this thesis, AI is exploited to develop advanced flow-sensing strategies for flow control. Chapter [3](#) contains a brief introduction to AI and more detailed explanations of the AI concepts employed in this thesis.

Although the data-driven approach presented in this thesis has been developed employing samples from a DNS, an experimental database would provide valuable validation, bringing the method closer to real-world applications and validating its robustness under more practical conditions, and potentially reveal additional insights. In spite of the additional challenges, an experimental campaign has been proposed to take synchronised measurements of the flow and of the wall fields. In chapter [4](#), the experimental techniques selected to perform both measurements are introduced and the experimental campaign is described. This procedure may be useful to get experimental datasets analogous to the computational datasets employed for the training of the machine learning (ML) models.

Chapter 2

Wall-bounded turbulence

This chapter covers some fundamental aspects of wall turbulence. A particular focus is on turbulent channel flows, which is the main test case used in this thesis.

The chapter starts with §2.1 covering an introduction to fundamental aspects of wall-bounded flows and their presence in common scenarios. This section also reports important fundamental research milestones of the 20th century that later brought studies establishing correlations within wall-bounded flows. Advancements in turbulence studies resulted in different models to explain turbulence, such as coherent structures. A review of these models is provided in §2.2. Then, different methodologies and mathematical models employed to study wall-bounded flows are covered in §2.3 and 2.4. This includes techniques to identify flow structures, and the approaches employed for flow estimation purposes, which started with linear models and evolved until recent literature, which makes extensive use of neural networks.

2.1 Wall-bounded flows

In a myriad of engineering and industrial applications fluid flows interact with solid boundaries. The geometry of the body and/or surface roughness significantly influence the flow characteristics and lead to complex interactions between the surface and the flow. As a direct consequence of the continuous behaviour of fluid flows, a velocity gradient is observed near the wall to comply with the no-slip condition. This leads to the formation of a boundary layer, a part of the flow field in which viscous forces are comparable to inertia forces. Within the boundary layer, the velocity gradient at the wall is responsible for the skin friction. One of the main properties of turbulent flows is their transport capability, which is responsible for high mixing and high mass, momentum, and energy transfer rates. In particular, the enhanced transport of momentum within a turbulent boundary layer leads to large velocity gradients and increased friction at the wall. This has a direct impact on the efficiency of all man-made applications where a solid body interacts with a flow, such as the energy expenditure to transport fluid in pipes, or the propulsive power to move a vehicle.

The classical canonical examples of wall-bounded flows include turbulent boundary layers, pipe flows, and channel flows. Furthermore, numerous examples can be found in multiple engineering devices, such as hydraulic pumps and valves, air conditioning units, heat exchangers, windmills, water treatment channels, internal combustion engines, and automobile exhaust pipes. Some examples are shown in figure 2.1. More specifically, the performance of aerospace devices relies highly on the behaviour of wall-bounded flows, with aircraft wings being a remarkable example. This motivates a strong interest in wall-bounded turbulence in the aerospace engineering community.

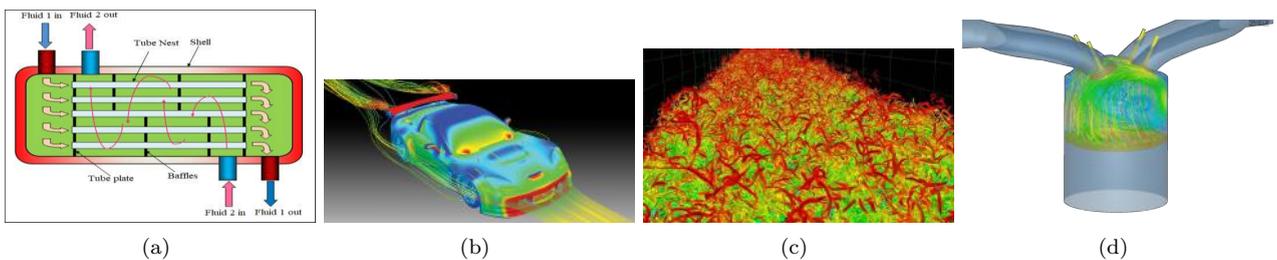


Figure 2.1: (a) Heat exchanger, reproduced from [59], (b) Flow over a car, (c) DNS of a turbulent boundary layer [60], (d) Engine flow simulation

Boundary layers were first described by Prandtl, who noticed a thin layer around the surface with significant viscous effects and an outer region with a nearly potential flow behaviour

where viscous effects are absent [8]. Given the no-slip boundary condition, the relative fluid velocity at the solid surface is zero and gradually increases up to the free stream velocity at the edge of the boundary layer. This theory revolutionised fluid dynamic studies, contributing to understanding turbulence, the transition from laminar to turbulent regime (see below) and the general improvement in aerodynamic designs.

Inertia and viscous forces play a relevant role in turbulent flows. Inertia forces tend to preserve velocity fluctuations, while viscous forces dissipate the turbulent kinetic energy of the flow into heat by the viscous shear stress. When viscous forces dominate, turbulence is damped and the flow remains in a laminar regime, while the prevalence of inertial forces produces turbulent regimes. The Reynolds number is a fundamental non-dimensional group in fluid dynamics that expresses the relation between these two forces:

$$Re = \frac{UL}{\nu} , \quad (2.1)$$

where ν is the kinematic viscosity of the fluid, and U and L are the characteristic speed and length of the problem respectively. Consequently, turbulent flows are characterised by $Re \gg 1$. Strict thresholds for categorisation might be difficult to impose and may be affected by the particular definition of characteristic length and speed, but we can follow this classification as a rule of thumb [1]:

- $Re < 10$: The flow is laminar. Perturbations are easily damped by viscous forces. There are no multi-scale motions. As an example, they can be found in pipe flow with small diameter and/or low bulk velocity.
- $10 < Re < 10^3$: The flow is at an intermediate state, commonly referred to as transitional. Inertia forces become stronger to produce perturbations in the flow although not a self-sustained turbulence. For instance, they can be experienced around a region with increased surface roughness in a pipe flow with a laminar regime.
- $Re > 10^3$: Viscous forces are not capable of damping the instabilities due to strong inertial forces and the flow becomes fully turbulent. The wake behind a boat is an example of turbulent flow.

Locally, the Reynolds number can be defined using the local streamwise coordinate x as characteristic length. If the problem domain is sufficiently large, laminar, transitional and turbulent regimes could be experienced in the flow over a surface with progressively increasing Re_x . The flow, originally laminar, becomes turbulent downstream, and in between we may find the tran-

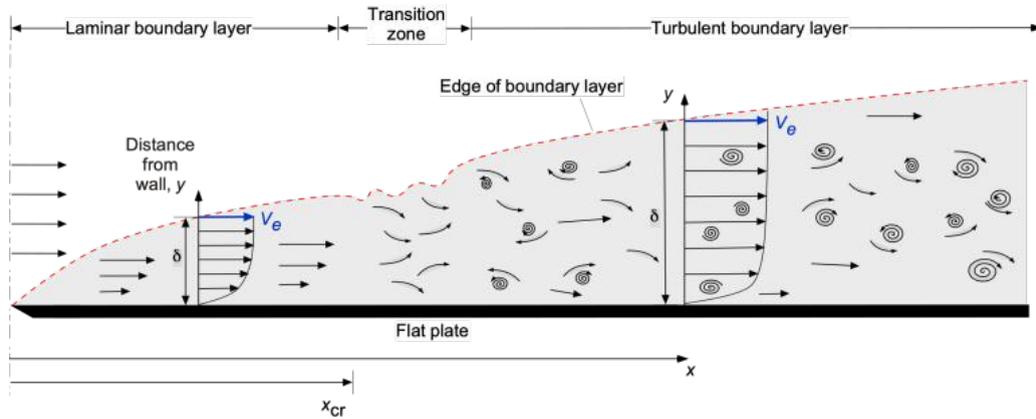


Figure 2.2: The development of a boundary layer over a flat plate. Reproduced with permission from [Introduction to Aerospace Flight Vehicles](#) Copyright © 2022–2024 by J. Gordon Leishman is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](#)

sition region where features of both of them coexist, as represented in figure [2.2](#). As mentioned, these thresholds may change among different problems. In the flat plate sketched in figure [2.2](#), transition might be experienced around $Re_x = \mathcal{O}(10^5)$ [\[61\]](#) or could be triggered earlier by some roughness effects.

In wall-bounded flows, the friction at the wall plays a remarkable role in the balance between inertia and viscous forces. The friction-based Reynolds number Re_τ is commonly used to characterise the state of the flow. Being u_τ the friction velocity, defined in terms of the wall-shear stress τ_w and the fluid density ρ as $\sqrt{\tau_w/\rho}$, Re_τ is given by :

$$Re_\tau = \frac{u_\tau h}{\nu}, \quad (2.2)$$

where h is the outer scale (half-channel size for channel flows, boundary layer thickness δ for boundary layers, pipe radius R for pipe flows). Considering that the speed of the fluid near the wall is significantly lower than the bulk/outer speed, and small scales prevail in it, scaling based on quantities relevant in the near-wall region is often more appropriate. This characteristic velocity and the fluid properties define the inner characteristic length scale as

$$l^* = \frac{\nu}{u_\tau}, \quad (2.3)$$

This quantity represents the order of the length scale of the smallest eddies in the flow. Alternatively, the friction-based Reynolds number can be written as the ratio between outer and inner characteristic length scales:

$$Re_\tau = \frac{h}{l^*}, \quad (2.4)$$

which is of utmost importance for studies on wall-bounded turbulence. It also reveals the relative size of the turbulent scales, meaning the variety of scales present in the flow. The complexity related to the wide range of length scales present in wall-bounded turbulent flows brings the Reynolds number to most practical flows.

The analysis conducted assessing the flow from different scaling conditions revealed the presence of characteristic scales and flow patterns. These can be identified as having time/space coherence, despite the chaotic behaviour of turbulent flows. This coherence opens up the possibility of investigating and characterising a flow regarding the space-time correlation of flow features.

Space-time correlations of velocity fluctuations were first evidenced by Favre et al. [38], [39] by means of hot-wire probe anemometers as in the setup represented in figure 2.3. These works introduced a framework to understand the flow dynamics and structure, strongly influencing both theoretical and practical advancements. The analyses of these space-time correlations were fundamental to the development of more precise turbulent models, providing useful information about how velocity fluctuations at different points of the domain are related in space and time. Examples can be found both for boundary layers [62] and channel flows [63], [64]. In particular, the studies on the structure and behaviour of large eddies in wall-bounded turbulent flows [65], and the detailed investigations focused on the viscous sublayer and near-wall turbulence in pipe flow [66] provided important insights. These studies advanced the understanding of the interaction between large-scale eddies and near-wall turbulence, revealing the critical role that these structures play in the overall dynamics of wall-bounded turbulence and helping to bridge the gap between near-wall and outer flow behaviour. This was even crucial for turbulent wall-bounded flow modelling such as Reynolds-averaged Navier-Stokes (RANS) [67] or large-eddy simulation (LES) [68].

The attached eddy model proposed by Townsend [69]–[71] is an example of a model to explain space-time correlation in wall-bounded flows. It proposes that self-similar wall-attached eddies populate the logarithmic layer of wall-bounded flows. This model does not yet describe the flow structure concerning how eddies are arranged but establishes that large wall-attached eddies are organised, persistent and contain the energetically dominant motions of a wall turbulent flow. This facilitated the interpretation of the contribution of these structures to the turbulence dynamics and energy distribution in the flow. However, observation of instantaneous fields failed to show evidence of the coherent structures that could explain the behaviour of the

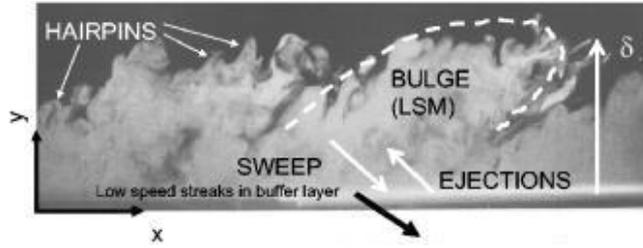


Figure 2.4: Flow visualisation of an LSM from Fig 2 (a) from reference [86] ‘Smoke visualisation of the streamwise-wall-normal plane in a turbulent boundary layer showing various eddy structures’, 2007 © AIP Publishing, reproduced with permission.

was supported by experimental works such as the flow visualisation of a zero-pressure-gradient turbulent boundary layer by Head & Bandyopadhyay [78], which characterised them as very elongated hairpin vortices that extend through most of the boundary layer at high Reynolds number, and as horseshoe vortices, not as elongated, at low Reynolds number. Similarly, a turbulent boundary layer densely populated with hairpin vortices was computationally obtained from the Navier-Stokes equations with a DNS by Wu and Moin [79].

Additionally, the models of large-scale motions (LSMs) and (VLSMs) have been gaining traction in recent decades. The LSMs rely on the hairpin structures seen in Ref. [78], and occur when such hairpins travel sequentially at a similar speed and arrange into a larger structure. Their dynamics and behaviour have been the object of special interest for various studies [80]–[86]. Through the presence of bulges, as seen in the flow visualisation in figure 2.4, a characteristic LSM streamwise size of the order of 2-3 δ (being δ the boundary layer thickness) [75] was estimated. The study of the flow dynamics of this type of event contributed to a better understanding of the role of turbulence in wall-bounded shear flows. A representation of the dynamics of a hairpin is given in figure 2.5. A region with strong velocity fluctuations in the second quadrant ($u < 0, v > 0$) is developed with a certain inclination with respect to the wall below the core connected to the two legs. One important feature of hairpin vortices is the low streamwise momentum region induced in between the two legs near the wall [83], [87], [88].

In relation to VLSMs, low streamwise momentum regions confined between high momentum regions were reported also in the logarithmic and wake regions [81], [84], [89]–[91]. Arrays of hot-wire anemometers were used to study the turbulent flow in pipes and channel flows, reporting streamwise lengths of the structures as long as 25-30 pipe radii R and half-channel heights h , respectively. Similarly, there is evidence of VLSM of the order of 15-20 δ in length in turbulent boundary layers [89]. Importantly, VLSMs seem to scale on outer variables. Hence their dimensions do not change with the Reynolds number [87]–[89], [92].

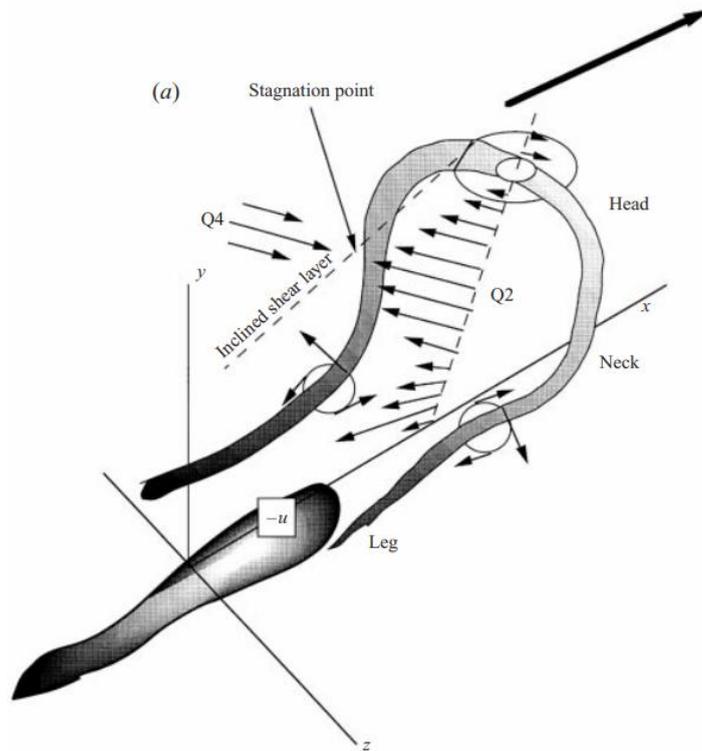


Figure 2.5: Representation of a hairpin vortex from Fig 10 (a) from reference [83] ‘Schematic of a hairpin vortex attached to the wall and the induced motion’, 2000 © Cambridge University Press, reproduced with permission.

LSMs and VLSMs are highly energetic carriers and contribute greatly to the production of turbulent kinetic energy and Reynolds stress. For example, they have been assessed in turbulent boundary layers, pipe flows and channel flows, highlighting the importance of their energy contribution and their interaction with the near-wall small scales as the Reynolds number increases [87]. The wavelengths of LSMs and VLSMs were characterised in pipe flows, and the contribution of VLSMs to the total turbulent streamwise kinetic energy was quantified in about 50%, while Reynolds shear stress due to VLSMs was more than half [84]. A similar study was developed on channel flows and zero pressure gradient turbulent boundary layers, reporting similarities supporting the idea that these motions are developed through similar mechanisms, and showed that the VLSMs may be responsible for the deceleration right above the region where the Reynolds shear stress is maximum [85]. As mentioned above in relation to Townsend’s attached-eddy model, the implications of LSMs and VLSMs seem accentuated with increasing Reynolds number [75]. Following POD-based approaches, it has been shown how the eigenmodes characterising the largest flow scales contain turbulent kinetic energy in very significant proportions [93]. Studies on this matter extend to the present days, frequently introducing AI-based approaches; for instance, explainable-AI has recently assessed the contribution of the

different scales and types of structures to the Reynolds stress [94].

Furthermore, the footprint of these structures at the wall is responsible for the turbulent kinetic energy in the sub-layer near the wall and shows consistency with the attached-eddy hypothesis by Townsend (1976) [89]. Furthermore, these motions that extend far from the wall, characterised by large scales and spatial wavelengths, were found to produce a modulation effect on the near-wall region, both by experimental and computational studies [89, 90, 92, 95–97]. Conversely, the outer layer dynamics and the coexistence of wall-attached and -detached families of vortices seem to be independent of the small structures in the near-wall flow dynamics according to the simulation with the artificial removal of the viscous layer in a channel flow in Ref. [98], and to be independent on the wall roughness [99]. The experiments developed in pipes in Ref. [100] highlighted the existence of detached LSMs mainly concentrate in the wake region and have varied spanwise sizes, while attached LSMs are smaller and more uniform in size, travel at smaller convection velocities and do not extend beyond the logarithmic layer. These attached LSMs may be the responsible ones for the aforementioned modulation effect of the near-wall cycle [95]. The fact that detached LSMs have a span comparable to that of VLSMs suggests that they could be involved in the formation of VLSMs, supporting their formation theory based on pseudo-streamwise alignment [81]. This might not be the case with the smaller wall-attached LSMs.

2.3 Strategies to discover patterns in wall-bounded flows

Wall-bounded flows contain a hierarchy of structures. The data generated by simulation and experiments of turbulent flows are high dimensional. However, the existence of coherent structures suggests the possibility of identifying low-dimensional approximations of turbulent flows to simplify the problem.

In the framework of providing low-rank approximations of turbulence, the theory of compact self-adjoint operators [101, 102] brings the advantage of enabling data compression through linear operators over not-necessarily linear data, and is one of the cornerstones of methods used to analyze wall-bounded flows, among others. Principal component analysis (PCA) decomposes complex systems in modes, preserving as much variability as possible and identifying the most important ones. These modes capture the relations defined by the linear combination of the original features. For example, a PCA-based model contributed to obtaining a more comprehensive model to predict the channel flow behaviour for different wall roughness [103]. Its

application to turbulence gave rise to the proper orthogonal decomposition (POD) methodology [104]. While PCA is used on matrices in a variety of fields, POD, with the same principles, is commonly applied in fluid dynamics with spatio-temporal series [105]. This enabled the identification of patterns within data sorted by their variance contribution, making POD very suitable for the study of the aforementioned LSMs and VLSMs as those are significantly energetic. The use of POD over the past decades has allowed the community to come up with advancements and a better understanding of wall-bounded turbulence [106].

Significant discoveries about patterns in wall-bounded flows were obtained through other data-driven techniques, such as dynamic mode decomposition (DMD). In this case the modal decomposition is applied on the frequency domain, with each dynamic mode being assigned to a specific frequency of the problem [107]. DMD enables the identification of non-linear features in flows with oscillatory elements and reconstruction of those fields [108]. Recent higher-order formulations have further extended the reach of the technique [109], [110]. DMD revealed important information about the cycle followed by certain types of structures in the near-wall region [111].

Alternative methodologies built on the AI framework have also studied the identification of patterns with low-order models. Manifold learning aims to uncover and represent the underlying low-dimensional structure of high-dimensional data, to facilitate analysis and understanding of complex systems, commonly non-linear. For instance, Isometric mapping has been employed to provide a nonlinear, low-dimensional representation of complex shear-flow data while preserving important geometric relationships, thus enabling more efficient and potentially more accurate low-order modelling [112]. As another example, a novel way to understand the effects of various actuation strategies on fluid behaviour was possible through its actuation manifold [113]. This approach has the potential to inform the development of more advanced control algorithms, ultimately leading to improved performance in fluid systems. Furthermore, the dimensionality of a problem can be reduced into a ‘latent space’ through the use of encoders, which are a key component of autoencoders [114], [115].

2.4 Estimation in wall-bounded flows

The aforementioned space-time correlations in wall-bounded flows by Favre et al. [38], [39] was a key milestone that opened the ground for the study of flow estimation from sensors. This covers approaches such as temporal predictions or spatial estimations from different points.

POD-based methodologies have been employed very extensively in this regard [106], [116], [117]. Extended-POD (EPOD) has been gaining momentum in the last two decades (although recently superseded by nonlinear estimators). This linear approach relies on the modal correlation between the decomposition of both elements to be correlated, under the hypothesis of the existence of a certain type of space-time correlation, depending on the particular configuration of the case under study. The extended modes report the correlation between flow quantities, such as the velocity, even with other possible quantities. These extended modes may provide eddy structures interactions or may link velocity modes to other quantities such as pressure [118]. Numerous examples of experimental applications of this linear methodology and derivatives to different fluid-mechanical configurations, such as jets [119], a wall-bounded flow with ribs [120], flow separation in boundary layers [121], pipe flows [122], the fluidic pinball problem or a channel flow [123], can be found in recent literature. A variety of experimental techniques can be found in these references, including particle image velocimetry (PIV) and infrared (IR) thermography.

Analogously, other linear approaches have been implemented with estimation purposes for wall-bounded flows [124]–[133], among which linear stochastic estimation must be highlighted. It can be shown that LSE is equivalent to the sum of the contribution of all extended modes [118]. Given the limitations of linear operators, the comparison between EPOD- and LSE-based methodologies has been discussed, for example, EPOD can be employed to select certain flow features, such as the most energetic ones, by mode truncation. Under the restriction of linearity dealing with higher-order physics, some of these approaches have been shown to successfully improve the estimation results thanks to different implementations intending to include non-linearities. For instance, we can find the extended Kalman filter based approach which introduces them in the Navier-Stokes equations as a forcing term [125], [126], [134]. Examples of works addressing other problems in turbulence and how to introduce non-linear terms in them can be found in the literature, such as the development of Eddy Viscosity Models [135]. An estimator based on POD and LSE efficiently estimates full flow fields from limited measurements by leveraging the POD representation, operating directly on POD bases rather than the full velocity fields [131].

Machine learning (ML) (see chapter 3) has recently shown its great potential to solve flow estimation problems [136]. For example, the heat transfer through a surface and the near-wall region has been estimated from wall measurements of pressure and shear stresses [137], and,

flow field resolution enhancement has been employed in the near-wall region to estimate wall shear stresses [138]. In particular, the flow estimation problem from wall measurements has been addressed with different neural network (NNs) architectures, such as convolutional neural networks (CNNs) [139], [140], recurrent neural networks (RNNs) [141] or fully-convolutional networks (FCNs) [142]. A superior estimation performance on this problem was shown employing generative adversarial networks (GANs) [143].

The importance of including non-linear terms in wall-bounded turbulence predictions has been the object of a longstanding debate. Both linear and non-linear transfer function based methods have been studied and compared, even with multiple inputs, [144]. The effect of non-linearities in a channel flow on the estimation of velocity fluctuations from wall measurements was addressed through multiple-time-delay estimation techniques, showing that non-linearities allow to extend the model into the buffer layer, where linear models were not capable of providing faithful results [145]. Also, Kalman filter estimators have been shown to provide better performances when non-linearities are considered within the approach [126], [134].

Recently, several flow estimation problems in wall-bounded flows have been approached in 3D employing varied methodologies. Velocities near the wall and in the buffer region have been successfully reconstructed from wall measurements using linear stochastic estimation [132]. An approach based on continuous data assimilation showed how to reconstruct the velocity field in a region from the velocity in a contiguous region [146]. Similarly, GANs have been employed for 3D flow reconstructions from 2D velocity measurements [147], or for super-resolution purposes from limited data [148]. Other 3D concepts of CNNs were employed to assess the dynamics of a channel flow [149] and to reconstruct free-surface flows [150]. 3D estimation of turbulent channel flows from wall measurements are still unexplored and are the object of the present thesis.

Chapter 3

Deep Learning for fluid flows

AI is a field of computing science seeking to develop systems capable of performing tasks that may require human intelligence. Typical examples include voice recognition, decision making or problem-solving.

Within AI, ML is particularly relevant, as it focuses on developing algorithms and models that enable machines to learn from data and improve over time with sufficient input. Other key branches of AI not covered in this thesis include natural language processing, which enables tasks like speech recognition and automatic translation; expert systems, which mimic human decision-making in specific domains; robotics, focusing on autonomous decision-making for real-world tasks; evolutionary algorithms, which simulate natural selection to optimize solutions; and computer vision, aimed at interpreting images and videos for tasks such as object recognition and image segmentation.

This chapter contains a discussion on ML, in §3.1 as a broad concept from a general perspective and in §3.2 with an utter focus on the main applications in the field of fluid dynamics. Then, §3.3 provides further details about deep neural networks (DNNs) and the elements employed in the GANs developed in the framework of this thesis. Finally, §3.4 focuses in more detail on GANs.

3.1 Overview on ML approaches

ML applications employ data to learn information, so the system can finally perform a certain task without being explicitly programmed to do so. Depending on how this data is employed, ML approaches are traditionally grouped in the following classes [151], [152]:

- Supervised learning methods employ labelled data and require external supervision. Decision trees, regressions (linear, logistic, ...), artificial NNs or support vector machines are examples of these algorithms.
- Unsupervised learning methods employ unlabelled data and do not require any external supervision. They are employed to identify patterns or to cluster data. Some example algorithms are K-means, K-Nearest Neighbours or Principal Component Analysis.
- Semi-supervised methods rely on partially labelled data or need partial supervision. Some of these algorithms are derived from variations of other supervised or unsupervised learning algorithms and are developed with the purpose of further improving their learning performance.
- Reinforcement learning does not need labelled or unlabelled data, but learns directly from experience, based on trial and error. An agent interacts with an environment and makes decisions to maximise a certain reward. Q-learning, dynamic programming and model-based value estimation are examples of reinforcement learning algorithms.

Recent disruptive advances in ML algorithm development have enabled their application in a wide range of sectors. To mention some examples beyond the scope of this thesis: in industry, they are used to optimise production processes and predictive maintenance; in medicine, for the early diagnosis of diseases and treatment customisation; in marketing and advertising, to customise advertisements and campaigns according to consumer behaviours and preferences; in the financial sector, for fraud detection and risk management. ML is revolutionising the way research and development is approached in multiple fields of knowledge, opening up new challenges and opportunities. Also, scientific research is being massively invested by the advances of ML. For example, ML is used in biology for analysing genomic sequences, in physics to simulate complex phenomena, in economics to predict market trends and optimise investment portfolios, and in meteorology to develop more precise predictive models.

3.2 Machine learning and its impact on fluid dynamics

At the end of section §1.1, some recent technological advancements with a direct beneficial impact on fluid dynamics research have been mentioned. This includes the exponential increase in the computational and processing capabilities of computers and the availability of large datasets from experiments and simulations. This allowed the implementation of ML techniques in fluid mechanics, unlocking new interesting research pathways [101].

A non-exhaustive list includes:

- Prediction of chaotic flows: for instance, RNNs have been employed to predict the evolution of chaotic dynamical systems, enabling arbitrarily long predictions [153].
- Flow estimation from sensors: this problem has been approached with different types of NNs [139], [142], [143], [147]. Convolutional layers are particularly interesting for this purpose due to their capabilities for pattern identification tasks. Further details about this application in the context of channel flows are contained more extensively in the remainder of this chapter.
- Low-order modelling: Manifold learning techniques perform nonlinear dimensionality reduction, with methods such as Isomap [112], [113]. Autoencoders are a type of neural network whose architecture enables the encoder to compress input data in the latent space and the decoder to retrieve information in the original decompressed format from the latent space [114], [115]. With a similar architecture to an autoencoder, but including skip connections, U-Nets also offer this capability [154]. Furthermore, some classical low-order modelling techniques have been integrated into varied ML techniques, such as POD, which have been employed in combination with CNNs [139]. Galerkin regression, which projects the governing equations, such as the Navier-Stokes equations, onto a reduced set of basis functions, is another case of approach combined with modern ML techniques. In Ref. [155], this approach enhances the physical consistency of reduced-order models while identifying simpler, more parsimonious modes.
- Super-resolution: different ML techniques have been implemented to reconstruct high-resolution turbulent flow fields from incomplete or low-resolution samples. This principle was used to increase the resolution of simulations [156] and experimental data [157], [158]. The implementation of these super resolution-oriented methodologies was seen successful thanks to CNNs [159], [160].

- Data repair: Deep learning (DL) techniques have been shown to successfully reconstruct flow databases with missing information [161].
- Acceleration of numerical simulations: the integration of ML techniques to accelerate computational fluid dynamics simulations has enhanced efficiency and enabled fast analysis of fluid flow phenomena [162]–[164].
- Flow control: different configurations and ML models can be found in the literature with active control purposes, leveraging deep reinforcement learning [165]–[167] or genetic algorithms [168]–[171].
- System identification: genetic programming has been used to compute conservation laws and governing equations of flow field configurations [172]. Other notable examples include the sparse identification of nonlinear dynamical systems [173] and the neural-network-based autoregressive models [174].
- Identification of closure models: ML-discovered models are offering new possibilities for closure [175].
- Data assimilation: some of its applications are the reconstruction of unknown regions in the flow [146], the reduction of error propagation in high-dimensional dynamical systems in fluid dynamics [176], and data enrichment/regularisation imposing compliance with first principles, as it is the case of physics informed neural networks (PINNs) where equations are combined in data-driven techniques through the loss function [177], [178].
- Sensor placement: data-driven techniques have been employed for smart sensor placement purposes, including techniques based on the exploitation of patterns in the data [179], [180].
- Uncertainty quantification: ML has contributed to overcoming certain limitations in the uncertainty quantification problem [176], [181].

3.3 Introduction to Deep Neural Networks

Among the several ML approaches, this thesis makes use of DNNs. In this section, the main concepts involving DNNs are introduced and described, paying special attention to the tools employed in this thesis.

Many fundamental concepts of modern DNN methods root their basis in applied mathematics

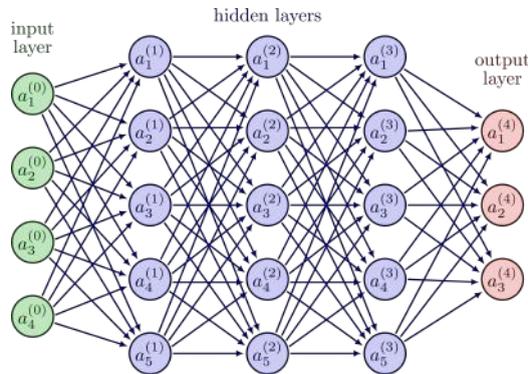


Figure 3.1: Example of Multilayer Perceptron network architecture. By Izaak Neutelings https://tikz.net/neural_networks/, 2021 © TikZ, reproduced with permission <https://creativecommons.org/licenses/by-sa/4.0/>

developments of several decades ago, even if their technical implementation was severely limited due to the lack of computational infrastructure [182]. The first works with NNs date back to the 1960s, with the work by Hubel and Weisel using individual neurons to study the activity in cat brains [183], for which they were awarded the Nobel Prize in Physiology of Medicine. Rosenblatt’s perceptron algorithm was presented in 1958 [184]. Inspired by human neurons, this model establishes a mapping between a vector of inputs \mathbf{x} and an output y . Rosenblatt’s perceptron imposes a linear transformation to these input values with the weight coefficients w_i that multiply the p entries in \mathbf{x} and the bias ζ :

$$y = s(\xi) = s\left(\sum_{i=1}^p w_i x_i + \zeta\right), \quad s(\xi) = \begin{cases} 1, & \text{if } \xi \geq 0 \\ 0, & \text{otherwise} \end{cases} \quad (3.1)$$

The linear transformation is followed by a threshold function that sets the output to 0 or 1, and is an example of what today is known as an activation function. The progressive combination of these units gives rise to NNs. NNs are a type of computational model that consists of a set of units, named artificial neurons, connected among them to transmit signals. The network receives data as input, going through the network follows a set of operations neuron after neuron, and produces some output. Perceptrons can be combined in layers that produce a vector rather than a single output. Then, layers can be combined sequentially in a particular type of artificial neural network known as multilayer perceptron (MLP). A diagram of an MLP is shown in figure 3.1.

The depth of the network is a concept commonly used to refer to the inclusion of multiple layers sequentially. Those layers that are not directly in contact with the input nor with the

output of the network are referred to as *hidden layers*. This brings DNNs and (DL) on stage, becoming very popular during the last decade thanks to their advanced capabilities and the increasing accessibility to the computational resources that allow us to develop them. These layers hierarchically process data and can learn very detailed and complex representations. For that to be effective, large amounts of data are typically needed. Hence, training DL models is computationally intensive and may require advanced hardware such as GPUs or TPUs—specialised hardware accelerators to efficiently handle parallel computing tasks—to handle the large number of computations required.

ML algorithms need to see data containing the information to be learnt. In this learning process, a key element is the loss function \mathcal{L} , intended to be minimised according to the purpose of the problem. This process is organised in epochs (one *epoch* is a complete pass of the full training dataset through the learning process), with the weights of the perceptrons being initialised and progressively updated until the network is trained and ready to work.

Data sets are typically organised into training, validation and testing sets. The training set is the one employed during the training process to adjust the weights recursively, epoch by epoch. The validation set is used during the training process to evaluate the model's performance on unseen data and help tune hyperparameters, such as learning rate, batch size, or regularisation factors. It provides an unbiased evaluation to prevent overfitting, ensuring that the model generalises well to new data. Finally, the testing set is reserved for the very end of the process, offering a final assessment of the model's performance. It is used to evaluate how well the trained model performs on entirely unseen data, ensuring that the model is reliable and effective in real-world applications.

Different loss function definitions may be convenient depending on the particular purpose, the type of neural network or the type of data employed. Some loss functions commonly used are Mean Squared Error (MSE), the Huber Loss, which combines MSE with Mean Absolute Error (MAE), Cross-Entropy, used for classification purposes with variants such as Binary Cross-Entropy, Categorical Cross-Entropy or Cosine Similarity Loss, among others.

The neural network learning process, the training, is based on backpropagation. The concept of backpropagation was formulated almost 40 years ago by Rumelhart (1986) [185]. It defines how the weights are updated from their current value, according to the gradient of the loss function \mathcal{L} with respect to the weights of the perceptron. The learning rate η , typically $\ll 1$, is a very important hyperparameter in the training process and can strongly influence the success

of the network since the weights are updated as:

$$\Delta w = -\eta \frac{\partial \mathcal{L}}{\partial w} . \quad (3.2)$$

Variations in the formulation introduced with (3.1) lead to different layers or elements to introduce in a NN. Activation functions are key elements that decide whether a neuron must be activated or not, introducing non-linearities in the model and hence allowing the network to learn more complex relations and patterns within the data. Different activation functions can have an influence on the learning process, as they have an impact on the gradient propagation. Activation functions can contribute to the efficiency of the model and can help in avoiding problems such as vanishing gradient. There are several commonly used activation functions [186]. The sigmoid function provides with a number in the 0–1 range, and is useful in the last layer in binary classification problems:

$$y = f(x) = \frac{1}{1 + e^{-x}} . \quad (3.3)$$

The Rectified Linear Unit, commonly known as ReLU, has become very popular for its simplicity and effectiveness and is defined as [187], [188]:

$$y = f(x) = \max(0, x) . \quad (3.4)$$

The hyperbolic tangent is also used as activation function, similar to the sigmoid function but with outputs centered at zero in the range -1–1:

$$y = f(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}} . \quad (3.5)$$

The softmax activation function provides a multiple output whose sum adds up to 1, which makes it very suitable for classification problems, to provide a score to each group as a matter of probability. It can be expressed as:

$$f(x_i) = \frac{e^{x_i}}{\sum_j e^{x_j}} . \quad (3.6)$$

There are other two activation functions employed in this thesis, defined on the basis of ReLU (3.4) among other possible variants. Leaky-ReLU provides the same value for positive entries as ReLU, while for negative ones provides a negative number proportional to the entry by a factor to be defined a , typically introducing a small gradient ($a < 1$) [189], [190]. It is

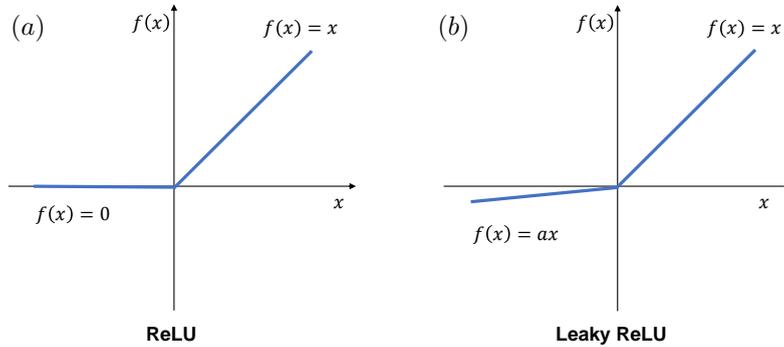


Figure 3.2: Plot of the ReLU (a) and PReLU/LeakyReLU (b) activation functions.

expressed as:

$$y = f(x) = \begin{cases} x, & \text{if } x \geq 0 \\ ax, & \text{otherwise} \end{cases}. \quad (3.7)$$

Parametric-ReLU (PReLU) adds the complexity that this parameter a is to be learnt in the training process, not initially set as for Leaky-ReLU [190]. It allows the introduction of further effects in the learning process while it does not deactivate all neurons receiving entries lower than the threshold. This fact makes PReLU very suitable to model fluctuations of centred data, as it is the case of the methods developed in this thesis, where the network receives the fluctuations of the wall measurements and must provide us with the fluctuations of the velocity field. The difference between these two functions with the classical ReLU definition can be visually seen in figure 3.2.

Other elements can be introduced in the network architecture beyond layers with multiple perceptrons and activation functions. Convolutional layers are indeed a very common feature. They were first defined and used in the 1980's with the NeoCognitron network (see figure 3.3) [191]–[193], showing a great potential particularly for digital imaging purposes. A convolutional layer weights inputs with the coefficients w_i as in (3.1), with the additional complexity that each output element is influenced by its neighbouring region. A kernel contains several weights that multiply the values of the input mapping, and the sum of these values provides the output. The kernel stride defines what step the kernel moves along the input image to map the output [194]. Great pattern and image recognition capabilities are claimed to CNNs [195]. Here the classical 2D convolution operation is defined and illustrated (see figure 3.4), for an input image

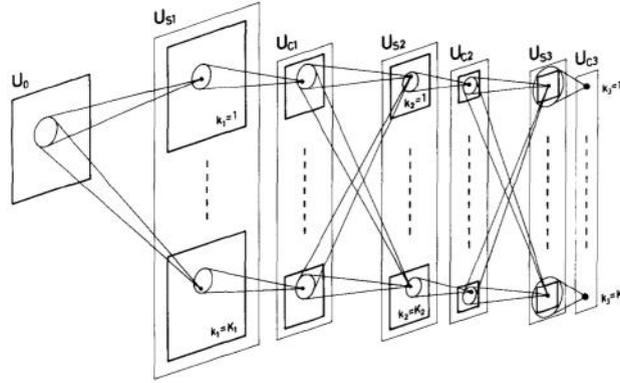


Figure 3.3: Sketch of the first CNN by Fukushima. Figure 2 from [191], 1980 © Springer-Verlag, reproduced with permission from SNCSC.

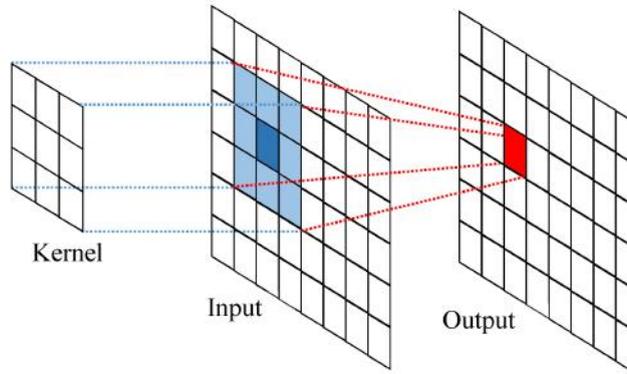


Figure 3.4: Diagram of a convolution operation. Figure 2 from [196], 2017 © by the authors, licensee MDPI, reproduced with permission.

I of size $[i, j]$ and a kernel K of size $[l, m]$:

$$C(i, j) = \sum_l^L \sum_m^M I(i-l, j-m) K(l, m) . \quad (3.8)$$

As was mentioned concerning the possible applications of ML in fluid mechanics, convolutional layers are highly effective in identifying patterns within data, particularly in images. With the kernel scanning the input data, these layers can detect local patterns such as edges, textures, and shapes. Hierarchically, simple patterns are captured in the early layers and more complex features in deeper layers. Each convolutional filter can detect patterns of certain length scales. In this thesis, which focuses on three-dimensional flow estimations, the main element of the NNs employed is the 3D convolution layer. It works exactly in the same way as a 2D layer, with the addition of one extra dimension. A 3D convolution layer receives a 3D feature as input mapping, to be convoluted with a 3D kernel with sizes $[l, m, n]$ that is strided along the domain in the three directions to produce a 3D output mapping. Following this procedure,

higher-order convolution layers could also be defined.

Kernels can act over multiple filter dimensions. The most common example can be a 2D RGB image, with three channels, one for each primary color. Then, the input image decomposed in channels is convoluted separately. This can be represented as an additional dimension of data in the problem. In the case of the CNNs employed in this work, the input filter dimension of the generator coincides with the different sensor inputs from the wall—wall pressure and streamwise and spanwise wall-shear stresses—while each output channel provides a different component of the velocity fluctuation. When multiple layers are included in the model architecture (hidden layers), the filters live in a hyperspace dimensionally different from the input and the output.

Convolutional layers offer different degrees of customisation that must be set during the definition of the model. They are of special importance as they directly affect the output size. The size of the kernel, already mentioned, is one of these parameters. Typically the same size is used along the different dimensions, but not necessarily. The striding defines how many elements the kernel must be displaced between the operations that lead to consecutive elements in the output. Padding encapsulates the input mapping in zeros so that the kernel acts over a slightly larger domain.

A certain region of an image may be affected by other regions as all of them contribute to setting the weights of the kernels in the learning process, but once set, it should not affect as the kernel only acts over the neighbouring pixels, ensuring bi-univocal mapping between the same input and output regions. This is not the case for fully connected layers, also used in this thesis. Those are defined by a size, which determines the number of weights to be used, and all the output elements are seen by all input elements. It does not act locally over data, as convolution operations do. It is used in the discriminator network presented in this thesis, and it is preceded by a flatten layer, which just rearranges a multidimensional tensor in a 1D array.

Furthermore, up-sampling layers are often employed in DNNs, which are also used in the architectures designed in this thesis. They are used to increase the size of the domain. They are commonly used for super-resolution purposes in 2D imaging problems [143], [160]. In the methodology concerning this thesis, they are needed to develop a 3D space from 2D. The original 2D wall images are rearranged in a 3D tensor with size 1 for the wall-normal direction. Then, each up-sampling layer increases the size of this dimension by a factor to be specified, for example, a factor 2 doubles the amount of data increasing the size of the domain in this direction while the rest of the dimensions are kept the same.

Residual blocks or skip connections are recurrently used in these architectures. They do not imply the presence of a specific layer, as convolutional or up-sampling layers, but refer to the way other layers are arranged. These skip connections allow information to pass over some layers, adding or concatenating the output of a layer to the output of a subsequent layer, skipping intermediate ones. One of their main advantages is that they contribute to mitigating the problem of gradient fading in DNNs. They can be easily recognised in the generator architecture in figure 3 of paper 1, with the residual blocks recurrently used and a big skip connection.

The elements mentioned above are examples of building blocks of NNs. The present overview covers the elements adopted in the architectures used in this thesis. However, there is a large dictionary of elements that can be used for different purposes, allowing NNs to deal with a variety of complex tasks and problems [197]. For example, pooling layers such as Max Pooling or Average Pooling layers are used to reduce the size of the input taking respectively the maximum or average value within a kernel. Moreover, normalisation layers such as Batch Normalisation normalise data with its mean and variance, which may contribute to reducing the covariate shift problem [198] and smooth the objective function [199], although gradient explosion issues during initialisation could be induced [200].

3.4 Generative Adversarial Networks

The previous section §3.3 briefly introduced some context about ML and described some important concepts about NNs and the architectures implemented in this work. These concepts are employed to develop a GAN, which deserves special attention and is detailed in this section.

Predictive AI models can employ historical and current data to establish future predictions, often in the form of values, categories or probabilities. On the other side, generative AI algorithms aim to create new samples in formats such as text, images or audio. Generative pre-trained transformers (GPTs) are a common approach for text generation. Other frequently used examples of generative AI algorithms are variational autoencoders, which work with probabilities in the latent space, and GANs [201].

GANs were introduced by Goodfellow and others in 2014 [202], who proposed them under the need to improve techniques for generating synthetic data, such as images, videos, text or music, that would be sufficiently realistic and useful for a variety of applications. Since then, GANs have revolutionised ML, enabling significant advances in the generation of new instances

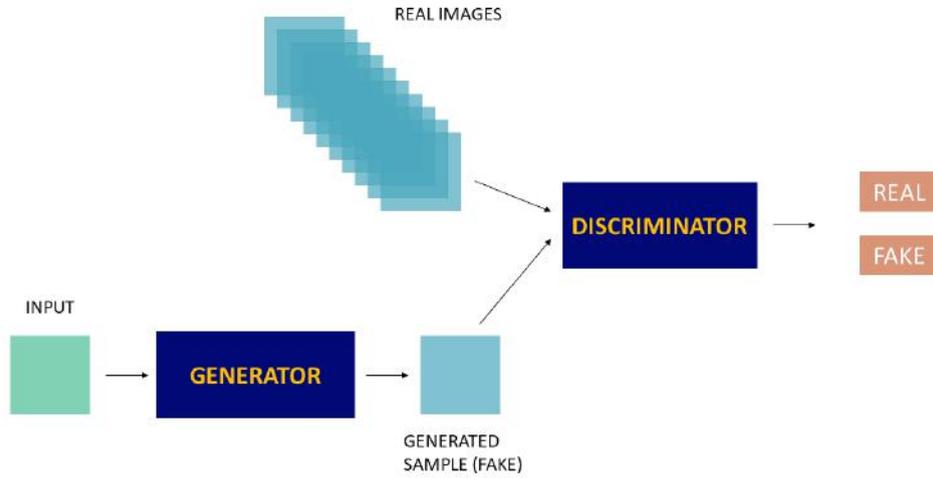


Figure 3.5: Diagram of a GAN.

of data not contained within the original data set according to the learnt patterns.

Based on the game theory, a GAN is composed of two networks, named generator and discriminator. Both networks are in competition against each other in a zero-sum game, with their loss functions evaluating their respective performance through the content loss and giving them a reward based on the bad performance of the opponent network, explaining the term ‘*adversarial*’:

$$\mathcal{L} = \mathcal{L}_{\text{content}} + \mathcal{L}_{\text{adversarial}} . \quad (3.9)$$

The generator is the one used after training, capable of performing the task of interest for which the GAN has been designed. The generator receives some random input and learns to deal with it to generate an output fulfilling certain requirements. The purpose of the discriminator is just to assist the generator during the training process and try to distinguish generated samples from original data samples.

At the beginning of the training process, none of the networks are experts in its task. The generator starts with difficulties in generating an output sample minimally similar to what it should be. The discriminator, even if not yet an expert on its categorisation task, can quickly distinguish original data samples from the poor-quality samples provided by a bad generator. The generator is penalised and learns to generate more realistic samples in an effort to minimize its loss function. The discriminator might start to find its task more difficult as the generator becomes more professional. This pushes the discriminator to improve its capabilities in classifying original and generated samples. At some point, the discriminator is confused in

distinguishing original samples from high-quality generated samples while the generator cannot further improve its performance. The loss functions report stable values and the training process can be concluded. The discriminator is not needed anymore, while the weights of the layers of the generator contain this learning feature capable of providing us with samples resembling the original ones.

The generator of the 3D-GAN implemented in this thesis does not receive random inputs. It is instead fed with wall measurements synchronised with the 3D velocity flow fields it must provide, similar to how other DNN concepts have been used for flow estimation purposes from the wall [139], [142], [143]. Both networks are trained on the basis of an adversarial loss, describing how successful are the generator and the discriminator in their task, and a content loss, reporting how accurate is the generator in preserving certain features of the original data (in our case, minimising the reconstruction error). The loss functions are completely described in equations (2.1)–(2.4) of paper 1. Furthermore, the tools aforementioned in section 3.3 are combined and integrated into a GAN to create a network that is a DNN, makes use mostly of 3D-convolutions, and has been designed and trained for the generation of a particular application in the field of wall-bounded turbulence.

3.5 3D-GANs for flow estimation

One of the main contributions of this thesis is the implementation of a methodology employing ML capable of making 3D flow field estimations from wall measurements. This methodology is covered in detail in paper 1. After sections §3.1–3.4 introducing the framework on which this approach is developed, this section covers certain details about the methodology that are not covered in paper 1. Furthermore, a set of cases with poorer wall data availability is proposed in paper 2. As the network is requested to do the same task, these cases are challenging the methodology.

This work has been done employing—for training, testing and validation—data from a DNS of a channel flow [203]. The solver is a pseudo-spectral code employing the wall-normal vorticity and the laplacian of the wall-normal velocity on the formulation. It uses a Fourier discretisation in the wall-parallel directions and a 7th-order compact finite differences in the wall-normal direction. Time integration is done with a semi-implicit 3-step Runge Kutta procedure and the non-linear terms are computed using a pseudo-spectral method. This algorithm is prepared to employ GPUs to accelerate the simulation.

This DNS simulates a channel flow at friction-based Reynolds number $Re_\tau = 200$. In papers 1 and 2 the dimensions of the channel are defined according to the half of the distance between the channel walls h , with a length of πh in the streamwise direction and $\pi/2h$ in the spanwise direction. The 3D-GAN works directly on data with the same resolution as the DNS, with 64 points in both wall-parallel directions and 64 points from the wall to the centre of the channel. Note that this grid sets equally-spaced points in the wall-parallel directions, but not in the wall-normal direction.

The decision on the size of this channel was the result of a compromise. The channel must be large enough to represent reasonably well the large scales. Large channels, on the other hand, may require high storage capabilities and simulation time. Data-driven techniques require a high number of samples, thus requiring running long simulations. The samples employed in this dataset were sampled about every 0.5 eddy turn-over times, to reduce the correlation between consecutive samples and reach statistical convergence. The eddy turn-over time is a characteristic timescale that represents the time it takes for a large eddy to rotate or ‘turn over’ once. It sets the timescale of the largest energy-containing turbulent structures. Given these limitations and constraints, the channel was defined as smaller than common channels with wall-parallel sizes (see for instance Ref. [204] where the channel has a length and a width of $8\pi h$ and $3\pi h$, respectively), but bigger than a minimal channel unit [205]. Figure 1 in paper 1 shows the characterisation of this channel, with the wall-normal profiles of the mean streamwise velocity and the standard deviation of the three velocity components. This figure offers a comparison with other larger channels from the literature and a minimal channel unit [205] at a similar Re_τ . The most notable difference is found in the peak of the standard deviation in the streamwise component, with an intermediate value between the minimal channel unit and the larger channels. This curve seems slightly shifted to higher y^+ values, mostly due to the small differences in Re_τ .

A key parameter for the choice of the dataset to be employed is the friction Reynolds number Re_τ . Different estimators, including linear estimators and DNNs, revealed that the reconstruction quality is acceptable up to a distance in wall-inner units that remains approximately unchanged with Re_τ [139], [143]. In fact, the capability to sense events at the wall is physically limited by their size and wall-normal distance. In relation to this, as discussed in section 2.2, wall-attached LSMs have a modulation effect of the near-wall cycle, while wall-detached LSMs may participate in the formation of VLSMs, whose size does not change with Re_τ . Thus, in a

high Re_τ channel flow ($Re_\tau \geq 1000$), a good reconstruction accuracy might be only expected very close to the wall. Hence, a low Re_τ simplifies the simulation and the amount of data to deal with, while its discussion would not differ much with respect to a high Re_τ channel flow.

The network we use for flow estimation in its simplest implementation requires pressure and both components of the wall-shear stresses at the wall. The three components of the velocity at the 3D fluid domain are needed for training. The network works on fluctuating quantities according to the Reynolds decomposition [206]. Furthermore, they are normalised with their standard deviation, providing all of them with a similar dispersion, which may facilitate the work of the network.

The details of the GANs used in this work are detailed in paper 1. The type of layers and neurons employed in the 3D-GAN are those described in the previous sections of this chapter. In particular, the generator network, shown in figure 3 of paper 1, employs convolutional layers, parametric-ReLU activation functions and up-sampling layers. This last type is essential for the extension of data in the wall-normal direction, from the wall planar input to the 3D domain, while it keeps the domain size in both wall-parallel directions. In turn, the discriminator network, described in figure 4 of paper 1, employs convolutional layers, Leaky-ReLU activation functions, a flatten layer, two dense layers and a sigmoid activation function. Further details about the training process are provided in paper 1.

Chapter 4

A technique to measure instantaneous convective heat transfer in wall-bounded flows

This chapter explores the use of IR thermography as a non-intrusive measurement technique for estimating wall-shear stresses in turbulent channel flows, which are essential for the DNNs and GANs developed in this thesis (chapter 3, papers 1 and 2). The thesis builds on DNS data but aims to demonstrate the feasibility of applying these algorithms to experimental data. Wall-shear stress is a critical parameter in the GANs presented, yet obtaining 2D time-resolved measurements non-intrusively remains challenging. To address this, the setup leverages IR thermography with a heated thin-foil sensor to measure heat transfer. The Stanton number, derived from heat transfer, correlates with the skin friction coefficient under certain conditions, making it relevant to the current study. In this chapter, we introduce the use of IR thermography to capture instantaneous temperature maps from a heated thin foil and apply energy balance principles to extract the convective heat transfer coefficient. Simultaneous wall and flow field measurements, such as those obtained via PIV, are essential for training and validating the GAN models.

This chapter begins with a discussion on the physical relationship between convective heat transfer and wall shear stress in §4.1, supporting the choice of this experimental approach. Section 4.2 highlights prior work on IR thermography in wall-bounded flows. The characteristics of the heated-thin-foil sensor, including its design, manufacturing, and data processing methods, are covered in §4.3. It also outlines the experimental setup and data processing procedures, with further details provided in paper 3. The utility of heat flux sensors to map the spatial distribution of convective heat transfer across the sensor surface is explained.

4.1 The link between wall-shear stress and convective heat transfer coefficient

In chapter §2 it was seen how different studies evidenced the correlation between certain wall quantities—such as the wall-shear stress—and other flow features near the wall. Hence, experimental techniques to measure wall-shear maps are of interest to obtain pictures of the footprint of the flow structures. However, the implementation of such an approach is extremely challenging with the available state-of-the-art technologies. In Ref. [207] a film wall-embedded sensor of shear stresses was designed. Although measurements were successfully conducted, certain important limitations were highlighted, such as the stability of mechanical properties, the difficulties in sensing low shear modulus or the limitations due to the film frequency response. Moreover, it is challenging to be manufactured. The liquid-crystal coating is a different technique for wall shear stress sensing. In Ref. [208], it was compared with a near-wall PIV measurements that provided more accurate results, although liquid-crystal coating offers full-field accessibility.

From dimensional analysis, comparing the momentum and energy balance equations it is possible to show that it exists a relation between the convective heat transfer coefficient and the wall-shear, referred to as the ‘Reynolds Analogy’ [209]. It is based on the similarity between momentum and internal energy profiles in a wall-bounded flow under steady state, zero-pressure-gradient and unitary Prandtl number (equal thermal and momentum diffusivity). Both transfer effects depend on similar equations, hence they should be represented with profiles with the same shape, analogous to each other.

The Stanton number St relates the heat transfer to a fluid with its thermal capacity in dimensionless terms. It can be expressed as

$$St = \frac{h_c}{\rho U c_p}, \quad (4.1)$$

where h_c is the convective heat transfer coefficient, ρ and c_p are respectively the density and the specific heat of the fluid.

The wall-shear stress can be expressed in dimensionless form using the friction coefficient $c_f = \frac{\tau_w}{1/2\rho U^2}$ and according to the Reynolds analogy, the Stanton number and c_f are related following [4.2].

$$St = \frac{c_f}{2}. \quad (4.2)$$

In this way, h_c (or St) could be directly linked to τ_w (or c_f). However, this approximation (4.2) is only strictly valid for time-averaged wall-bounded flows with zero-pressure gradient conditions ($dP/dx = 0$) and $Pr = 1$. If the last condition does not hold, the following correction should be introduced:

$$St = \frac{c_f}{2Pr^{0.6}}. \quad (4.3)$$

Away from the conditions for which the Reynolds analogy is valid, a strong relation still exists between wall-shear and convective heat transfer. For instance, a recent research work employed CNNs to estimate the heat transfer distributions, employing as input wall-shear stress maps in wall-bounded turbulent flows [137]. The results by Ref. [137] suggest that if the flow field can be estimated from wall shear measurements, this could also be possible from heat transfer measurements. As with other examples, different studies analysed similar correlations before the recent disruption of ML. Strong wall heat flux was found in correspondence of streamwise wall-shear stress fluctuations in a channel flow according to the simulations by [210]. Correlations in the near wall region of a channel with different quantities and derivatives were also a matter of study, including pressure, momentum and thermal streaks [211], further supporting the present choice of measuring the convective heat transfer coefficient h_c .

4.2 Experimental background on IR thermography for wall-bounded flows

Studies on wall-bounded flows involving heat transfer from the wall date back to before the application of IR thermography. For example, thermocouples were employed to measure the plate temperature over which the boundary layer developed, speeds were recorded with hot-wire anemometers, and correlations were established [212]. This type of correlation was also assessed by other experimental studies [213]. The first study measuring convective heat transfer in a wall-bounded flow, employing IR thermography was that by Thomann and Frisk [214] already in the sixties. If one considers time-resolved heat transfer measurement in a turbulent wall-bounded flow, IR cameras were first employed by Hetsroni & Rozenblit, who measured instantaneous heat transfer maps of a heated plate to study its thermal interaction with a particle-laden turbulent flow [215]. The IR instrumentation consisted of an IR scanner with an electronic control unit, a precursor system of the IR cameras that we use today for this type of experiment.

Similar problems have been addressed experimentally, making use of IR thermography to

read the temperature at the surface with synchronised PIV measurements for the flow velocity [216], [217]. This type of experiment may have certain limitations due to the sensor properties, the hardware and the noise affecting the acquisition. It may require the use of a low heat capacity heated sensor as thin as a few microns to make the temperature fluctuations evident, avoiding dampening by the thermal inertia of the material and lateral conduction [218], [219]. Despite this, these implementations might still be limited to low Reynolds numbers. Given the temperature measurements, the heat transfer coefficient can be obtained with a procedure including data filtering to clean noise and solving the energy balance problem for the heat transfer sensor [219]. A POD-based filter was shown to reduce the noise level from the measurements to retrieve the temperature fluctuations, indeed very small due to the high frequencies in turbulent flows, particularly in air flows [220]. A similar filtering approach was recently used in experiments synchronised with PIV to assess the correlation between heat transfer and turbulent structures on the turbulent flow developed over a ribbed surface [120]. At UC3M, recently, synchronised measurements have been developed to study the turbulent boundary layer in a water tunnel facility [221].

4.3 Sensor for IR: implementation and thermal model

The recent literature of IR thermography studies reviewed in section §4.2 highlights certain limitations and aspects to take into account when designing an experiment with this technology. A foil installed on the wall at the region of interest can be employed as a heat transfer sensor, enabling IR thermography to capture the foil temperature distribution needed to estimate the heat flux along the surface. In the heated thin foil sensor, a thermally thin foil is heated by a heat source and cooled by the flow [222].

The convective heat transfer coefficient is obtained from a local energy balance along the foil, considering all input and output thermal fluxes, so as heat sources or sinks within the foil, as sketched in figure 4.1. A simple mean to introduce the foil heating is through the Joule effect, connecting the foil to a power supply to discharge a constant current through it. The power introduced by the Joule effect VI can be considered to be uniformly generated throughout the foil surface area A , leading to a local heat input uniform in space and constant in time. To that end, it is convenient to employ a metal foil such as CuNi (constantan) [221] or CrNi-Steel alloys which ensure thermal conductivity and sufficient mechanical strength to keep shape and size while heated.

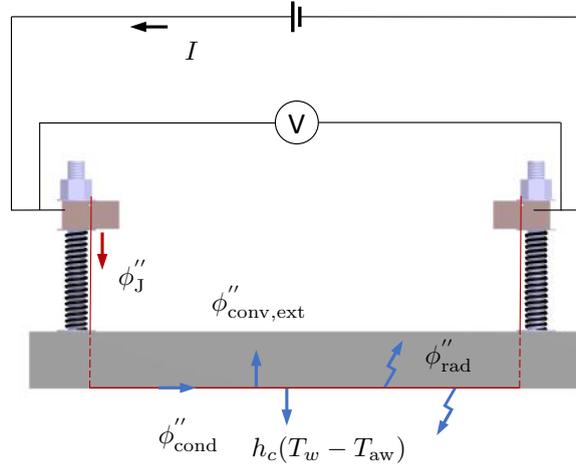


Figure 4.1: Heated Thin Foil thermal model

The fact that the foil is thermally thin is assessed through the Biot number, defined as $Bi = \frac{h_c a_f}{k_f}$, where a_f is the foil thickness and k_f is the foil thermal conductivity. The Biot number compares convective to conductive heat flux through the foil thickness. The foil is considered thermally thin if $Bi \ll 1$. It implies that the temperature is approximately constant through the foil thickness. The Fourier number Fo , defined as $\frac{\alpha t_{char}}{a_f^2}$, with t_{char} being the characteristic time scale, represents the ratio of heat conduction rate to heat storage rate in a material, indicating how quickly temperature changes propagate through it. While it is not a direct measure of thermal inertia, $Fo \gg 1$ suggests that heat conduction dominates over heat storage, making it useful for analysing unsteady heat transfer processes [222].

The convective heat transfer between the fluid and the foil can be quantified establishing an energy balance considering the different heat fluxes and sources. Following the sketch in figure 4.1, conduction through the foil span, convection and radiation on both sides of the foil, the unsteady variation of the foil temperature and the power input should be modelled:

$$c_{p,f} \rho_f a_f \frac{\partial T_w}{\partial t} = \phi_J'' - \phi_{cond}'' - 2\phi_{rad}'' - \phi_{conv,ext}'' - h_c(T_w - T_{aw}) . \quad (4.4)$$

Conduction occurs in a solid or fluid material with a temperature gradient. The tangential conduction heat flux q_{cond}'' along the foil is given by Fourier's Law:

$$q_{cond}'' = -k_f a_f \nabla T_w . \quad (4.5)$$

Note that it is expressed with the Laplacian formulation for a 2D conducting plane. Conduction in the thickness direction has been neglected under the thermally-thin-foil assumption. Radiation should be also considered on both sides of the foil. With this mechanism, energy is emitted as electromagnetic waves from all surfaces. It can be transferred between two surfaces that are not in contact even if there is no medium in between. The heat flux emitted by a surface is given by:

$$q''_{\text{rad}} = \varepsilon \sigma (T_w^4 - T_{\text{amb}}^4) , \quad (4.6)$$

being σ the radiation Stefan-Boltzmann constant and ε the surface's material emissivity, in the range from 0 (no emission) to 1 (blackbody, perfect emitter). Note that the radiation heat flux is added twice in equation (4.4), as it is emitted from both surfaces. This formulation is valid if both sides of the foil are coated (or uncoated) and hence have the same emissivity. Otherwise, different emissivity values for the two faces should be accounted for. Heat transfer by convection occurs between a solid surface and a fluid in contact with the surface. The energy is transferred by the diffusive random motion of particles and the advective motion of the fluid. In natural convection, the flow is due to the buoyancy forces due to the thermal gradient and the differences in the fluid density, while in forced convection the flow is powered by external means. Convection can be expressed with Newton's law of cooling:

$$q''_{\text{conv}} = h_c(T_w - T_{\text{aw}}) , \quad (4.7)$$

which is proportional to the temperature difference between the wall T_w and the adiabatic wall temperature T_{aw} , and to the convective heat transfer coefficient h_c , which may depend on the fluid properties, type of convection, flow field or surface geometry. Convection should be considered on both sides of the foil if exposed to a flow. The coefficient h_c on the internal side is the unknown of the energy balance. The (natural) convection on the foil external side should be studied according to the experimental setup. For example, with a thin foil horizontal on the top side of the channel, it could be modelled according to a natural convection cell developed on a heated plane facing up [223], [224], with:

$$\phi''_{\text{conv,ext}} = h_{\text{ext}}(\bar{T}_w - T_{\text{amb}}) , \quad (4.8)$$

where the average instantaneous temperature within the foil \bar{T}_w is taken and the external convection heat transfer coefficient h_{ext} is quantified as:

$$h_{\text{ext}} = \frac{Nu_{nc} \kappa_{\text{air}}}{\frac{WL}{2W+2L}} . \quad (4.9)$$

being W and L the geometrical dimensions of a rectangular foil. This quantity depends on the thermal conductivity of air κ_{air} and a characteristic length, defined as the ratio between the area of the thin foil exposed to this external convection $A = WL$ and its perimeter $2W + 2L$. The Grashof number is defined as:

$$Gr = g\zeta\bar{T}\frac{\left(\frac{WL}{2W+2L}\right)^3}{\nu^2}, \quad (4.10)$$

where g is the gravitational acceleration and ζ is the coefficient of thermal expansion. For this range of the value of the Pr Gr product, the Nusselt number of natural convection Nu_{nc} is estimated as [223]:

$$Nu_{nc} = 0.54(Gr Pr)^{0.25}. \quad (4.11)$$

Following equation [4.4] to estimate the convective heat transfer coefficient, two acquisitions must be performed, one with heating off, to evaluate the adiabatic wall temperature and one with heating on to measure a sequence of T_w maps. The above heat fluxes should be introduced in the energy balance, so as the Joule effect heat source $q_J'' = VI/A$, being V and I the voltage and current intensity provided to the foil. Concerning the foil thermal inertia, this leads to the foil being heated and/or cooled over time due to the temporal variation of the local convective heat transfer coefficient. Such quenching term can be modelled as $c_{p,f}\rho_f a_f \frac{\partial T_w}{\partial t}$ where $c_{p,f}$ is the foil heat capacity, ρ_f its density, and a_f its thickness.

The IR camera must be sensitive enough to capture these temperature oscillations related to the channel flow turbulence. This depends on the characteristic frequency of the phenomena f and the noise equivalent temperature difference NETD of the camera (hardware). This condition is expressed in [4.12], adapted from Ref. [225]:

$$\frac{\phi_J'' - \phi_{\text{cond}}'' - 2\phi_{\text{rad}}'' - \phi_{\text{conv,ext}}'' - h_c(T_w - T_{\text{aw}})}{\rho_f c_{p,f} a_f f \text{NETD}} \gg 1. \quad (4.12)$$

Several important considerations involving the setup might enhance the quality of the acquisition. All materials do not emit and absorb radiation power in the same way, even if they are at the same temperature—such as room temperature—because of their different emissivities. Avoiding having different objects or tools next to the experimental setup, in particular those whose position may enhance direct radiation over the foil, can help mitigate thermal reflections. The same happens with the IR camera and the lens. The IR camera actually measures radiation and obtains temperature maps from a calibration: focusing the camera with a small tilt angle, not directly in front of the foil, can be beneficial to avoid self-reflections. Furthermore, depend-

ing on how reflective the foil is, the foil might be coated. For example, if the IR experiment is combined with PIV and the foil is in the field of view of the PIV camera(s), a black-matte coating on that side of the foil may significantly improve the quality with which PIV particles are imaged. Also, a high emissivity paint coating on the face seen by the IR camera may facilitate the data acquisition. However, this may increase the effective heat capacity of the foil [219] leading to higher thermal inertia and an attenuation of temperature fluctuations. Whenever applied, the emissivity of the paint needs to be accounted for in the thermal inertia and the radiation modelling.

The implementation of a heated thin foil sensor for the measurement of time-resolved convective heat transfer maps in a turbulent channel are presented in paper 3 where the challenges related to temporal sampling, foil thickness and heating, and measurement uncertainty are discussed in detail.

Chapter 5

Main contributions and conclusions

The introductory part of this thesis discusses the importance of studying wall-bounded flows and their implications in different human-made devices. Particularly, there is an interest in designing active flow control systems to enhance the efficiency of those devices. New flow-sensing strategies need to be investigated to that end, a problem that has been the subject of study of many research works in the last decades. The technological availability of non-intrusive sensing techniques suitable for this kind of application makes it necessary to rely on measurements of other quantities, that fortunately are related to the velocity field. Hence, the flow field measurements for wall-bounded flow control applications might be indirect. A more complete and reliable fluid state could be estimated through more complex sensing systems and algorithms, which in turn could be used in a more efficient actuation system.

Furthermore, estimating the velocity field from these non-intrusive measurements can provide us with information about the physics involved, which is not a completely solved problem yet. The successful flow estimation to some extent indicates the existence of a certain relation between the patterns of the sensed quantities and the flow field. This information is crucial for achieving a more complete understanding of turbulence in wall-bounded flows. For example, this knowledge could be employed to get more efficient geometrical designs or actuation systems.

This thesis takes advantage of the recent disruption of data-driven techniques, aiming to provide us with a series of advancements with regard to AI-BASED SENSING OF TURBULENT WALL-BOUNDED FLOWS. These novel tools can be employed to develop a 3D flow field estimator based on instantaneous wall measurements, relying on the coexistence of wall and flow patterns sharing certain similarities to be seen by the data-driven algorithms.

5.1 Main highlights

Paper 1

One of the motivations of this thesis comes from the state-of-the-art research works addressing this type of flow estimation in wall-parallel planes. Successful flow reconstructions were possible by employing novel concepts of DNNs, such as GANs. However, these models presented some limitations, such as the lack of a 3D flow estimation. The 3D flow estimation can provide a more complete information set, not just because it has an additional dimension. It can be employed to assess a characterisation of the flow patterns as turbulent coherent structures and their footprint. The fundamental hypothesis of this thesis is that similar NN concepts working on a 3D domain may be capable of estimating the full velocity field.

The work covered in paper 1, making use of ML, is introduced in chapter [3](#) presenting the 3D-GAN methodology for this purpose. This work showed the capability to do such flow estimations at any wall-normal distance with a comparable accuracy to that of GANs for wall-parallel flow estimations. Computationally, it overcomes certain limitations. A single network can predict the 3D flow field, while multiple 2D networks are needed for a 3D characterisation. In terms of computational resources, such as the number of trainable parameters or training time, the 3D approach, of course, requires more resources, but fewer resources if one considers the amount of information being provided.

These results are coherent with the analogous 2D estimations and other approaches covered in the literature. The estimation accuracy changes with the wall distance, and the metrics for the streamwise velocity component are slightly better than for the other two components.

Different estimation volumes were defined, being the biggest one the volume comprised between the wall and the half-channel plane. This comparison reported different metrics for each volume case at the same wall-normal distance, suggesting that a footprint of the flow patterns on the wall may influence the quality of the estimator. The flow footprint, mainly composed of small-scale features in the near-wall region, makes the correlation difficult in the region populated with large scales far from the wall.

The flow fields must be also assessed concerning how precisely they respect the physics. The original and estimated mean squared velocity fluctuations and shear stress have been compared, showing a good agreement close to the wall, progressively lost towards the centre of the channel with the lack of capability of the network to estimate the velocity fluctuations. Another issue

compromising the physics is the continuity in the domain. Although the continuity equations are not included in the model, the 3D-convolution operations introduce an advantage with respect to the 2D approach, reporting a much lower standard deviation of divergence.

The main advantage of the 3D approach for the assessment of the reconstruction is the fact that coherent structures can be identified and completely located in space. The accuracy (or error) on the estimation at some wall-normal distance, previously mentioned, is not uniform for this distance. This is an average for all the elements comprised. While those belonging to a wall-attached event may have a strong footprint and have higher chances of being properly estimated, those wall-detached ones might be unseen by the network. Both effects are combined; not all wall-attached events are estimated with the same accuracy, being better reproduced if they remain within the near-wall region in all their extension. At the same time, wall-attached events with a deep penetration towards the outer region should be reproduced much better than detached events at a similar wall-normal distance.

These trends are general for all types of coherent structures. However, if the estimation of sweeps and ejections are evaluated separately and compared, some differences are found, coherent with the physics that explains them. The higher presence of wall-attached sweeps, generally smaller than ejections and hence concentrated in the near wall region, estimates sweeps less undistorted in this region. On the other hand, wall-attached ejections are less quantitative but much bigger and have more chances to populate regions far from the wall than sweeps. The data-driven algorithm learns better these patterns and estimates ejections more accurately than sweeps far from the wall.

Paper 2

Despite the potential of the 3D-GAN as a flow estimator, different aspects may limit a real implementation of the methodology in a laboratory environment, a step needed to design a prototype. This is, employing measured values, both for the velocities and for the wall-sensed quantities, instead of samples from a simulation. One of the main challenges would be to install that amount of sensors with the DNS resolution on the wall. The hypothesis introduced at this point states that employing fewer sensors, which certainly may simplify the setup, would make the estimation less accurate but the 3D-GAN would still be capable of estimating the flow to some extent.

Different approaches introduced in paper 2 intend to reduce the number of sensors. One of these approaches was to reduce the resolution of the sensor maps by increasing their spacing,

in both streamwise and spanwise directions. Moreover, it was done by employing only one of the three wall-sensed quantities, instead of the three of them. These two approaches were combined, bringing a set of scenarios. It can be concluded that the estimator capability is not only driven by the footprint of the structures, but also by how completely is that footprint being sensed with an alternative sensor disposition. The three wall quantities considered do not correlate with each of the velocity fluctuations in the same way, and as their patterns are of different characteristic sizes, some patterns may be more susceptible to not being properly sensed under a lower sensor resolution than other types of patterns.

Given this situation, a new hypothesis was brought, suggesting that an estimation improvement is possible with the same number of sensors if those could sense a higher proportion of the footprint under different streamwise and spanwise resolution configurations. Being patterns generally streamwise elongated, a lower streamwise resolution was proposed in exchange for a higher spanwise resolution. It was shown that the problem is sensitive to the sensor placement. However, the alternative configuration was not capable of improving the capabilities of the estimator, with an additional loss certainly marginal.

Another possible limitation for the 3D-GAN in an experimental environment and for other analogous estimators is the presence of noise in the measurements. Its effect has been investigated by introducing Gaussian noise on the simulated data. The noise effect is minimal in the viscous sublayer but further from the wall, it may introduce significant penalties, which are more significant when more sensors are employed. This is combined with other effects as the proportion of the footprint is being sensed with a given configuration. Closer to the centre of the channel, where the estimation capability is already poor, noise has minimal implications.

Paper 3 and synchronised measurements

Simulated data sets can be employed to design a model, as in this case, a data-driven estimator model. However, to achieve the real implementation of a model it is necessary to test it relying on experimental data, validating that the model can work with measurements taken with available acquisition technologies. This problem should overcome limitations due to this step from simulated to experimental data sets. For example, the resolution with which a certain parameter can be sampled might be lower, or the capability to take measurements of certain quantities easily available from a simulation can be limited experimentally, due to technical or technological reasons.

The wall shear stress is one of the wall quantities fed to the estimator, as in papers 1

and 2. This was possible from simulated data, however with the technological developments nowadays, the acquisition of shear stresses as mappings along a surface does not seem accessible. Fortunately, the estimation of heat transfer mappings is possible, while the correlation between wall shear stress and heat transfer has been studied long. This motivates the development of an experimental campaign to measure heat transfer at the wall through IR thermography.

This experimental approach is however non-trivial. The previous applications of this sensing technique in wall-bounded flows according to the literature reveal certain limitations and difficulties in this regard. Despite the high-frequency and small fluctuations aimed to be measured, another objective of this thesis is to develop a reliable wall heat transfer sensing strategy, suitable for wall-bounded flows, capable of overcoming the limitations there might be.

5.2 Future work

The main research direction for future development are outlined in the remainder of this section. Future work emerging from this thesis can build upon the foundational results and methodologies developed. A promising direction involves addressing the limitations encountered, such as the low Reynolds number used in the simulations. Refining computational models or experimental setups to handle higher Reynolds numbers could provide more intriguing results. Integrating advanced ML algorithms or improved sensor technologies may also offer new insights and enable the study of more complex systems. Expanding these methodologies to larger-scale, real-world applications could validate the findings on a broader scale and increase the impact of the research.

5.2.1 Development of more efficient estimation techniques

One of the main challenges was the high computational and memory demands required for working with higher Reynolds numbers, both for data generation via DNS and for implementing the 3D-GAN. This would be interesting to explore further, as it would allow testing the model in a more turbulent environment with a wider range of eddies and flow structures. Although the literature suggests that model accuracy may remain comparable under different channel flows in inner scaling, direct comparisons could highlight additional implications.

The rapid evolution of AI and ML continues to transform R&D in different fields, and this problem may benefit from alternative and newer NN architectures, which might offer new capabilities and advanced performances. Testing novel architectures or making adjustments

to the 3D-GAN, such as incorporating PINNs, could help overcome current limitations by introducing physical constraints during training. These modifications might enable the network to better handle the peculiarities of this flow-sensing problem.

This thesis is based on the physical relationship between the patterns exhibited by certain quantities in the wall and in the fluid. Under this pretext, a similar methodology could be used to study this flow estimation problem in different geometries, such as in pipe flows.

To develop new active flow control strategies, advanced flow-sensing sensing techniques might be needed and may offer great potential. The estimated flow field could be employed by the controller to determine the actuation to be commanded. An important feature of the 3D-GAN is the fact that it can characterise turbulent structures in its 3D domain, not just to identify the presence of a structure at a certain wall-normal distance with an unknown extent. In turn, the computational cost associated with this network is substantially larger than that of simpler models, as the amount of data is generated from each wall sample. All this would be very valuable information to compute the actuation required. Nevertheless, this amount of 3D data could be reduced, just by retaining the location and certain features of the patterns to be targeted by the actuation. To that end, there are certain ML techniques as discussed in chapter 3 that would allow us to obtain a low-dimensional manifold or a latent space with this information. It is important to obtain this set of reduced information from flow data estimated from wall measurements, as is the case of the 3D-GAN. If this reduced order model were directly obtained from the current velocity field, the footprint filtering effect of the estimator would be lost—the structures that are seen/unseen from wall samples. Finally, it might be interesting to study the estimation of this low-dimensional space with the information needed for actuation systems directly from wall measurements.

5.2.2 Addressing effects towards experimental flow sensing implementation

Another challenge that might be found in the real implementation of wall-sensing techniques for active flow control applications is the need for quick response calculations. If the calculation is not fast enough, the flow may undergo substantial changes when the actuation is computed compared to when the wall information was sensed. Studying flow estimation problems with a temporal shift between the sampling of the wall and the fluid information would allow us to mitigate this problem with some time margin for the calculation. The flow estimation of the 3D-GAN could incorporate this temporal shift, such that the flow fields embody also a temporal prediction.

In this thesis, some aspects of the resolution with which sensors can be installed on a wall are addressed. The setup can be simplified by employing fewer sensors, while their size or other installation restrictions can enforce a sensor resolution lower than that in the DNS. Some effects due to a poorer resolution have been studied from a general overview. However, it would be interesting to address this problem by attending to specific sensor technologies that could be mounted on the wall.

The potential of the 3D-GAN has been shown by employing different simulated wall-sensed quantities. To advance towards a real prototype, it is needed to validate the methodology employing experimental data. This includes the assessment of new challenges due to the shift from simulated to experimental data sets. For instance, experimental data might be obtained with poorer resolution, with a certain noise level, or without the periodic channel condition of the DNS data set.

5.2.3 Experimental validation with synchronized PIV and IR measurements

This section suggests an experimental setup for the acquisition of synchronised measurements of temperature and velocity fields. These data could be employed to validate data-driven-based estimators such as the 3D-GAN. To that end, the techniques employed are IR thermography, described in chapter 4, and PIV. The setup for the temperature measurements is certainly particular, and paper 3 is a dedicated article describing the design of the heat transfer measurement setup. This section covers the details of the synchronisation and an assessment of the correlation between heat transfer and velocity measurements.

This campaign has been conducted in the channel flow facility at Politecnico di Torino, with $Re_\tau = 220$ and a bulk airspeed of 4.95 m/s. The velocity measurements are taken at a frequency of 15 Hz, limited by the maximum frequency of the laser hardware employed. The temperature measurements are taken at 180 Hz since a high temporal resolution is needed to compute the temporal derivatives and quantify the heat transfer. Although this difference in frequencies, measurements are synchronised, with 12 phased temperature measurements for each velocity measurement. The temperature measurements were taken with an Infratec Camera ImageIR® 6300Z, with a resolution of 640×512 pixels. Detailed information about this acquisition, the heated thin foil system, the temperature filtering procedure and the heat transfer quantification is provided in paper 3.

The methacrylate channel walls allow optical access to the flow section for the PIV acqui-

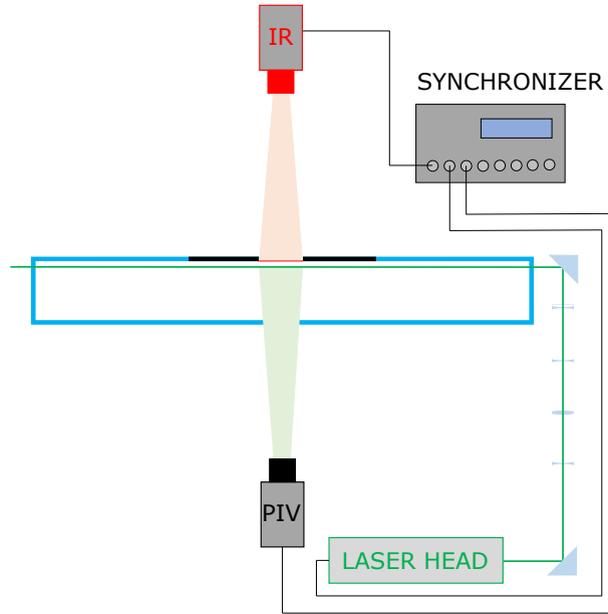


Figure 5.1: Side view of the experimental setup, with synchronised acquisitions with the IR camera (top) and the PIV camera (bottom). The black wall section represents the frame where the thin foil (red) is mounted

sition. The optics system develops a thin laser sheet that penetrates the channel through the sides to illuminate the region of interest during the PIV acquisition time. The PIV camera operates from the bottom part of the channel, with the laser sheet horizontally oriented, parallel to the channel walls. Simultaneously, the IR camera measures the temperature at a heated thin foil fitted on a frame flash mounted on the top wall, coincident with the position of the modular wall it replaces. A schematic representation of the process and the camera layout is shown in figure [5.1](#).

A camera Andor Zyla 5.5 has been used for planar PIV. It has a resolution of 2560×2160 pixels, with pixel size equal to $6.5 \mu\text{m}$. The mounted lens has a focal length of 60 mm. The flow is seeded with DEHS (Di(2-ethylhexyl)sebacate) particles for PIV purposes with $1.2 \mu\text{m}$ diameter. The flow region of interest for PIV is illuminated with a Litron Laser Nano L 200-15 PIV, with power supply LPU550, providing 532 nm green light up to 200 mJ per pulse. The maximum frequency is 15 Hz per laser head. A system of mirrors and lenses allows us to develop a thin sheet (with thickness smaller than 1 mm) parallel to the wall at the desired distance. Two cylindrical lenses with focal length $f = -50$ mm (divergent) expand the beam into a sheet. Two spherical lenses with focal lengths $f = 100$ (convergent) and $f = -75$ mm (divergent) make the sheet thin and uniform within the region of interest. The last mirror places the laser sheet at a distance of about 30 wall units from the channel wall and the thin foil.

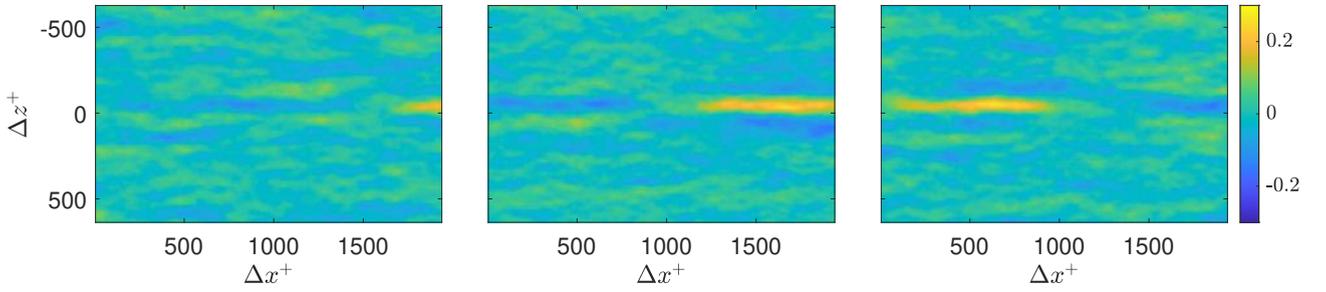


Figure 5.2: Correlation maps of the streamwise velocity maps to the Stanton number in the centre of the foil at different time phases. The central map corresponds to the same trigger for the velocity and the temperature acquisition. The map on the left corresponds to three IR trigger periods in advance. The map on the right corresponds to a delay of 5 IR trigger pulses.

Each image pair has been acquired with a time interval of $180 \mu\text{s}$ between the shooting of the first and the second laser. The original images have been pre-processed with a POD-based background removal algorithm [226]. These filtered images have been used for the PIV processing, following an iterative multipass algorithm [227], [228]. It starts with a window size of 64×64 pixels and follows with two additional iterations with a window size of 48×48 pixels, always with a window overlapping parameter equal to 0.75.

The correlation between both measurement sets has been evidenced with a temporal sequence of pairs of streamwise velocity and heat transfer synchronised maps. The temporal correlation has been computed for each point of the velocity map with respect to the Stanton number in the centre of the foil. This leads to a map showing a region with a high correlation with the central point. Besides, both adjacent regions in the spanwise direction report a negative correlation. Regions with negative correlation are also present upstream and downstream of the highly-correlated region. These patterns may also be useful for an estimator. The rest of the domain reports fluctuating, rather low correlations.

Interestingly, this acquisition contains 12 temperature time-resolved maps for each velocity map, each of them with a phasing of 180 Hz. Consequently, the correlation of a velocity map can be assessed for the temperature acquisition with a coincident trigger or for that with a slightly advanced or delayed trigger. Three examples are shown in figure 5.2, showing the shifting of the correlation patterns with the phasing. When the Stanton maps correspond to instants before the PIV image acquisition ($3 \times \frac{1}{180\text{Hz}}$ for the left map), the high correlation region is limited to a small region downstream. When both triggers coincide, this region elongates and moves upstream. It further moves upstream when the IR trigger is delayed ($5 \times \frac{1}{180\text{Hz}}$ for the right map).

The high temporal resolution requirement for the IR acquisition brought the possibility of addressing correlation effects with not necessarily time-coincident quantities. This can be interesting, for example, for flow control applications, as in a hypothetical case where the actuation is computed for the state of a flow some instants ahead. It also reveals a difference with respect to the data fed to the 3D-GAN in this thesis. The best correlation for instantaneous measurements is found at a certain distance downstream. If one aims to predict the flow field in time, it is further shifted downstream.

It has been shown that this setup configuration is capable of providing correlated measurements through synchronised acquisitions. It deals with limitations, such as mounting both camera systems providing them with their respective optical access without intercepting each other, or the different difficulties in sensing heat transfer fluctuations (which are very small and respond to high-frequency events) which require a careful experimental design as discussed in paper 3. This experimental configuration can be employed to advance towards prototypes in the flow estimation from wall measurements.

5.2.4 Going beyond the training dataset

Experimental measurements of the flow in 2D, as those described in section [5.2.3](#), could be employed in a 2D-GAN for planar estimation. This task may benefit from the use of transfer learning techniques from a pre-trained model with simulated data so that the model with experimental data begins the training process with the capability to recognise certain patterns similar to those it should learn to find. This model will learn discrepancies between the simulated and the experimental data, adapting to the experimental noise, measurement differences, and other specific details of the experimental data. The model would need to deal with data from a non-periodic channel flow. 2D flow fields acquired at various wall-normal distances could be combined to train a 3D estimator, which might necessitate integrating alternative NN approaches with the 3D-GAN. Additionally, a 3D estimator—such as the 3D-GAN—could employ experimental data acquired from tomographic PIV.

Analogously, alternative experimental campaigns could be proposed to work with other wall quantities than heat transfer. Those may have new challenges to be addressed for a successful acquisition procedure. For instance, pressure probes could be installed along the wall surface to obtain a synchronised data set employing a pressure scanner. The experimental setup proposed in this thesis concentrates most challenges in the heated thin foil sensor and its processing, while the foil coating facilitated the PIV acquisition. In this alternative campaign for pressure

sensing, a possible challenge would be found in the PIV acquisition, with reflections and an irregular surface with probes embedded.

Furthermore, the feasibility of conducting a similar experimental campaign on a channel at a higher Reynolds number could be assessed. It would be useful to test the applicability of the flow-sensing technique under turbulent scenarios closer to those encountered in human-made devices. This may bring further complexities, such as higher frequency events and a higher sampling frequency. With the experimental setup considered in the thesis, a higher flow speed would require a higher acquisition frequency. It might be limited by the hardware of the camera, although temporal resolution could be gained at the cost of some spatial resolution.

Part II

Papers

Paper 1

Three-dimensional generative adversarial neural networks for turbulent flow estimation from wall measurements

Antonio Cuéllar, Alejandro Güemes, Andrea Ianiro, Óscar Flores, Ricardo Vinuesa and Stefano Discetti, *Journal of Fluid Mechanics* (2024), vol. 991, A1, [doi:10.1017/jfm.2024.432](https://doi.org/10.1017/jfm.2024.432)

Different types of neural networks have been used to solve the flow sensing problem in turbulent flows, namely to estimate velocity in wall-parallel planes from wall measurements. Generative adversarial networks (GANs) are among the most promising methodologies, due to their more accurate estimations and better perceptual quality. This work tackles this flow sensing problem in the vicinity of the wall, addressing for the first time the reconstruction of the entire three-dimensional (3-D) field with a single network, i.e. a 3-D GAN. With this methodology, a single training and prediction process overcomes the limitation presented by the former approaches based on the independent estimation of wall-parallel planes. The network is capable of estimating the 3-D flow field with a level of error at each wall-normal distance comparable to that reported from wall-parallel plane estimations and at a lower training cost in terms of computational resources. The direct full 3-D reconstruction also unveils a direct interpretation in terms of coherent structures. It is shown that the accuracy of the network depends directly on the wall footprint of each individual turbulent structure. It is observed that wall-attached structures are predicted more accurately than wall-detached ones, especially at larger distances from the wall. Among wall-attached structures, smaller sweeps are reconstructed better than small ejections, while large ejections are reconstructed better than large sweeps as a consequence of their more intense footprint.

Keywords: Turbulent boundary layers, Channel Flow, Machine learning

1 Introduction

The ubiquitous nature of turbulent flows motivates the need for control to enhance the performance of a wide variety of devices. However, closed-loop control of turbulent flows (Choi, Moin & Kim 1994 [58]) requires continuous monitoring of their state. It is of utmost importance to be able to sense the flow state with minimal intrusiveness. Sometimes non-intrusive sensing is the only option available. This is the case of wall-bounded flows, making it possible to embed sensors within the wall. Non-intrusive sensing of turbulent flows has been the subject of several studies in the past decades. Machine learning has revolutionized the field of fluid mechanics (Brunton, Noack & Koumoutsakos 2020 [152]; Mendez et al. 2023 [229]), including both experiments (Discetti & Liu 2022 [230]; Vinuesa, Brunton & McKeon 2023 [231]) and simulations (Vinuesa & Brunton 2022 [164]). As such, the recent advances in machine learning and the wealth of available computational resources offer new interesting avenues for flow sensing.

The estimation of flow velocity solely on the basis of wall measurements was first explored using linear methods, such as linear stochastic estimation (LSE; Adrian 1996 [232]). The use of LSE was successful for the reconstruction of large-scale wall-attached eddies (Baars, Hutchins & Marusic 2016 [130]; Suzuki & Hasegawa 2017 [134]; Encinar, Lozano-Durán & Jiménez 2018 [233]; Illingworth, Monty & Marusic 2018 [234]; Encinar & Jiménez 2019 [132]). This methodology is capable of reconstructing a certain range of length scales of the structures populating the vicinity of the wall (i.e. the buffer layer) with reasonable accuracy. In the region farther from the wall, only large-scale motions are generally captured. These reconstructions can be more sophisticated by supplementing the methodology with further instruments in order to manipulate the filtering of scales, retaining and targeting the reconstruction over a broader spectrum. For example, in the work by Encinar & Jiménez (2019) [132] with a turbulent channel flow in a large computational domain at a high friction Reynolds number, the large-scale structures containing about 50 % of the turbulent kinetic energy and tangential Reynolds stresses are reconstructed successfully up to $y/h \approx 0.2$, while only attached eddies of sizes of the order of y are sensed in the logarithmic layer.

An alternative linear approach is the extended proper orthogonal decomposition (EPOD) (Borée 2003 [118]), which can establish a correlation between input and output quantities through the projection of their corresponding proper orthogonal decomposition (POD) modes (Lumley 1967 [104]). Despite leveraging only linear correlation, EPOD presents the advantage of being able to target specific significant features in a space of reduced dimensionality. A

non-exhaustive list of examples of EPOD applications to the reconstruction of turbulent flows includes the estimation of the low-dimensional characteristics of a transonic jet (Tinney, Ukeiley & Glauser 2008 [119]), wakes behind wall-mounted objects (Bourgeois, Noack & Martinuzzi 2013 [235]; Hosseini, Martinuzzi & Noack 2016 [236]), wing wakes (Chen, Raiola & Discetti 2022 [123]), turbulent channel flows (Discetti, Raiola & Ianiro 2018 [237]; Güemes, Discetti & Ianiro 2019 [139]) and even high-Reynolds-number pipe flows (Discetti et al. 2019 [238]). The limitations in terms of reconstruction capabilities and spectrum range found with EPOD are similar to those with LSE.

Lasagna et al. (2015 [145]) studied multiple-time-delay estimation techniques. Although linear methods provide accurate estimations in the viscous layer, nonlinearities must be considered to extend the reconstruction into the buffer layer. Also, Chevalier et al. (2006 [126]) and Suzuki & Hasegawa (2017 [134]) highlighted the importance of incorporating nonlinear terms to get a more accurate estimation with a Kalman filter. Following the seminal work by Milano & Koumoutsakos (2002 [136]), neural networks emerge as an alternative able to cope with nonlinear relations between sensor and flow features. Recently, deep-learning algorithms have been leveraged for flow reconstruction from sensors. For example, the laminar wake of a cylinder and the flow in a turbulent channel have been reconstructed successfully in two dimensions from coarse measurements with convolutional neural networks (CNNs) (Fukami, Fukagata & Taira 2019 [156], 2021 [239]). The performances of LSE and CNNs in estimation from wall measurements in a turbulent channel flow have been compared by Nakamura, Fukami & Fukagata (2022 [240]), reporting that linear models can provide comparable results at the cost of establishing a nonlinear framework to combine and provide the inputs to the system. Nevertheless, the use of nonlinearities through CNNs can be very effective and neural networks seem more robust against noise than LSE. Burst events in the near-wall region such as ejections and sweeps were studied by Jagodinski, Zhu & Verma (2023 [241]), with a three-dimensional (3-D) CNN capable of predicting their intensities, and also providing information about the dynamically critical phenomena without any prior knowledge. For the specific task of estimation of flow velocity from wall sensors, Güemes et al. (2019 [139]) proposed using CNNs to estimate temporal coefficients of the POD of velocity fields. This approach has shown to be superior to EPOD, achieving better accuracy at larger distances from the wall. Guastoni et al. (2021 [142]) compared the performances of estimators based on a fully convolutional network (FCN) to estimate the velocity fluctuations directly, or to estimate the field through POD modes (FCN-POD), using as input pressure and shear-stress fields at the wall. The FCN

and FCN-POD have shown remarkable accuracy for wall distances up to 50 wall units at a friction-based Reynolds number $Re_\tau = 550$. Recently, Guastoni et al. (2022 [242]) explored this FCN architecture, but using the convective heat flux at the wall, reporting a 50 % error reduction at 30 wall units.

An additional improvement has been achieved by Güemes et al. (2021 [143]) with an algorithm based on generative adversarial networks (GANs; Goodfellow et al. 2014 [202]). This architecture consists of two agents, a generator and a discriminator, which are trained to generate data from a statistical distribution and to discriminate real from generated data, respectively. Generator and discriminator networks compete in a zero-sum game during the training process, i.e. the loss of one agent corresponds to the gain of the other, and vice versa. These GANs have been applied for variegated tasks in fluid mechanics in the last years, including super-resolution (Deng et al. 2019 [157]; Güemes, Sanmiguel Vila & Discetti 2022 [158]; Yu et al. 2022 [148]) and field predictions (Chen et al. 2020 [243]; Li et al. 2023 [244]).

In the work by Güemes et al. (2021 [143]), GANs are used to generate wall-parallel velocity fields from wall measurements—pressure and wall-shear stresses. This architecture has shown better performances than the FCN and FCN-POD architectures proposed earlier (Guastoni et al. 2021 [142]); furthermore, it has shown remarkable robustness in the presence of coarse wall measurements. This aspect is particularly relevant for the practical implementation in experimental and real applications where the spatial resolution of the sensors might be a limitation.

The main shortcoming of the aforementioned studies is that the estimation is carried out with planar data, i.e. the velocity is estimated on wall-parallel planes. An ad hoc network must thus be trained for each wall-normal distance. However, turbulent boundary layers are characterized by the presence of 3-D coherent features (Jiménez 2018 [205]), a fact that was first realized in the visual identification of the near-wall streaks by Kline et al. (1967 [76]). These structures follow a process of lift-up, oscillation and bursting, referred to as the near-wall energy cycle (Hamilton, Kim & Waleffe 1995 [245]), responsible for maintaining turbulence near the wall (Jiménez & Pinelli 1999 [246]). A similar cycle, albeit more complex and chaotic, was later identified in the logarithmic layer (Flores & Jiménez 2010 [247]), involving a streamwise velocity streak with a width proportional to its height that bursts quasi-periodically.

The search for organized motions and coherent structures in wall-bounded turbulent flows has resulted in several families of structures. The definition of many of these structures is based on instantaneous velocity fields, like the hairpin packets of Adrian (2007 [86]), or the

more disorganized clusters of vortices of Del Álamo et al. (2006 [248]). Other structures, like the very large streaks of the logarithmic and outer region, have been described in terms of both two-point statistics (Hoyas & Jiménez 2006 [92]) and instantaneous visualizations (Hutchins & Marusic 2007 [89]). Of particular interest here are the Q-structures defined by Lozano-Durán, Flores & Jiménez (2012 [249]), which are based on a reinterpretation of the quadrant analysis of Willmarth & Lu (1972 [250]) and Lu & Willmarth (1973 [251]) to define the 3-D structures responsible of the turbulent transfer of momentum. These Q-structures are divided into wall-detached and wall-attached Qs events, in a sense similar to the attached eddies of Townsend (1961 [70]). As reported by Lozano-Durán et al. (2012 [249]), the detached Qs are background stress fluctuations, typically small and isotropic, without any net contribution to the mean stress. On the other hand, wall-attached Qs events are larger, and carry most of the mean Reynolds stress. Sweeps (Q4) and ejections (Q2) are the most common wall-attached Qs, appearing side by side in the logarithmic and outer regions.

It is reasonable to hypothesize that the nature of such coherent structures might have a relation with the capability of the GAN to reconstruct them or not. Employing state-of-the-art neural networks, the estimation of a full 3-D field from wall data requires the use of multiple networks targeting the reconstruction of wall-parallel planes at different wall distances. This implies bearing the computational cost of a cumbersome training of several networks, one for each of the desired wall-normal distances. Furthermore, each network is designed to reconstruct features at a certain distance from the wall, ignoring that the wall-shear stresses and pressure distributions depend also on scales located outside the target plane. The two-dimensional (2-D) reconstruction of an essentially 3-D problem complicates the interpretation of the actual scales that can be reconstructed in this process, while the continuity between adjoining layers—in terms of both absence of discontinuities within the field and mass conservation—is not necessarily preserved. Some recent works tackle similar problems, also from a 3-D perspective, such as the reconstruction of an unknown region of the flow through continuous assimilation technique by Wang & Zaki (2022 [146]), the reconstruction of fields from flow measurements by Yousif et al. (2023 [147]), and the reconstruction from surface measurements for free-surface flows through a CNN by Xuan & Shen (2023 [150]).

This work aims to overcome the aforementioned limitations by leveraging for the first time a full 3-D GAN architecture for 3-D velocity estimation from the wall. We employ a dataset of 3-D direct numerical simulations (DNS) of a channel flow. The reconstruction capabilities

Case	y/h range	N_y	N_x	N_z	Δy_{\min}^+	Δy_{\max}^+
A	[0-1]	64	64	64	0.48	6.5
B	[0-0.52]	48	64	64	0.48	5.0
C	[0-0.21]	32	64	64	0.48	2.6
D	[0.21-1]	32	64	64	2.8	6.5

Table 1. Details about the domain of the cases, as represented in figure 2.

of a 3-D GAN are assessed. Section 2 describes both the training dataset and the 3-D GAN networks employed in the present study, while § 3 reports and discusses the results both in terms of reconstruction error statistics and in terms of structure-specific reconstruction quality. The physical interpretation of the results is given in terms of the framework of quadrant analysis in three dimensions (Lozano-Durán et al. 2012 [249]). Finally, § 4 presents the conclusions of the study.

2 Methodology

2.1 Dataset description

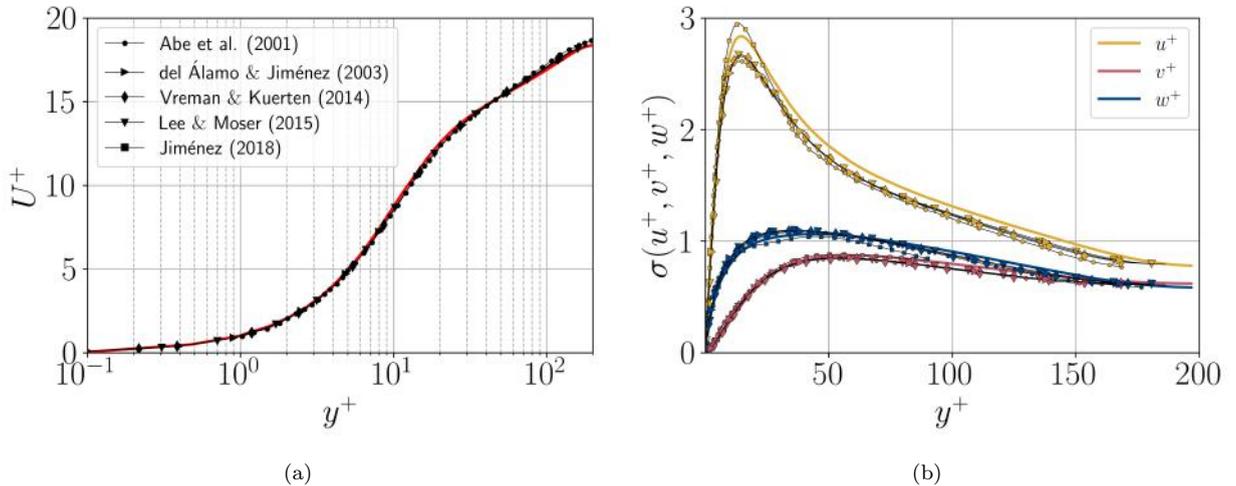


Figure 1. Wall-normal profiles of (a) the mean streamwise velocity and (b) standard deviation σ of the three velocity components. Data are presented in inner units and compared to other databases at a similar $Re_\tau \approx 180$, including a minimal channel unit (Jiménez 2018 [205]) and several bigger channels.

The dataset employed in this work consists of 3-D flow fields and shear and pressure fields at the wall of a minimal-flow-unit channel flow. Our numerical simulations are performed with a state-of-the-art pseudo-spectral code that uses a formulation based on the wall-normal vorticity and the Laplacian of the wall-normal velocity, and a semi-implicit Runge–Kutta for time integration (Vela-Martín et al. 2021 [203]). The solver uses a Fourier discretization

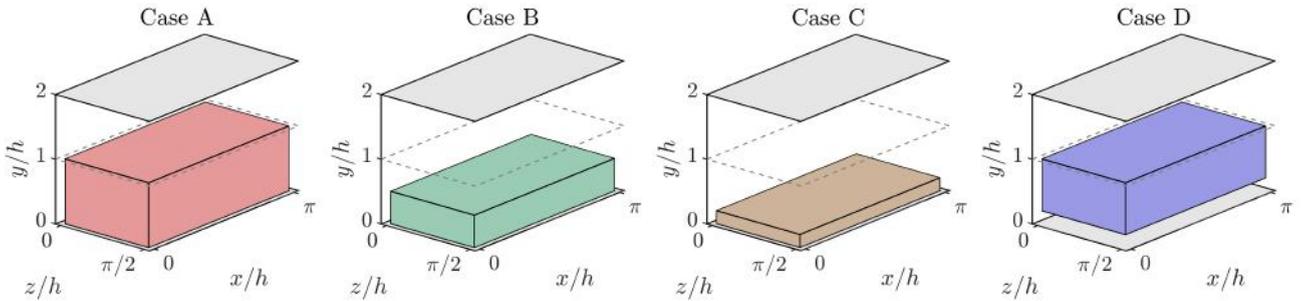


Figure 2. Representation of the reconstructed volume of the channel in each case, as defined in table 1.

in the wall-parallel directions and seventh-order compact finite differences in the wall-normal direction with spectral-like resolution (Lele 1992 [252]). The simulation domain is a periodic channel with two parallel walls located $2h$ apart, with sizes πh and $\pi/2h$ in the streamwise and spanwise directions, respectively. This small channel fulfils the conditions established in the work by Jiménez & Moin (1991 [253]), which defines the minimal channel unit able to sustain turbulence.

In this work, we indicate with x , y , and z the streamwise, wall-normal and spanwise directions, respectively, with their corresponding velocity fluctuations denoted by u , v and w . Simulations are performed at a friction-based Reynolds number $Re_\tau = u_\tau h/\nu \approx 200$, where ν refers to the kinematic viscosity, and $u_\tau = \sqrt{\tau_w/\rho}$ indicates the friction velocity, with τ_w the average wall-shear stress, and ρ the working-fluid density. The superscript $+$ is used to express a quantity in wall units. To ensure statistical convergence and to minimize the correlation between the fields employed, data were sampled every $\Delta t^+ \approx 98$, i.e. 0.5 eddy-turnover times. The mean streamwise profile and the standard deviation of the velocity components are shown in figure 1. Moreover, the mean-squared velocity fluctuations in inner units are plotted in figure 6, where they can be compared with those reported for similar channel flows at $Re_\tau \approx 180$ (Abe, Kawamura & Matsuo 2001 [254]; Del Álamo & Jiménez 2003 [255]; Vreman & Kuerten 2014 [256]; Lee & Moser 2015 [204]).

Both the wall pressure p_w and the wall-shear stress in the streamwise (τ_{wx}) and spanwise (τ_{wz}) directions are used for the flow field estimations. The data are fed into the proposed network on the same grid as the simulation. The streamwise and spanwise directions are discretized

Case	Residual block up-sampling scheme	Trainable parameters	
		G	D
A	0-6-12-18-24-30	9.0×10^6	18.2×10^6
B	0-5-10-15-20	8.0×10^6	18.2×10^6
C	6-12-18-24-30	8.0×10^6	23.8×10^6
D	0-7-14-21-28	8.0×10^6	23.8×10^6
2D-GAN		7.3×10^5	8.0×10^7

Table 2. Details on the implementation of the architectures. The column ‘Residual block up-sampling scheme’ indicates the indexes of the residual blocks that are followed by an up-sampling block. The number of trainable parameters of the generator (G) and discriminator (D) networks are also reported. Cases A–D are compared with the 2-D GAN from Güemes et al. (2021 [143]).

each with 64 equally spaced points, while the volume is discretized with 128 layers with variable spacing from the wall to the mid-plane. This discretization provides a set of grid points with a similar spacing to that found at $Re_\tau = 180$ in the work by Del Álamo & Jiménez (2003 [255]). The estimation capability of the 3-D fields was tested for volumes occupying the whole wall-parallel domain, but with different ranges in the wall-normal direction, giving rise to a set of test cases. These test cases are sketched in figure 2 and summarized in table 1, with N_x , N_y and N_z indicating the size of the mesh along each direction, and Δy_{\min}^+ and Δy_{\max}^+ max being respectively the minimum and maximum wall-normal lengths of each grid step within the domain of each of the cases. Starting from the wall, cases A to C progressively reduce their wall-normal top limit from $y/h < 1$ to $y/h < 0.21$ so as the number of $x - z$ layers. Case D is defined as the domain complementary to that of case C, i.e., covering wall-normal distances in the range $0.21 < y/h < 1$.

2.2 Generative adversarial networks

In this work, a GAN architecture is proposed to estimate 3-D velocity fields from wall measurements of pressure and shear stresses. The implementation details of the proposed architecture are presented below, being an extension to the 3-D space of the network proposed in the work of Güemes et al. (2021 [143]).

A schematic view of the generator network G can be found in figure 3. The network is similar to that proposed by Güemes et al. (2021 [143]), although with some modifications. It is fed with wall measurements and consists of 16 residual blocks, containing convolutional layers with batch normalization layers and parametric-ReLU. The classic ReLU activation function provides as output $f(x) = x$ for positive entries and $f(x) = 0$ (flat) for negative ones. Parametric-ReLU does the same on the positive input values, while for negative entries it is defined as $f(x) = ax$, where a is a parameter (He et al. 2015 [190]). In addition, sub-pixel convolution layers are

used at the end for super-resolution purposes, adding more or fewer layers depending on the resolution of the fed data. To deal with 3-D data, a third spatial dimension has been added to the kernel of the convolutional layers. Since the present dataset does not require the network to increase the resolution from the wall to the flow in the wall-parallel directions, the sub-pixel convolutional layers present after the residual blocks in Güemes et al. (2021 [143]) have been removed. Similarly, the batch-normalization layers were dropped since they were found to substantially increase the computational cost without a direct impact on the accuracy (He et al. 2016 [257]; Kim, Lee & Lee 2016 [258]).

The increase of the wall-normal thickness up to the desired output volume has been achieved by using blocks composed of up-sampling layers followed by convolutional layers with parametric-ReLU activation, which we will refer to as up-sampling blocks throughout this document. For cases A, B and D, the first block is placed before the residual blocks. They increase the size of the domain by a factor of 2 in all cases except for the first up-sampling block in case B, which increases the size of the domain by a factor of 3. The rest of the up-sampling blocks are applied after the residual blocks, whose indexes are specified in table 2, together with the number of trainable parameters of the networks. The number of residual blocks, which has been increased to 32 with respect to Güemes et al. (2021 [143]), and the criterion to decide when to apply the up-sampling blocks, are analysed with a parametric study, for which a summary can be found in Appendix A.

A schematic of the discriminator network D is presented in figure 4. This network is very similar to that proposed by Güemes et al. (2021 [143]). The main difference is the change of the convolutional kernel to the 3-D space, including this new dimension. It consists mainly of a set of convolutional layers that progressively reduce the size of the domain and increase the number of filters. Then, with a flatten layer and two fully-connected layers, the network provides a single output in the range 0–1. Further details can be found in Appendix A. Additionally, it is important to note that due to the wall-normal dimension of its input data, one discriminator block was removed from cases C and D, which led to the counter-effect of increasing the number of trainable parameters reported in table 2. This network makes use of the leaky-ReLU activation function (Maas, Hannun & Ng 2013 [189]; He et al. 2015 [190]).

The training process has been defined for 20 epochs, although the predictions are computed with the epoch where the validation loss stops decreasing, and the optimizer implements the Adam algorithm (Kingma & Ba 2014 [259]) with learning rate 10^{-4} . In total, 24 000 samples

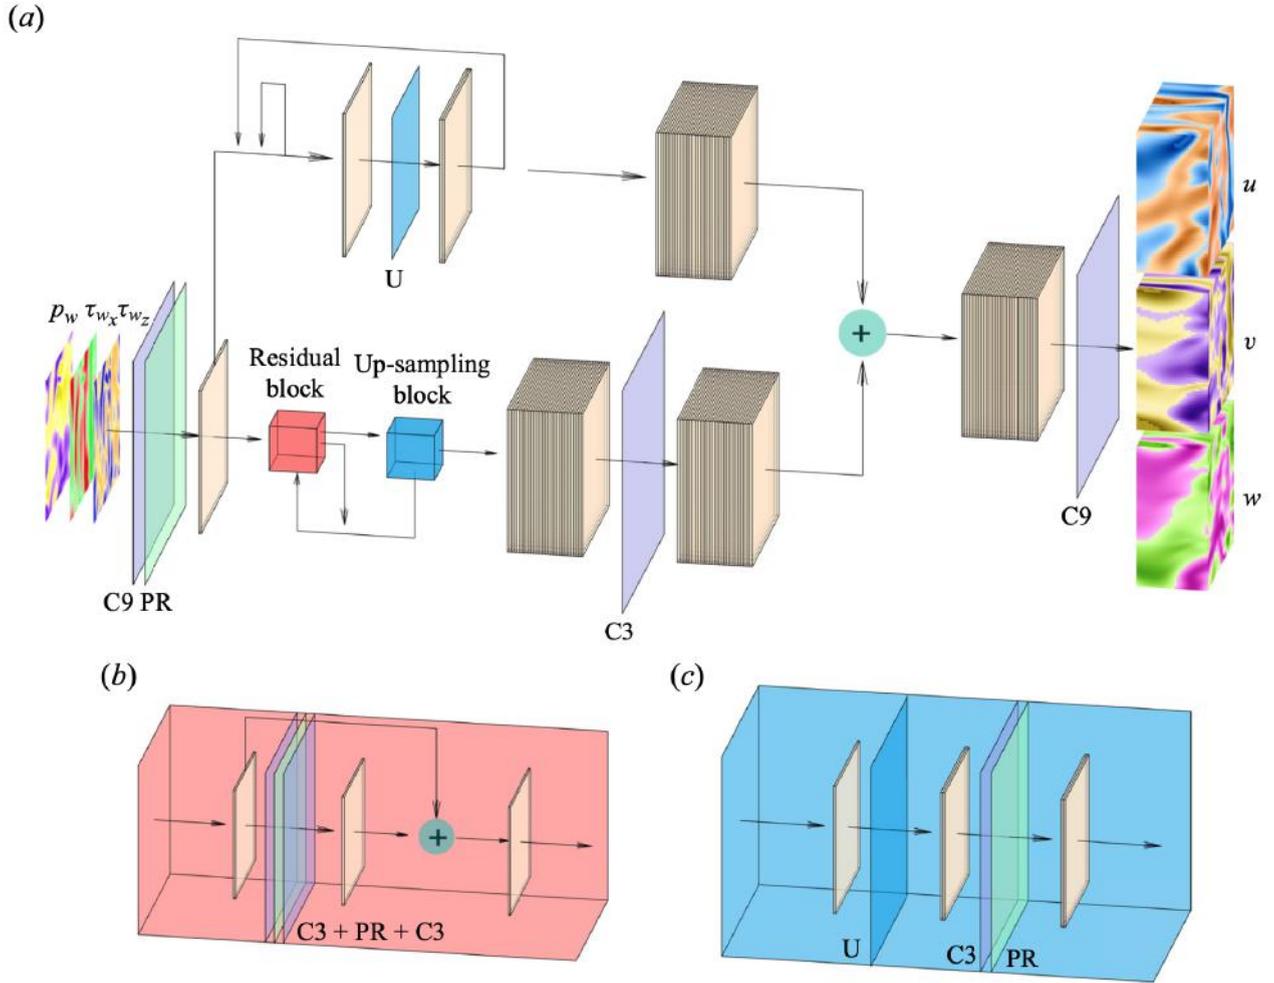


Figure 3. (a) Sketch of the generator network. (b) The residual block and (c) the up-sampling block sub-units, which are repeated recursively through network (a). The filter dimension is represented only in the network input $[p_w, \tau_{w_x}, \tau_{w_z}]$ and output $[u, v, w]$. All other layers work over 64 filters, except the last layer, which only has 3 filters coinciding with the output. The planar panels indicate the different layers of the network: up-sampling (U), parametric-ReLU (PR), and convolution layers with kernel sizes 3 (C3) and 9 (C9), respectively. Arrows indicate the flow of data through layers.

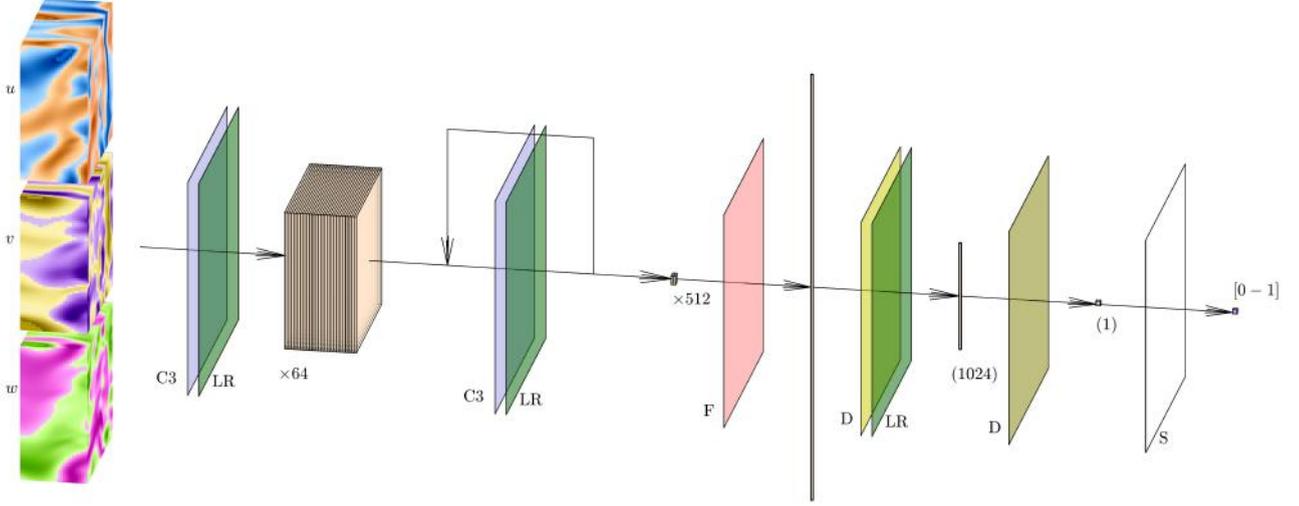


Figure 4. Sketch of the discriminator network. The network receives as input the velocity-fluctuation fields. The planar panels indicate the different layers of the network: the data passes through a set of convolutional (C3) and leaky-ReLU (LR) layers, reducing the dimension of the domain in the x , y and z coordinates as the number of filters increases progressively from 64 to 512. All these data are reshaped into a single vector with the flatten (F) layer. The dimensionality is reduced, first with a D fully connected layer with 1024 elements as output, and then with another D layer providing a single element, which is finally fed to a sigmoid (S) activation function.

have been used, keeping 4000 for validation and 4000 for testing. A random initial condition is set and evolved during about 100 eddy-turnover times to eliminate transient effects. Samples of the testing dataset are captured after approximately 100 eddy-turnover times from the last snapshot of the validation dataset to minimize correlation with the training data.

As mentioned above, the networks operate with the velocity fluctuations $[u, v, w]$. As there are significant differences in the mean values of the velocity components at the centre of the channel and in the vicinity of the wall, the mean values used to compute the field of fluctuation velocities have been obtained for each particular wall-normal distance y . In addition, to facilitate the training of the network, each fluctuating velocity component has been normalized with its standard deviation at each wall-normal layer (see figure 1). Similarly, the wall measurements $[p_w, \tau_{w_x}, \tau_{w_z}]$ provided to the network have been normalized with their mean value and standard deviation.

The training loss functions are defined as follows. The fluctuations of the velocity field can be represented as $\mathbf{u} = [u, v, w]$, such that \mathbf{u}_{DNS} is the original field, and \mathbf{u}_{GAN} is the field reconstructed by the generator network given its corresponding set of inputs $[p_w, \tau_{w_x}, \tau_{w_z}]$. With these two definitions, using the normalized velocity fields, the content loss based on the mean-squared error (MSE) is expressed as

$$\mathcal{L}_{MSE} = \frac{1}{3N_x N_y N_z} \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} \sum_{k=1}^{N_z} [(u_{\text{DNS}}(i, j, k) - u_{\text{GAN}}(i, j, k))^2 + (v_{\text{DNS}}(i, j, k) - v_{\text{GAN}}(i, j, k))^2 + (w_{\text{DNS}}(i, j, k) - w_{\text{GAN}}(i, j, k))^2] , \quad (2.1)$$

Using the binary cross-entropy, an adversarial loss is defined as:

$$\mathcal{L}_{adv} = -\mathbb{E}[\log D(\mathbf{u}_{\text{GAN}})] , \quad (2.2)$$

to quantify the ability of the generator to mislead the discriminator, with \mathbb{E} the mathematical expectation operator, and $D(\cdot)$ the output of the discriminator network when it receives a flow field as input—in this case, a GAN-generated flow field. This adversarial loss is combined with the content loss to establish the loss function of the generator network as

$$\mathcal{L}_G = \mathcal{L}_{MSE} + 10^{-3} \mathcal{L}_{adv} . \quad (2.3)$$

The loss function for the discriminator network, defined as

$$\mathcal{L}_D = -\mathbb{E}[\log D(\mathbf{u}_{\text{DNS}})] - \mathbb{E}[\log(1 - D(\mathbf{u}_{\text{GAN}}))] , \quad (2.4)$$

also uses the binary cross-entropy to represent its ability to label correctly the real and generated fields. To ensure stability during the training process, both the adversarial and discriminator losses are perturbed by subtracting a random noise in the range 0–0.2. This technique, referred to as label smoothing, makes it possible to reduce the vulnerability of the GAN by modifying the ideal targets of the loss functions (Salimans et al. 2016 [260](#)).

3 Results

3.1 Reconstruction accuracy

The reconstruction accuracy is assessed in terms of the MSE of the prediction. The contribution of each velocity component ($\mathbf{u} = [u, v, w]$) to the metric presented in (2.1) has been computed along the wall-normal direction, denoted as \mathcal{L}_u , \mathcal{L}_v and \mathcal{L}_w , respectively. In this case, the error is computed using only one component at a time, and the factor of 3 in the denominator is eliminated. As discussed in § 2.2, the training data have been normalized with their standard deviation for each wall-normal distance. This procedure allows us a straightforward comparison with the results of 2-D GAN architectures (Güemes et al. 2021 [143](#)). It must be remarked

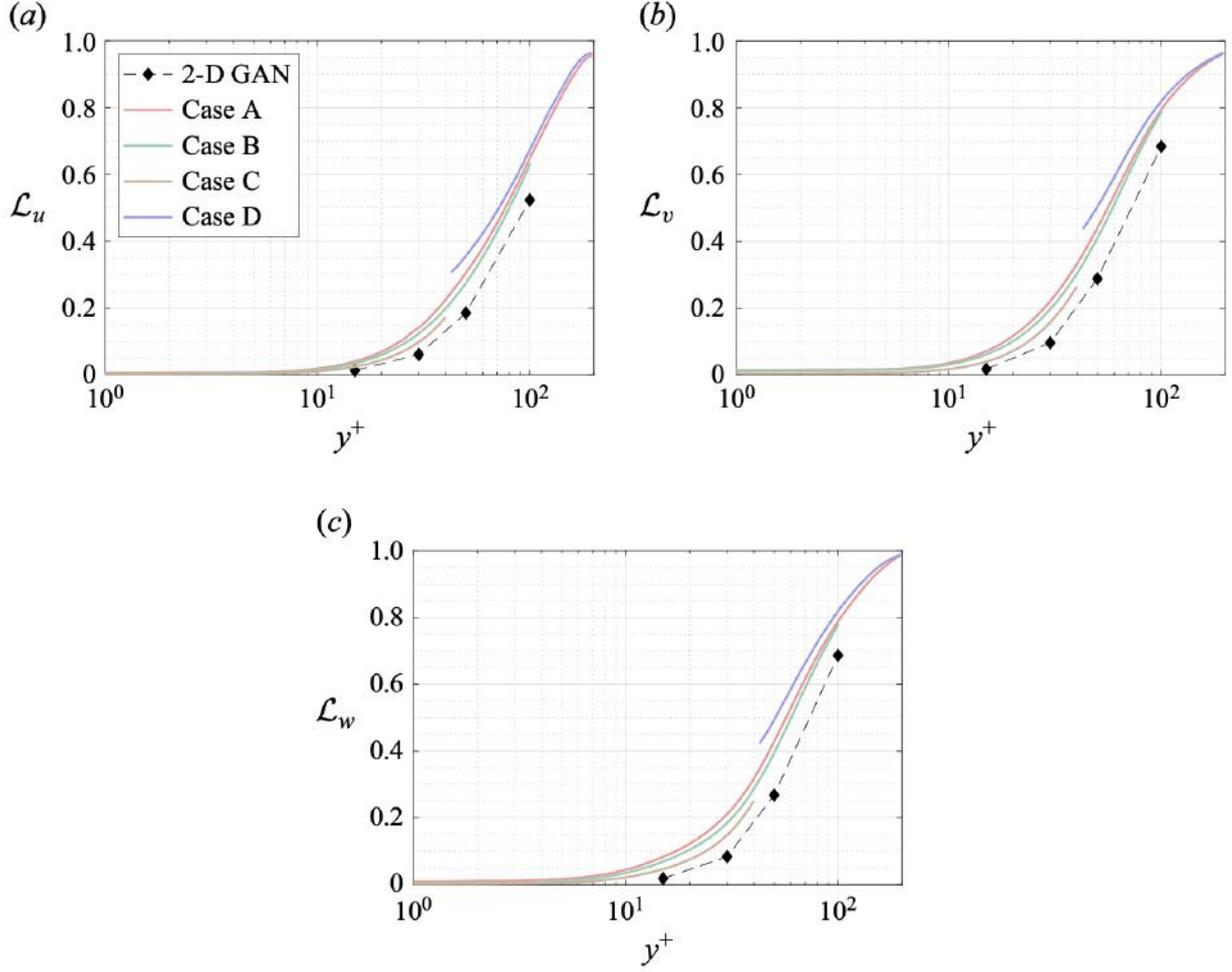


Figure 5. The MSE of the fluctuation velocity components (a) u , (b) v and (c) w for the 3-D GAN (continuous lines) and the 2-D GAN at $Re_\tau = 180$ (symbols with dashed lines) as implemented by Güemes et al. (2021 [143]). Velocity fluctuations are normalized with their standard deviation at each wall-normal coordinate y^+ .

that flow reconstruction by Güemes et al. (2021 [143]) is based on an open-channel simulation; nonetheless, the similar values of Re_τ numbers provide a quite accurate reference.

The MSE for each velocity component is plotted with respect to the wall-normal distance for the selected network architectures of each case A–D in figure 5, and numerical data of the error at the wall distances used in the 2-D approach are collected in table 3 for comparison. They have been computed for the velocity fluctuations normalized with their standard deviation at each wall-normal coordinate y^+ , allowing us to compare the accuracy of this network with the analogous 2-D study. Some general comments can be raised at first sight.

1. As expected, and also reported in the 2-D analysis (Güemes et al. 2021 [143]), the regions closer to the wall show a lower \mathcal{L} . This result is not surprising: at small wall distances the

y^+	Case	\mathcal{L}_u	\mathcal{L}_v	\mathcal{L}_w
15	A	0.043	0.076	0.088
	B	0.038	0.066	0.074
	C	0.027	0.044	0.050
	2D-GAN	0.013	0.018	0.019
30	A	0.137	0.214	0.205
	B	0.118	0.190	0.179
	C	0.095	0.152	0.143
	2D-GAN	0.061	0.097	0.084
50	A	0.306	0.440	0.429
	B	0.277	0.411	0.395
	D	0.356	0.505	0.494
	2D-GAN	0.185	0.289	0.268
100	A	0.639	0.788	0.782
	B	0.619	0.779	0.771
	D	0.665	0.815	0.813
	2D-GAN	0.524	0.684	0.687

Table 3. The MSE for the three velocity components and for each case, at different wall distances. The results are compared with the 2-D analysis by Güemes et al. (2021 [143]). These quantities correspond to velocity fluctuations normalized with their standard deviation at each wall-normal coordinate y^+ .

velocity fields show high correlation with the wall-shear and pressure distributions, thus simplifying the estimation task for the GAN, independently on the architecture.

2. The streamwise velocity fluctuation u always reports a slightly lower \mathcal{L} than v and w for all the tested cases. This is due to the stronger correlation of the streamwise wall-shear stress.
3. The 3-D GAN provides a slightly higher \mathcal{L} than the 2-D case. This was foreseeable: the 3-D architecture proposed here is establishing a mapping to a full 3-D domain, with only a slight increase in the number of parameters in the generator with respect to the 2-D architecture, as can be seen in table 2. Furthermore, there is a considerable reduction in the number of trainable parameters in the discriminator. If we consider $\mathcal{L}_u = 0.2$, then the reconstructed region with an error below this threshold is reduced from approximately 50 to slightly less than 40 wall units when switching from a 2-D to a 3-D GAN architecture.

In test case A, part of the effort in training is directed to estimating structures located far from the wall, thus reducing the accuracy of the estimation. For this reason, cases B and C were proposed to check whether reducing the wall-normal extension of the reconstructed domain would increase the accuracy of the network. Comparing the \mathcal{L} values of cases A, B and C in table 3, it is observed that there is some progressive improvement with these volume reductions, although it is only marginal. For example, the error \mathcal{L}_u at $y^+ = 100$ is 2% lower when switching from case A to case B, i.e. reducing by a factor of 2 the size of the volume

to be estimated. This fact can also be observed in figure 5(a), where the \mathcal{L}_u values for all the cases can be compared directly. Similar conclusions can be drawn from the other velocity components. The improvement between cases is marginal. The quality of the reconstruction of one region seems thus to be minimally affected by the inclusion of other regions within the volume to be estimated. This suggests that the quality of the reconstruction is driven mainly by the existence of a footprint of the flow in a certain region of the channel. In Appendix B, we have included a comparison with LSE, EPOD and a deep neural network that replicates the generator of case A and provides an estimation of the effect of the discriminator. The accuracy improvement of the 3-D GAN with respect to the LSE and the EPOD is substantial, while the effect of the adversarial training does not seem to be very significant. Nevertheless, previous works with 2-D estimations (Güemes et al. 2021 [143]) have shown that the superiority of the adversarial training is more significant if input data with poorer resolution are fed to the network. We hypothesize that a similar scenario might occur also for the 3-D estimation; nonetheless, exploring this aspect falls outside the scope of this work.

Case D, targeting only the outer region, is included to understand the effect of excluding the layers having a higher correlation with wall quantities from the reconstruction process. The main hypothesis is that during training, the filters of the convolutional kernels may focus on filtering small-scale features that populate the near-wall region. Comparing the plots for cases A and D in figure 5, it is found that the \mathcal{L} level in case D is even higher than in case A. While this might be surprising at first glance, a reason for this may reside in the difficulty of establishing the mapping from the large scale in the outer region to the footprint at the wall when such footprint is overwhelmingly populated by the imprint of near-wall small-scale features. Convolution kernels stride all along the domain, and when the flow field contains wall-attached events with a higher correlation, the performance far from the wall seems to be slightly enhanced. In case A, the estimator is able to establish a mapping for such small-scale features to 3-D structures, while for case D, such information, being uncorrelated with the 3-D flow features in the reconstruction target domain, is seen as random noise. This result reveals the importance of the wall footprint of the flow on the reconstruction accuracy.

Moreover, figure 6 shows the evolution of the mean-squared velocity fluctuations and the u^+v^+ shear stress of the reconstructed (case A) velocity fields with the wall-normal distance. As expected, far from the wall, the attenuation becomes more significant, while the accuracy is reasonable up to approximately 30 wall units, where the losses with respect to the DNS are

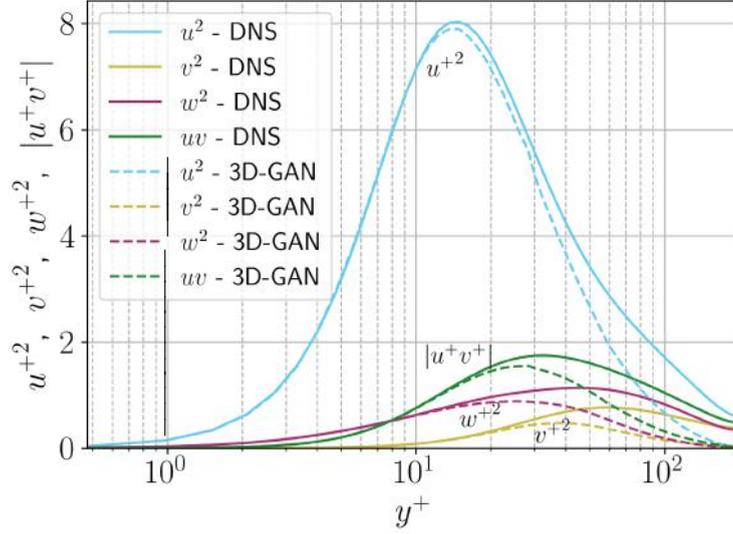


Figure 6. Mean-squared velocity fluctuations and shear-stress. Given in wall-inner units.

equal to 4.0% for u^2 , 9.7% for v^2 , 12.8% for w^2 and 6.9% for $|uv|$. It is important to remark that the network is not trained to reproduce these quantities, as the loss function is based on the MSE and the adversarial loss. The losses reported in these quantities also may explain the MSE trends in figure 5, where the error grows with the wall-normal distance as the network generates more attenuated velocity fields. The fact that the kernels in the convolutional layers progressively stride along the domain implies that although the continuity equation might not be imposed as a penalty to the training of the network, the 3-D methodology exhibits an advantage with respect to the 2-D estimation. To assess this point, we compared the divergence of the flow fields obtained with the 3-D GAN with the divergence obtained from the velocity derivatives of three neighbouring planes estimated with 2-D GANs following Güemes et al. (2021 [143]). The standard deviation of the divergence, computed at both $y^+ = 20$ and $y^+ = 70$, is approximately 6 times smaller when employing the 3-D GAN.

An additional assessment of the results is made by comparing the instantaneous flow fields obtained from the predictions with those from the DNS. As an example, figure 7 shows 2-D planes of \mathbf{u} at three different wall-normal distances, of an individual snapshot, according to case C. This case is selected for this example as it exhibits the best performance. In this test case, the attenuation of the velocity fluctuations is not significant. Up to the distances contained within case C, it is indeed possible to establish accurate correlations. On the contrary, the attenuation of the velocity field close to the centre of the channel is quite high (see figure 6). At $y^+ = 10$, it is difficult to find significant differences between the original and the reconstructed

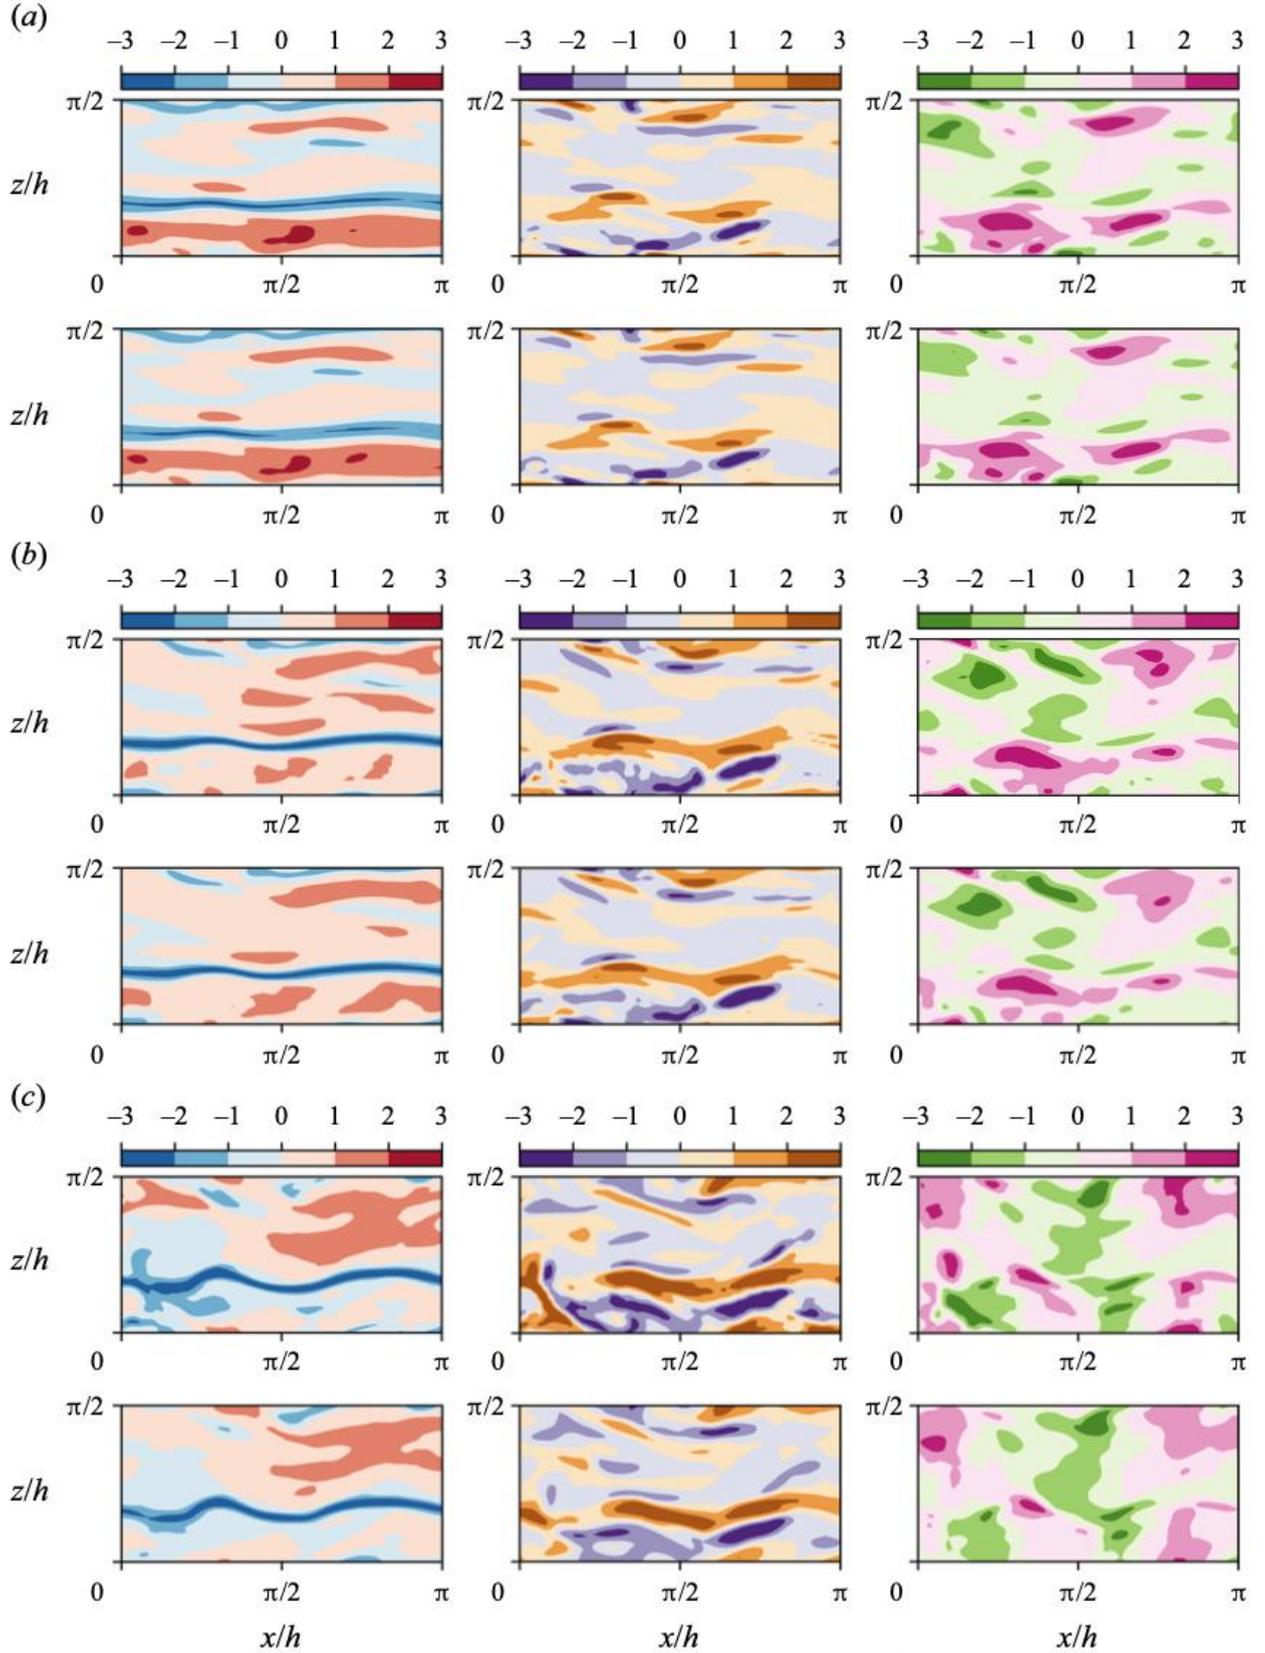


Figure 7. Instantaneous velocity fluctuations: u (left), v (centre) and w (right). From each pair of rows, the top row is the original field from the DNS, and the bottom row is the field reconstructed with the GAN, for case C. Different pairs of rows represent 2-D planes at different wall-normal distances, with (a) $y^+ = 10$, (b) $y^+ = 20$ and (c) $y^+ = 40$. Instantaneous values beyond $\pm 3\sigma$ are saturated for flow visualization purposes. Velocity fluctuations are normalized with their standard deviation at each wall-normal coordinate y^+ .

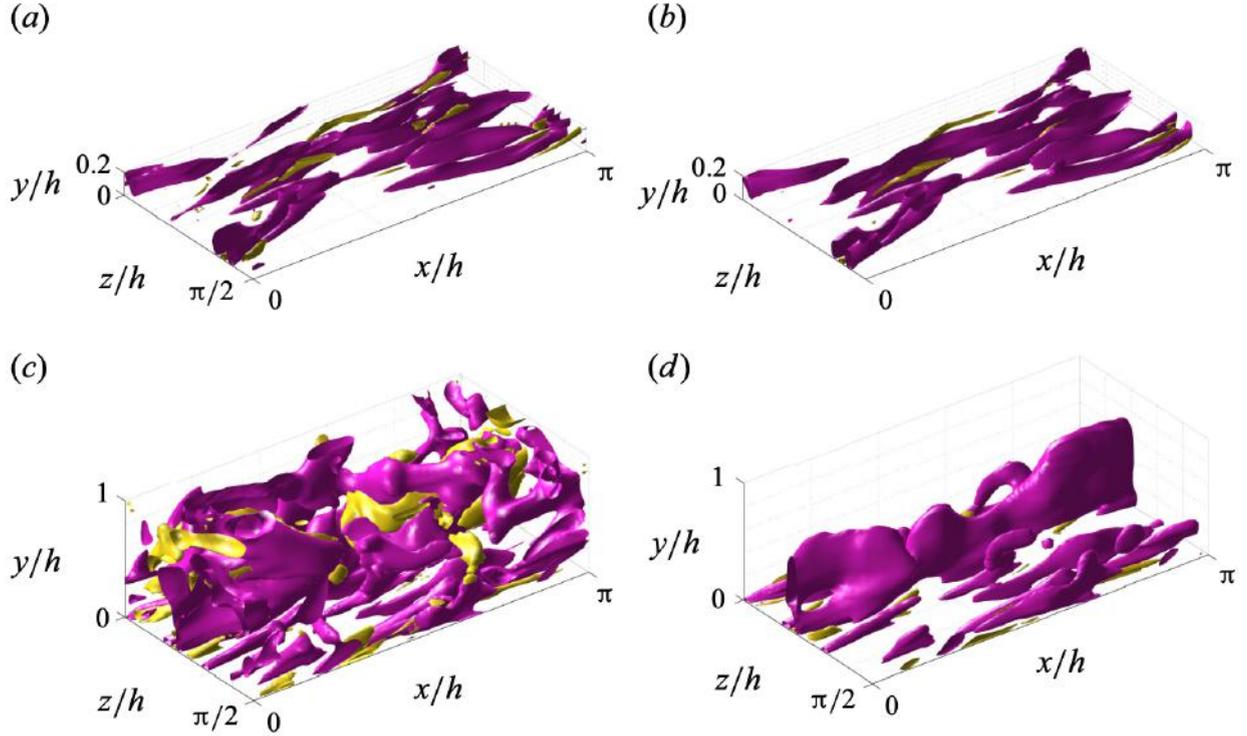


Figure 8. Instantaneous three-dimensional representation of the uv field, for (a) original and (b) prediction from case C and (c) original and (d) prediction from case A. Isosurfaces correspond to the 1.5 and -1.5 levels of uv respectively, in yellow for quadrants Q1 and Q3, and in pink for Q2 and Q4.

fields, with the smallest details of these patterns also being present. Farther away from the wall, at $y^+ = 20$, the estimation of the network is still very good, although small differences start to arise. At $y^+ = 40$, the large-scale turbulent patterns are well preserved, but the small ones are filtered or strongly attenuated.

In general, it can be observed that, regardless of the wall-normal location, the GAN estimator is able to represent well structures elongated in the streamwise direction (i.e. near-wall streaks), likely due to their stronger imprint at the wall. On the other hand, the u fields at $y^+ = 40$ are also populated by smaller structures that do not seem to extend to planes at smaller wall-normal distances, thus indicating that these structures are detached. From a qualitative inspection, the detached structures suffer stronger filtering in the reconstruction process. Analogous considerations can be drawn from observation of the v and w components.

3.2 Coherent structure reconstruction procedure

Further insight into the relation between reconstruction accuracy and features of the coherent structures is provided by observing isosurfaces of the product uv (the so-called uv sters; Lozano-

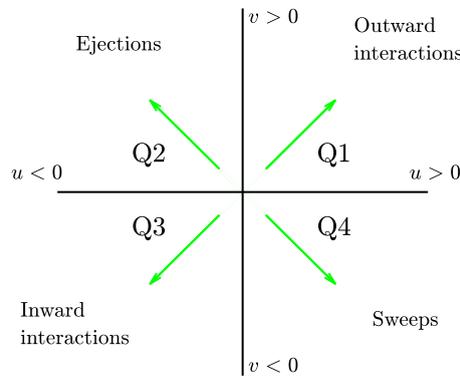


Figure 9. Quadrant map with the categorization of the turbulent motions as Qs events.

Durán et al. 2012 (249) reported in figure 8. Again, a sample of case C has been selected to compare the original structures (figure 8a) with the reconstructed ones (figure 8b). In both representations, structures of different sizes are observed, mainly aligned with the flow. The larger structures appear to be qualitatively well represented, while some of the smaller ones are filtered out or not well reproduced. Furthermore, it can be observed that the structures located farther from the wall are more intensively attenuated in the reconstruction process. Moreover, the majority of the structures and the volume identified within a structure are either sweeps or ejections.

The same procedure is followed for an instantaneous field reconstructed under case A, leading to figures 8(c) and 8(d), respectively. Similar phenomena are observed with some remarks. With a substantially larger volume, many more structures populate the original field. However, the reconstructed structures mainly appear close to the wall. Detached structures are filtered, while attached ones can be partially truncated in their regions farther away from the wall. Nevertheless, this instantaneous field also reveals that the filtering does not seem to be a simple function of the wall distance. Some of the attached structures recovered in the reconstructed field extend up far from the wall, and even up to the middle of the channel or beyond. This would not be possible if all types of structures were expected to be reconstructed up to a similar extent at any wall distance. The region far from the wall is depopulated using the same threshold for the isosurfaces, as the magnitude of the fluctuation velocities is strongly attenuated in this region.

For a quantitative assessment of the relation between coherent-structure features and reconstruction accuracy, here we follow an approach similar to that used by Lozano-Durán et al. (2012 (249)), based on a 3-D extension of the quadrant analysis (Willmarth & Lu 1972

[250]; Lu & Willmarth 1973 [251]), where turbulent structures are classified according to the quadrants defined in figure 9. A binary matrix on the same grid of the domain is built. The matrix contains ones in those points corresponding to a spatial position located inside a coherent structure, and zeros otherwise. Grid points where fluctuation velocities meet the following condition are within a structure

$$|\tau(x, y, z)| > H u'(y) v'(y) , \quad (3.1)$$

where $\tau(x, y, z) = -u(x, y, z)v(x, y, z)$, the prime superscripts (') indicate root-mean-squared quantities, and H is the hyperbolic hole size, selected to be equal to 1.75. This is the same structure-identification threshold as in Lozano-Durán et al. (2012 [249]), for which sweeps and ejections were reported to fill only 8 % of the volume of their channel, although these structures contained around 60 % of the total Reynolds stress at all wall-normal distances. Without any sign criterion, this condition is used for the identification of all types of structures. Moreover, the signs of $u(y)$ and $v(y)$ are to be considered to make a quantitative distinction between sweeps (Q4) and ejections (Q2) as essential multi-scale objects of the turbulent cascade model that produce turbulent energy and transfer momentum.

The cells activated by (3.1) are gathered into structures through a connectivity procedure. Two cells are considered to be within the same structure if they share a face, a side or a vertex (26 orthogonal neighbours), or if they are indirectly connected to other cells. Some of the structures are fragmented by the sides of the periodic domain. To account for this issue, a replica of the domain based on periodicity has been enforced on all sides. To avoid repetitions, only those structures whose centroid remains within the original domain are considered for the statistics. Besides, structures with volume smaller than $10^{-5}h^3$ have been removed from our collection, to concentrate the statistical analysis on structures with a significant volume. Further conditions have been set to remove other small structures that are not necessarily included in the previous condition. Structures that are so small that they occupy only one cell—regardless of their position along y and the cell size Δy^+ —and those whose centroid falls within the first wall-normal cell have been removed. Finally, structures that are contained within a bounding box of a size of one cell in any of the directions have been removed as well. For example, a structure comprising two adjoint cells delimited within a bounding box with two-cell size in one of the directions, but only one-cell size in the other two directions, would be discarded for the statistics. These restrictions still keep small-scale structures in the database, but eliminate those that are contaminated by the resolution errors of the simulation.

		Volume of structures		Number of structures	
		Absolute [h^3]	Relative %	Total	Attached %
Qs	Target	0.70	14.3%	14.5	68.9%
	Prediction	0.80	15.9%	13.3	77.4%
Q1s	Target	0.01	0.3%	3.85	90.3%
	Prediction	0.01	0.3%	3.80	90.7%
Q2s	Target	0.46	9.3%	5.24	81.0%
	Prediction	0.58	11.8%	4.59	89.4%
Q3s	Target	0.04	0.8%	3.10	68.4%
	Prediction	0.01	0.2%	3.12	69.5%
Q4s	Target	0.10	2.0%	7.41	81.5%
	Prediction	0.16	3.2%	7.11	85.2%

Table 4. Information about the structures identified with (3.1), (Qs) without any sign criterion on u and v and (Q1s-Q4s) for each quadrant. The DNS original data and the 3D-GAN prediction (case A) are compared with some statistics over the 4 000 testing snapshots, considering the average volume occupied by structures per snapshot, their proportion of volume over the domain, the average number of structures in each snapshot and the proportion of structures that are attached.

A statistical analysis on test case A has been carried out considering the different quadrants (see table 4). The figures reported by Lozano-Durán et al. (2012 [249]) can be used as a reference, although it must be remarked that discrepancies arise due to the differences in the fluid properties, the Reynolds number or the extension of the volume from the wall considered. The volume of a whole domain of case A is $4.93h^3$. The structures identified with (3.1) and the hyperbolic hole size used as threshold occupy only a small fraction of it, although they contain the most energetic part, able to develop and sustain turbulence. Note that the criterion established with this equation makes Qs (first row in table 4) not to be explicitly the sum of all individual Q1s, Q2s, Q3s, Q4s events—when no sign criterion is being applied over u' and v' , without any distinction among different type of events. As seen with the instantaneous snapshots in figure 8, these statistics from 4000 samples tell us that Q1s and Q3s (see figure 9) are less numerous than negative Qs, both in volume and in units – these structures account for just 2 % in volume and 7 % in units in the work by Lozano-Durán et al. (2012 [249]). In addition, this analysis reveals that most of the volume fulfilling (3.1) belongs to Q2 structures, with Q4s occupying significantly less volume, while in unit terms the population of Q4s is approximately 50 % higher than that of Q2s. Individual Q2s, although fewer in number, are much bigger than Q4s. Structures have been considered as attached if $y_{min}/h \leq 0.1$ (figures 10a,d,g), where y_{min} refers to the location of the closest point to the wall within a structure, and y_{max} to the farthest one. All types of structures are notably attached in more than 60% of the cases, with Q3s the most prone structures to be detached, and Q1s to be attached. Table 4 also offers a comparison between the target data from the DNS and the reconstruction from the 3-D

GAN. There are no large discrepancies between target and prediction, with all the comments already mentioned applying to both of them. However, statistics are better preserved for Q1s and Q3s than for Q2s and Q4s. As expected, the GAN tends to generate slightly fewer Q2 and Q4 structures, a difference due mainly to wall-detached structures that are not predicted. However, these generated structures are bigger and occupy a larger volume than the original ones. As discussed below, not all the volume in the predicted structures is contained within the original ones.

3.3 Statistical analysis methodology

A statistical analysis of the reconstruction fidelity of flow structures is carried out, with the previous condition (3.1) applied to the 4000 samples outside the training set. The structures found in the DNS- and GAN-generated domain pairs have been compared and matched. For each i th (or j th) structure, it is possible to compute its true volume $v_{T,i}$ (or predicted volume $v_{P,i}$) as

$$v_{T_i} = \sum_{x,y,z} \mathbf{T}_i \circ \mathbf{V} , \quad (3.2)$$

$$v_{P_j} = \sum_{x,y,z} \mathbf{P}_j \circ \mathbf{V} , \quad (3.3)$$

where the matrices \mathbf{T}_i and \mathbf{P}_i represent the target (DNS) and the prediction (GAN) domains for the i th structure, respectively, containing ones where the structure is present, and zeros elsewhere. Here, \mathbf{V} is a matrix of the same dimensions containing the volume assigned to each cell. In a similar way, combining these two previous expressions, the overlap volume of two structures i and j within the true and predicted fields, respectively, is

$$v_{T_i,P_j} = \sum_{x,y,z} \mathbf{T}_i \circ \mathbf{P}_j \circ \mathbf{V} . \quad (3.4)$$

It must be considered that during the reconstruction process, the connectivity of regions is not necessarily preserved. This gives rise to a portfolio of possible scenarios. For instance, an original structure could be split into two or more structures in the reconstruction; small structures, on the other hand, could be merged in the estimated flow fields. Moreover, the threshold in (3.1) is based on the reconstructed velocity fluctuations, thus it can be lower than in the original fields. Hence all possible contributions from different structures overlapping with

a single structure from the other dataset are gathered as follows:

$$\widehat{v}_{T_i,P} = \frac{\sum_j v_{T_i,P_j}}{v_{T_i}} , \quad (3.5)$$

$$\widehat{v}_{T,P_j} = \frac{\sum_i v_{T_i,P_j}}{v_{P_j}} . \quad (3.6)$$

With the hat marked used to indicate the ratio, these metrics give the overlapped volume proportion of each structure i from the target set or j from the prediction set and are defined in such a way as some structures either split or coalesce. These structures are classified into intervals according to their domain in the y direction, bounded by y_{min} and y_{max} . Given the matching proportion of all the structures of each target-DNS and prediction-GAN set falling in each interval (a, b) of y_{min} and each interval (c, d) of y_{max} , according to their individual bounds (respectively $y_{min,i}$ and $y_{max,i}$, or $y_{min,j}$ and $y_{max,j}$), their average matching proportions X_t and X_p are computed:

$$X_{t,(a,b - c,d)} = \overline{\widehat{v}_{T_i,P}} \quad \forall \quad i \text{ such that } a < y_{min,i}/h < b, c < y_{max,i}/h < d , \quad (3.7)$$

$$X_{p,(a,b - c,d)} = \overline{\widehat{v}_{T,P_j}} \quad \forall \quad j \text{ such that } a < y_{min,j}/h < b, c < y_{max,j}/h < d . \quad (3.8)$$

Additionally, out of all these categories onto which the structures are classified according to their minimum and maximum heights, those contained within the top 95 % of the joint probability density function (joint p.d.f.) have been identified with black dots in figures 10 and 11 to characterize the predominant structures in the flow.

3.4 Analysis of the joint probability density functions of reconstructed structures

The interpretation of the quantities defined in the previous subsection is as follows: X_t is the proportion of the volume of the structures from the target set represented within the reconstructed structures; X_p is the proportion of the volume of reconstructed structures matching the original ones. These quantities X_t and X_p are represented in figure 10 for each categorized bin and for cases A, B and C.

In figures 10(a,b,c) (for X_t), the joint p.d.f. is compiled for the target structures, and in figures 10(d,e,f) (for X_p), the joint p.d.f. is compiled for the predicted structures. The joint p.d.f. for the target structures indicates that the family of wall-attached structures (i.e. $y_{min} \leq 0.1h$) dominates the population, while wall-detached structures with $y_{min} \geq 0.1h$ do not extend far from the wall. Overall, the joint p.d.f.s are qualitatively similar to those obtained by Lozano-Durán et al. (2012 [249](#)) at higher Reynolds numbers, except for the wall-detached

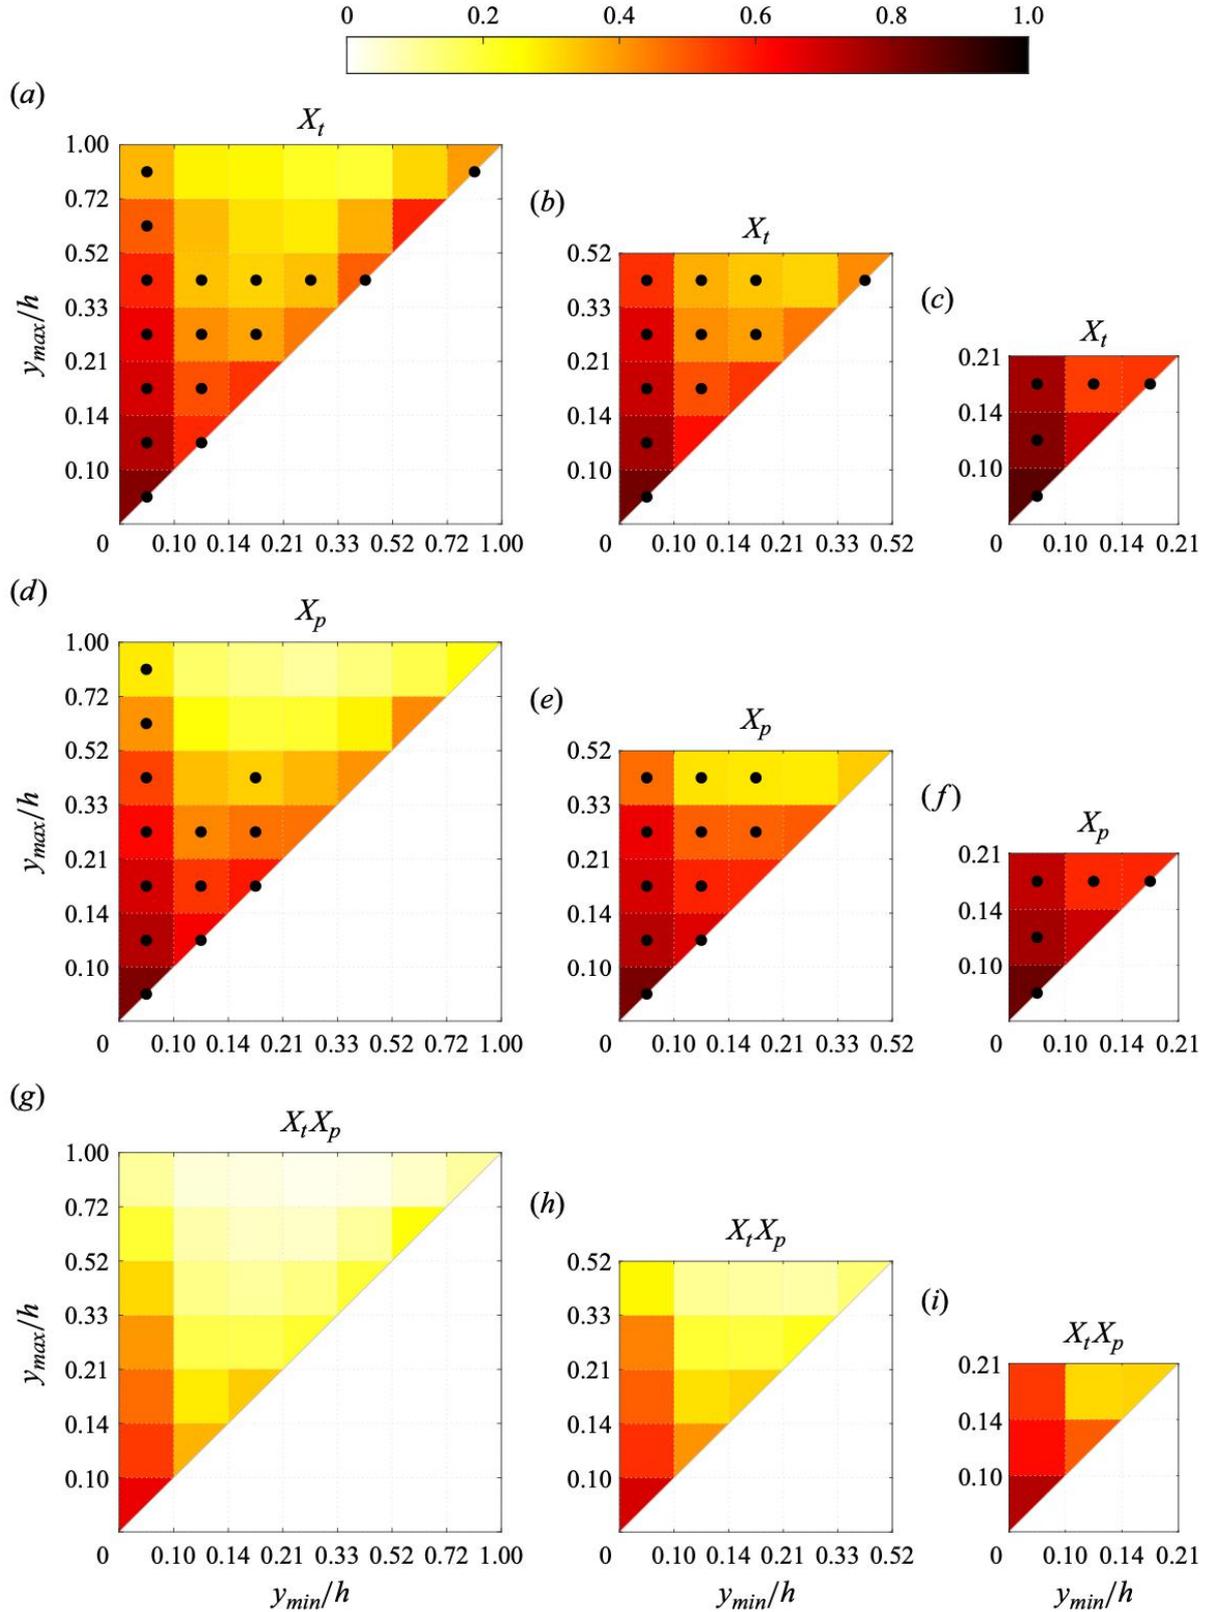


Figure 10. Maps of average density per bin of the matching quantities for cases (a,d,g) A, (b,e,h) B and (c,f,i) C, for the different metrics proposed. Dotted bins represent the top 95 % of the joint p.d.f. of the structures over that target or prediction set, respectively. Note that the scales in both axes are not uniform. Maps of average density per bin of the matching quantities for cases A (left column), B (centre column) and C (right column) for the different metrics proposed. Dotted bins represent the top 95% of the j.p.d.f. of the structures over that target or prediction set respectively. Note that the scales in both axes are not uniform.

structures, which in the latter have $y_{max} - y_{min}$ approximately independent of y_{min} . This difference could be explained as an effect of the low friction Reynolds number, especially since the detached structures in Lozano-Durán et al. (2012 [249](#)) are linked to the dissipation process in the logarithmic and outer regions. Note that the bin highlighted at the top-right corner for the original set of structures is likely linked not to detached structures, but to tall attached structures rising from the opposite wall and extending beyond the middle of the channel.

Compared to the joint p.d.f. of the predicted set, the wall-detached structures of the original set extend farther from the wall. This suggests that the 3-D GAN may be losing the farthest region from the wall of some of the reconstructed structures, consistent with the flow visualizations presented in figure 8. The best reconstruction is reported in all cases for the shortest wall-attached structures. The values of X_t and X_p for wall-attached structures reduce progressively as y_{max}/h increases. This trend is repeated for other columns of bins with $y_{min}/h \geq 0.10$, although the metrics are lower than for the wall-attached structures. The wall-detached structures that are contained within the top 95 % of the joint p.d.f. are reconstructed with modest values of X_t and X_p , approximately 0.5. The structures with poorer matching between the target and predicted sets (i.e. X_t and X_p smaller than 0.25) are relatively far from the wall, and do not belong to the 95 % of the joint p.d.f.s – suggesting that there are very few of them.

Several reasons may justify this performance: we expect a lower prediction ability from the 3-D GAN for wall-detached structures, for structures extending to higher y_{max}/h and for types of structures that are not particularly common, with the computational resources available in the training process having been used to target other patterns within the flow.

The joint p.d.f.s for figure 10(*b,e,h*) and figure 10(*c,f,i*), which do not consider those regions in case A with few structures and poorer reconstructions, show smaller differences between the original and predicted sets of structures. In these cases, the top rows of bins in figure 10(*b,c,e,f,h,i*) are expected to include structures extending beyond its y_{max}/h limit, cutting them and considering their respective reconstruction accuracy only up to its respective limit. As was also observed in figure 5 with the error trends, these metrics X_t and X_p also indicate a slightly improved prediction ability with the reductions in the volume of the domain considered.

With these distributions of average matching proportions X_t and X_p in each interval, a novel perspective on what the network is capable of reconstructing is given. The reconstructed volume of wall-attached structures is generally preserved (high X_t) and undistorted (high X_p) for wall-attached coherent structures. Even though the reconstruction precision in terms of volume

and shape of the coherent structure seems to reduce progressively for increasing y_{max}/h , the values of X_t and X_p for wall-attached structures are still higher than for any other bin with $y_{min}/h > 0.1$ if $y_{max}/h < 0.5$. It can be argued that the reduction in reconstruction accuracy for increasing wall-normal distance is due prevalently to the progressively decreasing number of wall-attached structures, which should have an impact on the training of the 3-D GAN.

The increase in reconstruction fidelity when reducing the wall-normal thickness of the volume of the domain (i.e. cases B and C in figure 10) is in line with the hypothesis that the estimator focuses its effort in reconstructing features extending down to the wall. A marginal increase in X_t is observed in bins corresponding to the same region for decreasing wall-normal thickness of the reconstructed volume. It can be hypothesized that the prevalence of wall-attached over detached—thus poorly correlated with flow quantities—structures in cases B and C simplifies the training of the network and improves its accuracy. Structures extending beyond the y_{max} limit of each case are collected within the top row of bins of each plot in figure 10, and their reconstruction accuracy is slightly increased when the volume of the domain is reduced y_{max} although they are cut and a part of them is not being considered.

The X_t and X_p distributions share some similarities, with the ideas mentioned above. The main difference between them is the fact that the average matching proportions are slightly higher for X_t than for X_p . From X_t , it is seen that with the behaviour just mentioned, the structures predicted by the network do not contain the whole volume of the original ones, denoting some loss of accuracy. Moreover, X_p tells us that the reconstructed structures contain not only sections within the original structures but also regions out of them. With both metrics and their physical meaning, the combined effect of these two losses together is shown as $X_t X_p$ in figure 10. The superior reconstruction of these wall-attached structures must be highlighted, with a progressive loss with the wall-normal size Δy . The small structures lying right over the diagonal report a lower overall score than the attached ones, but higher than other wall-detached structures.

According to (3.1), turbulent structures are defined independently of the sign of $u(y)$ and $v(y)$, but imposing signs on them, the structures can be classified following the quadrant analysis (Lozano-Durán et al. 2012 [249]), with sweeps and ejections being of special interest. The maps shown in figure 11 allow us to compare and establish further conclusions and differences in the performance of the 3-D GAN when reconstructing sweeps and ejections. In all the cases, the left column of wall-attached structures is fully contained within the top 95 % of the joint p.d.f.,

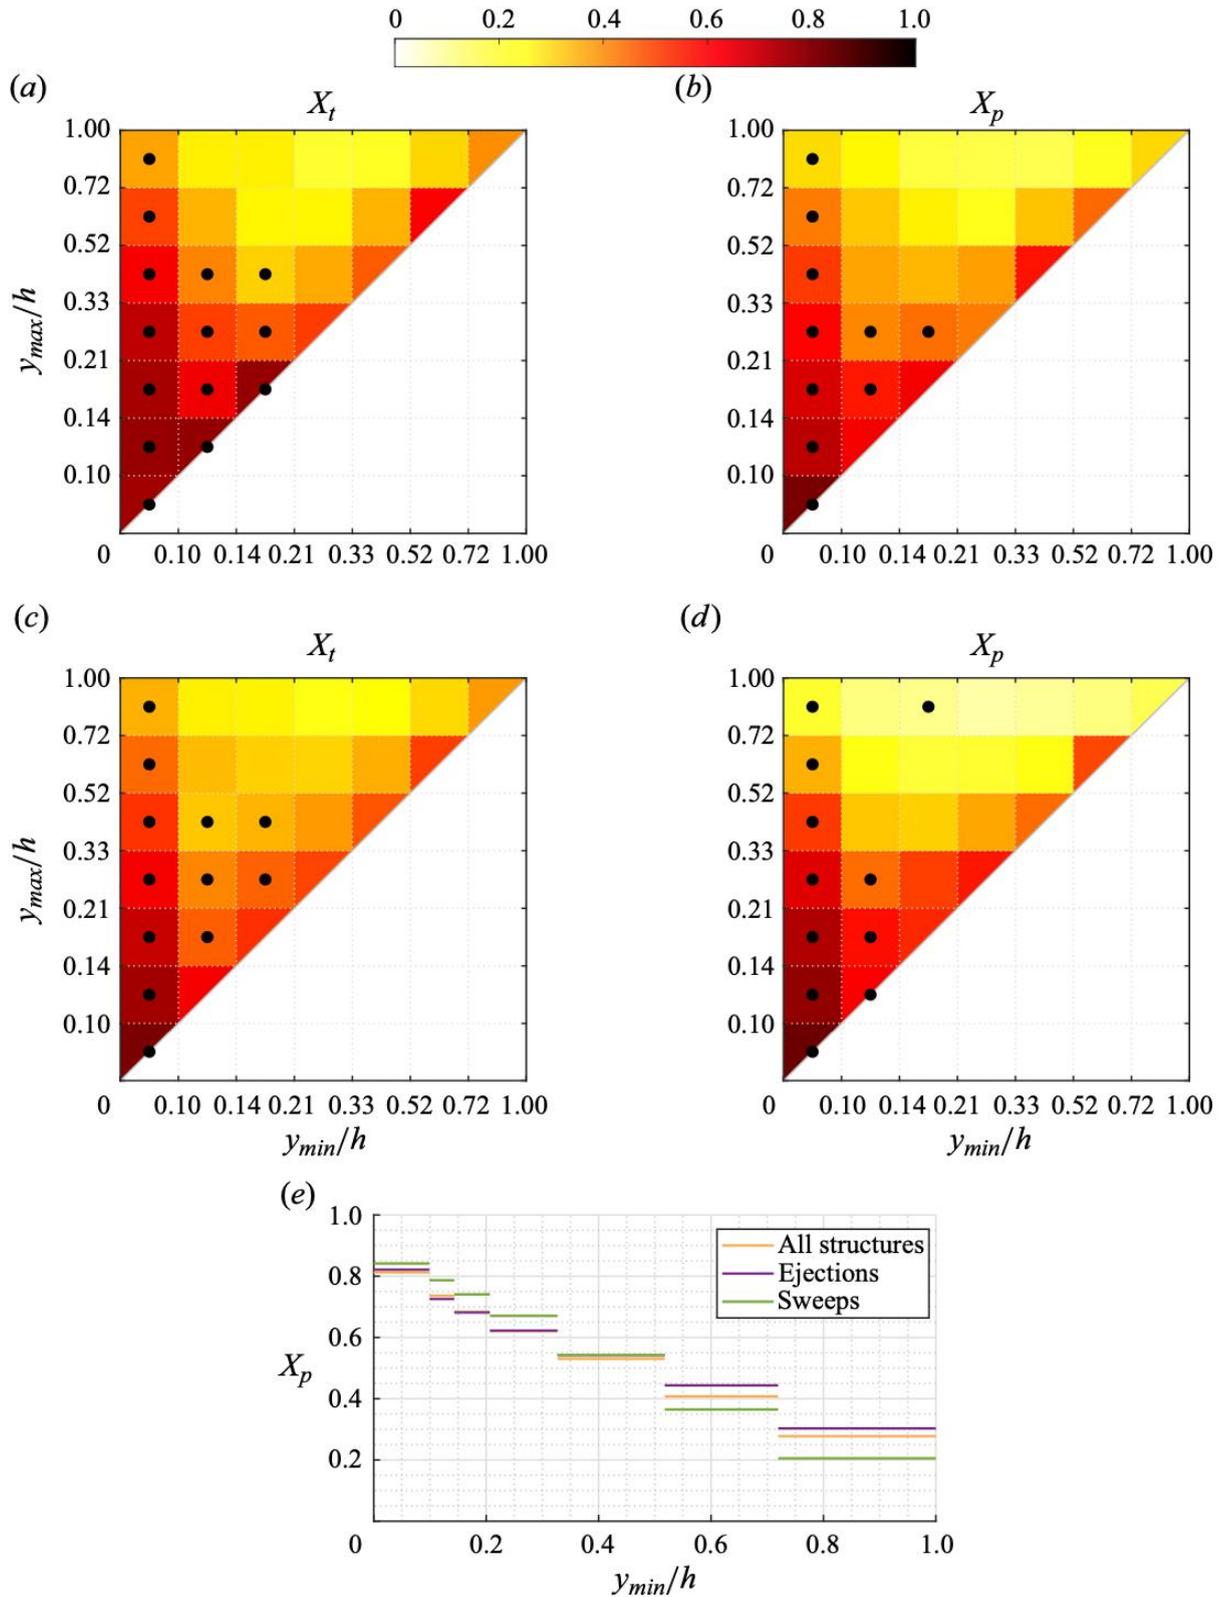


Figure 11. Maps of X_t and X_p for case A, considering only ejections-Q2 in (a,b) and only sweeps-Q4 in (c,d), analogous to those in figures 10(a) and 10(d) containing all structures together. Dotted bins represent the top 95 % of the joint p.d.f. of the structures over that target or prediction set, respectively. Profile (e) of average matching proportion X_p for case A, for the left column of bins ($y_{max}/h < 0.10$), comparing cases considering all the identified structures (figure 10d), only ejections (figure 11b) and only sweeps (figure 11d).

with their X_t or X_p magnitude decreasing progressively with increasing y_{max}/h , as expected. In none of them is the family of wall-detached structures with small Δy lying right above the diagonal a significant proportion of the population of structures. Hence these numerous structures highlighted in figure 10 may be mainly Q1 and Q3 structures. Moreover, at this threshold of the joint p.d.f., other structures that may be considered as attached although they are not in the leftmost column, are included in this top joint p.d.f. set. In terms of y_{min}/h , they extend from below 0.21 following this categorization in bins, while a very similar limit of 0.20 was set as the threshold to classify structures as attached or detached by (Lozano-Durán et al. 2012 [249]). However, in terms of y_{max}/h , the joint p.d.f. cut does not extend up to 1—which is the case for $y_{min}/h < 0.1$ bins. For both sweeps and ejections, these numerous structures do not extend beyond the 0.52 limit for X_t or beyond the 0.33 limit for X_p . This indicates that the network can reconstruct wall-attached sweeps and ejections with reasonable fidelity, although the part farther from the wall may be partially lost or even generated with poorer accuracy.

One of the main aspects reflected in figure 10 is the fact that more accurate reconstructions are reported for wall-attached structures, which are the most quantitative ones. The same is true for the distributions with sweeps and ejections as seen in figures 11(a)–11(d). For some of the bins right above the diagonal with wall-detached structures with small Δy , the quantities X_t and X_p reported are higher than for those bins of wall-attached structures with high Δy . Nevertheless, these wall-detached structures, as ejections (figures 11a,b) or sweeps (figures 11c,d), are not a big part of the population of structures, such that the training process may not focus substantially on them.

The distributions of X_t and X_p shown in figure 11 for ejections and sweeps are qualitatively similar, although some differences can be established. Considering X_t for those bins within the region highlighted by the joint p.d.f., the metric is always higher for ejections than for sweeps, although they follow the same trends. The 3-D GAN can reconstruct wall-attached structures, generally and statistically preserving ejections slightly better than sweeps. One possible reason would be that the correlation with the wall measurements used is stronger with ejections, emerging from the wall, than with sweeps, which travel towards the wall. This could not be explained with the pressure measurements, which may be quite antisymmetric for both types of structures, according to studies assessing this correlation, such as the one by Sanmiguel Vila & Flores (2018 [261]). Nevertheless, this is not the case for the wall-shear stresses, also used as input data.

Regarding X_p , it is difficult to establish such a distinction between Q2 and Q4 structures. To compare these variations in X_p along the structures with $y_{min}/h < 0.1$ more easily, figure 11(e) provides a different view. Wall-attached sweeps are less distorted than ejections and all structures overall when structures with short wall-normal heights are considered (up to $y_{max}/h < 0.3$). This trend is inverted for taller wall-attached sweeps ($y_{max}/h > 0.5$), reporting substantially lower X_p , with these sweeps being more distorted than ejections with similar y size. Although the quantities are very similar and follow the same trends, the footprint impact of sweeps over the wall is stronger for the short structures close to the wall.

In addition to this, the change of trend experienced might also be explained from a statistical point of view. Ejections are lower in number than sweeps (see table 4), but occupy a much larger volume, while most sweeps and ejections are wall-attached. If most sweeps are attached but much smaller than ejections, then there may be less volume far from the wall occupied by wall-attached sweeps than by wall-attached ejections. There are some big sweeps extending from near the wall to the mid-plane, but they are not very frequent in the dataset, or not as much as big wall-attached ejections. Hence the 3-D GAN may learn better the patterns of those ejections, which are much more common. On the other side, close to the wall, wall-attached sweeps might be much more common than wall-attached ejections, so that the opposite happens.

4 Conclusions

A direct 3-D reconstruction from wall quantities with 3-D generative adversarial networks (GANs) has been proposed and demonstrated. The flow estimator builds on the successful reconstruction by Güemes et al. (2021 [143](#)) using a mapping from wall-shear stress and wall-pressure to 2-D wall-parallel velocity fields, and extends it to a full 3-D estimation. This extension comes with an affordable increase in the number of parameters and computational cost of the training if compared to the 2-D architecture estimating a single plane. The main advantage is a direct full reconstruction of the flow topology, without the need for training multiple networks for planar reconstruction. The argument for the reduction of the number of network parameters is the concept of parameter sharing, according to which the filters of the convolutional layers are shared among different wall-normal distances. Besides, the reconstruction based on this methodology, in which the main element of the network is the 3-D convolutional layer, ensures continuity within the reconstructed domains. In contrast, a procedure based on merging independently reconstructed 2-D planar domains could give rise to

discontinuities.

The algorithm is tested on channel flow data at friction Reynolds number 200. The results in terms of reconstruction accuracy of the velocity fluctuations show a similar trend to the case of the 2-D single-plane estimators, with lower error on the streamwise velocity component with respect to the spanwise and wall-normal components. The error is in all cases slightly larger than in the reference case of the 2-D estimator with a similar Reynolds number. This was expected due to the comparably lower number of parameters per output node used in the 3-D estimator.

We also observe that a reduction in the target volume size does not always correspond to an improvement in accuracy. The estimator trained with test case D, which contained only the region with $y/h > 0.21$, performed worse in terms of reconstruction accuracy than the estimator of case A, whose target was the entire volume. This can be explained by the difficulty of the network to ignore the parts of the wall fields that were related only to structures located in $y/h < 0.21$ when trying to reconstruct the outer region of test case D. While in case A, a large portion of the network parameters is trained to establish the mapping between near-wall features and wall quantities, in case D, the estimator should learn to filter out the portion of wall quantities that is due to near-wall structures and at the same time is uncorrelated with the structures in the target volume. Due to the modulation effect of large scales on the near-wall cycle, there is an inevitable loss of accuracy in this process. While this might be frustrating in view of training neural networks that target the reconstruction of far-from-the-wall structures—which might be interesting for control purposes—the higher performance of 2-D estimation with respect to case D hints at the possibility of accuracy improvement by increasing the number of parameters for this task.

The estimators for each of the cases, regardless of the reconstruction domain, seem to target specific features of the flow. In particular, wall-attached structures are reproduced with high fidelity at least up to $y/h = 0.5$ (i.e. 100 wall units), while a significant fraction of detached structures is filtered out in the process. This is a desirable feature since wall-attached structures carry the bulk of Reynolds stresses, and it was somewhat foreseeable due to their stronger wall footprint. We thus envision higher difficulty in predicting detached structures.

Furthermore, there are some differences in the prediction quality depending on the type of structure in relation to the footprint that they produce. Among wall-attached structures extending only in the near-wall region, sweeps are estimated slightly better than ejections.

However, the opposite situation is found between wall-attached sweeps and ejections that extend up to the middle of the channel or close to it. The footprint of the structures to be reproduced is a key aspect in this process.

Sweeps and ejections dominate other types of structures. Although some of them, mainly wall-detached ones, are filtered by the GAN, the statistics of the identified structures in terms of volume and quantity are reasonably well preserved. There are more sweeps than ejections, but ejections are much bigger and occupy most of the volume identified as a structure according to the hyperbolic hole size employed.

Supplementary material

Data related to this work and trained models of the neural networks are openly available at <https://doi.org/10.5281/zenodo.11090713>. Codes developed in this work are openly available at: <https://github.com/erc-nextflow/3D-GAN>.

Funding

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No 949085, NEXTFLOW). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. A. C. M. acknowledges financial support from the Spanish Ministry of Universities under the Formación de Profesorado Universitario (FPU) program 2020. R. V. acknowledges financial support from ERC (grant agreement No. 2021-CoG-101043998, DEEPCONTROL).

Declaration of Interests

The authors report no conflict of interest.

Appendix A. Criteria for the design of the generator network

The networks proposed in this work have been trained on an NVIDIA RTX-3090 GPU. In addition to the necessary modifications to unlock mapping 2-D fields to 3-D domains, some changes to the network architecture were proposed in order to optimize it and obtain a more

accurate flow estimation.

An important degree of freedom in defining the generator network is the position of the up-sampling layers, a tool commonly used for super-resolution purposes to make it possible to match the lower-resolution input with the higher-resolution output. Often, these layers are placed at the beginning, in the first layer or before the bulk of residual blocks (Dong et al. 2015 [159]; Wang et al. 2015 [262]), which would allow the following convolution layers to operate in a wider domain. Other authors prefer to define a gradual positioning of these layers (Osendorfer, Soyer & van der Smagt 2014 [263]). These two approaches, and in particular the first one, have the inconvenience that they produce models with 3-D convolutional layers much heavier than if the up-sampling layers were placed at the end of the network, which is another alternative (Shi et al. 2016 [264]). This latter option is found in the studies by Ledig et al. (2017 [160]) and Güemes et al. (2021 [143]). Concerning convolution layers, these become more complex and computationally demanding, depending not only on the size of the domain but also on the number of filters included. If the up-sampling layers are gradually placed with a moderate number of convolution filters, and a large number of filters is used only at the end for a few convolution operations, then the required computational resources can be maintained or even substantially reduced.

Case A was first studied with a network comprising 16 residual blocks and all the up-sampling blocks after them. The convolutions in the residual blocks had 64 filters, and those in the up-sampling blocks had 256. As alternative architectures, several options have been tested, placing these up-sampling blocks not at the end, but right after specific residual blocks (as in figure 3 and table 2), which makes the convolution layers operate over broader domains in y . However, this change required convolution operations in up-sampling blocks to have 64 filters instead of 256, simplifying them and reducing the amount of trainable parameters. Indeed, this simplification is such that with the same machine and memory limitations, it allows us to increase the number of residual blocks further, from 16 to 32. This type of architecture is finally selected, as it reports a lower \mathcal{L} error (2.1) even using less trainable parameters and requiring a comparable training time. This arrangement is a balance between different aspects, placing them progressively to allow convolutions to operate over wider domains than if they were at the end, while maintaining an efficient use of the computational resources. Both models can be compared in figure 12(a) and in table 5.

In view of the results from case A and their physical analysis and interpretation, cases B

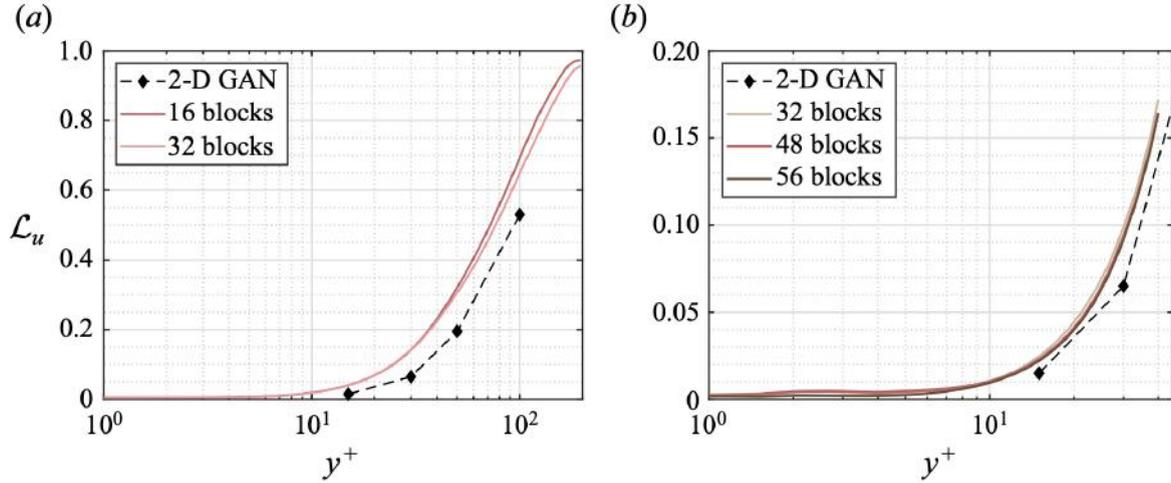


Figure 12. The MSE of the streamwise fluctuation velocity component u for test cases (a) A and (b) C, with alternative network architectures.

Residual Blocks	Filters per block		Residual block up-samp. scheme	Trainable parameters	
	Residual	Up-samp.		G	D
16	64	256	All at the end	13.6×10^6	18.2×10^6
32	64	64	0-6-12-18-24-30	9.0×10^6	18.2×10^6

Table 5. Details on the implementation of alternative architectures for case A. The ‘Residual block up-sampling scheme’ column indicates the indexes of the residual blocks that are followed by an up-sampling block.

Residual Blocks	Filters per block		Residual block up-samp. scheme	Trainable parameters	
	Residual	Up-samp.		G	D
32	64	64	6-12-18-24-30	8.0×10^6	23.8×10^6
48	64	64	8-16-24-32-40	11.6×10^6	23.8×10^6
56	64	64	10-20-30-40-50	13.3×10^6	23.8×10^6

Table 6. Details on the implementation of alternative architectures for case C. The ‘Residual block up-sampling scheme’ column indicates the indexes of the residual blocks that are followed by an up-sampling block.

and C were proposed. Case B needs a special implementation, as the network needs to provide output data with 48 layers in the wall-normal direction, which is not a power of 2. Hence the first up-sampling layer increases this size from 1 to 3, and the subsequent ones continue as powers of 2 up to 48. For the rest of the set-up, the selected architectures for these two cases follow the same structure, with 32 residual blocks in total. Also, case D was proposed as the second half of the layers originally in case A not included in case C, concentrating the resources in this region of the domain. Although the reconstructed volume in case D is larger than in case C, both have 32 layers in the wall-normal direction (see table 1).

As for the discriminator, the same structure is followed in all the cases, as depicted in figure 4. Pairs of 3D-convolution layers with increasing number of filters [64, 64, 128, 128, 256, 256, 512, 512] are used. From each pair, the first ones preserve the dimensions, and the second ones reduce the size of the data domain in the dimensions assigned to x , y and z , the first time by a factor of 4 and thereafter by a factor of 2.

An additional change was proposed. The depth of the network can be easily modified by setting more or fewer residual blocks, such as the aforementioned change from 16 to 32 which showed an improvement in reconstruction terms. In this sensitivity analysis, case C was further trained with 48 and with 56 residual blocks, with up-sampling blocks gradually placed every 8 and every 10 residual blocks respectively. These architectures and their performances can be compared in figure 12(b) and in table 6. They show a very moderate error reduction with respect to the previous situation with 32 blocks, with the curves practically overlapping each other, at the cost of using a substantially larger number of trainable parameters and taking longer to train. Without any remarkable improvements, these changes are discarded and the networks remain with 32 residual blocks for all cases.

Appendix B. Comparison of the 3D-GAN with other methodologies

In §1, GANs were set as the baseline for this work among other 2D estimation techniques. Along the discussion of the results of this work, the accuracy of the proposed methodology is compared to that of 2D-GANs in terms of MSE. In this appendix, we provide a comparison of the performance of this network (refer to case A, from the wall to the mid-plane of the channel) with alternative techniques, i.e. LSE and EPOD as linear techniques, and a deep neural network (DNN) in the ML framework. For the comparison with the 2D-GAN, it was convenient to compare the MSE of the velocity fluctuations normalised with their respective standard deviation at each wall-normal distance, as this normalization process is needed for both processes. For the comparisons we are providing here, this normalization is not needed as the performances are obtained from the same dataset, and the results expressed here in inner units could be more meaningful and easier to analyse. As well as for the 3D-GAN, 16 000 samples are used to establish the correlation and 4000 to test the capabilities of each methodology.

B.1 3D-GAN vs LSE

One of the most often used techniques in the literature for flow estimation from wall measurements and other flow estimation purposes is linear stochastic estimation. The 3D-GAN convolutional filters act over multi-dimensional matrices, guaranteeing domain continuity. In LSE, the estimator works independently for the estimation of different points, and in particular, must be different for each wall-normal distance and for each velocity component to be estimated.

In this case, the number of sensors n_s is 12288, as three quantities are measured on a 64×64 grid. Each of the quantities x' to be estimated can be computed through the projection of the vector $E[n_s \times 1]$ containing all the sensor measurements onto the vector $L[1 \times n_s]$ containing all the correlation coefficients:

$$x' = LE . \tag{B1}$$

Multiple samples are needed to compute the coefficients of L . To that end, multiple sensor entries are concatenated, defining a matrix where each column corresponds to a sample as follows:

$$\mathbf{E} = [E_1|E_2|\dots|E_t] . \tag{B2}$$

Given the known values of the quantities x_i assigned to each set of sensor data E_i , the coefficients

in L are obtained from this linear system:

$$L(\mathbf{E}\mathbf{E}^T) = [x_1, x_2, \dots, x_t] \mathbf{E}^T . \quad (\text{B3})$$

Once the coefficients that establish the correlation between the wall measurements and a velocity component at some point are known, the estimator (B1) can be used with new sensor inputs. It has been tested with 4 000 samples. The output has been compared with the fields of the DNS to compare its performance with the 3D-GAN. The MSE of the 3 velocity fluctuations $[u^+, v^+, w^+]$ is shown in figure 13.

It is evident that the 3D-GAN outperforms the LSE estimator. The errors are very low close to the wall. The three velocity components experience a peak with a maximum error around $y^+ \approx 80 - 90$ with the 3D-GAN. With the LSE, these error peaks are larger and are shifted towards the wall, particularly for u^+ , which has a substantially higher standard deviation (see figure 1). The errors of both techniques stabilize around similar values near the mid-plane of the channel.

The use of the 3D-GAN has a clear benefit on the estimation of u^+ , for which the highest errors are reported, with a maximum error almost 3 times lower with respect to the peak in LSE. The errors for the v^+ component, which has the lower standard deviation distribution along y^+ , are significantly lower than for u^+ and the benefit between the 3D-GAN and the LSE is not so remarkable, although still important. The error of w^+ and its error reduction observed between the two methodologies are in an intermediate position between u^+ and v^+ .

Note that in the present analysis, we are not using the spectral formulation of the linear stochastic estimation (SLES), which explicitly avoids spurious correlations between orthogonal Fourier modes (Encinar & Jiménez 2019 [132]). This choice is motivated by the interest in applying this technique to real-world applications, where the assumption of periodicity in the wall-parallel directions would be difficult to justify.

B.2 3D-GAN vs EPOD

The results obtained and previously discussed with the 3D-GAN are compared here with the performance of the EPOD estimator, whose methodology is described here.

The data of the velocity fluctuations, given in a multidimensional matrix, is rearranged in a 2D matrix X_U , with one row for each sample (16 000 in this case), with the velocity components assigned to each of the points in space along the columns. The same procedure is done with

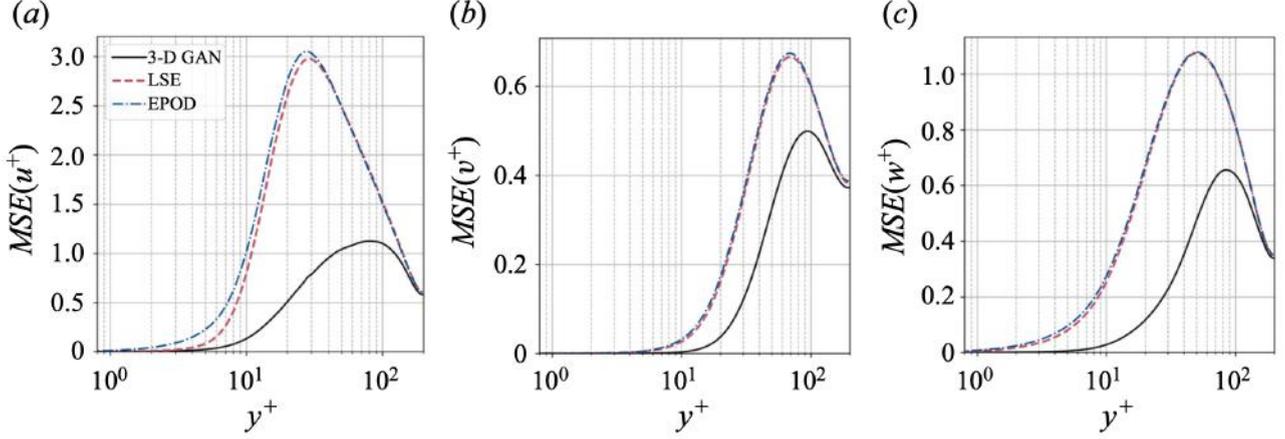


Figure 13. The MSE of the fluctuation velocity components (a) u^+ , (b) v^+ and (c) w^+ of the 3-D GAN case A (solid line), the LSE (dashed line) and the EPOD (dot-dashed line). Results are given in wall-inner units.

the data from the wall probes, cast in the matrix X_{pr} . All this information is reduced in modes following the singular value decomposition (SVD), with the temporal information of each sample in Ψ and the spatial information in Φ , leading to:

$$\left. \begin{aligned} X_U &= \Psi_U \Sigma_U \Phi_U^T ; \\ X_{pr} &= \Psi_{pr} \Sigma_{pr} \Phi_{pr}^T . \end{aligned} \right\} \quad (\text{B4})$$

These matrices have been cropped, retaining the most energetic 7 200 modes to remove those modes with low energy content that introduce noise in the problem. This threshold coincides with the 99% of the energy contained in X_U .

The final objective is to obtain the estimated (denoted with $*$) velocity fluctuations X_U^* for a different set of samples, given those probe measurements. To that end, the temporal coefficients Ψ_U^* associated with those samples are estimated and projected onto the spatial basis $\Sigma_U \Phi_U$ established previously as

$$X_U^* = \Psi_U^* \Sigma_U \Phi_U^T . \quad (\text{B5})$$

The temporal modes of the velocity field are obtained by projecting the temporal modes of the probes of the samples to be estimated onto the temporal correlation matrix Ξ of the probes and the velocity field:

$$\Psi_U^* = \Psi_{pr}^* \Xi = \Psi_{pr}^* \Psi_{pr}^T \Psi_U . \quad (\text{B6})$$

Again, the estimation of the 3-D GAN seems more faithful than that of the EPOD, as seen in figure 13. The behaviour of the error curve of the EPOD is very similar to that of the

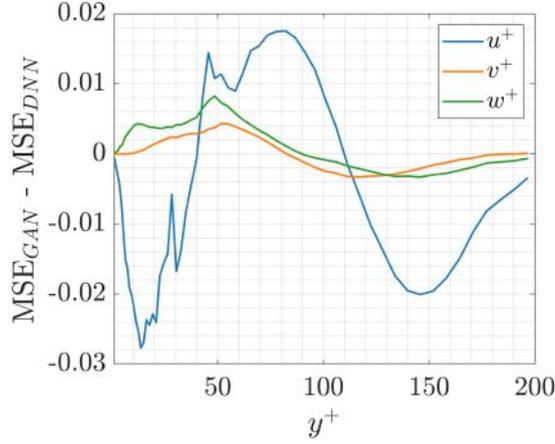


Figure 14. Difference of the MSEs of the fluctuation velocity components u^+ , v^+ and w^+ of the 3-D GAN and the DNN. Results are given in wall-inner units.

LSE, with the peak at a very similar y^+ and a small increment of error. The 3-D GAN, which incorporates nonlinearities in the problem, not only is capable of estimating the flow with a higher accuracy than these two linear techniques but also shifts the peak of maximum error away from the wall. This comparison shows the importance of performing the estimation with a nonlinear operator if accuracy farther from the wall is sought.

B.3 3D-GAN vs DNN

The performance of the 3D-GAN is also compared with a simpler concept of NN. To that end, this DNN replicates the generator network G while it neglects the discriminator D . Its loss function is solely based on the MSE, as in (2.1), with zero contribution from the adversarial loss (2.2).

The difference in the error between these two methodologies is in general moderately low (figure 14). The main benefit of the 3D-GAN is observed in the estimation of the streamwise velocity fluctuations. Although in some regions (approximately $y^+ \approx [40 - 110]$) the DNN estimates u^+ better than the 3D-GAN, the overall performance of the 3D-GAN is superior with a clear advantage in terms of accuracy in the near-wall region and in the outer region, *i.e.* the most challenging one. The improvement is much less relevant for the estimation of v^+ and w^+ . To provide a quantitative comparison of these results, the integrals of the area enclosed between these curves and the horizontal axis have been computed in the range $y^+ = [0 - 200]$. The biggest benefit is found for u^+ reporting -0.8027 . For v^+ there is a small benefit of -0.0131 , and for w^+ a small penalization of $+0.1424$.

Paper 2

Some effects of limited wall-sensor availability on flow estimation with 3D-GAN

Antonio Cuéllar, Andrea Ianiro, and Stefano Discetti, *Theoretical and Computational Fluid Dynamics* (2024), [doi:10.1007/s00162-024-00718-w](https://doi.org/10.1007/s00162-024-00718-w)

In this work we assess the impact of the limited availability of wall-embedded sensors on the full 3D estimation of the flow field in a turbulent channel with $Re_\tau = 200$. The estimation technique is based on a 3D generative adversarial network (3D-GAN). We recently demonstrated that 3D-GANs are capable of estimating fields with good accuracy by employing fully-resolved wall quantities (pressure and streamwise/spanwise wall shear stress on a grid with DNS resolution). However, the practical implementation in an experimental setting is challenging due to the large number of sensors required. In this work, we aim to estimate the flow fields with substantially fewer sensors. The impact of the reduction of the number of sensors on the quality of the flow reconstruction is assessed in terms of accuracy degradation and spectral length-scales involved. It is found that the accuracy degradation is mainly due to the spatial undersampling of scales, rather than the reduction of the number of sensors per se. We explore the performance of the estimator in case only one wall quantity is available. When a large number of sensors is available, pressure measurements provide more accurate flow field estimations. Conversely, the elongated patterns of the streamwise wall shear stress make this quantity the most suitable when only few sensors are available. As a further step towards a real application, the effect of sensor noise is also quantified. It is shown that configurations with fewer sensors are less sensitive to measurement noise.

Keywords Turbulent boundary layers, machine learning, channel flow, wall measurements

1 Introduction

In this work we challenge the methodology proposed in Ref. [265] for one-shot 3D flow estimation from wall quantities in scenarios with realistic limitations of practical applications. Limited number of sensors, choice of their distribution, measurement noise, and limited physical data information are the main aspects we address in this study.

Many studies have focused on flow control since the work by Prandtl in 1904, where the basics of the physics of boundary layer theory and separation were introduced [266], and great efforts are still being made in this area. Flow control could allow us to design more efficient devices aiming at different purposes, such as maximization of mixing, noise reduction or suppression, heat transfer, lift enhancement, delayed laminar-to-turbulent transition, or skin friction reduction [49].

The control of skin friction in wall-bounded flows is a key element for several energy-intensive industries (e.g. aviation), where the reduction of the viscous drag through turbulence control could lead to important savings. Several surface modifications such as LEBUs (large eddy break-up devices) and riblets have been shown to reduce the skin friction drag [52], [53]. However, passive devices cannot be tuned specifically for each operating condition; as such, active flow control is investigated as an alternative to passive flow control [54], [55]. Active flow control in a closed-loop arrangement allows to design actuation mechanisms specially designed for the control of a given state of the flow or specific flow features, e.g., coherent structures. With this regard, a more detailed comprehension of the flow allowed to shift the actuating techniques under study from energetically expensive brute-force approaches towards small amplitude forcing [56] to control coherent structures, with significant savings also in terms of mass and size. However, to build a closed-loop flow-control system we need the availability of sufficiently fast and accurate sensing and actuation devices. Hence, it is of utmost importance to research and develop sensing systems for flow control purposes.

The requirement of low intrusiveness sets the need to embed sensors within the wall. Space-time correlation within boundary layers can be exploited for this purpose. The existence of strong correlations among the motions within turbulent boundary layers was first discovered using point measurements with hot wires [267], [268]. In particular, linear techniques have been widely used to establish a correlation between the wall measurements and the turbulent flow field. The linear stochastic estimation (LSE) [130], [132], [232] searches for a linear transfer function from the wall to the flow fields. Other approaches focus on the existence of coherent

structures within the flow and aim at identifying a linear relation between coherent flow features and their wall signature. Since the development of the attached-eddy model by Townsend [72], large eddies attached to the wall are considered to be among the dominant energy-containing motions in wall-bounded turbulent flows. Some preliminary descriptions of the large eddies reported an anisotropic behaviour consistent with the presence of counter-rotating vortex pairs [65]. The presence of these characteristic events with certain coherence and spatio-temporal dependence amidst the randomness of turbulence made it possible to address this problem from a statistical point of view. One of the main tools for the identification of coherent structures in turbulence has been the proper orthogonal decomposition [POD, 269]. Reference [66] first employed POD for the decomposition of wall-bounded turbulence. Obtaining a low-order decomposition of wall-bounded turbulence allows us to have a compact description of the most energetic features of the flow which might have a distinct wall signature. A remarkable example of this approach for flow estimation is represented by the extended POD (EPOD) [118], [139], [237] which is found to have a comparable reconstruction accuracy with respect to LSE. Although faithful flow estimations in the vicinity of the wall can be obtained with these linear techniques, they are limited as far as the physics involved contains non-linear events, explaining why some studies coped with this problem by combining linear methodologies with the introduction of non-linearities. When dealing with multiple-time-delay estimation of a turbulent channel flow, faithful linear estimations were established within the viscous layer, while non-linearities enabled accurate estimations even in the buffer region [145]. With a similar reasoning, Kalman-filter-based estimators extended to account for non-linearities and introduced non-linear forcing were shown to outperform the traditional linear estimators [126], [134]. Recently, different resolvent-based approaches have been applied for estimation purposes, including the space-time flow statistics estimation from limited data [270]. The resolvent-based estimation of the non-linear forcing terms improved the estimation obtained with Kalman filter-based estimators [271]. This type of methodology was also used for velocity and pressure estimation in the flow from measurements of pressure and shear stress at the wall [272]. Likewise, resolvent-mode and resolvent-based approaches have been applied to estimate the flow from velocity sensors arranged in given planes of the domain [273].

The capability of machine learning to deal with non-linear problems leaves this branch of methodologies in an outstanding position to describe the non-linear relations needed to estimate the flow within a turbulent boundary layer. Convolutional neural networks (CNNs), thanks to the use of convolutional filters, can detect patterns within image data at multiple scales.

Analogously to POD, CNNs can identify patterns within the flow and have been exploited for instance for encoding purposes [114]. The use of convolutional filters within a CNN was shown to have better reconstruction capabilities than EPOD for the task of instantaneous field estimation [139]. While Ref. [139] employed CNNs to estimate the POD coefficients of the flow fields, later works successfully employed fully-convolutional networks (FCNs) [142], [242] to reconstruct the flow fields on planes at a certain distance from the wall. These neural networks are capable of making instantaneous reconstructions of the fluctuation velocity field in wall-parallel planes providing them with wall measurements of quantities such as the wall pressure, the wall shear stresses or the heat transfer through the wall. Alternatively, field predictions in other physical configurations have been targeted by CNNs with generative adversarial training, referred to as generative adversarial networks (GANs) [244]. GANs have shown to be very effective also for resolution enhancement purposes [148], [157], [158]. The implementation of GANs allowed us to estimate the turbulent field of a channel flow from coarse wall measurements at the same instant [143], reporting an outstanding performance under low-resolution wall-data input and better prediction capabilities than FCNs.

3D convolutional networks have been explored for the purpose of 3D reconstructions of free-surface flows [150] and turbulent channels [148] employing surface and flow field measurements, respectively. We recently extended the instantaneous flow estimation in wall-parallel planes through a GAN [143] to a full 3D estimation with 3D-GANs [265]. This alternative conceptualization of the problem overcomes the need to develop individual planar reconstructions at different wall-normal distances to get a full reconstruction, and produces benefits such as the reduction of trainable parameters in relative terms, or the generation of a 3D field without discontinuities. The availability of 3D fields could be directly used to study coherent turbulent structures. The analysis conducted over the predicted set of Q-events [265] highlighted that the capability of the network to reconstruct the field at a certain location depends both on the distance to the wall and on the wall footprint of the specific structure to reconstruct. Wall-attached sweeps and ejections that penetrate deep towards the centre of the channel can be reasonably well predicted by the 3D-GAN, while detached coherent events might be unseen. The capabilities of the methodology were also compared with traditional linear techniques, such as LSE or EPOD, showing clear advantages over them. The level of error was significantly reduced and the region from the wall to the point with maximum error was clearly extended. The interested reader can refer to Ref. [265] for an extensive performance assessment against linear techniques.

Another concern in flow control studies involves the number and positioning of the sensors. Fewer sensors would be preferred as the computations executed before any actuation or decision may be accelerated, while more sensors might provide a more detailed description of the state of the flow if they are placed efficiently. Even if the flow field to be controlled is a high-dimensional system, coherent structures could be represented on a latent low-dimensional attractor, enabling sparse sensing [179]. Recently developed systems for optimal placement of sensors and actuators have been demonstrated to provide substantially better performances than random placement approaches [274].

The practical implementation of this type of technology for the development of active control strategies might raise questions about its technical complexity or even its feasibility. To get an accuracy comparable to that in Ref. [265], two shear-stress and one pressure sensors must be embedded in the wall within a grid 64×64 points (streamwise \times spanwise) arranged within an area equal to πh (streamwise) \times $\pi h/2$ (spanwise), with h the half-channel height. Setting this case as a baseline, the objective of this work is to define alternative cases with a reduced amount of sensors and limited physical data information to assess how the performance of the 3D-GAN estimator degrades. Moreover, this work contains an assessment considering noise in measurements, a common issue in instrumental applications with sensors. After this introduction, §2 presents the methodology followed in this work, §3 reports and discusses the results obtained and §4 summarizes this manuscript highlighting the main conclusions and the most remarkable aspects.

2 Methodology

This work focuses on the reconstruction of the flow field within a turbulent channel. The dataset includes both wall and flow fields and it is the same dataset employed in Ref. [265], here briefly described in Sect. 2.1. To study the effect of limited sensor number, reduced physical quantity availability, and measurement noise, the information at the wall is subsampled and/or corrupted as described in Sects. 2.1.1 and 2.1.2.

2.1 Turbulent channel flow dataset and sensors in the wall

The dataset has been generated with a DNS [203] of a periodic turbulent channel flow with friction Reynolds number $Re_\tau = 200$, with dimensions $[\pi h \times 2h \times (\pi/2)h]$ —respectively in the streamwise, wall-normal, and spanwise directions, indicated respectively with x , y and z . It contains the three velocity components of the flow field with 64 equispaced points both along

the streamwise and spanwise directions, and 64 points in the wall-normal direction from one of the walls to the mid-plane (consider only half of the channel). Moreover, it contains wall measurements of pressure p_w and streamwise and spanwise wall shear stresses τ_{w_x} and τ_{w_z} . The configuration used in the present work is equivalent to that of Case A of Ref. [265], estimating the flow field of half channel, from the wall to the mid-plane, as represented in Fig. 1. The components of the velocity fluctuations at each wall-normal position are identified as $[u, v, w]$, using the + symbol to indicate quantities expressed in inner scaling.

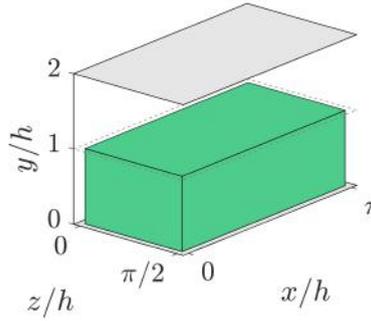


Fig. 1 Representation of the region from the bottom wall to the mid-plane that is reconstructed providing wall measurements of the bottom wall

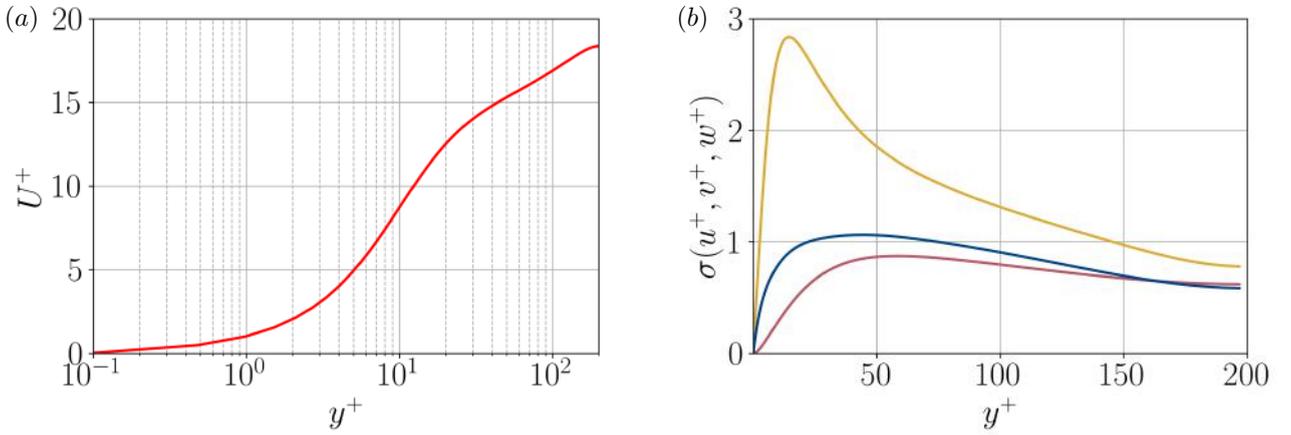


Fig. 2 Wall-normal profiles of (left) the mean streamwise velocity and (right) standard deviation σ of the three velocity components. This figure is adapted from Cuéllar et al., ‘Three-dimensional generative adversarial networks for turbulent flow estimation from wall measurements’, doi:10.1017/jfm.2024.432, licensed under CC BY 4.0, Ref. [265]

The wall-normal profiles of the channel are represented in Fig. 2. According to the literature, despite being small compared to some other databases, the present channel is big enough to contain self-sustained turbulence [253], [255]. Further details comparing this channel with similar channels at $Re_\tau = 180$ can be found in Ref. [265].

Table 1 Number of sensors and retained variance with respect to the DNS resolution on each sensor arrangement on the wall

Sensor arrangement	64×64	32×32	16×64	16×16	8×32	8×8
Number of sensors	$2^{12} \cdot 3$	$2^{10} \cdot 3$	$2^{10} \cdot 3$	$2^8 \cdot 3$	$2^8 \cdot 3$	$2^6 \cdot 3$
Retained variance p_w	100%	96.6%	90.0%	85.3%	69.8%	59.7%
Retained variance τ_{w_x}	100%	96.9%	98.3%	86.9%	92.0%	63.1%
Retained variance τ_{w_z}	100%	90.0%	89.9%	68.9%	70.1%	43.2%

2.1.1 Downsampling procedure

In this work, we are proposing different grid arrangements for the wall sensors used in the training and estimation process. Table 1 summarizes the different cases proposed in relation to the number of sensors used in each of them and the percentage of retained variance with respect to that of the DNS for each wall quantity. The variance is reduced because the downsampled fields employ bigger sensors which average the measurements of four contiguous sensors in the higher-resolution case, thus filtering smaller turbulent features. The final objective for all cases is to estimate the entire 3D flow field of the half channel, in its full resolution of $64 \times 64 \times 64$ gridpoints. The sensor resolution, on the other hand, is progressively downsampled in powers of two, leading to these three new cases: 32×32 , 16×16 and 8×8 (streamwise \times spanwise), as seen in Fig. 3. It represents the patterns' size in relation to the spacing between sensors and illustrates the effect of spatial averaging. Each downsampling step reduces the number of sensors by 75%. The case with full resolution of the wall sensors (64×64) is included for reference. It is worth remarking that each test case is trained separately with different random seeds of the weights. This ensures that the effect of sensor number reduction is properly isolated.

The corresponding downsampled sensor inputs are computed with an average pooling filter applied to the original DNS wall data. This is done to maintain full coverage of the wall but with progressively larger (and not overlapped) sensors. This introduces an additional challenge of progressive amplitude modulation of the smallest scales observed in the wall patterns.

Additional cases are proposed in which measurements from wall sensors are not downsampled by the same factor along x and z . This is particularly interesting as streamwise-elongated patterns are dominant. An increment in the spanwise sampling resolution could provide further information about the different elongated streaks that are seen along an array of sensors in this direction. On the contrary, increasing the sampling resolution along x might just provide more information about only one or a few streaks. Smart sensor placement techniques have been shown to use the same amount of sensors more efficiently [179], [274]. The methodology

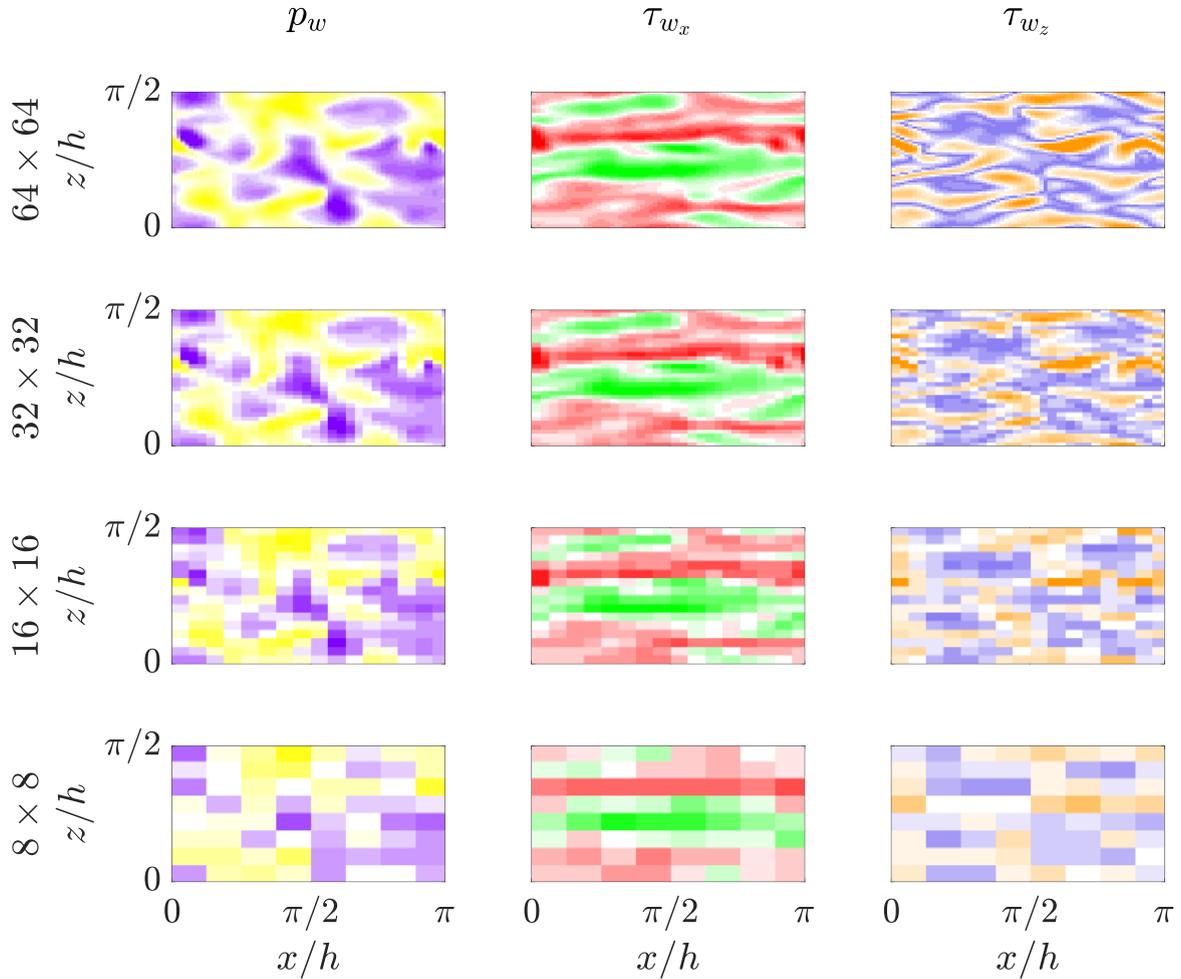


Fig 3 Instantaneous representation of wall quantities with different sampling levels. From left to right, wall pressure p_w , streamwise and spanwise wall shear stress τ_{w_x} , τ_{w_z} , respectively. From top to bottom, full resolution 64×64 , and progressive isotropic downsampling of factor 2, 4, and 8.

proposed here is not a smart sensor placement but tries to compare homogeneous sensor patterns with the same number of sensors providing full coverage of the domain considered in this problem. In line with this asymmetric downsampling, two additional cases are proposed, with corresponding baseline with the same number of sensors but symmetric distribution: 16×64 (being 32×32 as baseline) and 8×32 (with 16×16 as a reference for comparison). A visual representation of this downsampling scheme can be seen in Fig. 4.

To consider the case of limited physical data information we also explore the effect of employing sensors capable of measuring only one physical quantity—only τ_{w_x} , τ_{w_z} or p_w . Furthermore, we address jointly the effect of the lack of availability of high sensor resolution in the wall, together with that of different quantities.

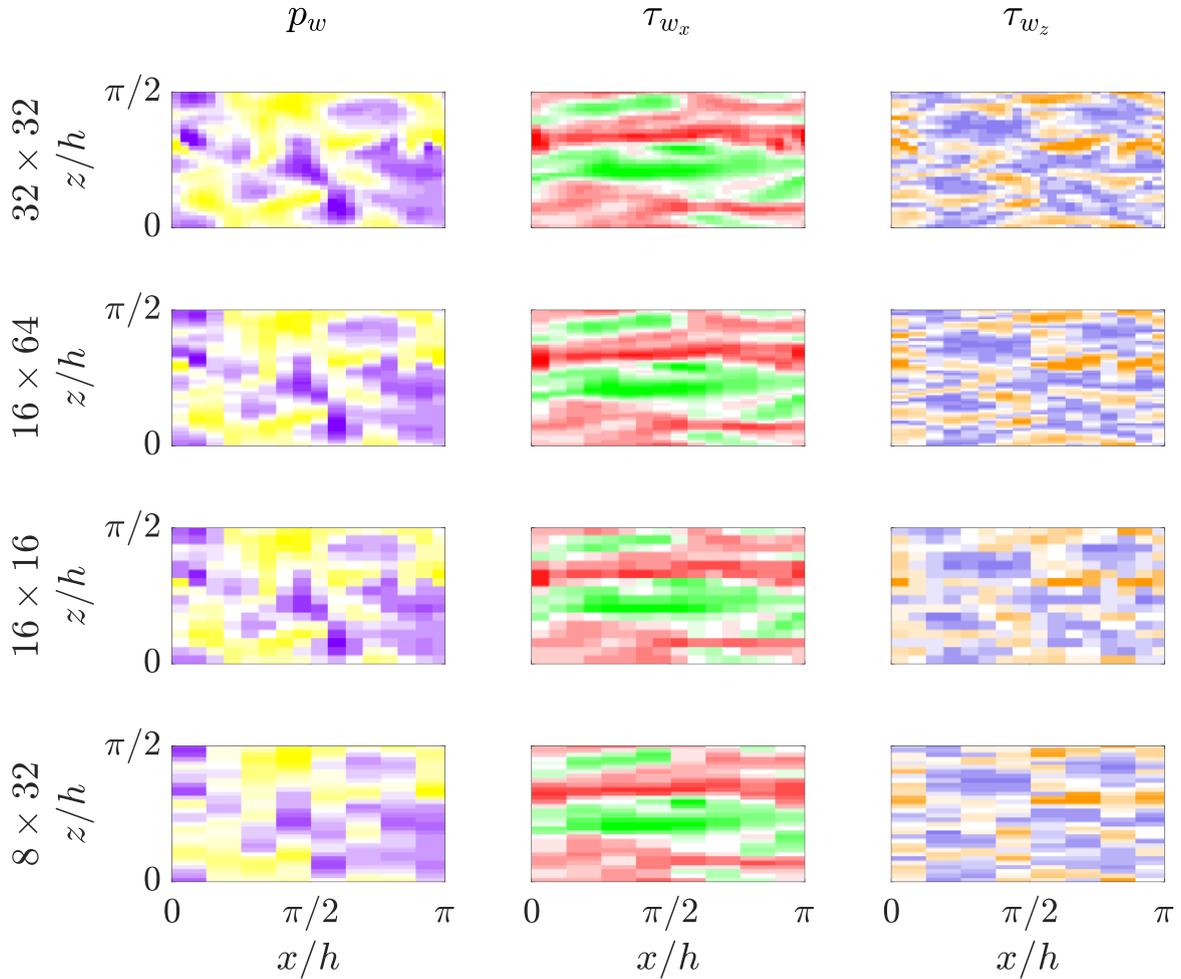


Fig. 4 Instantaneous representation of wall quantities with different sampling levels. From left to right, wall pressure p_w , streamwise and spanwise wall shear stress τ_{w_x} , τ_{w_z} , respectively. From top to bottom, on even lines the cases with asymmetric sensor distribution, and on odd lines their corresponding baseline cases

2.1.2 Noise modelling

In real applications, measurement noise in the sensors must be considered. To simulate this effect affecting wall measurements, random Gaussian noise is added to p_w , τ_{w_x} and τ_{w_z} . Two levels of noise are tested on the four cases with the same number of sensors in both directions. The standard deviation of the noise δ is set to 1% and 3% of the standard deviation of each of the wall-measured quantities. Recall that the downsampling is applied in terms of an average pooling layer directly applied to the sensed data. For these cases, the level of noise is thus increased according to \sqrt{N} (see Table 2) to account for the smoothing effect of the average pooling. No bias error is considered for simplicity.

To estimate the effect of noise on the reconstruction quality, isolating training and reconstruction processes, the networks, previously trained with the clean DNS data, are fed with the

Table 2 Number of sensors used per downsampled cell.

Sensor arrangement	64×64	32×32	16×16	8×8
N	1	4	16	64

wall fields with added noise. Results are reported in Sect. §3.4 where the performance is also compared with that of the reconstruction without noise in the input data.

2.2 Generative adversarial network

The flow estimation is carried out with a GAN. The architecture comprises two networks, the generator \mathcal{G} and the discriminator \mathcal{D} . The generator is fed with 2D instantaneous sensor fields and provides the three 3D fields of u , v and w at the same time instant. The discriminator is used solely during the training process. It is fed with flow fields—either the original ones from the DNS or those generated by \mathcal{G} —and must classify them accordingly. The two networks are trained adversarially. The loss function introduces a penalty based on the adversarial loss, quantifying how good is the performance of each network in the “game” it is playing against the other. At the end of the training process, \mathcal{G} should be capable of generating turbulent flow fields with a level of accuracy such that \mathcal{D} is confused, having difficulty in classifying fields as original or generated.

The GAN architecture we used is based on the field estimator in wall-parallel planes developed in Ref. [143]. It was extended to 3D in Ref. [265], with 64×64 sensor-grid inputs onto \mathcal{G} . In this work, the architecture is adapted to accommodate different input sizes (see sketch in Fig. 5). Upsampling layers are applied to make the size of the output of \mathcal{G} match the expected $64 \times 64 \times 64$ resolution of the half-channel field. Along the y direction, the domain needs to grow from a unique wall-parallel layer to 64, for which six upsampling layers of factor 2 are employed [265]. Instead, along both x and z , the sizes of the domain need to be widened by a factor defined as the ratio of the output (64) and the corresponding input sizes ($64 - 32 - 16 - 8$). This step is done with a single layer introduced after the first convolution filter of the 3D-GAN network in [265]. Alternatively, an additional procedure to increase the resolution of the wall measurements before feeding them to the 3D-GAN could be considered. However, CNNs can be used with super-resolution purposes [160], and the architecture employed in \mathcal{G} may jointly allow to increase the resolution and predict the flow field from wall measurements at once. The use of the same architecture for the different test cases, without the need of additional networks or any other type of implementation, and having used the same amount of trainable parameters,

further allows us to establish a fairer comparison between them.

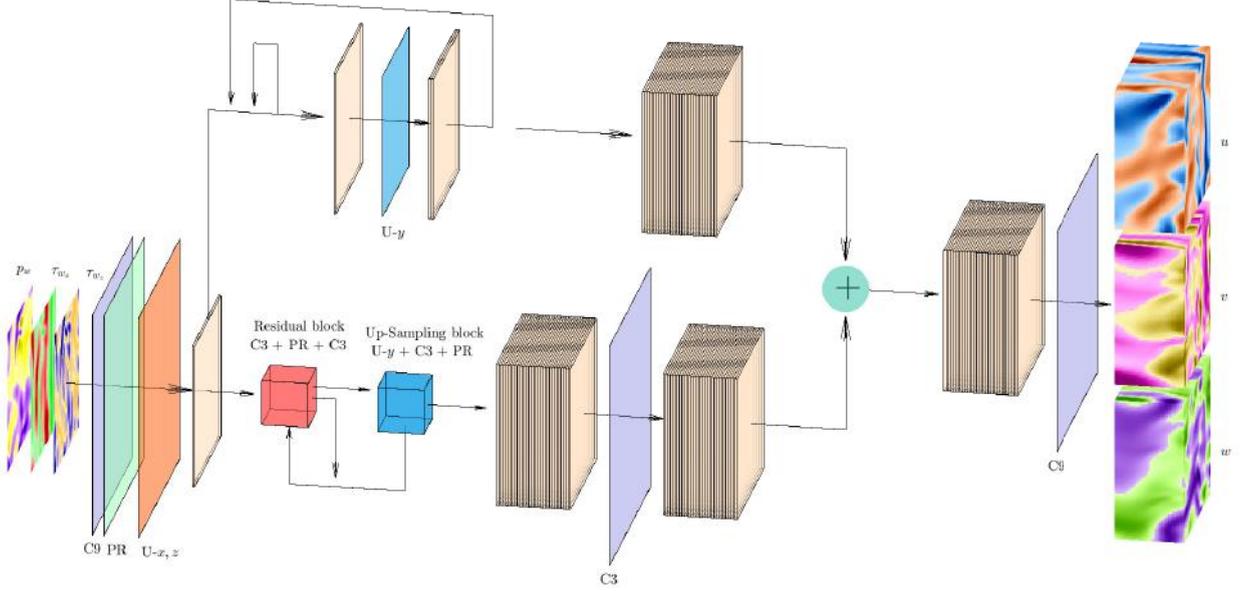


Fig. 5 Diagram of the generator network. The network contains convolutional layers (C3 and C9 with kernel size 3 or 9 respectively), parametric-ReLU (PR) and upsampling layers acting on the wall-normal direction (U-y) and in both wall parallel directions (U-x,z). The residual block and the up-sampling block are recursively repeated. This figure is adapted from Cuéllar et al., ‘Three-dimensional generative adversarial networks for turbulent flow estimation from wall measurements’, doi:10.1017/jfm.2024.432, licensed under CC BY 4.0 (<https://creativecommons.org/licenses/by/4.0/>), Ref. [265]

Moreover, the network is a deep neural network consisting of convolutional layers and PReLU activation functions [190]. They are arranged along 32 residual blocks with skip connections. The six upsampling layers to widen the domain in the y direction are spread along the depth of the network, and are followed by an additional convolutional filter. After all the residual blocks, there is a global skip connection and a final convolution layer with three output filters representing the $[u, v, w]$ velocity fluctuations respectively.

The discriminator receives the fields of $[u, v, w]$ with resolution $64 \times 64 \times 64$, and the size of the domain is progressively reduced, first with a set of convolutional filters with striding, and then with two fully-connected layers, with 1024 and 1 neuron as output. Finally, a sigmoid activation function is used to classify the field providing an output in the $[0 - 1]$ range. A sketch of the discriminator network \mathcal{D} is presented in figure 4 of reference [265].

The loss function of each network is composed of two terms. One aims to improve its own performance in terms of fitting the training data, while the other introduces a penalty based

on the performance of the other network (the so-called adversarial loss). The loss function of \mathcal{G} contains the MSE of the generated velocity field with respect to the original field from the DNS (\mathcal{L}_{MSE}), and the binary cross-entropy as adversarial loss:

$$\mathcal{L}_G = \mathcal{L}_{\text{MSE}} - 10^{-3} \mathbb{E}[\log D(\mathbf{u}_{\text{GAN}})] , \quad (1)$$

where \mathbb{E} is the expectation operator and $D(\mathbf{u}_{\text{GAN}})$ represents the output of \mathcal{D} when it receives a GAN-generated velocity field.

The loss function for \mathcal{D} is defined as the sum of two terms: its own expectation when it receives a DNS field as content loss and the expectation when it receives a generated field as adversarial loss:

$$\mathcal{L}_D = -\mathbb{E}[\log D(\mathbf{u}_{\text{DNS}})] - \mathbb{E}[\log(1 - D(\mathbf{u}_{\text{GAN}}))] . \quad (2)$$

During training, all DNS nodes are assigned the same weight in the loss, regardless of their corresponding volume (the mesh is finer close to the wall). We observed in our previous work [265] that the accuracy of the flow reconstruction is mainly driven by the existence of certain relations between wall and flow data. Consequently, introducing weights based on the volume of each DNS cell does not have a significant impact on training, and might even result in a detrimental effect on the accuracy, especially in the near-wall region.

3 Results

The mean-squared error (MSE) is used to evaluate the performance of \mathcal{G} in estimating the fluctuation velocity fields. It is computed independently for each component $[u, v, w]$ and for each y^+ coordinate and discussed in the following for the cases under study according to:

$$MSE(u_i^+, y^+) = \frac{1}{N_s N_x N_z} \sum_s \sum_x \sum_z (u_{i,\text{DNS}}^+(i, j, k) - u_{i,\text{GAN}}^+(i, j, k))^2 \quad (3)$$

where u_i^+ represents each possible velocity component $[u^+, v^+, w^+]$ with the subscript DNS or GAN to represent respectively the original and the reconstructed velocity fields, N_s is the number of samples in the testing set and N_x and N_z are the number of grid points along x and z , respectively.

Additionally, an integral error ε metric (4) based on the MSE has been defined to facilitate a quick comparison between the performance of different cases, merging in a single number the effect of each component and the distance to the wall. Table 3 collects the metric for the

Table 3 Integral error ε for the different cases proposed.

Grid	3 sensor types		1 sensor type			3 sensors + noise	
			p_w	τ_{w_x}	τ_{w_z}	1%	3%
64×64	0.207	0.224	0.224	0.232	0.242	0.219	0.234
32×32	0.211	0.233	0.233	0.241	0.254	0.222	0.239
16×64	0.214	-	-	-	-	-	-
16×16	0.227	0.272	0.2664	0.321	0.228	0.228	0.232
8×32	0.229	-	-	-	-	-	-
8×8	0.262	0.318	0.293	0.357	0.262	0.262	0.262

different cases considered in this work:

$$\varepsilon = \frac{1}{200} \int_{y^+=0}^{y^+=200} \frac{MSE(u) + MSE(v) + MSE(w)}{[rms(u) + rms(v) + rms(w)]^2} dy^+ . \quad (4)$$

Results are also discussed with a spectral analysis of the wall-sensed scales and the velocity scales.

3.1 Estimation with same downsampling factor in x and z

The performance of the four cases with the same downsampling factor in both directions can be compared in Fig. 6.

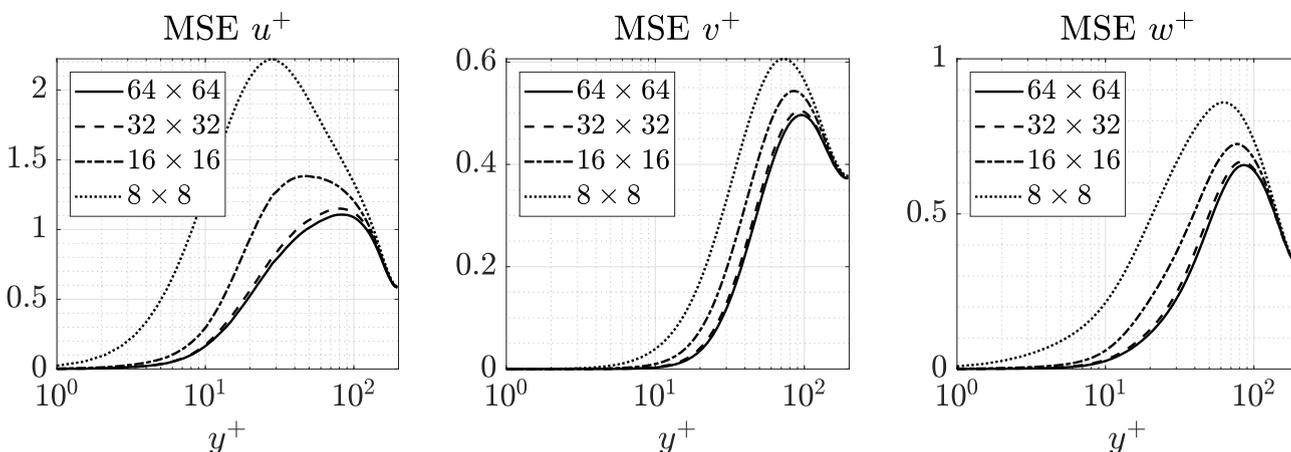


Figure 6 MSE reported for each velocity component and input-resolution case as a function of the wall-normal distance (—, 64×64 ; ---, 32×32 ; -.-, 16×16 ; ···, 8×8). The results are expressed in wall-inner units based on the friction velocity u_τ .

First of all, it should be remarked that all the curves observed in Fig. 6 have a similar behaviour. The error is very low in the vicinity of the wall, it reaches a maximum and then decreases. Practically no differences in error are observed at the centre of the channel. As the sensor resolution becomes coarser, the position of the maximum error shifts towards the wall, especially for u^+ . For all these cases, the MSE of u^+ is the highest, with maximum values of

about u_τ and even $2u_\tau$ for 8×8 . The MSE of w^+ is in an intermediate position, with the MSE of v^+ being the lowest. This is in line with the magnitudes of the standard deviation of the velocity fluctuations, as reported in Ref. [265].

As expected, the error increases progressively as the wall measurements become coarser—following cases with the same downsampling factor along both x and z directions. However, the additional error introduced between cases is not the same. From a qualitative observation of Fig. 3, one could notice that the smallest patterns observed with the 64×64 resolution are nearly preserved for the downsampled set with 32×32 sensors. Significant scale losses are instead observed for the 16×16 , and in particular, for the 8×8 resolutions. Figure 6 confirms that the 32×32 wall input generates a slightly higher error than that with 64×64 sensors, while these jumps increase when the resolution is further reduced. These jumps are in line with the behaviour of the integral error ε reported in Table 3. This effect should be assessed in relation with the distance between sensors (see Table 4) and the Power Spectral Density (PSD) of the measured patterns at the wall (Fig. 7). On the one hand, regarding the spacing between the sensors, the minimum characteristic sizes of the streamwise and spanwise patterns expected to be retained are twice the spacing distances reported in Table 4, according to the Nyquist-Shannon sampling theorem. Except for some of the smallest and low-energetic patterns observed in the 64×64 maps in Fig. 3, the sensor spacing for the 32×32 arrangement (see Table 4) fulfils this condition for most of the relevant flow scales. This might explain why the loss in accuracy reported in Fig. 6 between the 64×64 and 32×32 arrangements is minimum—despite the strong reduction in the number of sensors from 12288 to 3072. Instead, the 16×16 sensor spacing leaves an important proportion of the patterns out of the threshold; this effect is further exacerbated for 8×8 sensors. Similarly, Table 1 shows how the retained variance with respect to the original resolution, as an indicator of the contained energy, is quite similar for the cases with 64×64 and 32×32 sensors, while it is reduced more noticeably for 16×16 and 8×8 sensors. These reductions in the number of sensors determine an inadequate sampling of relevant flow scales, losing patterns that seem to be important for the reconstruction of the 3D flow field, in particular for the u component.

On the other hand, the spectrum maps shown in Fig. 7 (top row) also denote that large patterns—both along x and z —dominate for the three measured quantities, and those of τ_{w_x} are particularly elongated in the streamwise direction. The p_w and τ_{w_x} spectra are quite well conserved, with some degree of distortion growing with the downsampling factor, but still

Table 4 Spacing between gridpoints in the wall of the channel for each resolution configuration employed.

	64	32	16	8
$\Delta x/h$	0.049	0.098	0.196	0.393
Δx^+	9.82	19.63	39.27	78.54
$\Delta z/h$	0.025	0.049	0.098	0.196
Δz^+	4.91	9.82	19.63	39.27

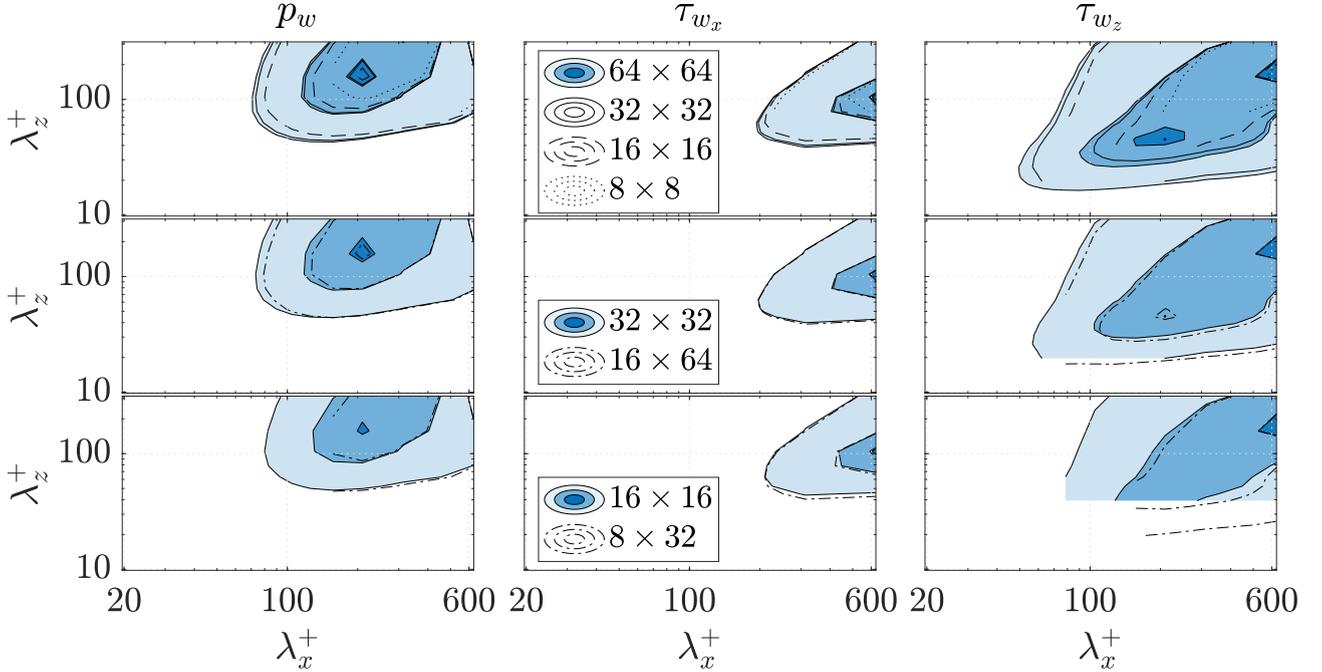


Fig. 7 Premultiplied PSD maps of the wall-sensed p_w (left), τ_{w_x} (centre) and τ_{w_z} (right). Cases with different sensor arrangements (top row: coloured —, 64×64 ; —, 32×32 ; - - - , 16×16 ; ·····, 8×8 ; central and bottom rows: coloured —, 64×64 and 32×32 ; - - - , 16×64 and 8×32) are compared along each row according to the legend. The contour levels correspond to 10%, 50% and 90% of the maximum PSD level in the original 64×64 map

moderate for 16×16 . The fact that both the 8×8 λ_x^+ and λ_z^+ minimum wavelengths penetrate within the contoured region (10% threshold) of the 64×64 map truly constraints the original spectral distribution for this case. The preservation of the τ_{w_z} maps is worse, mainly because the distribution of its characteristic scales reveals shorter and thinner patterns than for the other quantities. The downsampling has an effect on the minimum wavelength sampled. For what concerns the case with 32×32 resolution the effect is mainly felt in the streamwise direction. For the 16×16 resolution, the minimum wavelengths penetrate within the contoured region both in the streamwise and the spanwise direction. This effect is of course stronger for the 8×8 resolution. Beyond this limitation, the contoured lines within their respective limits look distorted to a greater extent than for the other two quantities. This discussion suggests that the penalty of these downsampled sensor arrangements affects the field reconstructions more

significantly with regard to the role of τ_{w_z} .

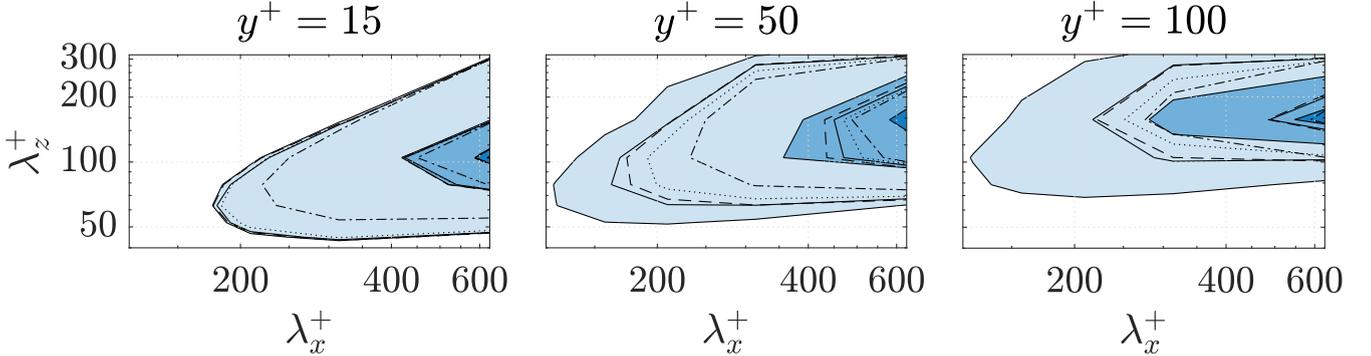


Fig. 8 Premultiplied PSD maps of the streamwise velocity component u^+ at three different wall-normal distances. The estimations from different number of sensors (contoured lines: —, 64×64 ; - - - -, 32×32 ; ·····, 16×16 ; - · - ·, 8×8) are compared with the spectrum of the DNS (coloured contour: —). The contour levels correspond to 10%, 50% and 90% of the maximum PSD level in the DNS u^+ field

Moreover, the spectral length-scales of the streamwise velocity component that the 3D-GAN is able to reconstruct from the different input sensor arrangements are reported in Fig. 8. In the different wall-parallel sections of the channel analyzed, the dominant patterns present in the original DNS field are as elongated as πh , with a width of about 100 – 150 wall inner units. At a distance of 15 wall units away from the wall, the spectra look quite well preserved, with a certain attenuation for the 8×8 arrangement. The spectrum for 16×16 is nearly coincident with both finer resolutions, even if its MSE (u^+) at this distance from the wall is substantially larger (see Fig. 6). At $y^+ = 50$ the spectral maps have a significant additional distortion, with an important part of the spectral power not recovered: no point in the map gets the 90% threshold even at the 64×64 resolution; this effect is more intense for coarser wall data. The attenuation is such that the 64×64 map loses almost 50% of the spectral power contained in the DNS data. The attenuation observed at $y^+ = 100$ is even strengthened, leading to a loss of more than 70%. A more detailed quantification of this attenuation is reported in Table 5. Being S the premultiplied power spectral density map of the predicted field, and \hat{S} that of the original field, the overall attenuation is $1 - \frac{\Sigma(S)}{\Sigma(\hat{S})}$, and the highest spectral intensity attenuation is $1 - \frac{\max(S)}{\max(\hat{S})}$. These results show the robustness of the 3D-GAN when using significantly fewer sensors than in the baseline configuration. The attenuation at 8×8 with respect to the full resolution is evident, but the reconstruction is certainly satisfactory if one considers that it employs only 1.5% the number of sensors in the 64×64 arrangement. It must be remarked, however, that for regions further from the wall the reconstruction error is mainly driven by the lower correlation of velocity fluctuations with wall quantities, thus the effect of reducing the

number of sensors is much less relevant.

Table 5 Attenuation of the u^+ field spectrum with respect to the DNS at $y^+ = [15, 50, 100]$.

Overall attenuation				Highest spectral intensity attenuation				
64×64	32×32	16×16	8×8		64×64	32×32	16×16	8×8
5.8%	4.3%	11.3%	39.6%	$y^+ = 15$	0.0%	0.0%	0.3%	8.2%
46.0%	47.0%	57.9%	70.3%	$y^+ = 50$	18.3%	11.8%	22.3%	24.7%
71.5%	72.2%	81.7%	83.0%	$y^+ = 100$	33.4%	30.7%	54.4%	44.9%

Additionally, the turbulent structures have been analyzed. Individual Q-events have been identified with the 3D-quadrant analysis [249] based on the planar quadrant analysis for the identification of turbulent structures [250], [251]. Structures are defined according to a hyperbolic hole size equal to $H=1.75$ following:

$$| -u(x, y, z)v(x, y, z) | > Hu'(y)v'(y) , \quad (5)$$

where $u'(y)$ and $v'(y)$ represent the root-mean-squared of each velocity component at a given wall distance y . These events have been bounded by their range in the wall-normal direction, being y_{min} and y_{max} the minimum and maximum distance from the wall of the points forming a structure, respectively. Wall-attached structures (setting the threshold at $y_{min}^+ \leq 20$) represent the vast majority of the population of turbulent structures. For each wall-attached structure in the DNS dataset, the proportion of its volume \mathcal{V} that is overlapped with analogous Q-events from the corresponding estimated (3D-GAN) field has been computed.

Structures have been classified in bins according to their y_{max} . For each bin, the average volume matching proportion X_t of all its structures has been computed as the average of the average overlap proportion \mathcal{V} of all these structures [265]. Figure 9 shows two main effects regarding X_t . Overall, wall-attached structures are preserved better the closer to the wall they remain, while those that extend towards the center of the channel have more chances to lose a bigger part of them in the reconstruction process. Besides, it shows the trend due to the effect of the number of sensors used. In line with some previous comments, the quality of the estimation becomes worse as fewer sensors are employed for the whole y_{max}/h span. However, the trend is not uniform, with the configuration with 32×32 sensors suffering a small penalty with respect to 64×64 , while Q-events generated with cases with further downsampling report more significant losses in terms of X_t . This is probably due to the aforementioned effect of the downsampling on the wall spectra and the length-scales that each sensor set is able to measure.

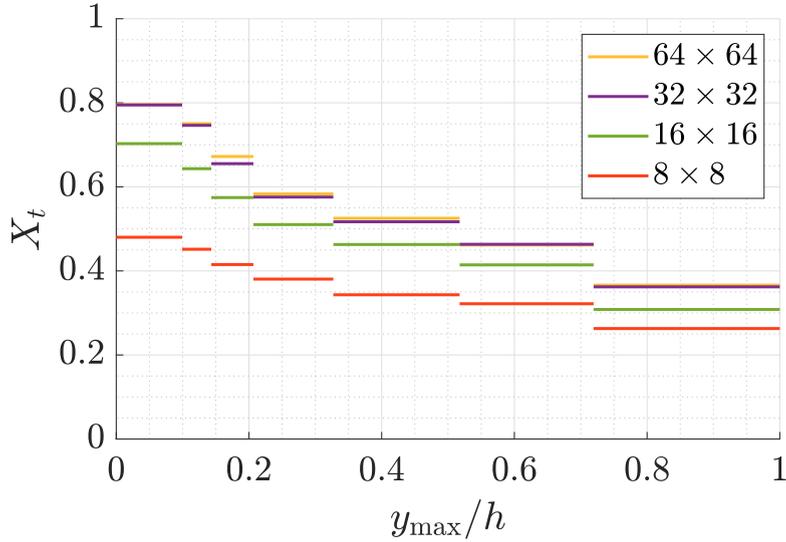


Fig. 9 Evolution of the overall matching proportion X_t of the wall-attached structures in the DNS fields with those from the generated fields. Structures are categorized into different groups according to their span y_{max}/h .

3.2 Estimation with different downsampling factor in x and z

Alternative cases were proposed applying different downsampling factors along directions x and z (see Fig. 4). The MSE of the two new cases can be observed in Fig. 10, where they are compared with their respective cases with the same number of sensors with equal downsampling factors in x and z . As seen in Table 1, the 16×64 arrangement has the same number of sensors as the 32×32 arrangement, while the same holds for the cases with 8×32 and 16×16 sensors.

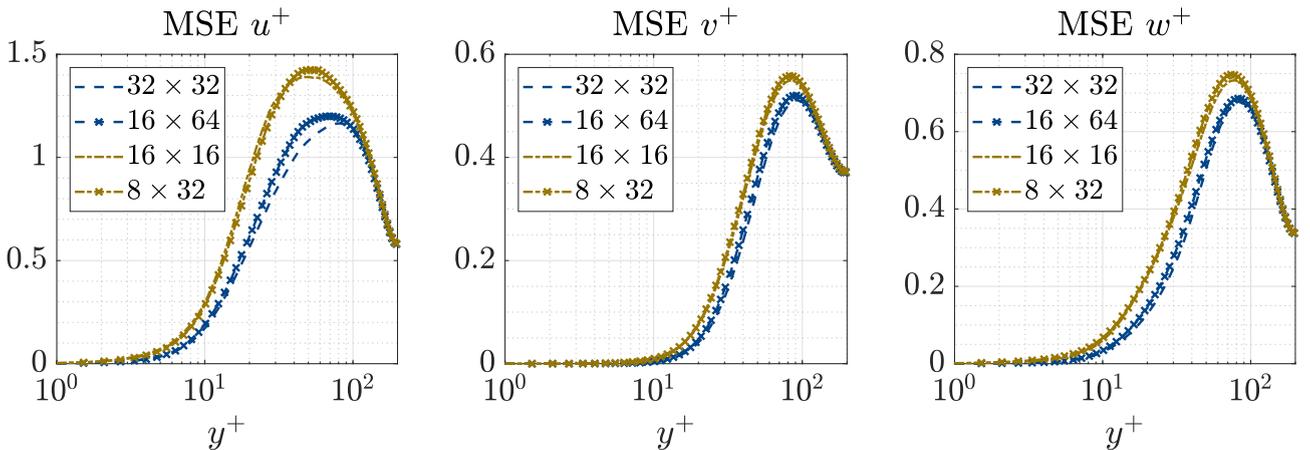


Fig. 10 MSE comparison for each velocity component as a function of the wall-normal distance for cases with the same number of sensors (---, 32×32 ; \times ---, 16×64 ; ---, 16×16 ; \times ---, 8×32). The results are expressed in wall-inner units based on the friction velocity u_τ .

The difference in error between these two pairs of cases is quite small. It is seen how the sensor arrangements with the same downsampling factor in x and z indeed work slightly better

than those with different downsampling factors, against the hypothesis introduced in §2—except for the estimation of u^+ with $2^{10} \cdot 33$ sensors in the range of $y^+ \approx [15 - 30]$. As in the DNS grid, all cases with the same downsampling factor already have twice as much resolution in z as in x . On the other hand, in the alternative arrangements, the resolution in z is 8 times as in x . There might be a ratio between the resolution along both directions for which the error is minimized. Although a ratio of 8 might seem a bit extreme for this optimum scenario, the penalty in terms of error is not very significant. Besides, given the quantization of this additional penalty as reported in Fig. 10, under the scenario of a physical implementation of this technology, these alternative arrangements could even be considered if they were preferred for any practical purpose against arrangements with same downsampling factor, at the cost of a slightly poorer estimation performance. Similar conclusions could be drawn from the integral error ε (see Table 3), where both cases report a very moderate increment with respect to the cases with the same number of sensors.

The spectrum of the three wall measurement quantities for these two pairs of cases can be compared in the central and bottom rows of Fig. 7. In both p_w spectra, the distortion of the contoured lines of the alternative arrangements with respect to those with same downsampling factor is not very significant. However, the downsampling factor of 8 applied on x in 8×32 highlights spectral losses of streamwise scales—as previously reported in the 8×8 arrangement. Regarding the τ_{w_x} , quantity for which scales are particularly elongated, none of these arrangements seems to be significantly penalized and the contoured lines are quite overlapped. Instead, the preservation of the scales in the τ_{w_z} map is not as good, as this spectrum contains smaller scales that might not be bounded by the minimum wavelengths. The increment in resolution in z in these two alternative arrangements with respect to those with same downsampling factor enables a reduction of the loss of the small τ_{w_z} spanwise scales previously reported in Sect. 3.1. Indeed, in the 16×64 case, the increment in the spanwise resolution with respect to 32×32 , allows us to partially recover a peak in the spectrum around $\lambda_x^+ \approx 200$ and $\lambda_z^+ \approx 50$ that was only reported in the original 64×64 spectrum map. Nevertheless, this is possible at the cost of losing further small streamwise patterns. In relation to the retained variance reported in Table 1, it is seen that the retained variance of the patterns of τ_{w_x} , particularly elongated, is better conserved with 16×64 and 8×32 than on their respective cases 32×32 and 16×16 . However, this difference is much smaller than the loss reported for the retained variance of p_w .

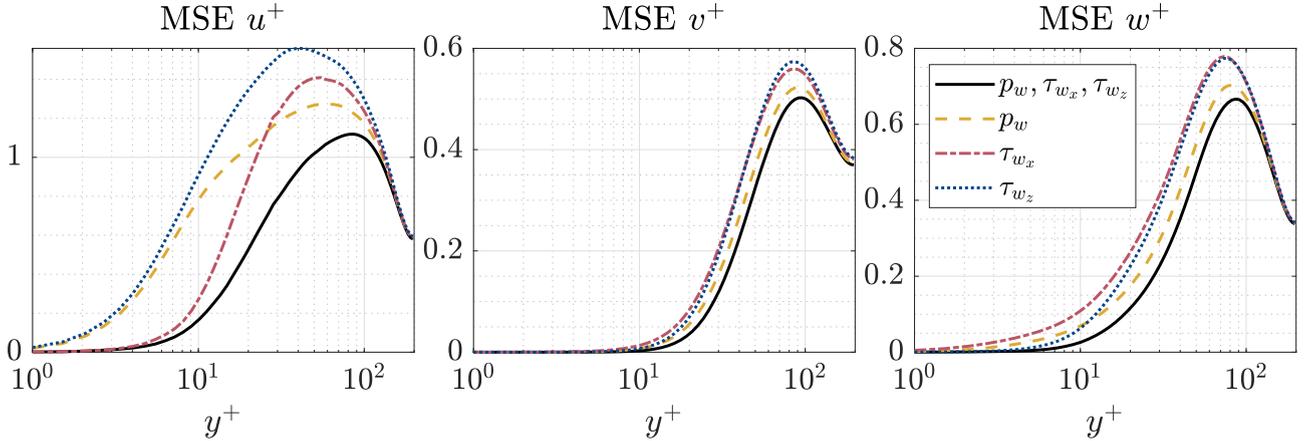


Fig. 11 MSE comparison of each velocity component between cases using only one type of sensor with the 64×64 resolution: —, $p_w, \tau_{w_x}, \tau_{w_z}$ (baseline); - - -, p_w ; - · - ·, τ_{w_x} ; ·····, τ_{w_z} .

3.3 Reconstruction with one type of sensor

Often, in realistic scenarios only one of the wall quantities might be accessible. In this section, we analyse the effect of reducing the availability of the wall-measured quantities. Figure 11 shows a comparison of the MSE reported when the 3D-GAN is fed only with pressure, streamwise or spanwise wall shear stress, arranged in the 64×64 resolution. The error of the baseline case with the three types of sensors is shown here for reference.

These results show that the relative importance of each of these three wall quantities for the flow field reconstruction is not the same. Besides, their relation with each velocity component is different. Close to the wall, the streamwise and spanwise velocity components show a stronger relation with the streamwise and spanwise wall shear stresses, respectively. However, the MSE corresponding to the wall pressure becomes lower beyond $y^+ \approx 25$ for u^+ and beyond $y^+ \approx 10$ for w^+ . Instead, p_w shows a slightly better capability to estimate v^+ throughout the domain.

A similar procedure is followed for the downsampled resolutions proposed in Fig. 3. To simplify it, we compare the MSE of the fluctuation velocity \mathbf{u}^+ , defined as the sum of the MSE of u^+ , v^+ , and w^+ , instead of for each of its components separately. The results for the four cases are plotted in Fig. 12.

For the 64×64 resolution, it is observed that the pressure sensors offer the best estimation performance throughout most of the domain. Only in the vicinity of the wall, up to $y^+ \approx 15$, τ_{w_x} provides a slightly smaller MSE. This finding is similar to that for the MSE of u^+ shown in Fig. 11. As the MSE of u^+ for the case using only τ_{w_x} reports a clear advantage close to the wall against the other two quantities, its influence prevails (64×64 in Fig. 12) when the

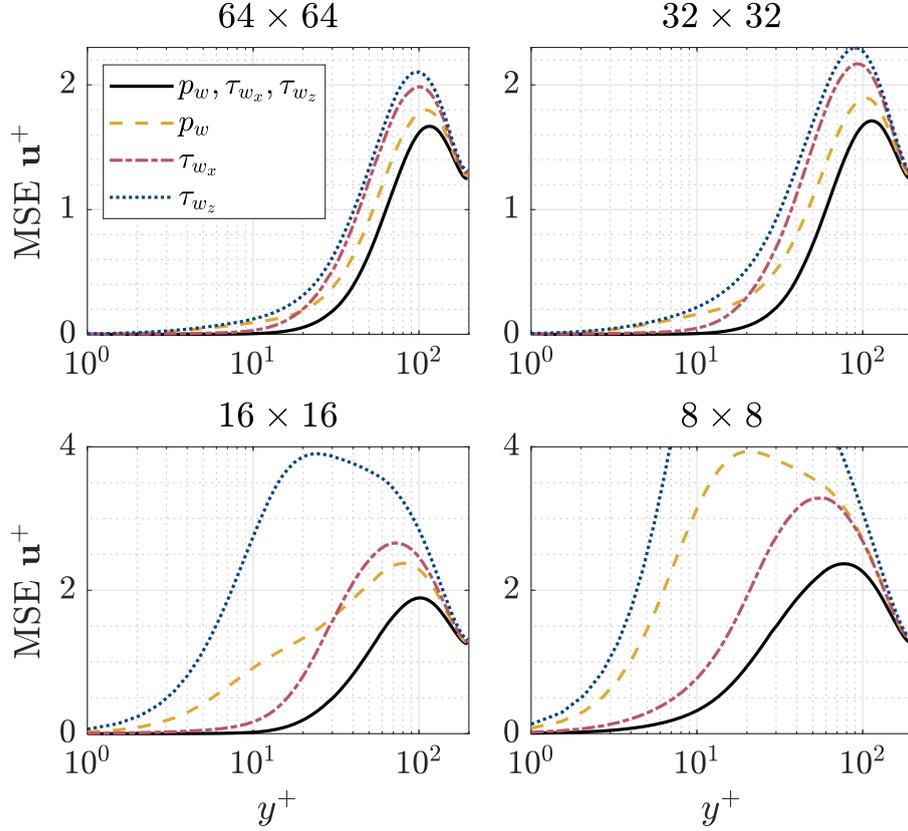


Fig. 12 MSE comparison between cases using only one type of sensor with different resolutions: —, $p_w, \tau_{w_x}, \tau_{w_z}$ (baseline); ----, p_w ; - - - - , τ_{w_x} ; ·····, τ_{w_z} .

MSE of the other velocity components are included (cases in which τ_{w_x} does not provide the lowest MSE but the differences between cases employing different wall-measured quantities are less significant).

No significant changes are reported for the 32×32 resolution, apart from a slight increase in the magnitude of the MSE. This effect was also shown in Fig. 6, where the three wall quantities were employed. However, differences are found for the cases with the lowest number of sensors. Having a limited number of sensors appear to be particularly detrimental for estimation using solely pressure. The estimation based on τ_{w_x} , on the other hand, is still reasonably good in the near-wall region. With 16×16 sensors, the pressure is not able to provide a lower MSE than τ_{w_x} until beyond $y^+ = 30$. With 8×8 sensors, τ_{w_x} provides the lower MSE throughout the way along from the wall to the mid-plane of the channel. As observed qualitatively in Fig. 3 and quantitatively in figure Fig. 7, τ_{w_x} exhibits large-scale patterns, which might explain why this quantity can provide us with better field estimations under a coarse grid arrangement when only one type of sensor is used.

The MSE of τ_{w_z} for these two coarser resolutions has very significant increments with respect to the two finer resolutions, even if with 64×64 sensors it already provided the worst estimation. This might be due to the fact that the patterns in τ_{w_z} are smaller, while the downsampling procedure preserves better the patterns of the other two quantities.

3.4 Effect of sensor noise

This section covers an analysis consisting on the quantification of the flow field estimation MSE when noise is applied over the three wall-sensed quantities.

Figure 13 shows the MSE profiles of the three velocity components (columns) for each resolution considered (rows), comparing this metric for the values of $\delta = [0\%, 1\%, 3\%]$. As expected, the MSE of the reconstructed fields increases with the noise added to the sensors. The curves follow a similar profile, with very low values in the vicinity of the wall, progressively increasing with the distance from the wall up to a maximum, and then reduced and stabilized towards the centre of the channel. For the three values of δ considered, the MSE levels are very similar in the closest and in the farthest regions to the wall, while the differences in the region in between these two are more evident. All this discussion applies to the three velocity components, although the differences in MSE for u^+ are more significant.

In relation with the input resolution, the main conclusion is the fact that the effect of δ is minimized with the downsampling factor. For the sensor arrangement 32×322 few differences are reported with respect to 64×64 ; the shifts between the curves of each δ are still very significant, and the level of error with 32×32 sensors is slightly higher than for 64×64 . With further downsampling, 16×16 sensors shows more moderate increments of the MSE curves with increasing δ . This effect is even strengthened at 8×8 , where the curves of each level of δ are nearly overlapped throughout the domain of the channel flow. In other words, all this discussion implies that, as the sensor arrangement becomes coarser, the MSE grows faster for $\delta = 0\%$ than for $\delta = 1\%$, and even more than for $\delta = 3\%$. As a remarkable example, comparing curves with $\delta = 3\%$ (more evident for u^+), the MSE for 16×16 is higher than for 32×32 and 64×64 only close to the wall and close to the centre of the channel, but in the rest of the channel and in particular, in the region of the peak, the MSE is significantly lower for 16×16 sensors. The same trends with the sensor resolution are reported in Table 3 according to the integral error ε .

4 Conclusion

The 3D-GAN developed in Ref. [265] has been tested under coarser grids of input sensors at the wall, willing to estimate the flow of half channel at the same resolution as in the original grid. This is a challenge for the methodology, allowing us to assess its performance when less data are available. This analysis quantifies the additional uncertainty introduced when fewer sensors are installed, which would definitely be a trade-off between simplicity and accuracy to analyze under a practical implementation scenario. The robustness of the 3D-GAN under these scenarios is demonstrated by the fact that the penalization reported under coarser wall measurements is not directly proportional to the reduction in the number of sensors employed, but rather depends on their capability to represent wall patterns.

When a progressive downsampling to the sensor grid is applied in both streamwise and spanwise directions in powers of 2, the three velocity components experience an increasing MSE metric over the whole half channel. Nevertheless, the error introduced with each downsampling step is not linearly dependent on the number of sensors. The quality of the reconstruction seems to be directly affected by the characteristic length-scales of the measured quantities, and which of them remain present with each of the sensor arrangements. The MSE for a wall input of 32×32 is almost equivalent to that of 64×64 sensors—the DNS resolution. Nevertheless, the smaller scales lost with any further downsampling seem to affect the accuracy of the estimator. The largest scales that remain in the wall after the coarsening process still carry very important information for the flow reconstruction. Although the attenuation experienced in the estimated fields becomes relevant, the largest scales present in the channel flow are still sensed.

The aforementioned cases are compared as they are downsampled in the same way along x and z . A comparison between this sensor arrangement setting and a different one with different downsampling in the two directions is established. The main hypothesis was that the structures are mostly elongated in the streamwise direction, thus it is more relevant to sample them properly in the spanwise direction. The results show, however, that no significant improvements are achieved with this alternative setting.

The reconstruction employing only one type of sensor revealed that each of the wall quantities considered does not contribute in the same way to the reconstruction of each velocity component. Moreover, the implications due to the lack of availability of sensor resolution are addressed. In general, pressure measurements would provide better flow reconstructions than the wall shear stresses if data could be sampled at sufficiently high spatial resolution. However,

this choice might bring an important penalization on the estimation of u^+ close to the wall, where τ_{w_x} performs remarkably better. Instead, when the resolution of the sensors is lower, τ_{w_x} is preferred. We address this to its wall patterns being significantly larger, thus a higher proportion of its energetic spectrum is preserved in the downsampled data.

Finally, the robustness of the 3D-GAN is assessed concerning the noise present in the sensing process at the wall, which would be present in a practical implementation of this methodology. As the amount of noise considered in the sensors increases, the accuracy of predictions is reduced. The reduction in accuracy is not the same for all the sensor arrangements considered. Finer resolutions are more affected by the noise introduced, while the 3D-GAN is more robust to noise under coarser wall inputs.

Author contribution All authors contributed to the study conception and design. Conceptualization: A.C., A.I., S.D.; Methodology: A.C., A.I., S.D.; Software: A.C.; Formal analysis: A.C.; Investigation: A.C.; Resources: A.I., S.D.; Data curation: A.C.; Writing - original draft preparation: A.C.; Writing - review and editing: A.C., A.I., S.D.; Visualization: A.C.; Supervision: A.I., S.D.; Project Administration: A.C., A.I., S.D.; Funding acquisition: A.C., S.D. All authors read and approved the final manuscript.

Funding A.C. acknowledges funding from the Spanish Ministry of Universities under the Formación de Profesorado Universitario (FPU) 2020 program. The project received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No 949085). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

Data and code availability Data related to this work and trained models of the neural networks are openly available at <https://doi.org/10.5281/zenodo.13587745>. Codes developed in this work are openly available at <https://github.com/erc-nextflow/3D-GAN-limited-sensors>.

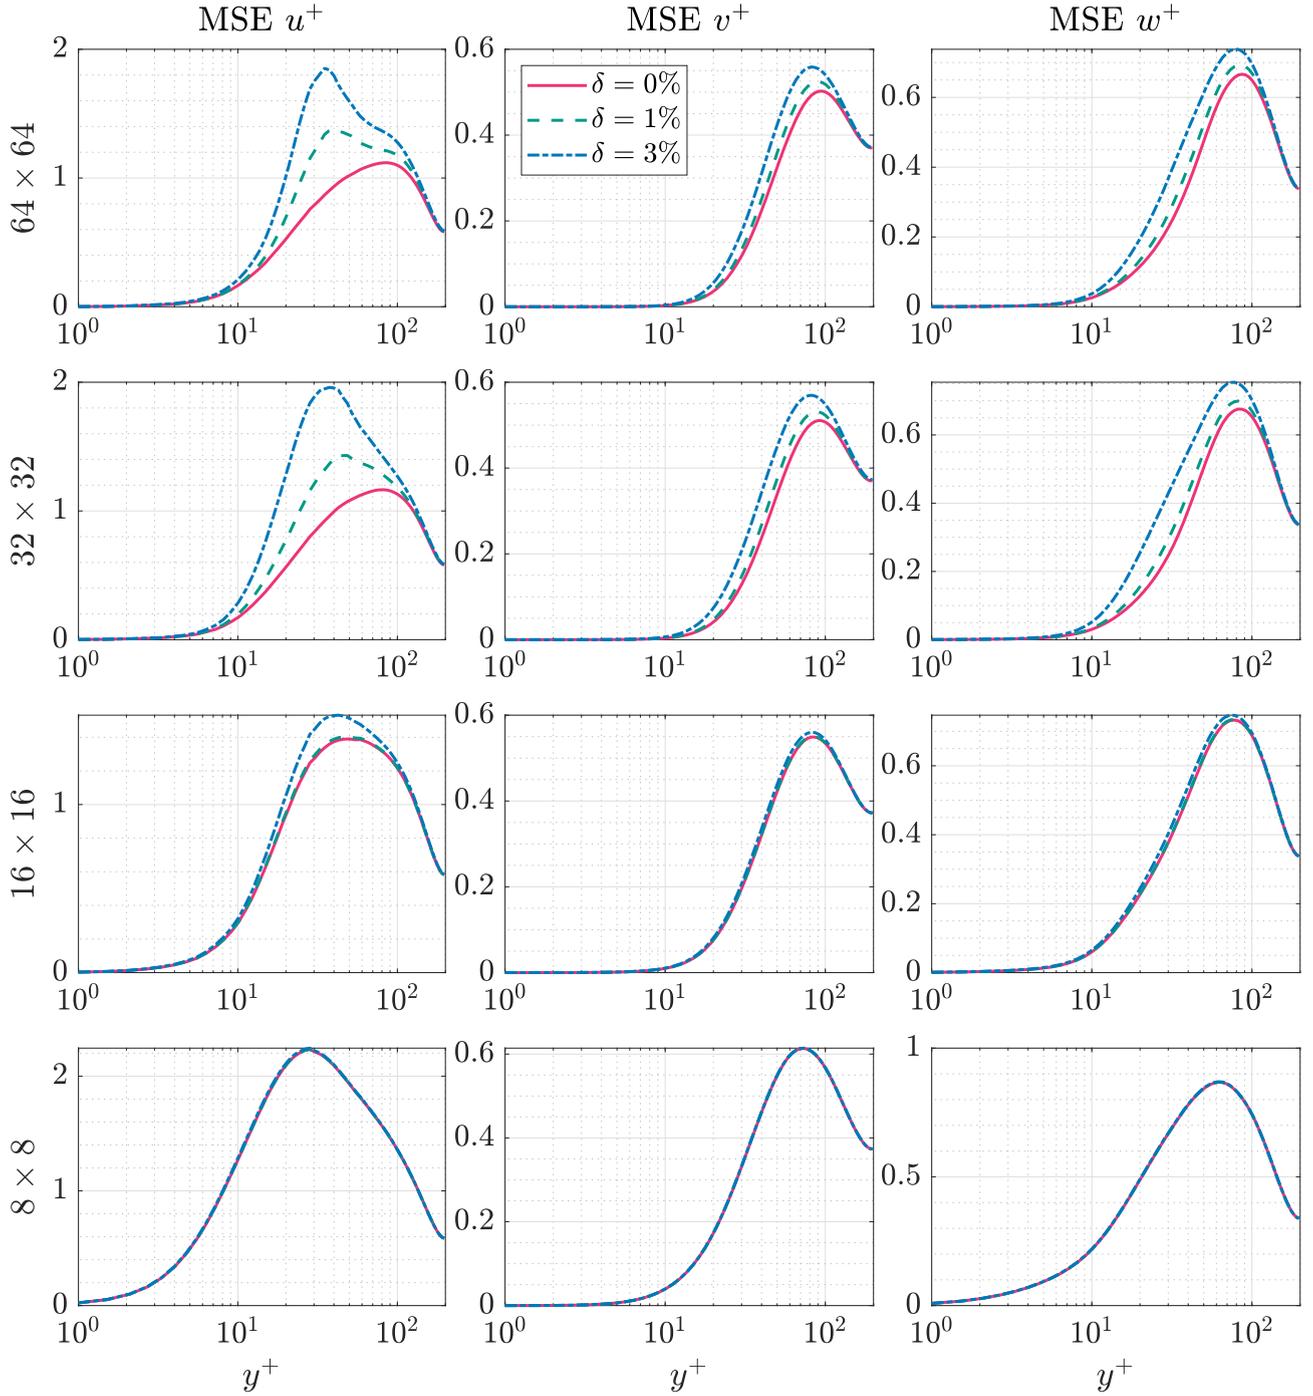


Fig. 13 MSE comparison between cases with different noise levels in the wall input —, $\delta = 0\%$; ---, $\delta = 1\%$; -·-·-, $\delta = 3\%$; for each velocity component and resolution case.

Paper 3

Measuring time-resolved heat transfer fluctuations on a heated-thin foil in a turbulent channel airflow

Antonio Cuéllar, Enrico Amico, Jacopo Serpieri, Gioacchino Cafiero, Woutijn J. Baars, Stefano Discetti and Andrea Ianiro, Under review (2024), [Preprint arXiv: 2410.12778 \[physics.flu-dyn\]](#)

We present an experimental setup to perform time-resolved convective heat transfer measurements in a turbulent channel flow with air as the working fluid. We employ a heated thin foil coupled with high-speed infrared thermography. The measurement technique is challenged by the thermal inertia of the foil, the high frequency of turbulent fluctuations, and the measurement noise of the infrared camera. We discuss in detail the advantages and drawbacks of all the design choices that were made, thereby providing a successful implementation strategy to obtain high-quality data. This experimental approach could be useful for experimental studies employing wall-based measurements of turbulence, such as flow control applications in wall-bounded turbulence.

Keywords Turbulent channel flow, Wall-based sensor, Infrared thermography

1. Introduction

In this work, we present an experimental setup for the acquisition of time-resolved measurements of the convective heat transfer coefficient on the wall of a turbulent channel. Measurements are based on the unsteady heated thin foil sensor [218], [219] coupled with infrared (IR) thermography as a temperature transducer. This technical design note responds to a challenge of measurements in air flows. Due to the foil thermal inertia and the characteristic frequencies of the problem, the measurement of the small amplitude temperature fluctuations requires careful image processing, including spatial, temporal and feature-based filtering. We critically discuss all the design choices made, together with the details of image processing to obtain high-quality instantaneous measurements of the Stanton number in air with this technique.

IR thermography has been applied in different experimental setups to study heat transfer by a wall-bounded grazing flow. Following the seminal work by Hetsroni and Rozenblit [215], Gurka et al. [216] conducted synchronised measurements of particle image velocimetry (PIV) and heat transfer on a hot foil applied to a turbulent wall-bounded water flow. A heated thin foil was employed on a similar boundary layer problem in water with high-frequency IR measurements in synchronization with PIV [221]. Similar experiments with air as the working fluid are challenged by lower fluctuations of the foil temperature and heat transfer, and higher frequency content of said fluctuations. Both of these aspects require a push towards thermally thinner and tailored processing of the images. The first challenge is merely technological. On the other hand, image conditioning requires careful consideration of the choice of the processing strategy, considering that the signal-to-noise ratio can easily be below 1.

The heated thin foil sensor is based on measuring the convection between the fluid and a heated foil through an energy balance (equation 1), as sketched in figure 1. It assesses the temporal temperature variation as an unsteady term with the contribution of the different instantaneous heat fluxes and sources, where c_p is the specific heat capacity, ρ is its density and a is its thickness.

$$c_p \rho a \frac{\partial T_w}{\partial t} = \phi_J'' - \phi_{\text{cond}}'' - 2\phi_{\text{rad}}'' - \phi_{\text{conv,ext}}'' - h_c(T_w - T_{\text{aw}}) . \quad (1)$$

The symbol ($''$) denotes that heat fluxes are expressed per unit area of the thin foil. These terms include the heat input provided to the sensor (in most cases produced by the Joule effect) ϕ_J'' , the conduction within the foil ϕ_{cond}'' , the radiation ϕ_{rad}'' emitted from both sides of the foil, and the convection, to be treated separately for the side exposed to the flow, equal to $h_c(T_w - T_{\text{aw}})$

and for the side not exposed to the flow $\phi''_{\text{conv,ext}}$, where natural convection may develop. The convective heat flux between the flow and the wall depends on the convection heat transfer coefficient h_c , which is the quantity to be determined, and on the local wall temperature T_w and the adiabatic wall temperature T_{aw} , which are both measured with the IR camera. All the terms in equation [1](#) are typically modelled following conduction and radiation laws, along with empirical correlations (see, e.g. those reported in [275](#)).

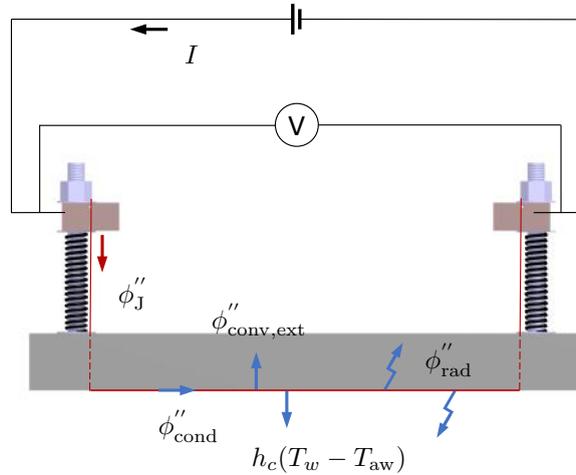


Figure 1. Sketch of the electric circuit mounted on the sensor and the heat fluxes on the energy balance of the thin foil.

The thermal model relies on the assumption that the foil is thermally thin, such that the temperature on the internal side of the foil (considered in the model) is equal to that on the external side (measured temperature). This approximation is valid if the Biot number, $Bi = \frac{h_c a}{\kappa} \ll 1$ (where κ is the foil thermal conductivity). Additionally, the Fourier number $Fo = \frac{\alpha t_{\text{char}}}{a^2}$ (with t_{char} being the characteristic time of the problem) compares the heat flux and the rate of thermal energy storage, requiring $Fo \gg 1$ to perform unsteady heat transfer measurements [222](#). Both requirements are met in the case of the experimental setup presented in our current work, by ensuring the film thickness a is small enough.

This experimental approach presents a series of challenges that must be carefully managed to ensure accurate results. Provided that the available IR camera is capable of an acquisition frequency large enough to temporally resolve the flow scales involved in the problem, several difficulties are still to be faced. First of all, the IR camera must be sensitive enough to detect the wall temperature fluctuations. While most of the terms are mainly steady, the temporal variations of h_c lead to temperature fluctuations of the foil, which are damped by the foil thermal

inertia on the left-hand side of equation (1). Secondly, the fluctuations to be measured must be compared to the Noise Equivalent Temperature Difference (NETD) of the camera to evaluate the IR sensor suitability [225]. The characteristic time—or the characteristic frequency—of the problem is a property of the flow and can be estimated. For instance, from dimensional analysis, it can be quantified as the ratio of the characteristic length and the characteristic velocity. Given such characteristic time, one may want to increase temperature fluctuations in case they are too small to be detected. Therefore, from equation 1 it is possible to argue that one can either increase the foil heating or choose a foil thin enough to amplify the temperature fluctuations. The thickness of the foil is a critical factor, however, it can not be chosen *ad libitum* as the foil needs to be manufacturable, robust enough to be implemented and sustain pressure and shear fluctuations without deforming, and ideally be cost-effective. In addition, the heating applied must be small enough to avoid perturbing the flow with undesired buoyancy forcing. Excessive heating could increase the Richardson number (ratio of natural to forced convection), potentially disturbing the results. Moreover, for IR measurements, a layer of paint is often applied to enhance the accuracy of temperature measurements by increasing the surface emissivity. This alters the thermal properties of the material and thus may distort the measurements if not properly accounted for [276]. A proper balance between these factors is essential for reliable outcomes.

Different models have been developed to address these complexities [218], [222], [276], [277]. However, for measurements in air flows, the temperature fluctuations are often of the same order (or smaller) than the NETD, thus making image preprocessing crucial. Techniques like Proper Orthogonal Decomposition (POD) or autoencoders can suppress noise in very low signal-to-noise ratio conditions [220], [278]. Additional techniques, such as bad pixel removal to eliminate outliers and detrending procedures to counteract increasing mean temperatures in the system, further improve data quality [221].

The remainder of this article contains a description of the experimental setup in section 2, a discussion of the results in section 3, the uncertainty is quantified in section 4 and the key findings are summarized in section 5.

2. Experimental setup

2.1 Channel Flow and Experimental Setup

The experiments of this work were performed in the Channel Flow facility of the aerodynamics laboratory at Politecnico di Torino. The walls of the channel are made of poly-methyl-methacrylate, defining an internal cross-section of $35 \times 420 \text{ mm}^2$ with a length of 10 m. The channel operates with air in an open-return wind tunnel configuration. The inlet section is composed of an electric pump followed by a small divergent section, a settling chamber and a convergent nozzle upstream of the channel entrance. To trigger transition to turbulence, two thin turbulator strips with zig-zag-shaped leading edges are placed at the end of the convergent on both channel walls along the spanwise direction. To ensure flow stability, a slightly divergent duct is installed at the outlet of the channel. The channel walls are made of modular sections which allow a custom layout to install suitable devices needed to carry out the experimental campaigns. A schematic representation of the channel is reported in figure 2.

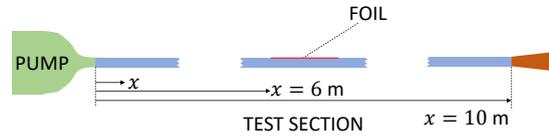


Figure 2. Sketch of the channel flow facility at Politecnico di Torino, with the thin foil (red) mounted on the top wall of the channel.

The experiments reported in this technical note have been conducted in a turbulent regime with bulk airspeed $U_\infty = 4.95 \text{ m/s}$. The Reynolds number based on the outer quantities is equal to $Re = \frac{U_\infty h}{\nu} = 5800$, where h is the half channel height, equal to 17.5 mm and ν the air kinematic viscosity. Accordingly, the friction-based Reynolds number $Re_\tau = \frac{u_\tau h}{\nu} = 220$, with $u_\tau = \sqrt{\frac{\tau_w}{\rho}}$ the friction velocity defined as the square root of the ratio between the wall skin friction τ_w and the fluid density ρ . The friction velocity is characterised by employing 16 static pressure ports distributed along the channel test section to quantify the pressure gradient. The static pressure measurements are carried out using a 16-channel DSA pressure scanner, with a maximum pressure range of 2500 Pa and a 0.05% full-scale accuracy. A linear fitting is applied to obtain the pressure gradient, which is then employed to determine the value of the wall friction [279].

2.2 Heat transfer measurements

A heated-thin-foil heat transfer sensor has been designed and mounted on the upper wall of the channel (see figure 3) by means of a frame embedded in the top wall, coincident with the position of the modular wall it replaces. This design choice minimizes the flow disturbance caused by buoyancy effects. Indeed, the intensity of natural convection for a heated plate facing downward is lower than when facing upward [280].

The sensor frame was 3D printed with PLA material. The thin foil is inserted on the frame through two grooves at its leading and trailing edge positions, 150 mm apart. The sensing area has a width $W = 100$ mm and it spans a length $L = 150$ mm along the channel wall. On the external part of the frame, four fixed bars are mounted, holding two copper block pairs, one for each end of the thin foil. Each pair of copper blocks forms a clamp that retains the thin foil in between. The thin foil is heated through Joule effect. For this purpose, a DC power supply is connected to the copper blocks. The higher electrical resistance of the CrNi-Steel alloy with respect to that of the copper, together with the small thickness of the thin foil ($5 \mu\text{m}$) with respect to that of the copper blocks (1 cm) ensures that the electric potential drop through the copper is negligible with respect to that through the thin foil. The copper blocks can thus be considered at practically constant voltage, ensuring uniform voltage and current in the spanwise direction. To minimize contact resistance between the copper blocks and the foil, a thin engraving in the copper block face in contact with the foil is filled with a 1 mm indium wire. The voltage differential applied to the foil is thus assumed to be equal to that between the copper blocks at the foil edges and it is measured with a voltmeter in contact with the copper blocks. The different parameters of the problem employed for this model are collected in table 1. Voltage and current are not included as several power levels are considered in this work.

To ensure foil tension, a spring is mounted around each bar, pushing apart the sensor frame and the copper blocks, thus tightening the thin foil. The frame adjusts the thin foil, subjecting it to tensile stresses to keep it flat and prevent any misalignment with the rest of the channel wall. The thin foil is made of 1.4310 CrNi-Steel alloy (Cr 16–18% and Ni 7–9%) able to withstand a nominal tensile stress $F > 1500$ N/mm². Kapton[®] tape is used for sealing around the sensor and avoiding any air leakage perturbing the flow. To assess a suitable foil thickness two foils are tested, with $a_f = 5 \mu\text{m}$ and $10 \mu\text{m}$, respectively.

An IR camera is employed as a temperature transducer to measure the temperature on the

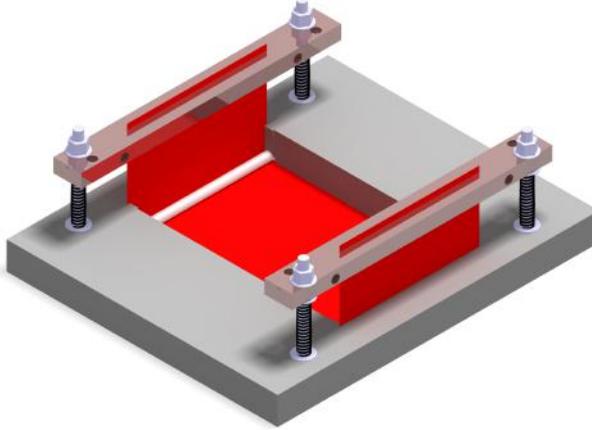


Figure 3. Isometric view of the external side of the sensor, with the thin foil (red) mounted on the frame.

Table 1. Values used in the thermal model of the sensor

Quantity	Symbol	Value	Units
Heat capacity (foil)	$c_{p,f}$	500 ± 5	J/(kg K)
Heat capacity (paint)	$c_{p,p}$	5000 ± 50	J/(kg K)
Density (foil)	ρ_f	7900 ± 50	kg/m ³
Density (paint)	ρ_p	1300 ± 50	kg/m ³
Therm. cond. (foil)	κ_f	17 ± 2	W/(m K)
Therm. cond. (paint)	κ_p	1.4 ± 0.1	W/(m K)
Area (Joule)	A^*	0.0254 ± 0.001	m ²
Foil length	L	0.15 ± 0.0005	m
Foil width	W	0.1 ± 0.0005	m
Emissivity (paint)	ϵ	0.95 ± 0.02	–
Ambient temperature	T_{amb}	19	°C

external side of this sensor. To enhance IR temperature measurements, given the low emissivity of steel, the external side of the thin foil is sprayed with high-emissivity matt black paint.

The IR camera used in this experiment as temperature detector is an Infratec Camera ImageIR[®] 6300Z, with a resolution of 640×512 pixels. The temperature resolution (NETD) of the IR camera at 30°C is 0.03 K, and its calibration accuracy is ± 2 K. The camera is mounted at a distance of 30 cm from the foil, with a focal length of 18 mm. This leads to a resolution of 0.21 mm/pixel along the sensor. An integration time equal to 2900 μs was set according to the configuration requirements of the camera hardware for the range of temperatures considered for this experiment. The IR camera sampling frequency was set to $f_s = 180$ Hz to ensure sufficient temporal resolution needed. This leads to a temporal separation between snapshots of less than $1/16^{\text{th}}$ of an eddy turnover time h/u_τ . Considering a convection velocity of near-wall streaks of about $11u_\tau$ [281] and the characteristic length of the sensor L , the sampling rate is sufficiently

large so that 12 snapshots cover the convective time of the flow over the streamwise length of the sensor.

As a standard procedure of the heated-thin-foil sensor to obtain the adiabatic wall temperature, an acquisition run without power input is carried out. Then, the electrical input is turned on to take the measurements of the temperature maps of the experiment.

The different heat fluxes in the thin foil are modelled to compute the heat transfer coefficient between the foil and the channel flow from the energy balance (1). The unsteady and conductive terms have been corrected to account for the effect of the paint (282), (283), as shown in equations (2) and (4) respectively, assuming a coating thickness $a_p = 20 \mu\text{m}$ on each side of the foil and using the foil (f) and paint (p) properties:

$$c_p \rho a \frac{\partial T_w}{\partial t} = (c_{p,f} \rho_f a_f + c_{p,p} \rho_p a_p) \frac{\partial T_w}{\partial t} . \quad (2)$$

The Joule effect term is given by:

$$\phi_J'' = \frac{VI}{A^*} , \quad (3)$$

where V and I , respectively, are the voltage and the intensity of the current applied, and A^* is the area through which the current is discharged. Note that this area does not coincide with the flow exposed area of the sensor, as the entire thin foil between the two pairs of copper blocks must be considered.

The conduction heat flux experienced through the foil is computed as:

$$\phi_{\text{cond}}'' = -(\kappa_f a_f + \kappa_p a_p) \nabla^2 T_w , \quad (4)$$

where κ_f and κ_a are the thermal conductivities. The nabla operator ∇ considers the second derivatives of the temperature maps along both principal directions of the plane of the foil.

For the quantification of the radiation effect, the emissivity ϵ of the surfaces of the thin foil should be introduced in the model. As both sides are covered with matt black paint, a high value (0.95) is taken—the external side is painted to improve its emissivity for IR acquisition purposes, and the internal side is painted to avoid reflections on synchronized PIV acquisition not described in this work. Applying Stefan-Boltzmann's law, radiation is modelled as:

$$\phi_{\text{rad}}'' = \sigma \epsilon (T_w^4 - T_{\text{amb}}^4) , \quad (5)$$

being σ the Stefan-Boltzmann constant and T_{amb} the ambient temperature in the laboratory

during the experimental acquisition used as a reference for radiation.

2.3 Image preprocessing

Due to the low signal-to-noise ratio, preprocessing of the images is paramount. The following steps are implemented:

- The effect of the heating of the foil frame and of the natural convection cells rising from the external side of the thin foil was visible on the sequence of original images, characterised by large patterns (characteristic length-scale of the order of the width of the thin foil) and low-frequency temperature oscillations. Due to the low-frequency nature of natural convection, a high-pass filter can be applied, removing it from the fields. To avoid filtering any actual feature within the turbulent flow, features whose characteristic time is at least 12 times larger than one eddy turnover time (h/u_τ) have been removed, thus, with a cutoff frequency of 0.9 Hz for this case. The filter cannot suppress characteristic patterns of the channel turbulence, as with a convection velocity of near-wall streaks of about $11u_\tau$ [281], the filter would suppress structures with a characteristic length greater than at least $130 h$, which is much larger than the size of the very large scale motions expected in the channel flow. Consequently, the term $\phi''_{\text{conv,ext}}$ in (1) should be neglected for the calculation of h_c . The order of this effect can be quantified according to a horizontal heated plate facing up.
- The temperature fluctuations expected in the thin foil are of a comparable order of magnitude to the IR camera noise, necessitating additional filtering. A 3D-Gaussian filtering (in the image plane and in time) with a smoothing kernel has been applied to smooth the data. In both directions in the image plane, the size of the kernel is 2 ($\Delta x^+ \approx 6$), contributing to dampening fluctuations with the size of a single pixel. The size of the kernel in the temporal dimension is 0.5, to evidence transitions and temporal fluctuations.
- Furthermore, a feature-based filtering, analogous to that done in Ref. [220] can be applied. The temperature maps can be decomposed in POD modes and a low-order reconstruction with only the most energetic modes can be employed to remove incoherent features. The number of the modes to be retained has been identified with the elbow method [284].
- To reduce field noise, due to minor non-uniformity in the gains the of IR camera sensing elements, an additional Gaussian filtering can be applied with an elongated kernel in the spanwise direction. This measure is taken to filter out residual striped patterns due to the acquisition mechanism of the IR camera, which scans the sensor in rows.

3. Results

The experimental campaign to test this experimental configuration has been done by acquiring datasets of 12000 samples at 180 Hz, tuning the Joule effect power input to different levels to explore the balance between sufficiently strong thermal fluctuations and spurious effects due to natural convection. Following the model from section 2.2 and the data filtering recipe of section 2.3, the convective heat transfer coefficient h_c has been quantified.

3.1 Results of the thermal model

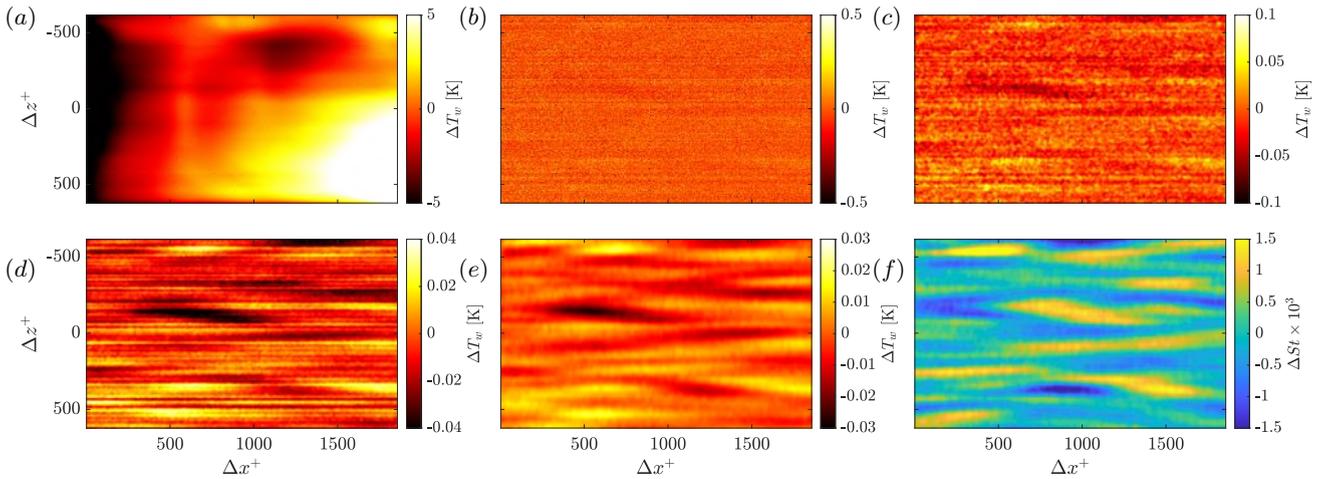


Figure 4. Maps of the thermal model at different filtering and processing stages in sequential order for an instantaneous field: (a) original, (b) high pass output, (c) 3D-Gaussian output, (d) POD output, (e) 3D-Gaussian output, (f) Stanton number map.

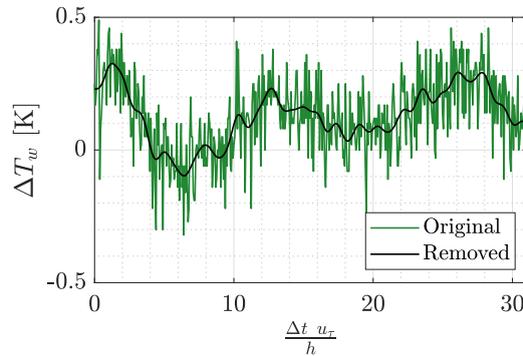


Figure 5. Original sequence of instantaneous temperature fluctuations of a pixel and removed low-frequency events.

The filtering sequence for an individual sample is depicted in figure 4. These results correspond to measurements with a foil with $5 \mu\text{m}$ thickness. The input current is set to 8.0 A with a voltage supply of 5.0 V, resulting in a foil temperature of about 35 K above room temperature. Although each filtering step may attenuate the temperature fluctuations, it is seen that the level

of noise is progressively reduced. The high-pass filter has successfully removed the convective cells and other constant or low-frequency events not related to the local channel turbulence, as shown in figure 5. However, the level of noise at this stage is still quite high. Some streamwise patterns start being observed after the 3D-Gaussian filter. For further smoothing, the POD filter cuts the noisy tail of modes with about 50% of the energy contained after the Gaussian filter. A Gaussian filter with actuation in the spanwise direction removes the remaining thin striped patterns. The resulting temperature map is used according to the energy balance in equation (1) to compute h_c , or the Stanton number $St = \frac{h_c}{\rho U_\infty c_p}$ in non-dimensional terms, employing the fluid density ρ and specific heat capacity c_p . As seen in figure 4 (f), the outcome of this procedure results in instantaneous heat transfer maps with patterns that show a relation with the wall-bounded turbulence expected in the channel. These patterns are elongated and nearly aligned in the streamwise direction, with lengths in the range $\Delta x^+ = [500 - 1000]$ and a span of $\Delta z^+ = [50 - 100]$, with the superscript $+$ indicating normalization with inner scaling. These values are similar to those reported in Ref. [255].

3.2 Effect of heating

A sensitivity analysis of the power input effect is conducted, heating the foil about 15, 25 and 35 K above room temperature. Stronger heating leads to the development of more pronounced convective cells. As a result, the magnitude removed with the high pass filter increases with the power input, as shown in the sequences in figure 6.

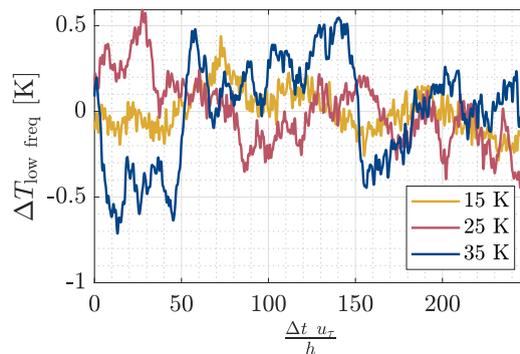


Figure 6. A sequence of low-frequency temperature fluctuations removed from a pixel with the high pass filter.

Examples of filtered T_w and St maps for these configurations are shown in figure 7. The temperature maps have been processed in the same manner for each power input level, with the POD filter threshold tailored to each case using the elbow method. Low heating levels make the temperature fluctuations less pronounced, while heating contributes to amplifying

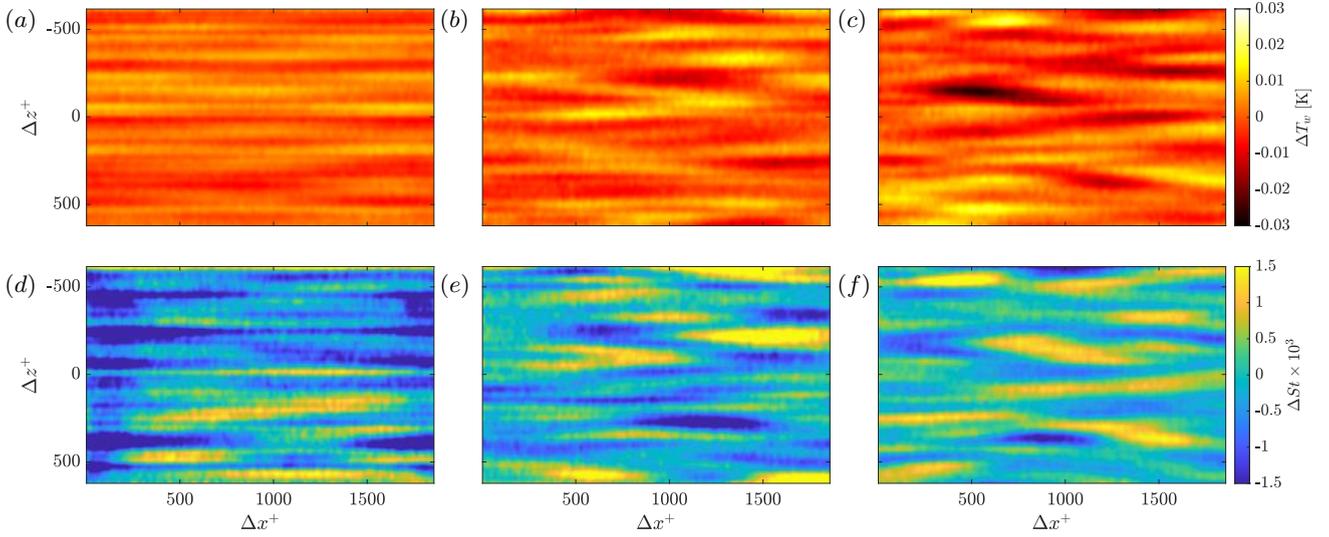


Figure 7. Instantaneous temperature map after the filtering process (top) and Stanton number map (bottom) at different heating conditions: from left to right (a) and (d) 3.2 V - 5.0 A (about 15 K), (b) and (e) 4.0 V - 6.5 A (about 25 K), (c) and (f) 5.0 V - 8.0 A (about 35 K, same as figure 4e - f).

the signal. Once filtered, the temperature sequence for the lowest heating level (as in figure 7 a) retains 53% of the variance than in the case with the highest power input (as in figure 7 c). The intermediate case retains 73% of its variance.

Beyond the differences in how pronounced those temperature fluctuations are, T_w maps show differences in their patterns that strongly influence the St patterns, as seen in the bottom row of figure 7. For the lowest heating level, the Stanton number patterns are still quite noisy and, from a qualitative inspection, seem less consistent with the physics of coherent structures in the near-wall region [281]. This sequence might necessitate further filtering, which may introduce more uncertainty. Further heating can magnify the temperature fluctuations, making the approach less susceptible to noise and facilitating obtaining patterns with physical meaning, as seen both in the T_w and St maps. When the foil is heated 25 K, the peaks are more evident than when it is heated 35 K. However, a higher level of heating shows less noisy Stanton number maps, with physical patterns being clearly identified. Quantitatively, the variance of these sequences can give a measure of the relative noise level among them. The intermediate and highest heating levels respectively retain 43% and 25% of the variance of the case with the lowest heating level. It must be considered though that natural convection is stronger in this latter case (although it is simple to filter with the strategy outlined in section 2.3) and that the thermal expansion of the foil might produce undesired deformations.

3.3 Effect of foil thickness

We also conducted tests with a foil of thickness $10\ \mu\text{m}$. This comparison is performed using a similar heating to the case with higher heating (30 K) and the same filtering procedure and settings. Although the requirements of small Biot and large Fourier numbers are still fulfilled, this sensor may have less sensitivity to the larger thermal inertia. A higher mass increases the thermal inertia and conduction, affecting the response time to temperature changes and distorting the temperature distribution on the surface. The map of filtered T_w and that of St at a given time instant of the sequence are shown in figure 8, with the foil heated to a middle point between the cases in figures 7 (centre and right). The temperature fluctuation map shows some patterns with shapes comparable to those with the $5\ \mu\text{m}$ foil with the highest heating level that reasonably might represent near-wall turbulence. However, the magnitude of the peaks of these fluctuations is lower. More importantly, although not observed in a single map, a visual inspection of a temporal sequence shows slower temperature changes. This may affect the quantification of heat transfer, as temporal derivatives are directly involved. The St map shows that the strongest fluctuations are well pronounced. However, other patterns, mainly representing noise rather than the wall-bounded turbulence physics in the near-wall region, remain stronger than those observed the $5\ \mu\text{m}$ foil. In light of these findings, one might expect that using a foil thinner than $5\ \mu\text{m}$ would further reduce those effects.

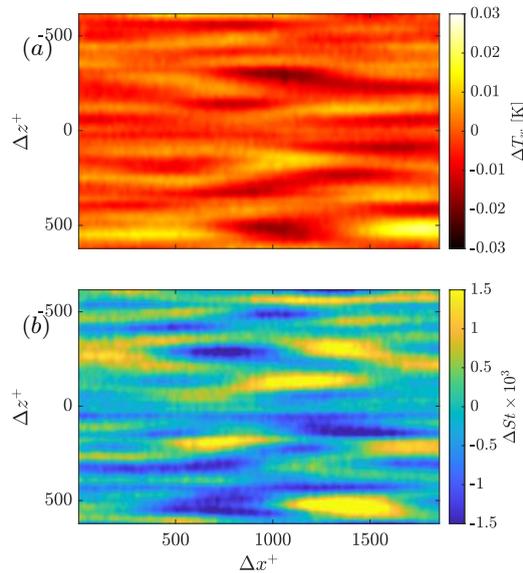


Figure 8. Instantaneous temperature map after the filtering process (a) and Stanton number map (b) from a $10\ \mu\text{m}$ foil heated about 30 K with 3.7 V - 10.0 A.

4. Uncertainty quantification

The uncertainty of the model has been quantified with a Monte Carlo approach. The uncertainties indicated in table 1 have been considered to define a 99% confidence interval on a Gaussian distribution to generate multiple random combinations of these inputs. For the temperature maps and the ambient temperature, the NETD has been considered to be uniformly distributed through the POD modes and the Gaussian filters, thus the actual noise variance has been assumed as the NETD² times the percentage of POD modes retained and the pixelwise contribution of Gaussian filters. As such, the measurement uncertainty of fluctuating St fields was quantified to be equal to 21.3%, 13.8%, 9.4%, respectively for the cases (a), (b), and (c), of figure 7. For the 10 μm foil, as reported in figure 8, it is equal to 11.6%. Despite having the same uncertainties for the voltage, current and thickness, the higher quantities introduce less relative uncertainty than for the 5 μm foil while the larger thickness magnifies the uncertainty contribution of the temporal derivatives. These results support the choice of increased foil heating and smaller foil thickness as a means of improving measurement quality.

5. Conclusions

The heat transfer fluctuations in an air channel flow are quantified through the energy balance of a heated thin foil embedded in the wall. Temperature measurements are conducted employing IR thermography. The main challenges in tackling this experiment are the presence of high-frequency events to be captured, the noise of the experimental equipment and the thermal inertia and conduction of the foil. These heat transfer measurements have been made possible through the balance between different aspects and components of the problem, showing also examples of less convenient configurations. Some critical factors are the black matt coating, the power input and the foil thickness. Additionally, a filtering process is capable of isolating the characteristic turbulent phenomena of the channel from other effects and removing the high level of noise contained in the original temperature acquisitions. The impact of some parameters of the problem on this filtering process has been analysed. Summarising, the main guidelines obtained from this paper are the following:

- It is recommended to remove spurious natural convection effects with high-pass filtering. It was shown that the natural convection effects could be suppressed for different heating levels, leading to similar results for all the conditions analysed in this note.
- Measurement noise can be successfully removed with a feature-based POD filter coupled

with a Gaussian filter to improve the data smoothness.

- Foil thickness should be minimized whenever possible; measurement results are found to be weakly dependent on foil thickness provided the heating is strong enough to ensure sufficient signal-to-noise ratio. Note that when dealing with very thin foils, the presence of a thin layer of paint on the foil (necessary to enable precise IR measurements) must be taken into account when estimating the foil thermal inertia.

Supplementary material Data and codes will be available upon publication.

Acknowledgements A.C. acknowledges financial support from the Spanish Ministry of Universities under the FPU programme 2020. This activity is part of the project ACCREDITATION (Grant No TED2021-131453B-I00), funded by MCIN/AEI/ 10.13039/501100011033 and by the “European Union NextGenerationEU/PRTR”. This activity is part of the project EXCALIBUR (Grant No PID2022-138314NB-I00), funded by MCIU/AEI/ 10.13039/501100011033 and by “ERDF A way of making Europe”. S.D. acknowledges funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement no. 949085, NEXTFLOW).

Bibliography

- [1] A. J. Smits and I. Marusic, “Wall-bounded turbulence,” *Phys. Today*, vol. 66, no. 9, pp. 25–30, 2013.
- [2] A. Tsiringakis, N. E. Theeuwes, J. F. Barlow, and G.-J. Steeneveld, “Interactions between the nocturnal low-level jets and the urban boundary layer: A case study over London,” *Bound.-Lay. Meteorol.*, vol. 183, no. 2, pp. 249–272, 2022.
- [3] A. Colagrossi, S. Marrone, P. Colagrossi, and D. Le Touzé, “Da Vinci’s observation of turbulence: A French-Italian study aiming at numerically reproducing the physics behind one of his drawings, 500 years later,” *Phys. Fluids*, vol. 33, no. 11, 2021.
- [4] J. J. Monaghan and J. B. Kajtar, “Leonardo da Vinci’s turbulent tank in two dimensions,” *Eur. J. Mech. B-Fluids*, vol. 44, pp. 1–9, 2014.
- [5] I. Marusic and S. Broomhall, “Leonardo da Vinci and fluid mechanics,” *Annu. Rev. Fluid Mech.*, vol. 53, no. 1, pp. 1–25, 2021.
- [6] A. N. Kolmogorov, “Equations of turbulent motion in an incompressible fluid,” in *Dokl. Akad. Nauk SSSR*, vol. 30, 1941, pp. 299–303.
- [7] W. Froude, *Rep. Br. Assoc. Adv. Sci.*, pp. 249–55, 1874.
- [8] P. L., “Über flüssigkeitsbewegung bei sehr kleiner reibung. verhandlungen,” in *III. Internationalen Mathematiker Kongresses, Heidelberg, 8-13 August 1904*, B. G. Teubner, Leipzig, 1905, pp. 485–491.
- [9] A. Gyr, W. Kinzelbach, and A. Tsinober, *Fundamental problematic issues in turbulence*. Springer Science & Business Media, 1999.
- [10] A. Chorin, “Lectures on turbulence theory, publish or perish,” *Mathematics Lecture Series*, vol. 5, 1975.
- [11] P. Constantin, “Geometric statistics in turbulence,” *SIAM rev.*, vol. 36, no. 1, pp. 73–98, 1994.
- [12] A. S. Monin, “On the nature of turbulence,” *Soviet Phys. Uspekhi*, vol. 21, no. 5, p. 429, 1978.
- [13] R. D. Moser, S. W. Haering, and G. R. Yalla, “Statistical properties of subgrid-scale turbulence models,” *Annu. Rev. Fluid Mech.*, vol. 53, no. 1, pp. 255–286, 2021.
- [14] J. Bec, K. Gustavsson, and B. Mehlig, “Statistical models for the dynamics of heavy particles in turbulence,” *Annu. Rev. Fluid Mech.*, vol. 56, no. 1, pp. 189–213, 2024.
- [15] R. M. McMullen, M. C. Krygier, J. R. Torczynski, and M. A. Gallis, “Navier-Stokes equations do not describe the smallest scales of turbulence in gases,” *Phys. Rev. Lett.*, vol. 128, no. 11, p. 114501, 2022.
- [16] T. Chen, T. Liu, Z.-Q. Dong, L.-P. Wang, and S. Chen, “Near-wall flow structures and related surface quantities in wall-bounded turbulence,” *Phys. Fluids*, vol. 33, no. 6, 2021.
- [17] J. P. Cooke, M. F. Campbell, E. B. Steager, I. Bargatin, M. H. Yim, and G. I. Park, “Numerical and experimental study on the addition of surface roughness to micro-propellers,” *Phys. Fluids*, vol. 35, no. 11, 2023.
- [18] A. Grava, M. Picillo, J. Serpieri, G. Iuso, L. Bernardos, and G. Cafiero, “Aerodynamic investigation of a drone propeller in cross-flow,” in *AIAA SCITECH 2024 Forum*, 2024, p. 0244.
- [19] L. Zhang, X. Shan, and T. Xie, “Active control for wall drag reduction: Methods, mechanisms and performance,” *IEEE Access*, vol. 8, pp. 7039–7057, 2020.
- [20] E. Amico, G. Cafiero, and G. Iuso, “Deep reinforcement learning for active control of a three-dimensional bluff body wake,” *Phys. Fluids*, vol. 34, no. 10, 2022.
- [21] C. Meneveau and J. Katz, “Scale-invariance and turbulence models for large-eddy simulation,” *Annu. Rev. Fluid Mech.*, vol. 32, no. 1, pp. 1–32, 2000.
- [22] J.-E. Schumann, S. Toosi, and J. Larsson, “Assessment of grid anisotropy effects on large-eddy-simulation models with different length scales,” *AIAA J.*, vol. 58, no. 10, pp. 4522–4533, 2020.

- [23] S. Heinz, R. Mokhtarpour, and M. Stoellinger, “Theory-based Reynolds-averaged Navier–Stokes equations with large eddy simulation capability for separated turbulent flow simulations,” *Phys. Fluids*, vol. 32, no. 6, 2020.
- [24] B. P. Brener, M. A. Cruz, R. L. Thompson, and R. P. Anjos, “Conditioning and accurate solutions of Reynolds average Navier–Stokes equations with data-driven turbulence closures,” *J. Fluid Mech.*, vol. 915, A110, 2021.
- [25] C. Mills, “Identifying the transition between laminar and turbulent flow using high-frequency pressure loss measurement,” Ph.D. dissertation, Coventry University Coventry, 2020.
- [26] J. Gomez, H. Yu, and Y. Andreopoulos, “Role of flow reversals in transition to turbulence and relaminarization of pulsatile flows,” *J. Fluid Mech.*, vol. 917, A27, 2021.
- [27] D. Feldmann, D. Morón, and M. Avila, “Spatiotemporal intermittency in pulsatile pipe flow,” *Entropy*, vol. 23, no. 1, p. 46, 2020.
- [28] K. R. Sreenivasan and V. Yakhot, “Dynamics of three-dimensional turbulence from Navier-Stokes equations,” *Phys. Rev. Fluids*, vol. 6, no. 10, p. 104604, 2021.
- [29] D. Livescu, “Turbulence with large thermal and compositional density variations,” *Annu. Rev. Fluid Mech.*, vol. 52, no. 1, pp. 309–341, 2020.
- [30] C. Caulfield, “Layering, instabilities, and mixing in turbulent stratified flows,” *Annu. Rev. Fluid Mech.*, vol. 53, no. 1, pp. 113–145, 2021.
- [31] K. A. Davis, G. Pawlak, and S. G. Monismith, “Turbulence and coral reefs,” *Annu. Rev. Marine Sci.*, vol. 13, no. 1, pp. 343–373, 2021.
- [32] S. Prants, “Transport barriers in geophysical flows: A review,” *Symmetry*, vol. 15, no. 10, p. 1942, 2023.
- [33] C. Octau, T. Dbouk, M. Watremez, *et al.*, “Liquid–solid two-phase jet in a turbulent crossflow: Experiments and simulations,” *Chem. Eng. Res. Des.*, vol. 155, pp. 156–171, 2020.
- [34] N. Jain, A. Le Moine, G. Chaussonnet, *et al.*, “A critical review of physical models in high temperature multiphase fluid dynamics: Turbulent transport and particle-wall interactions,” *Appl. Mech. Rev.*, vol. 73, no. 4, p. 040801, 2021.
- [35] R. J. Adrian and J. Westerweel, *Particle image velocimetry*. Cambridge University Press, 2011.
- [36] M. Raffel, C. E. Willert, F. Scarano, C. J. Kähler, S. T. Wereley, and J. Kompenhans, *Particle image velocimetry: a practical guide*. Springer, 2018.
- [37] J. Westerweel, G. E. Elsinga, and R. J. Adrian, “Particle image velocimetry for complex and turbulent flows,” *Annu. Rev. Fluid Mech.*, vol. 45, no. 1, pp. 409–436, 2013.
- [38] A. Favre, J. Gaviglio, and R. Dumas, “Space-time double correlations and spectra in turbulent boundary layer,” *J. Fluid Mech.*, vol. 2, no. 4, p. 313, 1957.
- [39] A. Favre, J. Gaviglio, and R. Dumas, “Further space-time correlations of velocity in a turbulent boundary layer,” *J. Fluid Mech.*, vol. 3, no. 4, p. 344, 1958.
- [40] R. Örlü, T. Fiorini, A. Segalini, G. Bellani, A. Talamelli, and P. H. Alfredsson, “Reynolds stress scaling in pipe flow turbulence—first results from CICLoPE,” *Philos. T. Roy. Soc. A: Math. Phys. Eng. Sci.*, vol. 375, no. 2089, p. 20160187, 2017.
- [41] J. Drahotský, P. Hanzelka, V. Musilová, M. Macek, R. Du Puits, and P. Urban, “Temperature profiles measurements in turbulent Rayleigh–Bénard convection by optical fibre system at the Barrel of II-menau,” in *EPJ Web of Conferences*, EDP Sciences, vol. 180, 2018, p. 02020.
- [42] D. Anderson, J. C. Tannehill, R. H. Pletcher, R. Munipalli, and V. Shankar, *Computational fluid mechanics and heat transfer*. CRC press, 2020.
- [43] J. Tu, G. H. Yeoh, C. Liu, and Y. Tao, *Computational fluid dynamics: a practical approach*. Elsevier, 2023.
- [44] J. Jiménez, “Computers and turbulence,” *Eur. J. Mech. B-Fluids*, vol. 79, pp. 1–11, 2020.
- [45] S. B. Pope, *Turbulent flows*. Cambridge University Press, 2000, p. 392.
- [46] J. Jiménez, “Computing high-Reynolds-number turbulence: Will simulations ever replace experiments?” *J. Turb.*, vol. 4, no. 1, p. 022, 2003.
- [47] United Nations General Assembly, *Resolution 70/1 (2015) - Transforming our world: the 2030 Agenda for Sustainable Development*, 2015.
- [48] Centre for Research on Energy and Clean Air (CREA), *EU’s CO2 emissions from fossil fuels drop 8% to reach lowest levels in 60 years*, https://energyandcleanair.org/wp/wp-content/uploads/2024/01/CREA_EU_2023_CO2_01.2024.pdf, 2024.

- [49] M. Gad-el-Hak, "Modern developments in flow control," *Appl. Mech. Rev.*, vol. 49, no. 7, pp. 365–379, 1996.
- [50] J. Jiménez, "Near-wall turbulence," *Phys. Fluids*, vol. 25, no. 10, 2013.
- [51] F. De Gregorio, D. Steiling, E. Benini, and R. Ponzà, "ERICA tiltrotor airframe wake characterization," 2015.
- [52] D. M. Bushnell and J. N. Hefner, *Viscous drag reduction in boundary layers*. Washington DC: American Institute of Aeronautics and Astronautics, 1990.
- [53] R. W. Barnwell and M. Y. Hussaini, *Natural laminar flow and laminar flow control*. New York: Springer Science & Business Media, 2012.
- [54] S. P. Wilkinson, "Interactive wall turbulence control," in *Viscous drag reduction in boundary layers*, D. M. Bushnell and J. N. Hefner, Eds., vol. 123, Washington, DC: AIAA, 1990, pp. 479–509.
- [55] P. Moin and T. Bewley, "Feedback control of turbulence," *App. Mech. Rev.*, vol. 47, no. 6 part 2, S3, 1994.
- [56] L. N. Cattafesta III and M. Sheplak, "Actuators for active flow control," *Annu. Rev. Fluid Mech.*, vol. 43, pp. 247–272, 2011.
- [57] I. Marusic, D. Chandran, A. Rouhi, *et al.*, "An energy-efficient pathway to turbulent drag reduction," *Nat. Commun.*, vol. 12, no. 1, p. 5805, 2021.
- [58] H. Choi, P. Moin, and J. Kim, "Active turbulence control for drag reduction in wall-bounded flows," *J. Fluid Mech.*, vol. 262, pp. 75–110, 1994.
- [59] S. Perumal, D. Sundaresan, R. Sivanraju, N. Tesfie, K. Ramalingam, and S. Thanikodi, "Heat transfer analysis in counter flow shell and tube heat exchanger using of design of experiments," *Therm. Sci.*, vol. 26, no. 2 Part A, pp. 843–848, 2022.
- [60] P. Schlatter and R. Örlü, "Assessment of direct numerical simulation data of turbulent boundary layers," *J. Fluid Mech.*, vol. 659, pp. 116–126, 2010.
- [61] A. C. Yunus, *Fluid Mechanics: Fundamentals And Applications (Si Units)*. Tata McGraw Hill Education Private Limited, 2010.
- [62] R. F. Blackwelder and L. S. Kovasznay, "Time scales and correlations in a turbulent boundary layer," *Phys. Fluids*, vol. 15, no. 9, pp. 1545–1554, 1972.
- [63] P. Moin and J. Kim, "Numerical investigation of turbulent channel flow," *J. Fluid Mech.*, vol. 118, pp. 341–377, 1982.
- [64] P. Moin and J. Kim, "The structure of the vorticity field in turbulent channel flow. Part 1. Analysis of instantaneous fields and statistical correlations," *J. Fluid Mech.*, vol. 155, pp. 441–464, 1985.
- [65] H. Grant, "The large eddies of turbulent motion," *J. Fluid Mech.*, vol. 4, no. 2, pp. 149–190, 1958.
- [66] J. Bakewell, P. Henry, and J. Lumley, "Viscous sublayer and adjacent wall region in turbulent pipe flow," *Phys. Fluids*, vol. 10, no. 9, pp. 1880–1889, 1967.
- [67] B. E. Launder, G. J. Reece, and W. Rodi, "Progress in the development of a Reynolds-stress turbulence closure," *J. Fluid Mech.*, vol. 68, no. 3, pp. 537–566, 1975.
- [68] J. Smagorinsky, "General circulation experiments with the primitive equations: I. the basic experiment," *Mon. weather rev.*, vol. 91, no. 3, pp. 99–164, 1963.
- [69] A. A. Townsend, "The structure of the turbulent boundary layer," in *Mathematical Proceedings of the Cambridge Philosophical Society*, Cambridge University Press, vol. 47, 1951, pp. 375–395.
- [70] A. Townsend, "Equilibrium layers and wall turbulence," *J. Fluid Mech.*, vol. 11, no. 1, pp. 97–120, 1961.
- [71] A. Townsend, *The structure of turbulent shear flow*. Cambridge University Press, 1976.
- [72] I. Marusic and P. M. Jason, "Attached eddy model of wall turbulence," *Annu. Rev. Fluid Mech.*, vol. 51, pp. 49–74, 2019.
- [73] J. Hwang and H. J. Sung, "Wall-attached structures of velocity fluctuations in a turbulent boundary layer," *J. Fluid Mech.*, vol. 856, pp. 958–983, 2018.
- [74] A. F. Hussain, "Role of coherent structures in turbulent shear flows," *Proceedings of the Indian Academy of Sciences Section C: Engineering Sciences*, vol. 4, pp. 129–175, 1981.
- [75] A. J. Smits, B. J. McKeon, and I. Marusic, "High-Reynolds number wall turbulence," *Annu. Rev. Fluid Mech.*, vol. 43, no. 1, pp. 353–375, 2011.
- [76] S. Kline, W. Reynolds, F. Schraub, and P. Runstadler, "The structure of turbulent boundary layers," *J. Fluid Mech.*, vol. 30, no. 4, pp. 741–773, 1967.
- [77] T. Theodorsen, "Mechanisms of turbulence," in *Proceedings of the 2nd Midwestern Conference on Fluid Mechanics*, 1952.

- [78] M. Head and P. Bandyopadhyay, “New aspects of turbulent boundary-layer structure,” *J. Fluid Mech.*, vol. 107, pp. 297–338, 1981.
- [79] X. Wu and P. Moin, “Direct numerical simulation of turbulence in a nominally zero-pressure-gradient flat-plate boundary layer,” *J. Fluid Mech.*, vol. 630, pp. 5–41, 2009.
- [80] J. Zhou, R. Adrian, and S. Balachandar, “Autogeneration of near-wall vortical structures in channel flow,” *Phys. Fluids*, vol. 8, no. 1, pp. 288–290, 1996.
- [81] K. Kim and R. Adrian, “Very large-scale motion in the outer layer,” *Phys. Fluids*, vol. 11, no. 2, pp. 417–422, 1999.
- [82] J. Zhou, R. J. Adrian, S. Balachandar, and T. Kendall, “Mechanisms for generating coherent packets of hairpin vortices in channel flow,” *J. Fluid Mech.*, vol. 387, pp. 353–396, 1999.
- [83] R. J. Adrian, C. D. Meinhart, and C. D. Tomkins, “Vortex organization in the outer region of the turbulent boundary layer,” *J. Fluid Mech.*, vol. 422, pp. 1–54, 2000.
- [84] M. Guala, S. Hommema, and R. Adrian, “Large-scale and very-large-scale motions in turbulent pipe flow,” *J. Fluid Mech.*, vol. 554, p. 521, 2006.
- [85] B. J. Balakumar and R. J. Adrian, “Large-and very-large-scale motions in channel and boundary-layer flows,” *Philos. T. Roy. Soc. A: Math. Phys. Eng. Sci.*, vol. 365, no. 1852, pp. 665–681, 2007.
- [86] R. J. Adrian, “Hairpin vortex organization in wall turbulence,” *Phys. Fluids*, vol. 19, no. 4, 2007.
- [87] B. Ganapathisubramani, E. K. Longmire, and I. Marusic, “Characteristics of vortex packets in turbulent boundary layers,” *J. Fluid Mech.*, vol. 478, pp. 35–46, 2003.
- [88] C. Tomkins and R. Adrian, “Spanwise structure and scale growth in turbulent boundary layers,” *J. Fluid Mech.*, vol. 490, pp. 37–74, 2003.
- [89] N. Hutchins and I. Marusic, “Evidence of very long meandering features in the logarithmic region of turbulent boundary layers,” *J. Fluid Mech.*, vol. 579, pp. 1–28, 2007.
- [90] N. Hutchins and I. Marusic, “Large-scale influences in near-wall turbulence,” *Philos. T. Roy. Soc. A: Math., Phys. and Eng. Sci.*, vol. 365, no. 1852, pp. 647–664, 2007.
- [91] J. Monty, J. Stewart, R. Williams, and M. Chong, “Large-scale features in turbulent pipe and channel flows,” *J. Fluid Mech.*, vol. 589, pp. 147–156, 2007.
- [92] S. Hoyas and J. Jiménez, “Scaling of the velocity fluctuations in turbulent channels up to $Re_\tau = 2003$,” *Phys. Fluids*, vol. 18, no. 1, 2006.
- [93] C. Sanmiguel Vila, R. Örlü, R. Vinuesa, P. Schlatter, A. Ianiro, and S. Discetti, “Adverse-pressure-gradient effects on turbulent boundary layers: Statistics and flow-field organization,” *Flow Turbul. Combust.*, vol. 99, no. 3, pp. 589–612, 2017.
- [94] A. Cremades, S. Hoyas, R. Deshpande, *et al.*, “Identifying regions of importance in wall-bounded turbulence through explainable deep learning,” *Nat. Commun.*, vol. 15, no. 1, p. 3864, 2024.
- [95] R. Mathis, N. Hutchins, and I. Marusic, “Large-scale amplitude modulation of the small-scale structures in turbulent boundary layers,” *J. Fluid Mech.*, vol. 628, pp. 311–337, 2009.
- [96] M. Tutkun, W. K. George, J. Delville, *et al.*, “Two-point correlations in high Reynolds number flat plate turbulent boundary layers,” *J. Turbul.*, no. 10, N21, 2009.
- [97] P. Schlatter, R. Örlü, Q. Li, *et al.*, “Turbulent boundary layers up to $Re_\theta = 2500$ studied through simulation and experiment,” *Phys. Fluids*, vol. 21, no. 5, 2009.
- [98] Ó. Flores and J. Jiménez, “Effect of wall-boundary disturbances on turbulent channel flows,” *J. Fluid Mech.*, vol. 566, pp. 357–376, 2006.
- [99] Ó. Flores, J. Jiménez, and J. C. Del Álamo, “Vorticity organization in the outer layer of turbulent channels with disturbed walls,” *J. Fluid Mech.*, vol. 591, pp. 145–154, 2007.
- [100] S. C. Bailey and A. J. Smits, “Experimental investigation of the structure of large-and very-large-scale motions in turbulent pipe flow,” *J. Fluid Mech.*, vol. 651, pp. 339–356, 2010.
- [101] S. L. Brunton and J. N. Kutz, *Data-Driven Science and Engineering*. University Of Washington: Cambridge University Press, 2019, ISBN: 978-1-108-42209-3.
- [102] I. Gohberg, S. Goldberg, and M. A. Kaashoek, “Spectral theory of compact self adjoint operators,” in *Basic Classes of Linear Operators*. Basel: Birkhäuser Basel, 2003, pp. 171–191.
- [103] S. Wu, K. T. Christensen, and C. Pantano, “Modelling smooth-and transitionally rough-wall turbulent channel flow by leveraging inner–outer interactions and principal component analysis,” *J. Fluid Mech.*, vol. 863, pp. 407–453, 2019.
- [104] J. Lumley, “The structure of inhomogeneous turbulent flows,” *Atmospheric Turbulence and Wave Propagation*, pp. 166–178, 1967.

- [105] M. Ilak and C. W. Rowley, “Modeling of transitional channel flow using balanced proper orthogonal decomposition,” *Phys. Fluids*, vol. 20, no. 3, 2008.
- [106] G. Berkooz, P. Holmes, and J. Lumley, “The proper orthogonal decomposition in the analysis of turbulent flows,” *Annu. Rev. Fluid Mech.*, vol. 25, pp. 539–75, 1993.
- [107] P. J. Schmid, “Dynamic mode decomposition and its variants,” *Annu. Rev. Fluid Mech.*, vol. 54, no. 1, pp. 225–254, 2022.
- [108] H. Jang, C. Ozdemir, J.-H. Liang, and M. Tyagi, “Oscillatory flow around a vertical wall-mounted cylinder: Dynamic mode decomposition,” *Phys. Fluids*, vol. 33, no. 2, 2021.
- [109] S. Le Clainche and J. M. Vega, “Higher order dynamic mode decomposition,” *SIAM J. Appl. Dyn. Syst.*, vol. 16, no. 2, pp. 882–925, 2017.
- [110] S. Le Clainche and J. M. Vega, “Higher order dynamic mode decomposition to identify and extrapolate flow patterns,” *Phys. Fluids*, vol. 29, no. 8, 2017.
- [111] Y. Mizuno, D. Duke, C. Atkinson, and J. Soria, “Investigation of wall-bounded turbulent flow using dynamic mode decomposition,” in *J. Phys. Conf. Ser.*, IOP Publishing, vol. 318, 2011, p. 042040.
- [112] E. Farzamnig, A. Ianiro, S. Discetti, *et al.*, “From snapshots to manifolds—a tale of shear flows,” *J. Fluid Mech.*, vol. 955, A34, 2023.
- [113] L. Marra, G. Y. Cornejo Maceda, A. Meilán-Vila, *et al.*, “Actuation manifold from snapshot data,” *J. Fluid Mech.*, vol. 996, A26, 2024.
- [114] A. Solera-Rico, C. Sanmiguel Vila, M. Gómez-López, *et al.*, “ β -Variational autoencoders and transformers for reduced-order modelling of fluid flows,” *Nat. Commun.*, vol. 15, no. 1, p. 1361, 2024.
- [115] R. Vinuesa, “Perspectives on predicting and controlling turbulent flows through deep learning,” *Phys. Fluids*, vol. 36, no. 3, 2024.
- [116] M. Sen, K. Bhaganagar, and V. Juttijudata, “Application of proper orthogonal decomposition (POD) to investigate a turbulent boundary layer in a channel with rough walls,” *J. Turbul.*, no. 8, N41, 2007.
- [117] S. D. Muralidhar, B. Podvin, L. Mathelin, and Y. Fraigneau, “Spatio-temporal proper orthogonal decomposition of turbulent channel flow,” *J. Fluid Mech.*, vol. 864, pp. 614–639, 2019.
- [118] J. Borée, “Extended proper orthogonal decomposition: A tool to analyse correlated events in turbulent flows,” *Exp. Fluids*, vol. 35, no. 2, pp. 188–192, 2003.
- [119] C. Tinney, L. Ukeiley, and M. N. Glauser, “Low-dimensional characteristics of a transonic jet. Part 2. Estimate and far-field prediction,” *J. Fluid Mech.*, vol. 615, pp. 53–92, 2008.
- [120] F. Mallor, M. Raiola, C. Sanmiguel Vila, R. Örlü, S. Discetti, and A. Ianiro, “Modal decomposition of flow fields and convective heat transfer maps: An application to wall-proximity square ribs,” *Exp. Therm. Fluid Sci.*, vol. 102, pp. 517–527, 2019.
- [121] C. Hoarau, J. Borée, J. Laumonier, and Y. Gervais, “Analysis of the wall pressure trace downstream of a separated region using extended proper orthogonal decomposition,” *Phys. Fluids*, vol. 18, no. 5, 2006.
- [122] R. K. Schlander, S. Rigopoulos, and G. Papadakis, “Analysis of wall mass transfer in turbulent pipe flow combining extended proper orthogonal decomposition and Fukagata-Iwamoto-Kasagi identity,” *Phys. Rev. Fluids*, vol. 7, no. 2, p. 024603, 2022.
- [123] J. Chen, M. Raiola, and S. Discetti, “Pressure from data-driven estimation of velocity fields using snapshot PIV and fast probes,” *Exp. Therm. Fluid Sci.*, vol. 136, p. 110647, 2022.
- [124] F. Nicoud, J. Baggett, P. Moin, and W. Cabot, “Large eddy simulation wall-modeling based on suboptimal control theory and linear stochastic estimation,” *Phys. Fluids*, vol. 13, no. 10, pp. 2968–2984, 2001.
- [125] J. Hoepffner, M. Chevalier, T. R. Bewley, and D. S. Henningson, “State estimation in wall-bounded flow systems. Part 1. Perturbed laminar flows,” *J. Fluid Mech.*, vol. 534, pp. 263–294, 2005.
- [126] M. Chevalier, J. Hoepffner, T. R. Bewley, and D. S. Henningson, “State estimation in wall-bounded flow systems. Part 2. Turbulent flows,” *J. Fluid Mech.*, vol. 552, pp. 167–187, 2006.
- [127] I. Marusic, R. Mathis, and N. Hutchins, “Predictive model for wall-bounded turbulent flow,” *Sci.*, vol. 329, no. 5988, pp. 193–196, 2010.
- [128] D. Lasagna, M. Orazi, and G. Iuso, “Multi-time delay, multi-point linear stochastic estimation of a cavity shear layer velocity from wall-pressure measurements,” *Phys. Fluids*, vol. 25, no. 1, 2013.
- [129] I. Marusic, W. J. Baars, and N. Hutchins, “An extended view of the inner-outer interaction model for wall-bounded turbulence using spectral linear stochastic estimation,” *Procedia Eng.*, vol. 126, pp. 24–28, 2015.

- [130] W. J. Baars, N. Hutchins, and I. Marusic, “Spectral stochastic estimation of high-Reynolds-number wall-bounded turbulence for a refined inner-outer interaction model,” *Phys. Rev. Fluids*, vol. 1, no. 5, p. 054406, 2016.
- [131] B. Podvin, “On combining linear stochastic estimation and proper orthogonal decomposition for flow reconstruction,” *Exp. Fluids*, vol. 58, 2018.
- [132] M. P. Encinar and J. Jiménez, “Logarithmic-layer turbulence: A view from the wall,” *Phys. Rev. Fluids*, vol. 4, no. 11, p. 114603, 2019.
- [133] V. Gupta, A. Madhusudanan, M. Wan, S. J. Illingworth, and M. P. Juniper, “Linear-model-based estimation in wall turbulence: Improved stochastic forcing and eddy viscosity terms,” *J. Fluid Mech.*, vol. 925, A18, 2021.
- [134] T. Suzuki and Y. Hasegawa, “Estimation of turbulent channel flow at based on the wall measurement using a simple sequential approach,” *J. Fluid Mech.*, vol. 830, p. 760, 2017.
- [135] W. Bauer, O. Haag, and D. Hennecke, “Accuracy and robustness of non-linear eddy viscosity models,” in *Engineering Turbulence Modelling and Experiments 4*, W. Rodi and D. Laurence, Eds., Oxford: Elsevier Science Ltd, 1999, pp. 113–124.
- [136] M. Milano and P. Koumoutsakos, “Neural network modeling for near wall turbulent flow,” *J. Comp. Phys.*, vol. 182, no. 1, pp. 1–26, 2002.
- [137] J. Kim and C. Lee, “Prediction of turbulent heat transfer using convolutional neural networks,” *J. Fluid Mech.*, vol. 882, A18, 2020.
- [138] E. Lagemann, S. L. Brunton, and C. Lagemann, “Uncovering wall-shear stress dynamics from neural-network enhanced fluid flow measurements,” *P. Roy. Soc. A-Math. Phys.*, vol. 480, no. 2292, p. 20230798, 2024.
- [139] A. Güemes, S. Discetti, and A. Ianiro, “Sensing the turbulent large-scale motions with their wall signature,” *Phys. Fluids*, vol. 31, no. 12, p. 125112, 2019.
- [140] L. Guastoni, M. P. Encinar, P. Schlatter, H. Azizpour, and R. Vinuesa, “Prediction of wall-bounded turbulence from wall quantities using convolutional neural networks,” in *J. Phys. Conf. Ser.*, IOP Publishing, vol. 1522, 2020, p. 012022.
- [141] L. Guastoni, P. A. Srinivasan, H. Azizpour, P. Schlatter, and R. Vinuesa, *On the use of recurrent neural networks for predictions of turbulent flows*, 2020. arXiv: [2002.01222 \[physics.flu-dyn\]](https://arxiv.org/abs/2002.01222).
- [142] L. Guastoni, A. Güemes, A. Ianiro, *et al.*, “Convolutional-network models to predict wall-bounded turbulence from wall quantities,” *J. Fluid Mech.*, vol. 928, A27, 2021.
- [143] A. Güemes, S. Discetti, A. Ianiro, B. Sirmacek, H. Azizpour, and R. Vinuesa, “From coarse wall measurements to turbulent velocity fields through deep learning,” *Phys. Fluids*, vol. 33, no. 7, 2021.
- [144] K. Sasaki, R. Vinuesa, A. V. Cavalieri, P. Schlatter, and D. S. Henningson, “Transfer functions for flow predictions in wall-bounded turbulence,” *J. Fluid Mech.*, vol. 864, pp. 708–745, 2019.
- [145] D. Lasagna, L. Fronges, M. Orazi, and G. Iuso, “Nonlinear multi-time-delay stochastic estimation: Application to cavity flow and turbulent channel flow,” *AIAA J.*, vol. 53, no. 10, pp. 2920–2935, 2015.
- [146] M. Wang and T. A. Zaki, “Synchronization of turbulence in channel flow,” *J. Fluid Mech.*, vol. 943, A4, 2022.
- [147] M. Z. Yousif, L. Yu, S. Hoyas, R. Vinuesa, and H. Lim, “A deep-learning approach for reconstructing 3D turbulent flows from 2D observation data,” *Sci. Rep.*, vol. 13, no. 1, p. 2529, 2023.
- [148] L. Yu, M. Z. Yousif, M. Zhang, S. Hoyas, R. Vinuesa, and H.-C. Lim, “Three-dimensional ESRGAN for super-resolution reconstruction of turbulent flows with tricubic interpolation-based transfer learning,” *Phys. Fluids*, vol. 34, no. 12, 2022.
- [149] E. Jagodinski, X. Zhu, and S. Verma, “Inverse identification of dynamically important regions in turbulent flows using three-dimensional convolutional neural networks,” *Phys. Rev. Fluids*, vol. 8, no. 9, p. 094605, 2023.
- [150] A. Xuan and L. Shen, “Reconstruction of three-dimensional turbulent flow structures using surface measurements for free-surface flows based on a convolutional neural network,” *J. Fluid Mech.*, vol. 959, A34, 2023.
- [151] S. K. Chinnamgari, *R Machine Learning Projects: Implement supervised, unsupervised, and reinforcement learning techniques using R 3.5*. Packt Publishing Ltd, 2019.
- [152] S. L. Brunton, B. R. Noack, and P. Koumoutsakos, “Machine learning for fluid mechanics,” *Annu. Rev. Fluid Mech.*, vol. 52, no. 1, pp. 477–508, 2020.

- [153] H. Fan, J. Jiang, C. Zhang, X. Wang, and Y.-C. Lai, “Long-term prediction of chaotic systems with machine learning,” *Phys. Rev. Res.*, vol. 2, no. 1, p. 012080, 2020.
- [154] P. Pant, R. Doshi, P. Bahl, and A. Barati Farimani, “Deep learning for reduced order modelling and efficient temporal evolution of fluid simulations,” *Phys. Fluids*, vol. 33, no. 10, 2021.
- [155] J. C. Loiseau and S. L. Brunton, “Constrained sparse Galerkin regression,” *J. Fluid Mech.*, vol. 838, pp. 42–67, 2018.
- [156] K. Fukami, K. Fukagata, and K. Taira, “Super-resolution reconstruction of turbulent flows with machine learning,” *J. Fluid Mech.*, vol. 870, pp. 106–120, 2019.
- [157] Z. Deng, C. He, Y. Liu, and K. C. Kim, “Super-resolution reconstruction of turbulent velocity fields using a generative adversarial network-based artificial intelligence framework,” *Phys. Fluids*, vol. 31, no. 12, 2019.
- [158] A. Güemes, C. Sanmiguel Vila, and S. Discetti, “Super-resolution generative adversarial networks of randomly-seeded fields,” *Nat. Mach. Int.*, vol. 4, no. 12, pp. 1165–1173, 2022.
- [159] C. Dong, C. C. Loy, K. He, and X. Tang, “Learning a deep convolutional network for image super-resolution,” in *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part IV 13*, Springer, 2014, pp. 184–199.
- [160] C. Ledig, L. Theis, F. Huszár, *et al.*, “Photo-realistic single image super-resolution using a generative adversarial network,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4681–4690.
- [161] P. Díaz-Morales, A. Corrochano, M. López-Martín, and S. Le Clainche, “Deep learning combined with singular value decomposition to reconstruct databases in fluid dynamics,” *Expert Syst. Appl.*, vol. 238, p. 121924, 2024.
- [162] B. Kim, V. C. Azevedo, N. Thuerey, T. Kim, M. Gross, and B. Solenthaler, “Deep fluids: A generative network for parameterized fluid simulations,” in *Comput. graph. forum*, Wiley Online Library, vol. 38, 2019, pp. 59–70.
- [163] D. Kochkov, J. A. Smith, A. Alieva, Q. Wang, M. P. Brenner, and S. Hoyer, “Machine learning-accelerated computational fluid dynamics,” *P. Natl. Acad. Sci.*, vol. 118, no. 21, e2101784118, 2021.
- [164] R. Vinuesa and S. L. Brunton, “Enhancing computational fluid dynamics with machine learning,” *Nat. Comp. Sci.*, vol. 2, no. 6, pp. 358–366, 2022.
- [165] J. Rabault, M. Kuchta, A. Jensen, U. Réglade, and N. Cerardi, “Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control,” *J. Fluid Mech.*, vol. 865, pp. 281–302, 2019.
- [166] T. Lee, J. Kim, and C. Lee, “Turbulence control for drag reduction through deep reinforcement learning,” *Phys. Rev. Fluids*, vol. 8, no. 2, p. 024604, 2023.
- [167] L. Guastoni, J. Rabault, P. Schlatter, H. Azizpour, and R. Vinuesa, “Deep reinforcement learning for turbulent drag reduction in channel flows,” *Eur. Phys. J. E*, vol. 46, no. 4, p. 27, 2023.
- [168] N. Benard, J. Pons-Prats, J. Periaux, *et al.*, “Turbulent separated shear flow control by surface plasma actuator: Experimental optimization by genetic algorithm approach,” *Exp. Fluids*, vol. 57, pp. 1–17, 2016.
- [169] R. Li, B. R. Noack, L. Cordier, J. Borée, and F. Harambat, “Drag reduction of a car model by linear genetic programming control,” *Exp. Fluids*, vol. 58, pp. 1–20, 2017.
- [170] G. Minelli, T. Dong, B. Noack, and S. Krajnović, “Upstream actuation for bluff-body wake control driven by a genetically inspired optimization,” *J. Fluid Mech.*, vol. 893, A1, 2020.
- [171] R. Castellanos, A. Ianiro, and S. Discetti, “Genetically-inspired convective heat transfer enhancement in a turbulent boundary layer,” *Appl. Therm. Eng.*, vol. 230, p. 120621, 2023.
- [172] M. Schmidt and H. Lipson, “Distilling free-form natural laws from experimental data,” *Sci.*, vol. 324, no. 5923, pp. 81–85, 2009.
- [173] S. L. Brunton, J. L. Proctor, and J. N. Kutz, “Sparse identification of nonlinear dynamics with control (sindyc),” *IFAC-PapersOnLine*, vol. 49, no. 18, pp. 710–715, 2016.
- [174] J. Dandois, E. Garnier, and P.-Y. Pamart, “NARX modelling of unsteady separation control,” *Exp. Fluids*, vol. 54, pp. 1–17, 2013.
- [175] B. Sanderse, P. Stinis, R. Maulik, and S. E. Ahmed, *Scientific machine learning for closure models in multiscale problems: A review*, 2024. arXiv: [2403.02913 \[math.NA\]](https://arxiv.org/abs/2403.02913)
- [176] S. Cheng, C. Quilodrán-Casas, S. Ouala, *et al.*, “Machine learning with data assimilation and uncertainty quantification for dynamical systems: A review,” *IEEE/CAA J. Automatica Sinica*, vol. 10, no. 6, pp. 1361–1387, 2023.

- [177] M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *J. Comp. Phys.*, vol. 378, pp. 686–707, 2019.
- [178] S. Cai, Z. Wang, F. Fuest, Y. J. Jeon, C. Gray, and G. E. Karniadakis, “Flow over an espresso cup: Inferring 3-D velocity and pressure fields from tomographic background oriented schlieren via physics-informed neural networks,” *J. Fluid Mech.*, vol. 915, A102, 2021.
- [179] K. Manohar, B. W. Brunton, J. N. Kutz, and S. L. Brunton, “Data-driven sparse sensor placement for reconstruction: Demonstrating the benefits of exploiting known patterns,” *IEEE Contr. Syst. Mag.*, vol. 38, no. 3, pp. 63–86, 2018.
- [180] K. Manohar, J. N. Kutz, and S. L. Brunton, “Optimal sensor and actuator selection using balanced model reduction,” *IEEE T. Automat. Contr.*, vol. 67, no. 4, pp. 2108–2115, 2021.
- [181] Y. Liu, D. Wang, X. Sun, N. Dinh, and R. Hu, “Uncertainty quantification for multiphase-cfd simulations of bubbly flows: A machine learning-based bayesian approach supported by high-resolution experiments,” *Reliab. Eng. Sys. Safe.*, vol. 212, p. 107636, 2021.
- [182] X. Glorot, A. Bordes, and Y. Bengio, “Deep sparse rectifier neural networks,” in *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, JMLR Workshop and Conference Proceedings, 2011, pp. 315–323.
- [183] D. H. Hubel and T. N. Wiesel, “Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex,” *J. Physiol.*, vol. 160, no. 1, pp. 106–154, 1962.
- [184] F. Rosenblatt, “The perceptron: A probabilistic model for information storage and organization in the brain.,” *Psychol. Rev.*, vol. 65, no. 6, p. 386, 1958.
- [185] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, “Learning representations by back-propagating errors,” *Nat.*, vol. 323, no. 6088, pp. 533–536, 1986.
- [186] P. Ramachandran, B. Zoph, and Q. Le, *Searching for activation functions*, 2017. arXiv: [1710.05941 \[cs.NE\]](https://arxiv.org/abs/1710.05941).
- [187] J. Brownlee, “A gentle introduction to the rectified linear unit (ReLU),” *Machine learning mastery*, vol. 6, 2019.
- [188] *What is rectified linear unit (ReLU)? — introduction to ReLU activation function*, <https://www.mygreatlearning.com/blog/relu-activation-function/>, Accessed: 2021-05.
- [189] A. L. Maas, A. Y. Hannun, A. Y. Ng, *et al.*, “Rectifier nonlinearities improve neural network acoustic models,” in *International conference on machine learning*, Atlanta, GA, vol. 30, 2013, p. 3.
- [190] K. He, X. Zhang, S. Ren, and J. Sun, “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1026–1034.
- [191] K. Fukushima, “Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position,” *Biol. Cybern.*, vol. 36, no. 4, pp. 193–202, 1980.
- [192] K. Fukushima and S. Miyake, “Neocognitron: A self-organizing neural network model for a mechanism of visual pattern recognition,” in *Competition and cooperation in neural nets*, Springer, 1982, pp. 267–285.
- [193] K. Fukushima, “Neocognitron: A hierarchical neural network capable of visual pattern recognition,” *Neural networks*, vol. 1, no. 2, pp. 119–130, 1988.
- [194] *2D convolution in image processing*, <https://www.allaboutcircuits.com/technical-articles/two-dimensional-convolution-in-image-processing/>, Accessed: 2024-08, 2018.
- [195] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” *Adv. Neur. In.*, vol. 25, 2012.
- [196] A. Boudjelal, Z. Messali, and A. Elmoataz, “A novel kernel-based regularization technique for pet image reconstruction,” *Technologies*, vol. 5, no. 2, p. 37, 2017.
- [197] J. Collis, “Glossary of deep learning: Batch normalisation,” *Medium*, 2017.
- [198] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *International conference on machine learning*, PMLR, 2015, pp. 448–456.
- [199] S. Santurkar, D. Tsipras, A. Ilyas, and A. Madry, “How does batch normalization help optimization?” *Adv. Neur. In.*, vol. 31, 2018.
- [200] G. Yang, J. Pennington, V. Rao, J. Sohl-Dickstein, and S. S. Schoenholz, *A mean field theory of batch normalization*, 2019. arXiv: [1902.08129 \[cs.NE\]](https://arxiv.org/abs/1902.08129).

- [201] *The History and Evolution of Generative AI: From GAN to GPT*, <https://medium.com/@ibtissam.essadik/the-history-and-evolution-of-generative-ai-from-gan-to-gpt-4-96c1fa69f4e7>, Accessed: 2024-10, 2024.
- [202] I. Goodfellow, J. Pouget-Abadie, M. Mirza, *et al.*, “Generative adversarial nets,” *Adv. Neur. In.*, vol. 27, 2014.
- [203] A. Vela-Martín, M. P. Encinar, A. García-Gutiérrez, and J. Jiménez, “A low-storage method consistent with second-order statistics for time-resolved databases of turbulent channel flow up to $Re_\tau = 5300$,” *J. Comput Sci-Neth.*, vol. 56, p. 101 476, 2021.
- [204] M. Lee and R. D. Moser, “Direct numerical simulation of turbulent channel flow up to,” *J. Fluid Mech.*, vol. 774, pp. 395–415, 2015.
- [205] J. Jiménez, “Coherent structures in wall-bounded turbulence,” *J. Fluid Mech.*, vol. 842, P1, 2018.
- [206] O. Reynolds, “IV. On the dynamical theory of incompressible viscous fluids and the determination of the criterion,” *Philos. T. Roy. Soc. A: Math. Phys. Eng. Sci.*, no. 186, pp. 123–164, 1895.
- [207] O. Amili and J. Soria, “A film-based wall shear stress sensor for wall-bounded turbulent flows,” *Exp. Fluids*, vol. 51, no. 1, pp. 137–147, 2011.
- [208] N. Fujisawa, Y. Oguma, and T. Nakano, “Measurements of wall-shear-stress distribution on an NACA0018 airfoil by liquid-crystal coating and near-wall particle image velocimetry (PIV),” *Meas. Sci. Tech.*, vol. 20, no. 6, p. 065 403, 2009.
- [209] O. Reynolds, “On the extent and action of the heating surface of steam boilers,” *Proc. Lit. Soc. Manchester*, vol. 14, pp. 7–12, 1874.
- [210] H. Abe, H. Kawamura, and Y. Matsuo, “Surface heat-flux fluctuations in a turbulent channel flow up to $Re_\tau = 1020$ with $Pr = 0.025$ and 0.71 ,” *Int. J. Heat Mass Tran.*, vol. 25, no. 3, pp. 404–419, 2004.
- [211] H. Abe and R. A. Antonia, “Near-wall similarity between velocity and scalar fluctuations in a turbulent channel flow,” *Phys. Fluids*, vol. 21, no. 2, 2009.
- [212] N. Bagheri and B. R. White, “Experimental measurements of large-scale temperature fluctuation structures in a heated incompressible turbulent boundary layer,” *Int. J. Heat Mass Tran.*, vol. 36, no. 4, pp. 907–918, 1993.
- [213] R. Antonia, L. Krishnamoorthy, and L. Fulachier, “Correlation between the longitudinal velocity fluctuation and temperature fluctuation in the near-wall region of a turbulent boundary layer,” *Int. J. Heat Mass Tran.*, vol. 31, no. 4, pp. 723–730, 1988.
- [214] H. Thomann and B. Frisk, “Measurement of heat transfer with an infrared camera,” *Int. J. Heat Mass Tran.*, vol. 11, no. 5, pp. 819–826, 1968.
- [215] G. Hetsroni and R. Rozenblit, “Heat transfer to a liquid—solid mixture in a flume,” *Int. J. Multiphase Flow*, vol. 20, no. 4, pp. 671–689, 1994.
- [216] R. Gurka, A. Liberzon, and G. Hetsroni, “Detecting coherent patterns in a flume by using PIV and IR imaging techniques,” *Exp. Fluids*, vol. 37, pp. 230–236, 2004.
- [217] S. Yamada and H. Nakamura, “Construction of 2D-3C PIV and high-speed infrared thermography combined system for simultaneous measurement of flow and thermal fluctuations over a backward facing step,” *Int. J. Heat Mass Tran.*, vol. 61, pp. 174–182, 2016.
- [218] H. Nakamura, “Frequency response and spatial resolution of a thin foil for heat transfer measurements using infrared thermography,” *Int. J. Heat Mass Tran.*, vol. 52, no. 21-22, pp. 5040–5045, 2009.
- [219] H. Nakamura and S. Yamada, “Quantitative evaluation of spatio-temporal heat transfer to a turbulent air flow using a heated thin-foil,” *Int. J. Heat Mass Tran.*, vol. 64, pp. 892–902, 2013.
- [220] M. Raiola, C. S. Greco, M. Contino, S. Discetti, and A. Ianiro, “Towards enabling time-resolved measurements of turbulent convective heat transfer maps with IR thermography and a heated thin foil,” *Int. J. Heat Mass Tran.*, vol. 108, pp. 199–209, 2017.
- [221] F. Foroozan, A. Güemes, M. Raiola, R. Castellanos, S. Discetti, and A. Ianiro, “Synchronized measurement of instantaneous convective heat flux and velocity fields in wall-bounded flows,” *Meas. Sci. Tech.*, vol. 34, no. 12, p. 125 301, 2023.
- [222] T. Astarita and G. M. Carlomagno, *Infrared thermography for thermo-fluid-dynamics*. Springer, Berlin, Heidelberg, 2013.
- [223] J. R. Lloyd and W. R. Moran, “Natural Convection Adjacent to Horizontal Surface of Various Planforms,” *J. Heat Tran.*, vol. 96, no. 4, pp. 443–447, Nov. 1974.
- [224] J. Holman, “Heat transfer,” in (Mechanical engineering series), 7th, Mechanical engineering series. London: McGraw-Hill, 1990, pp. 333–352.

- [225] C. S. Greco, A. Ianiro, and G. Cardone, “Time and phase average heat transfer in single and twin circular synthetic impinging air jets,” *Int. J. Heat Mass Tran.*, vol. 73, pp. 776–788, 2014.
- [226] M. A. Mendez, M. Raiola, A. Masullo, *et al.*, “POD-based background removal for particle image velocimetry,” *Exp. Therm. Fluid Sci.*, vol. 80, pp. 181–192, 2017.
- [227] T. Astarita and G. Cardone, “Analysis of interpolation schemes for image deformation methods in PIV,” *Exp. Fluids*, vol. 38, pp. 233–243, 2005.
- [228] T. Astarita, “Analysis of weighting windows for image deformation methods in PIV,” *Exp. Fluids*, vol. 43, pp. 859–872, 2007.
- [229] M. A. Mendez, A. Ianiro, B. R. Noack, and S. L. Brunton, *Data-Driven Fluid Mechanics: Combining First Principles and Machine Learning*. Cambridge University Press, 2023.
- [230] S. Discetti and Y. Liu, “Machine learning for flow field measurements: A perspective,” *Meas. Sci. Technol.*, vol. 34, no. 2, p. 021 001, 2022.
- [231] R. Vinuesa, S. L. Brunton, and B. J. McKeon, “The transformative potential of machine learning for experiments in fluid mechanics,” *Nat. Rev. Phys.*, vol. 5, no. 9, pp. 536–545, 2023.
- [232] R. J. Adrian, “Stochastic estimation of the structure of turbulent fields,” in *Eddy structure identification*, Springer, 1996, pp. 145–195.
- [233] M. P. Encinar, A. Lozano-Durán, and J. Jiménez, “Reconstructing channel turbulence from wall observations,” in *Proceedings of the Summer Program. Center for Turbulence Research, Stanford University*, 2018.
- [234] S. J. Illingworth, J. P. Monty, and I. Marusic, “Estimating large-scale structures in wall turbulence using linear models,” *J. Fluid Mech.*, vol. 842, pp. 146–162, 2018.
- [235] J. A. Bourgeois, B. R. Noack, and R. J. Martinuzzi, “Generalized phase average with applications to sensor-based flow estimation of the wall-mounted square cylinder wake,” *J. Fluid Mech.*, vol. 736, pp. 316–350, 2013.
- [236] Z. Hosseini, R. J. Martinuzzi, and B. R. Noack, “Modal energy flow analysis of a highly modulated wake behind a wall-mounted pyramid,” *J. Fluid Mech.*, vol. 798, pp. 717–750, 2016.
- [237] S. Discetti, M. Raiola, and A. Ianiro, “Estimation of time-resolved turbulent fields through correlation of non-time-resolved field measurements and time-resolved point measurements,” *Experimental Thermal and Fluid Science*, vol. 93, pp. 119–130, 2018.
- [238] S. Discetti, G. Bellani, R. Örlü, *et al.*, “Characterization of very-large-scale motions in high-Re pipe flows,” *Exp. Therm. Fluid Sci.*, vol. 104, pp. 1–8, 2019.
- [239] K. Fukami, K. Fukagata, and K. Taira, “Machine-learning-based spatio-temporal super resolution reconstruction of turbulent flows,” *J. Fluid Mech.*, vol. 909, A9, 2021.
- [240] T. Nakamura, K. Fukami, and K. Fukagata, “Identifying key differences between linear stochastic estimation and neural networks for fluid flow regressions,” *Sci. Reps.*, vol. 12, no. 1, p. 3726, 2022.
- [241] E. Jagodinski, X. Zhu, and S. Verma, *Uncovering dynamically critical regions in near-wall turbulence using 3D convolutional neural networks*, 2020. arXiv: [2004.06187 \[physics.flu-dyn\]](https://arxiv.org/abs/2004.06187).
- [242] L. Guastoni, A. G. Balasubramanian, A. Güemes, *et al.*, *Non-intrusive sensing in turbulent boundary layers via deep fully-convolutional neural networks*, 2022. arXiv: [2208.06024 \[physics.flu-dyn\]](https://arxiv.org/abs/2208.06024).
- [243] D. Chen, X. Gao, C. Xu, *et al.*, “FlowGAN: A conditional generative adversarial network for flow prediction in various conditions,” in *2020 IEEE 32nd International Conference on Tools with Artificial Intelligence (ICTAI)*, IEEE, 2020, pp. 315–322.
- [244] T. Li, M. Buzdicotti, L. Biferale, and F. Bonaccorso, “Generative adversarial networks to infer velocity components in rotating turbulent flows,” *Eur. Phys. J. E*, vol. 46, no. 5, p. 31, 2023.
- [245] J. M. Hamilton, J. Kim, and F. Waleffe, “Regeneration mechanisms of near-wall turbulence structures,” *J. Fluid Mech.*, vol. 287, pp. 317–348, 1995.
- [246] J. Jiménez and A. Pinelli, “The autonomous cycle of near-wall turbulence,” *J. Fluid Mech.*, vol. 389, pp. 335–359, 1999.
- [247] Ó. Flores and J. Jiménez, “Hierarchy of minimal flow units in the logarithmic layer,” *Phys. Fluids*, vol. 22, no. 7, p. 071 704, 2010.
- [248] J. C. Del Álamo, J. Jiménez, P. Zandonade, and R. D. Moser, “Self-similar vortex clusters in the turbulent logarithmic region,” *J. Fluid Mech.*, vol. 561, pp. 329–358, 2006.
- [249] A. Lozano-Durán, Ó. Flores, and J. Jiménez, “The three-dimensional structure of momentum transfer in turbulent channels,” *J. Fluid Mech.*, vol. 694, pp. 100–130, 2012.

- [250] W. W. Willmarth and S. S. Lu, “Structure of the Reynolds stress near the wall,” *J. Fluid Mech.*, vol. 55, no. 1, pp. 65–92, 1972.
- [251] S. S. Lu and W. W. Willmarth, “Measurements of the structure of the Reynolds stress in a turbulent boundary layer,” *J. Fluid Mech.*, vol. 60, no. 3, pp. 481–511, 1973.
- [252] S. K. Lele, “Compact finite difference schemes with spectral-like resolution,” *J. Comp. Phys.*, vol. 103, no. 1, pp. 16–42, 1992.
- [253] J. Jiménez and P. Moin, “The minimal flow unit in near-wall turbulence,” *J. Fluid Mech.*, vol. 225, pp. 213–240, 1991.
- [254] H. Abe, H. Kawamura, and Y. Matsuo, “Direct numerical simulation of a fully developed turbulent channel flow with respect to the reynolds number dependence,” *J. Fluids Eng.*, vol. 123, no. 2, pp. 382–393, 2001.
- [255] J. C. Del Álamo and J. Jiménez, “Spectra of the very large anisotropic scales in turbulent channels,” *Phys. Fluids*, vol. 15, no. 6, pp. L41–L44, 2003.
- [256] A. W. Vreman and J. G. M. Kuerten, “Comparison of direct numerical simulation databases of turbulent channel flow at $Re_\tau = 180$,” *Phys. Fluids*, vol. 26, no. 1, 2014.
- [257] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [258] J. Kim, J. K. Lee, and K. M. Lee, “Deeply-recursive convolutional network for image super-resolution,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 1637–1645.
- [259] D. P. Kingma and J. Ba, *Adam: A method for stochastic optimization*, 2014. arXiv: [1412.6980 \[cs.LG\]](https://arxiv.org/abs/1412.6980).
- [260] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, *Improved Techniques for Training GANs*, 2016. arXiv: [1606.03498 \[cs.LG\]](https://arxiv.org/abs/1606.03498).
- [261] C. Sanmiguel Vila and Ó. Flores, “Wall-based identification of coherent structures in wall-bounded turbulence,” *J. of Phys.: Conference Series*, vol. 1001, no. 1, p. 012007, 2018.
- [262] Z. Wang, D. Liu, J. Yang, W. Han, and T. Huang, “Deep networks for image super-resolution with sparse prior,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 370–378.
- [263] C. Osendorfer, H. Soyer, and P. v. d. Smagt, “Image super-resolution with fast approximate convolutional sparse coding,” in *International Conference on Neural Information Processing*, Springer, 2014, pp. 250–257.
- [264] W. Shi, J. Caballero, F. Huszar, *et al.*, “Real-Time Single Image and Video Super-Resolution Using an Efficient Sub-Pixel Convolutional Neural Network,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [265] A. Cuéllar, A. Güemes, A. Ianiro, Ó. Flores, R. Vinuesa, and S. Discetti, “Three-dimensional generative adversarial networks for turbulent flow estimation from wall measurements,” *J. Fluid Mech.*, vol. 0, A1, 2024.
- [266] L. Prandtl, “Über flüssigkeitsbewegung bei sehr kleiner reibung,” *Verhandl. 3rd Int. Math. Kongr. Heidelberg, Germany*, pp. 484–491, 1904.
- [267] A. J. Favre, J. J. Gaviglio, and R. J. Dumas, “Space-time double correlations and spectra in a turbulent boundary layer,” *J. Fluid Mech.*, vol. 2, no. 4, pp. 313–342, 1957.
- [268] A. J. Favre, J. J. Gaviglio, and R. J. Dumas, “Further space-time correlations of velocity in a turbulent boundary layer,” *J. Fluid Mech.*, vol. 3, no. 4, pp. 344–356, 1958.
- [269] P. Holmes, J. L. Lumley, G. Berkooz, and C. W. Rowley, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge: Cambridge University Press, 2012.
- [270] A. Towne, A. Lozano-Durán, and X. Yang, “Resolvent-based estimation of space-time flow statistics,” *J. Fluid Mech.*, vol. 883, A17, 2020.
- [271] E. Martini, A. V. G. Cavalieri, P. Jordan, A. Towne, and L. Lesshafft, “Resolvent-based optimal estimation of transitional and turbulent flows,” *J. Fluid Mech.*, vol. 900, A2, 2020.
- [272] F. R. Amaral, A. V. G. Cavalieri, E. Martini, P. Jordan, and A. Towne, “Resolvent-based estimation of turbulent channel flow using wall measurements,” *J. Fluid Mech.*, vol. 927, A17, 2021.
- [273] R. Arun, H. J. Bae, and B. J. McKeon, “Towards real-time reconstruction of velocity fluctuations in turbulent channel flow,” *Phys. Rev. Fluids*, vol. 8, no. 6, p. 064612, 2023.
- [274] K. Manohar, J. N. Kutz, and S. L. Brunton, “Optimal sensor and actuator selection using balanced model reduction,” *IEEE Transactions on Automatic Control*, vol. 67, no. 4, pp. 2108–2115, 2022.
- [275] I. F. P. Bergman T. L. Lavine A. S. and D. D. P., *Fundamentals of Heat and Mass Transfer*. John Wiley & Sons, 2020.

- [276] W. E. Strafford J. and E. V., "Characterizing convective heat transfer using infrared thermography and the heated-thin-foil technique," *Meas. Sci. Tech.*, vol. 20, no. 10, p. 105401, 2009.
- [277] I. A. D. S. Torre A.F.M. and C. G.M., "Evaluation of anisotropic tangential conduction in printed-circuit-board heated-thin-foil heat flux sensors," *Int. J. Heat Mass Tran.*, vol. 127, pp. 1138–1146, 2018.
- [278] D. Á. J. C. and J. Jiménez, "Estimation of turbulent convection velocities and corrections to taylor's approximation," *J. Fluid Mech.*, vol. 640, pp. 5–26, 2009.
- [279] C. G. Serpieri J. and I. G., "Conditioning turbulent channel flows with wall plasma jets," in *AIAA aviation forum and ascend 2024*, 2024, p. 4384.
- [280] W. H. Y., *Handbook of essential formulae and data on heat transfer for engineers*. Longman, London, 1977.
- [281] J. C. Del Álamo and J. Jiménez, "Estimation of turbulent convection velocities and corrections to taylor's approximation," *J. Fluid Mech.*, vol. 640, pp. 5–26, 2009.
- [282] R. O. and P. J., "Thermal properties of paint coatings on different backings using a scanning photo acoustic technique," *Meas. Sci. Tech.*, vol. 17, no. 11, p. 2945, 2006.
- [283] M. Susa, X. Maldague, S. Svaic, I. Boras, and A. Bendada, "The influence of surface coatings on the differences between numerical and experimental results for samples subject to a pulse thermography examination," in *Proc. 9th Int. Conf. Quant. Infr. Thermogr.(QIRT)*, 2008, pp. 491–497.
- [284] C. R. B., "The scree test for the number of factors," *Multivar. Behav. Res.*, vol. 1, no. 2, pp. 245–276, 1966.