



HAL
open science

Autoencoder-based dimensionality reduction of turbulent channel flow under spanwise wall oscillations

Lou Guérin, Thomas Fisk, Laurent Cordier, Cédric Flageul, Lionel Agostini

► **To cite this version:**

Lou Guérin, Thomas Fisk, Laurent Cordier, Cédric Flageul, Lionel Agostini. Autoencoder-based dimensionality reduction of turbulent channel flow under spanwise wall oscillations. *Neural Computing and Applications*, 2026, 38 (6), pp.148. <10.1007/s00521-026-11856-z>. <hal-05588434>

HAL Id: hal-05588434

<https://hal.science/hal-05588434v1>

Submitted on 16 Apr 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Autoencoder-Based Dimensionality Reduction of Turbulent Channel Flow Under Spanwise Wall Oscillations

Lou Guérin, Thomas Fisk, Laurent Cordier, Cédric Flageul, Lionel
Agostini

Pprime institute, Curiosity team, CNRS, Université de Poitiers,
Poitiers, France.

Abstract

This study investigates the efficacy of AutoEncoder (AE) neural networks in reducing the dimensionality of a heated turbulent channel flow subjected to Spanwise Wall Oscillations (SWO) at $Re_\tau = 200$. The dataset, generated using implicit Large Eddy Simulations (LES), comprises velocity and temperature fields at various oscillation amplitudes. Convolutional AutoEncoders (CAEs) are compared to the benchmark Proper Orthogonal Decomposition (POD) method to assess flow complexity and quantify improvements in reconstruction accuracy. β -Variational AutoEncoders (β -VAEs) are employed to enhance the orthogonality of the latent samples, aiming to approach the uncorrelated modes of POD. The results demonstrate that CAEs consistently outperform POD in identifying a more compact coordinate system, capturing smaller flow scales with greater accuracy, especially for the wall-normal velocity component. Despite significant pruning of the latent variables, the β -VAE achieves a fine balance between the orthogonality of POD and the reconstruction capabilities of the CAE. Compared to POD, autoencoders provide superior separability of latent samples in two distinct aspects: (i) more effective amplitude separation, and (ii) more physically coherent phase separation. The CAE in particular leads to the best performance in both respects. This research highlights the potential of autoencoders in dimensionality reduction for complex turbulent flows and provides insights into the trade-off between reconstruction accuracy and latent-space interpretability.

Keywords: Dimensionality reduction, autoencoders, proper orthogonal decomposition, large eddy simulations, turbulent channel flow, spanwise wall oscillations.

31 1 Introduction

32 Societal Context

33 Wall-bounded turbulent flows are prevalent across engineering systems, from aero-
34 dynamic applications to heat exchangers, where their behaviour fundamentally
35 determines system performance. In these flows, turbulence manifests through enhanced
36 momentum mixing in the near-wall region, resulting in increased friction drag along
37 solid surfaces. This drag force becomes more pronounced at higher Reynolds numbers,
38 significantly impacting energy consumption in numerous applications [1]. The trans-
39 portation sector, in particular, faces substantial efficiency losses due to friction drag,
40 affecting both external aerodynamics and internal flows [2]. As environmental con-
41 cerns intensify and regulations on transport-related emissions become more stringent,
42 the mitigation of friction drag has emerged as a critical engineering challenge [3]. Par-
43 allel to drag reduction, the enhancement of heat transfer efficiency presents another
44 crucial engineering objective, particularly in thermal management systems. While tur-
45 bulent flows naturally promote heat transfer through increased near-wall mixing [4],
46 the traditional coupling between momentum and thermal transport often results in a
47 compromise: attempts to reduce drag typically lead to decreased heat transfer rates,
48 and vice versa. This intricate relationship necessitates innovative control strategies
49 that can effectively manage both phenomena.

50 Spanwise wall oscillations strategy

51 The near-wall region of turbulent flows exhibits distinct coherent structures that sig-
52 nificantly influence momentum mixing and heat transfer processes. Understanding
53 these structures and their interactions is fundamental for developing effective flow
54 control strategies. Spanwise wall oscillations (SWO) have emerged as a promising
55 control technique, capable of modifying these coherent structures [1]. By generating
56 a spanwise strain layer, SWO disrupts the formation of near-wall streaks, attenu-
57 ates quasi-streamwise vortices, and increases viscous sublayer thickness [5–7]. The
58 effectiveness of this control approach depends on multiple parameters, including oscil-
59 lation amplitude (W^+) and period (T^+). Numerical and experimental studies have
60 demonstrated drag reduction of up to 40-50% at low Reynolds numbers [8, 9].

61 A study by Gu erin et al. [10] has demonstrated that particular oscillation param-
62 eters ($T^+ = 500$, $W^+ = 30$) can produce heat-transfer enhancement at approximately
63 double the rate of drag increase at $Pr = 1$, indicating a breakdown of the Reynolds
64 analogy. The recent work of Rouhi et al. [11] systematically examined the influence
65 of oscillation frequency and Prandtl number for both SWO and streamwise travelling
66 waves, at fixed $Re_\tau = 590$ and $W^+ = 12$, with particular focus on simultaneous reduc-
67 tion of both heat transfer and drag. Their results reveal that for $Pr > 1$, heat transfer
68 reduction can exceed drag reduction significantly, for example reaching a 10 % supe-
69 rior reduction at $Pr = 7.5$. The authors employ the Stokes layer protrusion height
70 ($\ell_{0.01}^*$), originally introduced by Rouhi et al. [12] as the elevation where \overline{w}^{2*} drops to
71 about 1 % of the spanwise turbulent stress, rather than the conventional Stokes layer
72 thickness [13]. This choice is fundamentally motivated by the protrusion height's near

73 proportionality with the maximum Stokes layer production - the sole external source
74 term that directly injects energy into the turbulent stress budgets. As the oscillation
75 frequency decreases, the Stokes layer becomes more protrusive. In this case, when
76 $Pr > 1$, the conductive sublayer is thinner, thus exposing the energetic small scales
77 of $\Theta'v'$ to larger Stokes layer production. Consequently, these temperature fluctua-
78 tions gain more energy compared to their velocity counterparts, ultimately resulting
79 in greater reduction of drag than heat transfer.

80 **Reduced-order modelling**

81 The investigation of turbulent flows often relies on Direct Numerical Simulation
82 (DNS), which resolves all ranges of spatio-temporal structures populating the flow.
83 However, for wall-bounded flows, DNS computational resources requirements scale
84 with the Reynolds number Re as Re^3 , making it prohibitively expensive, even
85 untenable, for practical applications, particularly when exploring multiple control
86 parameters in flow control strategies such as SWO [6]. This computational limitation
87 has motivated the development of more efficient numerical approaches. Among them,
88 Reduced-Order Models (ROMs), which address this challenge through a two-step
89 process:

- 90 • The first one consists of identifying a more efficient coordinate system where the
91 flow dynamics can be represented with fewer degrees of freedom. These purely
92 *data-driven* methods have demonstrated remarkable efficiency in capturing essential
93 dynamics within a low-dimensional manifold. This dimensionality reduction phase
94 aims to capture the essential flow physics within a compact representation. Traditional
95 approaches such as Proper Orthogonal Decomposition (POD) achieve this
96 through linear projections [14], whereas neural network architectures like autoen-
97 coders can capture non-linear relationships, potentially leading to more compact
98 representations [15].
- 99 • The second step leverages this compact representation to determine the underlying
100 physics governing the flow evolution. This involves identifying relationships between
101 elements in the reduced coordinate system, effectively establishing a simplified
102 dynamical model.

103 ROMs have become particularly attractive for flow control applications, where
104 computational efficiency is paramount for determining optimal control strategies. For
105 instance, a data-driven ROM, called DManD ROM [16], employs hybrid autoencoders
106 to find a more compact coordinate system for a turbulent minimal Couette flow,
107 reducing the dimension from 10^5 to 18 DOFs. The combination of DManD ROM
108 with the soft actor-critic reinforcement learning algorithm [17] has achieved significant
109 speed-up compared to DNS while maintaining control strategy effectiveness. The dis-
110 tribution of parameters used to create a data-driven ROM plays an essential role in
111 ensuring robustness, both in the dimensionality reduction phase and when predicting
112 flow dynamics [18].

113 Recent Advances in Neural Network Architectures

114 Recent advances in neural-network architectures have significantly enhanced ROM
115 capabilities. Traditionally, POD has been the primary technique for dimensionality
116 reduction, generating compact coordinate systems through orthogonal modes that
117 capture the most energetic flow structures [14]. While POD ensures interpretability
118 by providing orthonormal spatial modes and temporal coefficients, its reliance on
119 linear projections presents limitations when addressing complex, multi-scale, nonlinear
120 turbulent systems.

121 Autoencoders have emerged as a promising alternative [19], with their non-linear
122 activation functions proving particularly effective in capturing coherent structures in
123 turbulent flows [20]. Convolutional autoencoders demonstrate notable effectiveness in
124 detecting spatial patterns and correlations. Studies comparing POD and autoencoder
125 performance [21, 22] consistently demonstrate that autoencoders typically require
126 fewer modes to capture equivalent system energy, though their non-linear nature often
127 results in non-orthogonal modes that reduce model interpretability.

128 β -Variational autoencoders present a balanced approach between the interpretabil-
129 ity of POD and the compact representation of autoencoders [23]. Through the use of
130 a probabilistic encoder and the incorporation of a Kullback-Leibler (KL) divergence
131 term in the loss function, weighted by parameter β , these models encourage mode
132 decorrelation by minimising the divergence between the latent distribution and a stan-
133 dard Gaussian [24]. The latent distribution is thus encouraged to remain close to a
134 centered isotropic distribution, favouring a denser latent space, beneficial for coher-
135 ent sample generation via latent interpolation [25]. Eivazi et al. [24], Kang et al. [25]
136 each propose an algorithm that ranks latent variables with regards to their contribu-
137 tion to reconstruction, providing a reliable tool for enhancing model interpretability.
138 More recently, there has been a strong focus on enhancing the physical coherence of
139 AEs. For instance, Mohan et al. [26] incorporated the Helmholtz decomposition and
140 fixed neural network weights to enforce incompressibility in the autoencoder recon-
141 struction. Similarly, Fukami and Taira [27] introduced a lift-augmented autoencoder,
142 where the latent space directly encodes lift-related information, achieved by integrat-
143 ing a multi-layer-perceptron that is simultaneously trained to predict lift from the
144 latent representation.

145 The broader landscape of modern dimensionality reduction encompasses meth-
146 ods such as Koopman operator-based approaches, which seek to linearize nonlinear
147 dynamics [28], as well as sparse regression techniques like the Sparse Identification of
148 Nonlinear Dynamics (SINDy) [29] and its hybrid architecture, SINDy-SHRED [30],
149 which derive interpretable, parsimonious governing equations directly from data.

150 Manuscript Structure

151 This study presents a comprehensive investigation of dimensionality reduction tech-
152 niques, comparing traditional POD approaches with advanced neural network archi-
153 tectures. The analysis focuses on a heated turbulent channel flow described in Section

154 2. The dataset comprises two-dimensional slices of velocity components and temper-
155 ature fields, with flow control implemented through spanwise wall oscillations at a
156 period of $T^+ = 500$.

157 The investigation employs 21 numerical simulations with systematic variations in
158 oscillation amplitude W^+ , ensuring robust validation across a range of control param-
159 eters. The manuscript structure follows a systematic progression through methodology
160 and results:

- 161 • Section 2 presents the numerical framework, detailing flow conditions, simulation
162 parameters and the dataset derived from the simulations.
- 163 • Section 3 describes the implementation of POD and autoencoder variants. Particular
164 attention is given to the non-linear capabilities of neural architectures in capturing
165 coherent structures and their potential advantages and disadvantages over linear
166 POD decomposition.
- 167 • In Section 4, the performance of CAEs is evaluated against POD, with a focus
168 on reconstruction accuracy, the capacity to resolve different flow scales, and the
169 minimal number of modes required to represent dominant flow structures.
- 170 • Section 5 investigates β -VAE architectures, focusing on their ability to provide
171 latent variable pruning and to generate uncorrelated modes through KL-divergence
172 regularisation.
- 173 • In Section 6, the aptitude of each AE variants and POD to effectively separate latent
174 samples based on their oscillations amplitudes and phases is assessed.

175 The primary focus of this study is to establish optimal dimensionality reduction
176 strategies rather than a dynamical model with temporal evolution prediction. This
177 investigation aims to develop compact, physically meaningful representations of tur-
178 bulent flows under controlled conditions, with particular emphasis on maintaining
179 interpretability in the reduced-order space.

Table 1 Table of Notations (Part 1)

Symbol	Description
General Mathematical Notation	
$\langle \bullet \rangle$	Ensemble averaging operator
$\overline{\bullet}$	Time and periodic directions averaged quantity ($\overline{\bullet} = \langle \bullet \rangle_{t,x,z}$)
χ'	Fluctuating component of a variable χ ($\chi' = \chi - \overline{\chi}$)
$\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$	Multivariate Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$
Flow Variables & Parameters	
u, v, w	Streamwise, wall-normal, and spanwise velocity components
p	Pressure
Θ	Temperature
Re	Reynolds number (based on bulk velocity U_b)
Re_τ	Friction Reynolds number ($Re_\tau = Hu_\tau/\nu$)
H	Channel half-height
u_τ	Friction velocity ($u_\tau = \sqrt{\frac{1}{Re} \frac{\partial \overline{u_w}}{\partial y}}$)
Θ_τ	Friction temperature ($\Theta_\tau = \frac{\alpha}{u_\tau} \frac{\partial \overline{\Theta_w}}{\partial y}$)
ν, α	Kinematic viscosity, thermal diffusivity
Pr	Prandtl number
C_f	Friction coefficient ($C_f = \frac{9}{2Re} \frac{\partial \langle u_w \rangle_{x,z}}{\partial y}$)
Nu	Nusselt number ($Nu = 2 \frac{\partial \langle \Theta_w \rangle_{x,z}}{\partial y}$)
$\Delta C_f, \Delta Nu$	Percentage variation of C_f and Nu with actuation
W^+	Oscillation amplitude in wall units
T^+	Oscillation period in wall units
φ	Oscillation phase
χ_w	Quantity at the wall ($y = 0, y = 2H$)
χ^+	Quantity χ normalized in wall units
χ_{nom}^+	Quantity χ normalized in wall units using unactuated friction quantities
χ_{ac}^+	Quantity χ normalized in wall units using actuated friction quantities
Performance Metrics	
MSE	Mean Squared Error loss function (10)
$E_r(\chi)$	Energy reconstruction percentage for field χ (16)
$D_{\text{KL}}[\bullet \bullet]$	Kullback-Leibler divergence over two probability distributions
M	Latent variable correlation matrix
$M_{m,n}$	Correlation coefficient between latent variables m and n (17)
Mean _c	Mean Correlation (average off-diagonal of $ M $) (19)
MeanMax _c	Mean-Max Correlation (20)
D_B	Davies-Bouldin clustering index (21)
$D_B(\mathcal{C}_n, \mathcal{C}_m)$	Davies-Bouldin clustering index between cluster \mathcal{C}_n and \mathcal{C}_m (22)
$C(\phi, LV(i))$	Pearson correlation between physical quantity ϕ and latent variable i (24)

Table 2 Table of Notations (Part 2)

Symbol	Description
Dataset & Indices	
$\chi_{j,k,l}^{(i)}$	Data point for field i , simulation j , snapshot k , grid point l
$i \in \{1, 2, 3, 4\}$	Index for fields: u, v, w, Θ
$j \in \{0, \dots, 20\}$	Index for simulation / oscillation amplitude W^+
$k \in \{0, \dots, 7680\}$	Index for snapshots: subdomain (12) / actuation cycle (20) / actuation phase (32)
$l \in \{1, \dots, 2048\}$	Index for spatial grid point: $(n_x, n_z) = (64, 32)$
$\langle \chi_{j,k,l}^{(i)} \rangle_{k,l}$	Average of dataset samples with field i and simulation j over snapshots k , and grid points l
$\chi_{j,k,l}^{(i) \prime}$	Fluctuations of $\chi_{j,k,l}^{(i)}$ over snapshots k , and grid points l $(\chi_{j,k,l}^{(i) \prime} = \chi_{j,k,l}^{(i)} - \langle \chi_{j,k,l}^{(i)} \rangle_{k,l})$
$\widetilde{\chi_{j,k,l}^{(i)}}$	Model reconstruction of $\chi_{j,k,l}^{(i)}$
μ_i, σ_i	Mean and standard deviation of field i , over simulations j , snapshots k , and gridpoints l ($\mu_i = \langle \chi_{j,k,l}^{(i)} \rangle_{j,k,l}$, $\sigma_i = \sqrt{\langle (\chi_{j,k,l}^{(i)} - \mu_i)^2 \rangle_{j,k,l}}$)
$\chi_{\text{std},j,k,l}^{(i)}$	Field wise standardized data (6)
$\chi_{\text{norm},j,k,l}^{(i)}$	Final normalized version data (7)
Dimensionality Reduction Models	
L_d	Latent space dimension
\mathbf{z}	Latent space vector
$\boldsymbol{\mu}, \boldsymbol{\sigma}$	Latent mean and standard deviation vectors (VAE)
$z_{j,k}(n)$	n -th coordinate of the latent vector with simulation index j and snapshot index k
$\langle z_{j,k} \rangle_{j,k}$	Average of the latent vectors over simulations j and snapshots k
$z'_{j,k}$	Fluctuations of latent sample $z_{j,k}$ over simulations j and snapshots k $(z'_{j,k} = z_{j,k} - \langle z_{j,k} \rangle_{j,k})$
U, Σ, V	SVD matrices (left singular vectors, singular values, right singular vectors)
$q_{\psi}(\mathbf{z} \mathbf{x})$	Encoder's approximate posterior distribution (VAE)
$q_{\psi}(\mathbf{z})$	Marginal latent distribution (VAE)
$p_{\theta}(\mathbf{x} \mathbf{z})$	Decoder's likelihood distribution (VAE)
$p_{\theta}(\mathbf{z})$	Prior distribution (standard Gaussian $\mathcal{N}(\mathbf{0}, \mathbf{I})$)
$p_{\theta}(\mathbf{x})$	Underlying distribution

181 2 Heated turbulent channel flow under spanwise 182 wall oscillation

183 This section provides detailed information on the flow conditions and numerical meth-
184 ods used to generate the dataset. Furthermore, it outlines the pre-processing steps and
185 explains the rationale behind the decisions that shaped the dataset structure.

186 2.1 Flow conditions

The flow conditions for the present study are comprehensively described in Guérin et al. [10, 31], and are briefly outlined here for convenience. Large Eddy Simulations of




Fig. 1 Schematic of the flow configuration.

a canonical channel flow are carried out at $\text{Re}_\tau = \frac{H u_\tau}{\nu} \approx 200$ where u_τ is the friction velocity and ν the kinematic viscosity. The numerical domain is defined as:

$$\Omega = \left\{ \mathbf{x} = (x \ y \ z)^\top \in \mathbb{R}^3 \mid x \in [0, L_x], \ y \in [0, 2H], \ z \in [0, L_z] \right\},$$

187 where x , y , and z represent the streamwise, wall-normal, and spanwise directions,
 188 respectively (see Fig. 1). The velocity, pressure, and temperature fields are denoted
 189 by $\mathbf{u}(\mathbf{x}, t) = (u(\mathbf{x}, t) \ v(\mathbf{x}, t) \ w(\mathbf{x}, t))^\top$, $p(\mathbf{x}, t)$, and $\Theta(\mathbf{x}, t)$.

190 The evolution of the velocity \mathbf{u} is governed by the incompressible Navier-Stokes
 191 equations, while the evolution of the temperature Θ follows the passive transport
 192 equation of a scalar by the velocity field \mathbf{u} . Assuming a flow density equal to 1 and
 193 adopting the non-dimensional framework of Gu erin et al. [10], the following set of
 194 governing equations is obtained. All quantities are expressed in dimensionless form,
 195 and thus no physical units are associated with them.

$$\begin{cases} \nabla \cdot \mathbf{u} = 0 \\ \frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \frac{1}{\text{Re}} \Delta \mathbf{u} + \mathbf{f} \\ \frac{\partial \Theta}{\partial t} + \mathbf{u} \cdot \nabla \Theta + f_\Theta = \frac{1}{\text{Re Pr}} \Delta \Theta, \end{cases} \quad (1)$$

196 where $\text{Re} = \frac{H U_b}{\nu}$, with U_b with the bulk velocity. Pr denotes the Prandtl number,
 197 which quantifies the ratio of momentum to heat diffusion. To simplify the problem and
 198 promote initial similarity between temperature and streamwise velocity, the Prandtl
 199 number is set to 1. To further promote this unactuated similarity between thermal and
 200 velocity components, the mixed boundary conditions [32] are employed for the temper-
 201 ature. By definition, χ_w represents the boundary conditions imposed on any variable χ
 202 on the walls, *i.e.* $\chi_w = \{\chi(\mathbf{x}) \mid x \in [0, L_x], \ z \in [0, L_z] \text{ and } y = 0 \text{ or } y = 2H\}$. Period-
 203 icity is assumed in the x and z directions, while the top and bottom wall satisfy no-slip

204 boundary conditions. As mentioned in Section 1, in the actuated case, the spanwise
 205 velocity is set to an in-phase periodic oscillation, varying in time. This control action
 206 depends on two parameters, the oscillation period T , and the amplitude W :

$$\begin{aligned}\Theta_w &= 0, \\ \mathbf{u}_w(t) &= (0, 0, W \sin(\frac{2\pi t}{T}))^\top.\end{aligned}\tag{2}$$

207 Variables denoted with a superscript $+$ indicates values scaled by wall units, *i.e.*
 208 non dimensionalized by the friction velocity $u_\tau = \sqrt{\frac{1}{\text{Re}} \frac{\partial \overline{u_w}}{\partial y}}$, where $\frac{\partial \overline{u_w}}{\partial y}$ is the wall-
 209 normal derivative of the averaged streamwise velocity at the wall, and the friction
 210 temperature $\Theta_\tau = \frac{\alpha}{u_\tau} \frac{\partial \overline{\Theta_w}}{\partial y}$, where α is the thermal diffusivity and $\frac{\partial \overline{\Theta_w}}{\partial y}$ is the derivative
 211 of the averaged wall temperature. Due to the periodic time-dependence introduced
 212 by the wall oscillations, two distinct variants of wall units are employed for the non-
 213 dimensionalization of the actuated flow. These are distinguished by the friction velocity
 214 (or temperature) used for scaling: χ_{nom}^+ uses the nominal friction velocity from the
 215 unactuated flow ($u_{\tau,\text{nom}}$), providing a fixed reference, while χ_{ac}^+ uses the instantaneous
 216 friction velocity of the actuated flow ($u_{\tau,\text{ac}}$), which captures the modulated near-wall
 217 dynamics.

218 The friction coefficient C_f and Nusselt number Nu characterize drag and heat
 219 transfer, respectively. In this configuration, these quantities are defined by:

$$C_f(t) = \frac{9}{2\text{Re}} \frac{\partial \langle u_w \rangle_{x,z}}{\partial y}, \quad \text{Nu}(t) = 2 \frac{\partial \langle \Theta_w \rangle_{x,z}}{\partial y}.\tag{3}$$

220 In Guérin et al. [10], the parameters $T^+ = 500$ and $W^+ = 30$ led to a preferential
 221 increase in Nu over C_f . The variation of these quantities induced by SWO can be
 222 expressed as a percentage in the following quantity:

$$\Delta C_f = 100 \times \left(\frac{\langle C_{f,\text{ac}} \rangle_t}{\langle C_{f,\text{unac}} \rangle_t} - 1 \right),\tag{4}$$

223 where $C_{f,\text{ac}}$ and $C_{f,\text{unac}}$ account for the friction coefficient of the controlled and
 224 unactuated flow, respectively. ΔNu is defined analogously, using Nu_{ac} and Nu_{unac} .

225 2.2 Numerical methods

226 In this study, the numerical simulations are carried out using the open-source Xcom-
 227 pact3d framework [33–35]. The spatial discretization employs sixth-order compact
 228 finite difference schemes, while the time integration is handled by a third-order explicit
 229 Runge-Kutta method. To ensure zero velocity divergence, a fractional step method is
 230 used, which solves a Poisson equation for the pressure gradient through 3D fast fourier
 231 transforms [33]. Given the extensive parameter space investigation required for W^+
 232 variations in this study, implicit LES are employed instead of DNS to lower computa-
 233 tional costs. This approach also effectively reduces the dimensionality of the original
 234 flow space representation. Details of the coarse mesh utilized in the LES can be found
 235 in Table 3, along with DNS comparison references [10, 36, 37].

Table 3 Simulations computational conditions, against references [10, 36, 37].

Cases	(Lx, Ly, Lz)	Δx^+	Δy_{\min}^+	Δy_{\max}^+	Δz^+	Δt^+	t_{collect}^+	Re_τ
[36]	(12.8, 2, 6.4)	3.0	0.20	5.93	3.0	NA	3960	180
[37]	(6.4, 2, 3.2)	1.1	0.05	0.97	1.1	NA	1677	180
[10]	(24, 2, 6)	10.68	0.43	6.16	5.34	$1.56 \cdot 10^{-2}$	> 10000	179
S_1	(8, 2, 6)	11.0	2.0	15.7	11.0	$7.3 \cdot 10^{-2}$	10000	196

236 The DNSs conducted by Gu erin et al. [10] employ a significantly finer mesh
237 compared to the present S_1 configuration. Specifically, the mesh resolution in the
238 streamwise (Δx), spanwise (Δz), and the two wall-normal directions (Δy_{\min} and
239 Δy_{\max}) is enhanced by factors of 1.03, 2.06, 4.65, and 2.55, respectively. Additionally,
240 the computational domain extends three times further in the streamwise direction,
241 providing a more detailed resolution of the turbulent structures. Consistent with the
242 methodology of Flageul et al. [38], the S_1 mesh incorporates uniform numerical dissipa-
243 tion across all spatial directions. This dissipation is implemented using a fourth-order
244 accurate spectral vanishing viscosity scheme [39]. The integration of added numeri-
245 cal dissipation, combined with the coarser mesh refinement of the S_1 configuration,
246 classifies these simulations as implicit LES. As validated in Section 2.3, the S_1 mesh
247 configuration leads to results comparable to those obtained from the higher-resolution
248 grid employed in the reference DNS studies.

249 Each simulation starts with a warm-up period of five oscillation cycles, equivalent
250 to 2500 nominal wall-unit time steps, ensuring that an equilibrium state is achieved.
251 Following equilibrium, an additional 20 oscillation cycles are conducted to gather data
252 snapshots. During the data collection process, 32 snapshots are recorded at regular
253 intervals within each actuation cycle, capturing the flow dynamics comprehensively
254 across 32 distinct phases (see Section 6 for phase separability analysis). All simulations
255 are initiated from a common set of input snapshots, derived from the steady-state of
256 an unactuated simulation using the S_1 configuration.

257 The current numerical strategy balances computational efficiency with the need
258 to accurately capture essential flow features, particularly the interactions between
259 turbulent structures and spanwise wall oscillations.

260 2.3 Validation

261 The validation of numerical simulations entails a systematic comparison of unactuated
262 flow statistics against established reference data. Validation of the velocity statistics is
263 already conducted in Gu erin et al. [31], whereas the temperature statistics are evalu-
264 ated against the numerical results of Seki et al. [37]. The assessment of controlled flow
265 characteristics, specifically heat transfer and drag coefficients, is conducted through
266 comparison with the data reported by Gu erin et al. [10]. The validation procedure
267 necessitates a precise definition of the averaging methodology and associated flow
268 fluctuations. For this purpose, a generalised decomposition is applied to any flow vari-
269 able χ (which may represent velocity components, temperature, or pressure) into its
270 mean and fluctuating components. This decomposition is expressed mathematically
271 as follows:

$$\chi(x, y, z, t) = \bar{\chi}(y) + \chi'(x, y, z, t), \quad (5)$$

272 where $\bar{\chi} = \langle \chi \rangle_{x,z,t}$ represents the time, streamwise and spanwise-averaged value at
 273 a specific wall-normal position y , and χ' denotes the fluctuations around this mean.




Fig. 2 Wall-normal distribution of (a) mean temperature $\bar{\Theta}$, (b) turbulent heat flux $-\overline{\Theta'v'}$. Statistics obtained from simulations on S_1 are represented by the plain blue line. Reference results from Seki et al. [37] are depicted by the black line with crosses.

274 Fig. 2 presents a comparison between the unactuated simulation statistics and
 275 those reported by Seki et al. [37]. The results demonstrate a high level of agreement,
 276 with minor discrepancies primarily attributed to differences in Reynolds numbers. It is
 277 noteworthy that the discrepancies are more pronounced for the mean temperature $\bar{\Theta}$
 278 compared to the mean streamwise velocity \bar{u} [31], in the vicinity of its peak. However,
 279 the turbulent heat flux $-\overline{\Theta'v'}$, which is directly related to the Nusselt number Nu ,
 280 exhibits minor discrepancies. Specifically, in the unactuated case, the DNS resolution
 281 from Guérin et al. [10] at $Re_\tau = 200$ shows a 4.7% and 4% difference in C_f and Nu ,
 282 respectively, when compared to the S_1 resolution. These results demonstrate sufficient
 283 accuracy, particularly considering the computational efficiency gained through the
 284 S_1 mesh configuration (see Guérin et al. 31), thus validating the robustness of the
 285 numerical approach.

Table 4 Comparison of the percentage variation in the Nusselt number (ΔNu) and friction coefficient (ΔC_f) between the S_1 mesh configuration and the reference DNS by Guérin et al. [10] for oscillation parameters $T^+ = 500$, $W^+ = 30$.

Configuration	ΔNu	ΔC_f
S_1	18.11%	9.79%
Guérin et al. [10]	14.93%	7.72%

286 Table 4 presents a comparative analysis of the variations in averaged Nusselt
 287 number (Nu) and friction coefficient (C_f) induced by SWO between the S_1 mesh
 288 configuration and the DNS results from Gu erin et al. [10]. The comparison focuses
 289 on results for oscillation parameters given by $T^+ = 500$ and $W^+ = 30$. The S_1
 290 configuration demonstrates strong correlation with the reference DNS data, albeit
 291 with marginally enhanced variations in both Nu and C_f . The equilibrium state and
 292 transitory characteristics maintain consistency with the reference study.

293 The LES accuracy was verified for two boundary conditions: the unactuated state
 294 ($W^+ = 0$) and maximum actuation amplitude ($W^+ = 30$), for which the Stokes layer
 295 protrusion height is at its peak, exerting a more pronounced influence on the flow
 296 dynamics, particularly for an actuation period of $T^+ = 500$. In this investigation, sim-
 297 ulations will be performed using oscillation amplitudes within the range $W^+ \in [0, 30]$.
 298 The demonstrated agreement between LESs and DNSs solutions at these extrema
 299 substantiates the capability of the S_1 mesh configuration to accurately resolve the
 300 fundamental flow physics under both unactuated and controlled conditions, thereby
 301 ensuring reliable predictions of the wall-bounded turbulent flow behaviour throughout
 302 the parametric space.

303 2.4 Dataset

304 The present investigation utilizes a comprehensive dataset generated through 21
 305 implicit LESs using the S_1 configuration detailed in Table 3. The oscillation period
 306 is fixed at $T^+ = 500$. The restriction to amplitude variations whilst maintaining a
 307 fixed oscillation period ($T^+ = 500$) represents a deliberate choice grounded in the
 308 underlying physics of the flow-control mechanism. As shown by Gu erin et al. [10], the
 309 combination of $W^+ = 30$ and $T^+ = 500$ yielded a preferential heat-transfer enhance-
 310 ment nearly twice that of the corresponding drag increase. This configuration is not
 311 arbitrary but reflects the requisite spatial and temporal scales necessary for effective
 312 modulation of near-wall coherent structures.

313 The oscillation period T^+ governs two fundamental aspects of the control mecha-
 314 nism: the spatial penetration of the Stokes layer and the temporal synchronisation with
 315 the self-sustaining streak-formation process, populating the buffer layer. In contrast,
 316 the oscillation amplitude W^+ directly controls the intensity of the forcing, i.e. the ener-
 317 getic magnitude of the perturbation applied to the near-wall streaks. Variations in W^+
 318 modulate the strength of streak disruption and reorganisation, which in turn governs
 319 the degree of momentum and thermal mixing enhancement. As the amplitude param-
 320 eter represents a more straightforward control of the forcing intensity acting upon the
 321 physically targeted zone already optimised through the selection of $T^+ = 500$, the
 322 present parametric investigation is focused on W^+ variations within the range $[0, 30]$
 323 (using a regular subdivision of this interval: $W^+ \in \{x = 1.5 \times i \mid i \in \{0, \dots, 20\}\}$).
 324 This choice permits systematic examination of the relationship between forcing
 325 strength and flow modification whilst maintaining spatial and temporal targeting
 326 established by the fixed oscillation period.

327 Furthermore, the amplitude parameter exhibits more predictable and systematic
 328 behaviour than the period, rendering it a more suitable focus for a preliminary investi-
 329 gation into model robustness and interpolation capabilities. The decision to vary solely

330 the oscillation amplitude is thus rooted in the physics of the flow-control mechanism
331 and represents a methodologically sound approach to exploring the parameter space.

332 Each simulation generated 640 snapshots comprising velocity components u , v , w ,
333 and temperature Θ . The data reduction methodology employs a systematic hierar-
334 chical approach to efficiently compress the high-dimensional flow field data. Initially,
335 the three-dimensional computational snapshots, encompassing approximately 800,000
336 grid points per field, undergo a primary dimensionality reduction by extracting two-
337 dimensional slices at the wall-normal position $y_{ac}^+ \approx 15$ from both channel boundaries.
338 This specific wall-normal location is selected because the small-scale coherent struc-
339 tures, which constitute the predominant turbulent features at the present Reynolds
340 number, exhibit maximum intensity at this distance from the wall. The strategic selec-
341 tion of these wall-normal positions is fundamentally important for characterising the
342 enhancement mechanisms of thermal and momentum transport induced by spanwise
343 wall oscillations. At these locations, the manifestation of important turbulent phenom-
344 ena—including streak intensity modulation, structural distortions, and alterations in
345 sweep and ejection events—reaches peak intensity [10]. Consequently, predicated on
346 the well-established premise that the two-dimensional slices at $y_{ac}^+ \approx 15$ encapsulate
347 the essential physical information of the flow field, the compression methodology is
348 specifically applied to these planar sections rather than the complete three-dimensional
349 channel domain. This targeted approach confers substantial computational efficiency,
350 significantly reducing the processing requirements associated with the compression of
351 the full three-dimensional flow field. This initial dimensionality reduction diminishes
352 the data volume to $128 \times 96 = 12,288$ grid points per field.

353 To further enhance computational tractability, a subsequent spatial decomposi-
354 tion partitions the computational domain into six distinct subdomains: two sections
355 along the streamwise direction (x -axis) and three sections along the spanwise direction
356 (z -axis). As evidenced in Fig. 3, this strategic partitioning preserves the integrity of
357 coherent structures within each subdomain, thereby maintaining the essential physical
358 characteristics of the flow. Following this decomposition, the resultant snapshots com-
359 prise $64 \times 32 = 2,048$ grid points per field. The four primary flow fields (u , v , w and Θ)
360 are subsequently consolidated into a unified three-dimensional array with dimensions
361 $64 \times 32 \times 4 = 8,192$ degrees of freedom, wherein the third dimension represents the
362 stacked field variables. This integrated approach to dimensionality reduction demon-
363 strates superior efficiency compared to individual field compression methodologies, as
364 comprehensively documented in Section 4.

365 The complete dataset comprises 161,280 samples, derived from 7,680 samples
366 per simulation. This 2D field extraction and domain decomposition strategy, com-
367 bined with the implicit-LES approach, constitutes the initial phase of dimensionality
368 reduction while maintaining sufficient physical information within the dataset.

369 The test subset consists of samples from simulations with $W^+ \in \{6, 12, 18, 24\}$,
370 whereas the training subset encompasses the remaining samples, establishing an
371 approximate 80/20 distribution. This partitioning strategy enables assessment of the
372 model’s interpolation capabilities across intermediate oscillation amplitudes, thus eval-
373 uating its robustness to variations in the control parameter. The dataset structure
374 follows a four-dimensional indexing system where $\chi_{j,k,l}^{(i)}$ represents individual data

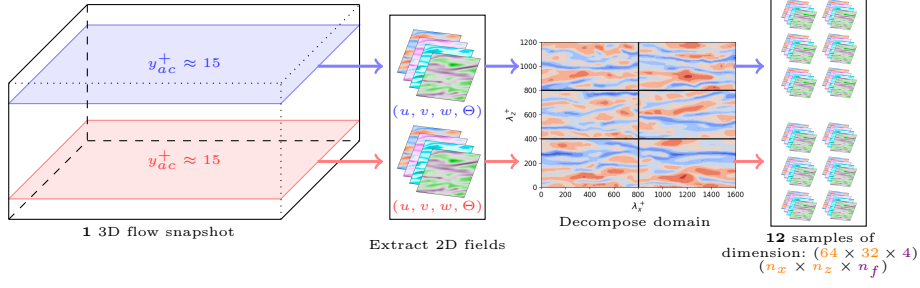


Fig. 3 Processing of 3D snapshots into 2D fields. 2D fields extraction: Slices are taken at $y_{ac}^+ \approx 15$ and its upper channel symmetric. Domain decomposition: the x-axis is divided into two segments, and the z-axis is split into three parts. λ_x^+ and λ_z^+ here represent the length (in wall-units) of the domain in the streamwise and spanwise directions respectively.

375 points. The indices correspond to: field type $i = 1, \dots, 4$ for u, v, w, Θ , simulation
 376 $j \in \{0, \dots, 20\}$, snapshot number $k = 1, \dots, n_{\text{samples}}$ (taking into account both time
 377 indexes and sub-domains), and spatial position $l = 1, \dots, 2048$ within each snapshot.
 378 The preprocessing stage addresses the inherent scale disparities among the flow fields,
 379 as illustrated in Fig. 4(a). The distributions of u and Θ exhibit positive shifts relative
 380 to the zero-centred distributions of v and w . The v component displays a particularly
 381 concentrated distribution with minimal variance, necessitating careful consideration
 382 during the reconstruction process to maintain appropriate weighting among variables.




Fig. 4 Dataset Scaling: (a) Distributions Prior to Scaling, (b) Distributions Following Scaling.
 — : u distribution, - - - : v distribution, - . - . : w distribution, : Θ distribution.

383 The data preprocessing implements a two-stage normalisation procedure. The first
 384 one applies field-specific standardisation:

$$\chi_{\text{std},j,k,l}^{(i)} = \frac{\chi_{j,k,l}^{(i)} - \mu_i}{\sigma_i}, \quad (6)$$

385 where $\mu_i = \langle \chi_{j,k,l}^{(i)} \rangle_{j,k,l}$ denotes the mean and $\sigma_i = \sqrt{\langle (\chi_{j,k,l}^{(i)} - \mu_i)^2 \rangle_{j,k,l}}$ represents
386 the standard deviation for each field i , computed across all simulations from the training
387 subset j , snapshots k and grid points l . The second stage implements a global
388 MinMax normalisation:

$$\chi_{\text{norm},j,k,l}^{(i)} = \frac{\chi_{\text{std},j,k,l}^{(i)} - \min_{i,j,k,l}(\chi_{\text{std},j,k,l}^{(i)})}{\max_{i,j,k,l}(\chi_{\text{std},j,k,l}^{(i)}) - \min_{i,j,k,l}(\chi_{\text{std},j,k,l}^{(i)})}. \quad (7)$$

389 The resultant distributions, illustrated in Fig. 4(b), exhibit centrality around 0.55
390 with harmonised variance and bounded values within $[0, 1]$. This normalisation strategy
391 enhances the performance of both POD and AE methods, particularly improving
392 the reconstruction accuracy of the wall-normal velocity component. The normalisation
393 parameters derived from the training dataset are applied consistently across all
394 model implementations. The inverse transformations restore the physical scaling when
395 computing performance metrics. This methodology ensures robust preprocessing while
396 preventing statistical contamination between training and test datasets.

397 3 Models

398 This section provides a comprehensive overview of the theoretical foundations and
399 practical implementation details for all the models employed in this study, including
400 POD and all AE variants.

401 3.1 Proper Orthogonal Decomposition

402 As outlined in Section 1, POD serves as a benchmark model for this study. POD is a
403 technique used to extract dominant patterns from high-dimensional data by identifying
404 a set of orthogonal modes ranked by their energy contribution [40]. The process
405 involves computing the eigenvalue decomposition of mean-centered data's covariance
406 matrix. Equivalently, the modes can be obtained by singular value decomposition of
407 the data matrix. The singular values of this decomposition represent the energy associated
408 with each mode, while the left singular vectors form spatial basis functions.
409 These modes are ordered by decreasing singular value magnitude, ensuring that the
410 most energetic structures are captured first.

411 The POD approximation of a function X depending on space x and time t is
412 expressed as:

$$X(x, t) \approx \sum_{i=1}^{L_d} \sigma_i \alpha_i(t) \rho_i(x), \quad (8)$$

413 where the right hand side represents the L_d -mode approximation of the function X .
414 Here, $\rho_i(x)$ are the spatial basis functions (orthonormal left singular vectors), σ_i are
415 the singular values and $\alpha_i(t)$ are the corresponding temporal coefficients, which also
416 form an orthonormal basis.

417 The orthogonality of both its spatial modes and temporal coefficients is a key
418 feature of POD, giving a clear interpretation of their individual contributions to the

419 overall reconstruction. This property is a specific instance of a broader goal in dimen-
 420 sionality reduction: achieving a statistically independent latent space. In this context,
 421 a model’s *latent samples orthogonality* refers to the orthogonality of the learned latent
 422 representations, which implies a statistically uncorrelated structure. A latent space
 423 where variables are uncorrelated is often more interpretable and efficient, as each
 424 dimension is encouraged to capture a distinct, independent factor of variation. For
 425 POD, this is guaranteed by construction through the orthogonality of the temporal
 426 coefficients $\alpha_i(t)$.

427 An other notable advantage of POD is the ranking of the modes with regards to
 428 their energy contributions using σ_i which yields additional interpretability. However,
 429 achieving convergence may require a large number of modes in scenarios involving
 430 turbulent flows or highly nonlinear governing equations. This is particularly relevant
 431 for the data considered in this study, where the fields u , v , w , and Θ exhibit complex
 432 spatial and temporal dynamics. To apply POD, each sample of shape $64 \times 32 \times 4$ is
 433 flattened into an 8192-dimensional vector, representing the spatial domain x in (8).
 434 To replicate the testing process used in the AE models, the extended POD approach
 435 [41] is employed in this study. This involves projecting the test data onto the POD
 436 spatial basis derived from the training data, *i.e.*

$$V_{\text{test}}^{\text{T}} = \Sigma_{\text{train}}^{-1} U_{\text{train}}^{\text{T}} X_{\text{test}}, \quad (9)$$

437 where V represents the right singular vectors (temporal coefficients), U the left singular
 438 vectors (spatial basis), Σ the matrix of singular values and X is the input data matrix.
 439 This process compromises the orthogonality of the test data temporal coefficients
 440 which no longer form a strict orthonormal basis, as attested in Section 5.

441 3.2 Convolutional AutoEncoder (CAE)

442 This section provides a concise overview of AutoEncoders (AEs) and Convolutional
 443 AutoEncoders (CAEs), introducing key concepts of convolutional layers, such as ker-
 444 nels and strides. The neural network architecture employed in this study is then
 445 detailed, and serves as a foundational framework for all autoencoder models discussed
 446 in this paper. The section also examines additional model hyperparameters, including
 447 activation functions, learning rate and batch size. An autoencoder is a neural network
 448 designed to compress high-dimensional data into a lower-dimensional representation
 449 (encoding) and reconstruct the original input from this compressed representation
 450 (decoding), see Fig. 5. The compressed representation, referred as the latent variable
 451 and denoted by the symbol \mathbf{z} , is the coordinate vector of the input data in the low-
 452 dimensional space. This representation is essential for the reconstruction of the input
 453 during the decoding process. During training, the encoding and decoding components
 454 are optimized collaboratively using backpropagation algorithms to identify the best
 455 low-dimensional representation. This is achieved by minimizing via gradient descent
 456 algorithms the difference between the input and its reconstruction, typically quantified
 457 using the Mean Squared Error (MSE) loss function.




Fig. 5 Autoencoder representation for the reconstruction of the (u, v, w, Θ) 2D fields. z is the latent representation of (u, v, w, Θ) , and $(\tilde{u}, \tilde{v}, \tilde{w}, \tilde{\Theta})$, its reconstruction through the autoencoder.

458 In this paper, the training process minimizes the MSE loss function. Adopting the
 459 same notation as in Section 2.4, where $\widetilde{\chi}_{j,k,l}^{(i)}$ represents the autoencoder reconstruction
 460 of grid point $\chi_{j,k,l}^{(i)}$, the MSE loss for a dataset χ is expressed as:

$$\text{MSE}(\chi) = \sum_{i=1}^4 \langle (\chi_{j,k,l}^{(i)} - \widetilde{\chi}_{j,k,l}^{(i)})^2 \rangle_{j,k,l}. \quad (10)$$

461 A convolutional autoencoder builds upon the traditional autoencoder architecture
 462 by using convolutional layers for encoding and transposed convolutional layers for
 463 decoding. In this study, 2D convolutions are applied, treating each input field as an
 464 individual channel. In the encoding phase, a 2D convolutional layer applies a small
 465 matrix called a kernel to a 2D input field using a sliding window approach. The
 466 kernel moves across the input in discrete steps, performing element-wise multiplication
 467 with the input region it overlaps, and then summing the results to produce a single
 468 value in the output. This operation detects local features such as edges, textures, or
 469 patterns in the input data. The kernel typically slides with a specified stride (the step
 470 size between applications) and may involve padding (extra values added around the
 471 input) to preserve the spatial dimensions of the data (see Fig. 6). The decoder uses
 472 transposed convolutional layers, which reverse the convolution operation. This process
 473 is often referred as up-sampling and allows the model to reconstruct the spatial details
 474 of the original input while maintaining the learned features. By using convolutional
 475 layers in both encoding and decoding, CAEs are particularly effective for processing
 476 spatial data, such as images or flow fields, e.g. velocity and temperature distributions
 477 [22, 42], as they capture hierarchical spatial patterns while efficiently reducing the
 478 dimensionality of the input data.

479 Using zero padding with a stride of 1 and a 3×3 kernel ensures that the center
 480 of the kernel passes over every element of the input field, resulting in an output field
 481 with the same dimensions as the input. This strategy is applied in layers 1, 2, and 3
 482 of the present encoder architecture (see Fig. 7), where the field size remains fixed at
 483 64×32 , while the number of channels (or filters) increases from 4 in the input to 24
 484 in the third layer. The stride of 1 in the first three convolutional layers is designed
 485 to capture the fundamental patterns of the flow, without reducing its dimensionality
 486 at this stage. In contrast, setting the stride to 2 causes the kernel to skip every other
 487 row and column, reducing the spatial dimensions of the output to one-fourth the size
 488 of the input. Fig. 6 illustrates this effect, where white blocks are skipped due to the




Fig. 6 Example of a convolutional layer with a 3×3 kernel (red matrix), with zero padding (yellow blocks) and stride = 2 on a 4×4 input matrix. Orange blocks on the output correspond to locations where the convolutions are performed on the input.

489 stride of 2, leading to a 2×2 output field. This downsampling strategy is used in
 490 layers 4, 5, and 6 of the current model to progressively reduce the spatial dimensions,
 491 while capturing the more complex features of the flow. After layer 6, the output tensor
 492 is flattened into a 4096-dimensional vector and mapped to a 500-dimensional latent
 493 vector via a fully connected layer (see Fig. 7). This aggressive dimensionality reduction
 494 accelerates convergence while preserving accuracy, outperforming architectures with
 495 intermediate dense layers.




Fig. 7 Architecture of the encoder component of the CAE. Orange blocks represent convolutional layers, while solid arrows indicate batch normalization. Dotted lines represent the flattening layer, and solid lines denote the fully connected (dense) layers.

496 The decoder architecture implements a hierarchical expansion of the compressed
 497 latent representation through a sequence of precisely configured layers. The recon-
 498 struction process begins with a fully-connected layer that transforms the latent vector
 499 into a 4096-dimensional space, which is subsequently reshaped into a $(8, 4, 128)$
 500 tensor format. This initial transformation is followed by a series of transposed con-
 501 volutional operations that progressively restore the spatial dimensions of the flow field
 502 through controlled upsampling procedures. The decoder structure exhibits partial
 503 symmetry with the encoder (Fig. 5), but introduces a modification in the filter dis-
 504 tribution. The implemented architecture employs an asymmetric filter arrangement
 505 of $(128, 128, 64, 32, 12, 4)$ across its layers, diverging from the traditional symmetric
 506 configuration of $(128, 64, 32, 24, 12, 8, 4)$. This architectural adaptation has demon-
 507 strated slightly superior reconstruction capabilities, particularly for the wall-normal

508 velocity component. This architecture creates an information bottleneck at the latent
 509 space, compelling the encoder-decoder pair to operate synergistically in extracting and
 510 preserving the most salient flow features necessary for accurate reconstruction.

511 The neural network architecture incorporates exponential linear unit activa-
 512 tion functions [43] throughout its structure, facilitating non-linear transformations
 513 within the normalised $[0, 1]$ range established in Section 2.4. The training protocol
 514 encompasses:

- 515 • 200 epochs with 64 samples per batch
- 516 • 3×3 convolution kernels
- 517 • Rectified adaptive moment estimation optimisation algorithm [see 44]
- 518 • Initial learning rate: 10^{-3}
- 519 • Adaptive rate reduction: 50% decrease after 10 epochs without improvement
- 520 • Minimum learning rate threshold: 10^{-6}

521 The latent dimension (L_d) is established at 500 components, ensuring a minimum
 522 80% energy reconstruction threshold (as given by (16)) across all flow fields. This
 523 configuration emerges from systematic parameter exploration, demonstrating optimal
 524 performance metrics in the reconstruction task.

525 3.3 β -Variational AutoEncoder (VAE)

526 Although conventional autoencoders demonstrate robust capabilities in reconstruct-
 527 ing complex fluid dynamics, their latent space representations often exhibit significant
 528 intermodal correlations, which impede the formation of interpretable basis. The tradi-
 529 tional POD method maintains strict orthogonality, but its inherent linear framework
 530 limits its compression efficiency. The β -VAE architecture represents a sophisticated
 531 solution to the challenge of balancing these competing objectives. The fundamental
 532 distinction of VAEs lies in their probabilistic approach to latent space encoding. Rather
 533 than implementing deterministic mappings, VAEs encode input data into probabil-
 534 ity distributions from which latent variables (LVs) are sampled, as depicted in Fig. 8.
 535 This probabilistic framework introduces a regulated stochastic element to the encod-
 536 ing process. The present VAE implementation maintains architectural consonance
 537 with the previously described convolutional autoencoder (Fig. 7), preserving identical
 538 hyperparameter configurations. The distinguishing elements manifest in three critical
 539 aspects: (i) the probabilistic nature of the latent distribution, (ii) the stochastic sam-
 540 pling mechanism employed for latent variable generation, and (iii) the modified loss
 541 function formulation. This adaptation enables the VAE to maintain reconstruction
 542 fidelity whilst encouraging the emergence of more interpretable latent representations.

543 The encoder models the approximate posterior distribution $q_\psi(\mathbf{z}|\mathbf{x})$, where \mathbf{x} is
 544 the input, as a multivariate Gaussian $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, where $\boldsymbol{\Sigma} = \text{diag}(\boldsymbol{\sigma}^2)$. For each sample
 545 in the dataset, the encoder produces a mean vector $\boldsymbol{\mu}$ and a standard deviation vector
 546 $\boldsymbol{\sigma}$, representing the diagonal elements of the covariance matrix $\boldsymbol{\Sigma}$.

547 Directly sampling \mathbf{z} from a Gaussian distribution would render both the latent
 548 variables and the encoder weights non-differentiable, preventing the use of backpropa-
 549 gation. To address this, the reparameterization trick is employed. Instead of sampling \mathbf{z}

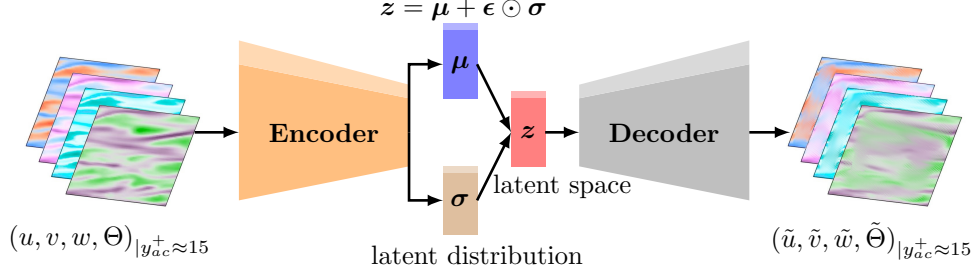


Fig. 8 VAE representation. ϵ is sampled from a standard Gaussian distribution of zero mean and identity matrix for covariance. \mathbf{z} is obtained via a linear combination of the predicted mean $\boldsymbol{\mu}$, and an element-wise product of the sampled ϵ and the predicted standard deviation $\boldsymbol{\sigma}$.

550 directly, the encoder outputs the mean $\boldsymbol{\mu}$ and standard deviation $\boldsymbol{\sigma}$, and \mathbf{z} is computed
 551 as:

$$\mathbf{z} = \boldsymbol{\mu} + \epsilon \odot \boldsymbol{\sigma}, \quad (11)$$

552 where ϵ is sampled from a standard normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$. This approach shifts
 553 the randomness to ϵ , making the loss function differentiable with respect to $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$,
 554 thus allowing effective backpropagation to optimize these parameters.

555 VAEs are based on the assumption that the input \mathbf{x} is a sample from an
 556 unknown underlying distribution, denoted as $p_{\theta}(\mathbf{x})$, which represents the probability
 557 of generating real samples \mathbf{x} . Introducing conditional probability, $p_{\theta}(\mathbf{x})$ is given by

$$p_{\theta}(\mathbf{x}) = \int p_{\theta}(\mathbf{x}|\mathbf{z})p_{\theta}(\mathbf{z}) d\mathbf{z}, \quad (12)$$

558 where the prior distribution $p_{\theta}(\mathbf{z})$ represents the assumption of the model concerning
 559 the distribution of the latent space, while $p_{\theta}(\mathbf{x}|\mathbf{z})$ describes the probabilistic decoder
 560 distribution that generates the data \mathbf{x} given the latent variable \mathbf{z} . The objective is to
 561 maximize $p_{\theta}(\mathbf{x})$ or equivalently its log-marginal likelihood $\log(p_{\theta}(\mathbf{x}))$.

562 Using Bayes' theorem and the logarithmic properties, the following inequality can
 563 be derived [45]:

$$\log(p_{\theta}(\mathbf{x})) \geq -D_{\text{KL}}[q_{\psi}(\mathbf{z}|\mathbf{x})||p_{\theta}(\mathbf{z})] + \mathbb{E}_{q_{\psi}(\mathbf{z}|\mathbf{x})}[\log(p_{\theta}(\mathbf{x}|\mathbf{z}))], \quad (13)$$

564 where D_{KL} denotes the KL divergence, a measure of dissimilarity between two probab-
 565 ility distributions. This equation, known as the Evidence Lower Bound (ELBO),
 566 provides a computable lower bound on the log-marginal likelihood of the data. By
 567 maximizing the ELBO (the right hand side), the log-marginal likelihood of generat-
 568 ing real samples, is indirectly maximized. To achieve this, one must minimize the
 569 KL divergence between the approximate posterior $q_{\psi}(\mathbf{z}|\mathbf{x})$ and the prior distribution
 570 $p_{\theta}(\mathbf{z})$, while simultaneously maximizing the expected log-likelihood of generating \mathbf{x}
 571 given \mathbf{z} . The second term in the ELBO is equivalent to minimizing the reconstruction
 572 error (MSE) used in the CAE loss function (10).

573 The prior distribution $p_{\theta}(\mathbf{z})$, which regularizes the latent space and is used in the
574 KL divergence term, is here set to a standard Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$. Not only
575 is this prior commonly used due to its simplicity and mathematical tractability, but it
576 also enforces the latent space to be dense and orthogonal, thereby reducing correlations
577 between the latent variables. Indeed, minimizing the KL divergence between $p_{\theta}(\mathbf{z})$ and
578 $q_{\psi}(\mathbf{z}|\mathbf{x})$ indirectly minimizes the KL divergence between $p_{\theta}(\mathbf{z})$ and the marginal latent
579 distribution $q_{\psi}(\mathbf{z})$ (see Hoffman and Johnson 46). Strong regularization thus ensures
580 the latent distribution remains close to the isotropic standard Gaussian, encouraging
581 each dimension to encode more independent factors of variation. This reduced corre-
582 lation between latent variables thus enhances model interpretability. A dense latent
583 space also facilitates smoother interpolations, which is particularly valuable for sample
584 generation tasks. The combination of these characteristics makes the standard Gaus-
585 sian prior especially suitable for VAEs in both representation learning and generative
586 modeling.

587 Given the chosen distributions, the KL divergence term in the ELBO simplifies to:

$$D_{\text{KL}}[q_{\psi}(\mathbf{z}|\mathbf{x})||p_{\theta}(\mathbf{z})] = -\frac{1}{2} \left\langle \sum_{i=1}^{L_d} (\log(\sigma_{j,k}(i)^2) + 1 - \mu_{j,k}(i)^2 - \sigma_{j,k}(i)^2) \right\rangle_{j,k}, \quad (14)$$

588 where $\mu_{j,k}(i)$ denotes the mean of the i -th coordinate of the latent distribution, for
589 simulation j and snapshot k . The same notations apply to the standard deviation
590 $\sigma_{j,k}(i)$.

591 The β -VAE introduces a real number $\beta \in \mathbb{R}^+$, which acts as a parameter to balance
592 the trade-off between orthogonality of the latent samples and reconstruction quality.
593 Given the MSE loss defined in (10), the following loss function is obtained for β -VAEs:

$$\text{Loss}_{\text{VAE}}(\chi, \boldsymbol{\mu}, \boldsymbol{\sigma}, \beta) = \text{MSE}(\chi) + \beta \times D_{\text{KL}}[q_{\psi}(\mathbf{z}|\mathbf{x})||p_{\theta}(\mathbf{z})]. \quad (15)$$

594 This formulation allows for optimization of both the reconstruction accuracy and
595 latent samples orthogonality.

596 Beyond their improved orthogonality and interpolation properties, β -VAEs often
597 exhibit an automatic mode pruning effect [25]. The regularization term in (15)
598 forces some of the latent distributions to collapse toward the standard Gaussian
599 prior, effectively over-regularizing these variables and minimizing their contribution
600 to reconstruction. Building on this observation, Kang et al. [25] developed a two-stage
601 approach: (i) a ranking algorithm that assesses latent variable contribution through
602 KL divergence analysis, and (ii) a pruning procedure that eliminates non-informative
603 LVs while preserving reconstruction quality. This combined framework applied to β -
604 VAEs can provide a method for determining optimal latent dimensionality, which is
605 particularly valuable for dimension reduction tasks.

606 3.4 Metrics

607 To compare the performance of various models with different loss functions, it is
608 necessary to define metrics that allow meaningful evaluations. These metrics can be

609 separated in three categories: (i) reconstruction of the four fields, (ii) latent samples
610 orthogonality, and (iii) latent samples separability with regards to flow parameters.

611 For the first category, the metric of interest is the per-field energy reconstruction
612 [23, 24], which measures the energy percentage captured by the model. It is particularly
613 relevant to evaluate the reconstruction of each variable separately, as they are not
614 all reconstructed with the same accuracy, a point that will be discussed in detail in
615 Section 4.

616 Mathematically, the metric is defined as the squared difference between the fluctu-
617 ations of input and reconstructed fields, normalized by the squared norm of the input
618 fluctuations. This metric for a field $\chi^{(i)}$, using the same notation as in Section 2.4, is
619 expressed as:

$$E_r(\chi^{(i)}) = 100 \times \left(1 - \left\langle \frac{\sum_{l=1}^{2048} [\chi_{j,k,l}^{(i)'} - \widetilde{\chi_{j,k,l}^{(i)}}']^2}{\sum_{l=1}^{2048} [\chi_{j,k,l}^{(i)'}]^2} \right\rangle_{j,k} \right), \quad (16)$$

620 where $\chi_{j,k,l}^{(i)'} = \chi_{j,k,l}^{(i)} - \langle \chi_{j,k,l}^{(i)} \rangle_{k,l}$ and $\widetilde{\chi_{j,k,l}^{(i)}}$ represents the reconstruction of the grid
621 point $\chi_{j,k,l}^{(i)}$ obtained by the model. The means utilized to calculate the fluctuations
622 are therefore specific to each field i and simulation j .

623 To evaluate the orthogonality of the latent samples, a correlation matrix M is
624 constructed. The correlation coefficient between two components m and n of the latent
625 vector \mathbf{z} is defined as:

$$M_{m,n} = \frac{C_{mn}}{\sqrt{C_{mm}C_{nn}}} \quad \text{for } 1 \leq n, m \leq L_d. \quad (17)$$

626 C_{mn} , the covariance between modes m and n , is given by:

$$C_{mn} = \langle z'_{j,k}(m)z'_{j,k}(n) \rangle_{j,k}, \quad (18)$$

627 where $z_{j,k}(m)$ denotes the m -th latent coordinate of simulation j and snapshot k , and
628 $z'_{j,k} = z_{j,k} - \langle z_{j,k} \rangle_{j,k}$ represents the corresponding fluctuation, obtained by averaging
629 over all dataset latent samples.

630 Technically, the correlation matrix contains all the information about the orthog-
631 onality (or correlation) of each mode with the others. However, extracting a scalar
632 metric from this matrix to represent orthogonality is often more practical, especially
633 for comparing different models. A common approach is to calculate the determinant
634 of the correlation matrix [23, 24], which equals 1 in the case of perfect orthogonality,
635 corresponding to an identity matrix. However, with high-dimensional latent spaces,
636 this metric becomes less reliable, as even modes that are nearly orthogonal can result
637 in a determinant close to zero. This limitation is evident in the current study, affecting
638 all AE models as well as the extended POD method [41].

639 To address this, two alternative metrics are introduced: the mean correlation
640 (Mean_c) and the mean-max correlation (MeanMax_c). These quantities are defined as
641 follows.

642 1. Mean Correlation (Mean_c): The mean correlation measures the average off-diagonal
643 coefficients of the correlation matrix in absolute value, providing an overall indica-
644 tion of how correlated the latent modes are with each other. Based on the definition
645 of the covariance coefficients (17) taken in absolute value, it is reasonable to express
646 this metric as a percentage:

$$\text{Mean}_c = 100 \times \langle |M_{i,j}| \rangle_{1 \leq i, j \leq L_d, i \neq j}. \quad (19)$$

647 While Mean_c provides a good overview of the orthogonality of the latent sam-
648 ples, it can sometimes dilute the impact of a few highly correlated modes due to
649 averaging.

650 2. Mean-Max Correlation (MeanMax_c): To complement the Mean_c metric, the mean-
651 max correlation is introduced. This metric calculates the average of the maximum
652 correlation coefficient (in absolute value) for each mode, helping to assess whether
653 modes are uncorrelated with all other modes in the system.
654

$$\text{MeanMax}_c = 100 \times \left\langle \max_{1 \leq j \leq L_d, j \neq i} (|M_{i,j}|) \right\rangle_{1 \leq i \leq L_d}. \quad (20)$$

655 In a well-structured and orthogonal latent space, both the Mean_c and MeanMax_c
656 metrics should exhibit consistency, with both being close to 0. Together, these two
657 metrics provide a comprehensive view of the degree of correlation in the latent space,
658 ensuring that no significant dependencies among the latent modes are overlooked.

659 Finally, there remains to evaluate the separability of the latent samples with
660 regards to flow parameters. This refers to the global structure of the latent set, namely,
661 how well it can be partitioned into clusters that correspond to distinct generative fac-
662 tors. In simple flows, for which the data can be compressed into a space of dimension
663 2-3, visualizing the latent space allows for straightforward evaluation of this separa-
664 bility. However, with the present turbulent flow dataset requiring much larger latent
665 dimensions, direct visualization is impossible, necessitating the use of quantitative
666 metrics to assess this aspect effectively.

667 The chosen approach consists of selecting all samples corresponding to a specific
668 amplitude (or phase) within the test dataset, and to create a cluster with these sam-
669 ples, defined by its centroid and contained data samples. From these clusters, one can
670 easily apply any clustering metric, taking into account intra or inter cluster vari-
671 ances. A relevant metric that takes into account both of these aspects is the Davies-Bouldin
672 indicator [47], which is calculated as the maximal ratio of the averaged distance
673 between a point and its centroid and the distance between two centroids, averaged
674 over all clusters, *i.e.*

$$D_B = \left\langle \max_{m \neq n} \left(\frac{\delta_n + \delta_m}{\text{dist}(\mathbf{c}_n, \mathbf{c}_m)} \right) \right\rangle_n. \quad (21)$$

675 Here, δ_n is the average distance of all the points contained in \mathcal{C}_n , cluster of
676 index n , to the centroid \mathbf{c}_n . This cluster intra-variance is estimated as $\delta_n =$

677 $\frac{1}{|\mathcal{C}_n|} \sum_{\mathbf{z}_{j,k} \in \mathcal{C}_n} \text{dist}(\mathbf{z}_{j,k}, \mathbf{c}_n)$ where $\mathbf{z}_{j,k}$ corresponds to a latent sample and $\text{dist}(\mathbf{x}, \mathbf{y})$
 678 measures the euclidean distance between vectors \mathbf{x} and \mathbf{y} .

679 The perfect Davies-Bouldin score is 0 which means clusters are well separated,
 680 clearly defined and dense. Higher scores mean clusters are poorly defined and overlap.

681 When evaluating the separability of only two clusters, the Davies-Bouldin indicator
 682 reduces to:

$$D_B(\mathcal{C}_n, \mathcal{C}_m) = \frac{\delta_n + \delta_m}{\text{dist}(\mathbf{c}_n, \mathbf{c}_m)}. \quad (22)$$

683 Using this metric, one can estimate a matrix D such that $D_{n,m} = D_B(\mathcal{C}_n, \mathcal{C}_m)$,
 684 describing how each cluster is separated with the others.

685 4 Comparing CAE with POD

686 This section conducts a comparative evaluation of CAE and POD methodolo-
 687 gies, focusing on three core objectives: assessing their reconstruction performance,
 688 examining their interpolation capabilities at unseen oscillation amplitudes W^+ , and
 689 investigating the advantages of compressing all fields simultaneously versus processing
 690 them individually. Unless otherwise noted, all results are derived from test dataset.

Table 5 Comparison of energy reconstruction for each field, where the AE is trained either to reconstruct all fields simultaneously with $L_d = 500$ (Together) or to reconstruct them individually using 125 modes for each (Separately).

	Together ($L_d=500$)	Separately ($L_d=125$)
$E_r(u)$	92%	84%
$E_r(v)$	81%	65%
$E_r(w)$	92%	90%
$E_r(\Theta)$	90%	78%

691 Table 5 highlights the superior reconstruction accuracy achieved by the CAE when
 692 trained on all fields simultaneously within a 500-dimensional latent space, compared to
 693 independent field training using 125-dimensional latent spaces per field. This compar-
 694 ison is dimensionally consistent, as the cumulative latent dimensionality for individual
 695 field processing also totals 500. The simultaneous compression approach improves
 696 reconstruction accuracy for the u , v , and Θ fields by 8%, 16% and 12% respectively.
 697 This enhancement leverages the strong physical correlations between fields, particu-
 698 larly between the streamwise velocity (u) and temperature (Θ) fields, as expected
 699 from the Reynolds analogy. Notably, the wall-normal velocity component (v) consis-
 700 tently exhibits lower reconstruction accuracy across both approaches, a limitation that
 701 warrants further investigation. Considering these results, dimensionality reduction
 702 models are systematically applied to the four fields together instead of separately for
 703 the rest of this study.

704 Fig. 9 compares the reconstructed-energy percentage of POD and CAE across var-
 705 ious latent dimensions, for each field. The results show that the CAE consistently




Fig. 9 Energy reconstruction scores (E_r) as a function of the latent dimension (L_d) for extended POD (blue) and CAE (green). Solid lines indicate the training data results, while dotted lines highlight scores on the test data. The fields correspond to the following: (a) Streamwise velocity u , (b) Wall-normal velocity v , (c) Spanwise velocity w , and (d) Temperature Θ .

706 outperforms extended POD for all L_d -values. Extended POD only matches the per-
707 formance of CAE on v when $L_d \geq 1000$, with both methods performing similarly for
708 larger dimensions. For each field, the performance gap narrows as L_d increases, with
709 the greatest difference observed at the lowest value of L_d analysed, *i.e.* 250. Notably,
710 this gap is most pronounced for v , where a difference of over 50% in E_r is observed
711 at $L_d = 250$, highlighting the advantage of the non-linear CAE approach, particularly
712 for smaller latent dimensions. Another major result is the number of modes for which
713 each model reaches the 90% reconstruction accuracy benchmark. Excluding the v -
714 field, where both models reach this target around $L_d = 1200$, extended POD requires
715 2 to 3 times more modes than the CAE to achieve 90% accuracy. This underscores the
716 CAE's superior ability to compress information, as it reaches comparable or better
717 accuracy with significantly fewer modes, making it better suited for this dataset.

718 As noted earlier, the v -field consistently achieves lower E_r scores than the other
719 three fields. Fig. 9 shows that for $L_d = 250$, $E_r(u)$, $E_r(w)$, and $E_r(\Theta)$ all exceed 80%
720 for the CAE, while $E_r(v)$ remains around 65%. As mentioned previously, this issue

721 is even more pronounced for POD, further highlighting the challenges specific to the
722 v -field. This is at least partially attributed to the smaller turbulent scales contained
723 in the v -field compared to the others, which require more modes to be captured. The
724 v -component is also initially less energy-dense compared to the other fields, and the
725 applied scaling may not have fully remedied this discrepancy. Additionally, the long
726 tail observed in the scaled v -distribution (Fig. 4(b)) does not contribute significantly to
727 this disparity, as a pre-scaling truncation of this tail did not improve the reconstruction
728 accuracy. Another possible explanation is that SWO enhance sweeps and ejections in
729 the wall-normal direction [10], further expanding the v -distribution in the actuated
730 case.

Table 6 Energy reconstruction scores ($E_r(\chi)$ for $\chi = [u, v, w, \Theta]$), mean-correlation (Mean_c) and mean max-correlation (MeanMax_c) on the test dataset, with $L_d = 500$, on all models. Comparison of results for extended POD, CAE and β -VAE for $\beta = 5 \times 10^{-7}$ (see Section 5).

	CAE	β -VAE	POD
$E_r(\chi)$	[92,81,92,90]	[89,74,89,85]	[85,65,84,81]
Mean_c	2.1	1.1	0.7
MeanMax_c	9.7	6.9	4.1

731 The selection of latent dimension $L_d = 500$ as the principal case investigated is
732 established to guarantee that the CAE maintains a minimum reconstruction accuracy
733 of 80% across all four target fields when assessed against the testing dataset. The
734 comparative analysis at $L_d = 500$ reveals the superior performance of the CAE relative
735 to extended POD when utilizing an identical number of modes. Specifically, the CAE
736 demonstrates enhanced reconstruction efficiency with improvements of 7%, 16%, 8%,
737 and 9% in E_r for the u , v , w , and Θ fields, respectively (see Table 6). These quantitative
738 results provide substantial evidence for the enhanced capabilities of the non-linear
739 autoencoder methodology compared to its linear counterpart.

740 Both models demonstrate robust generalization capabilities, as evidenced by the
741 comparable behaviour of training and testing curves across all flow fields in Fig. 9.
742 This characteristic was specifically evaluated through the train-test partition detailed
743 in Section 2.4. The observed higher test scores compared to training scores for the
744 w field are likely attributable to the disparity in dataset sizes, with the training set
745 containing simulations at 17 distinct W^+ values, while the test set comprises only
746 4 W^+ values. This imbalance in sampling density may lead to better generalization
747 performance on the sparser test dataset.

748 Fig. 10 offers enhanced insight into the interpolation performance of both CAE
749 and POD (with $L_d = 500$) by disaggregating the reconstruction errors across the four
750 distinct oscillation amplitudes ($W^+ \in \{6, 12, 18, 24\}$) in the test dataset. Notably,
751 reconstruction accuracy varies significantly with amplitude for both methods, with
752 performance gaps in E_r reaching $\approx 7.5\%$ for the CAE compared to $\approx 12\%$ for POD,
753 across all fields except w , where the observed gap is slightly lower. This smaller




Fig. 10 Energy reconstruction for each W^+ case in the test dataset, with $L_d = 500$. (a): CAE results, (b): extended POD results.

754 performance variation demonstrates the CAE’s superior robustness to changes in actu-
755 ation amplitude. A clear difference is observed between the reconstructed fields: while
756 reconstruction of u , v , and Θ progressively deteriorates with increasing W^+ , w recon-
757 struction conversely improves. This dichotomy aligns with physical expectations, as
758 higher oscillation amplitudes transform near-wall w fluctuations into smoother peri-
759 odic motions, whereas lower amplitudes retain more stochastic behavior characteris-
760 tic of the unactuated case. The degradation in reconstructing other fields with increas-
761 ing W^+ suggests that despite the more periodic nature, larger amplitudes introduce
762 additional flow complexity, consistent with the rising Re_τ observed by Gu erin et al.
763 [10] at $W^+ = 30$.

764 The spatial reconstructions depicted in Fig. 11 present comparative visualisations
765 across the complete LES domains, incorporating sliding averages to maintain subdo-
766 main continuity (see Fig. 3). The comparative analysis encompasses Ground Truth
767 representations (left column), CAE reconstructions (centre column), and extended
768 POD reconstructions (right column) for each flow field. These snapshots are picked
769 from the simulation of the test dataset corresponding to the highest oscillation
770 amplitude: $W^+ = 24$. This simulation is characterized by its important small-scale
771 fluctuations and by the largest gap in performances when comparing CAE and POD
772 (see Fig. 10). The Ground Truth data exhibits pronounced, spatially intermittent fluc-
773 tuations of significant amplitude, characteristic of turbulent small-scale structures.
774 These fine-scale turbulent manifestations present significant reconstruction challenges,
775 necessitating the implementation of 500 modes to achieve satisfactory accuracy.

776 The v -field exhibits particularly pronounced small-scale structural elements, pre-
777 senting substantial reconstruction challenges for both methodologies. Its distribution
778 demonstrates a notable centralised tendency (see Fig. 4), which manifests in reduced
779 reconstruction emphasis from both approaches. The extended POD reconstruction
780 demonstrates adequate representation of large-scale flow features. However, the recon-
781 structed flow structures exhibit considerable “smoothing” relative to the Ground
782 Truth, with numerous small-scale features absent from the reconstruction.

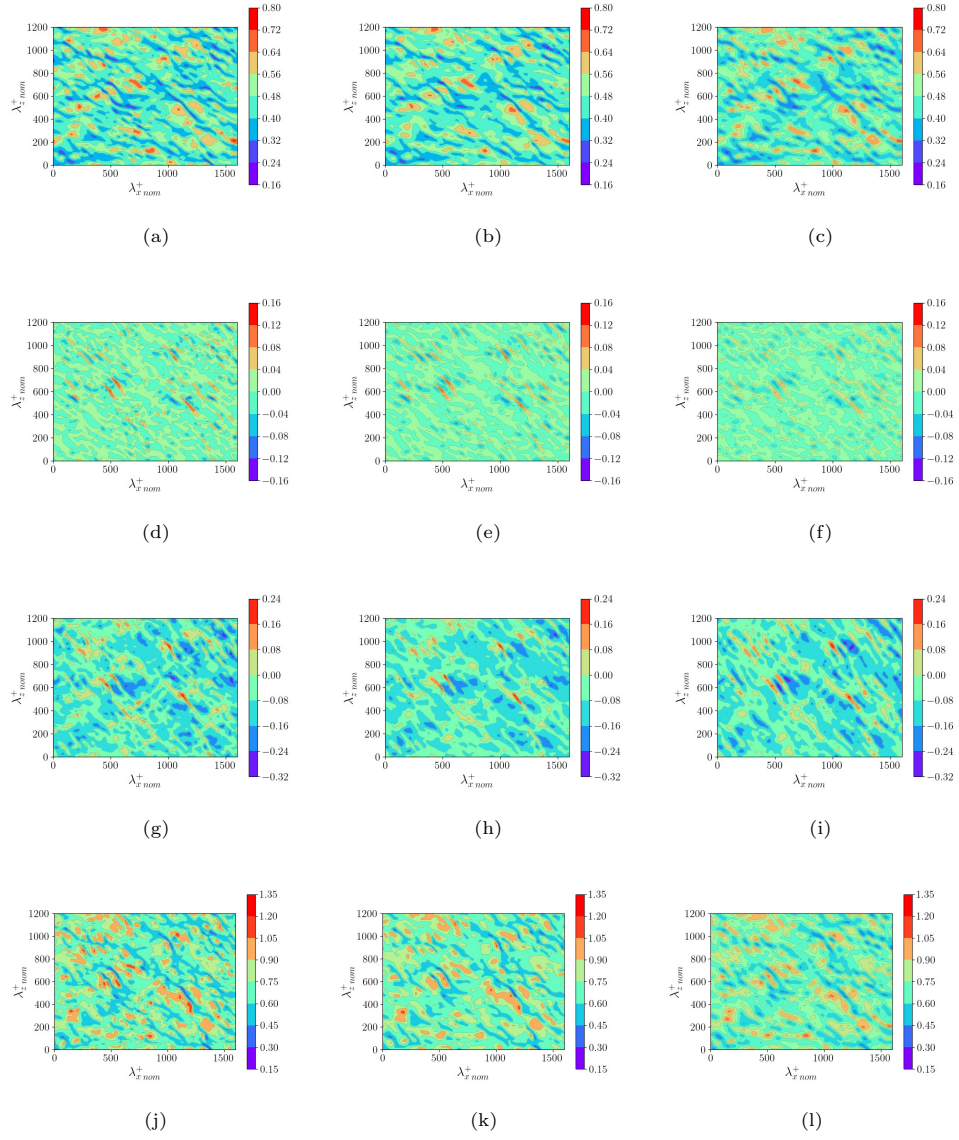


Fig. 11 Reconstruction of snapshots with $L_d = 500$ for the maximal amplitude simulation within the test dataset ($W^+ = 24$): (Left) Ground truth, (Middle) CAE reconstruction, and (Right) extended POD reconstruction. The four fields are displayed from top to bottom: u , v , w , and Θ . Each snapshot represents the full LES domain.

783 Visual evidence of the ongoing challenge faced by POD to adequately reconstruct
784 small-scale fluctuations in the flow is provided by the Power Spectral Densities (PSDs)
785 in the spanwise direction represented in Fig. 12 for all components. Pre-multiplied
786 PSDs are compared for the ground truth data (black solid lines) of the test dataset




Fig. 12 Pre-multiplied Power spectral density $k_z \Phi_{\chi^{(i)} \chi^{(i)}}(\lambda_z)$ of (a) u , (b) v , (c) w , (d) Θ , for the maximal amplitude simulation within the test dataset ($W^+ = 24$). The spectras are determined on the subdomains of Fig. 3 and not on the full LES-domain. Black solid lines account for the Ground Truth data, green dashed lines for the CAE estimation, blue dotted lines for the extended POD predictions, and purple dash-dotted lines for the β -VAE predictions.

787 simulation with oscillation amplitude $W^+ = 24$, CAE reconstruction (green dashed
788 lines) and extended POD reconstruction (blue dotted lines). For all fields, the blue
789 curve remains beneath the black curve for all flow scales, with the gap narrowing for
790 the larger scales of the flow, resulting in POD not capturing the total energy of the
791 system. The model particularly fails to capture the smaller scales of the flow, which
792 is further confirmed by the flat left tail of the blue spectras. As anticipated, this phe-
793 nomenon is even more pronounced for the wall-normal velocity v (see Fig. 12(b)),
794 which exhibits a significantly larger energy gap compared to the other fields. This
795 behavior is expected, as a substantial portion of the energy in v is concentrated in
796 its smaller-scale structures. This energy distribution likely explains the comparatively

797 lower reconstruction scores for the v -field, as accurately capturing these fine-scale fea-
798 tures presents a greater challenge. This aligns with the theoretical foundation of POD,
799 where truncating the modes inherently excludes the lower-energy modes associated
800 with these fine scales. Furthermore, the reconstructed snapshots reveal that extended
801 POD struggles to capture extreme values within flow structures, demonstrating a
802 tendency towards mean-value approximation. These constraints highlight the inher-
803 ent limitations of linear methods for representing turbulent flows, where small-scale
804 dynamics are crucial for an accurate reconstruction.

805 The CAE reconstruction methodology demonstrates enhanced capability in captur-
806 ing localised variations (see Fig. 11), preserving these features with increased fidelity.
807 The reconstruction quality is substantial across the majority of fields, with superior
808 representation of small-scale non-linear features compared to the POD approach. This
809 difference in model performances is confirmed by the spectral analysis of Fig. 12, in
810 which the CAE consistently reconstructs all the flow scales up to $\lambda_{z_{nom}}^+ \approx 200$ more
811 accurately than POD. Intriguingly, the POD performances are closer to those of the
812 CAE on the temperature field, as they provide similar estimations for wavelengths
813 $\lambda_{z_{nom}}^+ > 50$. The CAE demonstrates significantly lower performance in the v field in
814 comparison to the other components, primarily due to the inherent small-scale char-
815 acteristics of this field. Nevertheless, the wall-normal velocity is characterised by the
816 largest discrepancy between the POD and CAE spectra, further demonstrating the
817 superiority of the latter approach in reconstructing turbulent flows. These improve-
818 ments are attributed to the architectural implementation of convolutional layers (see
819 Section 3.2) in conjunction with non-linear activation functions, specifically designed
820 for the detection and preservation of localised, non-linear features in spatial data
821 structures. Despite the generally satisfactory reconstruction performance, regions char-
822 acterised by significant noise, continue to present reconstruction challenges for the
823 CAE methodology.

824 The CAE demonstrates superior performance in reconstruction relative to extended
825 POD. However, it exhibits suboptimal results in terms of latent samples orthogonal-
826 ity (see Table 6). Although extended POD is not strictly orthogonal (as explained in
827 Section 3.1), it achieves mean correlation scores (Mean_c) that are three times lower
828 than those of the CAE, and its modes are, at most, less than half as correlated as
829 the CAE latent variables (MeanMax_c). An ideal model should achieve the high recon-
830 struction accuracy of the CAE while producing correlation scores closer to those of
831 extended POD, allowing for more interpretable modes. This balance will be explored
832 further in the next section.

833 5 Balancing orthogonality of the latent samples and 834 reconstruction accuracy with β -VAE

835 In this section, the orthogonality of the latent samples is examined by sampling latent
836 variables in a manner that encourages independence and proximity to the standard
837 Gaussian distribution, as opposed to using a deterministic, unstructured approach (see
838 Section 3.3). More precisely, β -VAEs are employed to determine whether a suitable

839 compromise can be achieved between the orthogonality of extended POD and the
 840 reconstruction accuracy of the CAE.

841 A parametric study on β showed that the value 5×10^{-7} led to the best trade-
 842 off between reconstruction and orthogonality. This value was subsequently retained.
 843 Table 6 highlights the obtained trade-off. Even with a relatively small β value, the
 844 reconstruction accuracy across all fields experiences a slight decrease, underscoring
 845 the significant impact of the latent space structure on energy reconstruction and the
 846 difficulty of finding an appropriate balance. The orthogonality of the latent samples
 847 is largely enhanced comparatively to the CAE, with a decrease of 1% and 2.8% in
 848 Mean_c and MeanMax_c , respectively. Interestingly, this reduction of high-correlation
 849 relationships in the latent space impacts the reconstruction of v and Θ fields the
 850 most, as attested by Table 6. This suggests that the improvements in latent samples
 851 orthogonality may come at the expense of reconstruction accuracy for these fields,
 852 which is likely linked to the high correlation among all four fields.

853 Fig. 12 demonstrates that, like the CAE, the β -VAE (purple dash-dotted line)
 854 outperforms extended POD in reconstructing the flow, particularly for the smaller
 855 turbulent scales and at the energy peak. However, this is not the case for the tempera-
 856 ture field, where extended POD achieves better accuracy at the energy peak, aligning
 857 with the previously noted decline in β -VAE performance for this field. While the β -
 858 VAE performs marginally worse than the CAE in flow scale reconstruction overall, it
 859 matches the CAE's performance in the v field for this specific simulation.

860 The mode ranking algorithm developed by Kang et al. [25] was applied follow-
 861 ing the completion of model training. As detailed in Section 3.3, latent variables
 862 are sorted according to their Kullback-Leibler divergence from the one-dimensional
 863 standard Gaussian prior, wherein higher divergence values are indicative of greater
 864 information content. Through the regularisation term in (15), latent variables with
 865 minimal KL divergence are forced towards the standard Gaussian distribution. As a
 866 direct consequence, these over-regularised variables are rendered functionally redun-
 867 dant, contributing negligibly to the reconstruction process. Dimensional reduction
 868 through mode pruning is thus mathematically justified, whereby a systematic excision
 869 of redundant modes is performed based on their information content. In this pruning
 870 procedure, the most informative latent variables (classified as physics-aware in Kang
 871 et al. 25) are preserved within the reduced-order representation, while the less infor-
 872 mative components (physics-unaware) are systematically eliminated. This dimensional
 873 truncation is implemented by fixing the physics-unaware modes to their respective
 874 estimated mean values (23a), effectively neutralising their contribution to the flow
 875 field reconstruction.

876 The KL divergence is computed using (14) (applied individually to each latent
 877 variable) and the empirically estimated parameters $\hat{\mu}_n$ and $\hat{\sigma}_n$ determined as

$$\hat{\mu}_n = \langle z_{j,k}(n) \rangle_{j,k}, \quad (23a)$$

$$\hat{\sigma}_n = \sqrt{\langle (z'_{j,k}(n))^2 \rangle_{j,k}}. \quad (23b)$$

879 This procedure allows the mode ranking and pruning algorithm to be applied to the
 880 CAE, despite its deterministic nature.




Fig. 13 KL divergence of each latent variable with the 1D standard Gaussian prior, for the CAE (green crosses) and β -VAE (purple Ys), with $L_d = 500$ and ranking of the modes with regards to their KL divergence score.

881 Fig. 13 displays the KL divergence of each latent variable, ranked in descending
 882 order, with CAE results marked by green crosses and β -VAE results by purple Ys.
 883 The CAE modes exhibit a KL divergence roughly four times higher than those of
 884 the β -VAE in average, highlighting the efficacy of the regularization term in (15). A
 885 sharp decline in KL divergence beyond latent variable 424 of the β -VAE separates the
 886 physics-aware modes from the physics-unaware ones, suggesting the presence of the
 887 pruning effect commonly observed in VAEs. Here, 76 latent distributions are effectively
 888 reduced to a standard Gaussian, contributing negligibly to reconstruction. A similar,
 889 though less pronounced, KL drop occurs for the CAE at mode 480. However, even the
 890 least informative CAE latent variables retain a KL divergence of approximately 0.04,
 891 preventing it from being classified as physics-unaware.

892 Following the methodology of [25], a physics-informed β -VAE can be derived by
 893 retaining only the 424 modes with the highest KL divergence. Table 7 confirms that
 894 the pruned model achieves reconstruction performance identical to the full model,
 895 validating the pruning effect induced by the β -VAE regularization term.

Table 7 Energy reconstruction scores ($E_r(\chi)$ for $\chi = [u, v, w, \Theta]$), mean-correlation (Mean_c) and mean max-correlation (MeanMax_c) on the test dataset, for the CAE with $L_d = 400$, and the pruned β -VAE, using 424 modes.

	CAE ₄₀₀	β -VAE _{pruned}
$E_r(\chi)$	[90,75,91,87]	[89,74,89,85]
Mean_c	1.8	1.3
MeanMax_c	9.4	6.5

896 Before comparing the models in terms of latent samples orthogonality, it is impor-
897 tant to note that the mean correlation (Mean_c) should be interpreted with caution, as
898 it may underestimate the presence of highly correlated modes. Fig. 14(a) shows the
899 correlation matrix of the CAE, with LVs ranked according to their KL divergence.
900 The last 10 to 20 modes, corresponding to the KL divergence drop of Fig. 13, exhibit
901 nearly perfect correlation with each other and are significantly correlated with the
902 remaining modes compared to the average behavior. However, these LVs have a lim-
903 ited effect on the overall Mean_c score, which remains low (2.1%, as shown in Table 6).
904 This highlights the potential of Mean_c to obscure redundancies within the latent space.
905 By contrast, the mean-max correlation (MeanMax_c) offers a more sensitive perspec-
906 tive. This metric captures the maximum correlation for each mode, making it better
907 suited to identify hidden redundancies. For the CAE, the MeanMax_c score (9.7%)
908 reveals significant entanglement among modes that the Mean_c score fails to reflect.
909 The CAE latent space demonstrates clear entanglement, especially for the least con-
910 tributing LVs, a significant limitation of this architecture. This entanglement implies
911 that features detected by its convolutional layers lack independence and interpretabil-
912 ity. Consequently, the CAE risks encoding redundant features, where similar patterns
913 are repeatedly captured across different modes.

914 Although the extended POD spatial basis is orthogonal, when validating the results
915 on the test data, the temporal coefficients do not necessarily form an orthogonal basis,
916 as illustrated in Fig. 14(c), which is generated from the temporal coefficients of each
917 sample. The ranking algorithm of Kang et al. [25] is not applied in this case, as POD
918 already organizes its modes by decreasing energy contribution. This ranking leads to
919 modes with similar energy levels being more likely to exhibit correlation due to their
920 tendency to describe similar flow structures. This pattern is reflected in the correlation
921 matrix, where coefficients closer to the diagonal show stronger correlations, while those
922 further away tend to be less correlated, as indicated by the light green regions in the
923 bottom-left and top-right corners. Interestingly, a few secondary blue diagonal lines,
924 both near and distant from the matrix main diagonal, indicate higher than average
925 correlations between modes with significantly different energy contributions. Unlike
926 the CAE, the extended POD does not exhibit dark blue dots outside the diagonal,
927 implying that no mode is highly correlated with another. This is consistent with
928 the MeanMax_c score of 4.1% (see Table 6). As a result, the temporal coefficients of
929 extended POD can be considered near-orthogonal.

930 The β -VAE correlation matrix (Fig. 14(b)) demonstrates a notable improvement
931 over that of the CAE. Overall, the matrix is characterized by larger green regions,
932 indicating that the correlation coefficients of these modes are significantly damped.
933 This improvement is further validated in Table 6, where the mean correlation (Mean_c)
934 drops from 2.1% for the CAE to 1.1% for the β -VAE. The 76 physics-unaware LVs
935 are clearly visible in the correlation matrix, represented by light green lines that high-
936 light their purely stochastic behavior. Similar to the CAE correlation matrix, the
937 least contributing modes (those just above the pruned threshold) show high mutual
938 correlation, though to a much lesser degree. This observation suggests that orthogo-
939 nality improvement is not merely a global average effect but occurs at multiple levels.

940 Supporting this conclusion, the β -VAE’s mean max correlation (MeanMax_c) drops to
 941 6.9%, nearly 3 percentage points lower than the CAE’s MeanMax_c of 9.7%.

942 However, one might question to what extent the β -VAE’s superior orthogonal-
 943 ity stems from its physics-unaware modes, exhibiting very low correlations (see Fig.
 944 14(b)). To investigate this, Table 7 compares the pruned β -VAE with a CAE using
 945 $L_d = 400$ (a relevant comparison since the pruned β -VAE utilizes 424 latent variables).
 946 While both models achieve comparable reconstruction quality (with the CAE showing
 947 slightly better scores), the pruned β -VAE maintains significantly better orthogonality
 948 scores, particularly in terms of MeanMax_c . This demonstrates that the β -VAE’s regu-
 949 larization achieves more than simple mode pruning, establishing an effective trade-off
 950 between reconstruction fidelity and latent samples orthogonality.

951 A comparison of the β -VAE correlation matrix with the extended POD matrix
 952 reveals that the latter model continues to demonstrate superior performance with
 953 respect to latent samples orthogonality. The results presented in Tables 6 and 7 sub-
 954 stantiate this conclusion. As with the extended POD correlation matrix, discernible
 955 patterns are evident in the β -VAE matrix, albeit of a markedly different nature. Indeed,
 956 when excluding the physics-unaware LVs, the overall correlation of each remaining LV
 957 with other modes appears to decrease as its KL divergence increases.




Fig. 14 Correlation matrix obtained on the test dataset, with $L_d = 500$, for models: (a) CAE, (b) β -VAE, (c) extended POD. Absolute values are applied to all coefficients. Colors range from light green (uncorrelated) to dark blue (highly correlated).

958 The results demonstrate that the β -VAE achieves a balanced compromise between
 959 latent samples orthogonality, effective W^+ interpolation, and field reconstruction,
 960 while not outperforming either the CAE or extended POD in their specific strengths.
 961 By applying the mode ranking algorithm from [25], the β -VAE can be successfully
 962 reduced to a model with $L_d = 424$, while maintaining a fair trade-off between recon-
 963 struction quality and orthogonality. These results highlight the inherent challenges in
 964 simultaneously achieving both orthogonal latent variables and precise reconstruction
 965 accuracy for complex datasets requiring high latent dimensions.

6 Separability of generative factors within the latent space

The interpretability of the latent space extends beyond the orthogonality of the latent samples and the ranking of latent variables based on reconstruction contribution. An ideal latent space should exhibit effective separability between samples derived from different generative factors of the flow, a property particularly important for generating new data through latent interpolation. Clear distinctions in the latent space enable smoother and more meaningful transitions between interpolated values. In this section, the latent space organisation of each proposed dimensionality reduction model is analysed with regard to the separation of samples with different oscillation amplitudes W^+ or different oscillation phases φ (see Section 2.2 for their definitions).

Table 8 Global Davies-Bouldin indicator as given by (21) applied to clusters of oscillation amplitudes ($D_B(W^+)$) and phases ($D_B(\varphi)$), using the test dataset, with $L_d = 500$, on all models. Comparison of results for extended POD, CAE and β -VAE for $\beta = 5 \times 10^{-7}$ (see Section 5).

	CAE	β -VAE	POD
$D_B(W^+)$	23.84	28.35	41.10
$D_B(\varphi)$	27.70	33.90	27.67

Table 8 presents the Davies-Bouldin indicator (21) for the latent spaces of each model, evaluated for both W^+ and φ . Based on this metric, the CAE achieves the best separability with respect to W^+ , marginally outperforming the β -VAE, while extended POD produces the weakest clustering performance. For oscillation phases (φ), both extended POD and CAE exhibit better separability than the β -VAE. These results indicate that the latent spaces of each model are structured differently depending on these oscillation parameters, with the CAE demonstrating the most effective overall separation.

The cluster separation of oscillation amplitudes W^+ is investigated in Fig. 15 through the Davies-Bouldin indicator for two clusters W_i^+ and W_j^+ as given by (22).

Higher D_B values, indicating worse separability, are observed for clusters with similar W^+ values across all considered models. This effect is particularly pronounced in the CAE and β -VAE results, where the D_B score decreases dramatically when comparing cases with W^+ gaps of 12 and even more so for gaps of 18. For the autoencoders, the most challenging differentiations occur for $D_B(6, 12)$ and $D_B(18, 24)$, while distinguishing $W^+ = 12$ from $W^+ = 18$ proves relatively easier. This behaviour can be explained through physical interpretation: $W^+ = 12$ represents conditions closer to the unactuated case ($W^+ = 0$) characterised by stochastic fluctuations, while $W^+ = 18$ approaches the fully actuated case ($W^+ = 30$), characterised by smoother periodic fluctuations. Although both autoencoder approaches show similar trends, consistently




Fig. 15 Davies-Bouldin indicator for two clusters W_i^+ and W_j^+ ($D_B(W_i^+, W_j^+)$) for the test dataset, with $L_d = 500$, for models: (a) CAE, (b) β -VAE, (c) extended POD. Only the lower triangular part of the matrix is shown due to the symmetry of the operator.

997 lower coefficients are featured in the CAE matrix, indicating superior overall W^+
 998 separability, a result consistent with Table 8.

999 A similar yet weaker dependence on W^+ gaps is revealed in the extended POD
 1000 matrix. For instance, comparable magnitudes are demonstrated by $D_B(6, 12)$ and
 1001 $D_B(6, 18)$. Different behaviour is shown by extended POD compared to the autoen-
 1002 coders, with $D_B(12, 18)$ proving the most difficult case to differentiate. This further
 1003 confirms that the latent space is organised fundamentally differently by POD with
 1004 respect to this generative factor. Overall, significantly inferior W^+ separability is
 1005 demonstrated by POD compared to both autoencoder approaches.




Fig. 16 Davies-Bouldin indicator for two clusters φ_i and φ_j ($D_B(\varphi_i, \varphi_j)$) for the test dataset, with $L_d = 500$, for models: (a) CAE, (b) β -VAE, (c) extended POD. Only the lower triangular part of the matrix is shown due to the symmetry of the operator.

1006 As detailed in Section 2.2, each actuation cycle is decomposed into 32 distinct
 1007 phases in the dataset. All simulations in the test dataset feature actuation ($W^+ > 0$),
 1008 exhibiting periodic behaviour that justifies phase-based analysis. The dependence on
 1009 the phase of the oscillation cycle becomes more prominent with higher values of W^+ .
 1010 The separability analysis protocol is applied in Fig. 16 to evaluate phase separability
 1011 across models.

1012 For the autoencoder approaches, similar trends are exhibited by both models,
1013 with the CAE consistently achieving substantially enhanced separability. A pattern
1014 analogous to the W^+ case is observed, where poorer separability is demonstrated by
1015 temporally adjacent phases compared to distant phases. Notably, the highest intra-
1016 group D_B scores are exhibited by phases in the ranges $\varphi \in [\frac{T}{4} + \delta_\varphi, \frac{T}{2} - \delta_\varphi]$ and their
1017 symmetric counterparts $\varphi \in [\frac{3T}{4} + \delta_\varphi, T - \delta_\varphi]$, where δ_φ represents a small phase
1018 offset. These ranges, corresponding to the ascending drag regime (transitioning from
1019 $C_{f,\min}$ to $C_{f,\max}$), are distinguished by steeper temporal gradients in wall friction
1020 compared to other phases of the cycle, as identified in Gu erin et al. [10]. Additionally,
1021 the phases $\varphi = \frac{T}{4}$ and $\varphi = \frac{3T}{4}$ correspond to the extremal values of the control
1022 signal $w_w(t)$, where $w_w(\frac{T}{4}) = W$ and $w_w(\frac{3T}{4}) = -W$, which may explain the poor
1023 separability scores observed around these phases.

1024 While the symmetric nature of friction variation arising from the sinusoidal wall
1025 oscillation [10] might suggest difficulty in distinguishing phases separated by $\frac{T}{2}$,
1026 efficient separation is actually demonstrated in the autoencoder latent spaces for
1027 these cases. This suggests that other physical distinctions, such as Stokes strain sign
1028 reversal or streak orientation (see Gu erin et al. 10), are effectively captured by the
1029 autoencoders.

1030 Distinct behaviour is exhibited by the extended POD approach, as evidenced by
1031 elevated D_B scores for phases φ near $\frac{T}{8}$ with their symmetric counterparts at $\varphi \approx \frac{5T}{8}$.
1032 This observation suggests that difficulties are encountered by the extended POD latent
1033 space in distinguishing these phase pairs, which correspond to similar drag and heat
1034 transfer characteristics but opposing streak orientations. More broadly, the presence
1035 of the corresponding non-diagonal line in Fig. 16(c) reveals that phases $\varphi_1 \in [0, \frac{T}{4}]$
1036 are not well separated from phases $\varphi_2 \approx \frac{3T}{4} - \varphi_1$. This implies that the extended
1037 POD methodology struggles to differentiate between friction-descending phase pairs
1038 φ_1/φ_2 originating from opposite halves of the actuation cycle and displaying nearly
1039 identical distances from $C_{f,\min}/C_{f,\max}$ phases. These phase pairs are characterised by
1040 substantially different streak orientations and opposite Stokes strain directions in the
1041 actuated case, which confirms the difficulty of extended POD to effectively capture
1042 related flow features within its latent space.

1043 Beyond these distinctions, phase separability trends similar to the autoencoders
1044 are demonstrated by extended POD, albeit with better performance for phase pairs
1045 outside the challenging regions with ascending drag.

1046 Three primary results are obtained from the separability analysis of generative
1047 factors of the flow:

- 1048 1. Superior performance in separating different W^+ amplitudes is consistently
1049 achieved by autoencoder architectures compared to extended POD, with the CAE
1050 achieving optimal capability. This suggests that amplitude-dependent flow features
1051 are better captured in the latent representations of autoencoders.
- 1052 2. Similar phase separation characteristics are exhibited by both CAE and β -VAE,
1053 showing:
 - 1054 • Reduced separability in ascending friction phases.
 - 1055 • Superior overall performance achieved by the CAE.

1056 • Effective discrimination of phases with different streak orientation or opposing
1057 Stokes strain direction.

1058 3. The best global $D_B(\varphi)$ score is achieved by extended POD (see Table 8) and excel-
1059 lent separability is shown for certain phase combinations. However, weak separation
1060 of part of the descending friction phase-pairs is provided by this method, which
1061 differ significantly in terms of streak patterns and Stokes strain.

1062 While extended POD is not matched by autoencoders in terms of latent samples
1063 orthogonality, more physically meaningful characterisations of the generative factors of
1064 the flow are provided by their learned representations, offering enhanced interpretabil-
1065 ity. The best performance overall was achieved by the CAE, especially with regard to
1066 oscillation phases. The approach to investigate these separabilities remains basic, con-
1067 sidering one cluster for each phase or amplitude. Deeper insights into the separability
1068 of SWO amplitudes and phases, as well as mode correlations, may be provided by
1069 employing Cluster-based Reduced-Order Modelling (CROM) as proposed by Kaiser
1070 et al. [48]. Additionally, new information about the interactions between C_f and Nu,
1071 particularly concerning the breakdown of the Reynolds analogy, may be revealed. This
1072 extended cluster-based analysis remains a topic for future investigations.

1073 7 Latent variable interpretability

1074 This study has demonstrated that autoencoders outperform POD in both reconstruc-
1075 tion quality and the separability of the flow’s generative factors within the latent space.
1076 To further assess interpretability, a preliminary exploratory analysis is conducted
1077 wherein correlations are examined between the latent variables and four physical quan-
1078 tities: the friction coefficient C_f , the Nusselt number Nu, the oscillation phases φ
1079 and the amplitudes W^+ . These quantities are of particular interest, as the present
1080 actuation strategy resulted in preferential heat-transfer enhancement at $W^+ = 30$ [10].

1081 7.1 Methodological Limitations

1082 It is essential to acknowledge a fundamental constraint of this analysis: the Pear-
1083 son correlation coefficient employed herein measures exclusively linear relationships,
1084 whereas the principal advantage of autoencoders over POD resides in their capacity
1085 to capture non-linear projections. This limitation is particularly pronounced for the
1086 CAE, whose training objective (minimisation of reconstruction error) imposes no con-
1087 straints favouring disentangled representations. Consequently, physical information
1088 is expected to be distributed across multiple latent variables in complex, non-linear
1089 patterns that linear correlation analysis is inherently unable to detect. The absence
1090 of strong linear correlations thus does not indicate the absence of encoded physical
1091 information, but rather reflects the inadequacy of the linear probing methodology for
1092 non-linear latent spaces.

1093 To quantify linear relationships, the Pearson correlation coefficient is computed
1094 across all dataset samples. For a physical quantity ϕ (W^+ , φ , C_f , or Nu), the
1095 correlation between ϕ and the i^{th} latent variable is given by:

$$C(\phi, LV(i)) = \frac{\langle z'_{j,k}(i)\phi'_{j,k} \rangle_{j,k}}{\sqrt{\langle z'_{j,k}(i)z'_{j,k}(i) \rangle_{j,k} \langle \phi'_{j,k}\phi'_{j,k} \rangle_{j,k}}}, \quad (24)$$

1096 where the averaging and fluctuations over simulations j and snapshots k are defined
1097 as in Section 3.4.

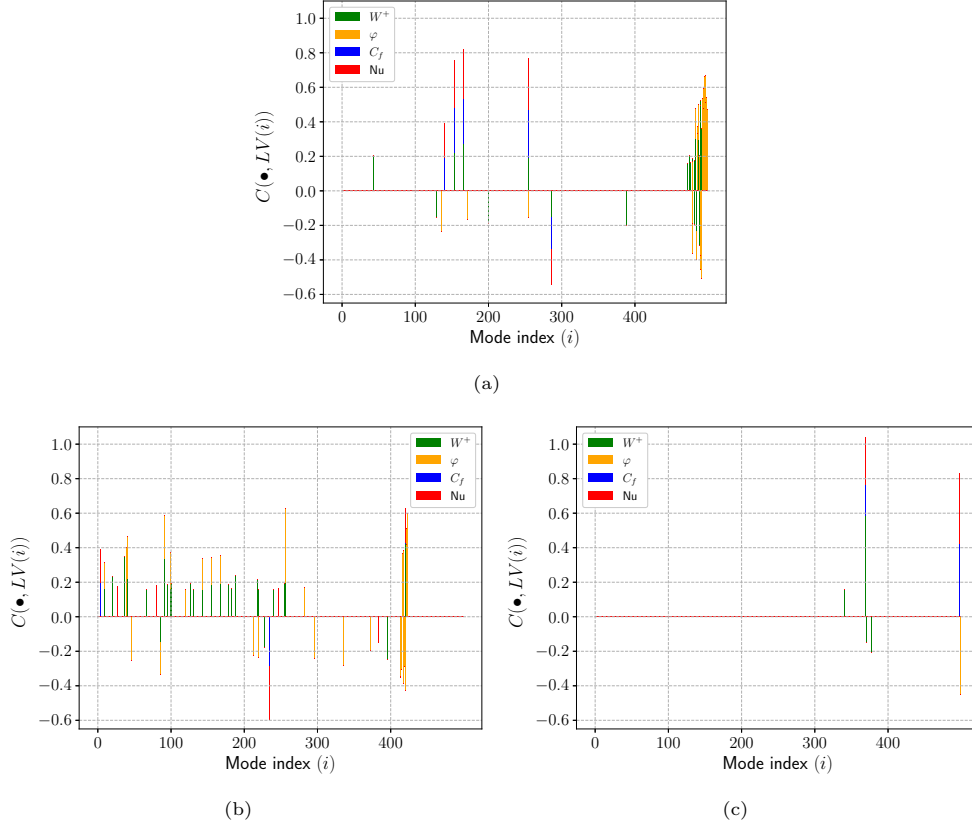


Fig. 17 Pearson correlation between latent variables and physical quantities: oscillation amplitudes W^+ in green, oscillation phases φ in orange, friction coefficient C_f in blue, and Nusselt number Nu in red. For each latent variable, correlations with the physical quantities are displayed as stacked bars, with each subsequent bar beginning at the endpoint of the previous one to enable direct comparison. Latent variables are ranked according to Section 5, and correlations with $|C(\bullet, LV(i))| < 0.15$ are omitted for clarity. Results are presented on the test data for (a): CAE, (b): β -VAE, and (c): extended POD.

1098 **7.2 Results and Interpretation**

1099 Fig. 17 presents the correlation results for the CAE, β -VAE, and extended-POD
 1100 models, with latent variables ranked using the algorithm of Kang et al. [25] for autoen-
 1101 coders, whilst POD modes retain their inherent energy-based ordering. Correlations
 1102 below 0.15 in absolute value are omitted. Across all models, no significant correlation
 1103 (above 0.7) is observed, and the majority of latent variables exhibit correlations below
 1104 the threshold, indicating predominantly weak or absent linear relationships.

1105 The CAE results (Fig. 17(a)) demonstrate negligible correlations for the most
 1106 reconstruction-contributing latent variables (first 100 LVs), consistent with the
 1107 expectation that unconstrained non-linear optimisation produces highly entangled rep-
 1108 resentations wherein physical information is encoded in complex, non-linear patterns
 1109 inaccessible to linear probing. Intermediate modes (indices 100–300) exhibit modest
 1110 correlations (0.2–0.3), whilst the least contributing modes display stronger correlations
 1111 with W^+ and φ (up to 0.6).

1112 The β -VAE exhibits a substantially greater number of detectable linear correlations
 1113 (Fig. 17(b)), including amongst its most reconstruction-contributing variables. This
 1114 enhancement is directly attributable to the regularisation term in equation (15), which
 1115 penalises complex, entangled structures by minimising the KL divergence between the
 1116 latent distribution and a standard Gaussian prior. This constraint forces simpler, more
 1117 disentangled representations wherein relationships between latent variables and phys-
 1118 ical quantities manifest in more direct, linearly detectable forms. The decorrelation of
 1119 latent variables imposed by the regularisation thus renders the encoded physical infor-
 1120 mation more amenable to linear probing techniques, particularly amongst the most
 1121 reconstruction-contributing modes. Importantly, this apparent improvement in “inter-
 1122 pretability” represents a trade-off: the β -VAE sacrifices some reconstruction efficiency
 1123 and optimal non-linear encoding capacity to achieve a simpler latent structure that
 1124 is more accessible to linear analysis. The enhanced detectability of correlations thus
 1125 reflects regularisation-imposed simplification rather than superior information encod-
 1126 ing per se. Deep entanglement persists despite regularization, as numerous distinct
 1127 latent variables demonstrate correlations to the same physical quantities.

1128 Extended-POD demonstrates the weakest linear correlations, with only six modes
 1129 exceeding the threshold (Fig. 17(c)), all corresponding to higher-order (low-energy)
 1130 modes. This observation is consistent with POD’s disentanglement (as quantified by
 1131 the correlation metrics in Table 6), yet also reflects a fundamental limitation of the
 1132 POD methodology: modes are ordered by energy contribution rather than by physical
 1133 significance. Consequently, there is no *a priori* reason for specific modes (especially the
 1134 most energetic ones) to encode particular physical quantities. This encoding strategy,
 1135 whilst maintaining strong orthogonality properties, does not prioritise alignment with
 1136 specific physical generative factors or quantities.

1137 A common result across all models is that latent variables correlated with C_f also
 1138 exhibit a similar correlation with Nu , consistent with the Reynolds analogy. For a
 1139 limited number of these modes, a concurrent correlation with W^+ is observed, likely
 1140 due to the gradual breakdown of the Reynolds analogy as W^+ increases. Furthermore,
 1141 the least contributing modes in both AE models (excluding the pruned ones for the
 1142 β -VAE) demonstrate strong correlations with φ or W^+ . This highlights the significant

1143 entanglement among these modes, as previously illustrated in Fig. 14, and suggests
1144 that they contribute to capturing the phase and amplitude behavior of the flow.

1145 7.3 Connection to Latent Space Separability

1146 The correlation patterns observed herein may be related to the separability analysis
1147 presented in Section 6, wherein the Davies-Bouldin index was employed to quantify
1148 cluster separation with respect to oscillation parameters. It is pertinent to recognise,
1149 however, that these two metrics probe fundamentally distinct aspects of latent space
1150 organisation. The Pearson correlation coefficient assesses local, linear relationships
1151 between individual latent variables and physical quantities, thereby measuring the
1152 extent to which a single mode varies linearly with a given parameter. In contrast, the
1153 Davies-Bouldin index evaluates global geometric structure, quantifying the separa-
1154 tion between cluster centroids relative to intra-cluster variance across the entire latent
1155 space. As such, strong linear correlations are neither necessary nor sufficient for achiev-
1156 ing superior cluster separability, as effective geometric organisation may be realised
1157 through complex, non-linear patterns distributed across multiple latent dimensions.

1158 This conceptual distinction is demonstrated by the CAE results. The model
1159 achieves superior separability for amplitude-based clustering (lowest $D_B(W^+)$ score in
1160 Table 6) despite exhibiting weak linear correlations between individual latent variables
1161 and W^+ . This apparent paradox is resolved by recognising that physical distinctions
1162 between amplitudes are embedded through complex, non-linear encodings distributed
1163 across multiple latent dimensions. Whilst individual modes exhibit negligible linear
1164 correlation with W^+ , the collective geometric structure of the latent space permits
1165 effective amplitude-based clustering. The Davies-Bouldin metric, being sensitive to
1166 centroid separation in high-dimensional space, successfully detects this global organi-
1167 sation, whereas the Pearson coefficient, probing only linear relationships in individual
1168 dimensions, fails to capture the distributed, non-linear encoding mechanism.

1169 The β -VAE demonstrates moderately degraded separability relative to the CAE
1170 (intermediate D_B scores in Table 6), yet exhibits substantially enhanced linear corre-
1171 lation, particularly amongst its most reconstruction-contributing modes. This inverse
1172 relationship between separability and linear detectability represents the direct conse-
1173 quence of the regularisation strategy: the KL divergence term in equation (15)
1174 constrains the latent distribution to approximate a standard Gaussian prior, thereby
1175 encouraging simpler, more linearly structured encodings. Whilst this regularisation
1176 enhances the linear detectability of physical relationships (rendering them accessible
1177 to correlation-based probing), it simultaneously constrains the model's capacity to
1178 construct optimally separated geometric clusters through complex, non-linear trans-
1179 formations. The trade-off is thus explicit: the β -VAE sacrifices some degree of optimal
1180 geometric organisation to achieve a more linearly interpretable latent structure.

1181 Extended-POD exhibits particularly poor separability for amplitude-based clus-
1182 tering (highest $D_B(W^+)$ score in Table 6) alongside intermediate performance for
1183 phase-based organisation, despite maintaining superior orthogonality (as quantified in
1184 Table 6). This behaviour is directly attributable to POD's fundamental organisational
1185 principle: modes are ordered by energy contribution rather than by alignment with

1186 physical generative factors. The orthogonal basis constructed through POD is opti-
1187 mised exclusively for capturing variance in descending order, with no mechanism to
1188 ensure that specific modes encode particular physical parameters such as W^+ or φ .
1189 Samples characterised by different amplitudes or phases are thus not systematically
1190 organised within distinct geometric regions of the latent space, resulting in poorly
1191 separated clusters and elevated Davies-Bouldin scores. The autoencoders, by contrast,
1192 learn latent representations through data-driven non-linear optimisation, thereby con-
1193 structing structures that inherently reflect the underlying generative factors. This
1194 fundamental difference in organisational principle explains why the autoencoders
1195 achieve substantially superior physical separability despite the CAE’s employment of
1196 highly non-linear, entangled encodings and despite POD’s superior orthogonality.

1197 8 Conclusions and perspectives

1198 This study comprehensively evaluated the effectiveness of proper orthogonal decom-
1199 position and various autoencoder architectures in reducing the dimensionality of
1200 turbulent channel flow data at $\text{Re}_\tau = 200$ and $\text{Pr} = 1$. The dataset, comprising
1201 velocity and temperature fields, was generated using Large Eddy Simulations under
1202 spanwise wall oscillations with a fixed oscillation period and varying amplitudes. Spe-
1203 cific attention was set towards small turbulent scales reconstruction, orthogonality of
1204 the latent samples, and latent separability of generative factors. The main results from
1205 the investigation are summarised as follows:

- 1206 • CAEs consistently outperformed POD in terms of reconstruction accuracy across
1207 all flow fields, particularly at lower latent dimensions. The most significant improve-
1208 ment was observed in the wall-normal velocity component (v), highlighting the
1209 advantage of non-linearity offered by AE architectures in capturing complex tur-
1210 bulent dynamics. Furthermore, the CAE demonstrated superior generalization
1211 capability across the range of oscillation amplitudes.
- 1212 • Compressing all four fields (u , v , w , Θ) simultaneously resulted in superior recon-
1213 struction performance compared to individual field compression. This enhancement
1214 leverages the inherent physical correlations between the variables, especially between
1215 the streamwise velocity (u) and temperature (Θ) fields.
- 1216 • Both POD and AE methodologies encountered difficulties in accurately recon-
1217 structing small-scale turbulent structures, particularly within the v -field. POD, con-
1218 strained to 500 modes, was limited to capturing only large-scale flow characteristics,
1219 and exhibited a tendency towards mean-value approximation.
- 1220 • β -VAEs achieved a balance between the reconstruction performance of CAEs and
1221 latent samples orthogonality of extended POD. Although there was a decrease in
1222 accuracy for the v and Θ fields, correlation metrics indicated improved latent space
1223 structure compared to CAEs.
- 1224 • The β -VAE maintains optimal performance with only 424 latent variables ($L_d =$
1225 424), preserving an effective balance between reconstruction fidelity and orthogo-
1226 nality. Notably, in this pruned architecture, latent variables exhibiting greater KL
1227 divergence tend to show reduced inter-mode correlations.

- 1228 • Both CAEs and β -VAEs demonstrate superior capability in distinguishing
1229 amplitude-related features (W^+), compared to extended POD. Autoencoders
1230 additionally incorporate key physical insights by differentiating phases (φ) with sub-
1231 stantially different Stokes strain or streak orientation. Despite similar trends, the
1232 CAE consistently outperformed the β -VAE in terms of latent separability.
- 1233 • The β -VAE sacrifices some degree of optimal geometric organisation to achieve a
1234 more linearly interpretable latent structure. The CAE's weak linear correlations for
1235 its dominant modes do not indicate absence of physical information, but rather its
1236 encoding in complex, non-linear patterns. Extended-POD's weak correlations and
1237 poor physical separability arise from its energy-based rather than physics-based
1238 mode ordering. For all models, latent variables correlated with C_f consistently
1239 exhibit similar correlations with Nu, reflecting the Reynolds analogy.

1240 Based on the current results, future research directions are proposed to further
1241 enhance the understanding and application of dimensionality reduction techniques in
1242 turbulent flows or to provide a reduced order dynamical model for the present dataset:

- 1243 • Employ CROM techniques [48] to elucidate the dynamics encoded within the latent
1244 spaces of each model. This approach, which is more complex than that of the present
1245 study, may reveal more meaningful information regarding how each model structures
1246 its latent space with regards to physical parameters.
- 1247 • Investigate the application of three-dimensional convolutional layers to compress
1248 the entire 3D channel, as opposed to the current two-dimensional slice approach.
1249 This holistic compression strategy may capture spatial correlations more effectively,
1250 improving reconstruction accuracy and model generalisation. However, compressing
1251 the 3D space, characterized by a much larger dimension, may prove a challenging
1252 task.
- 1253 • Given the inherent limitations of linear probing in non-linear latent spaces, future
1254 investigations should employ non-linear probing techniques (such as mutual informa-
1255 tion analysis, supervised classification tasks, or non-linear dimensionality reduction
1256 of the latent space) to thoroughly evaluate the physical information encoded within
1257 these high-dimensional representations.

1258 A primary and promising direction is to extend the learned latent representations
1259 for temporal forecasting and control. This could be achieved by combining the AE-
1260 based dimensionality reduction with long short term memory networks or transformers
1261 [24] to develop an efficient predictive model of the flow, thereby facilitating real-time
1262 flow control and optimisation in engineering applications.

1263 Appendix

1264 Generalised Stokes Layer Solution

1265 For an arbitrary periodic spanwise wall motion, the laminar Stokes layer solution
1266 is constructed through harmonic superposition [13]. Consider a wall velocity signal
1267 decomposed into its Fourier series:




Fig. 18 Phase averaged wall-normal velocity $\tilde{w}(y, \varphi)$ for oscillation parameters $T^+ = 500$, $W^+ = 30$, at $Re_\tau = 180$. Blue contours indicate the laminar analytical case (27), while red contours correspond to the turbulent case. Solid contours correspond to positive values while dashed contours indicate negative values.

$$w_w(t) = W_m \sum_{n=1}^{\infty} A_n \sin\left(\frac{2\pi nt}{T} + \phi_n\right). \quad (25)$$

1268 where W_m represents the maximum wall velocity, A_n are the normalised harmonic
 1269 amplitudes satisfying $\sum_{n=1}^{\infty} |A_n|^2 = 1$, and ϕ_n are the phase angles. The corresponding
 1270 laminar Stokes layer velocity profile is given by:

$$w_{St}(y, t) = W_m \sum_{n=1}^{\infty} A_n \exp\left(-\sqrt{n} \frac{y}{\delta_{s,1}}\right) \sin\left(\frac{2\pi nt}{T} + \phi_n - \sqrt{n} \frac{y}{\delta_{s,1}}\right), \quad (26)$$

1271 where $\delta_{s,1} = \sqrt{2\nu T/(2\pi)}$ represents the characteristic Stokes layer thickness for
 1272 the fundamental frequency.

1273 **Sinusoidal Specialisation:**

1274 For the specific case of sinusoidal wall motion $w_w(t) = W \sin(2\pi t/T)$, we have
 1275 $W_m = W$, $A_1 = 1$, $A_n = 0$ for $n > 1$, and $\phi_1 = 0$, reducing equation (26) to:

$$w_{St}(y, t) = W \exp\left(-\frac{y}{\delta_s}\right) \sin\left(\frac{2\pi t}{T} - \frac{y}{\delta_s}\right), \quad (27)$$

1276 where $\delta_s = \delta_{s,1}$ is the classical Stokes layer thickness. Fig. 18 illustrates the differ-
 1277 ence between the laminar and turbulent phase-averaged profiles, where the analytical
 1278 solution results in a deeper penetration of the Stokes layer into the buffer layer.

1279 **Statements and Declarations**

1280 **Conflict of interest**

1281 The authors report no conflict of interest.

1282 **Data availability statement**

1283 The dataset used in this study is publicly available on the following link:
1284 <https://doi.org/10.57745/GU68MY>.

1285 **References**

- 1286 [1] Ricco, P., Skote, M., Leschziner, M.A.: A review of turbulent skin-friction drag
1287 reduction by near-wall transverse forcing. *Progress in Aerospace Sciences* **123**,
1288 100713 (2021). Publisher: Elsevier
- 1289 [2] Abdulbari, H.A., Zulkifli, R., Salleh, M.Z., Yusoff, M.Z.: Skin-Friction Drag
1290 Reduction Using Riblets, Dimples, and Oscillating Wall. *Journal of Mechanical*
1291 *Engineering and Sciences* **4**, 452–464 (2013)
- 1292 [3] Asidin, S., Kuntjoro, W., Salleh, M.Z.: A Review on Flow Control Methods in
1293 Turbine Internal Cooling. *Applied Sciences* **9**(16), 3233 (2019)
- 1294 [4] Zhang, X., Li, T., Jiang, N.: Flow Control Strategies for Improving Performance of
1295 Turbomachinery: A Review. *Progress in Aerospace Sciences* **115**, 100595 (2020)
- 1296 [5] Agostini, L., Touber, E., Leschziner, M.: The turbulence vorticity as a window
1297 to the physics of friction-drag reduction by oscillatory wall motion. *International*
1298 *Journal of Heat and Fluid Flow* **51**, 3–15 (2015)
- 1299 [6] Quadrio, M., Ricco, P.: Critical assessment of turbulent drag reduction through
1300 spanwise wall oscillations. *Journal of Fluid Mechanics* **521**, 251–271 (2004)
- 1301 [7] Choi, K., Clayton, B.R.: The mechanism of turbulent drag reduction with wall
1302 oscillation. *International Journal of Heat and Fluid Flow* **22**(1), 1–9 (2001) [https://doi.org/10.1016/S0142-727X\(00\)00070-9](https://doi.org/10.1016/S0142-727X(00)00070-9)
- 1303
- 1304 [8] Viotti, C., Quadrio, M., Luchini, P.: Streamwise oscillation of spanwise velocity
1305 at the wall of a channel for turbulent drag reduction. *Physics of fluids* **21**(11),
1306 115109 (2009)
- 1307 [9] Touber, E., Leschziner, M.: Near-wall streak modification by spanwise oscillatory
1308 wall motion and drag-reduction mechanisms. *Journal of Fluid Mechanics* **693**,
1309 150–200 (2012)
- 1310 [10] Guérin, L., Flageul, C., Cordier, L., Grieu, S., Agostini, L.: Preferential enhance-
1311 ment of convective heat transfer over drag via near-wall turbulence manipulation
1312 using spanwise wall oscillations. *International Journal of Heat and Fluid Flow*
1313 **110**, 109564 (2024) <https://doi.org/10.1016/j.ijheatfluidflow.2024.109564>
- 1314 [11] Rouhi, A., Hultmark, M., Smits, A.J.: Spanwise wall forcing can reduce turbulent
1315 heat transfer more than drag. *Journal of Fluid Mechanics* **1010**, 59 (2025) <https://doi.org/10.1017/jfm.2025.310>
- 1316

- 1317 [12] Rouhi, A., Fu, M.K., Chandran, D., Zampiron, A., Smits, A.J., Marusic, I.: Tur-
1318 bulent drag reduction by spanwise wall forcing. Part 1. Large-eddy simulations.
1319 Journal of Fluid Mechanics **968**, 6 (2023) <https://doi.org/10.1017/jfm.2023.499>
- 1320 [13] Cimarelli, A., Frohnapfel, B., Hasegawa, Y., Angelis, E.D., Quadrio, M.: Pre-
1321 diction of turbulence control for arbitrary periodic spanwise wall movement.
1322 Physics of Fluids **25**(7) (2013) <https://doi.org/10.1063/1.4813807> . Publisher:
1323 AIP Publishing
- 1324 [14] Berkooz, G., Holmes, P., Lumley, J.L.: The proper orthogonal decomposition in
1325 the analysis of turbulent flows. Annual review of fluid mechanics **25**(1), 539–575
1326 (1993)
- 1327 [15] Agostini, L.: Exploration and prediction of fluid dynamical systems using auto-
1328 encoder technology. Physics of Fluids **32**(6), 067103 (2020) [https://doi.org/10.](https://doi.org/10.1063/5.0012906)
1329 [1063/5.0012906](https://doi.org/10.1063/5.0012906)
- 1330 [16] Linot, A.J., Graham, M.D.: Dynamics of a data-driven low-dimensional model
1331 of turbulent minimal Couette flow. Journal of Fluid Mechanics **973**, 42 (2023)
1332 <https://doi.org/10.1017/jfm.2023.720>
- 1333 [17] Linot, A.J., Zeng, K., Graham, M.D.: Turbulence control in plane Couette flow
1334 using low-dimensional neural ODE-based models and deep reinforcement learn-
1335 ing. International Journal of Heat and Fluid Flow **101**, 109139 (2023) <https://doi.org/10.1016/j.ijheatfluidflow.2023.109139>
- 1337 [18] Gruber, A., Gunzburger, M., Ju, L., Wang, Z.: A comparison of neural net-
1338 work architectures for data-driven reduced-order modeling. Computer Methods
1339 in Applied Mechanics and Engineering **393**, 114764 (2022) [https://doi.org/10.](https://doi.org/10.1016/j.cma.2022.114764)
1340 [1016/j.cma.2022.114764](https://doi.org/10.1016/j.cma.2022.114764)
- 1341 [19] Bank, D., Koenigstein, N., Giryes, R.: Autoencoders. arXiv. arXiv:2003.05991 [cs]
1342 (2021)
- 1343 [20] Agostini, L., Leschziner, M.: Auto-encoder-assisted analysis of amplitude and
1344 wavelength modulation of near-wall turbulence by outer large-scale structures
1345 in channel flow at friction Reynolds number of 5200. Physics of Fluids **34**(11),
1346 115142 (2022) <https://doi.org/10.1063/5.0123119>
- 1347 [21] Du, X., Qu, X., He, Y., Guo, D.: Single Image Super-Resolution Based on Multi-
1348 Scale Competitive Convolutional Neural Network. Sensors **18**(3), 789 (2018)
1349 <https://doi.org/10.3390/s18030789>
- 1350 [22] Csala, H., Dawson, S.T.M., Arzani, A.: Comparing different nonlinear dimen-
1351 sionality reduction techniques for data-driven unsteady fluid flow modeling.
1352 Physics of Fluids **34**(11) (2022) <https://doi.org/10.1063/5.0127284> . Publisher:
1353 AIP Publishing

- 1354 [23] Solera-Rico, A., Vila, C.S., Gómez-López, M., Wang, Y., Almashjary, A., Dawson,
1355 S.T.M., Vinuesa, R.: β -Variational autoencoders and transformers for reduced-
1356 order modelling of fluid flows. *Nature Communications* **15**(1), 1–15 (2024) <https://doi.org/10.1038/s41467-024-45578-4> . Number: 1 Publisher: Nature Publishing
1357 Group
1358
- 1359 [24] Eivazi, H., Clainche, S.L., Hoyas, S., Vinuesa, R.: Towards extraction of orthog-
1360 onal and parsimonious non-linear modes from turbulent flows. *Expert Systems*
1361 *with Applications* **202**, 117038 (2022). Publisher: Elsevier
- 1362 [25] Kang, Y., Yang, S., Yee, K.: Physics-aware Reduced-order Modeling of Transonic
1363 Flow via β -Variational Autoencoder. Publisher: arXiv Version Number: 2 (2022).
1364 <https://doi.org/10.48550/ARXIV.2205.00608>
- 1365 [26] Mohan, A.T., Lubbers, N., Chertkov, M., Livescu, D.: Embedding hard phys-
1366 ical constraints in neural network coarse-graining of three-dimensional tur-
1367 bulence. *Physical Review Fluids* **8**(1), 014604 (2023) [https://doi.org/10.1103/](https://doi.org/10.1103/PhysRevFluids.8.014604)
1368 [PhysRevFluids.8.014604](https://doi.org/10.1103/PhysRevFluids.8.014604) . Publisher: American Physical Society
- 1369 [27] Fukami, K., Taira, K.: Grasping extreme aerodynamics on a low-dimensional
1370 manifold. *Nature Communications* **14**(1), 6480 (2023) [https://doi.org/10.1038/](https://doi.org/10.1038/s41467-023-42213-6)
1371 [s41467-023-42213-6](https://doi.org/10.1038/s41467-023-42213-6)
- 1372 [28] Bevanda, P., Sosnowski, S., Hirche, S.: Koopman operator dynamical models:
1373 Learning, analysis and control. *Annual Reviews in Control* **52**, 197–212 (2021)
1374 <https://doi.org/10.1016/j.arcontrol.2021.09.002> . Accessed 2025-10-19
- 1375 [29] Brunton, S.L., Proctor, J.L., Kutz, J.N.: Discovering governing equations from
1376 data by sparse identification of nonlinear dynamical systems. *Proceedings of the*
1377 *National Academy of Sciences* **113**(15), 3932–3937 (2016) [https://doi.org/10.](https://doi.org/10.1073/pnas.1517384113)
1378 [1073/pnas.1517384113](https://doi.org/10.1073/pnas.1517384113) . Accessed 2025-10-19
- 1379 [30] Williams, J.P., Zahn, O., Kutz, J.N.: Sensing with shallow recurrent decoder
1380 networks. *Proceedings of the Royal Society A: Mathematical, Physical and Engi-*
1381 *neering Sciences* **480**(2298), 20240054 (2024) [https://doi.org/10.1098/rspa.2024.](https://doi.org/10.1098/rspa.2024.0054)
1382 [0054](https://doi.org/10.1098/rspa.2024.0054) . Accessed 2025-10-19
- 1383 [31] Guérin, L., Cordier, L., Flageul, C., Grieu, S., Agostini, L.: Policy-based opti-
1384 mization for drag reduction via spanwise wall oscillations. *Neural Computing and*
1385 *Applications* (2025) <https://doi.org/10.1007/s00521-025-11067-y>
- 1386 [32] Kasagi, N., Tomita, Y., Kuroda, A.: Direct Numerical Simulation of Passive Scalar
1387 Field in a Turbulent Channel Flow. *Journal of Heat Transfer* **114**(3), 598–606
1388 (1992) <https://doi.org/10.1115/1.2911323>
- 1389 [33] Laizet, S., Lamballais, E.: High-order compact schemes for incompressible flows:

- 1390 A simple and efficient method with quasi-spectral accuracy. *Journal of Computa-*
 1391 *tional Physics* **228**(16), 5989–6015 (2009) [https://doi.org/10.1016/j.jcp.2009.05.](https://doi.org/10.1016/j.jcp.2009.05.010)
 1392 [010](https://doi.org/10.1016/j.jcp.2009.05.010)
- 1393 [34] Laizet, S., Li, N.: Incompact3d: A powerful tool to tackle turbulence problems
 1394 with up to $O(10^5)$ computational cores. *International Journal for Numerical*
 1395 *Methods in Fluids* **67**(11), 1735–1757 (2011) <https://doi.org/10.1002/fld.2480>
- 1396 [35] Bartholomew, P., Deskos, G., Frantz, R.A., Schuch, F.N., Lamballais, E., Laizet,
 1397 S.: Xcompact3D: An open-source framework for solving turbulence problems on
 1398 a Cartesian mesh. *SoftwareX* **12**, 100550 (2020) [https://doi.org/10.1016/j.softx.](https://doi.org/10.1016/j.softx.2020.100550)
 1399 [2020.100550](https://doi.org/10.1016/j.softx.2020.100550)
- 1400 [36] Abe, H., Antonia, R.A., Kawamura, H.: Correlation between small-scale velocity
 1401 and scalar fluctuations in a turbulent channel flow. *Journal of Fluid Mechanics*
 1402 **627**, 1–32 (2009) <https://doi.org/10.1017/S0022112008005569>
- 1403 [37] Seki, Y., Iwamoto, K., Kawamura, H.: Prandtl number Effect on Turbulence
 1404 Quantities through High Spatial Resolution DNS of Turbulent Heat Transfer in
 1405 a Channel Flow. In: *Proceedings of the International Symposium on Turbulence,*
 1406 *Heat and Mass Transfer*, pp. 301–304. Begellhouse, Dubrovnik, Croatia (2006).
 1407 <https://doi.org/10.1615/ICHMT.2006.TurbulHeatMassTransf.560>
- 1408 [38] Flageul, C., Benhamadouche, S., Lamballais, E., Laurence, D.: DNS of turbulent
 1409 channel flow with conjugate heat transfer: Effect of thermal boundary conditions
 1410 on the second moments and budgets. *International Journal of Heat and Fluid*
 1411 *Flow* **55**, 34–44 (2015) <https://doi.org/10.1016/j.ijheatfluidflow.2015.07.009>
- 1412 [39] Lamballais, E., Fortuné, V., Laizet, S.: Straightforward high-order numerical dis-
 1413 sipation via the viscous term for direct and large eddy simulation. *Journal of*
 1414 *Computational Physics* **230**(9), 3270–3275 (2011) [https://doi.org/10.1016/j.jcp.](https://doi.org/10.1016/j.jcp.2011.01.040)
 1415 [2011.01.040](https://doi.org/10.1016/j.jcp.2011.01.040)
- 1416 [40] Cordier, L., Bergmann, M.: Proper Orthogonal Decomposition: an overview. In:
 1417 *Lecture Series 2002-04, 2003-03 and 2008-01 on Post-processing of Experimental*
 1418 *and Numerical Data*, pp. 1–46. Von Kármán Institute for Fluid Dynamics, ???
 1419 (2008). ISBN 978-2-930389-80-X
- 1420 [41] Borée, J.: Extended proper orthogonal decomposition: a tool to analyse correlated
 1421 events in turbulent flows. *Experiments in Fluids* **35**(2), 188–192 (2003) <https://doi.org/10.1007/s00348-003-0656-3>. Company: Springer Distributor: Springer
 1422 Institution: Springer Label: Springer Number: 2 Publisher: Springer-Verlag
 1423
- 1424 [42] Fukami, K., Nakamura, T., Fukagata, K.: Convolutional neural network based
 1425 hierarchical autoencoder for nonlinear mode decomposition of fluid field data.
 1426 *Physics of Fluids* **32**(9) (2020) <https://doi.org/10.1063/5.0020721>. Publisher:
 1427 AIP Publishing

- 1428 [43] Rasamoelina, A.D., Adjailia, F., Sinčák, P.: A review of activation function for
1429 artificial neural network. In: 2020 IEEE 18th World Symposium on Applied
1430 Machine Intelligence and Informatics (SAMI), pp. 281–286 (2020). IEEE
- 1431 [44] Liu, L., Jiang, H., He, P., Chen, W., Liu, X., Gao, J., Han, J.: On the variance of
1432 the adaptive learning rate and beyond. arXiv preprint arXiv:1908.03265 (2019)
- 1433 [45] Kingma, D.P., Welling, M.: An Introduction to Variational Autoencoders. Founda-
1434 tions and Trends® in Machine Learning **12**(4), 307–392 (2019) [https://doi.
1435 org/10.1561/22000000056](https://doi.org/10.1561/22000000056)
- 1436 [46] Hoffman, M.D., Johnson, M.J.: Elbo surgery: yet another way to carve up the
1437 variational evidence lower bound. In: Workshop in Advances in Approximate
1438 Bayesian Inference, NIPS, vol. 1 (2016)
- 1439 [47] Xu, D., Tian, Y.: A Comprehensive Survey of Clustering Algorithms. Annals of
1440 Data Science **2**(2), 165–193 (2015) <https://doi.org/10.1007/s40745-015-0040-1>
- 1441 [48] Kaiser, E., Noack, B.R., Cordier, L., Spohn, A., Segond, M., Abel, M., Daviller,
1442 G., Östh, J., Krajnović, S., Niven, R.K.: Cluster-based reduced-order modelling
1443 of a mixing layer. Journal of Fluid Mechanics **754**, 365–414 (2014). Publisher:
1444 Cambridge University Press