THÈSE

Pour l'obtention du grade de
DOCTEUR DE L'UNIVERSITÉ DE POITIERS
UFR des sciences fondamentales et appliquées
Pôle poitevin de recherche pour l'ingénieur en mécanique, matériaux et énergétique
(Diplôme National - Arrêté du 25 mai 2016)

École doctorale : Mathématiques, informatique, matériaux, mécanique, énergétique, Poitiers
Secteur de recherche : Mécanique des fluides

Présentée par :
Lou Guérin

Improving solar-receiver efficiency by manipulating turbulence in heat-collecting, fluid-conveying channels: from physical understanding to machine-learning-based optimization

Directeur(s) de Thèse :
Laurent Cordier, Lionel Agostini, Stéphane Grieu

Soutenue le 09 décembre 2025 devant le jury

Jury :

Président	**Michel Bergmann**	Directeur de recherche INRIA, Université de Bordeaux
Rapporteur	**Elie Hachem**	Professeur, CEMEF, Mines ParisTech PSL, Sophia Antipolis
Rapporteur	**Taraneh Sayadi**	Professeure, M2N, Conservatoire National des Arts et Métiers, Paris
Membre	**Cédric Flageul**	Maître de conférences, Institut Pprime, Université de Poitiers
Membre	**Lionel Mathelin**	Chargé de recherche CNRS, LISN, Université Paris-Saclay
Membre	**Adrien Toutant**	Maître de conférences HDR, PROMES, Université de Perpignan Via Domitia

Pour citer cette thèse :
Lou Guérin. *Improving solar-receiver efficiency by manipulating turbulence in heat-collecting, fluid-conveying channels: from physical understanding to machine-learning-based optimization* [En ligne]. Thèse Mécanique des fluides. Poitiers : Université de Poitiers, 2025. Disponible sur Internet <http://theses.univ-poitiers.fr>

UNIVERSITY OF POITIERS

DOCTORAL SCHOOL MIMME (ED 651)

Mathématiques, Informatique, Matériaux, Mécanique, Energétique

P H D T H E S I S

to obtain the title of

Doctor

of the University of Poitiers

Specialty : Fluid Mechanics

Defended by

Lou Guérin

Improving solar-receiver efficiency by manipulating turbulence in heat-collecting, fluid-conveying channels: *From physical understanding to Machine-Learning-based optimization*

Thesis Advisors: Laurent Cordier, Lionel Agostini, Stéphane Grieu

prepared at Pprime Institute, Curiosity Team

defended on December 9, 2025

Jury :

Reviewers :	Taraneh Sayadi	Professor	M2N, CNAM, Paris
	Elie Hachem	Professor	CEMEF, Mines Paris PSL, Sophia Antipolis
Examinators :	Lionel Agostini	CNRS Researcher	Pprime Institute, Univ. Poitiers
	Michel Bergmann	INRIA Research Director	MEMPHIS, Univ. Bordeaux
	Laurent Cordier	CNRS Research Director	Pprime Institute, Univ. Poitiers
	Cédric Flageul	Assistant Professor	Pprime Institute, Univ. Poitiers
	Lionel Mathelin	CNRS Researcher	LISN, Univ. Paris-Saclay
	Adrien Toutant	Assistant Professor	PROMES, Univ. Perpignan via Domitia

Acknowledgements

Je tiens tout d'abord à exprimer ma profonde gratitude à Taraneh Sayadi et Elie Hachem pour avoir accepté d'être les rapporteurs de cette thèse, malgré leurs emplois du temps très chargés. Je mesure la chance d'avoir pu bénéficier de leur expertise et de leur temps précieux. Je tiens à remercier également les examinateurs, Lionel Mathelin, Michel Bergmann, Adrien Toutant et Cédric Flageul, pour la qualité des discussions et des remarques lors de ma soutenance. Leur attention rigoureuse et leurs retours éclairés ont été une réelle source d'enrichissement pour ce travail. Un merci tout particulier va à Adrien Toutant pour son écoute attentive et sa bienveillance tout au long des CSI.

Je souhaite ensuite témoigner toute ma reconnaissance aux personnes qui ont encadré cette thèse: Laurent Cordier, Lionel Agostini et Stéphane Grieu. Votre accompagnement m'a permis d'apprendre énormément, et j'ai particulièrement apprécié l'équilibre entre l'autonomie que vous m'avez accordée et votre disponibilité dès que j'en avais besoin. Laurent, un grand merci pour ta disponibilité et ta curiosité intellectuelle, qui ont stimulé ma réflexion et m'ont incité à approfondir certains sujets avec plus de rigueur. Lionel, je te suis extrêmement reconnaissant pour ta bienveillance et ton approche décontractée, qui ont rendu ce parcours bien plus agréable que je ne l'aurais imaginé au départ. Ton franc-parler m'aura aussi permis de mieux comprendre comment fonctionne le monde de la recherche. Merci aussi pour ta réactivité, tes conseils précieux, et la liberté que tu m'as laissée dans le choix de mes sujets de recherche. Je n'oublie pas Cédric Flageul, dont la contribution, bien qu'informelle, a été déterminante pour cette thèse. Merci pour ces nombreuses discussions dès ma première année et merci pour ton aide, notamment avec Xcompact3D et la thermo. Enfin, un grand merci à Rodriguo Vincente-Cruz de m'avoir guidé dans la prise en main du code et à Michael Leschziner pour ses précieux conseils sur la turbulence proche-paroi.

Pendant cette thèse, j'ai aussi eu la chance de découvrir l'enseignement, une expérience très enrichissante qui m'a souvent sorti la tête du guidon quand la recherche stagnait. Un grand merci donc à l'ENSMA et en particulier à Nassim Razaaly et Guillaume Lehnasch de m'avoir fait confiance et offert cette opportunité dans un cadre parfait pour mes premiers pas dans l'enseignement. Cette thèse a également été l'occasion pour moi d'encadrer mon premier stagiaire. Je tiens donc à remercier mes encadrants de m'avoir accordé cette opportunité si formatrice. Thomas, un grand merci pour ton implication importante dans ce projet. Cette thèse m'a également permis de vivre une expérience internationale enrichissante, avec deux mois passés au laboratoire KTH de Stockholm. Je remercie donc Ricardo Vinuesa pour son invitation, ainsi que l'EUR INTREE, l'université de Poitiers et l'ED MIMME pour leur soutien financier. Un grand merci aussi à Marcial, Roko, Rubens, Piero, Makoto et Roger pour les excellents moments partagés, que ce soit au laboratoire,

sur le tatami, ou sous le ciel du cercle polaire!

Bien que je n'aie pas passé autant de temps sur place que je l'aurais souhaité, chaque moment au laboratoire a été marqué par des rencontres qui ont rendu cette expérience particulièrement agréable. Je tiens à remercier chaleureusement toutes les personnes qui y ont contribué, et notamment mon co-bureau Niccolo, ainsi que les doctorants et post-docs de l'ENSMA pour leur accueil, leur bonne humeur, et bien sûr, ces nombreuses parties de baby-foot. Une pensée spéciale pour Tom et Tarun, avec qui j'ai partagé ma première année de thèse. Vos conseils et nos échanges m'ont vraiment ouvert les yeux sur les hauts (et les bas !) de la vie de doctorant. Enfin, un grand merci à Antoine pour ses conseils et son soutient en cette fin de thèse.

La vie à Poitiers ne se résume pas au laboratoire, merci donc à Geoffrey, Alex et Johan pour toutes les séances de sport, les soirées jeux, Spy ou Spaceforce. Merci aussi aux amis que j'ai pu rencontrer plus tôt dans mes études. En commençant par les anciens, Théo et Benoît, avec qui l'aventure a commencé l'avant-veille du brevet. Merci d'être toujours là après toutes ces années, et pour tous ces moments partagés à mes côtés. Benoît, on se sera beaucoup ratés de peu dans nos études-sup. J'ai beaucoup apprécié tes conseils à ce sujet, et ils ont été précieux, notamment lors de cette thèse. Merci aussi d'avoir trouvé l'offre de cette thèse et à peu près toutes mes offres de stages avant ça. Merci aussi à Valentin et Geoffrey, présents depuis la licence. Je suis très content d'être encore en contact avec vous. Geoffrey, merci pour tous ces fous rires et ces pranks. Valentin, merci d'être parti en licence avec moi : ce semestre a changé ma vie. On aura découvert la montagne ensemble et fait de super rencontres. Encore merci pour cette journée à Flåm, très chargée en avalanches.

Ce serait mentir de dire que le sport a juste été un outil pour décompresser pendant cette thèse. Un immense merci à Théo, qui m'a poussé à relever le défi d'un premier marathon, une aventure qui m'a ensuite ouvert les portes du trail. Merci surtout à mon frère Noé, pour s'être lancé à mes côtés dans ce marathon, puis dans notre premier trail. Ces deux courses restent parmi les moments les plus forts pour moi durant ces trois dernières années. Un très grand merci aussi à vous, Lucas et Teemo, de m'avoir suivi dans les défis de trail qui ont suivi avec autant de panache. On se souviendra des week-ends chocs à Saint-Jean avec des Mondarrain *3 à 12 % de batterie corporelle. Sincèrement, merci d'être des personnes pas prises de tête et qui ne reculent devant rien, ça a été un plaisir de courir avec vous et je ne doute pas qu'on aura d'autres occasions de se mettre dans le dur à l'avenir! Cette thèse aura aussi été l'occasion de traverser pas mal d'endroits à pattes, notamment la Corse, les îles Lofoten et le tour du Mont-Blanc. Teemo, merci de t'être lancé dans chacun de ces défis avec moi, c'étaient à chaque fois des moments très marquants, même si parfois un peu éprouvants. Merci d'avoir partagé de nombreuses galères pendant ces voyages : dormir sur une aire d'autoroute, traverser des crêtes sous des rafales de vent, ou encore se faire prendre en stop par des bateaux et des bus de retraités finlandais. Il faudra faire attention à l'avenir à ne pas se faire attraper par la police du kiff. Merci aussi à Noé, Alex et Bioj pour les nombreuses randos, on

se souviendra notamment du prank du feu, des sources chaudes, ou du hors-piste à Ulvik.

Je tiens bien évidemment à remercier Gaïa, qui m'a accompagnée sur une longue partie de cette thèse. Tu es sûrement celle qui a le plus subi mes moments difficiles, un grand merci donc pour ton soutien et ta patience. Merci pour tous ces moments passés ensemble : la cuisine, les nombreuses parties de Hanabi le mercredi, les balades dans le parc avec les chats, les sorties trails, ou encore les baignades dans de l'eau pas toujours très chaude. Merci de partager autant de passions avec moi, mais aussi de m'avoir appris plein de nouvelles choses (notamment le code couleur pour les éponges). Merci de m'avoir accompagnée en trek à de nombreuses reprises (bientôt le record du nombre de nuits en tente) et surtout en Norvège cette année. On se souviendra de l'interminable étape Torrehytten-Tyssevasbu en plein cagnard. Merci de supporter mes bad-moods et mes jeux de mots pourris depuis tout ce temps. Plus sérieusement, merci de me rendre aussi heureux, c'est sans aucun doute grâce à toi que cette seconde partie de thèse s'est aussi bien passée.

Enfin, je tiens à remercier ma famille. Merci de m'avoir laissé une grande liberté pendant toute mon enfance, et merci à vous trois de m'avoir inculqué des valeurs qui font de moi la personne que je suis aujourd'hui. Papa, merci de m'avoir éduqué avec une approche très relax, sans jamais me mettre trop la pression, mais aussi de m'avoir initié au sport (même si ça aura pris du temps). J'ai encore quelques années pour battre ton temps au marathon. Noé, merci pour tous ces moments complices en grandissant et de continuer à partager plein de moments avec moi aujourd'hui, malgré la distance. Maman, merci de m'avoir soutenu durant mes études et durant cette thèse, ainsi que pour tous ces moments passés en cohabitation pendant ces trois années. Merci aussi de m'avoir encouragé à faire des études. Peu importe où ça me mènera, ça a été des années très enrichissantes, et si c'était à refaire, je le referais sans hésiter.

Nomenclature

Fields, Coordinates and Symbols:

$\boldsymbol{x} = (x, y, z)$	Cartesian spatial coordinates
t	Time coordinate
$\Delta\boldsymbol{x}$, Δt	Spatial/temporal discretization scales ($\Delta\boldsymbol{x} = (\Delta x, \Delta y, \Delta z)$)
t_{collect}	Period used for statistics collection
$\boldsymbol{u} = (u, v, w)^\intercal$	Velocity field (streamwise, wall-normal, spanwise)
p	Pressure field
$\boldsymbol{\omega}$	Vorticity field ($\boldsymbol{\omega} = \nabla \times \boldsymbol{u}$)
ρ	Fluid density
μ	Dynamic viscosity
ν	Kinematic viscosity ($\nu = \mu/\rho$)
δ	Characteristic flow length
δ_ν	Viscous length-scale ($\delta_\nu = \nu/u_\tau$)
H	Channel half-height
U	Characteristic flow velocity
U_∞	Free-stream velocity
U_b	Bulk velocity ($U_b = \frac{1}{\lvert\Omega\rvert}\int_\Omega \langle u\rangle_t \ \mathrm{d}\Omega$)
a	Local speed of sound
τ_ν	Viscous time-scale ($\tau_\nu = \nu/u_\tau^2$)
τ_w	Wall shear stress ($\tau_w = \mu \cdot \partial\overline{u_{\text{w}}}/\partial y$)
u_τ	Friction velocity ($u_\tau = \sqrt{\tau_w/\rho}$)
κ	Von Kármán constant for velocity ($\kappa = 0.41$)
B	Velocity log-law intercept ($B = 5.0$)
$\boldsymbol{f}$	Momentum forcing term ($\boldsymbol{f} = \left(-\frac{\partial \overline{p}}{\partial x}, \quad 0, \quad 0\right)^\intercal$)
T, Θ	Temperature field (dimensional/nondimensional)
α	Thermal diffusivity
λ	Thermal conductivity
h	Convective heat transfer coefficient
Θ_τ	Friction temperature ($\Theta_\tau = \alpha/u_\tau \cdot \partial\overline{\Theta_{\text{w}}}/\partial y$)
T_b	Bulk temperature ($T_b = \frac{1}{\lvert\Omega\rvert}\int_\Omega \langle T\rangle_t \ \mathrm{d}\Omega$)
T_{w}	Wall temperature
q_w	Wall heat flux ($q_w = h(\langle T_{\text{w}}\rangle_t - T_b)$)
κ_Θ	Von Kármán constant for temperature ($\kappa_\Theta = 0.44$)
B_Θ	Thermal log-law intercept function
f_Θ	Thermal forcing term for global heat balance ($f_\Theta = Au$)
δ_α	Thermal/Conductive sublayer thickness

$R_{\alpha,\nu}$	Ratio of thermal/viscous sublayer thinning
$\boldsymbol{z}$	Latent variable or latent space vector
$\boldsymbol{\mu}$	Mean vector of the latent distribution (VAE)
$\boldsymbol{\sigma}$	Standard deviation vector of the latent distribution (VAE)
L_d	Latent space dimension
$p_{\boldsymbol{\theta}}(\boldsymbol{x})$	Underlying dataset distribution
$q_{\psi}(\boldsymbol{z}\|\boldsymbol{x})$	Approximate posterior distribution (encoder)
$p_{\boldsymbol{\theta}}(\boldsymbol{x}\|\boldsymbol{z})$	Likelihood distribution (decoder)
$p_{\boldsymbol{\theta}}(\boldsymbol{z})$	Prior distribution over latent variables ($p_{\boldsymbol{\theta}}(\boldsymbol{z}) = \mathcal{N}(\mathbf{0}, \boldsymbol{I})$)
r_t	Reward received at time step t
$\boldsymbol{s_t}$	State of the environment at time step t
$\boldsymbol{a_t}$	Action taken by the agent at time step t
Π_θ	DRL policy parameterized by θ
π_θ	SRL policy parameterized by θ
$\boldsymbol{\Sigma_t}$	PBO covariance matrix at time instant t
$\boldsymbol{\mu_t}$	PBO mean vector at time instant t
$CI_{95\%}$	95% confidence interval for optimized parameters

Dimensionless Numbers and Control Parameters:

Re	Reynolds number ($\mathrm{Re} = U\delta/\nu$)
Re_τ	Friction Reynolds number ($\mathrm{Re}_\tau = u_\tau\delta/\nu$)
Pr	Prandtl number ($\mathrm{Pr} = \nu/\alpha$)
Pe	Péclet number ($\mathrm{Pe} = \mathrm{Re}\,\mathrm{Pr}$)
Ma	Mach number ($\mathrm{Ma} = U/a$)
C_f	Skin friction coefficient ($C_f = 2\tau_w/\rho U^2$)
P_{sav}	Percentage of power saved through drag reduction
P_{req}	Percentage of power required for actuation
P_{net}	Net power saving ($P_{\mathrm{net}} = P_{\mathrm{sav}} - P_{\mathrm{req}}$)
Nu	Nusselt number ($\mathrm{Nu} = 2Hh/\lambda$)
St	Stanton number ($\mathrm{St} = \mathrm{Nu}/\mathrm{Pe}$)
Nu_r	Relative Nusselt number enhancement ($\mathrm{Nu}_r = \overline{\mathrm{Nu}_{\mathrm{ac}}}/\overline{\mathrm{Nu}_{\mathrm{unac}}}$)
A_n	Analogy factor ($A_n = (\mathrm{Nu}_{\mathrm{ac}}/\mathrm{Nu}_{\mathrm{unac}})/(C_{f,\mathrm{ac}}/C_{f,\mathrm{unac}})$)
w_{w}	Spanwise wall velocity signal
T	Oscillation period
W	Oscillation amplitude
k_x	Streamwise wavenumber
λ_x	Streamwise wavelength ($\lambda_x = 2\pi/k_x$)
c	Wave speed ($c = 1/Tk_x$)
φ	Oscillation phase
a_1, a_2	Waveform control parameters
$\{(t_i^1, y_i^1)\}_{i=1,\ldots,5}$	Control points coordinates for first half period of w_{w}
$\{(t_i^2, y_i^2)\}_{i=1,\ldots,5}$	Control points coordinates for second half period of w_{w}

S_{Stokes}	Stokes strain ($S_{Stokes} = \partial\widetilde{w}/\partial y$)
SSR	Stokes strain rate (SSR $= \partial^2\widetilde{w}/\partial y \partial t$)
$\ell_{0.01}$	Stokes layer protrusion height

Statistics, Scalings, and Operators

$\overline{\bullet}$	Time and homogeneous directions average ($\overline{\bullet} = \langle\bullet\rangle_{x,z,t}$)
$\bullet'$	Temporal fluctuations ($\bullet' = \bullet - \overline{\bullet}$)
$\widetilde{\bullet}$	Phase-average ($\widetilde{\bullet} = \langle\bullet\rangle_{x,z,\varphi}$)
$\bullet''$	Stochastic fluctuations ($\bullet'' = \bullet - \widetilde{\bullet}$)
$\widehat{\bullet}$	Periodic fluctuations ($\widehat{\bullet} = \widetilde{\bullet} - \overline{\bullet}$)
$\langle\bullet\rangle_t$	Time average
$\bullet^+_{\text{nom}}$	Quantity normalized in wall-units using unactuated flow scaling
$\bullet^+_{\text{ac}}$	Quantity normalized in wall-units using actuated flow scaling
$\bullet_{\text{unac}}$	Unactuated flow quantity
$\bullet_{\text{ac}}$	Actuated flow quantity
$\bullet_{\text{w}}$	Quantity at the wall
$\langle\bullet,\bullet\rangle$	Scalar product
$R^{x,z}_{\bullet}$	2D streamwise-spanwise velocity correlation
$\Phi_{\bullet\bullet}$	Power spectral density
$D_{\text{KL}}(\bullet,\bullet)$	Kullback-Leibler divergence
$E_r(\bullet)$	Per-field energy reconstruction metric
$\text{Mean}_c(\bullet)$	Mean Correlation metric for latent space orthogonality
$\text{MeanMax}_c(\bullet)$	Mean-Max Correlation metric for latent space orthogonality
$D_B(\bullet)$	Davies-Bouldin clustering indicator

Acronyms

CSP	Concentrated Solar Power
QSV	Quasi-Streamwise Vortices
DNS	Direct Numerical Simulation
LES	Large Eddy Simulation
RANS	Reynolds-Averaged Navier-Stokes
CFR	Constant Flow Rate
CBT	Constant Bulk Temperature
CFL	Courant-Friedrichs-Lewy number
SWO	Spanwise Wall Oscillations
DR	Drag Reduction
STW	Streamwise Traveling Waves
DHT	Dissimilar Heat Transfer ($A_n > 1$)
MBC	Mixed Boundary Conditions
UHG	Uniform Heat Generation boundary condition
IF	IsoFlux boundary condition

CTD	Constant Temperature Difference boundary condition
DHT	Dissimilar Heat Transfer
ML	Machine Learning
ROM	Reduced Order Model
CROM	Cluster-Based Reduced Order Model
POD	Proper Orthogonal Decomposition
PCA	Principal Component Analysis
AE	AutoEncoder
CAE	Convolutional AutoEncoder
β-VAE	β Variational Autoencoder
A-CAE	Augmented Convolutional AutoEncoder
MSE	Mean Squared Error
ELBO	Evidence Lower Bound
LSTM	Long Short-Term Memory network
ELU	Exponential Linear Unit activation function
RAdam	Rectified Adam optimization algorithm
RL	Reinforcement Learning
DRL	Deep Reinforcement Learning
SRL	Simplified Reinforcement Learning
PBO	Policy-Based Optimization
PPO	Proximal Policy Optimization
PSO	Particle Swarm Optimization
BO	Bayesian Optimization
CMA-ES	Covariance Matrix Adaptation Evolution Strategy
GAE	Generalized Advantage Estimation
KDE	Kernel Density Estimation

Contents

Acknowledgements **i**

Nomenclature **iv**

Contents **xi**

1 Introduction **1**

1.1 Context and Objectives 1

1.1.1 Industrial context: Concentrated Solar Power plants 1

1.1.2 Objectives 3

1.2 Turbulent flows 4

1.2.1 Wall-bounded turbulence 4

1.2.2 Heat transfer in channel-flows 12

1.2.3 Numerical simulations and modeling of turbulent flows 15

1.2.4 Spanwise wall oscillations 16

1.3 Reduced-order modeling 20

1.3.1 Identifying a more compact coordinate system 21

1.3.2 Dynamical Modeling in Latent Space 26

1.4 Deep Reinforcement Learning for control systems 27

1.4.1 Single-step reinforcement learning for open-loop control . . . 29

Chapter summary 30

2 Preferential enhancement of convective heat transfer over drag **31**

2.1 Introduction 32

2.2 Flow conditions and numerical procedure 34

2.2.1 Physical modelling 34

2.2.2 Non-dimensionalisation 35

2.2.3 Fluctuating quantities 37

2.2.4 Details of numerical simulations 37

2.3 Validation of uncontrolled numerical simulations 38

2.4 Control strategy: Spanwise wall oscillations 42

2.4.1 Numerical simulation 42

2.4.2 Results for optimal drag reduction . . . 44
2.5 Breaking the Reynolds Analogy: Dissimilar Enhancement of Convective Heat Transfer . . . 45
2.5.1 Dissimilar heat transfer . . . 45
2.5.2 Near-wall streak rotation and distortion under oscillatory conditions . . . 51
2.5.3 FIK identity component analysis of transport phenomena . . . 54
2.5.4 Actuation induced enhancement of momentum and thermal mixing . . . 57
2.6 Perspectives on mechanisms and correlations driving the Reynolds-Analogy breakdown . . . 58
2.7 Impact of Prandtl and Reynolds numbers . . . 59
2.8 Conclusion . . . 61
Chapter summary . . . 63

3 Autoencoder-based dimensionality reduction 67
3.1 Introduction . . . 68
3.2 Heated turbulent channel flow under spanwise wall oscillation . . . 69
3.2.1 Flow conditions . . . 69
3.2.2 Numerical methods . . . 70
3.2.3 Validation . . . 71
3.2.4 Datasets . . . 72
3.3 Models . . . 75
3.3.1 Proper Orthogonal Decomposition . . . 75
3.3.2 Convolutional AutoEncoder (CAE) . . . 76
3.3.3 β-Variational AutoEncoder (β-VAE) . . . 79
3.3.4 Metrics . . . 82
3.4 Comparing CAE with POD . . . 84
3.5 Balancing latent orthogonality and reconstruction accuracy with β-VAE 90
3.6 Separability of generative factors within the latent space . . . 95
3.7 Friction augmented CAE for enhanced latent physical coherence . . . 98
3.7.1 Improvements in Nusselt number and friction coefficient predictions . . . 100
3.7.2 Latent separability analysis . . . 102
3.8 Conclusions and perspectives . . . 105
Chapter summary . . . 107

4 Policy-based optimization of spanwise wall oscillations 111
4.1 Introduction . . . 112
4.2 Flow conditions and numerical simulation details . . . 115
4.2.1 Details of numerical simulations . . . 116
4.2.2 Validation of numerical simulations . . . 117
4.3 Control-law optimization using reinforcement Learning . . . 118
4.3.1 Reinforcement Learning . . . 118
4.3.2 Policy-based optimization . . . 121
4.3.3 PBO configuration . . . 127
4.4 Optimal drag reduction . . . 128
4.5 Optimal net energetic balance . . . 132
4.6 Conclusions & Perspectives . . . 136
Extended Results . . . 139
4.7 Application to heat-transfer enhancement . . . 139
4.7.1 PBO-LES configuration . . . 139
4.7.2 Results . . . 140
4.7.3 DNS validation and comparative analysis . . . 143
4.7.4 Conclusion . . . 145
4.8 Improving Heat transfer with minimum drag penality . . . 146
4.8.1 Results . . . 147
4.8.2 Investigating optimal DHT period at fixed amplitude . . . 149
4.8.3 DNS validation and comparative analysis . . . 150
4.8.4 Conclusion . . . 154
Chapter summary . . . 155

5 Optimizing the waveform of the spanwise wall signal 159
5.1 Introduction . . . 160
5.2 Methods . . . 161
5.2.1 Waveform parameterization . . . 161
5.2.2 PBO-LES configuration . . . 164
5.2.3 Validation of waveform parameterization . . . 165
5.3 Optimal signal for drag reduction and net energetic gain . . . 166
5.3.1 Drag reduction . . . 167
5.3.2 Net energetic gain . . . 171

5.3.3 Conclusion 174
5.4 Optimal signal for dissimilar heat transfer 175
5.4.1 DNS validation and analysis 178
5.4.2 Conclusion 179
Chapter summary 181

6 Conclusions and perspectives **185**
6.1 Conclusions 185
6.2 Perspectives 188

Appendices **191**

A Maximizing heat-transfer enhancement with control cost penalty **193**

B Injectivity of $\mathcal{L}$ **197**

Bibliography **199**

Abstract/Résumé **216**

Chapter 1

Introduction

Contents

1.1 Context and Objectives 1
1.1.1 Industrial context: Concentrated Solar Power plants 1
1.1.2 Objectives 3
1.2 Turbulent flows 4
1.2.1 Wall-bounded turbulence 4
1.2.2 Heat transfer in channel-flows 12
1.2.3 Numerical simulations and modeling of turbulent flows 15
1.2.4 Spanwise wall oscillations 16
1.3 Reduced-order modeling 20
1.3.1 Identifying a more compact coordinate system 21
1.3.2 Dynamical Modeling in Latent Space 26
1.4 Deep Reinforcement Learning for control systems 27
1.4.1 Single-step reinforcement learning for open-loop control . . . 29
Chapter summary 30

1.1 Context and Objectives

1.1.1 Industrial context: Concentrated Solar Power plants

The urgent need to mitigate climate change, as underscored by the *IPCC Sixth Assessment Report* (Lee and Romero, 2023), has made the reduction of CO_2 equivalent emissions a global imperative. The decarbonization of electricity generation represents a critical pathway, with renewables needing to displace fossil fuels that still accounted for approximately 61% of global electricity production in 2024 despite record renewable growth (Climate, 2025). This transition is essential to meet the Paris Agreement's goal of limiting global warming to 1.5°C above pre-industrial levels (Lee and Romero, 2023).

Among renewable technologies, solar energy has emerged as a key solution, with solar photovoltaic (PV) and wind power collectively contributing 15% of global electricity generation in 2024 (Enerdata, 2024). Solar PV alone represented about

12% of global electricity production, demonstrating remarkable versatility through rooftop installations, agricultural integration (agrivoltaics), and utility-scale deployments.

However, electricity storage with batteries is still expensive, which makes it difficult to deal with solar energy variability of days and weeks. Concentrated Solar Power (CSP) plants, on the other hand, convert sunlight into thermal energy that can be stored and then further converted to electricity, which is a unique advantage in the renewable energy landscape. Indeed, thermal energy storage presents both efficiency and environmental advantages over electrical storage in the PV case (Rodat and Thonig, 2024). Future developments may even enable solar energy conversion into storable solar fuels for extended storage periods and long-range transport. Despite this major storage advantage, CSP represented only 0.06% of global electricity generation in 2018 (Drawdown, 2025). As of 2023, Spain and the United States maintain the highest cumulative installed capacity for CSP globally, with China following closely behind. Collectively, these three nations account for approximately 70% of the world's total CSP capacity, demonstrating their dominant position in this sector. On the other hand, new projects are rising mainly in China and the United Arab Emirates (Rodat and Thonig, 2024).

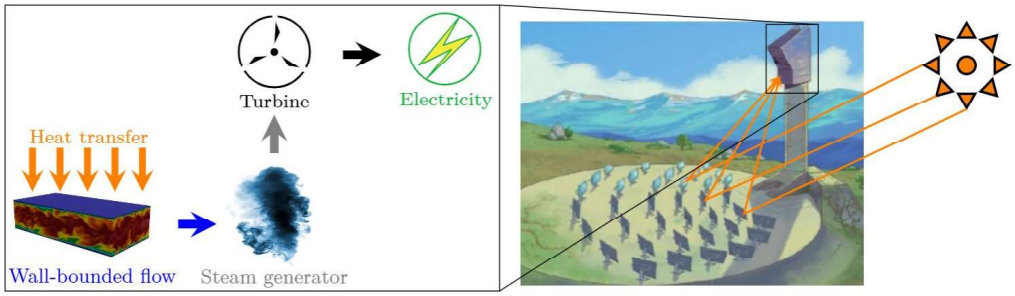

Figure 1.1: Simplified scheme of concentrated solar power plants.

As illustrated in Figure 1.1, this technology operates by concentrating sunlight via mirrors onto a receiver, containing heat transfer fluids (mainly oil, water and molten salts). This thermal energy can either drive turbines immediately (through a steam generator) or be stored for later use, addressing the intermittency challenges of other renewable energies (Rodat and Thonig, 2024). The molten salt circulates from a cold tank to the receiver, where it undergoes heating before being stored in a hot tank during daylight hours. Following sunset, the system utilizes the stored hot salt to generate steam for turbine operation, after which the cooled salt returns to the cold tank. Molten salts predominantly serve as sensible heat storage media, regardless of whether they also function as the heat transfer fluid in the system. More details regarding recent CSP implementations are provided in Ding and Bauer (2021).

CSP plants can be categorized into four main types based on their sunlight con-

centration methods: (i) Linear Fresnel reflectors, (ii) Parabolic troughs, (iii) Solar towers, and (iv) Parabolic dishes (see Rodat and Thonig 2024). Notably, for configurations (ii), (iii), and (iv), solar radiation reaches the receiver from multiple directions, potentially leading to symmetric heating. This contrasts with configuration (i), where radiation affects only one wall surface of the flow boundary, resulting in asymmetric heating. Current deployment statistics show parabolic troughs dominate the CSP market with approximately 73% share, followed by solar towers at 23% (Rodat and Thonig, 2024). Among these technologies, Fresnel systems demonstrate optimal land use efficiency, whereas tower and trough configurations provide the highest overall energy-conversion efficiency.

One of the most important element in these systems is the receiver, which concentrates solar radiation to heat a working fluid. Modern CSP systems typically use molten salt storage at 300-570°C (Rodat and Thonig, 2024), enabling up to 17.5 hours of storage capacity (at a 110 MW scale). Modern designs focus on achieving higher operating temperatures to boost energy output. This drives research into considering alternative heat-transfer working fluids. Indeed, a new generation of CSP plants aims to improve on molten-salt by using novel heat transfer fluids (Ding and Bauer, 2021) (for example liquid metals or pressurized air), which could reach higher temperatures, potentially above 1000°C (REN21, 2024).

Pressurized air offers particular advantages as it does not have an inherent temperature limit like molten salts. The maximum temperature is thus instead determined by the receiver materials, with ceramic components typically withstanding 900 K to 1300 K (David et al., 2023). Moreover, air avoids pollution risks associated with fluid leakage.

In addition to the adoption of pressurized air as a working fluid, the incorporation of flow control strategies may achieve additional efficiency gains in next-generation CSP plants. The current work focuses specifically on evaluating the potential of Spanwise Wall Oscillations (SWO) as a mechanism for thermal performance enhancement in wall-bounded flows.

1.1.2 Objectives

The primary objective of this thesis is to evaluate the potential of SWO to be utilized for heat-transfer enhancement in wall-bounded flows, while ensuring that any associated increase in drag is minimized. Particular attention is given to the implications for next-generation CSP plants, where the effects of the control strategy on pressurized air needs to be compared with those on molten salts.

First, in Chapter 2, numerical simulations of a heated turbulent channel flow with spanwise wall oscillations are conducted, with the aim of examining how the control method affects heat transfer performance and near-wall turbulence. A simplified numerical configuration is employed compared to the one used by David et al. (2023), which, although further removed from actual CSP configurations, is expected to

yield more fundamental insights into near-wall turbulence.

Second, in Chapter 3, an attempt to develop a Reduced-Order Model (ROM) is provided, with the objective of reproducing the predictive capability of numerical simulations at reduced computational cost. Special emphasis is placed on the use of non-linear AutoEncoder (AE) neural networks for flow field dimensionality reduction (Agostini and Leschziner, 2022).

Finally, in Chapter 4, the amplitude and period of the control signal are optimized to achieve two objectives: (i) reducing drag, and (ii) maximizing heat transfer enhancement while mitigating any associated increase in drag. Chapter 5 extends these results by introducing waveform optimization of the wall signal through polynomial interpolation. To this end, Deep Reinforcement Learning (DRL) methods will be employed, leveraging their demonstrated success in diverse flow control applications, as highlighted in Vignon et al. (2023); Garnier et al. (2019).

1.2 Turbulent flows

1.2.1 Wall-bounded turbulence

Laminar flows represent the most elementary fluid motion regime, characterized by ordered motion, viscous dominance, and deterministic behavior. Most industrial flows, however, are turbulent, meaning inertial forces dominate viscous forces. This state of fluid motion is characterized by:

- **Stochasticity**: Chaotic, three-dimensional velocity fluctuations generated by a wide range of scales, making precise prediction challenging. Unlike highly viscous fluids, where applying an opposite action can often restore the system to its initial state, deformations in stochastic systems are typically irreversible.
- **Diffusivity**: Enhanced mixing of momentum and mass.
- **Vorticity**: Intense rotational motions.
- **High dimensionality**: The dimension required to accurately describe turbulent flows increases significantly with the ratio between inertial and viscous forces, resulting in a vast number of degrees of freedom. This high dimensionality significantly complicates analysis, modeling, and prediction, requiring advanced computational approaches.
- **Non-linearity**: The development and characteristics of flow structures at different scales are deeply interconnected, with smaller and larger scales influencing one another. This mutual dependence leads to complex, non-linear dynamics that amplify the challenges in modeling and predicting turbulent behavior.

Despite apparent stochasticity, turbulent flows exhibit a certain coherence, expressed as statistically persistent patterns.

Turbulent flows in engineering systems are predominantly wall-bounded, with pipe, channel, and boundary layer flows representing canonical configurations of both theoretical and practical importance. These flows exhibit unique characteristics due to the presence of solid boundaries, where viscous effects dominate in the immediate wall vicinity. The Turbulent Boundary Layer (TBL), defined as the thin fluid region adjacent to a surface where velocity transitions from zero (no-slip condition) to the freestream value U_∞, plays a crucial role in determining both drag and heat transfer characteristics (Figure 1.2).

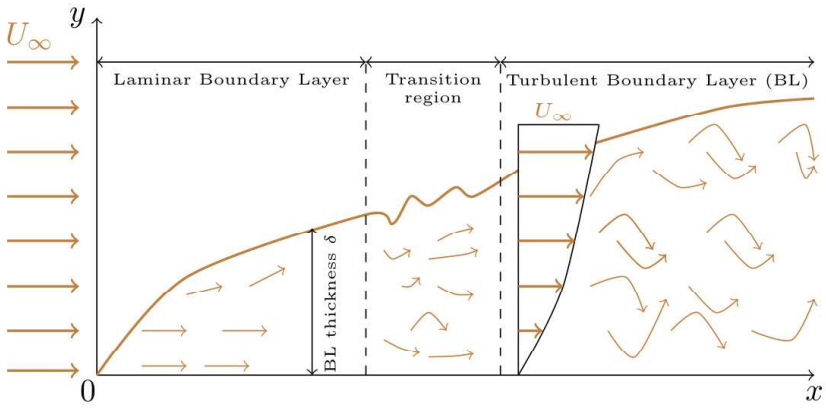

Figure 1.2: Schematic of the boundary layer development over a flat plate.

The shear layer, induced by the no-slip condition enforced by viscosity, drives the emergence of structures with distinct features:

- **Inhomogeneity**: Statistical properties vary with wall-normal distance.
- **Anisotropy**: Turbulent fluctuations become directionally dependent, with streamwise components typically exceeding wall-normal ones (Banerjee et al., 2007).
- **Coherent structures**: Quasi-organized motions govern near-wall momentum transfer dynamics (Jimenez, 2017).

The velocity field $\boldsymbol{u}(\boldsymbol{x} = (x, y, z), t)$ represents the kinematic state of a fluid, defined as:

$$\boldsymbol{u}(\boldsymbol{x}, t) = \begin{pmatrix} u \\ v \\ w \end{pmatrix}, \tag{1.1}$$

denoting the instantaneous velocity components in the x (streamwise), y (wall-normal), and z (spanwise) directions, respectively, at time instant t and position $\boldsymbol{x}$. For turbulent flows, flow fields can be decomposed into the following components:

$$\boldsymbol{u} = \overline{\boldsymbol{u}} + \boldsymbol{u}', \quad \overline{\boldsymbol{u}'} = 0, \tag{1.2}$$

where $\overline{\boldsymbol{u}}$ is the average (in time and homogeneous directions) velocity field and $\boldsymbol{u}'$ are the corresponding fluctuations.

The pressure field $p(\boldsymbol{x}, t)$ is dimensionally equivalent to both a force distributed over a surface and an energy distributed within a volume. In fluid and solid mechanics, pressure is often interpreted as the isotropic component of the stress tensor. For a Newtonian fluid, the stress tensor can be expressed simply in terms of the spatial derivatives of the velocity field.

The vorticity $\boldsymbol{\omega}$ is the field describing local rotation of the flow, defined as the curl of the velocity field:

$$\boldsymbol{\omega} = \nabla \times \boldsymbol{u} = \left(\omega_x = \frac{\partial w}{\partial y} - \frac{\partial v}{\partial z},\ \omega_y = \frac{\partial u}{\partial z} - \frac{\partial w}{\partial x},\ \omega_z = \frac{\partial v}{\partial x} - \frac{\partial u}{\partial y} \right)^{\mathsf{T}}. \tag{1.3}$$

Mathematically, the Navier-Stokes equations for incompressible flows describe the motion of viscous fluid substances, expressing conservation of mass and momentum:

$$\begin{aligned} &\text{Continuity:} \quad \nabla \cdot (\rho \boldsymbol{u}) = 0, \\ &\text{Momentum:} \quad \rho \left(\frac{\partial \boldsymbol{u}}{\partial t} + (\boldsymbol{u} \cdot \nabla)\, \boldsymbol{u} \right) = -\nabla p + \mu \Delta \boldsymbol{u} + \boldsymbol{f}, \end{aligned} \tag{1.4}$$

where μ is the constant viscosity of the fluid, ρ is its density, and $\boldsymbol{f}$ accounts for external forces.

The Mach number $\mathrm{Ma} = \frac{U}{a}$, where U is a characteristic speed of the flow (for example U_∞ for TBLs), and a represents the local speed of sound, quantifies flow compressibility effects. Incompressible flows assume constant density ($\rho = \text{const}$), which is valid for low Mach numbers ($\mathrm{Ma} \ll 1$), simplifying the continuity equation to the divergence free velocity condition: $\nabla \cdot \boldsymbol{u} = 0$. Compressible flows account for density variations ($\partial\rho/\partial t \neq 0$), essential at high Mach numbers ($\mathrm{Ma} \gtrsim 0.3$), requiring the full continuity equation $\partial\rho/\partial t + \nabla \cdot (\rho \boldsymbol{u}) = 0$ and coupling with energy/pressure via an equation of state.

A non-dimensionalization of the incompressible equations (as described in Section 2.2 for the channel flow) results in the the following:

$$\begin{cases} \nabla \cdot \boldsymbol{u} = 0 \\ \dfrac{\partial \boldsymbol{u}}{\partial t} + (\boldsymbol{u} \cdot \nabla)\, \boldsymbol{u} = -\nabla p + \dfrac{1}{\mathrm{Re}} \Delta \boldsymbol{u} + \boldsymbol{f}. \end{cases} \tag{1.5}$$

The flow regime is characterized by the Reynolds number, which is the ratio between inertial and viscous forces:

$$\mathrm{Re} = \frac{U\delta}{\nu}, \tag{1.6}$$

where $\nu = \frac{\mu}{\rho}$ is the kinematic viscosity, and δ represents a characteristic length (the boundary layer thickness in the case of a turbulent boundary layer). Laminar flows therefore correspond to low Reynolds number, while at high Reynolds numbers, flows tend to be turbulent.

At elevated Re, enhanced turbulent mixing improves momentum transfer but simultaneously increases skin-friction drag (Ricco et al., 2021), which can be described as the resistance that occurs when the fluid rubs against the wall. This drag is quantified by the friction coefficient:

$$C_f = \frac{\tau_w}{\frac{1}{2}\rho U^2}, \tag{1.7}$$

where this quantity directly results from the wall shear stress:

$$\tau_w = \mu \frac{\partial \overline{u_w}}{\partial y}, \tag{1.8}$$

where $\overline{u}$ is the averaged wall-parallel (streamwise) velocity, y is the wall-normal direction, and the subscript w denotes wall quantities (at $y = 0$ or $y = 2H$ for the channel-flow).

Scaling and Structure of Wall-Bounded Turbulence

The dominance of viscous forces near solid boundaries necessitates the use of wall-derived scaling parameters. The fundamental friction velocity u_τ is derived from the wall shear stress τ_w and fluid density ρ:

$$u_\tau = \sqrt{\frac{\tau_w}{\rho}}. \tag{1.9}$$

The associated viscous length scale δ_ν, representing the near-wall region where viscous effects prevail, is defined through:

$$\delta_\nu = \frac{\nu}{u_\tau}. \tag{1.10}$$

This leads to the characteristic viscous time scale:

$$\tau_\nu = \frac{\nu}{u_\tau^2}. \tag{1.11}$$

Quantities normalized by these scales are denoted in wall units with the '+' superscript, such as the wall-normal distance $y^+ = y/\delta_\nu$ and velocity $\boldsymbol{u}^+ = \boldsymbol{u}/u_\tau$.

A rigorous treatment of all scaling laws for near-wall turbulence is omitted here, with comprehensive discussions available in Kim (2012); Tang and Antonia (2023, 2022).

The friction Reynolds number characterizes the scale separation between global and viscous effects:

$$Re_\tau = \frac{u_\tau \delta}{\nu} = \frac{\delta}{\delta_\nu} = \delta^+. \tag{1.12}$$

This dimensionless parameter reflects the ratio between the length scale δ and the viscous scale δ_ν thus quantifying the separation between inner and outer scaling regimes.

The connection between viscous wall effects and turbulent dynamics organizes wall-bounded flows into four stratified regions:

- **Viscous sublayer** ($y^+ \lessapprox 10$): Viscosity-dominated region which has a strong influence over the flow structures and corresponds to a linear velocity profile $u^+ = y^+$, i.e. mean flow quantities scale with u_τ and δ_ν (Jimenez, 2017). It generates a large portion of skin friction.
- **Buffer layer** ($10 \lessapprox y^+ \lessapprox 30$): Transition zone where peak turbulence production occurs.
- **Logarithm layer** ($30 \lessapprox y^+ \lessapprox 0.2\,\mathrm{Re}_\tau$): Turbulence-dominated region with universal velocity profile $u^+ = \frac{1}{\kappa}\log(y^+) + B$ ($\kappa \approx 0.41$, $B \approx 5.0$) (Kim, 2012). Statistically, this zone exhibits momentum transfer mechanisms where large-scale motions in the outer flow (responsible for production) progressively transfer momentum toward the small scales in the wall-region (responsible for dissipation), mirroring the turbulent energy cascade that directs kinetic energy (Jiménez, 2012).
- **Outer layer** ($0.2\,\mathrm{Re}_\tau \lessapprox y^+ \lessapprox \mathrm{Re}_\tau$): Region where the viscous effects are mostly negligible and the length scale is delimited by δ.

These flow regions are illustrated in Figure 1.3, where the characteristic velocity profiles (linear in the viscous sublayer ($\overline{u}^+ = y^+$) and logarithmic in the logarithm layer ($\overline{u}^+ = \frac{1}{\kappa}\log(y^+)+B$)) exhibit exact correspondence with the theoretical scaling laws described above.

Coherent Structures and Momentum Transport Mechanisms

Despite stochasticity, turbulent flows are characterized by coherent structures, which are organized recurring patterns present in flow motions. These structures are dominant in the transport of momentum and energy processes (Jimenez, 2017). They can be separated into different categories, including:

- **Streaks**: Elongated low/high speed velocity regions located at the edge of the viscous sublayer and oriented in the streamwise direction. Their role is very

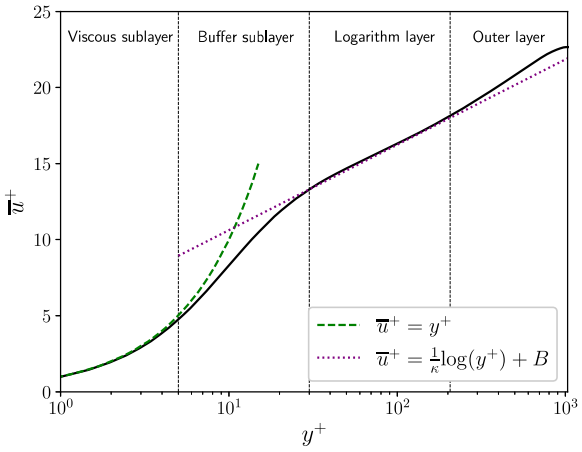

Figure 1.3: Wall normal variation of streamwise mean velocity profile $\overline{u}^+$ for a Direct Numerical Simulation (DNS) of a turbulent channel flow at $\mathrm{Re}_\tau = 1020$ (see Section 2.7 for more details). Each region of the flow is highlighted. Green dashed line corresponds to $\overline{u}^+ = y^+$ while the purple dotted line illustrates the $\overline{u}^+ = \frac{1}{\kappa}\log(y^+) + B$ curve.

significant relative to momentum transport in the near-wall region (scale with δ_ν). These streaks are combined to **Quasi-Streamwise Vortices (QSV)**, near-wall vortices roughly aligned with the flow.

- **Super streaks**: Large-scale streaks appearing in the outer layer with increasing Re_τ (scale with δ).

- **Attached eddies** : Self-similar eddies populating the flow region between streaks and super streaks.

The evolution of structural topology in wall-bounded flows as the Reynolds number increases is pictured in Figure 1.4. The streaks are the first coherent structures to appear at the edge of the viscous sublayer as turbulence develops. With rising Reynolds number, the thickness of the viscous sublayer (δ_ν) decreases significantly compared to the boundary layer thickness (δ). This leads to viscous effects being concentrated in a rapidly shrinking zone near the wall surface, while turbulent shear stress ($\overline{u'v'}$) becomes more dominant. Furthermore, under elevated Reynolds number conditions, the position of the streaks migrates toward the wall, and their characteristic length scales contract substantially relative to the domain size (Agostini and Leschziner, 2019). These characteristic length scales are $\lambda_x^+ \approx 1000$, and $\lambda_z^+ \approx 100$ in the streamwise and spanwise direction respectively. Concurrently, with increasing Re, large scale structures appear in the outer region (super streaks), as well as wall-attached eddies (Agostini and Leschziner, 2019) which appear in the region delimited by δ and δ_ν. As the large scale motions become progressively stronger, their influence on the near-wall structures is enhanced. This effect manifests through two distinct mechanisms: (i) the immediate modification of skin friction via the so-called "footprint" phenomenon, and (ii) the alteration of the streaks' features through an

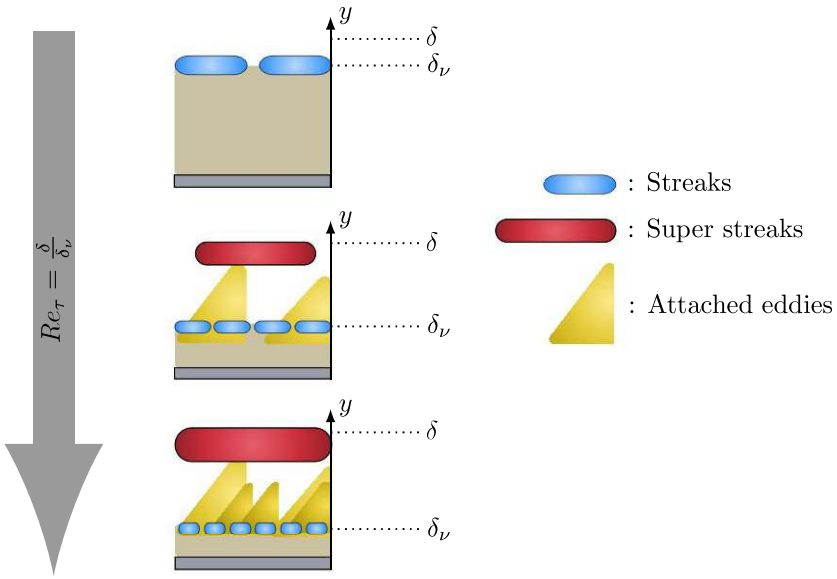

Figure 1.4: Illustration of coherent structure variations as Re_τ increases.

indirect process referred to as "modulation" (Mathis et al., 2011). Despite extensive research, comprehension of these structures, their interactions, and their impact on momentum transport remains not fully understood.

The production of turbulent energy is expressed as:

$$P = -\overline{u'v'}\frac{\partial \overline{u}}{\partial y}. \tag{1.13}$$

The turbulent shear stress $\overline{u'v'}$ is closely associated to momentum transport, thus to the friction drag. It can be decomposed into four quadrants based on the signs of u' and v' (Wallace, 2016):

- **Q1 (Outward interaction)**: $u' > 0$, $v' > 0$: High-speed fluid moving away from the wall.
- **Q2 (Ejection)**: $u' < 0$, $v' > 0$: Low-speed fluid ejected outward. These events are a dominant contributor to $\overline{u'v'}$ and friction drag.
- **Q3 (Inward interaction)**: $u' < 0$, $v' < 0$: Low-speed fluid moving toward the wall.
- **Q4 (Sweep)**: $u' > 0$, $v' < 0$: High-speed fluid sweeping toward the wall (also significant for momentum transport).

The QSVs are closely associated with sweep and ejection events, while inward

and outward interactions are the dominant quadrants for $\overline{\omega_x'\omega_x'}$ and $\overline{\omega_y'\omega_y'}$ (Wallace, 2016).

1.2.1.1 Turbulent channel-flow

Channel flows represent one of the three canonical wall-bounded flow configurations, offering a physically consistent analog for solar receiver geometries (David et al., 2023).

Although all canonical wall-bounded flows are known to share many similarities, particularly in terms of wall-region characteristics, the channel flow is distinguished by several unique features.

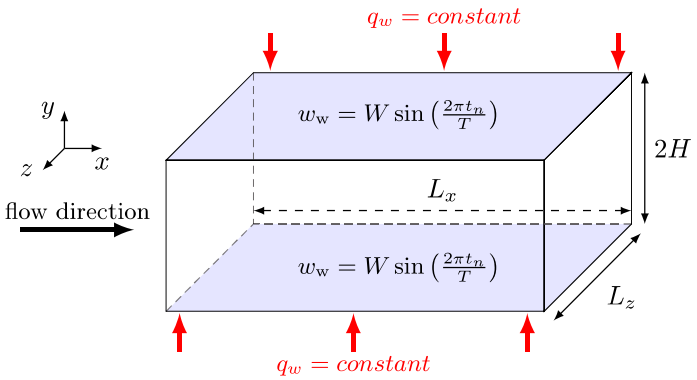

Figure 1.5: Schematic representation of the channel-flow configuration.

The channel domain $\Omega \subset \mathbb{R}^3$ is defined by:

$$x \in [0, L_x], \quad y \in [0, 2H], \quad z \in [0, L_z], \tag{1.14}$$

where H is the characteristic length δ of the channel flow, corresponding to its half height. The lengths of the domain in the streamwise and spanwise directions are referred to as L_x and L_z respectively (see Figure 1.5).

In the current study, the reference velocity U required for the definition of C_f (Eq. (1.7)) and Re (Eq. (1.6)) is established as the bulk velocity:

$$U_b = \frac{1}{|\Omega|} \int_\Omega \langle u \rangle_t \ \mathrm{d}\Omega, \tag{1.15}$$

where $|\Omega|$ is the fluid inner volume of the channel and $\langle \bullet \rangle_t$ the time average. The mean wall-normal momentum conservation in channel flows reveals a unidirectional and constant pressure gradient force:

$$\left(-\frac{\partial \overline{p}}{\partial x}, \quad 0, \quad 0\right)^\intercal, \quad \text{with} \quad \frac{\partial \overline{p}}{\partial x} = \text{const.} \tag{1.16}$$

This implies a linear mean pressure drop along the channel. The streamwise momentum balance shows an equilibrium between pressure and viscous effects:

$$\tau_w = -H\frac{\partial \overline{p}}{\partial x}. \tag{1.17}$$

These friction losses are balanced by the forcing term of Eq. (1.5) $\boldsymbol{f} = \left(-\frac{\partial \overline{p}}{\partial x},\ 0,\ 0\right)^{\mathsf{T}}$ representing, for example, a pump's action to maintain the flow's statistically steady state. The relationship between the forcing term $\boldsymbol{f}$ and wall shear stress τ_w enables direct computation of the skin friction coefficient C_f without requiring explicit wall-shear stress evaluation.

There are three main strategies utilized to address the forcing:

- **CPG (Constant Pressure Gradient)**: A fixed mean pressure gradient $-\frac{\partial \overline{p}}{\partial x} = \text{const}$ is imposed, driving the flow.
- **CFR (Constant Flow Rate)**: The bulk velocity U_b is held constant. This mimics experimental setups with fixed flow rate.
- **CPI (Constant Power Input)**: The total energy input (power) is kept constant, allowing both pressure gradient and flow rate to vary. This approach is particularly relevant for industrial applications where energy efficiency is constrained.

The present numerical implementation employs the CFR strategy by enforcing a fixed bulk velocity $U_b = \frac{2}{3}$ at every time step. Periodic boundary conditions are implemented in both the streamwise (x) and spanwise (z) directions, facilitating statistical averaging of flow quantities along these homogeneous dimensions. In the present configuration, this allows the mean operator to be defined as $\overline{\bullet} = \langle \bullet \rangle_{x,z,t}$, combining spatial and temporal averaging. For the wall-normal direction (y), the baseline configuration adopts no-slip boundary conditions ($\boldsymbol{u} = \boldsymbol{0}$) at both top and bottom walls ($y = 0$ and $y = 2H$).

1.2.2 Heat transfer in channel-flows

Turbulent flows, and channel flows in particular, have been thoroughly described in the previous sections. However, the behavior of the temperature field in such flows remains to be addressed.

The temperature, typically referred to as a scalar quantity, is governed by a passive equation when its effects on momentum are neglected. When treated as an active scalar (more representative of real flows), these effects are taken into account, for example accounting for buoyancy forces using the Boussinesq approximation (Sebilleau et al., 2018).

In the present study, these coupling effects are neglected, resulting in the following temperature equation, which corresponds to the passive transport of a scalar T by the velocity $\boldsymbol{u}$:

$$\frac{\partial T}{\partial t} + \boldsymbol{u} \cdot \nabla T = \alpha \Delta T, \tag{1.18}$$

where T is the temperature field and α is the thermal diffusivity.

Mixed boundary conditions

Numerous thermal boundary condition options exist, each representing different physical approximations. This includes, Uniform Heat Generation (UHG), IsoFlux (IF), Constant Temperature Difference (CTD), heat exchange coefficient (Robin), and Mixed Boundary Conditions (MBC) (see Flageul et al., 2015; Kasagi et al., 2011; Cruz, 2021; Alcántara Ávila, 2021).

The present study implements Mixed Boundary Conditions (Kasagi et al., 1992), which combine two assumptions: (i) constant heat flux at both walls, and (ii) zero temperature fluctuations at walls ($T_\mathrm{w} = \langle T_\mathrm{w} \rangle_t$).

Section 2.2.1 details the nondimensionalization procedure and global heat balance, yielding the nondimensional temperature equation:

$$\frac{\partial \Theta}{\partial t} + \boldsymbol{u} \cdot \nabla \Theta + f_\Theta = \frac{1}{\mathrm{Re}\,\mathrm{Pr}} \Delta \Theta, \tag{1.19}$$

where $f_\Theta = Au$ maintains global heat balance (see Section 2.2.1). $\mathrm{Pr} = \nu/\alpha$ represents the Prandtl number, i.e. the ratio of momentum diffusivity to thermal diffusivity which varies with each fluid. The Péclet number $\mathrm{Pe} = \mathrm{Re}\,\mathrm{Pr}$ represents the ratio of convective to conductive heat transport. For $\mathrm{Pe} > 100$, axial heat conduction becomes negligible compared to convective transport (Cruz, 2021), an assumption satisfied in all present simulations.

MBC results in a simplified implementation and a strong enforcement of the Reynolds analogy, which establishes a direct coupling between momentum and thermal transport. This coupling complicates the design of control strategies, as efforts to enhance heat transfer inevitably produce a comparable effect on drag, thereby raising the difficulty of achieving an optimal trade-off compared to configurations where the Reynolds analogy is weaker (Kasagi et al., 2011). Consequently, MBC provides a rigorous and demanding benchmark for assessing the performance of heat transfer control methods.

Reynolds analogy

The Nusselt number characterizes the wall temperature gradient thus the heat transfer and is defined as:

$$\mathrm{Nu} = \frac{2Hh}{\lambda}, \tag{1.20}$$

where h denotes the convective heat-transfer coefficient and λ the thermal conductivity. Nu can be directly estimated from the forcing term f_Θ, avoiding wall-normal derivative evaluation.

The Reynolds analogy establishes a fundamental relationship between momentum and heat transfer in turbulent flows. Mathematically, it expresses the similarity between the dimensionless temperature Θ^+ and streamwise velocity u^+ profiles through the Stanton number St (which is proportional to the Nusselt number):

$$\mathrm{St} \equiv \frac{\mathrm{Nu}}{\mathrm{Re}\,\mathrm{Pr}} = \frac{C_f}{2}. \tag{1.21}$$

Since the Reynolds analogy inherently couples heat and momentum transfer, any control strategy aiming to preferentially enhance the Nusselt number (Nu) over the skin friction coefficient (C_f) must actively decouple these transport mechanisms.

A comprehensive discussion of Reynolds analogies and their modifications for variable fluids property can be found in Mahulikar and Herwig (2008).

Thermal Wall Scaling and Log-Law Behavior

The friction temperature scale Θ_τ, defined as:

$$\Theta_\tau = \frac{\alpha}{u_\tau}\frac{\partial\overline{\Theta_\mathrm{w}}}{\partial y}, \tag{1.22}$$

characterizes near-wall thermal gradients and serves as the fundamental scaling parameter for thermal wall units, analogous to the friction velocity u_τ for velocity scaling. Within the thermal (or conductive) sublayer, which is the thermal analog of the viscous sublayer, the dimensionless temperature profile follows the linear relationship:

$$\Theta^+ = Pr\, y^+. \tag{1.23}$$

For $\mathrm{Pr} = 1$, thermal and viscous sublayer lengths become comparable. In this case, the temperature equation (Eq. (1.19)) closely resembles the streamwise momentum equation (Eq. (1.5)) (excluding pressure and forcing terms) which favours Reynolds analogy in the baseline case.

The thermal log-law region exhibits analogous behavior to the momentum log-law, but with Prandtl number modifications. For moderate Reynolds numbers, the temperature profile follows:

$$\Theta^+ = \frac{1}{\kappa_\Theta}\log(y^+) + B_\Theta(\mathrm{Pr}), \tag{1.24}$$

where $\kappa_\Theta \approx 0.44$ and $B_\Theta(\mathrm{Pr})$ is the thermal intercept that varies with Prandtl number (Alcántara Ávila, 2021).

1.2.3 Numerical simulations and modeling of turbulent flows

The incompressible Navier-Stokes equations (Eq. (1.5)) serve as a highly effective framework for modeling fluid dynamics, offering accurate predictions of turbulent behavior across all relevant spatial and temporal scales. Nevertheless, the model presents significant challenges due to its mathematical complexity. Current analytical methods yield exact solutions only for laminar scenarios.

Under appropriate conditions, however, turbulence can be treated as a deterministic process with intricate temporal evolution. This perspective forms the foundation of Computational Fluid Dynamics (CFD), where the continuous partial differential equations governing fluid motion are approximated using discrete spatial and temporal representations (Cruz, 2021). Through this discretization process, time and space are divided into finite elements with characteristic dimensions $\Delta \boldsymbol{x}$ and Δt. The original continuous differential operators are then replaced by discrete counterparts, implemented through various numerical methods (including spectral approaches, finite-difference schemes, finite volume techniques, and finite element analysis) to simulate the physical phenomena described by the Navier-Stokes equations. Furthermore, given the demonstrated accuracy of these equations in representing fluid behavior, discrete time evolution of flow fields can be computed from varying initial and boundary conditions.

In CFD, turbulent flow simulations are typically approached through three principal strategies (Alcántara Ávila, 2021):

- **Reynolds-Averaged Navier-Stokes (RANS)**: Solves time-averaged equations with turbulence closure models. It is the most computationally efficient approach, but provides only statistical information about the flow fields. RANS are suitable for industrial applications where detailed turbulence data is not required or too expensive computationally.

- **Large Eddy Simulation (LES)**: Resolves large, energy-containing eddies directly while modeling small scales. This method offers a balance between the accuracy of DNS and computational efficiency of RANS. Explicit LES computes subgrid-scales stresses through a defined model, while implicit LES relies solely on numerical dissipation from discretization schemes to mimic unresolved turbulence effects.

  Direct Numerical Simulation (DNS): Solves the full Navier-Stokes equations without modeling, capturing all turbulent scales. This method provides complete physical accuracy but becomes prohibitively expensive at high Reynolds numbers. It is thus primarily used for fundamental research and model development.

The present study combines DNS and implicit LES, both conducted using the Xcompact3d Fortran framework (Bartholomew et al., 2020). The complete numerical configuration for the DNS, including spatial discretization schemes, temporal

integration methods, grid resolution, and heat transfer and friction coefficient calculations, is presented in Section 2.2. The corresponding implementation details for the implicit LES approach are documented in Section 3.2.2.

1.2.4 Spanwise wall oscillations

Fundamentals of streamwise traveling waves

The control of turbulent boundary layers through spanwise wall motions represents a significant advancement in flow control strategies, with streamwise traveling waves (STWs) (Quadrio et al., 2009) emerging as one of the most effective techniques for substantial Drag Reduction (DR). The fundamental formulation of STWs establishes an oscillatory spanwise wall velocity field described by:

$$w_{\mathrm{w,st}}(x,t) = W \sin\left(k_x x + \frac{2\pi t}{T}\right), \tag{1.25}$$

where W represents the forcing amplitude, k_x denotes the streamwise wavenumber, and T is the oscillation period. The subscript st explicitly denotes a spatio-temporal forcing. This forcing creates complex interactions with near-wall turbulence structures. The wave speed is defined as $c = \frac{1}{Tk_x}$. Comprehensive parameter space mapping has revealed that downstream traveling waves ($c > 0$) achieve peak drag reduction (up to 45%) when the dimensionless wave speed $c^+ = c/u_\tau$ is approximately 10-12. This optimal wave speed corresponds to the convection velocity of near-wall turbulent structures, suggesting that maximal drag reduction occurs when the forcing is tuned to the natural time scales of the turbulent boundary layer (Knoop et al., 2025).

Implementation strategies: spatial vs. temporal forcing

STWs can also be implemented through two distinct alternative forcing strategies: spatial forcing and temporal forcing. The temporal variant (Jung et al., 1992) (see Figure 1.5), known as Spanwise Wall Oscillations (SWO) is defined as:

$$w_{\mathrm{w}}(t) = W \sin\left(\frac{2\pi t}{T}\right), \tag{1.26}$$

whereas the spatial variant is defined as $w_{\mathrm{w,s}}(x) = W \sin(k_x x)$, in which the subscript s corresponds to a spatial forcing. Temporal forcing has historically received more extensive research attention (Touber and Leschziner, 2012; Agostini et al., 2015; Quadrio and Ricco, 2004; Choi and Clayton, 2001; Ricco et al., 2021). However, Viotti et al. (2009) demonstrated that both approaches share fundamental similarities in their drag reduction mechanisms, with spatial forcing exhibiting slightly superior performance, achieving 45% DR at $W^+ = 12$ and $Re_\tau = 200$. The optimal streamwise wavelength $\lambda_x = \frac{2\pi}{k_x}$ for DR is found at $\lambda_x^+ \approx 1000$, which

corresponds to $T^+ \approx 100$ under a convective transformation based on the nominal near-wall convection velocity. From an implementation perspective, spatial forcing offers notable practical advantages. Its simpler mechanical requirements facilitate realization through various active and passive devices. A compelling example is the experimental setup by Knoop et al. (2025), which achieved square-wave spatial forcing using alternating spanwise-moving belts, building upon their earlier methodology presented in Knoop et al. (2024). This contrasts with the more complex mechanisms typically required for temporal forcing. The broader experimental landscape for SWO studies has been significantly advanced through wind tunnel investigations, usually employing Particle Image Velocimetry (PIV) (Marusic et al., 2021; Chandran et al., 2023; Bird et al., 2018). These studies have collectively enhanced the understanding of both forcing modalities while establishing robust experimental protocols for drag reduction research.

The effectiveness of SWO is critically dependent on the careful selection of control parameters. The key dimensionless parameter for the temporal forcing is the oscillation period $T^+ = Tu_\tau^2/\nu$. Extensive numerical investigations have demonstrated that the optimal oscillation period for maximum drag reduction is approximately $T^+ \approx 100$ (Quadrio and Ricco, 2004), with this value showing remarkable consistency across a wide range of Reynolds numbers from $Re_\tau = 200$ up to $Re_\tau = 6\,000$ (Marusic et al., 2021). The Reynolds number dependence of drag reduction effectiveness follows an approximate scaling law of $DR \sim Re_\tau^{-0.2}$ (Gatti and Quadrio, 2013), reflecting the decreasing relative influence of viscous effects at higher Reynolds numbers and presenting a significant challenge for practical implementation in high Reynolds number applications. The drag reduction exhibits a monotonic increase with the forcing amplitude $W^+ = W/u_\tau$ up to an optimal threshold value (Quadrio and Ricco, 2004). Furthermore, maintaining equivalent drag reduction performance at elevated Reynolds numbers necessitates proportionally larger forcing amplitudes.

Vorticity dynamics and near-wall turbulence modulation

The actuation triggers an initial transient phase characterized by rapidly developing drag reduction effects (Figure 1.6). After a few oscillation cycles, the flow transitions to a statistically steady (here low drag) state where turbulence properties exhibit periodic modulation synchronized with the forcing cycle (Agostini et al., 2014). This periodicity enables phase resolved analysis of turbulence dynamics. Hereafter, the notation $\widetilde{\bullet}$ represents phase averaged quantities, with the complete mathematical definition provided in Section 2.2. This phase averaging operation isolates the periodic component of the flow variables synchronized with the actuation cycle.

SWO can be systematically analyzed through two fundamental kinematic quantities (Akhavan et al., 1993): the Stokes strain $S_{Stokes} = \partial\widetilde{w}/\partial y$ and its temporal gradient, the Stokes strain rate $SSR = \partial^2\widetilde{w}/\partial y\partial t$. These quantities govern the modification of near-wall turbulence and subsequent drag variations through several interconnected physical mechanisms. The instantaneous thickness of the Stokes layer

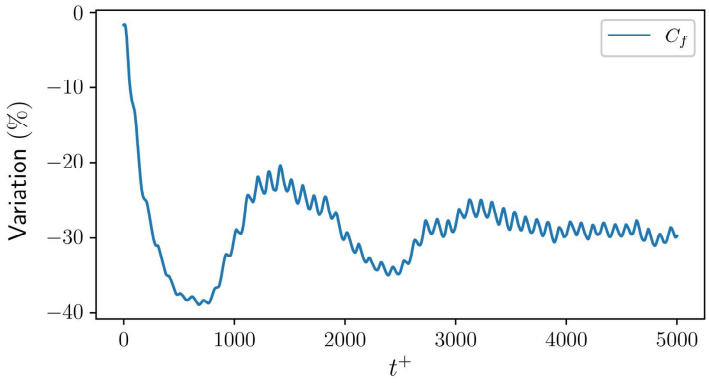

Figure 1.6: Temporal evolution of drag reductions obtained by spanwise oscillatory wall actuation for control parameters $T^+ = 200$ and $W^+ = 12$, at $\mathrm{Re}_\tau = 180$. C_f is the friction drag coefficient. χ^+ for any variable χ corresponds to a normalization based on the mean friction velocity of the unactuated flow, the so-called nominal configuration. Time is scaled in wall units. Results arise from the S_2 configuration of Table 2.1.

(the flow region dominated by the Stokes strain and its dynamics) varies throughout the oscillation cycle, reaching its maximum extent during phases of peak wall velocity and contracting during reversal phases. This dynamic behavior creates time-dependent modifications that are crucial for drag reduction. Rouhi et al. (2023) defined the Stokes layer protrusion height ($\ell^\star_{0.01}$), expressed in actuated flow wall units (where the superscript $\star$ denotes normalization by the actuated flow's friction velocity and viscosity), as the elevation at which $\widetilde{w^2}^\star$ decays to 1% of its peak value. This metric is specifically chosen because it exhibits near-proportionality with the maximum of the Stokes layer production term, which represents the sole external energy injection mechanism in the turbulent stress budgets. The approach differs fundamentally from conventional Stokes layer thickness definitions (Cimarelli et al., 2013) by directly linking the geometric length scale to the dominant production mechanism in actuated flows.

The spanwise wall motion induces a cascade of vorticity modifications, where changes in ω_y significantly alter ω_x and ω_z. Specifically, variations in ω_x directly govern drag fluctuations (Agostini et al., 2014). The periodic generation of streamwise vorticity $\omega_x \approx \partial w/\partial y$ through the Stokes strain leads to subsequent tilting and stretching of vortex lines, producing a net spanwise vorticity ω_z that modifies the mean velocity profile (Agostini et al., 2015). This process is particularly effective near $y^+ \approx 15$, where it creates (i) a reduction in mean shear near the wall, (ii) a corresponding increase in buffer layer momentum, and (iii) a weakening of quasi-streamwise vortices through vorticity cancellation. The effectiveness of these vorticity modifications is directly related to the Stokes strain rate (Agostini et al., 2015). High $\partial^2\widetilde{w}/\partial y\partial t$ magnitudes occur during phases of rapid strain variation and are particularly effective at disrupting near-wall turbulence. During these high-$\partial^2\widetilde{w}/\partial y\partial t$ phases, the ω_y vortices undergo tilting in the spanwise direction, thereby attenuating wall-normal vorticity while enhancing ω_z. As ω_y, alongside shear strain,

constitutes a fundamental component of the self-sustaining cycle of near-wall turbulence, its reduction disrupts the generation of streaks (Waleffe, 1997). Consequently, suppressing ω_y inhibits streak formation, diminishes momentum mixing, and ultimately reduces drag. Conversely, during low-$\partial^2\widetilde{w}/\partial y\partial t$ phases, termed "lingering" phases, the turbulence structures show partial recovery (Knoop et al., 2025), leading to the characteristic bimodal variation in drag reduction effectiveness observed throughout the oscillation cycle (see Figure 1.6).

Heat transfer considerations and dissimilar control strategies

While the majority of research on spanwise wall oscillations has focused on momentum transfer and drag reduction, the concomitant effects on heat transfer represent an equally important consideration for practical applications. The fundamental relationship between momentum and heat transfer in turbulent boundary layers is described by the Reynolds analogy, which suggests that modifications to friction drag should produce corresponding changes in heat transfer.

Through LES of weakly compressible flow (Mach = 0.5) at $\mathrm{Re}_\tau = 180$ and $\mathrm{Pr} = 0.72$, Fang et al. (2009) demonstrated persistent thermal-velocity streak alignment under spanwise oscillations ($6.35 \leq W^+ \leq 19.05$, $T^+ \approx 100$). Subsequent work (Fang et al., 2010) confirmed strong momentum and heat transport coupling, while Fang and Lu (2010) revealed transient drag amplification using in-phase actuation, with heat transfer consistently mirroring drag trends. These findings were extended to Mach = 2.9 by Ni et al. (2016), who developed a Stokes-corrected wall heat flux exhibiting drag like behavior. More recently, Rouhi et al. (2025) investigated SWO and STW effects at $\mathrm{Re}_\tau = 590$ and $W^+ = 12$, showing that for $\mathrm{Pr} > 1$, heat transfer reduction surpasses drag reduction for near-optimal periods (by 10% at $\mathrm{Pr} = 7.5$). Their analysis reveals that supra-optimal periods enhance Stokes layer protrusion, which, combined with thinner conductive sublayers at $\mathrm{Pr} > 1$, disproportionately amplifies $\Theta' v'$ fluctuations compared to $u'v'$ through increased production in the small turbulent scales. This mechanism leads to stronger drag reduction relative to heat transfer. Jung et al. (1992) demonstrated that drag amplification can be achieved using SWO with an oscillation period of $T^+ = 500$, which generates a more protrusive Stokes layer compared to optimal drag-reducing configurations. Under the Reynolds analogy, this suggests potential for concurrent heat transfer enhancement, though the possible decoupling of momentum and thermal transport mechanisms warrants further investigation.

More sophisticated control strategies have been developed to deliberately decouple heat and momentum transfer, achieving what is termed *dissimilar heat transfer* (Uchino et al., 2017). This approach aims to maximize heat transfer enhancement while minimizing frictional drag penalties, effectively breaking the Reynolds analogy. The dissimilarity is quantified by the analogy factor $A_n = (\mathrm{St}/\mathrm{St}_0)/(C_f/C_{f0}) = (\mathrm{Nu}/\mathrm{Nu}_0)/(C_f/C_{f0})$ (Uchino et al., 2017), with subscript 0 denoting unactuated values. Values of $A_n > 1$ indicate preferential enhancement of heat transfer rela-

tive to drag. Wall blowing and suction techniques have been shown to achieve up to 24% drag reduction while simultaneously increasing heat transfer by over 50% (Yamamoto et al., 2013). The choice of thermal boundary conditions significantly influences the potential for achieving dissimilar heat transfer, with comparative studies evaluating UHG, IF, and CTD boundary conditions (Kasagi et al., 2011). These analyses reveal that CTD boundary conditions combined with Prandtl numbers significantly different from unity provide the most favorable conditions for dissimilar heat transfer, while UHG conditions with $Pr = 1$ present the greatest challenge due to identical boundary conditions and forcing terms for streamwise velocity and temperature fields.

Despite significant progress in understanding spanwise wall oscillations, several important challenges remain. High Reynolds number scaling requires further investigation, particularly for $Re_\tau > 10\,000$, as the precise Reynolds number dependence of both drag reduction and heat transfer modification is not fully understood. Many studies focus on gross drag reduction without fully accounting for the energy required to sustain the wall motion (Quadrio and Ricco, 2004), making comprehensive energy balance analyses necessary to evaluate practical feasibility. The vast majority of studies consider idealized channel flow configurations, leaving the influence of complex geometries and pressure gradients insufficiently explored (although some investigations have extended these analyses to complex geometries, including transonic airfoils Quadrio et al. 2022). Additionally, strategies for simultaneously optimizing both drag reduction and heat transfer enhancement in practical systems remain underdeveloped, representing an important area for future research.

This work tackles these challenges through comprehensive numerical analysis, with three primary objectives: (i) quantifying heat-transfer modifications under controlled forcing conditions, (ii) assessing potential net energy efficiency gains across the parameter space, and (iii) evaluating the potential for achieving dissimilar heat transfer, which is a particularly relevant consideration for CSP applications. These investigations aim to provide a more complete understanding of the fundamental physics governing SWO while developing guidelines for their implementation in real-world applications.

1.3 Reduced-order modeling

As discussed in the preceding section, numerical simulations become increasingly computationally expensive at higher Reynolds numbers. To address this challenge, various approaches have been developed to reduce computational costs while maintaining reasonable accuracy. These methods fall under the broader framework of Reduced-Order Modeling (ROM), which approximates flow dynamics using a low dimension space.

Classical turbulence modeling approaches like RANS and LES inherently perform model reduction by significantly reducing the number of degrees of freedom

required to represent the flow compared to DNS. While these methods lower computational expenses, they inevitably introduce trade-offs in accuracy due to their inherent modeling assumptions.

An alternative approach leverages data-driven techniques to construct efficient ROMs (Csala et al., 2022). The implementation typically follows a two stage framework:

1. **Reduced-order space identification:** The first stage, detailed in Section 1.3.1, involves the extraction of a reduced-order coordinate system. This process identifies a low-dimensional space capable of capturing the dominant flow features while minimizing information loss.

2. **Dynamics Learning:** The second stage, presented in Section 1.3.2, consists of developing a predictive dynamical model within the reduced-order space. The model is trained to reproduce the essential flow dynamics at significantly reduced computational cost while maintaining fidelity to the full order system behavior.

1.3.1 Identifying a more compact coordinate system

Conventional data-driven dimensionality reduction approaches typically involve learning two transformations: (i) from the high-dimensional input space to a reduced-order latent space, and (ii) from the latent space back to the original space. By comparing the reconstructed data with the original inputs, this framework permits quantitative evaluation of the reduction quality.

Various techniques can be used to obtain such latent representations, including kernel principal component analysis (KPCA), locally linear embedding (LLE), Laplacian eigenmaps (LEM), isometric mapping (Isomap), multidimensional scaling (MDS), self-organizing maps (SOM), uniform manifold approximation and projection (UMAP), and curvilinear component analysis (CCA) (Csala et al., 2022).

The present study concentrates on two specific approaches: Proper Orthogonal Decomposition (POD) and non-linear AutoEncoder (AE) neural networks, which will be described in detail subsequently.

Proper Orthogonal Decomposition

POD (Berkooz et al., 1993), also known as Principal Component Analysis (PCA) in statistical contexts, is a mathematically rigorous dimensionality reduction technique that extracts dominant coherent structures from high-dimensional fluid flow data. The method constructs an optimal low-rank orthogonal basis that maximizes captured energy for any given basis dimension, making it particularly valuable for analyzing turbulent flows where organized motions coexist with random fluctuations (Agostini, 2020).

Given an ensemble of mean-centered flow snapshots $\{\boldsymbol{u}(\boldsymbol{x}_n, t_m)\}_{n=1,m=1}^{N,M} \in \mathbb{R}^{N\times M}$ arranged as columns of the data matrix $\boldsymbol{X} = [\boldsymbol{u}(\bullet, t_1), \cdots, \boldsymbol{u}(\bullet, t_M)] \in \mathcal{M}_{N\times M}(\mathbb{R})$, POD utilizes the correlation matrix of the input data:

$$\boldsymbol{C} = \boldsymbol{X}\boldsymbol{X}^\top \in \mathcal{M}_{N\times N}(\mathbb{R}), \quad \boldsymbol{C}\boldsymbol{U} = \boldsymbol{U}\boldsymbol{\Sigma}^2, \tag{1.27}$$

where the second equation represents the eigenvalue decomposition of a symmetric positive semi definite matrix, with $\boldsymbol{\Sigma}^2$ containing the positive eigenvalues and $\boldsymbol{U}$ being an unitary matrix ($\boldsymbol{U}\boldsymbol{U}^\top = \boldsymbol{U}^\top\boldsymbol{U} = \boldsymbol{I}$).

The same operation can be performed on the transpose of the correlation matrix (or temporal correlation matrix):

$$\boldsymbol{C}^\intercal = \boldsymbol{X}^\top\boldsymbol{X} \in \mathcal{M}_{M\times M}(\mathbb{R}), \quad \boldsymbol{C}^\intercal\boldsymbol{V} = \boldsymbol{V}\boldsymbol{\Sigma}^2, \tag{1.28}$$

which yields identical eigenvalues $\boldsymbol{\Sigma}^2$ through its eigenvalue decomposition.

The singular value decomposition of the input data matrix $\boldsymbol{X}$ naturally emerges from these eigenvalue decompositions:

$$\boldsymbol{X} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^\top = \sum_{i=1}^{\min(N,M)} \sigma_i \boldsymbol{\eta}_i \boldsymbol{\alpha}_i^\top, \tag{1.29}$$

where:

- $\boldsymbol{U} = [\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_N]$ contains the spatial POD modes, which form an orthonormal basis, as $\boldsymbol{U}$ is unitary ($\boldsymbol{\eta}_i \cdot \boldsymbol{\eta}_j = \delta_{ij}$). These spatial modes often correspond to coherent structures within turbulent flows.
- $\boldsymbol{V}^\intercal = \left[\boldsymbol{\alpha}_1, \cdots, \boldsymbol{\alpha}_M\right]^\intercal$ contains the temporal coefficients, which also form an orthonormal basis (as $\boldsymbol{V}$ is unitary). This means the coefficients are uncorrelated, i.e. they evolve independently and represent separate factors of variation within the dataset.
- $\boldsymbol{\Sigma} \in \mathcal{M}_{N\times M}(\mathbb{R})$ contains the singular values $(\sigma_1, \ldots, \sigma_{\min(N,M)})$, which are positive and allow a ranking of the modes with regards to their energy contributions σ_k^2.

Because the POD modes are hierarchically ordered by their energy contributions $\sigma_1^2 \geq \sigma_2^2 \geq \cdots \geq \sigma_{\min(N,M)}^2$, they enable optimal low rank approximations through modal truncation (Berkooz et al., 1993). The rank-r approximation of the input matrix is given by:

$$\boldsymbol{X}_r = \boldsymbol{U}_r\boldsymbol{\Sigma}_r\boldsymbol{V}_r^\top = \sum_{i=1}^{r} \sigma_i \boldsymbol{\eta}_i \boldsymbol{\alpha}_i^\top, \tag{1.30}$$

where $\boldsymbol{U}_r \in \mathcal{M}_{N\times r}(\mathbb{R})$ contains the first r spatial modes $\{\boldsymbol{\eta}_i\}_{i=1}^r$, $\boldsymbol{\Sigma}_r = \mathrm{diag}(\sigma_1,\ldots,\sigma_r) \in \mathcal{M}_{r\times r}(\mathbb{R})$, and $\boldsymbol{V}_r^\intercal \in \mathcal{M}_{r\times M}(\mathbb{R})$ contains the first r temporal coefficients $\{\boldsymbol{\alpha}_i\}_{i=1}^r$.

This truncated reconstruction satisfies the Eckart-Young theorem (Berkooz et al., 1993), providing the optimal rank-r approximation:

$$\boldsymbol{X_r} = \underset{\tilde{X},\ \mathrm{rank}(\tilde{\boldsymbol{X}})=r}{\arg\min} \left(\|\boldsymbol{X} - \tilde{\boldsymbol{X}}\|_2 \right). \tag{1.31}$$

Once the matrices $\boldsymbol{U}$ (spatial basis) and $\boldsymbol{V}$ (temporal coefficients) have been computed, any data sample $\boldsymbol{u}(\bullet, t_m)$ can be fully characterized by its corresponding temporal coefficients alone, as $\boldsymbol{u}(\bullet, t_m) \approx \sum_{i=1}^r \sigma_i \boldsymbol{\eta}_i \boldsymbol{\alpha}_i(m)$ and the spatial basis functions $\{\eta_i\}_{i=1,\ldots,r}$ are held constant. The latent representations of each sample are thus contained in the columns of $\boldsymbol{V}^\intercal$.

The cumulative energy captured by the truncated POD expansion is given by:

$$E_r = \frac{\sum_{k=1}^r \sigma_k^2}{\sum_{i=1}^{\min(N,M)} \sigma_i^2}, \tag{1.32}$$

where r is typically chosen such that $E_r \geq 1 - \varepsilon$ for some prescribed tolerance $\varepsilon \ll 1$.

The snapshot method (Hilberg et al., 1994) involves solving the reduced eigenvalue problem on the temporal correlation matrix $\boldsymbol{X}^\top \boldsymbol{X} \in \mathcal{M}_{M\times M}(\mathbb{R})$ rather than the spatial correlation matrix $\boldsymbol{X}\boldsymbol{X}^\top \in \mathcal{M}_{N\times N}(\mathbb{R})$, which is computationally advantageous when $M \ll N$. This avoids explicit computation of the spatial correlation matrix, as the POD modes can be recovered as:

$$\boldsymbol{U} = \boldsymbol{X}\boldsymbol{V}\boldsymbol{\Sigma}^{-1}. \tag{1.33}$$

When the spatial dimension is smaller than the number of snapshots ($N \ll M$), an alternative implementation computes the POD modes directly through eigenvalue decomposition of the spatial correlation matrix $\boldsymbol{C}$, avoiding the larger $M \times M$ eigenvalue problem of the snapshot method.

Autoencoder neural networks

A neural network defines a parametric mapping $f_\theta : \mathbb{R}^d \to \mathbb{R}^k$ through a series of layered transformations. For an L-layer network with input $\boldsymbol{x} \in \mathbb{R}^d$, the forward pass computes:

$$\boldsymbol{h}^{(l)} = \sigma\left(\boldsymbol{W}^{(l)}\boldsymbol{h}^{(l-1)} + \boldsymbol{b}^{(l)}\right), \quad l = 1,\ldots,L, \tag{1.34}$$

where:

- $\boldsymbol{W}^{(l)} \in \mathbb{R}^{n_l \times n_{l-1}}$ is the weight matrix at layer l.
- $\boldsymbol{b}^{(l)} \in \mathbb{R}^{n_l}$ is the bias vector at layer l.
- $\sigma(\bullet)$ denotes an element-wise activation function (e.g., ReLU, tanh, or sigmoid (Rasamoelina et al., 2020)). Non-linear activation functions result in a non-linear mapping.
- $\boldsymbol{h}^{(0)} = \boldsymbol{x}$ and $\boldsymbol{h}^{(L)} = f_\theta(\boldsymbol{x})$ represent the network's input and output, respectively.

In supervised learning, the network's performance is quantified through a loss function $\mathcal{L}(\theta, \boldsymbol{x})$, which represents the optimization objective and usually measures a discrepancy between predictions $\tilde{\boldsymbol{x}} = f_\theta(\boldsymbol{x})$ and target ground-truth outputs ($\boldsymbol{y}$). Common choices include:

- **Mean Squared Error (MSE)**: $\mathcal{L}(\theta, \boldsymbol{x}) = \frac{1}{M}\sum_{i=1}^{M} \|\tilde{\boldsymbol{x}}_i - \boldsymbol{y}_i\|_2^2$ for regression tasks (predictions in a continuous space).
- **Cross-Entropy**: $\mathcal{L}(\theta, \boldsymbol{x}) = -\sum_{i=1}^{M} \boldsymbol{y}_i \log \tilde{\boldsymbol{x}}_i$ for classification tasks (predictions in a discrete space).

The weights $\theta = \{\boldsymbol{W}^{(l)}, \boldsymbol{b}^{(l)}\}_{l=1}^{L}$ are optimized via gradient-based methods, e.g., gradient descent as in Eq. (1.35) or Adam (Kingma and Ba, 2014), which iteratively update θ to minimize $\mathcal{L}(\theta, \boldsymbol{x})$:

$$\theta_{t+1} = \theta_t - \eta \nabla_\theta \mathcal{L}(\theta_t, \boldsymbol{x}), \tag{1.35}$$

where η is the learning rate and $\nabla_\theta \mathcal{L}$ is computed via backpropagation (Rumelhart et al., 1986). Due to the non-convex nature of the optimization landscape, these methods are only guaranteed to converge to a local optimum rather than global. Non-linear activation functions enable learning complex mappings, but their derivatives critically affect gradient flow during optimization.

The accuracy and the generalization capability of the network are determined by three interconnected factors. First, the batch size controls how many samples are processed before each gradient update. Second, the number of epochs specifies how many complete passes are made through the entire dataset. Third, careful validation monitoring detects overfitting, which occurs when the model memorizes training patterns but fails to generalize to new data. The standard practice of dividing the dataset into training and validation subsets (usually a 80/20% split) provides an estimate of generalization error, allowing for early stopping when validation performance deteriorates. Additionally, maintaining a separate test dataset is often used to obtain an unbiased final evaluation of model performance.

Autoencoders (Bank et al., 2021) are specialized neural networks (see Figure 1.7) that learn compact latent representations through two components: (i) an encoder

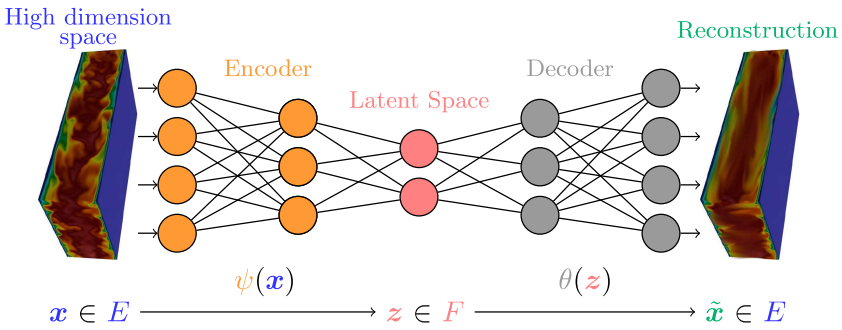

Figure 1.7: Autoencoder neural network representation.

$\psi(\bullet)$ that maps the input space in $\mathbb{R}^d$ to a latent space in $\mathbb{R}^{L_d}$ ($\psi(\boldsymbol{x}) = \boldsymbol{z}$), and (ii) a decoder $\theta(\bullet)$ that reconstructs the input from its latent representation ($\theta(\boldsymbol{z}) = \tilde{\boldsymbol{x}}$).

By selecting a latent dimension $L_d \ll d$, the architecture creates an information bottleneck that forces the network to learn compressed representations. This is achieved by optimizing the network's weights and biases to minimize the reconstruction error, typically measured using the MSE loss (Bank et al., 2021):

$$\mathcal{L}_{\text{MSE}}(\theta, \psi, \boldsymbol{X}) = \frac{1}{M} \sum_{i=1}^{M} \|\boldsymbol{x}_i - \theta(\psi(\boldsymbol{x}_i))\|_2^2, \tag{1.36}$$

where M is the number of training samples and $\boldsymbol{X} = [\boldsymbol{x}_1, \ldots, \boldsymbol{x}_M]$ is the training dataset (or equivalently $[\boldsymbol{u}(\bullet, t_1), \ldots, \boldsymbol{u}(\bullet, t_M)]$ using the POD notations).

When linear activations are used, the optimal autoencoder solution coincides with POD :

$$\tilde{\boldsymbol{X}} = \boldsymbol{X}_r, \quad \boldsymbol{Z} = \boldsymbol{V}^\intercal, \tag{1.37}$$

where $\tilde{\boldsymbol{X}}$ is composed of the AE reconstruction of dataset $\boldsymbol{X}$: $\tilde{\boldsymbol{X}} = [\tilde{\boldsymbol{x}}_1, \ldots, \tilde{\boldsymbol{x}}_M] = [\theta(\psi(\boldsymbol{x}_1)), \ldots, \theta(\psi(\boldsymbol{x}_M))]$, and $\boldsymbol{Z}$ contains the corresponding latent samples: $\boldsymbol{Z} = [\boldsymbol{z}_1, \ldots, \boldsymbol{z}_M] = [\psi(\boldsymbol{x}_1), \ldots, \psi(\boldsymbol{x}_M)]$. Non-linear activation functions allow autoencoders to learn complex, non-linear manifolds, giving them a theoretical advantage over the linear projections of POD for compressing non-linear phenomena like turbulent flows (Agostini, 2020).

However, this comes at the expense of certain properties inherent to POD. Notably, the latent variables in non-linear AEs are not guaranteed to be uncorrelated, meaning they do not necessarily represent independent factors of variation within the dataset. Furthermore, unlike in POD, the latent variables obtained from AEs are not inherently ordered by their energy contribution, which also reduces their interpretability (Kang et al., 2022).

To partially address these limitations, an algorithm proposed by Kang et al. (2022) ranks the latent variables according to their reconstruction contribution, thereby improving interpretability (see Section 3.3.3 for details). Additionally, β-Variational AEs (β-VAEs) (Eivazi et al., 2022; Solera-Rico et al., 2024) can mitigate

the loss of orthogonality by learning a latent distribution and constraining it to approximate a standard Gaussian. This regularization promotes decorrelation among latent variables, partially recovering orthogonality-like properties. Further discussion of β-VAEs is provided in Section 3.3.3.

Beyond β-VAEs, the autoencoder framework has evolved through numerous specialized architectures, which are proposed to enhance reconstruction accuracy or physical interpretability. Among these, convolutional AEs (Gruber et al., 2022) exploit spatial correlations, graph AEs (Gruber et al., 2022) preserve topological relationships in non-Euclidean domains, hierarchical AEs (Fukami et al., 2020) capture multi-scale dependencies, lift-augmented AEs (Fukami and Taira, 2023) enhance lift related latent information through an additional neural network, and constrained AEs (Mohan et al., 2023) enforce physical constraints (e.g., incompressibility) directly in the latent space.

1.3.2 Dynamical Modeling in Latent Space

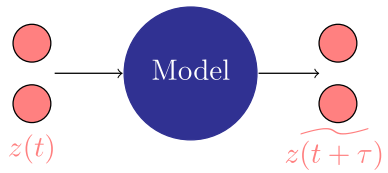

Figure 1.8: Latent dynamics prediction scheme.

Once a compact representation has been identified, the next challenge consists of constructing an efficient dynamical model within the latent space (see Figure 1.8). The main approaches for this task include:

- **Galerkin Projection for POD**: The classical approach projects the governing equations onto the POD basis, yielding a reduced-order system (Holmes et al., 2012). For a dynamical system $\frac{\partial \boldsymbol{u}}{\partial t} = \mathcal{F}(\boldsymbol{u})$, the Galerkin projection gives:

$$\frac{da_i}{dt} = \langle \boldsymbol{\eta}_i, \mathcal{F}\left(\sum_{j=1}^{r} a_j \boldsymbol{\eta}_j\right)\rangle, \quad i = 1, \ldots, r, \tag{1.38}$$

  where $\boldsymbol{\eta}_i$ are the orthonormal POD modes, a_i are the corresponding temporal coordinates, and $\langle \bullet, \bullet \rangle$ denotes the inner product.

- **CROM (Cluster-based Reduced-Order Modeling)**: A data-driven alternative that discretizes the latent space into clusters and models transitions via Markov chains (Kaiser et al., 2014). The dynamics are approximated as:

$$\boldsymbol{p}(t + \Delta t) = \boldsymbol{P}\boldsymbol{p}(t), \tag{1.39}$$

where $\boldsymbol{P}$ is the transition matrix between clusters and $\boldsymbol{p}$ the probability distribution. This approach is highly efficient to further understand the dynamics of the flow through cluster-based analysis (Xu and Tian, 2015).

- **Time Series Prediction**: Modern data-driven approaches employ neural networks, e.g., Long Short Term Memory networks (LSTMs) (Abadía-Heredia et al., 2022), transformers (Solera-Rico et al., 2024), or neural Ordinary Differential Equations (ODEs) (Linot and Graham, 2023), to learn latent dynamics directly:

$$\boldsymbol{z}_{t+1} = f_\theta(\boldsymbol{z}_t, ..., \boldsymbol{z}_{t-k}) + \varepsilon_t, \tag{1.40}$$

where f_θ is a learned mapping and ε_t represents modeling error.

Each method presents distinct trade-offs between physical interpretability (Galerkin), statistical consistency (CROM), and flexibility (time series prediction). The choice thus depends on the application requirements and available data quality.

1.4 Deep Reinforcement Learning for control systems

Control theory governs how systems achieve desired behaviors through input modulation. Control strategies can be passive, relying on inherent system dynamics, e.g., riblets (Abdulbari et al., 2013) or active, using external energy to enforce desired states, e.g., SWO (Agostini et al., 2014). Within active control, open-loop approaches execute predetermined inputs without feedback, e.g., a scheduled trajectory, while closed-loop methods dynamically adjust inputs based on real-time measurements, e.g., thermostats maintaining temperature via error correction. Open-loop methods prioritize stability but lack adaptability, whereas closed-loop control trades computational complexity for robustness and precision (Vignon et al., 2023).

Reinforcement Learning (RL) formalizes closed-loop control through discrete-time interactions between an agent and its environment (numerical simulations or ROMs in the present context). The sequential process unfolds as follows: At each timestep t, the agent observes the current flow state $\boldsymbol{s_t} \in \mathcal{S}$, e.g., 2D velocity fields at $y^+ \approx 15$, selects a control action $\boldsymbol{a_t} \in \mathcal{A}$, e.g., actuation amplitude adjustment, and receives a scalar reward $r_t(\boldsymbol{s_t}, \boldsymbol{a_t})$ quantifying immediate performance. The environment then transitions to a new state $\boldsymbol{s_{t+1}}$ according to the flow dynamics, and the cycle repeats. This iterative process, illustrated in Figure 1.9, continues until the agent reaches a terminal state, with the complete sequence of state-action transitions constituting an episode. Through these episodes, the agent learns a policy $\Pi_\theta(\boldsymbol{a}|\boldsymbol{s})$ that maximizes the expected discounted return $\mathbb{E}_{(\boldsymbol{s},\boldsymbol{a})}\left[\sum_{k=0}^{\infty} \gamma^k r_{t+k}\right]$, where $\gamma \in [0, 1]$ is the discount rate which governs tradeoffs between short-term and long-term control objectives.

A fundamental challenge in RL is the exploration-exploitation dilemma (Vignon et al., 2023). Exploration involves strategic randomization of actions to discover

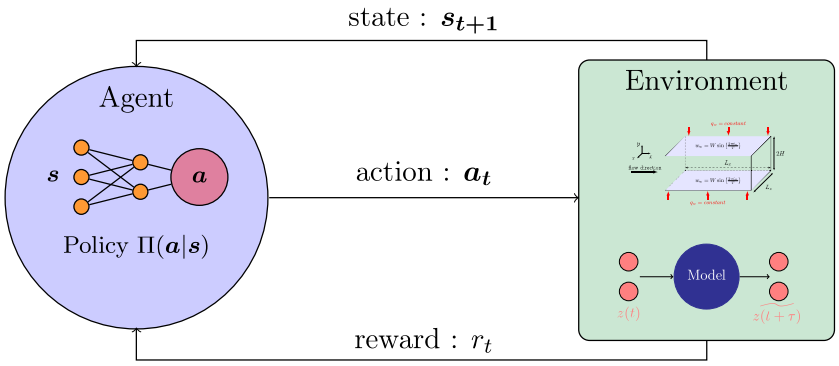

Figure 1.9: Deep Reinforcement Learning scheme.

high reward states, e.g., through ε-greedy policies where actions are random with probability ε or via noise injection in continuous domains, while exploitation leverages current knowledge to maximize immediate rewards. This tradeoff is particularly critical in physical systems where exhaustive exploration may be costly while greedy exploitation risks converging to suboptimal policies that miss globally superior control strategies.

For discrete, low-dimensional action spaces, value-based methods (Garnier et al., 2019) provide a foundational approach by recursively estimating the optimal state-value function $V(\boldsymbol{s})$ through the Bellman equation:

$$V(\boldsymbol{s_t}) = r_t + \gamma V(\boldsymbol{s_{t+1}}), \tag{1.41}$$

where the policy Π selects the action that maximizes expected returns to obtain $\boldsymbol{s_{t+1}}$.

Q-learning (Garnier et al., 2019) extends this concept by estimating the value of action-state pairs $Q(\boldsymbol{s}, \boldsymbol{a})$ through:

$$Q(\boldsymbol{s_t}, \boldsymbol{a_t}) = r_t + \gamma \max_a \left(Q(\boldsymbol{s_{t+1}}, \boldsymbol{a}) \right). \tag{1.42}$$

These approaches show limitations in continuous or high-dimensional spaces due to the need for exhaustive state evaluation.

Deep Reinforcement Learning (DRL) addresses this limitation by leveraging neural networks to approximate either the value function $V(\boldsymbol{s})$, the Q-function $Q(\boldsymbol{s}, \boldsymbol{a})$, or directly parameterizing the policy $\Pi_\theta(\boldsymbol{a}|\boldsymbol{s})$ (Vignon et al., 2023). These function approximators enable generalization across high-dimensional or continuous state and action spaces. DRL algorithms can be systematically categorized based on their approach to value estimation, environment dynamics modeling, or policy optimization, each of which is thoroughly presented in Section 4.3.1.

DRL methods demonstrate versatility beyond flow control, finding applications across diverse domains such as robotics (Sumanas et al., 2022), video games (Shao et al., 2019), and economics (Mosavi et al., 2020). Nian et al. (2020) offers a comprehensive overview of RL methods and their applications in industrial processes.

1.4.1 Single-step reinforcement learning for open-loop control

Open-loop control reduces to parameter optimization, e.g., tuning SWO period and amplitude for drag reduction, but in some cases, faces black-box, non-convex cost functions, losing analytical form and requiring expensive evaluations (e.g., CFD simulations). Meta-heuristic methods such as Particle Swarm Optimization (PSO) (Kennedy and Eberhart, 1995) and Covariance Matrix Adaptation Evolution Strategy (CMA-ES) (Hansen, 2023), as well as model-based approaches like Bayesian Optimization (BO) (Frazier, 2018) or the ensemble-variational (EnVar) algorithm (Jahanbakhshi and Zaki, 2019) overcome this by balancing parameter space exploration with targeted sampling in order to respect computational constraints.

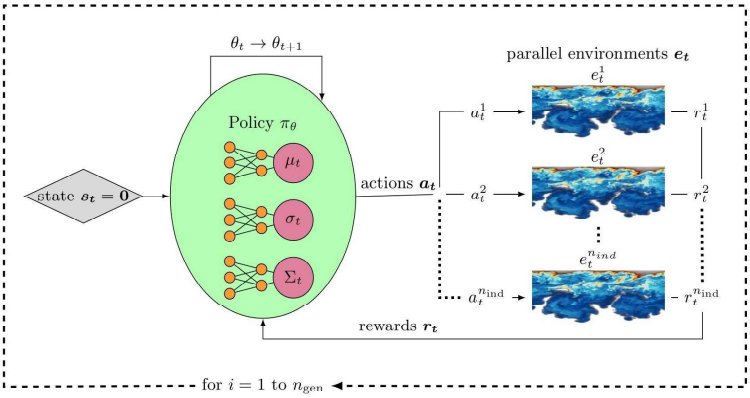

Figure 1.10: Principle of Policy-Based Optimization scheme. μ_t (mean), σ_t (standard deviation), and Σ_t (non-diagonal covariance coefficients) define the action distribution $\pi_\theta(\boldsymbol{a})$. Training parameters include n_{ind} (parallel environments per iteration) and n_{gen} (total training iterations).

A recent extension of DRL to open-loop control, termed Simplified (or Single-step) Reinforcement Learning (SRL), has emerged as a promising parameter optimization approach (Ghraieb et al., 2021; Lai et al., 2021). Initial implementations adapted a truncated version of the Proximal Policy Optimization (PPO) DRL algorithm (Schulman et al., 2017) for open-loop scenarios, where each action automatically results in a terminal state. This framework requires multiple environment evaluations at each iteration (executed in parallel or sequentially) to maintain the exploration-exploitation balance. This algorithm has evolved into Policy-Based Optimization (PBO) (Viquerat et al., 2023), which innovatively integrates concepts from CMA-ES by representing the policy distribution through a covariance matrix. As illustrated in Figure 1.10, PBO employs three specialized neural networks to parameterize the mean, standard deviation, and covariance structure of the distribution. The optimization procedure samples a batch of control actions from the current policy distribution (e.g., SWO period and amplitude parameters). These parameters are evaluated in parallel across multiple environment instances (typically numerical simulations), with each evaluation yielding a scalar reward (e.g., drag reduction). The collective reward outcomes then drive the iterative refine-

ment of the stochastic policy through gradient-based updates to the distribution parameters. Unlike closed-loop control, state observations become unnecessary in this open-loop framework, as actions are optimized purely through their reward performance (Ghraieb et al., 2021).

Benchmark studies demonstrate PBO's competitive performance against established optimization methods such as BO (Wang et al., 2023), CMA-ES (Viquerat et al., 2023), and PSO (Mohammadikalakoo et al., 2024). The complete theoretical foundation and implementation details for both SRL and PBO are systematically presented in Section 4.3.

Chapter summary

This chapter has established the foundational context for this thesis, beginning with the industrial relevance of CSP systems and their associated thermal-fluid challenges. Section 1.2 provided an overview of the governing physics, including incompressible wall-bounded turbulence, coupled heat transfer mechanisms, SWO control strategy, and CFD methodologies. The specific flow configuration adopted for this work was delineated, highlighting its link with CSP applications. Furthermore, key machine learning techniques were introduced (particularly ROM and DRL) that underpin the control strategy optimization framework developed in subsequent chapters. Detailed theoretical and implementation aspects of these methods will be expanded in Chapters 3 and 4.

Version française

Ce chapitre esquisse le cadre industriel et technologique dans lequel cette thèse s'insère. Les systèmes à concentration solaires actuels et de nouvelle génération sont présentés, ainsi que les défis technologiques associés. La Section 1.2 détaille les processus physiques sous-jacents, notamment la turbulence pariétale, le transfert thermique turbulent, les oscillations transverses pariétales et les principes de la mécanique des fluides numérique. La maquette numérique adoptée dans cette étude est décrite en mettant en valeur les liens qu'elle entretient avec les systèmes à concentration solaire. Enfin, les algorithmes d'apprentissage automatique mis en oeuvre dans cette étude sont présentés. Une analyse plus détaillée de ces algorithmes est présentée aux Chapitres 3 et 4.

Chapter 2

Preferential enhancement of convective heat transfer over drag

Contents

2.1 Introduction 32
2.2 Flow conditions and numerical procedure 34
2.2.1 Physical modelling 34
2.2.2 Non-dimensionalisation 35
2.2.3 Fluctuating quantities 37
2.2.4 Details of numerical simulations 37
2.3 Validation of uncontrolled numerical simulations 38
2.4 Control strategy: Spanwise wall oscillations 42
2.4.1 Numerical simulation 42
2.4.2 Results for optimal drag reduction 44
2.5 Breaking the Reynolds Analogy: Dissimilar Enhancement of Convective Heat Transfer 45
2.5.1 Dissimilar heat transfer 45
2.5.2 Near-wall streak rotation and distortion under oscillatory conditions 51
2.5.3 FIK identity component analysis of transport phenomena . . 54
2.5.4 Actuation induced enhancement of momentum and thermal mixing 57
2.6 Perspectives on mechanisms and correlations driving the Reynolds-Analogy breakdown 58
2.7 Impact of Prandtl and Reynolds numbers 59
2.8 Conclusion 61
Chapter summary 63

This chapter presents an investigation of heat transfer response to spanwise wall oscillation actuation through DNS simulations at distinct parameter sets. The study serves four primary objectives: (i) establishing the numerical simulation framework, (ii) demonstrating the potential of SWO for achieving dissimilar heat transfer, (iii) investigating the physical mechanisms governing both heat-transfer enhancement

and thermal-momentum transport decoupling, and (iv) assessing the impact of Pr and Re variations on the control effect.

This chapter integrates a published research article spanning Sections 2.1 to 2.6, which addresses objectives (i) to (iii) at $\mathrm{Re}_\tau = 180$ and $\mathrm{Pr} = 1$. Additionally, Section 2.7 examines objective (iv), offering analysis of the influence of Prandtl and Reynolds numbers on flow control performance.

2.1 Introduction

Turbulent flows play an important role in determining the performance characteristics of a wide range of industrial equipment and environmental applications. One significant consequence of turbulence is the increased mixing of momentum, leading to high friction drag on surfaces, which intensifies as the Reynolds number increases. Friction drag can significantly affect the efficiency and effectiveness of a device or process in a broad range of applications (Ricco et al., 2021). This is particularly relevant in the transport sector, encompassing both self-propelling bodies moving through fluids and fluids being transported via ducts or pipes for example (Abdulbari et al., 2013). Due to global warming issues, there is an increasing pressure to reduce transport-related emissions, with friction drag being a significant contributor to these emissions (Asidin et al., 2019). On the other hand, improving heat exchanger performance is a primary technological challenge for a wide range of engineering systems that involve cooling or heating processes. The goal is to achieve greater efficiency, in line with the industrial and societal imperatives for cost-effective and sustainable energy transfer. Enhancing the turbulent fluxes within the wall-bounded region is generally beneficial for improving heat transfer (Zhang et al., 2020). Therefore, a balance needs to be found between drag-induced losses and heat transfer.

Controlling the boundary layer to decrease drag has been an active area of research for many years. In this regard, imposing spanwise wall oscillations (SWO) has emerged as a promising technique as extensively reviewed in Section 1.2.4.

A parametric study conducted by Quadrio and Ricco (2004) using Direct Numerical Simulations (DNS) at $\mathrm{Re}_\tau = 200$ demonstrated a maximum drag reduction of 44.7% with an oscillation period $T^+ = 100$ and amplitude $W^+ = 27$. It is noteworthy that within the range of oscillation parameters investigated ($0 \leq W^+ \leq 30$, $0 \leq T^+ \leq 300$), an overall drag increase was not observed. To the best of the authors' knowledge, the only documented instance of a net drag increase at low Reynolds numbers is reported in the simulations of Jung et al. (1992), specifically for an extended oscillation period of $T^+ = 500$ with $W = 0.8Q_x$, where Q_x denotes the fixed streamwise flow rate, at $\mathrm{Re}_\tau = 200$. This higher period value was found to induce significantly more pronounced drag variations in the periodic equilibrium state compared to shorter periods.

Recent research has also explored using spanwise wall oscillations to control heat

transfer, either by reducing it simultaneously with drag (Rouhi et al., 2023) or by achieving dissimilar heat transfer (Fang et al., 2009; Fang and Lu, 2010; Fang et al., 2010; Ni et al., 2016), as detailed in Section 1.2.4.

The studies by Rouhi et al. (2022); Kuwata (2022) have demonstrated that passive control methods can be used to create Kelvin-Helmholtz (K-H) rollers, which can weaken drag while strengthening heat transfer. These rollers play an important role in disrupting the Reynolds analogy. Indeed, by introducing inflection point in the wall-normal distribution of the streamwise velocity gradient, passive control devices initiate the formation of K-H rollers with a clockwise rotation in the near-wall region. The direction of rotation of these rollers has a significant impact on the variation of heat transfer and drag. For heat transfer, the mixing induced by the rollers is beneficial regardless of the direction of rotation, as it facilitates the transport of heat away from the wall, leading to improved heat-transfer rates. In contrast, the effect on drag is more complex. The clockwise rotation of the K-H rollers reduces the velocity gradient at the wall, weakening skin friction. The K-H rollers thus selectively reduce drag while enhancing heat transfer, underpinning the favourable breakdown of the Reynolds analogy. This insight highlights the critical role of flow structures, particularly K-H rollers, in achieving disparate improvements in heat transfer and drag. By manipulating these flow structures through passive methods, such as high-aspect-ratio longitudinal ribs, efficient heat transfer surfaces with minimal drag penalties can be achieved.

This study, performed by DNS at $\mathrm{Re}_\tau = 180$ and $\mathrm{Pr} = 1$, aims to investigate the effect of spanwise wall oscillations on achieving a dissimilar increase in convective heat transfer compared to drag. In Section 2.2, we describe the flow conditions and detail the numerical procedure. Validation of uncontrolled numerical simulations is carried out in Section 2.3 by comparing the obtained DNS statistics with references at $\mathrm{Re}_\tau = 180$ and analyzing the influence of mesh sizes. Section 2.4 also introduces and validates the control strategy. The results, presented in Section 2.5, initially focus on identifying parameters outside the range studied in Quadrio and Ricco (2004) that enable a heat transfer increase. The parameters $T^+ = 500$ and $W^+ = 30$ serve as the starting point for investigation, as this period has been shown in Jung et al. (1992) to cause an overall drag increase. The FIK identity decompositions (Fukagata et al., 2002) are then investigated to identify the different components contributing to heat transfer and drag variations. Structural analysis of the streaks and quadrant analysis are then performed to gain a better understanding of the thermal/momentum enhancements. While this work focuses on validating the configuration and extending spanwise wall oscillations to control the disproportionate increase of heat transfer over drag, determining the optimal parameters to maximize this effect will be treated in Chapter 4. Instead, the focus is set on fundamentally characterizing the mechanisms responsible for the greater increase in heat transfer compared to drag, induced by the oscillatory wall forcing. To the best of the authors' knowledge, this study represents the first instance of thermal transport enhancement exceeding the friction increase using spanwise wall oscillations in an

incompressible flow, effectively breaking the Reynolds analogy.

2.2 Flow conditions and numerical procedure

2.2.1 Physical modelling

The evolution of the velocity $\boldsymbol{u}$ is given by the incompressible Navier-Stokes equations, and that of the temperature T corresponds to the passive transport of a scalar by the velocity $\boldsymbol{u}$, as described in Section 1.2. In most applications, significant heat transfer leads to buoyancy forces which would result in an active scalar temperature. In this study, these forces are neglected in order to simplify the analysis. Assuming a flow density equal to 1, the following set of governing equations is obtained:

$$\begin{cases} \nabla \cdot \boldsymbol{u} = 0 \\ \dfrac{\partial \boldsymbol{u}}{\partial t} + (\boldsymbol{u} \cdot \nabla)\, \boldsymbol{u} = -\nabla p + \nu \Delta \boldsymbol{u} \\ \dfrac{\partial T}{\partial t} + \boldsymbol{u} \cdot \nabla T = \alpha \Delta T, \end{cases} \tag{2.1}$$

Hereafter, bulk quantities are defined for any field χ as $\chi_b = \dfrac{1}{|\Omega|} \displaystyle\int_\Omega \langle \chi \rangle_t \ \mathrm{d}\Omega$.

The flow rate U_b is therefore obtained for $\chi = u$ and the bulk temperature T_b for $\chi = T$.

In numerical simulations, periodic boundary conditions are imposed in the streamwise and spanwise directions of the flow, as well as homogeneous Dirichlet boundary conditions on the walls for both velocity and temperature (see below).

Concerning the thermal boundary condition, several options are available. Here, mixed boundary conditions are imposed, as described in Section 1.2.2 (see Kasagi et al. (1992) for a more complete presentation). When MBC boundary conditions are specified, an averaged constant uniform heat flux q_w is applied on the walls (see Figure 1.5) while assuming that the temperature fluctuations at the wall are null. Therefore, the temperature at the walls is not time-dependent ($T_\mathrm{w} = \langle T_\mathrm{w} \rangle_t$). Using the global heat balance for constant heat flux, it can be shown that $\langle T_\mathrm{w} \rangle_t$ increases linearly in the streamwise direction ($\langle T_\mathrm{w} \rangle_t = Ax + T_0$), as stated in Kasagi et al. (1992). For this type of thermal boundary condition, an appropriate non-dimensional form of the temperature is defined as:

$$\Theta = \frac{T - \langle T_\mathrm{w} \rangle_t}{T_b - \langle T_\mathrm{w} \rangle_t}. \tag{2.2}$$

It can also be shown that the scaling temperature $\langle T_\mathrm{w} \rangle_t - T_b$ is constant by applying Newton's law of cooling, $q_w = h(\langle T_\mathrm{w} \rangle_t - T_b)$ where the heat transfer coefficient h is constant, and using the constant flux assumption. Furthermore, the assumption

of zero fluctuations at the wall can be directly incorporated in the prescription of the boundary condition:

$$\Theta_\mathrm{w} = \Theta(T = T_\mathrm{w}) = \frac{T_\mathrm{w} - \langle T_\mathrm{w}\rangle_t}{T_b - \langle T_\mathrm{w}\rangle_t} = \frac{T'_\mathrm{w}}{T_b - \langle T_\mathrm{w}\rangle_t} = 0. \tag{2.3}$$

Finally, the bulk temperature is constant ($\Theta_b = \Theta(T = T_b) = 1$) ensuring the thermal stationary condition.

The Mixed Boundary Conditions exhibit a similar analogy (in between streamwise velocity and temperature) to UHG, albeit with a different source term for temperature, and benefit from extensive documentation with numerous available databases on the unactuated case at various Reynolds numbers (Abe et al., 2009; Seki et al., 2006). Moreover, the present work establishes that the source term difference does not significantly contribute to the achieved dissimilarity, rendering MBC analogous to UHG for this specific purpose. The relevance of working with such a configuration, characterized by initial similarities between temperature and streamwise velocity, lies in its ability to limit the potential factors of initial dissimilarity, thereby simplifying the subsequent analysis of dissimilar heat transfer. Consequently, the current study employs a passive scalar temperature with MBC and unity Prandtl number to facilitate a focused investigation of the dissimilarity mechanisms.

2.2.2 Non-dimensionalisation

The dimensionless Navier-Stokes equations are defined as:

$$\begin{cases} \nabla^\star \cdot \boldsymbol{u}^\star = 0 \\ \dfrac{\partial \boldsymbol{u}^\star}{\partial t^\star} + (\boldsymbol{u}^\star \cdot \nabla^\star)\,\boldsymbol{u}^\star = -\nabla^\star p^\star + \dfrac{1}{\mathrm{Re}}\Delta^\star \boldsymbol{u}^\star + \boldsymbol{f}^\star \\ \dfrac{\partial \Theta}{\partial t^\star} + \boldsymbol{u}^\star \cdot \nabla^\star \Theta + A^\star u^\star = \dfrac{1}{\mathrm{Pe}}\Delta^\star \Theta, \end{cases} \tag{2.4}$$

with

$$\Theta = \frac{T - \langle T_\mathrm{w}\rangle_t}{T_b - \langle T_\mathrm{w}\rangle_t},\ \boldsymbol{x}^\star = \frac{\boldsymbol{x}}{H},\ t^\star = t\frac{U_b}{H},\ p^\star = \frac{p}{U_b^2},\ \boldsymbol{u}^\star = \frac{\boldsymbol{u}}{U_b},\ A^\star = A\frac{H}{T_b - \langle T_\mathrm{w}\rangle_t}.$$

Introducing the variable Θ given by Eq. (2.2) into the temperature equation generates a forcing term $f_\Theta = Au$. The approach to implement this forcing term follows the constant bulk temperature (CBT) method, wherein $\Theta_b = 1$ is mandated at each iteration. For velocity, a constant flow rate (CFR) strategy is imposed, fixing $U_b = \frac{2}{3}$ at each iteration. The corresponding forcing term $\boldsymbol{f}^\star = \begin{pmatrix} f_x^\star & 0 & 0 \end{pmatrix}^\mathsf{T}$ arises from the mean pressure gradient driving the flow. This compensates for viscous friction to achieve steady state (see Section 1.2.1). The CFR procedure is common and will not be detailed here. It is worth mentioning that the source term f_Θ has

u as a factor and therefore, contrary to the forcing term in the UHG case, is not uniform in space, which differs from the forcing term $f_x^\star$.

Subsequently, for simplicity, notations without superscript $\bullet^\star$ will refer to dimensionless values. Notations with superscript $\bullet^+$ will denote values scaled by wall units, i.e. non-dimensionalised by u_τ, Θ_τ and Re_τ, the friction velocity, friction temperature and friction Reynolds number, respectively.

The two key dimensionless numbers characterizing the flow are the Reynolds number $\mathrm{Re} = \frac{U_b H}{\nu}$ and the Péclet number $\mathrm{Pe} = \mathrm{Re} \times \mathrm{Pr}$, where Pr is the Prandtl number defined as $\mathrm{Pr} = \frac{\nu}{\alpha}$. For this study, the Prandtl number is set to $\mathrm{Pr} = 1$ to obtain higher similarity between the velocity and temperature equations as $\mathrm{Pe} = \mathrm{Re}$. In this case, identical boundary conditions are utilized for both velocity and temperature. This enables simpler analysis, as differences in the variations of drag and heat transfer induced by actuation will not originate from disparities in the unactuated boundary conditions or Prandtl number effects. This channel flow configuration was selected for its simplicity and ability to induce substantial similarity between the streamwise velocity and temperature fields.

In this configuration, the sole differences between the streamwise velocity and temperature fields that could potentially facilitate attaining dissimilar heat transfer are:

- The source term difference, f_x is uniform in space while f_Θ is not.
- The pressure term in the streamwise velocity equation which indirectly applies the divergence-free condition of the continuity equation. This term is absent in the temperature equation.
- The linearity of the convective term in the temperature equation does not apply to streamwise velocity.

The friction coefficient, C_f, and Nusselt number, Nu, characterizing drag and heat transfer, respectively (see Section 1.2.1) are defined by the following equations in the current configuration:

$$C_f\,(t) = \frac{\tau_w}{\frac{1}{2}\rho U_b^2} = \frac{2\nu}{U_b^2}\frac{\partial\langle u\rangle_{x,z}}{\partial y}_{\mathrm{w}} = \frac{9\nu}{2}\frac{\partial\langle u\rangle_{x,z}}{\partial y}_{\mathrm{w}}, \tag{2.5}$$

$$\mathrm{Nu}(t) = \frac{2Hh}{\lambda} = \frac{2}{T_b - T_{\mathrm{w}}}\frac{\partial\langle T\rangle_{x,z}}{\partial y}_{\mathrm{w}} = 2\frac{\partial\langle\Theta\rangle_{x,z}}{\partial y}_{\mathrm{w}}. \tag{2.6}$$

The first equation simplifies by injecting the shear stress $\tau_w = \mu\frac{\partial\langle u\rangle_{x,z}}{\partial y}_{\mathrm{w}}$ and the imposed bulk velocity $U_b = \frac{2}{3}$. The Nusselt number definition uses the full height of the channel $2H$ as a reference length. Its simplification arises from Newton's law of cooling for internal flows $h = \frac{q_w}{T_{\mathrm{w}} - T_b}$, with $q_w = -\lambda\frac{\partial\langle T\rangle_{x,z}}{\partial y}_{\mathrm{w}}$ and the above non-dimensionalization of H and Θ. Computation of both dimensionless parameters utilizes the forcing terms, f_x and f_Θ.

Cases	(Lx, Ly, Lz)	Δx^+	Δy^+_{min}	Δy^+_{max}	Δz^+	Δt^+	t^+_{collect}	Re_τ
Abe et al. (2009)	$(12.8, 2, 6.4)$	3.0	0.20	5.93	3.0	NA	3960	180
Seki et al. (2006)	$(6.4, 2, 3.2)$	1.1	0.05	0.97	1.1	NA	1677	180
S_2	$(24, 2, 6)$	10.68	0.43	6.16	5.34	$3.41\ 10^{-2}$	3790	179
$S_2{}'$	$(8, 2, 4)$	4.96	0.29	4.13	3.53	$1.1\ 10^{-2}$	5613	178

Table 2.1: DNS unactuated computational conditions at $\text{Re}_\tau \approx 180$.

2.2.3 Fluctuating quantities

For validation, current simulations are compared to reference uncontrolled cases (see Section 2.3). Correspondence between simulations and canonical cases provides a measure of implemented model and numerical method veracity.

To enable forthcoming validations, the definition of the time-wise and phase-wise averages must be clearly established. The time-wise average used in the unactuated case is clearly defined in Eq. (1.2). For the actuated field, the stochastic fluctuations are noted as χ'' and are obtained by removing the phase average of the raw field:

$$\chi''(x, y, z, t) = \chi(x, y, z, t) - \widetilde{\chi}(\varphi, y), \tag{2.7}$$

with the phase average $\widetilde{\chi}(\varphi, y)$ defined as:

$$\widetilde{\chi}(\varphi, y) = \frac{1}{N} \sum_{n=1}^{N} \langle \chi(x, y, z, \varphi + nT) \rangle_{x,z}, \tag{2.8}$$

and where $\varphi \in [0, T]$ is the phase and N is the number of cycles over which the averaging is performed.

Finally, the periodic fluctuations are denoted as:

$$\widehat{\chi}(\varphi, y) = \widetilde{\chi}(\varphi, y) - \overline{\chi}(y). \tag{2.9}$$

2.2.4 Details of numerical simulations

The open source in-house Xcompact3d framework (Laizet and Lamballais, 2009; Laizet and Li, 2011; Bartholomew et al., 2020) is used to perform numerical simulations. Sixth-order compact finite difference schemes are used for spatial discretisation and a third-order explicit Runge-Kutta scheme is chosen for time integration. The condition of zero velocity divergence is ensured using a fractional step method, where a Poisson equation for the pressure gradient is solved with 3D FFTs (Laizet and Lamballais, 2009). For validation, two mesh resolutions, namely S_2 and $S_2{}'$, are investigated and compared with reference simulations (see Table 2.1).

The first reference simulation (Abe et al., 2009) is utilized to validate the velocity statistics, while the second reference simulation (Seki et al., 2006) verifies the temperature statistics.

In a previous study, using also the Xcompact3D code, it was demonstrated that increasing the numerical dissipation in the streamwise direction could, in certain cases, produce results equivalent to those obtained by increasing the spatial resolution (Flageul et al., 2015). This numerical dissipation is introduced using a fourth-order accurate spectral vanishing viscosity method (Lamballais et al., 2011). Consequently, additional dissipation is applied in the streamwise direction for the S_2 mesh only. Section 2.3 demonstrates that this configuration enables the S_2 mesh to yield results comparable to those obtained using the higher resolution $S_2{}'$ grid, despite the reduced spatial resolution of S_2.

After a warm-up period of approximately 1 500 (in t^+), stationarity was verified using the friction velocity and temperature.

2.3 Validation of uncontrolled numerical simulations

Figures 2.1(a) and (b) represent the wall-normal distributions of the mean streamwise velocity and temperature, respectively. Simulations utilizing the $S_2{}'$ mesh are denoted by the solid blue lines, while the orange dashed lines indicate simulations on the S_2 mesh. The black solid lines with crosses represent the reference data of Abe et al. (2009) providing the benchmark velocity profile and Seki et al. (2006) supplying the temperature profile. For both variables, the current simulation results exhibit strong agreement with the reference data, with only negligible discrepancies in the channel center that are statistically inconsequential.

Figures 2.2 and 2.3 convey the second-order statistical moments for velocity and temperature, respectively. Figure 2.2 shows the Reynolds stresses while Figure 2.3 displays the temperature variance and turbulent heat fluxes. For the Reynolds stresses, both simulations exhibit excellent agreement with the reference data, with minor deviations mainly for the $S_2{}'$ mesh. Specifically, the $S_2{}'$ discretization slightly overpredicts the peak of $\overline{u'u'}$ (Figure 2.2(a)) and underpredicts $\overline{w'w'}$ (Figure 2.2(d)). The profiles of the wall shear stress $\overline{u'v'}$ (Figure 2.2(b)) are nearly superimposed for both discretizations, indicating the accuracy of the prediction of this quantity.

Regarding temperature statistics in Figure 2.3, unlike to the Reynolds stresses,

Cases	$\langle C_f \rangle_t$	$\langle \mathrm{Nu} \rangle_t$
Abe et al. (2009)	8.05 10^{-3}	NA
Seki et al. (2006)	8.43 10^{-3}	21.07
Vreman and Kuerten (2014)	8.11 10^{-3}	NA
S_2	8.17 10^{-3}	21.03
$S_2{}'$	8.05 10^{-3}	20.65

Table 2.2: Mean friction coefficient and Nusselt number on mesh S_2 and $S_2{}'$ against reference data.

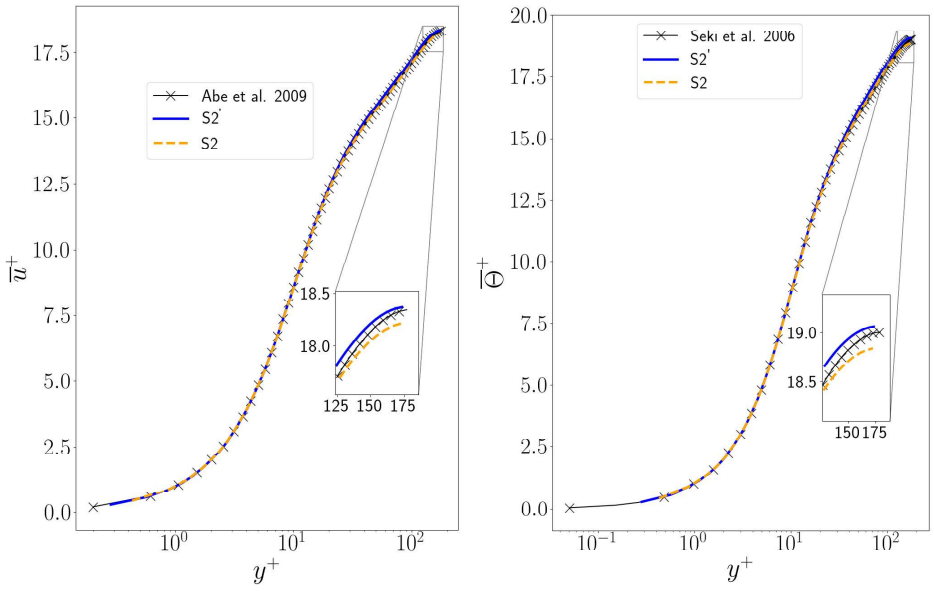

Figure 2.1: Wall-normal distribution of the mean: (a) streamwise velocity and (b) temperature. Statistics obtained from simulations on $S_2{}'$ and S_2 are represented by plain blue line and dashed orange line, respectively. Reference results (Abe et al., 2009; Seki et al., 2006) are shown by the black line with crosses.

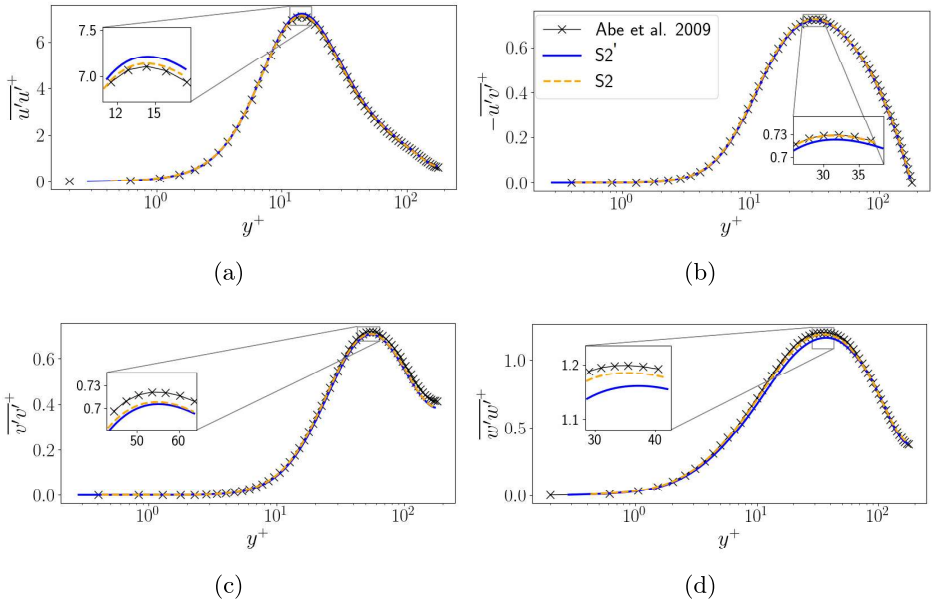

Figure 2.2: Wall-normal distribution of the Reynolds stress components: (a) $\overline{u'u'}$, (b) $-\overline{u'v'}$, (c) $\overline{v'v'}$ and (d) $\overline{w'w'}$ in wall units. See caption of Figure 2.1 for legend information.

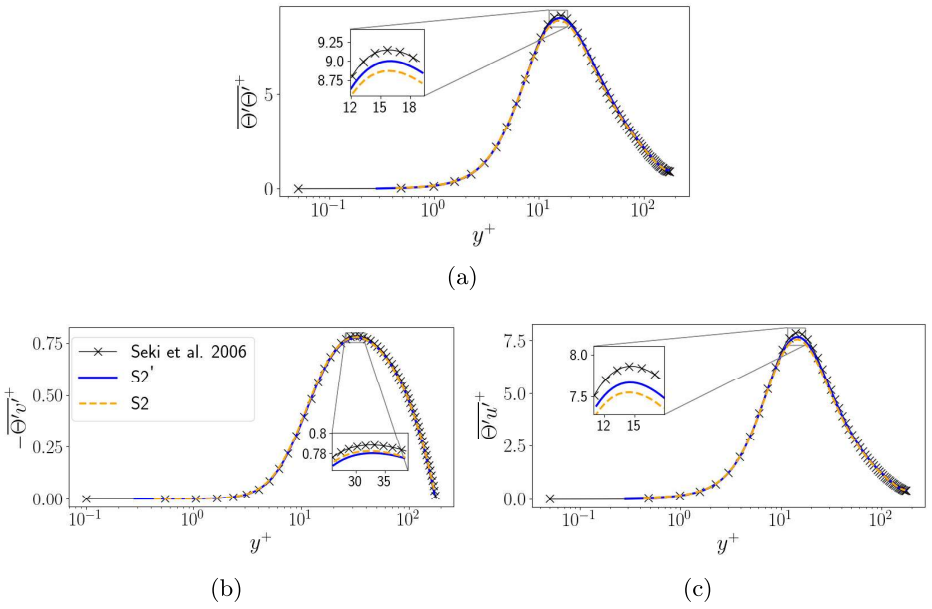

Figure 2.3: Wall-normal distribution of the turbulent heat transport: (a) $\overline{\Theta'\Theta'}$, (b) $-\overline{\Theta'v'}$ and (c) $\overline{\Theta'u'}$ in wall units for Pr = 1. See caption of Figure 2.1 for legend information.

the $S_2{}'$ mesh results show marginally better agreement with the reference data compared to the S_2 mesh. However, both simulations slightly underestimate the peaks of $\overline{u'\Theta'}$ and $\overline{\Theta'\Theta'}$. The discrepancy with the simulations of Seki et al. (2006) could result from their finer mesh resolution. However, their data were collected over a more limited spatial and temporal domain compared to the present study. This smaller sampling could cause statistics that are not fully converged. Overall, Figures 2.2 and 2.3 highlight robust agreement with the reference statistics (Abe et al., 2009; Seki et al., 2006). Slight differences are likely attributable to statistical convergence.

Table 2.2 presents the temporal average of the friction coefficient $\langle C_f \rangle_t$ and Nusselt number $\langle \mathrm{Nu} \rangle_t$, which characterize drag and heat transfer, respectively. Comparison against three distinct sets of reference database serves to validate these quantities. Examining the S_2 mesh results, $\langle C_f \rangle_t$ lies between the reference values, with discrepancies of 3.2% with Seki et al. (2006), 1.4% against Abe et al. (2009), and 0.9% with Vreman and Kuerten (2014). The $\langle C_f \rangle_t$ value from the simulation on the $S_2{}'$ mesh is equal to the one measured by Abe et al. (2009). Regarding the Nusselt number, the S_2 configuration exhibits a discrepancy of 0.2% with the results of Seki et al. (2006), while the $S_2{}'$ configuration shows a larger deviation of 2.0%. Both discretizations accurately predict the relevant statistics in the present study, with the additional observation that the S_2 mesh demonstrates a lower error in computing the Nusselt number.

The pre-multiplied power spectral densities and the wall-normal distribution of the variance of the streamwise velocity and temperature are conveyed by Figure 2.4.

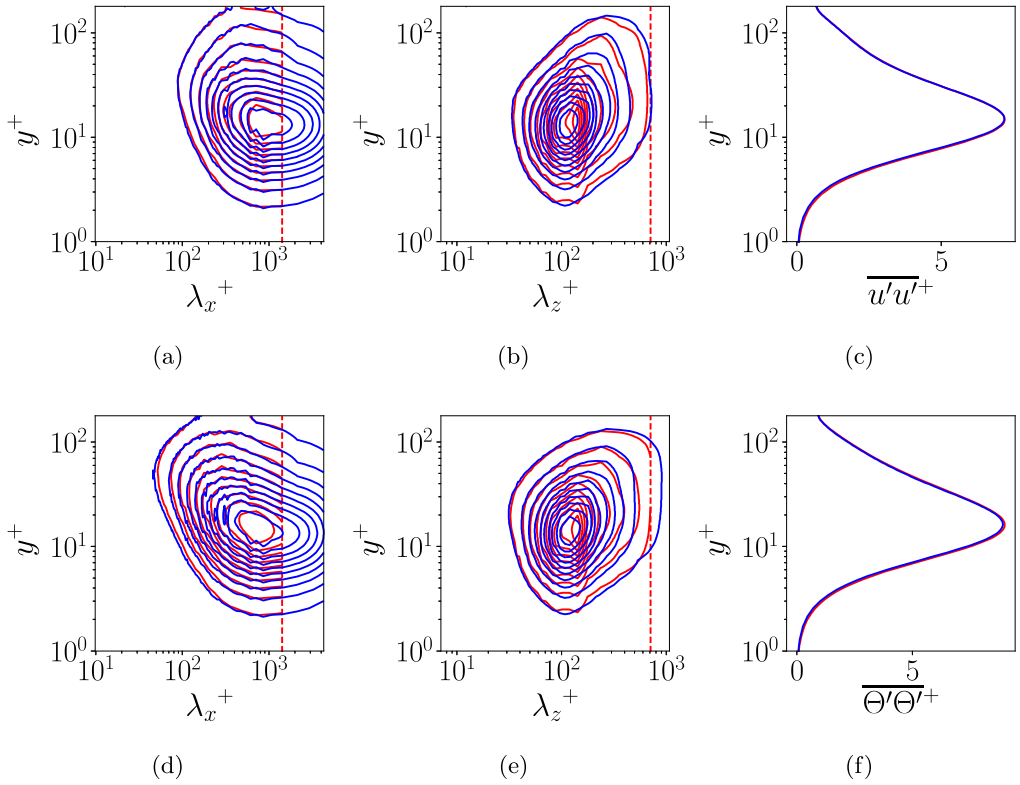

Figure 2.4: (a,b): Pre-multiplied Power Spectral Densities of the streamwise velocity field $[k_\xi \Phi_{uu}(\lambda_\xi, y)]^+$: (a) in the streamwise direction, $\xi = x$, (b) in the spanwise direction, $\xi = z$, and (c) $\overline{u'u'}^+$ wall-normal distribution. (d,e,f): same quantities for the temperature field. Red isolines are the results obtained on mesh $S_2{}'$ and blue isolines, on mesh S_2. Isoline levels span linearly between the minimal and maximal values. Red vertical dashed lines show the limit in domain size of mesh $S_2{}'$.

The left column shows spectra in the streamwise direction while the middle column shows spectra in the spanwise direction. The right column displays the wall-normal variation of $\overline{\chi'\chi'}$ obtained by integrating the pre-multiplied power spectral density $\Phi_{\chi'\chi'}$ along the wavelength:

$$\overline{\chi'\chi'} = \int_0^\infty \Phi_{\chi'\chi'} \, \mathrm{d}f = \int_0^\infty \lambda \Phi_{\chi'\chi'} \, \mathrm{d}\log\lambda. \tag{2.10}$$

These figures show the impact of the domain sizes of meshes $S_2{}'$ and S_2. The isolines and second order stresses match closely between the two meshes, indicating that an increased resolution on S_2 would not necessarily improve the precision of results. This underpins the earlier hypothesis that the reference statistics from Seki et al. (2006) may not be fully converged. As expected, velocity spectrum peaks occur at $\lambda_x^+ \approx 1\,000$ and $\lambda_z^+ \approx 100$. The small $S_2{}'$ domain size does not capture low-frequency content. See Tiselj (2014); Abe et al. (2004a) for a deeper discussion on large scale structures and their impact on statistics.

Analysis of the streamwise velocity and temperature spectra reveals that the streamwise domain length of mesh S_2 is marginally adequate to capture the relevant flow physics. Quantitative assessment of the second-order statistics shows mesh S_2 provides slightly improved precision for predicting the turbulence quantities relevant for this study. Based on these assessments, mesh S_2 is selected for the remainder of this work, examining the effects of spanwise wall oscillations on variations in drag and convective heat transfer. The marginally enhanced domain size and sufficient resolution of S_2 are expected to provide more accurate quantification of the oscillation-induced changes in wall shear stress and Nusselt number under the actuation conditions considered.

2.4 Control strategy: Spanwise wall oscillations

2.4.1 Numerical simulation

The values of T, the control period, and W, the velocity amplitude, being fixed, the spanwise wall oscillations are enforced via the following time-dependent boundary condition:

$$w_{\mathrm{w}}(t_n) = W \sin\left(\frac{2\pi}{T} t_n\right), \tag{2.11}$$

where $t_n = n\Delta t$ denotes the discrete time at step index n, and Δt is the constant time step size of the numerical simulation. This formulation imposes an in-phase harmonic oscillation of the spanwise velocity w at the top and bottom walls.

To enable examination of the actuation effects, phase-averaged statistics are defined per Eq. (2.7), Eq. (2.8) and Eq. (2.9). These leverage the inherent periodicity of the oscillatory forcing to delineate the statistical evolution across discrete phases within the actuation cycle. Such detailed interrogation is necessary to elucidate the fundamental physical phenomena governing the fluctuations in drag and heat transfer.

Due to the imposed oscillations and their periodic time-dependence, two variants of wall units may be employed for non-dimensionalization in the actuated flow:

- χ^+_{nom} where the normalization is based on the mean friction velocity for the unactuated flow ($\overline{u_{\tau,\mathrm{nom}}}$), the so-called nominal configuration;
- χ^+_{ac} where the mean friction velocity for the actuated flow ($\overline{u_{\tau,\mathrm{ac}}}$) is used for scaling.

For both computational meshes, the Courant-Friedrichs-Lewy (CFL) number was constrained under 0.5. Table 2.3 provides details on the actuated flow simulations performed in Section 2.5. To satisfy the CFL limit with Mesh S_2, the time step was reduced to compensate for CFL increase induced by the actuation (compare

Cases	Δt^+_{nom}	t^+_{collect}	number of cycles	phases per cycle: snapshots / stats
S_2	$1.56\ 10^{-2}$	11 500	23	32 / 400
$S_2{}'$	$1.1\ 10^{-2}$	11 500	23	32 / 288

Table 2.3: Computational parameters of the actuated DNS for $T^+_{\text{nom}} = 500$ and $W^+_{\text{nom}} = 30$, values used in Section 2.5.

the values reported in Table 2.1 and Table 2.3 for S_2). All fields were initialised using snapshots from the unactuated case taken in the statistically steady state. The warm-up period was set to 5 actuation cycles for the chosen set of oscillation parameters.

Relative to the unactuated case, attaining statistical convergence necessitates extended integration times, as precise representations are critical for each phase of the actuation cycle. Thirty-two snapshots were acquired evenly during each oscillation period, providing adequate resolution to compute power spectral densities and probability distribution functions across all actuation phases. The mesh utilized for the phase analysis of the first and second order moments is much more refined than the snapshot sampling. This furnishes the phase-resolved data essential for describing the physical processes governing the evolution of turbulence statistics, including drag and heat transfer, within each portion of the periodic actuation cycle.

Bradshaw and Pontikos (1985) were the first to recognize the crucial importance of the rate of change of the spanwise velocity gradient dW/dy, rather than the gradient itself, in reducing the intensity of near-wall turbulent structures and consequently the drag. They found that the spanwise shear needed to vary in space to sustain the drag reduction effect. However, it is important to note that the spanwise velocity gradient does not directly contribute to the drag experienced by the fluid moving along the channel. The spanwise gradient opposes the spanwise wall motion but does not hinder the fluid's streamwise motion. Numerous studies have conducted numerical simulations of turbulent channel flows subjected to spanwise wall oscillations while maintaining a constant streamwise pressure gradient (Quadrio and Ricco, 2004; Dong et al., 2019). Depending on the set of oscillation parameters, these simulations have demonstrated an increase in the flow rate, indicating a reduction in drag. This drag reduction occurs despite the presence of the spanwise velocity gradient, further emphasizing that the spanwise gradient does not directly participate in the overall drag experienced by the flow. Therefore, while the rate of change of the spanwise velocity gradient plays a crucial role in modulating the near-wall turbulence and consequently influencing the drag, it does not directly contribute to the drag itself. The drag reduction observed in numerical simulations with constant pressure gradient conditions underscores this point, as the flow rate increases despite the presence of the spanwise gradient. For these reasons, in the actuated case, the friction coefficient definition remains the same as stated in Eq. (2.5), based on the streamwise shear only.

2.4.2 Results for optimal drag reduction

Parametric studies at low Reynolds numbers have identified optimal dimensionless oscillation amplitude (W^+) and period (T^+) that maximize drag decrease. Significantly, Quadrio and Ricco (2003) demonstrated approximately 40% drag reduction (after the transient phase) at $\mathrm{Re}_\tau = 200$ utilizing control parameters equal to $W^+_{\mathrm{nom}} = 18$ and $T^+_{\mathrm{nom}} = 125$. In Figure 2.5, the temporal evolution of the drag reduction obtained by Quadrio and Ricco (2003) is represented by the black line. The blue line depicts the analogous temporal drag variation at $\mathrm{Re}_\tau = 180$, defined as:

$$\mathrm{Variation}(t) = \left(\frac{C_{f\,\mathrm{ac}}(t)}{\langle C_{f\,\mathrm{unac}} \rangle_t} - 1 \right) \times 100, \tag{2.12}$$

where the subscript "unac" refers to the uncontrolled configuration.

The obtained trend closely resembles that of Quadrio and Ricco (2003), with the drag reduction converging to $\sim 40\%$, validating the implemented wall actuation methodology.

Additional code verification simulations were conducted using oscillation parameters of $T^+_{\mathrm{nom}} = 200$ and $W^+_{\mathrm{nom}} = 12$ on the S_2 mesh. The resulting mean drag reduction (not shown here) exhibited a 2.10% deviation compared to the reference data of Quadrio and Ricco (2004). This minor deviation is statistically inconsequential and could potentially arise from lack of convergence or the documented Reynolds number effects described by Touber and Leschziner (2012). Specifically, they demonstrated that actuation efficacy diminishes with increasing Re_τ.

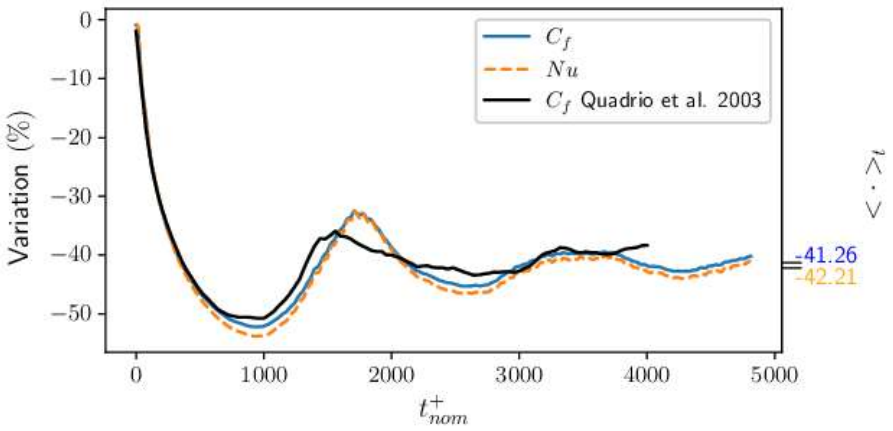

Figure 2.5: Temporal evolution of drag and heat transfer reductions obtained by spanwise oscillatory wall actuation for control parameters $T^+_{\mathrm{nom}} = 125$ and $W^+_{\mathrm{nom}} = 18$. The black curve denotes data from Quadrio and Ricco (2003) for reference.

In Figure 2.5, the temporal evolution of the Nusselt number is superimposed on the drag variation, revealing a strong correlation between the two quantities. The heat transfer reduction closely follows the drag reduction, with both reaching a decrease of approximately 40%. These findings emphasize the capability of spanwise

wall oscillations to substantially reduce not only drag but also heat transfer. The tight coupling between the modulated heat transfer and drag aligns with previous observations by Fang and Lu (2010), who investigated active spanwise wall fluctuations in a compressible flow using LES. It is important to note, however, that their study employed a more complex form of wall motion, specifically out-of-phase and in-phase active spanwise wall fluctuations, which differs from the oscillatory forcing applied in the present work. Nonetheless, the qualitative agreement in the heat transfer and drag reduction trends underscores the robustness of the phenomenon across various flow regimes and actuation strategies.

2.5 Breaking the Reynolds Analogy: Dissimilar Enhancement of Convective Heat Transfer

The disruption of near-wall streak formation by short-period oscillations has been demonstrated to reduce both drag and heat transfer, which aligns with expectations, as attenuated streaks diminish ejection and sweep events, thereby reducing mixing momentum. However, as oscillation periods increase, the potential for streak generation process can resume, which may intensify turbulent transport, leading to an increased drag and potentially heat transfer. This study investigates an actuation scenario that amplifies drag to characterize the associated heat transfer response. To the authors' knowledge, as discussed in Section 2.1, the only reported case of drag amplification at low Reynolds numbers employed an oscillation period of $T^+_{\mathrm{nom}} = 500$ (Jung et al., 1992). In contrast to the well-established drag-reduction case, where the decrease in drag is attributed to the reduced magnitude of streaks due to the interference of the Stokes layer with their formation process (Touber and Leschziner, 2012; Agostini et al., 2014), the mechanisms driving the drag increase for $T^+_{\mathrm{nom}} = 500$ remain unknown. Results demonstrating a strong coupling between drag and heat transfer responses for the drag-reducing scenario may not be directly extended to the drag-increase scenario. Therefore, before determining an optimal control scenario for enhancing heat transfer (which is beyond the scope of this chapter), it is necessary to ascertain whether heat transfer increases concurrently with drag for $T^+_{\mathrm{nom}} = 500$. If this is indeed the case, it is necessary to determine whether the Reynolds analogy holds or breaks down, with either drag or heat transfer being more significantly amplified.

2.5.1 Dissimilar heat transfer

Figure 2.6 illustrates the temporal evolution of spatially-averaged drag and heat transfer under the influence of spanwise wall oscillations with control parameters $T^+_{\mathrm{nom}} = 500$ and $W^+_{\mathrm{nom}} = 30$. This amplitude value has been selected so that the Stokes layer penetrates the initial region of the buffer layer, where the streaks originate and develop. During the first half of the oscillation cycle, both heat transfer and

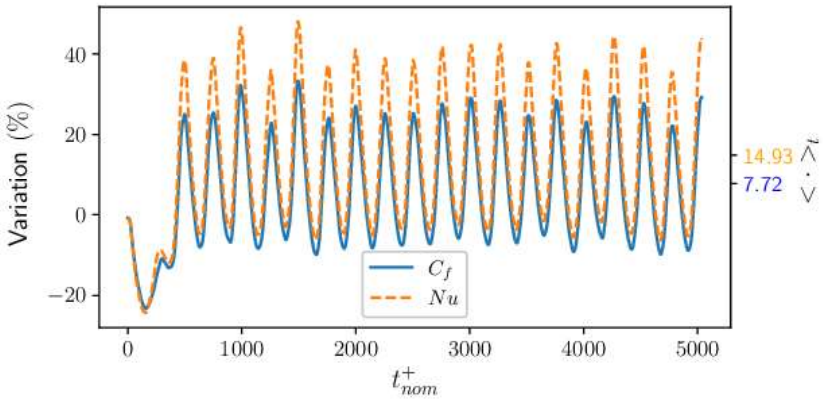

Figure 2.6: Temporal evolution of drag and heat transfer augmentations obtained by spanwise oscillatory wall actuation for control parameters $T^+_{\mathrm{nom}} = 500$ and $W^+_{\mathrm{nom}} = 30$.

skin friction experience a notable decrease of up to 25%. This initial reduction upon the onset of actuation aligns with previous findings and can be primarily attributed to the gradual penetration of the oscillatory Stokes layer into the viscous sublayer and buffer region, which disrupts the formation of streaks. Following this initial phase, the heat transfer and friction exhibit periodic oscillations at the actuation frequency corresponding to $T^+_{\mathrm{nom}} = 500$. As observed in Agostini et al. (2015), the effects on drag and heat transfer are consistent between the first and second halves of the cycle, despite rotating the streaks in opposite directions. The friction drag fluctuates between 5% below and 30% above the baseline, resulting in an average enhancement of 7.72%. In contrast, the heat transfer minima are slightly higher, while the maxima reach approximately 45% above the baseline. This disproportionate amplification of heat transport leads to an average thermal intensification of 14.93%, effectively doubling the increase in friction. These results emphasize the effectiveness of spanwise wall oscillations in enhancing turbulent heat convection. Furthermore, they reveal a significant asymmetry between the modulated heat transfer and drag, with the heat transfer oscillations exhibiting considerably larger amplitudes and gains compared to the friction fluctuations.

An analogy factor A_n is introduced to quantify differences between heat transfer and drag responses, as previously described in Section 1.2.4. This factor, defined as the ratio of Nusselt number to friction coefficient for the actuated and unactuated configurations, is given by

$$A_n(t) = \frac{\mathrm{Nu_{ac}}(t)/\overline{\mathrm{Nu_{unac}}}}{C_{f\,\mathrm{ac}}(t)/\overline{C_{f\,\mathrm{unac}}}}. \tag{2.13}$$

Its phase and time averages can be deduced. They are given by

$$\widetilde{A_n} = \frac{\widetilde{\mathrm{Nu_{ac}}}/\overline{\mathrm{Nu_{unac}}}}{\widetilde{C_{f\,\mathrm{ac}}}/\overline{C_{f\,\mathrm{unac}}}}, \quad \text{and} \quad \overline{A_n} = \frac{\overline{\mathrm{Nu_{ac}}}/\overline{\mathrm{Nu_{unac}}}}{\overline{C_{f\,\mathrm{ac}}}/\overline{C_{f\,\mathrm{unac}}}}. \tag{2.14}$$

Control parameters resulting in $A_n > 1$ indicate greater heat-transfer enhancement

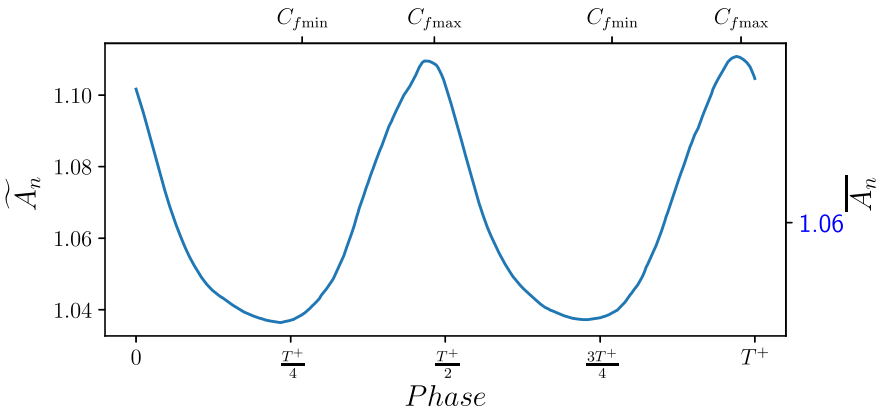

Figure 2.7: Phase-wise averaged evolution of the analogy factor $\widetilde{A_n}$ obtained by spanwise oscillatory wall actuation for control parameters $T^+_{\text{nom}} = 500$ and $W^+_{\text{nom}} = 30$.

relative to the drag increase.

Through the phase-averaged analogy factor $\widetilde{A_n}$, Figure 2.7 distinctly shows the pronounced friction and heat-transfer asymmetry over oscillation phases. $\widetilde{A_n}$ reaches a minimum of slightly below 1.04 during friction minima. Meanwhile, it strengthens to 1.11 coinciding with peak friction coefficient. The shorter duration of $\widetilde{A_n}$ intensification compared to attenuation might indicate that distinct mechanisms drive the reduction and augmentation stages, and also lead to an average value of $\overline{A_n}$ equal to 1.06, highlighting a greater heat-transfer enhancement.

A comparison can be made to the streamwise travelling wave-like wall deformation used in Uchino et al. (2017) at $\text{Re}_\tau = 180$ with constant temperature difference (CTD) boundary conditions. They obtained an average analogy factor of $\overline{A_n} = 1.13$ with optimal parameters. Even though the control method is different, a comparison with the current study can be performed considering the similarities in Reynolds and Prandtl numbers. The current results using spanwise wall oscillations give a lower time-averaged value of $\overline{A_n} = 1.06$ with mixed boundary conditions. This reduced performance is likely because dissimilarity is more difficult to achieve using mixed versus CTD conditions, as shown in Kasagi et al. (2011). Additionally, only one parameter set was tested here, so higher analogy factors may be achievable by optimizing the parameters. Nevertheless, the mixed boundary conditions pose an inherent challenge for maximizing dissimilarity that steady CTD does not encounter, as stated in Section 1.2.4.

The first-order moments obtained for the actuated flow are compared to the unactuated case in Figure 2.8. The black and green curves represent the total average values for the actuated and unactuated flows, respectively. In addition to these curves, the wall-normal distributions obtained for the phase-averaged values corresponding to $C_{f\max}$ and $C_{f\min}$ events are represented by the red and blue dotted lines, respectively. The wall-normal distributions of the time and phase averages of Θ and u are plotted in Figures 2.8(a) and 2.8(b). A slight increase in $\overline{\Theta}$ is observed

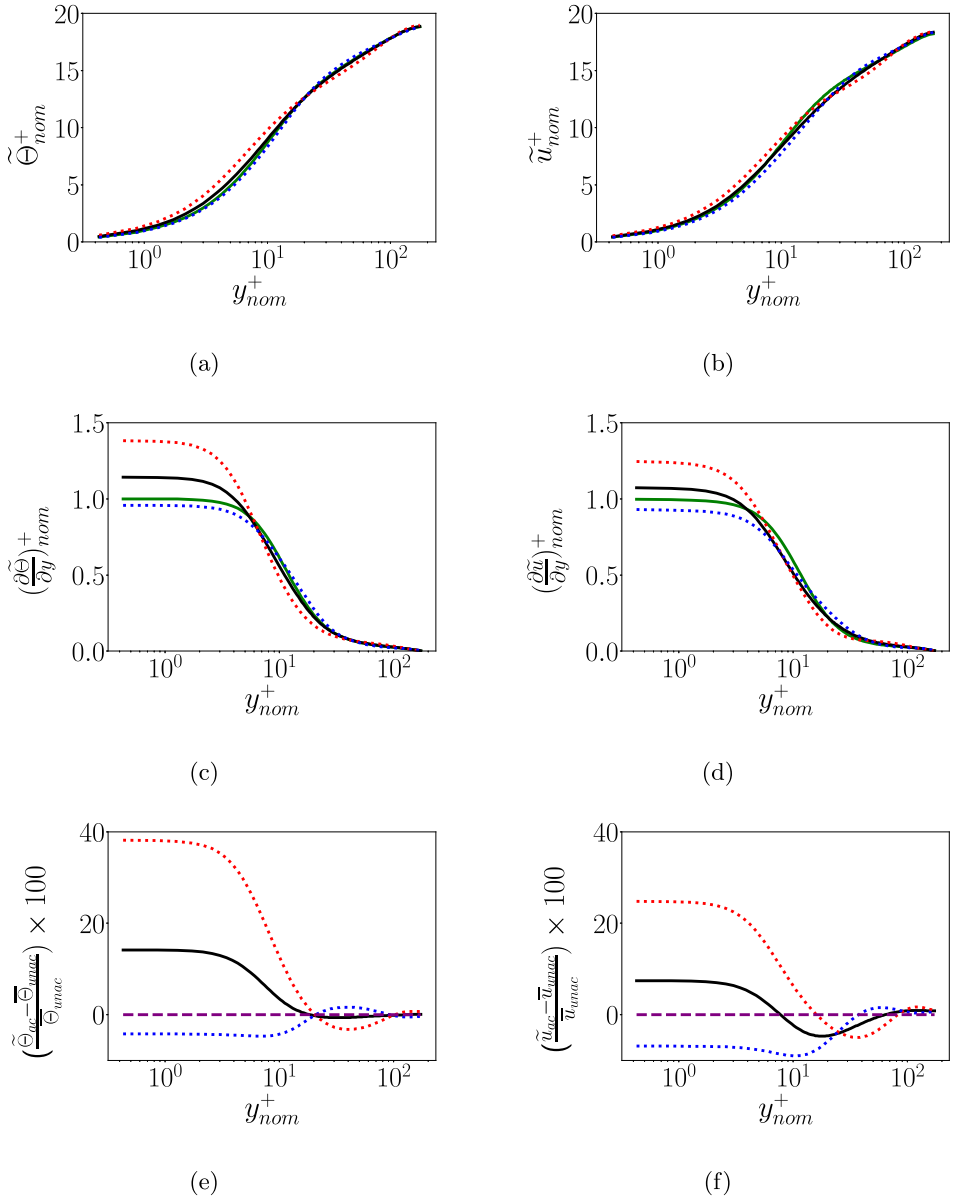

Figure 2.8: Wall-normal distribution of the mean temperature (left column) and mean streamwise velocity (right column). (a,b): in nominal wall units, (c,d): derivative in the wall normal direction in nominal wall units, (e,f): variation with regards to the unactuated case (purple dashed lines indicate zero variation). Blue dotted lines indicate quantities at $C_{f\,\mathrm{min}}$ phases, red dotted lines during $C_{f\,\mathrm{max}}$ phases, black full lines the time averaged **actuated** case and green full lines the time averaged unactuated case.

under the effect of the control, while no change is visible in the curves of $\overline{u}$. However, the curves showing the phase averages reveal that the actuation induces relatively moderate fluctuations, especially at the edge of the viscous sublayer.

To highlight the differences induced by the forcing, the wall-normal derivatives

of Θ and u are shown in Figures 2.8(c) and 2.8(d). These derivatives can be directly associated with the Nusselt and friction coefficients. The figures reveal that the actuation induces an increase for both variables, with a stronger effect on Θ compared to u. Furthermore, the phase variations are more significant for Θ, with the curve corresponding to $C_{f\,\mathrm{min}}$ being very close to the profile obtained for the unactuated case. This emphasizes that the actuation affects Θ more than u, leading to a slightly greater increase in Θ.

To quantify the increase, Figures 2.8(e) and 2.8(f) show the relative increase in percentage compared to the unactuated flow. The results highlight that Θ is indeed more affected than u by the control, with stronger phase fluctuations. The curves for $C_{f\,\mathrm{min}}$ are slightly negative, but the value reached for $C_{f\,\mathrm{max}}$ events is almost twice as strong for Θ than for u, increasing by nearly 40%. The values obtained at the wall for the mean average show an increase of around 15% and 7.5% for Θ and u, respectively, in agreement with the results reported in Figure 2.6. These figures shed light on the fact that the actuation affects Θ more than u, with a stronger amplification of Θ. Moreover, it can be highlighted from Figures 2.8(c) and 2.8(d) that during $C_{f\,\mathrm{max}}$ events, the thermal sublayer thickness reduces slightly more significantly than the viscous sublayer's. The thickness of these sublayers can be represented by the extent of the plateau region at the wall. This observation concurs with more precise measurements comparing the reduction in the sublayer thicknesses from the unactuated to the actuated case at $C_{f\,\mathrm{max}}$ phases. Indeed, this comparison reveals a 26 % decrease for the viscous sublayer against 31 % for the thermal one.

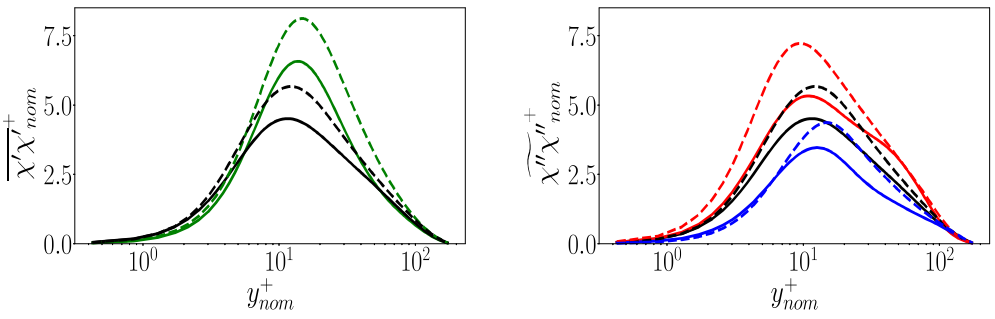

(a) (b)

Figure 2.9: Wall normal distributions of u and Θ variances, (solid): $\chi = u$, (dashed): $\chi = \Theta$. (a): unactuated against actuated, (b): all actuated cases. Black indicates quantities for the time averaged **actuated** case, green for the time averaged unactuated case, blue at $C_{f\,\mathrm{min}}$ phases and red during $C_{f\,\mathrm{max}}$ phases.

The second-order statistics for Θ and u are investigated in Figure 2.9, with their wall-normal distributions represented by dashed and solid lines, respectively. As in the previous figures, the values for the actuated and unactuated cases are represented by black and green lines, respectively. Figure 2.9(a) shows that the variance for

both variables strongly decreases when actuation is applied, with a slightly more significant weakening observed for Θ. These observations are quite surprising, as streamwise velocity fluctuations are typically associated with the magnitude of the streaks. The weakening of the variance suggests that the streaks are attenuated by the actuation, and thus, the drag increase is not directly related to the streak magnitude. However an actuation induced reorientation of the streaks (see Agostini et al. (2014)) could result in the decrease in $\overline{u'u'}$ being compensated by an increase of $\overline{w'w'}$. It is also noteworthy that despite the Prandtl number being fixed at unity, the Θ fluctuations are stronger than those of u.

The phase-conditional wall-normal distributions of the variance are shown in Figure 2.9(b), using the same colour code as previously used in Figure 2.8. Several interesting observations can be made from these distributions. Firstly, as expected and in contrast to the previous observations, the maximum variance corresponds to the case where the drag is maximum. This suggests that the turbulent fluctuations are most energetic during the high-drag phase of the oscillation cycle. Furthermore, if the maximum value reached for each profile corresponds to the edge of the viscous or thermal sublayer, a more significant decrease in the thickness of the thermal sublayer can once more be observed during $C_{f_{\max}}$ phases. This thinning of the thermal sublayer has important implications for the drag and heat-transfer enhancement. The thickness of the sublayer is a crucial parameter to consider when examining the drag and heat-transfer enhancement, as the mixing of momentum primarily occurs at the edge of the sublayer. If mixing takes place closer to the wall, its effect on the drag and heat transfer is expected to be substantial and drive their enhancement. To shed light on the magnitude of the mixing, structural changes of the streaks are first explored, as their rotation may induce stronger mixing. The cross-correlations between Θ and v, and between u and v, are then investigated in the following parts by measuring their contributions using the FIK identity and by analysing their joint Probability Density Functions (PDFs).

The FIK identity, named after Fukagata, Iwamoto, and Kasagi (Fukagata et al., 2002), provides a framework for decomposing the friction coefficient and Nusselt number into contributions from different dynamical effects, such as the laminar, turbulent, and inhomogeneous terms. By evaluating the terms related to the Reynolds shear stresses $\overline{u'v'}$ and the turbulent heat flux $\overline{\Theta'v'}$, the relative importance of the momentum and thermal mixing can be quantified.

Moreover, the joint PDFs of u'-v' and Θ'-v' offer insights into the statistical structure of the turbulence and its phase-dependent behaviour. Previous studies have shown that drag-reducing actuation can lead to a narrowing of the joint PDFs and a reduction in the intensity of ejection and sweep events. Examining the joint PDFs at different phases of the oscillation cycle will reveal how the actuation modulates the turbulent mixing and its association with the drag and heat-transfer enhancement.

2.5.2 Near-wall streak rotation and distortion under oscillatory conditions

In this section, the near-wall streaks' response to the spanwise oscillatory wall motion will be investigated. Structural changes will be discussed with a special focus set towards dissimilarities in between the temperature and momentum streaks.

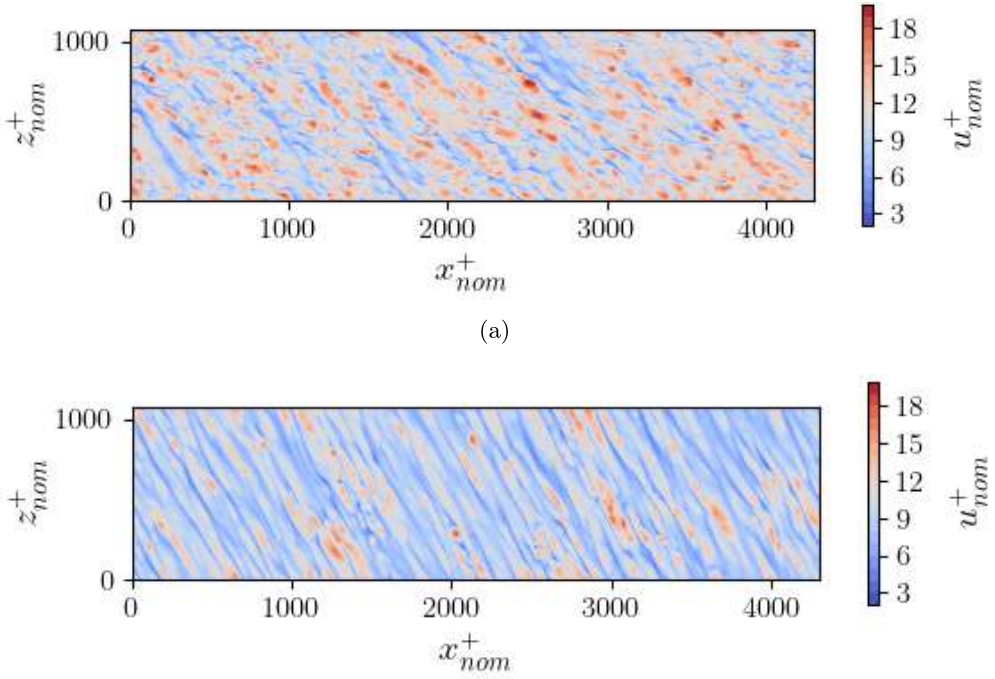

Figure 2.10: Streamwise velocity field at $y^+_{\mathrm{nom}} \approx 11.5$ in the actuated case with $T^+_{\mathrm{nom}} = 500$ and $W^+_{\mathrm{nom}} = 30$, at the phase corresponding to: (a) first $C_{f\,\mathrm{max}}$ and (b) first $C_{f\,\mathrm{min}}$ event.

Snapshots of the streamwise velocity and thermal fields, parallel to the wall and taken at the wall-normal locations corresponding to the peak of their variances, are presented in Figures 2.10 and 2.11, respectively. The lower plots correspond to the first $C_{f\,\mathrm{min}}$ phase, while the upper plots correspond to the first $C_{f\,\mathrm{max}}$ phase. A few comments can be made.

- The imposed spanwise wall oscillations induce substantial structural modifications compared to canonical turbulent boundary layers, evident in the pronounced tilting and distortion of near-wall structures. The structures appear fundamentally distinct from the elongated high- and low-speed streaks inherent to unactuated turbulent flows.

- Fields corresponding to the $C_{f\,\mathrm{min}}$ phase show a strong tilting of the streaks, which is the most pronounced when compared to all other available phases. The streaks also appear more elongated. These observations are confirmed

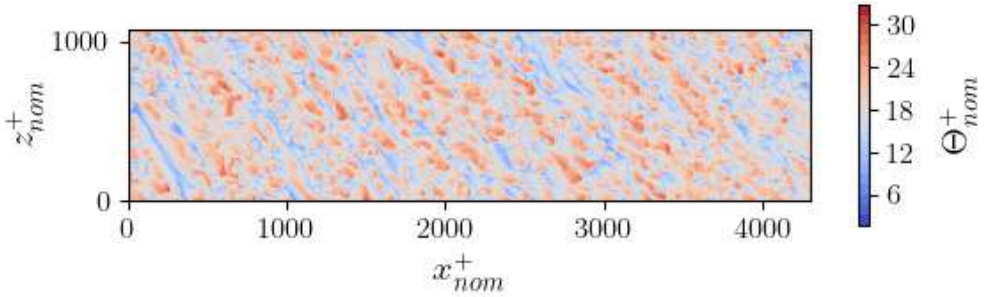

(a)

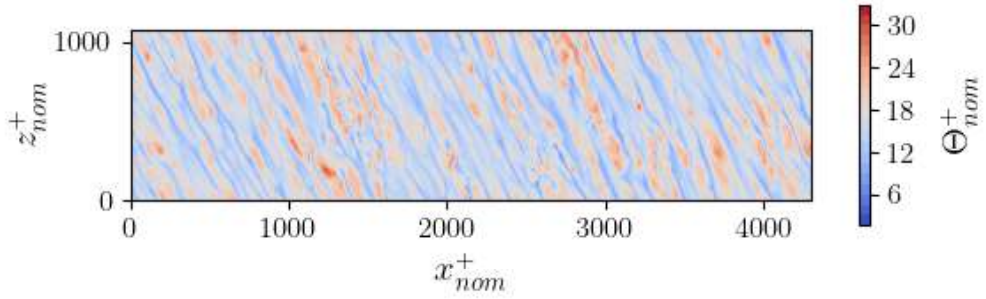

(b)

Figure 2.11: Temperature field at $y^+_{\text{nom}} \approx 13.5$ in the actuated case with $T^+_{\text{nom}} = 500$ and $W^+_{\text{nom}} = 30$, at the phase corresponding to: (a) first $C_{f\,\text{max}}$ and (b) first $C_{f\,\text{min}}$ event.

by performing a 2D correlation map for u in the streamwise and spanwise directions as given by

$$R_u^{x,z}(d_x, d_z) = \frac{\langle u'(\boldsymbol{x}, t) u'(\boldsymbol{x} + d_x \boldsymbol{e_x} + d_z \boldsymbol{e_z}, t) \rangle_{t,x,z}}{[\langle u'^2(\boldsymbol{x}, t) \rangle_{t,x,z}]^{\frac{1}{2}} [\langle u'^2(\boldsymbol{x} + d_x \boldsymbol{e_x} + d_z \boldsymbol{e_z}, t) \rangle_{t,x,z}]^{\frac{1}{2}}}. \tag{2.15}$$

Figure 2.12(b) conveys the map, revealing a streak reorientation angle of 61.8°.

- During the $C_{f\,\text{max}}$ phases, the structures exhibit a stronger magnitude and more complex patterns. Visual comparison between the fields obtained for $C_{f\,\text{min}}$ and $C_{f\,\text{max}}$ reveals that the weakest structures during the $C_{f\,\text{max}}$ phases have a similar tilting angle to those in the $C_{f\,\text{min}}$ phases. However, the strongest structures are reoriented in the flow direction, with a tilting angle that appears to be half as large.

- The 2D correlation map presented in Figure 2.12(a) confirms this observation. The isolines of correlation around the maximum value are oriented at an angle of 36.7° (see yellow line), while the isolines describing the tails of the correlation map show a stronger tilting at an angle of 51.8° (see purple line), closer to that observed at $C_{f\,\text{min}}$ (blue line). The actuated 2D correlation map in Figure 2.12(c) illustrates a reorientation angle very close to that corresponding to the $C_{f\,\text{min}}$ phases for all its structures.

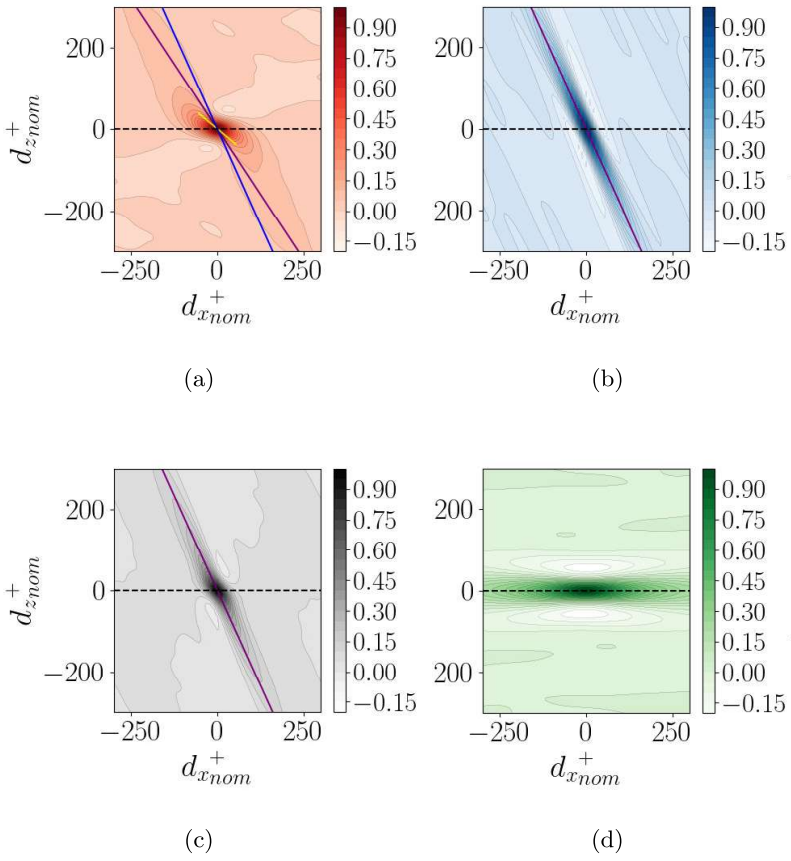

Figure 2.12: 2D correlation maps of the streamwise velocity $R_u^{x,z+}{}_{\text{nom}}(d_x, d_z)$ at $y^+_{\text{nom}} \approx 11.5$, in the actuated case with $T^+_{\text{nom}} = 500$ and $W^+_{\text{nom}} = 30$, at (a): $C_{f\,\text{max}}$ phases, (b): $C_{f\,\text{min}}$ phases, (c): **all phases**, and (d): the unactuated case. The purple line links the maximums of the tail of the correlation, while the yellow line does it for the strongest structures at the center. The blue line replicates the purple line from (b).

- Observations on the temperature 2D correlation maps (not shown here) lead to results analogous to the streamwise velocity. However, the tilting angle is approximately 3° higher for all cases except for the strongest structures at $C_{f\,\text{max}}$ phases, where it is around 6° higher. This difference highlights the non-analogous distortion of the streaks in the velocity and temperature fields.

This analysis underscores the actuation-induced tilting of the streaks. Intriguingly, the strongest structures during the maximum friction phases are characterized by a lower angle of rotation. Although the temperature and velocity fields react to the actuation in a similar manner, differences in their streak distortion are evident, which is consistent with the observed analogy breaking.

2.5.3 FIK identity component analysis of transport phenomena

This section aims to gain some insights on the mechanisms leading to the breakdown of the Reynolds analogy under imposed spanwise wall oscillations. To achieve this, the influence of the forcing terms appearing in the governing equations will be scrutinized. A secondary objective is to gain a more profound understanding of the connection between frictional quantities and turbulent stresses. In pursuit of these goals, a comprehensive decomposition of the friction coefficient and Nusselt/Stanton numbers into their constituent components will be performed by employing the FIK identities, as described in Gomez et al. (2009), and their thermal counterpart introduced in Hasegawa and Kasagi (2011).

The friction coefficient $\widetilde{C_f}$ is known to be directly linked to the turbulent shear stress $\widetilde{u''v''}$, while the Nusselt number is connected to the turbulent heat flux $\widetilde{\Theta''v''}$. This link will be investigated in the case of the actuated case using the FIK analysis, where it is expected that the more pronounced enhancement in heat transfer should also be discernible when comparing the turbulent shear stress and heat flux.

Adapting the FIK identities to the current configuration reveals the decompositions obtained for the Nusselt number and the friction coefficient averaged in time and phase, as detailed in the following equations:

$$
\begin{aligned}
\overline{\mathrm{Nu}} &= \underbrace{6}_{\mathrm{Nu_l}} + \underbrace{3\mathrm{Pe}\int_0^2 (y-1)\overline{\Theta''v''}\,\mathrm{d}y}_{\mathrm{Nu_t}} - \underbrace{\frac{3\mathrm{Pe}A}{2}\int_0^2 (y-1)^2(\overline{u}_b-\overline{u})\,\mathrm{d}y}_{\mathrm{Nu_f}} \\
\widetilde{\mathrm{Nu}} &= \overline{\mathrm{Nu}} + \underbrace{3\mathrm{Pe}\int_0^2 (y-1)\widetilde{\Theta''v''}\,\mathrm{d}y}_{\mathrm{Nu_p}} \\
\overline{C_f} &= \underbrace{\frac{9}{\mathrm{Re}}}_{C_{f_\mathrm{l}}} + \underbrace{\frac{27}{4}\int_0^2 (y-1)\overline{u''v''}\,\mathrm{d}y}_{C_{f_\mathrm{t}}} \\
\widetilde{C_f} &= \overline{C_f} + \underbrace{\frac{27}{4}\int_0^2 (y-1)\widetilde{u''v''}\,\mathrm{d}y}_{C_{f_\mathrm{p}}}.
\end{aligned}
\tag{5.2}
$$

The first constant component (C_{f_l} or $\mathrm{Nu_l}$) is referred to as the laminar component, the second (C_{f_t} or $\mathrm{Nu_t}$) represents the turbulent component, and the third (C_{f_p} or $\mathrm{Nu_p}$) is the periodic component. For the Nusselt number, the term $\mathrm{Nu_f}$ corresponds to the forcing term component, which is absent from the decomposition of the friction coefficient due to its negligible value. The turbulent and periodic components contain a weighting function $y-1$, which weights the shear stresses as a function of proximity to the wall. Consequently, shear stresses close to the wall contribute more significantly than those closer to the channel centre. Strategies that alter the magnitude or very near-wall distribution of turbulent shear stress and heat

flux can therefore have a substantial impact on drag and heat transfer. For instance, a differential thinning of the thermal sublayer relative to the viscous sublayer could partially explain the observed dissimilarity in drag and heat-transfer enhancement. Coincidentally, it was shown previously that the thinning of the thermal sublayer is more significant than the one of the viscous sublayer during maximum friction phases, where the dissimilarity is at its peak.

These formulations establish and clarify the direct link between drag/heat transfer and turbulent shear stress/heat flux while also distinguishing between turbulent components and their periodic fluctuations. Moreover, these decompositions elucidate the contribution of forcing terms to drag and heat transfer.

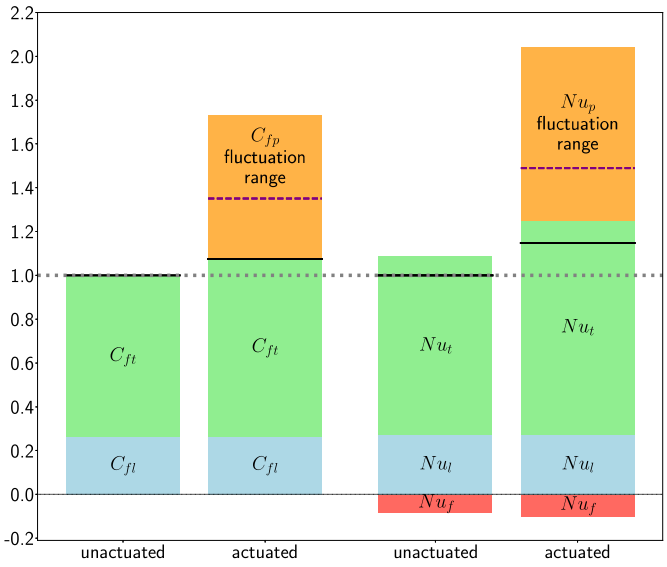

Figure 2.13: FIK identity contributions of $\overline{C_f}/\overline{C_{f\,\mathrm{unac}}}$ and $\overline{\mathrm{Nu}}/\overline{\mathrm{Nu}_{\mathrm{unac}}}$. The actuated case corresponds to parameters $T^+ = 500$, $W^+ = 30$. Blue bins correspond to the laminar component, green bins to the turbulent component, red bins to the forcing term component and orange bins to the range of fluctuations of the periodic component. ..... : total unactuated value, ▬ : total value, - - - : separation of negative and positive fluctuations.

Figure 2.13 shows the results derived from the FIK identity decomposition of the friction coefficient and Nusselt number under conditions of $T^+_{\mathrm{nom}} = 500$ and $W^+_{\mathrm{nom}} = 30$, relative to the unactuated scenario. Each component represented is normalised with respect to the total unactuated value. A notable distinction is observed in between each turbulent component where $\mathrm{Nu_t}$ shows more significant enhancement from which the dissimilarity arises. The greater positive periodic fluctuations highlights a more pronounced difference between the average and high intensity stresses, compared with the difference in between the low intensity and average values. This observed gap is far more pronounced for $\mathrm{Nu_p}$ than it is for C_{f_p}, highlighting once more how the phases of maximum friction carry the highest dissimilarity. Negative values of the $\mathrm{Nu_f}$ component weaken heat transfer, as shown by the red bins. However, the small differences between actuated and unactuated

Nu$_f$ suggest that the forcing term has a negligible effect on increasing heat transfer, therefore on the dissimilarity. Given the insignificance of the forcing term, the only two remaining possible causes of dissimilarity are the solenoidal velocity condition (applied by the pressure term in the streamwise velocity equation) or the linear temperature equation (see Section 2.2).

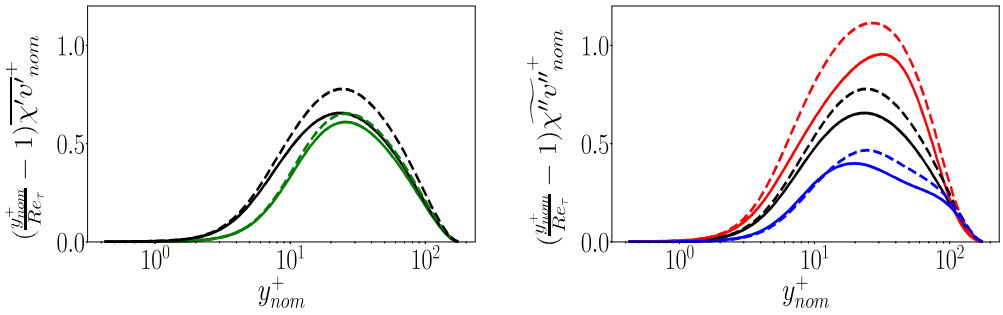

(a) (b)

Figure 2.14: Wall normal distributions of weighted turbulent shear stress and heat flux, (solid): $\chi = u$, (dashed): $\chi = \Theta$. (a): unactuated against actuated, (b): all actuated cases. Blue indicates quantities at $C_{f\,\mathrm{min}}$ phases, red during $C_{f\,\mathrm{max}}$ phases, black the time averaged **actuated** case and green the time averaged unactuated case.

Figure 2.14 shows the wall-normal distribution of the weighted actuated turbulent shear stress and heat flux based on the FIK identity turbulent components. The red lines correspond to the maximum drag phases, the blue lines to the minimum drag phases and the black lines represent the time average, while the green lines, for reference, show the time-averaged distributions obtained for the unactuated case.

From Figure 2.14(a), it is clear that the increase in drag is due to the actuation-induced increase in the weighted turbulent shear stress near the wall, which extends to the peak. This influence on drag is particularly pronounced, given the weight function in the turbulent terms of the FIK identities. A similar conclusion can be drawn for the turbulent heat flux, albeit with a more substantial increase in the weighted stress around the peak, characterising the stronger heat-transfer enhancement. As visible in Figure 2.14(b), during minimum friction phases, this gap starts further away from the wall, implying lower dissimilarity. During maximum friction phases, the stronger enhancement of the heat flux relative to the shear stress starts closer to the wall and is more significant, as expected.

Indeed, the dissimilarity in between the increases in drag and heat transfer becomes more apparent when examining the profiles of turbulent shear stress and heat flux, which results in the enhancement in heat transfer exceeding that of drag, leading to the observed dissimilarity with $\overline{A_n} = 1.06$. A detailed comparison between the $C_{f\,\mathrm{min}}$ and $C_{f\,\mathrm{max}}$ phases has shown the importance of the weighted stresses' influence on the dissimilarity. In addition, this analysis reaffirms the observation that phases marked by pronounced fluctuations in velocity and temperature are further

away from time-averaged quantities, even more so for temperature.

2.5.4 Actuation induced enhancement of momentum and thermal mixing

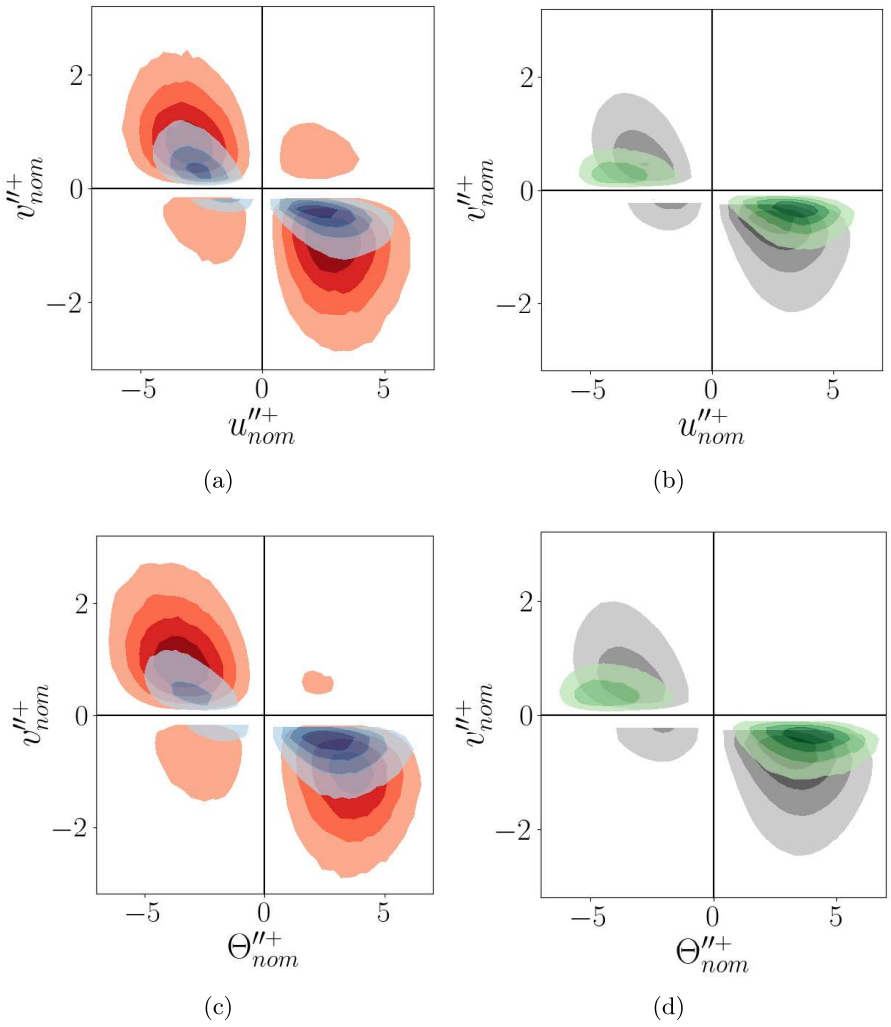

Figure 2.15: Normalised quadrant-conditioned joint PDF (a,b): $|u''v''|\ P(u'', v'')$ at $y^{+}_{\mathrm{nom}} = 11.5$, (c,d): $|\Theta''v''|\ P(\Theta'', v'')$ at $y^{+}_{\mathrm{nom}} = 13.5$. Blue indicates quantities at $C_{f\,\mathrm{min}}$ phases, red during $C_{f\,\mathrm{max}}$ phases, black the **actuated** case and green the unactuated case.

In the previous section, it was demonstrated how the increase in near-wall turbulent shear stress and heat flux results in better mixing, implying drag and heat-transfer enhancement. In this section, it is attempted to obtain a deeper understanding of the mixing enhancement through quadrant analysis.

Figure 2.15 highlights the probability density functions of the wall normal velocity fluctuations joint with the streamwise velocity/temperature fluctuations. They

are weighted by the absolute values of the fluctuations product in order to separate each quadrants. The dominant quadrants for these quantities are the ejection and sweeps, as is shown in the unactuated case. Comparing the actuated to the unactuated case highlights a pronounced enhancement of the wall normal velocity fluctuations, which is strongest during $C_{f_{\max}}$ phases and weakest during $C_{f_{\min}}$ phases. The $C_{f_{\max}}$ phases also show more enhanced ejections when comparing with the $C_{f_{\min}}$ phases, balancing the outer layer sweeps. The magnitude of the temperature/streamwise velocity streaks was previously shown to reduce with the oscillation (see Figure 2.9), even during $C_{f_{\max}}$ phases. This is a major difference when comparing with the drag reduction mechanism where the decrease of the structure's intensity plays a substantial part. Therefore, this decrease is compensated by a better mixing in between the near-wall and outer layers, highlighted by the enhanced sweep and ejection events in the wall normal direction. This improvement of momentum/thermal mixing results in the observed drag/heat transfer increases. Unfortunately, the probability density functions do not highlight any significant signs of dissimilarity.

2.6 Perspectives on mechanisms and correlations driving the Reynolds-Analogy breakdown

The present study highlights notable differences in the mechanisms driving the breakdown of the Reynolds analogy compared to previous research on Kelvin-Helmholtz (K-H) rollers by Rouhi et al. (2022); Kuwata (2022). Unlike K-H rollers, spanwise-wall motion fails to generate strong vortical motion that creates separation regions, which typically decrease drag while increasing heat transfer. This study finds that the breakdown occurs only when drag is increased, suggesting that different mechanisms may be at play. However, analysis of velocity correlations reveals a stronger correlation between v and Θ than between v and u, indicating that sweep and ejection motions have a more significant impact on heat transfer than on drag. This effect is amplified when drag reaches its maximum, resulting in a 40% increase in heat transfer compared to a 20% increase in drag. These results suggest that the mechanism driving the breakdown of the Reynolds analogy in this study might share some similarities with those observed in K-H rollers. Given the contrasting indications from these observations, it remains unclear whether the mechanisms at play are entirely different from or share similarities with those observed with K-H rollers. Further investigation is required to clarify this ambiguity.

Furthermore, the pressure-velocity correlation term appears to weaken Reynolds stresses while having minimal impact on the v-Θ correlation. This allows the v-Θ term to remain efficient and dominate the heat transfer process, potentially explaining the breakdown of the Reynolds analogy in this case. The dichotomy between the effects on drag and heat transfer may be a key factor in the observed phenomenon. To gain a deeper understanding of the mechanisms involved, further research is es-

sential. Future studies should focus on exploring the hypothesis surrounding the role of the pressure-velocity correlation term, shedding light on the underlying mechanisms, and investigating potential similarities and differences with the mechanisms observed in K-H rollers.

2.7 Impact of Prandtl and Reynolds numbers

The demonstrated capability of SWO to preferentially enhance heat transfer over drag in configurations where breaking the Reynolds analogy is non-trivial represents a fundamentally significant finding. However, practical applications involve varying Prandtl numbers (depending on working fluids) and substantially higher Reynolds numbers than the current $\mathrm{Re}_\tau = 180$, which barely sustains turbulent conditions.

This warrants further investigation into how Re and Pr influence actuation-induced drag and heat transfer modifications. To maintain the validity of the S_2 configuration, the moderate range $\mathrm{Pr} \in \{0.5, 0.71, 1, 2\}$ is examined.

This verifies:

- $Pr \leq 2$ to avoid mesh refinement requirements due to thermal sublayer thinning as Pr increases.
- $Pr \geq 0.5$ to prevent extension to an even larger domain.
- The inclusion of $\mathrm{Pr} = 0.71$, corresponding to air, which is relevant for CSP plant applications (see Section 1.1.1).

The $T^+ = 500$, $W^+ = 30$ simulation at $\mathrm{Re}_\tau = 180$ already included scalar equations for these four Pr values. Additional simulations are thus required for both unactuated and actuated conditions at higher Re. The value $\mathrm{Re}_\tau = 1020$ is selected (matching reference DNS Abe et al. 2004b), which introduces larger turbulent scales compared to the baseline $\mathrm{Re}_\tau = 180$.

Cases	(Lx, Ly, Lz)	Δx^+	$\Delta y^+_{\mathrm{min}}$	$\Delta y^+_{\mathrm{max}}$	Δz^+	Δt^+	t^+_{collect}	Re_τ
S_3	$(24, 2, 6)$	10.68	0.5	5.24	5.31	1.1×10^{-2}	7080	1020

Table 2.4: DNS unactuated computational conditions at $\mathrm{Re}_\tau \approx 1020$.

The corresponding S_3 configuration, presented in Table 2.4 maintains:

- Same domain size and similar resolution to S_2 (Table 2.1).
- Adapted timestep and statistical sampling.
- Four scalar equations (all Pr cases).
- Additional streamwise dissipation (see details in Section 2.2).

The implementation utilizes a coarser mesh to accelerate the initial transient in the unactuated case. The actuated simulation is performed at $T^+ = 500$, $W^+ = 30$, totalling 14 oscillation cycles. Additionally time advancement is handled via a second-order Adams-Bashforth scheme, chosen for its computational efficiency over the third-order Runge-Kutta method employed in the S_2 configuration.

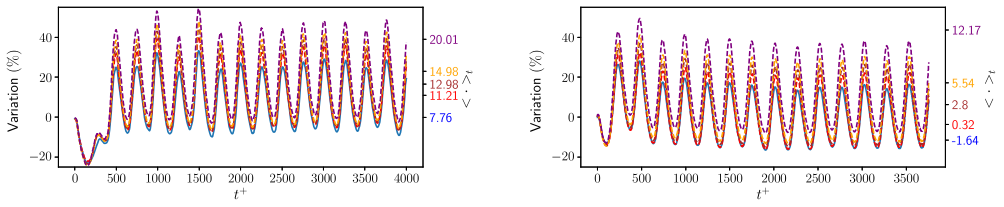

(a) (b)

Figure 2.16: Temporal evolution of drag and heat-transfer variations obtained by spanwise oscillatory wall actuation for control parameters $T^+_{\text{nom}} = 500$ and $W^+_{\text{nom}} = 30$. (a): at $Re_\tau = 180$ (S_2), (b): at $Re_\tau = 1020$ (S_3). Blue solid lines indicate C_f variations, red, brown, orange and purple dashed lines correspond to Nu variations at Pr = 0.5, 0.71, 1 and 2, respectively.

Figure 2.16 displays the variations in C_f and Nu across all investigated Prandtl numbers for the oscillation parameters $T^+ = 500$, $W^+ = 30$, with Figure 2.16(a) and Figure 2.16(b) presenting the results at $\text{Re}_\tau = 180$ and $\text{Re}_\tau = 1020$ respectively. The data in Figure 2.16(a) reveal that all Prandtl numbers yield greater heat-transfer enhancement than drag ($\overline{A_n} > 1$), demonstrating dissimilarity occurrence for both Pr < 1 and Pr ≥ 1 cases. The Pr < 1 values examined remain close to unity, potentially explaining this consistent behavior. Among the tested Prandtl numbers, higher Pr values produce stronger heat-transfer enhancement (thus dissimilarity), reaching a maximum 20% increase at Pr $= 2$. This trend aligns with the findings of Rouhi et al. (2025) for $T^+ > 250$, where they observed greater dissimilarities ($\overline{A_n} > 1$) for higher Prandtl numbers up to Pr $= 20$, suggesting the current results may extend beyond the current Pr $= 2$ limit. If confirmed, this would imply molten salts (Pr $= 20$) respond more favorably to this control strategy than air (Pr $= 0.71$) (see Section 1.1.1 for CSP context). The enhanced heat transfer at larger Pr likely stems from thermal sublayer thinning, which exposes temperature fluctuations to stronger Stokes layer production and consequently amplifies small-scale energy, consistent with the mechanism described in Rouhi et al. (2025). While the transient Nu evolution remains similar across all cases, equilibrium state variations grow more pronounced with increasing Pr. For all investigated Prandtl numbers, the disparity between heat transfer and drag variations peaks during maximum friction phases.

Regarding Reynolds number effects, comparison with Figure 2.16(b) (showing equivalent variations at $\text{Re}_\tau = 1020$) reveals analogous Prandtl number trends. However, the transient phase shortens at $\text{Re}_\tau = 1020$ and exhibits approximately 10% weaker friction reduction. Both drag and heat-transfer variations shift downward, yielding lower mean enhancements, which is consistent with the work of Marusic et al. (2021), demonstrating that higher Re requires larger amplitudes for com-

parable effects. Specifically, a 1.64% C_f decrease is observed, coupled with Nu increases ranging from 0.32% (Pr = 0.5) to 12.17% (Pr = 2). This simultaneous drag reduction and heat-transfer enhancement presents particular practical interest for industrial applications.

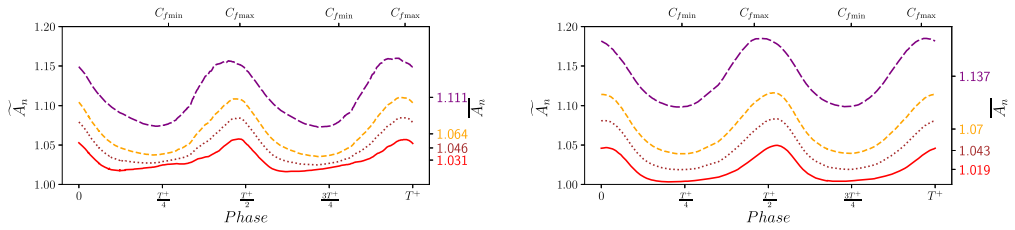

Figure 2.17: Phase-wise averaged evolution of the analogy factor $\widetilde{A_n}$ obtained by spanwise oscillatory wall actuation for control parameters $T^+_{\text{nom}} = 500$ and $W^+_{\text{nom}} = 30$. (a): at $Re_\tau = 180$ (S_2), (b): at $Re_\tau = 1020$ (S_3). Red solid, brown dotted, orange dashed and purple dash-dotted lines correspond to $\widetilde{A_n}$ variations at Pr = 0.5, 0.71, 1 and 2, respectively.

To determine whether the dissimilarity observed at $\text{Re}_\tau = 180$ persists at higher Reynolds numbers, Figure 2.17 presents the analogy factor variations across both Re and Pr values. The Prandtl number effects on $\widetilde{A_n}$ follow predictable patterns from prior analysis, with the largest disparities occurring during maximum friction phases. Notably, the Pr = 2 case exhibits consistent dissimilarity improvement throughout all actuation phases, accompanied by a slightly flattened peak. These trends remain qualitatively similar at $\text{Re}_\tau = 1020$, though with two important distinctions: (i) the analogy factor peak undergoes a small rightward shift, and (ii) the curves for Pr < 1 and Pr ≥ 1 shift downward and upward, respectively, indicating a modest amplification of the Pr influence on dissimilarity with increasing Re. This finding is particularly encouraging, as it demonstrates that dissimilar heat transfer not only persists but shows slight enhancement for Pr ≥ 1 at higher Re. For air (Pr = 0.71), the dissimilarity effects remain nearly consistent across Re_τ values ($\overline{A_n} = 1.046$ versus $\overline{A_n} = 1.043$), suggesting compatibility with this control strategy for industrial applications, albeit weaker than for higher Pr fluids.

2.8 Conclusion

The efficacy of spanwise wall oscillations as a versatile technique for manipulating both drag and heat transfer in turbulent channel flow is elucidated by the present chapter. Through direct numerical simulations conducted at $\text{Re}_\tau = 180$ and Pr = 1, it is demonstrated that:

- Heat transfer can be attenuated by employing oscillation parameters known to reduce drag, such as $T^+ = 125$ and $W^+ = 18$. Under these conditions,

the reduction in heat transfer and drag go hand in hand, thus preserving the Reynolds analogy. This behaviour is anticipated, as these control parameters suppress streak formation and the associated ejection-sweep motions, thereby weakening the turbulent mixing that governs both momentum and heat transport.

- In contrast, the results demonstrate that heat transfer can be enhanced by increasing the oscillation period. By setting $T^+ = 500$ and $W^+ = 30$, a substantial enhancement in heat transfer is observed, underscoring the potential of spanwise wall oscillations as a mean to intensify thermal transport.

- Remarkably, under the extended oscillation period and amplitude, the heat transfer intensification of 15% markedly surpasses the modest 7% drag increase. This dissimilarity between heat transfer and drag points to a breakdown of the Reynolds analogy, highlighting the potential of spanwise wall oscillations for optimising heat exchanger performance.

- The oscillations induce a spanwise tilting of the streaks which is most significant during $C_{f_{\min}}$ phases and not as important for the strongest structures during the $C_{f_{\max}}$ phases. Furthermore, this rotation is shown to be slightly superior for the temperature streaks, which aligns with the analogy breaking.

- The FIK identity analysis shed light on the physical mechanisms governing the observed heat-transfer enhancement exceeding drag rise under spanwise wall oscillations. The negligible impact of the forcing terms showed, by elimination, that the perceived dissimilarity arises from the pressure term in the streamwise velocity equation and the linearity of the temperature equation. To strengthen the dissimilarity, the turbulent heat flux has to be amplified over the turbulent shear stress in the near-wall region, which was achieved in the present work.

- The current study provides compelling evidence that mixing momentum associated with sweep/ejection events is amplified for the actuated case, despite the apparent diminution of the streak amplitude. This enhanced mixing leads to a thinner thermal sublayer in comparison to the viscous sublayer, during the maximum drag phases, where the dissimilarity is the highest. As a result, the intensification of heat transfer surpasses the drag increase, potentially influenced by the pressure redistribution term. However, further investigation is necessary to confirm this hypothesis. These results highlight the complexity of the interactions between the oscillatory Stokes layer and the turbulent flow structures, which extend beyond a simple dependence on streak strength. Further research is essential to elucidate the intricate mechanisms governing these interactions and their profound impact on thermal transport in actuated turbulent boundary layers.

The foundation for future research aimed at exploiting the dissimilarity between heat transfer and drag through spanwise wall oscillations is laid by the present

work. By gaining a deeper understanding of the fundamental mechanisms driving this phenomenon, the development of optimised flow control strategies designed to enhance heat exchanger efficiency can be guided. Consequently, future work should systematically explore the parameter space to identify optimal oscillation conditions that maximize heat transfer augmentation while minimizing drag penalties. This targeted investigation forms the core of Chapters 4 and 5.

Furthermore, complementary simulations across a range of Reynolds ($\mathrm{Re}_\tau \in \{180, 1020\}$) and Prandtl ($\mathrm{Pr} \in \{0.5, 0.71, 1, 2\}$) numbers, using fixed actuation parameters ($T^+ = 500$, $W^+ = 30$), reveal that:

- Heat transfer enhancement and dissimilarity increase with Pr, suggesting superior receptivity of high-Pr fluids to this control strategy.
- Similar trends are observed at higher Re, albeit with attenuated variations of Nu and C_f, resulting in the favorable scenario of heat transfer increase combined with drag reduction. The dissimilarity effect is also preserved across Reynolds numbers, showing slight improvement with increasing Reynolds for $\mathrm{Pr} \geq 1$ and a moderate reduction for $\mathrm{Pr} < 1$.

The robustness of these results is encouraging for industrial applications. However, the idealized conditions of the numerical configuration (symmetric heating, passive scalar temperature) may limit direct translation to industrial settings. Additionally, this translation is subject to practical challenges such as mechanical constraints and the energy cost of sustained oscillations at relevant scales.

Chapter summary

This chapter establishes a fundamental benchmark for subsequent investigations by identifying SWO parameters that preferentially enhance heat transfer relative to drag. The study introduces and validates the numerical framework, including the computational methodology, simulation meshes, and particularly the thermal boundary condition. Through snapshot visualizations, two-dimensional autocorrelations, quadrant analysis, and FIK identity decomposition, the work provides the first insights into Reynolds analogy breakdown under SWO actuation, laying the groundwork for the following chapters.

The study reveals two distinct thermal-momentum transport regimes under SWO actuation at $\mathrm{Re}_\tau = 180$ and $\mathrm{Pr} = 1$. When employing drag-reducing parameters ($T^+ = 125$, $W^+ = 18$), heat transfer attenuation accompanies drag reduction, preserving the Reynolds analogy through suppressed streak formation and weakened turbulent mixing. Conversely, extended oscillation periods ($T^+ = 500$, $W^+ = 30$) produce disproportionate thermal enhancement, with heat transfer intensification (15%) significantly exceeding the modest drag increase (8%), demonstrating effec-

tive Reynolds analogy breaking. This dissimilarity highlights the potential of SWO for heat exchanger optimization, notably for concentrated solar power plants.

Analysis of 2D auto-correlations shows that oscillations induce preferential spanwise tilting of thermal streaks over velocity structures. This rotation of the streaks is most pronounced during minimum-friction phases. The FIK identity analysis reveals that the dissimilarity stems primarily from the pressure term in the streamwise velocity equation rather than from the forcing terms of the temperature and momentum equations. While streak amplitudes diminish, enhanced sweep/ejection events promote more efficient thermal and momentum mixing, resulting in drag and heat transfer increase. A slightly higher thinning of the thermal sublayer compared to the viscous one may also characterize the dissimilarity. These findings underscore the complex interplay between oscillatory forcing and turbulent transport mechanisms, suggesting that thermal-momentum decoupling depends on subtler flow modifications than simple streak strength variation. This warrants further investigation of the underlying mechanisms, as well as a parametric investigation aiming to identify the optimal SWO conditions for dissimilar heat transfer, which is the focus of Chapters 4 and 5.

A Prandtl and Reynolds number investigation revealed that heat-transfer enhancement grows with Pr, likely due to thermal sublayer thinning. The results at higher Re maintain similar trends with attenuated effects, providing a favorable drag reduction/heat transfer increase scenario. The effect of Pr on dissimilarity is shown to be slightly more pronounced at elevated Re. These results suggest molten salts could respond better than air to the present control strategy, while the Re-robustness supports industrial applicability. However, the idealized conditions of the numerical configuration and potential practical challenges could narrow the scope of direct applications to industrial environments.

Version Française

Ce chapitre introduit des paramètres d'oscillation transverse pariétale qui serviront de référence tout au long de cette thèse. Ces paramètres présentent la particularité d'augmenter le transfert thermique de façon plus significative que la traînée. La validation des méthodes de simulation numérique associées est d'abord présentée, couvrant la méthodologie de calcul, les maillages utilisés et les conditions aux limites thermiques. Ce travail propose ensuite une analyse préliminaire de la rupture de l'analogie de Reynolds sous l'effet des oscillations, en s'appuyant sur des visualisations de champs instantanés, des auto-corrélations bidimensionnelles et la décomposition FIK.

L'étude, à $\mathrm{Re}_\tau = 180$ et $\mathrm{Pr} = 1$, révèle deux régimes distincts de transport turbulent sous oscillations. Avec des paramètres réducteurs de traînée ($T^+ = 125$, $W^+ = 18$), l'atténuation du transfert thermique accompagne celle de la traînée, préservant l'analogie de Reynolds via la suppression des structures et l'affaiblissement du mélange turbulent. À l'inverse, des périodes d'oscillation plus longues ($T^+ = 500$,

$W^+ = 30$) produisent un gain thermique disproportionné, avec une intensification de 15% du transfert thermique dépassant largement l'augmentation modérée de 8% de la traînée. Ces résultats démontrent une rupture effective de l'analogie de Reynolds. Cette dissimilarité entre la réponse thermique et dynamique souligne le potentiel des oscillations transverses pariétales pour l'optimisation d'échangeurs thermiques, notamment les centrales solaires à concentration.

L'analyse des auto-corrélations montre que les oscillations induisent une inclinaison préférentielle des structures thermiques par rapport aux structures de vitesse. Cette rotation est plus marquée pendant les phases de friction minimale. L'analyse FIK révèle que la dissimilarité obtenue provient principalement du terme de pression dans l'équation de vitesse longitudinale, plutôt que des termes de forçage. Bien que l'intensité des structures diminue, l'amplification des événements de type balayage/éjection favorise un mélange thermique et dynamique plus efficace, conduisant à une augmentation de la traînée et du transfert thermique. Un amincissement légèrement plus prononcé de la sous-couche thermique par rapport à la sous-couche visqueuse pourrait également être à l'origine de la dissimilarité. Ces résultats mettent en lumière l'interaction complexe entre le forçage oscillatoire et les mécanismes de transport turbulent, suggérant que le découplage thermique-dynamique dépend de modifications d'écoulement plus subtiles qu'une simple variation d'intensité des structures. Cela justifie une étude plus approfondie des mécanismes sous-jacents, ainsi qu'une analyse paramétrique visant à identifier les conditions optimales d'oscillations pour un transfert de chaleur dissymétrique. Ce sujet sera traité dans les Chapitres 4 et 5.

L'étude de l'impact des nombres de Prandtl et Reynolds révèle que l'augmentation du transfert thermique croît avec Pr, probablement due à l'amincissement de la sous-couche thermique associé. Les résultats à Re élevé conservent des tendances similaires, toutefois avec des augmentations atténuées du transfert de chaleur et de la traînée. Cette atténuation permet d'obtenir un scénario favorable de réduction de traînée/amplification du transfert de chaleur. L'influence de Pr sur la dissymétrie s'avère légèrement plus marquée aux Re élevés. Ces résultats suggèrent que les sels fondus seraient plus réceptifs que l'air à cette stratégie de contrôle, tandis que la robustesse en Re soutient son applicabilité industrielle. Cependant, les conditions idéalisées de la configuration numérique et d'éventuelles contraintes pratiques pourraient limiter la portée des applications industrielles.

Chapter 3

Autoencoder-based dimensionality reduction

Contents

3.1 Introduction . . . 68
3.2 Heated turbulent channel flow under spanwise wall oscillation 69
3.2.1 Flow conditions . . . 69
3.2.2 Numerical methods . . . 70
3.2.3 Validation . . . 71
3.2.4 Datasets . . . 72
3.3 Models . . . 75
3.3.1 Proper Orthogonal Decomposition . . . 75
3.3.2 Convolutional AutoEncoder (CAE) . . . 76
3.3.3 β-Variational AutoEncoder (β-VAE) . . . 79
3.3.4 Metrics . . . 82
3.4 Comparing CAE with POD . . . 84
3.5 Balancing latent orthogonality and reconstruction accuracy with β-VAE . . . 90
3.6 Separability of generative factors within the latent space . 95
3.7 Friction augmented CAE for enhanced latent physical coherence . . . 98
3.7.1 Improvements in Nusselt number and friction coefficient predictions . . . 100
3.7.2 Latent separability analysis . . . 102
3.8 Conclusions and perspectives . . . 105
Chapter summary . . . 107

This chapter builds on the results of Chapter 2 and investigates various dimensionality reduction techniques for a heated turbulent channel flow implicit LES dataset at $\mathrm{Re}_\tau = 200$. The study evaluates each technique's performance in terms of compression efficiency, interpretability, and physical coherence. Successful dimensionality reduction could enable creation of a reduced-order dynamical model, offering significant computational cost savings compared to DNS or LES for parametric studies. A critical requirement for such models is robustness to actuation

parameters, as they must interpolate accurately on unknown parameters to be profitable for parametric exploration.

The chapter first follows the structure of an article recently submitted to *Neural Computing and Applications* (from Section 3.1 to Section 3.6). This study focuses on the comparison of diverse dimensionality reduction techniques, namely convolutional autoencoders (Agostini and Leschziner, 2022) (AE), β-variational autoencoders (Eivazi et al., 2022; Solera-Rico et al., 2024), and proper orthogonal decomposition (Borée, 2003). The dataset builds on the results of Chapter 2, consisting of implicit LES at $\mathrm{Re}_\tau = 200$, $\mathrm{Pr} = 1$, with SWO actuation at fixed $T^+ = 500$ and varying amplitudes $W^+ \in [0, 30]$. The focus remains strictly on dimensionality reduction rather than reduced-order dynamical modeling.

An additional study attempting to enhance latent space structure coherence with regards to physical quantities while predicting C_f and Nu from latent samples is also provided in Section 3.7.

3.1 Introduction

The investigation of turbulent flows often relies on Direct Numerical Simulation (DNS), which resolves all ranges of spatio-temporal structures populating the flow. However, for wall-bounded flows, DNS computational resources requirements scale with the Reynolds number Re as Re^3, making it prohibitively expensive, even untenable, for practical applications, particularly when exploring multiple control parameters in flow control strategies such as SWO (Quadrio and Ricco, 2004). This computational limitation has motivated the development of more efficient numerical approaches. Among them, reduced-order models, which were thoroughly described in Section 1.3.

ROMs have become particularly attractive for flow control applications, where computational efficiency is paramount for determining optimal control strategies. For instance, a data-driven ROM, called DManD ROM (Linot and Graham, 2023), employs hybrid autoencoders to find a more compact coordinate system for a turbulent minimal Couette flow, reducing the dimension from 10^5 to 18 degrees of freedom. The combination of DManD ROM with the Soft Actor-Critic (SAC) reinforcement learning algorithm (Linot et al., 2023) has achieved significant speed-up compared to DNS while maintaining control strategy effectiveness. The distribution of parameters used to create a data-driven ROM plays an essential role in ensuring robustness, both in the dimensionality reduction phase and when predicting flow dynamics (Gruber et al., 2022).

This study presents a comprehensive investigation of dimensionality reduction techniques, comparing traditional POD approaches with advanced neural network architectures. The analysis focuses on a heated turbulent channel flow described in Section 3.2. The dataset comprises two-dimensional slices of velocity components and temperature fields, with flow control implemented through spanwise wall

oscillations at a period of $T^+ = 500$.

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

- Section 3.2 presents the numerical framework, detailing flow conditions, simulation parameters and the dataset derived from the simulations.
- Section 3.3 describes the implementation of POD and autoencoder variants. Particular attention is given to the non-linear capabilities of neural architectures in capturing coherent structures and their potential advantages and disadvantages over linear POD decomposition.
- In Section 3.4, the performance of CAEs is evaluated against POD, with a focus on reconstruction accuracy, the capacity to resolve different flow scales, and the minimal number of modes required to represent dominant flow structures.
- Section 3.5 investigates β-VAE architectures, focusing on their ability to provide latent variable pruning and to generate uncorrelated modes through KL-divergence regularisation.
- In Section 3.6, the aptitude of each AE variants and POD to effectively separate latent samples based on their oscillations amplitudes and phases is assessed.

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

3.2 Heated turbulent channel flow under spanwise wall oscillation

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

3.2.1 Flow conditions

The flow conditions for the present study are comprehensively described in Section 2.2 and are briefly outlined here for convenience. Large Eddy Simulations of a

canonical channel flow are carried out at $\mathrm{Re}_\tau \approx 200$.

In order to promote initial similarity between temperature and streamwise velocity, the Prandtl number is set to 1. To further promote this unactuated similarity between thermal and velocity components, the Mixed Boundary Conditions (MBC) (Kasagi et al., 1992) are employed for the temperature, as in Chapter 2.

In Chapter 2, the parameters $T^+ = 500$ and $W^+ = 30$ led to a preferential increase in Nu over C_f. The variation of these quantities induced by SWO can be expressed as a percentage in the following quantity:

$$\Delta C_f = 100 \times \left(\frac{\langle C_{f,\mathrm{ac}} \rangle_t}{\langle C_{f,\mathrm{unac}} \rangle_t} - 1 \right). \tag{3.1}$$

3.2.2 Numerical methods

In this study, the numerical simulations are carried out using the open-source Xcompact3d framework (Laizet and Lamballais, 2009; Laizet and Li, 2011; Bartholomew et al., 2020). Given the extensive parameter space investigation required for W^+ variations in this study, implicit LES are employed instead of DNS to lower computational costs. This approach also effectively reduces the dimensionality of the original flow space representation. Details of the coarse mesh utilized in the LES can be found in Table 3.1, along with DNS comparison references (Abe et al., 2009; Seki et al., 2006; Guérin et al., 2024).

Cases	(Lx, Ly, Lz)	Δx^+	$\Delta y^+_{\mathrm{min}}$	$\Delta y^+_{\mathrm{max}}$	Δz^+	Δt^+	$\mathrm{t}^+_{\mathrm{collect}}$	Re_τ
Abe et al. (2009)	$(12.8, 2, 6.4)$	3.0	0.20	5.93	3.0	NA	3960	180
Seki et al. (2006)	$(6.4, 2, 3.2)$	1.1	0.05	0.97	1.1	NA	1677	180
S_2 **(Guérin et al., 2024)**	$(24, 2, 6)$	10.68	0.43	6.16	5.34	$1.56\ 10^{-2}$	$> 10\ 000$	179
S_1	$(8, 2, 6)$	11.0	2.0	15.7	11.0	$7.3\ 10^{-2}$	10 000	196

Table 3.1: Simulations computational conditions, against references (Abe et al., 2009; Seki et al., 2006), and S_2 (Guérin et al., 2024).

The DNSs conducted with the S_2 configuration of Section 2.2 employ a significantly finer mesh compared to the present S_1 configuration. Specifically, the mesh resolution in the streamwise (Δx), spanwise (Δz), and the two wall-normal directions (Δy_{min} and Δy_{max}) is enhanced by factors of 1.03, 2.06, 4.65, and 2.55, respectively. Additionally, the computational domain extends three times further in the streamwise direction, providing a more detailed resolution of the turbulent structures. Consistent with the methodology of Flageul et al. (2015), the S_1 mesh incorporates uniform numerical dissipation across all spatial directions. This dissipation is implemented using a fourth-order accurate spectral vanishing viscosity scheme (Lamballais et al., 2011). The integration of added numerical dissipation, combined with the coarser mesh refinement of the S_1 configuration, classifies these simulations as implicit LES. As demonstrated in the subsequent validation, the S_1

mesh configuration delivers results comparable to those from higher-resolution reference DNS grids.

Each simulation starts with a warm-up period of five oscillation cycles, equivalent to 2 500 nominal wall-unit time steps, ensuring that an equilibrium state is achieved. Following equilibrium, an additional 20 oscillation cycles are conducted to gather data snapshots. During the data collection process, 32 snapshots are recorded at regular intervals within each actuation cycle, capturing the flow dynamics comprehensively across 32 distinct phases (see Section 3.6 for phase separability analysis). All simulations are initiated from a common set of input snapshots, derived from the steady-state of an unactuated simulation using the S_1 configuration.

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

3.2.3 Validation

The validation of numerical simulations entails a systematic comparison of unactuated flow statistics against established reference data. Validation of the velocity statistics is conducted in Section 4.2.2, whereas the temperature statistics are evaluated against the numerical results of Seki et al. (2006). The assessment of controlled flow characteristics, specifically heat transfer and drag coefficients, is conducted through comparison with the data reported in Section 2.5.

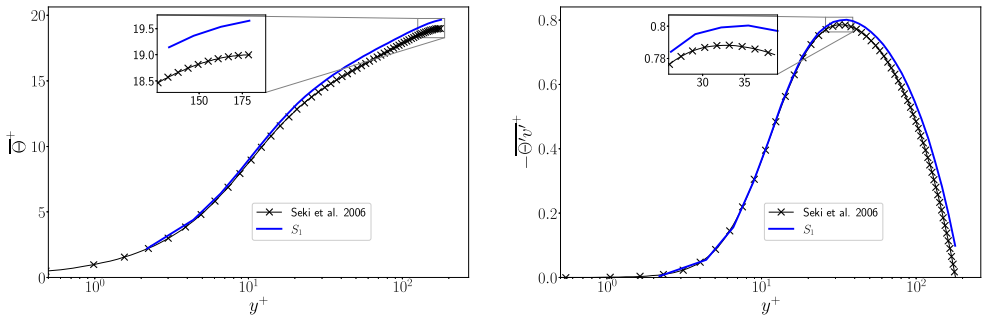

(a) (b)

Figure 3.1: Wall-normal distribution of (a) mean temperature $\overline{\Theta}$, (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. (2006) are depicted by the black line with crosses.

Figure 3.1 presents a comparison between the unactuated simulation statistics and those reported by Seki et al. (2006). The results demonstrate a high level of agreement, with minor discrepancies primarily attributed to differences in Reynolds numbers. It is noteworthy that the discrepancies are more pronounced for the mean temperature $\overline{\Theta}$ compared to the mean streamwise velocity $\overline{u}$ (see Figure 4.1(a)), in the vicinity of its peak. However, the turbulent heat flux $-\overline{\Theta' v'}$, which is directly

related to the Nusselt number Nu, exhibits minor discrepancies. Specifically, in the unactuated case, the S_2 DNS resolution at $\mathrm{Re}_\tau = 200$ shows a 4.7% and 4% difference in C_f and Nu, respectively, when compared to the S_1 resolution. These results demonstrate sufficient accuracy, particularly considering the computational efficiency gained through the S_1 mesh configuration (see Table 4.1 for more details), thus validating the robustness of the numerical approach.

Configuration	$\%\Delta\mathrm{Nu}$	$\%\Delta C_f$
S_1	18.11	9.79
S_2 (Guérin et al., 2024)	14.93	7.72

Table 3.2: Comparison of the percentage variation in the Nusselt number (ΔNu) and friction coefficient (ΔC_f) between the S_1 and S_2 (Guérin et al., 2024) mesh configurations for oscillation parameters $T^+ = 500$, $W^+ = 30$.

Table 3.2 presents a comparative analysis of the variations in averaged Nusselt number (Nu) and friction coefficient (C_f) induced by SWO between the S_1 mesh configuration and the DNS S_2 results. The comparison focuses on results for oscillation parameters given by $T^+ = 500$ and $W^+ = 30$. The S_1 configuration demonstrates strong correlation with the reference DNS data, albeit with marginally enhanced variations in both Nu and C_f. The equilibrium state and transitory characteristics maintain consistency with the reference study.

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

3.2.4 Datasets

The present investigation utilises a comprehensive dataset generated through 21 implicit LESs using the S_1 configuration detailed in Table 3.1. The simulations incorporate spanwise wall oscillations with a fixed period of $T^+ = 500$ and systematically varied oscillation amplitudes W^+ in the range $[0, 30]$, using a regular subdivision of this interval: $W^+ \in \{x = 1.5 \times i \mid i \in \{0, \ldots, 20\}\}$ (see Figure 3.2). Each simulation generated 640 snapshots comprising velocity components u, v, w, and temperature Θ. The data reduction methodology employs a systematic hierarchical approach to efficiently compress the low-dimensional flow field data. Initially, the three-dimensional computational snapshots, encompassing approximately

800 000 grid points per field, undergo a primary dimensionality reduction by extracting two-dimensional slices at the wall-normal position $y^+_{ac} \approx 15$ from both channel boundaries. This specific wall-normal location is selected because the small-scale coherent structures, which constitute the predominant turbulent features at the present Reynolds number, exhibit maximum intensity at this distance from the wall. The strategic selection of these wall-normal positions is fundamentally important for characterising the enhancement mechanisms of thermal and momentum transport induced by spanwise wall oscillations. At these locations, the manifestation of important turbulent phenomena reaches peak intensity, including streak intensity modulation, structural distortions, and alterations in sweep and ejection events (see Section 2.5). Consequently, predicated on the well-established premise that the two-dimensional slices at $y^+_{ac} \approx 15$ encapsulate the essential physical information of the flow field, the compression methodology is specifically applied to these planar sections rather than the complete three-dimensional channel domain. This targeted approach confers substantial computational efficiency, significantly reducing the processing requirements associated with the compression of the full three-dimensional flow field. This initial dimensionality reduction diminishes the data volume to $128 \times 96 = 12\ 288$ grid points per field.

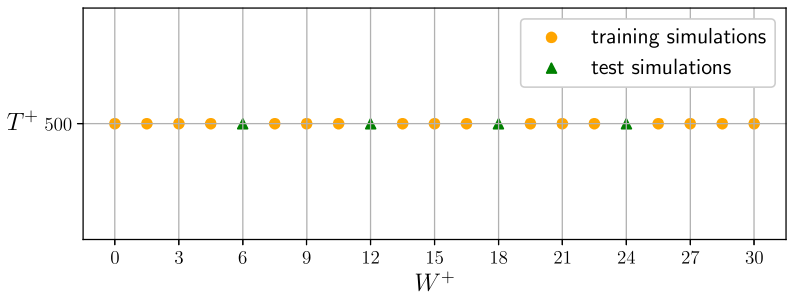

Figure 3.2: Distribution of oscillation amplitudes W^+ within training and test datasets.

To further enhance computational tractability, a subsequent spatial decomposition partitions the computational domain into six distinct subdomains: two sections along the streamwise direction (x-axis) and three sections along the spanwise direction (z-axis). As evidenced in Figure 3.3, this strategic partitioning preserves the integrity of coherent structures within each subdomain, thereby maintaining the essential physical characteristics of the flow. Following this decomposition, the resultant snapshots comprise $64 \times 32 = 2\ 048$ grid points per field. The four primary flow fields (u, v, w and Θ) are subsequently consolidated into a unified three-dimensional array with dimensions $64 \times 32 \times 4 = 8\ 192$ degrees of freedom, wherein the third dimension represents the stacked field variables. This integrated approach to dimensionality reduction demonstrates superior efficiency compared to individual field compression methodologies, as comprehensively documented in Section 3.4.

The complete dataset comprises 161 280 samples, derived from 7 680 samples

per simulation. This 2D field extraction and domain decomposition strategy, combined with the implicit-LES approach, constitutes the initial phase of dimensionality reduction while maintaining sufficient physical information within the dataset.

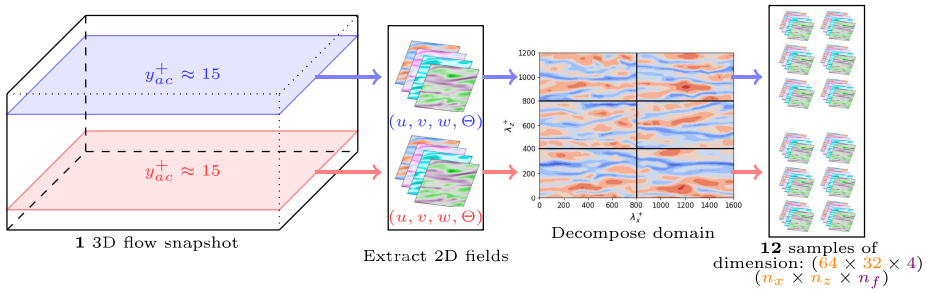

Figure 3.3: Processing of 3D snapshots into 2D fields. 2D fields extraction: Slices are taken at $y^+_{\text{ac}} \approx 15$ and its upper channel symmetric. Domain decomposion: the x-axis is divided into two segments, and the z-axis is split into three parts. λ^+_x and λ^+_z here represent the length (in wall-units) of the domain in the streamwise and spanwise directions respectively.

The test subset consists of samples from simulations with $W^+ \in \{6, 12, 18, 24\}$, whereas the training subset encompasses the remaining samples, establishing an approximate 80/20 distribution (see Figure 3.2). This partitioning strategy enables assessment of the model's interpolation capabilities across intermediate oscillation amplitudes, thus evaluating its robustness to variations in the control parameter. The dataset structure follows a four-dimensional indexing system where $\chi^{(i)}_{j,k,l}$ represents individual data points. The indices correspond to: field type $i = 1, \ldots, 4$ for u, v, w, Θ, simulation $j \in \{0, \ldots, 20\}$, snapshot number $k = 1, \ldots, n_{\text{samples}}$ (taking into account both time indexes and sub-domains), and spatial position $l = 1, \ldots, 2048$ within each snapshot. The preprocessing stage addresses the inherent scale disparities among the flow fields, as illustrated in Figure 3.4(a). The distributions of u and Θ exhibit positive shifts relative to the zero-centred distributions of v and w. The v component displays a particularly concentrated distribution with minimal variance, necessitating careful consideration during the reconstruction process to maintain appropriate weighting among variables.

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

$$\chi^{(i)}_{j,k,l} \leftarrow \frac{\chi^{(i)}_{j,k,l} - \mu_i}{\sigma_i}, \tag{3.2}$$

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

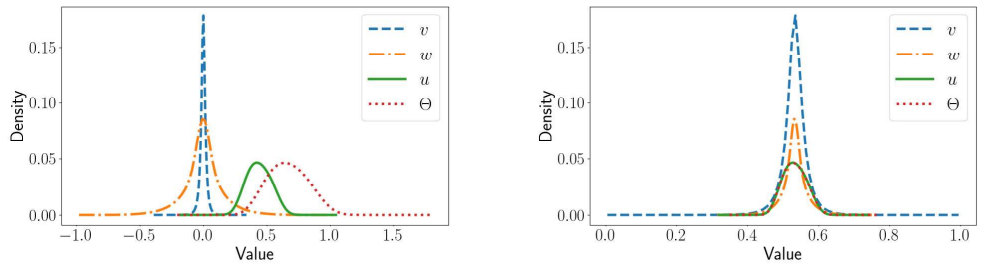

(a) (b)

Figure 3.4: Dataset Scaling: (a) Distributions Prior to Scaling, (b) Distributions Following Scaling. ▬ : u distribution, - - - : v distribution, -·-· : w distribution, ····· : Θ distribution.

$$\chi^{(i)}_{j,k,l} \leftarrow \frac{\chi^{(i)}_{j,k,l} - \min\limits_{i,j,k,l}(\chi^{(i)}_{j,k,l})}{\max\limits_{i,j,k,l}(\chi^{(i)}_{j,k,l}) - \min\limits_{i,j,k,l}(\chi^{(i)}_{j,k,l})}. \tag{3.3}$$

The resultant distributions, illustrated in Figure 3.4(b), exhibit centrality around 0.55 with harmonised variance and bounded values within $[0, 1]$. This normalisation strategy enhances the performance of both POD and AE methods, particularly improving the reconstruction accuracy of the wall-normal velocity component. The normalisation parameters derived from the training dataset are applied consistently across all model implementations. The inverse transformations restore the physical scaling when computing performance metrics. This methodology ensures robust preprocessing while preventing statistical contamination between training and test datasets.

3.3 Models

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

3.3.1 Proper Orthogonal Decomposition

As outlined in Section 3.1, POD serves as a benchmark model for this study. POD is a technique used to extract dominant patterns from low-dimensional data by identifying a set of orthogonal modes ranked by their energy contribution (Cordier and Bergmann, 2008). The process involves computing the eigenvalue decomposition of mean-centered data's covariance matrix (see Section 1.3.1 for more details). Equivalently, the modes can be obtained by Singular Value Decomposition (SVD) of the

data matrix. The singular values of this decomposition represent the energy associated with each mode, while the left singular vectors form spatial basis functions. These modes are ordered by decreasing singular value magnitude, ensuring that the most energetic structures are captured first.

The orthogonality of the modes is a key feature of POD, giving a clear interpretation of their individual contributions to the overall reconstruction. An other notable advantage of POD is the ranking of the modes with regards to their energy contributions which yields additional interpretability. However, achieving convergence may require a large number of modes in scenarios involving turbulent flows or highly non-linear governing equations. This is particularly relevant for the data considered in this study, where the fields u, v, w, and Θ exhibit complex spatial and temporal dynamics. To apply POD, each sample of shape $64 \times 32 \times 4$ is flattened into an 8192-dimensional vector, representing the spatial domain. To replicate the testing process used in the AE models, the extended POD approach (Borée, 2003) is employed in this study. This involves projecting the test data onto the POD spatial basis derived from the training data, i.e.

$$V_{\text{test}}^{\intercal} = \Sigma_{\text{train}}^{-1} U_{\text{train}}^{\intercal} X_{\text{test}}. \tag{3.4}$$

This process compromises the orthogonality of the test data temporal coefficients which no longer form a strict orthonormal basis, as attested in Section 3.5.

3.3.2 Convolutional AutoEncoder (CAE)

This section provides a concise overview of AutoEncoders (AEs) and Convolutional AutoEncoders (CAEs), introducing key concepts of convolutional layers, such as kernels and strides. The neural network architecture employed in this study is then detailed, and serves as a foundational framework for all autoencoder models discussed in this chapter. The section also examines additional model hyperparameters, including activation functions, learning rate and batch size.

As described in Section 1.3.1, an autoencoder is a neural network designed to compress low-dimensional data into a lower dimensional representation (encoding) and reconstruct the original input from this compressed representation (decoding), see Figure 3.5. The compressed representation, referred to as the latent variable and denoted by the symbol $\boldsymbol{z}$, is the coordinate vector of the input data in the low-dimensional space. This representation is essential for the reconstruction of the input during the decoding process (see Section 1.3.1 for more details).

In this chapter, the training process minimizes the MSE loss function. Adopting the same notation as in Section 3.2.4, where $\widetilde{\chi_{j,k,l}^{(i)}}$ represents the autoencoder reconstruction of grid point $\chi_{j,k,l}^{(i)}$, the MSE loss (Eq. (1.36)) for dataset χ is expressed as:

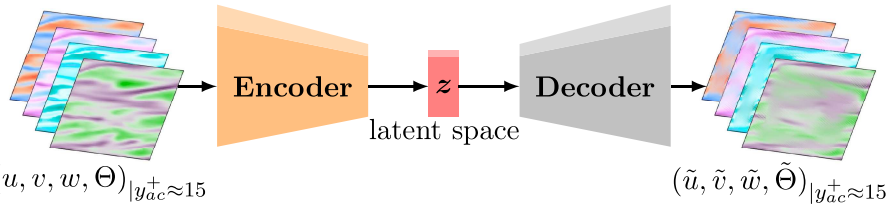

Figure 3.5: Autoencoder representation for the reconstruction of the (u, v, w, Θ) 2D fields. z is the latent representation of (u, v, w, Θ), and ($\widetilde{u}$, $\widetilde{v}$, $\widetilde{w}$, $\widetilde{\Theta}$), its reconstruction through the autoencoder.

$$\mathrm{MSE}(\chi) = \sum_{i=1}^{4} \langle (\chi^{(i)}_{j,k,l} - \widetilde{\chi^{(i)}_{j,k,l}})^2 \rangle_{j,k,l}. \tag{3.5}$$

A convolutional autoencoder builds upon the traditional autoencoder architecture by using convolutional layers for encoding and transposed convolutional layers for decoding. In this study, 2D convolutions are applied, treating each input field as an individual channel. In the encoding phase, a 2D convolutional layer applies a small matrix called a kernel to a 2D input field using a sliding window approach. The kernel moves across the input in discrete steps, performing element-wise multiplication with the input region it overlaps, and then summing the results to produce a single value in the output. This operation detects local features such as edges, textures, or patterns in the input data. The kernel typically slides with a specified stride (the step size between applications) and may involve padding (extra values added around the input) to preserve the spatial dimensions of the data (see Figure 3.6). The decoder uses transposed convolutional layers, which reverse the convolution operation. This process is often referred to as up-sampling and allows the model to reconstruct the spatial details of the original input while maintaining the learned features. By using convolutional layers in both encoding and decoding, CAEs are particularly effective for processing spatial data, such as images or flow fields, e.g. velocity and temperature distributions (Csala et al., 2022; Fukami et al., 2020), as they capture hierarchical spatial patterns while efficiently reducing the dimensionality of the input data.

Using zero padding with a stride of 1 and a 3×3 kernel ensures that the center of the kernel passes over every element of the input field, resulting in an output field with the same dimensions as the input. This strategy is applied in layers 1, 2, and 3 of the present encoder architecture (see Figure 3.7), where the field size remains fixed at 64×32, while the number of channels (or filters) increases from 4 in the input to 24 in the third layer. The stride of 1 in the first three convolutional layers is designed to capture the fundamental patterns of the flow, without reducing its dimensionality at this stage. In contrast, setting the stride to 2 causes the kernel to skip every other row and column, reducing the spatial dimensions of the output to one-fourth the size of the input. Figure 3.6 illustrates

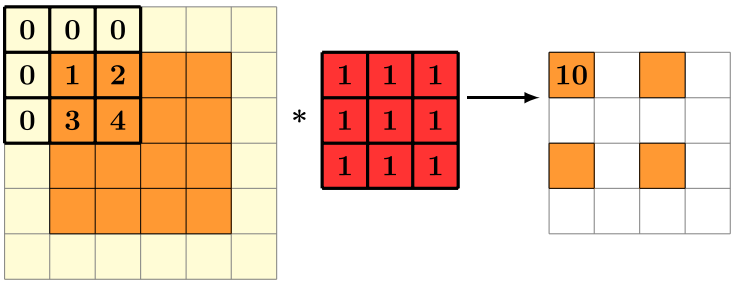

Figure 3.6: Example of a convolutional layer with a 3×3 kernel (red matrix), with zero padding (yellow blocks) and stride $= 2$ on a 4×4 input matrix. Orange blocks on the output correspond to locations where the convolutions are performed on the input.

this effect, where white blocks are skipped due to the stride of 2, leading to a 2×2 output field. This downsampling strategy is used in layers 4, 5, and 6 of the current model to progressively reduce the spatial dimensions, while capturing the more complex features of the flow. After layer 6, the output tensor is flattened into a 4096-dimensional vector and mapped to a 500-dimensional latent vector via a fully connected layer (see Figure 3.7). This aggressive dimensionality reduction accelerates convergence while preserving accuracy, outperforming architectures with intermediate dense layers.

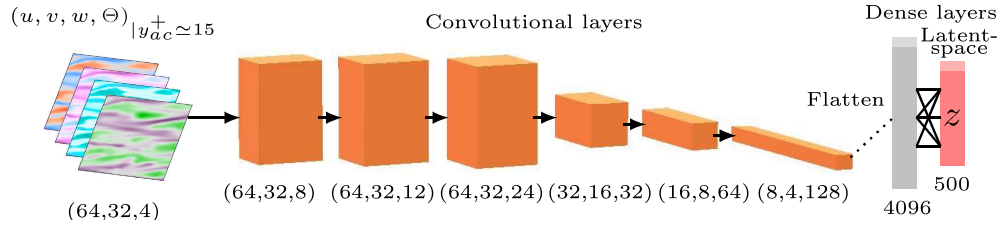

Figure 3.7: Architecture of the encoder component of the CAE. Orange blocks represent convolutional layers, while solid arrows indicate batch normalization. Dotted lines represent the flattening layer, and solid lines denote the fully connected (dense) layers.

The decoder architecture implements a hierarchical expansion of the compressed latent representation through a sequence of precisely configured layers. The reconstruction process begins with a fully-connected layer that transforms the latent vector into a 4096-dimensional space, which is subsequently reshaped into a (8,4,128) tensor format. This initial transformation is followed by a series of transposed convolutional operations that progressively restore the spatial dimensions of the flow field through controlled upsampling procedures. The decoder structure exhibits partial symmetry with the encoder (Figure 3.5), but introduces a modification in the filter distribution. The implemented architecture employs an asymmetric filter arrangement of $(128, 128, 64, 32, 12, 4)$ across its layers, diverging from the traditional symmetric configuration of $(128, 64, 32, 24, 12, 8, 4)$. This architectural adaptation has demonstrated slightly superior reconstruction capabilities, particularly for the wall-normal velocity component ($\approx 2\%$ gain in $E_r(v)$ (Eq. (3.11))). This architec-

ture creates an information bottleneck at the latent space, compelling the encoder-decoder pair to operate synergistically in extracting and preserving the most salient flow features necessary for accurate reconstruction.

The neural network architecture incorporates Exponential Linear Unit (ELU) activation functions (Rasamoelina et al., 2020) throughout its structure, facilitating non-linear transformations within the normalised $[0, 1]$ range established in Section 3.2.4. The training protocol encompasses:

- 200 epochs with 64 samples per batch
- 3×3 convolution kernels
- RAdam optimisation algorithm (Liu et al., 2019)
- Initial learning rate: 10^{-3}
- Adaptive rate reduction: 50% decrease after 10 epochs without improvement
- Minimum learning rate threshold: 10^{-6}
- TensorFlow[1] implementation

The latent dimension (L_d) is established at 500 components, ensuring a minimum 80% energy reconstruction threshold (as given by Eq. (3.11)) across all flow fields. This configuration emerges from systematic parameter exploration, demonstrating optimal performance metrics in the reconstruction task.

3.3.3 β-Variational AutoEncoder (β-VAE)

Although conventional autoencoders demonstrate robust capabilities in reconstructing complex fluid dynamics, their latent space representations often exhibit significant intermodal correlations, which impede the formation of interpretable basis. The traditional POD method maintains strict orthogonality, but its inherent linear framework limits its compression efficiency. The β-VAE architecture represents a sophisticated solution to the challenge of balancing these competing objectives. The fundamental distinction of VAEs lies in their probabilistic approach to latent space encoding. Rather than implementing deterministic mappings, VAEs encode input data into probability distributions from which latent variables (LVs) are sampled, as depicted in Figure 3.8. This probabilistic framework introduces a regulated stochastic element to the encoding process. The present VAE implementation maintains architectural consonance with the previously described convolutional autoencoder (Figure 3.7), preserving identical hyperparameter configurations. The distinguishing elements manifest in three critical aspects: (i) the probabilistic nature of the latent distribution, (ii) the stochastic sampling mechanism employed for latent variable

[1] TensorFlow `2.11.0` with CUDA `12.8` acceleration

generation, and (iii) the modified loss function formulation. This adaptation enables the VAE to maintain reconstruction fidelity whilst encouraging the emergence of more interpretable latent representations.

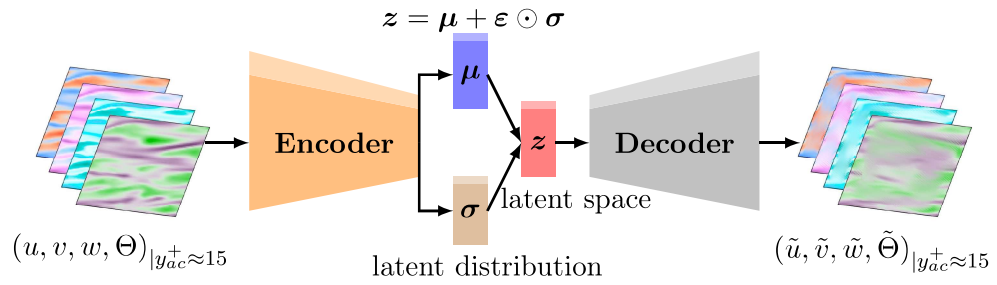

Figure 3.8: VAE representation. $\boldsymbol{\varepsilon}$ is sampled from a standard Gaussian distribution of zero mean and identity matrix for covariance. $\boldsymbol{z}$ is obtained via a linear combination of the predicted mean $\boldsymbol{\mu}$, and an element-wise product of the sampled $\boldsymbol{\varepsilon}$ and the predicted standard deviation $\boldsymbol{\sigma}$.

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

Directly sampling $\boldsymbol{z}$ from a Gaussian distribution would render both the latent variables and the encoder weights non-differentiable, preventing the use of backpropagation. To address this, the reparameterization trick is employed. Instead of sampling $\boldsymbol{z}$ directly, the encoder outputs the mean $\boldsymbol{\mu}$ and standard deviation $\boldsymbol{\sigma}$, and $\boldsymbol{z}$ is computed as:

$$\boldsymbol{z} = \boldsymbol{\mu} + \boldsymbol{\varepsilon} \odot \boldsymbol{\sigma}, \tag{3.6}$$

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

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

$$p_{\boldsymbol{\theta}}(\boldsymbol{x}) = \int p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{z}) p_{\boldsymbol{\theta}}(\boldsymbol{z})\, d\boldsymbol{z}, \tag{3.7}$$

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

Using Bayes' theorem and the logarithmic properties, the following inequality can be derived (Kingma and Welling, 2019):

$$\log(p_{\boldsymbol{\theta}}(\boldsymbol{x})) \geq -D_{\mathrm{KL}}[q_{\boldsymbol{\psi}}(\boldsymbol{z}|\boldsymbol{x})||p_{\boldsymbol{\theta}}(\boldsymbol{z})] + \mathbb{E}_{q_{\boldsymbol{\psi}}(\boldsymbol{z}|\boldsymbol{x})}[\log(p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{z}))], \tag{3.8}$$

where D_{KL} denotes the Kullback-Leibler (KL) divergence, a measure of dissimilarity between two probability distributions. This equation, known as the Evidence Lower Bound (ELBO), provides a computable lower bound on the log-marginal likelihood of the data. By maximizing the ELBO (the right hand side), the log-marginal likelihood of generating real samples, is indirectly maximized. To achieve this, one must minimize the KL divergence between the approximate posterior $q_{\boldsymbol{\psi}}(\boldsymbol{z}|\boldsymbol{x})$ and the prior distribution $p_{\boldsymbol{\theta}}(\boldsymbol{z})$, while simultaneously maximizing the expected log-likelihood of generating $\boldsymbol{x}$ given $\boldsymbol{z}$. The second term in the ELBO is equivalent to minimizing the reconstruction error (MSE) used in the CAE loss function (Eq. (3.5)).

The prior distribution $p_{\boldsymbol{\theta}}(\boldsymbol{z})$, which regularizes the latent space and is used in the KL divergence term, is here set to a standard Gaussian distribution $\mathcal{N}(\mathbf{0}, \boldsymbol{I})$. Not only is this prior commonly used due to its simplicity and mathematical tractability, but it also enforces the latent space to be dense and orthogonal, thereby reducing correlations between the latent variables. Indeed, minimizing the KL divergence between $p_{\boldsymbol{\theta}}(\boldsymbol{z})$ and $q_{\boldsymbol{\psi}}(\boldsymbol{z}|\boldsymbol{x})$ indirectly minimizes the KL divergence between $p_{\boldsymbol{\theta}}(\boldsymbol{z})$ and the marginal latent distribution $q_{\boldsymbol{\psi}}(\boldsymbol{z})$ (see Hoffman and Johnson (2016)). Strong regularization thus ensures the latent distribution remains close to the isotropic standard Gaussian, encouraging each dimension to encode more independent factors of variation. This reduced correlation between latent variables thus enhances model interpretability. A dense latent space also facilitates smoother interpolations, which is particularly valuable for sample generation tasks. The combination of these characteristics makes the standard Gaussian prior especially suitable for VAEs in both representation learning and generative modeling.

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

$$D_{\mathrm{KL}}[q_{\boldsymbol{\psi}}(\boldsymbol{z}|\boldsymbol{x})||p_{\boldsymbol{\theta}}(\boldsymbol{z})] = -\frac{1}{2}\left\langle \sum_{i=1}^{L_d} \left(\log\left(\sigma_{j,k}(i)^2\right) + 1 - \mu_{j,k}(i)^2 - \sigma_{j,k}(i)^2\right) \right\rangle_{j,k}, \tag{3.9}$$

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

The β-VAE introduces a real number $\beta \in \mathbb{R}^+$, which acts as a parameter to balance the trade-off between latent space orthogonality and reconstruction quality. Given the MSE loss defined in Eq. (3.5), the following loss function is obtained for β-VAEs:

$$\mathrm{Loss}_{\mathrm{VAE}}(\chi,\ \boldsymbol{\mu},\ \boldsymbol{\sigma},\ \beta) = \mathrm{MSE}(\chi) + \beta \times D_{\mathrm{KL}}[q_{\boldsymbol{\psi}}(\boldsymbol{z}|\boldsymbol{x})||p_{\boldsymbol{\theta}}(\boldsymbol{z})]. \tag{3.10}$$

This formulation allows for optimization of both the reconstruction accuracy and the latent space orthogonality.

Beyond their improved orthogonality and interpolation properties, β-VAEs often exhibit an automatic mode pruning effect (Kang et al., 2022). The regularization term in Eq. (3.10) forces some of the latent distributions to collapse toward the standard Gaussian prior, effectively over-regularizing these variables and minimizing their contribution to reconstruction. Building on this observation, Kang et al. (2022) developed a two-stage approach: (i) a ranking algorithm that assesses latent variable contribution through KL divergence analysis, and (ii) a pruning procedure that eliminates non-informative LVs while preserving reconstruction quality. This combined framework applied to β-VAEs can provide a method for determining optimal latent dimensionality, which is particularly valuable for dimension reduction tasks.

3.3.4 Metrics

To compare the performance of various models with different loss functions, it is necessary to define metrics that allow meaningful evaluations. These metrics can be separated in three categories: (i) reconstruction of the four fields, (ii) latent variables orthogonality, and (iii) latent samples separability with regards to flow parameters.

For the first category, the metric of interest is the per-field energy reconstruction (Eivazi et al., 2022; Solera-Rico et al., 2024), which measures the energy percentage captured by the model. It is particularly relevant to evaluate the reconstruction of each variable separately, as they are not all reconstructed with the same accuracy, a point that will be discussed in detail in Section 3.4.

Mathematically, the metric is defined as the squared difference between the fluctuations of input and reconstructed fields, normalized by the squared norm of the input fluctuations. This metric for a field $\chi^{(i)}$, using the same notation as in Section 3.2.4, is expressed as:

$$E_r(\chi^{(i)}) = 100 \times \left(1 - \left\langle \frac{\sum_{l=1}^{2048} [\chi^{(i)}_{j,k,l}{}' - \widetilde{\chi^{(i)}_{j,k,l}}']^2}{\sum_{l=1}^{2048} [\chi^{(i)}_{j,k,l}{}']^2} \right\rangle_{j,k} \right), \tag{3.11}$$

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

To evaluate the orthogonality of the latent space, a correlation matrix M is constructed. The correlation coefficient between two components m and n of the latent vector $\boldsymbol{z}$ is defined as:

$$M_{m,n} = \frac{C_{mn}}{\sqrt{C_{mm} C_{nn}}} \quad \text{for } 1 \leq n, m \leq L_d. \tag{3.12}$$

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

$$C_{mn} = \left\langle z'_{j,k}(m) z'_{j,k}(n) \right\rangle_{j,k} , \tag{3.13}$$

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

Technically, the correlation matrix contains all the information about the orthogonality (or correlation) of each mode with the others. However, extracting a scalar metric from this matrix to represent orthogonality is often more practical, especially for comparing different models. A common approach is to calculate the determinant of the correlation matrix (Eivazi et al., 2022; Solera-Rico et al., 2024), which equals 1 in the case of perfect orthogonality, corresponding to an identity matrix. However, with high-dimensional latent spaces, this metric becomes less reliable, as even modes that are nearly orthogonal can result in a determinant close to zero. This limitation is evident in the current study, affecting all AE models as well as the extended POD method (Borée, 2003).

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

1. Mean Correlation (Mean_c): The mean correlation measures the average off-diagonal coefficients of the correlation matrix in absolute value, providing an overall indication of how correlated the latent modes are with each other. Based on the definition of the covariance coefficients (Eq. (3.12)) taken in absolute value, it is reasonable to express this metric as a percentage:

$$\text{Mean}_c = 100 \times \langle |M_{i,j}| \rangle_{1 \leq i,j \leq L_d,\ i \neq j}. \tag{3.14}$$

   While Mean_c provides a good overview of the latent space orthogonality, it can sometimes dilute the impact of a few highly correlated modes due to averaging.

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

$$\text{MeanMax}_c = 100 \times \Big\langle \max_{1 \leq j \leq L_d,\ j \neq i} (|M_{i,j}|) \Big\rangle_{1 \leq i \leq L_d}. \tag{3.15}$$

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

Finally, there remains to evaluate the separability of the latent samples with regards to flow parameters. In simple flows, for which the data can be compressed into a space of dimension 2-3, visualizing the latent space allows for straightforward

evaluation of this separability. However, with the present turbulent flow dataset requiring much larger latent dimensions, direct visualization is impossible, necessitating the use of quantitative metrics to assess this aspect effectively.

The chosen approach consists of selecting all samples corresponding to a specific amplitude (or phase) within the test dataset, and to create a cluster with these samples, defined by its centroid and contained data samples. From these clusters, one can easily apply any clustering metric, taking into account intra or inter cluster variances. A relevant metric that takes into account both of these aspects is the Davies-Bouldin indicator (Xu and Tian, 2015), which is calculated as the maximal ratio of the averaged distance between a point and its centroid and the distance between two centroids, averaged over all clusters, i.e.

$$D_B = \left\langle \max_{m \neq n} \left(\frac{\delta_n + \delta_m}{\mathrm{dist}(\boldsymbol{c}_n, \boldsymbol{c}_m)} \right) \right\rangle_n . \tag{3.16}$$

Here, δ_n is the average distance of all the points contained in $\mathcal{C}_n$, cluster of index n, to the centroid $\boldsymbol{c}_n$. This cluster intra-variance is estimated as $\delta_n = \frac{1}{|\mathcal{C}_n|} \sum_{\boldsymbol{z_{j,k}} \in \mathcal{C}_n} \mathrm{dist}(\boldsymbol{z_{j,k}}, \boldsymbol{c}_n)$ where $\boldsymbol{z_{j,k}}$ corresponds to a latent sample and $\mathrm{dist}(\boldsymbol{x}, \boldsymbol{y})$ measures the euclidean distance between vectors $\boldsymbol{x}$ and $\boldsymbol{y}$.

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

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

$$D_B(\mathcal{C}_n, \mathcal{C}_m) = \frac{\delta_n + \delta_m}{\mathrm{dist}(\boldsymbol{c}_n, \boldsymbol{c}_m)}. \tag{3.17}$$

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

All pre- and post-processing operations were handled using NumPy[2] and scikit-learn[3].

3.4 Comparing CAE with POD

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

[2] NumPy `1.24.2`
[3] scikit-learn `1.2.0`

	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%

Table 3.3: Comparison of energy reconstruction for each field, where the CAE is trained either to reconstruct all fields simultaneously with $L_d = 500$ (Together) or to reconstruct them individually using 125 modes for each (Separately).

Table 3.3 highlights the superior reconstruction accuracy achieved by the CAE when trained on all fields simultaneously within a 500-dimensional latent space, compared to independent field training using 125-dimensional latent spaces per field. This comparison is dimensionally consistent, as the cumulative latent dimensionality for individual field processing also totals 500. The simultaneous compression approach improves reconstruction accuracy for the u, v, and Θ fields by 8%, 16% and 12% respectively. This enhancement leverages the strong physical correlations between fields, particularly between the streamwise velocity (u) and temperature (Θ) fields, as expected from the Reynolds analogy. Notably, the wall-normal velocity component (v) consistently exhibits lower reconstruction accuracy across both approaches, a limitation that warrants further investigation. Considering these results, dimensionnality reduction models are systematically applied to the four fields together instead of separately for the rest of this study.

Figure 3.9 compares the reconstructed-energy percentage of POD and CAE across various latent dimensions, for each field. The results show that the CAE consistently outperforms extended POD for all L_d-values. Extended POD only matches the performance of CAE on v when $L_d \geq 1\,000$, with both methods performing similarly for larger dimensions. For each field, the performance gap narrows as L_d increases, with the greatest difference observed at the lowest value of L_d analysed, i.e. 250. Notably, this gap is most pronounced for v, where a difference of over 50% in E_r is observed at $L_d = 250$, highlighting the advantage of the nonlinear CAE approach, particularly for smaller latent dimensions. Another major result is the number of modes for which each model reaches the 90% reconstruction accuracy benchmark. Excluding the v-field, where both models reach this target around $L_d = 1\,200$, extended POD requires 2 to 3 times more modes than the CAE to achieve 90% accuracy. This underscores the CAE's superior ability to compress information, as it reaches comparable or better accuracy with significantly fewer modes, making it better suited for this dataset.

As noted earlier, the v-field consistently achieves lower E_r scores than the other three fields. Figure 3.9 shows that for $L_d = 250$, $E_r(u)$, $E_r(w)$, and $E_r(\Theta)$ all exceed 80% for the CAE, while $E_r(v)$ remains around 65%. As mentioned previously, this issue is even more pronounced for POD, further highlighting the challenges specific to the v-field. This is at least partially attributed to the smaller turbulent

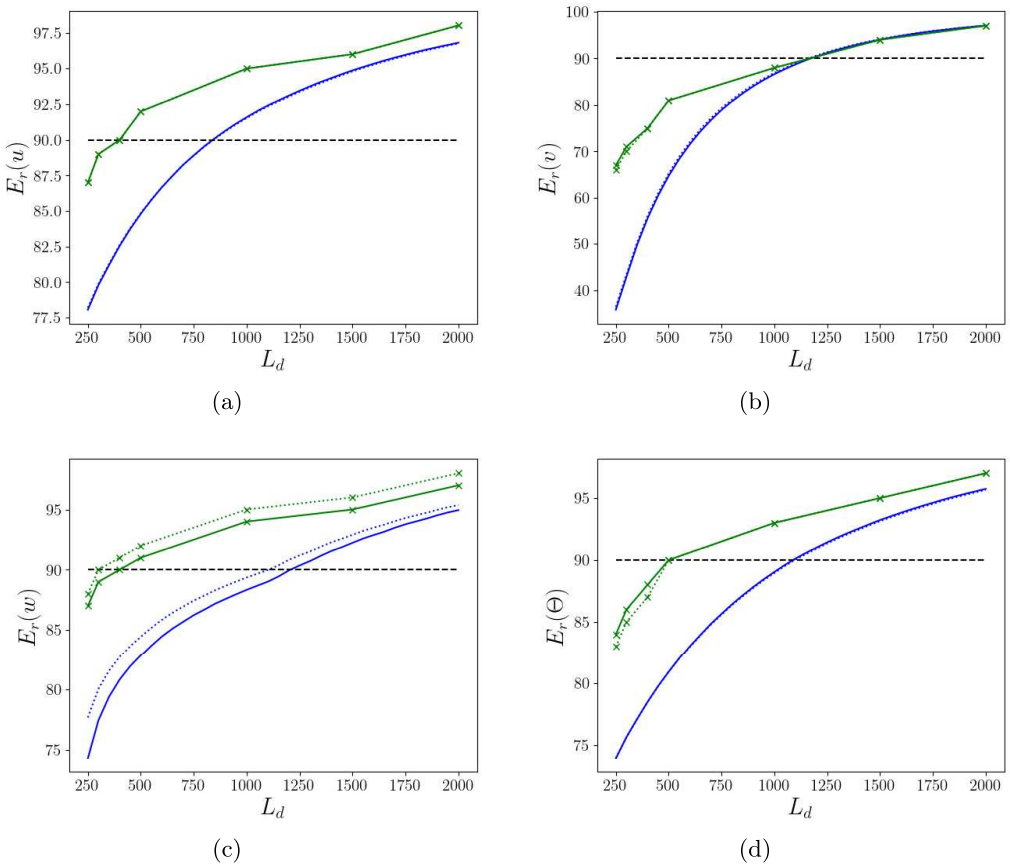

Figure 3.9: Energy reconstruction scores (E_r) as a function of the latent dimension (L_d) for extended POD (blue) and CAE (green). Solid lines indicate the training data results, while dotted lines highlight scores on the test data. The fields correspond to the following: (a) Streamwise velocity u, (b) Wall-normal velocity v, (c) Spanwise velocity w, and (d) Temperature Θ.

scales contained in the v-field compared to the others, which require more modes to be captured. The v-component is also initially less energy-dense compared to the other fields, and the applied scaling may not have fully remedied this discrepancy. Additionally, the long tail observed in the scaled v-distribution (Figure 3.4(b)) does not contribute significantly to this disparity, as a pre-scaling truncation of this tail did not improve the reconstruction accuracy. Another possible explanation is that SWO enhance sweeps and ejections in the wall-normal direction, as attested in Section 2.5, further expanding the v-distribution in the actuated case.

Both models demonstrate robust generalization capabilities, as evidenced by the comparable behaviour of training and testing curves across all flow fields. This characteristic was specifically evaluated through the train-test partition detailed in Section 3.2.4. The observed higher test scores compared to training scores for the w field are likely attributable to the disparity in dataset sizes, with the training set containing simulations at 17 distinct W^+ values, while the test set comprises only 4

W^+ values (see Figure 3.2). This imbalance in sampling density may lead to better generalization performance on the sparser test dataset.

	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

Table 3.4: 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 3.5).

The selection of latent dimension $L_d = 500$ as the principal case investigated is established to guarantee that the CAE maintains a minimum reconstruction accuracy of 80% across all four target fields when assessed against the testing dataset. The comparative analysis at $L_d = 500$ reveals the superior performance of the CAE relative to extended POD when utilizing an identical number of modes. Specifically, the CAE demonstrates enhanced reconstruction efficiency with improvements of 7%, 16%, 8%, and 9% in E_r for the u, v, w, and Θ fields, respectively (see Table 3.4). These quantitative results provide substantial evidence for the enhanced capabilities of the non-linear autoencoder methodology compared to its linear counterpart.

The spatial reconstructions depicted in Figure 3.10 present comparative visualisations across the complete LES domains, incorporating sliding averages to maintain subdomain continuity (see Figure 3.3). The comparative analysis encompasses Ground Truth representations (left column), CAE reconstructions (centre column), and extended POD reconstructions (right column) for each flow field. These snapshots are picked from the simulation of the test dataset corresponding to the highest oscillation amplitude: $W^+ = 24$. This simulation is characterized by its important small-scale fluctuations and by the largest gap in performances when comparing CAE and POD. The Ground Truth data exhibits pronounced, spatially intermittent fluctuations of significant amplitude, characteristic of turbulent small-scale structures. These fine-scale turbulent manifestations present significant reconstruction challenges, necessitating the implementation of 500 modes to achieve satisfactory accuracy.

The v-field exhibits particularly pronounced small-scale structural elements, presenting substantial reconstruction challenges for both methodologies. Its distribution demonstrates a notable centralised tendency (see Figure 3.4), which manifests in reduced reconstruction emphasis from both approaches. The extended POD reconstruction demonstrates adequate representation of large-scale flow features. However, the reconstructed flow structures exhibit considerable "smoothing" relative to the Ground Truth, with numerous small-scale features absent from the reconstruction.

Visual evidence of the ongoing challenge faced by POD to adequately reconstruct small-scale fluctuations in the flow is provided by the Power Spectral Densi-

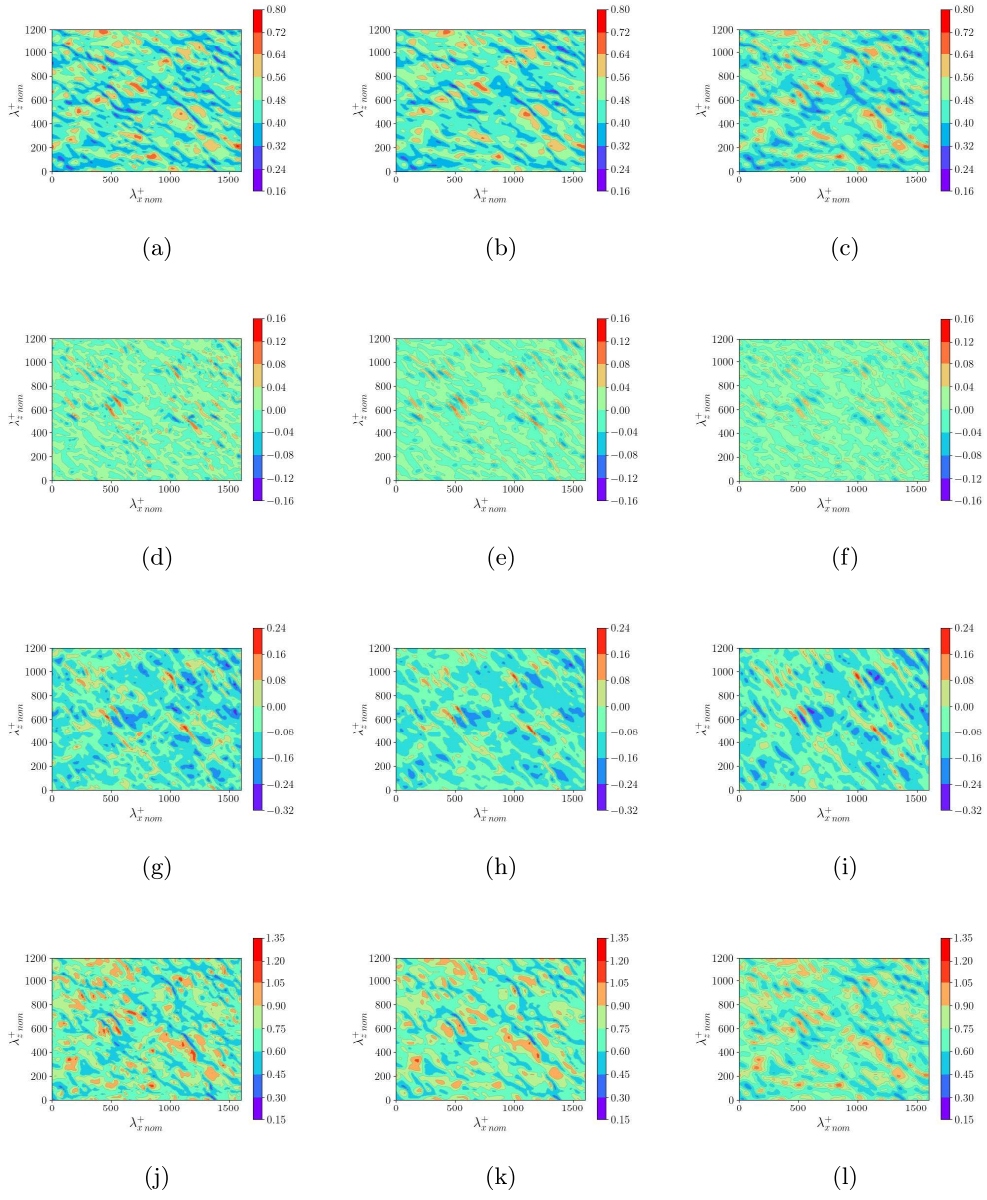

Figure 3.10: Reconstruction of snapshots with $L_d = 500$ for the maximal amplitude simulation within the test dataset ($W^+ = 24$): (Left) Ground truth, (Middle) CAE reconstruction, and (Right) extended POD reconstruction. The four fields are displayed from top to bottom: u, v, w, and Θ. Each snapshot represents the full LES domain.

ties (PSDs) in the spanwise direction represented in Figure 3.11 for all components. Pre-multiplied PSDs are compared for the ground truth data (black solid lines) of the test dataset simulation with oscillation amplitude $W^+ = 24$, CAE reconstruction (green dashed lines) and extended POD reconstruction (blue dotted lines). For all fields, the blue curve remains beneath the black curve for all flow scales, with the gap narrowing for the larger scales of the flow, resulting in POD not capturing

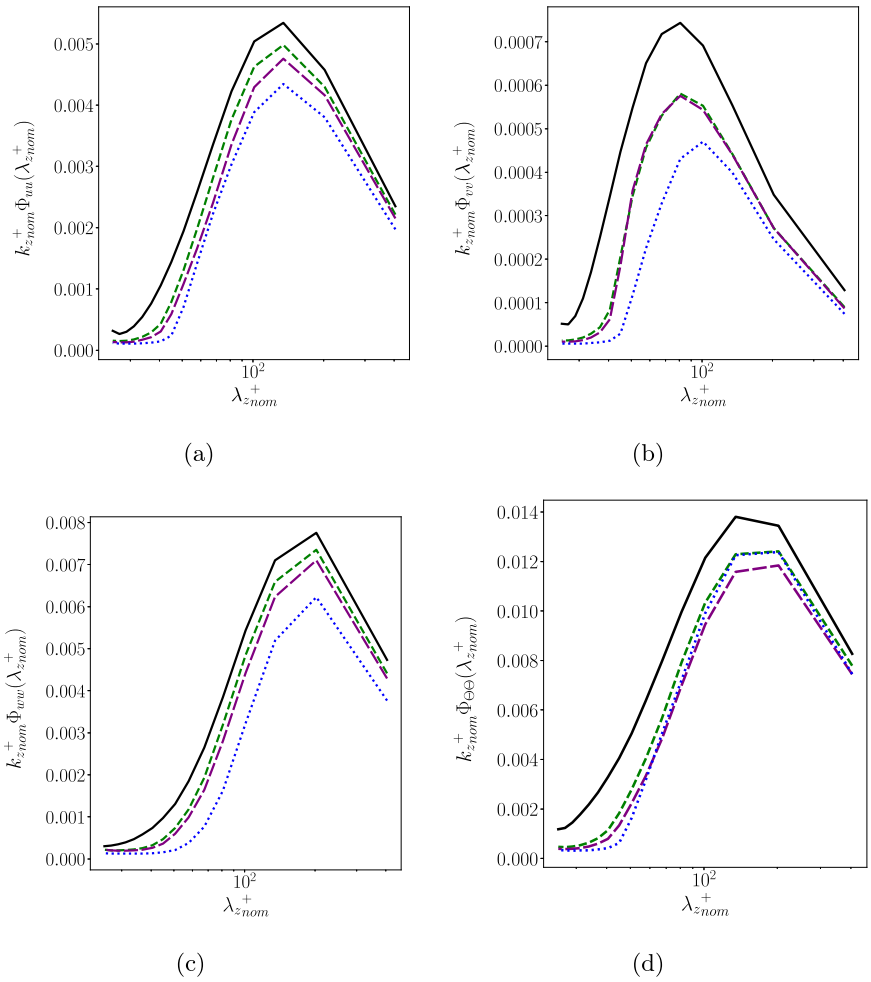

Figure 3.11: Pre-multiplied Power spectral density $k_z\Phi_{\chi^{(i)}\chi^{(i)}}(\lambda_z)$ of (a) u, (b) v, (c) w, (d) Θ, for the maximal amplitude simulation within the test dataset ($W^+ = 24$). The spectras are determined on the subdomains of Figure 3.3 and not on the full LES-domain. Black solid 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.

the total energy of the system. The model particularly fails to capture the smaller scales of the flow, which is further confirmed by the flat left tail of the blue spectras. As expected, this trend is even more pronounced on v (see Figure 3.11(b)), which is characterized by a much larger energy gap compared to the other fields. This behaviour is expected on the wall-normal velocity field, as an important percentage of its energy is concentrated within its smaller structures. These results are consistent with the theoretical framework of POD, as the lower-energy modes associated with smaller scales are excluded when restricting the reconstruction to the first modes. Additionally, the reconstructed snapshots indicate that extended POD demonstrates limitations in reproducing extreme values within the reconstructed structures, exhibiting a tendency towards mean-value approximation. These constraints show the

inherent limitations of linear methodologies when applied to turbulent flows, where small-scale dynamics are fundamental to accurate flow representation.

The CAE reconstruction methodology demonstrates enhanced capability in capturing localised variations (see Figure 3.10), preserving these features with increased fidelity. The reconstruction quality is substantial across the majority of fields, with superior representation of small-scale non-linear features compared to the POD approach. This difference in model performances is confirmed by the spectral analysis of Figure 3.11, in which the CAE consistently reconstructs all the flow scales up to $\lambda_{z^+_{\text{nom}}} \approx 200$ more accurately than POD. Intriguingly, the POD performances are closer to those of the CAE on the temperature field, as they provide similar estimations for wavelengths $\lambda_{z^+_{\text{nom}}} > 50$. The CAE demonstrates significantly lower performance in the v field in comparison to the other components, primarily due to the inherent small-scale characteristics of this field. Nevertheless, the wall-normal velocity is characterised by the largest discrepancy between the POD and CAE spectra, further demonstrating the superiority of the latter approach in reconstructing turbulent flows. These improvements are attributed to the architectural implementation of convolutional layers (see Section 3.3.2) in conjunction with non-linear activation functions, specifically designed for the detection and preservation of localised, non-linear features in spatial data structures. Despite the generally satisfactory reconstruction performance, regions characterised by significant noise, continue to present reconstruction challenges for the CAE methodology.

The CAE demonstrates superior performance in reconstruction relative to extended POD. However, it exhibits suboptimal results in terms of latent orthogonality (see Table 3.4). Although extended POD is not strictly orthogonal (as explained in Section 3.3.1), it achieves mean correlation scores (Mean_c) that are three times lower than those of the CAE, and its modes are, at most, less than half as correlated as the CAE latent variables (MeanMax_c). An ideal model should achieve the high reconstruction accuracy of the CAE while producing correlation scores closer to those of extended POD, allowing for more interpretable modes. This balance will be explored further in the next section.

3.5 Balancing latent orthogonality and reconstruction accuracy with β-VAE

In this section, the orthogonality of the latent space is examined by sampling latent variables in a manner that encourages independence and proximity to the standard Gaussian distribution, as opposed to using a deterministic, unstructured approach (see Section 3.3.3). More precisely, β-VAEs are employed to determine whether a suitable compromise can be achieved between the orthogonality of extended POD and the reconstruction accuracy of the CAE.

A parametric study on β showed that the value 5×10^{-7} led to the best trade-off between reconstruction and orthogonality. This value was subsequently retained.

Table 3.4 highlights the obtained trade-off. Even with a relatively small β value, the reconstruction accuracy across all fields experiences a slight decrease, underscoring the significant impact of the latent space structure on energy reconstruction and the difficulty of finding an appropriate balance. The latent orthogonality is largely enhanced comparatively to the CAE, with a decrease of 1% and 2.8% in Mean_c and MeanMax_c, respectively. Interestingly, this reduction of high-correlation relationships in the latent space impacts the reconstruction of v and Θ fields the most, as attested by Table 3.4. This suggests that the improvements in latent space orthogonality may come at the expense of reconstruction accuracy for these fields, which is likely linked to the high correlation among all four fields.

Figure 3.11 demonstrates that, like the CAE, the β-VAE (purple dash-dotted line) outperforms extended POD in reconstructing the flow, particularly for the smaller turbulent scales and at the energy peak. However, this is not the case for the temperature field, where extended POD achieves better accuracy at the energy peak, aligning with the previously noted decline in β-VAE performance for this field. While the β-VAE performs marginally worse than the CAE in flow scale reconstruction overall, it matches the CAE's performance in the v field for this specific simulation.

The mode ranking algorithm developed by Kang et al. (2022) was applied following the completion of model training. As detailed in Section 3.3.3, latent variables are sorted according to their Kullback-Leibler divergence from the one-dimensional standard Gaussian prior, wherein higher divergence values are indicative of greater information content. Through the regularisation term in Eq. (3.10), latent variables with minimal KL divergence are forced towards the standard Gaussian distribution. As a direct consequence, these over-regularised variables are rendered functionally redundant, contributing negligibly to the reconstruction process. Dimensional reduction through mode pruning is thus mathematically justified, whereby a systematic excision of redundant modes is performed based on their information content. In this pruning procedure, the most informative latent variables (classified as physics-aware in Kang et al. (2022)) are preserved within the reduced-order representation, while the less informative components (physics-unaware) are systematically eliminated. This dimensional truncation is implemented by fixing the physics-unaware modes to their respective estimated mean values (Eq. (3.18a)), effectively neutralising their contribution to the flow field reconstruction.

The KL divergence is computed using Eq. (3.9) (applied individually to each latent variable) and the empirically estimated parameters $\widehat{\mu}_n$ and $\widehat{\sigma}_n$ determined as

$$\widehat{\mu}_n = \langle z_{j,k}(n) \rangle_{j,k}, \tag{3.18a}$$

$$\widehat{\sigma}_n = \sqrt{\langle (z'_{j,k}(n))^2 \rangle_{j,k}}. \tag{3.18b}$$

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

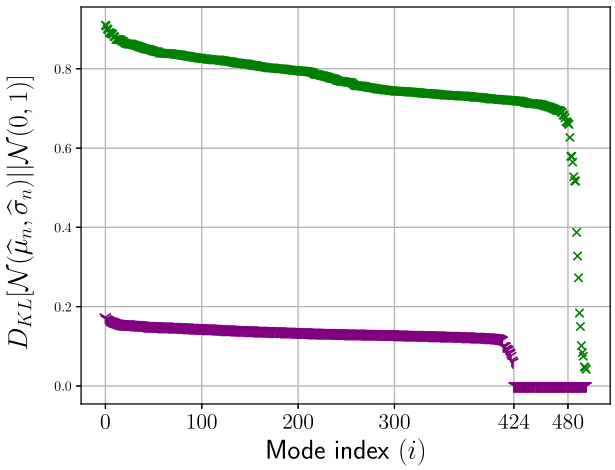

Figure 3.12: KL divergence of each latent variable with the 1D standard Gaussian prior, for the CAE (green crosses) and β-VAE (purple Ys), with $L_d = 500$ and ranking of the modes with regards to their KL divergence score.

Figure 3.12 displays the KL divergence of each latent variable, ranked in descending order, with CAE results marked by green crosses and β-VAE results by purple Ys. The CAE modes exhibit a KL divergence roughly four times higher than those of the β-VAE in average, highlighting the efficacy of the regularization term in Eq. (3.10). A sharp decline in KL divergence beyond latent variable 424 of the β-VAE separates the physics-aware modes from the physics-unaware ones, suggesting the presence of the pruning effect commonly observed in VAEs. Here, 76 latent distributions are effectively reduced to a standard Gaussian, contributing negligibly to reconstruction. A similar, though less pronounced, KL drop occurs for the CAE at mode 480. However, even the least informative CAE latent variables retain a KL divergence of approximately 0.04, preventing it from being classified as physics-unaware.

Following the methodology of Kang et al. (2022), a physics-informed β-VAE can be derived by retaining only the 424 modes with the highest KL divergence. Table 3.5 confirms that the pruned model achieves reconstruction performance identical to the full model, validating the pruning effect induced by the β-VAE regularization term.

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

Table 3.5: 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.

Before comparing the models in terms of latent orthogonality, it is important to note that the mean correlation (Mean_c) should be interpreted with caution, as it may underestimate the presence of highly correlated modes. Figure 3.13(a) shows the correlation matrix of the CAE, with LVs ranked according to their KL divergence. The last 10 to 20 modes, corresponding to the KL divergence drop of Figure 3.12, exhibit nearly perfect correlation with each other and are significantly correlated with the remaining modes compared to the average behavior. However, these LVs have a limited effect on the overall Mean_c score, which remains low (2.1%, as shown in Table 3.4). This highlights the potential of Mean_c to obscure redundancies within the latent space. By contrast, the mean-max correlation (MeanMax_c) offers a more sensitive perspective. This metric captures the maximum correlation for each mode, making it better suited to identify hidden redundancies. For the CAE, the MeanMax_c score (9.7%) reveals significant entanglement among modes that the Mean_c score fails to reflect. The CAE latent space demonstrates clear entanglement, especially for the least contributing LVs, a significant limitation of this architecture. This entanglement implies that features detected by its convolutional layers lack independence and interpretability. Consequently, the CAE risks encoding redundant features, where similar patterns are repeatedly captured across different modes.

Although the extended POD spatial basis is orthogonal, when validating the results on the test data, the temporal coefficients do not necessarily form an orthogonal basis, as illustrated in Figure 3.13(c), which is generated from the temporal coefficients of each sample. The ranking algorithm of Kang et al. (2022) is not applied in this case, as POD already organizes its modes by decreasing energy contribution. This ranking leads to modes with similar energy levels being more likely to exhibit correlation due to their tendency to describe similar flow structures. This pattern is reflected in the correlation matrix, where coefficients closer to the diagonal show stronger correlations, while those further away tend to be less correlated, as indicated by the light green regions in the bottom-left and top-right corners. Interestingly, a few secondary blue diagonal lines, both near and distant from the matrix main diagonal, indicate higher than average correlations between modes with significantly different energy contributions. Unlike the CAE, the extended POD does not exhibit dark blue dots outside the diagonal, implying that no mode is highly correlated with another. This is consistent with the MeanMax_c score of 4.1% (see Table 3.4). As a result, the temporal coefficients of extended POD can be considered near-orthogonal.

The β-VAE correlation matrix (Figure 3.13(b)) demonstrates a notable improvement over that of the CAE. Overall, the matrix is characterized by larger green regions, indicating that the correlation coefficients of these modes are significantly damped. This improvement is further validated in Table 3.4, where the mean correlation (Mean_c) drops from 2.1% for the CAE to 1.1% for the β-VAE. The 76 physics-unaware LVs are clearly visible in the correlation matrix, represented by light green lines that highlight their purely stochastic behavior. Similar to the CAE correlation matrix, the least contributing modes (those just above the pruned

threshold) show high mutual correlation, though to a much lesser degree. This observation suggests that orthogonality improvement is not merely a global average effect but occurs at multiple levels. Supporting this conclusion, the β-VAE's mean max correlation (MeanMax_c) drops to 6.9%, nearly 3 percentage points lower than the CAE's MeanMax_c of 9.7%.

However, one might question to what extent the β-VAE's superior orthogonality stems from its physics-unaware modes, exhibiting very low correlations (see Figure 3.13(b)). To investigate this, Table 3.5 compares the pruned β-VAE with a CAE using $L_d = 400$ (a relevant comparison since the pruned β-VAE utilizes 424 latent variables). While both models achieve comparable reconstruction quality (with the CAE showing slightly better scores), the pruned β-VAE maintains significantly better orthogonality scores, particularly in terms of MeanMax_c. This demonstrates that the β-VAE's regularization achieves more than simple mode pruning, establishing an effective trade-off between reconstruction fidelity and latent space orthogonality.

A comparison of the β-VAE correlation matrix with the extended POD matrix reveals that the latter model continues to demonstrate superior performance with respect to mode orthogonality. The results presented in Tables 3.4 and 3.5 substantiate this conclusion. As with the extended POD correlation matrix, discernible patterns are evident in the β-VAE matrix, albeit of a markedly different nature. Indeed, when excluding the physics-unaware LVs, the overall correlation of each remaining LV with other modes appears to decrease as its KL divergence increases.

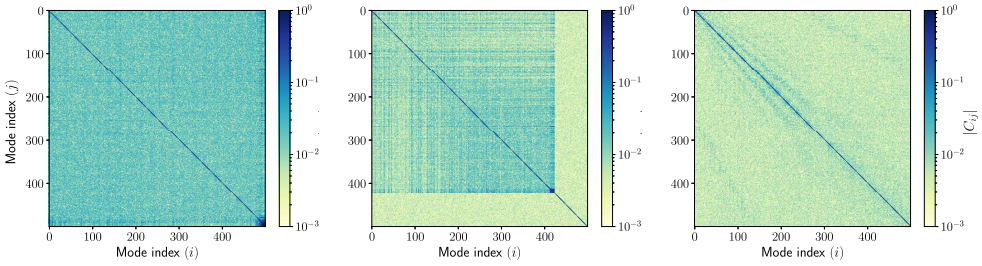

(a) (b) (c)

Figure 3.13: 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).

The results demonstrate that the β-VAE achieves a balanced compromise between mode orthogonality, effective W^+ interpolation, and field reconstruction, while not outperforming either the CAE or extended POD in their specific strengths. By applying the mode ranking algorithm from Kang et al. (2022), the β-VAE can be successfully reduced to a model with $L_d = 424$, while maintaining a fair trade-off between reconstruction quality and orthogonality. These results highlight the inherent challenges in simultaneously achieving both orthogonal latent variables and precise reconstruction accuracy for complex datasets requiring high latent dimensions.

3.6 Separability of generative factors within the latent space

The interpretability of the latent space extends beyond the orthogonality of modes and their ranking 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 2.4 for their definitions).

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

Table 3.6: Global Davies-Bouldin indicator as given by Eq. (3.16) 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 3.5).

Table 3.6 presents the Davies-Bouldin indicator (Eq. (3.16)) 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 Figure 3.14 through the Davies-Bouldin indicator for two clusters W_i^+ and W_j^+ as given by Eq. (3.17).

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 lower coefficients are featured in the CAE matrix, indicating superior overall W^+ separability, a result consistent with Table 3.6.

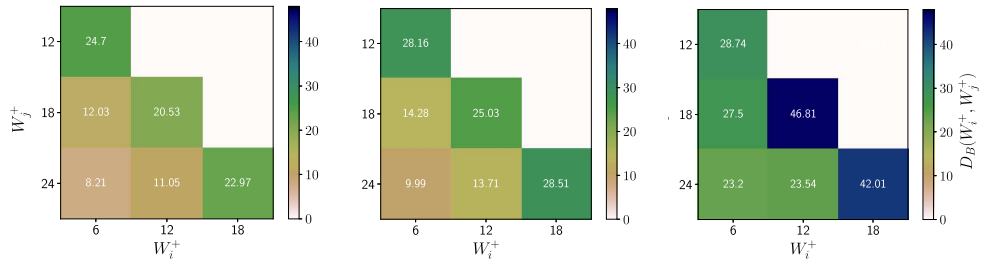

(a) (b) (c)

Figure 3.14: 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.

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

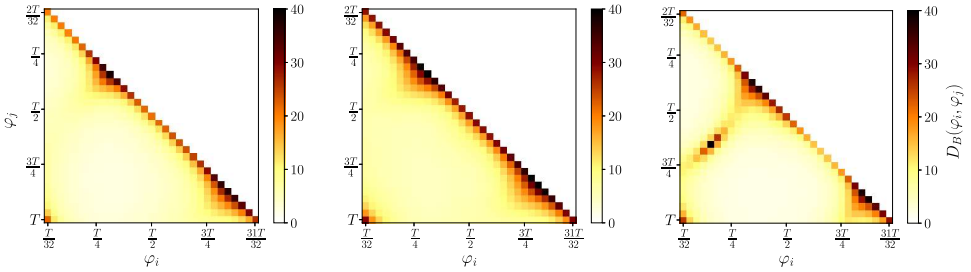

(a) (b) (c)

Figure 3.15: 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.

As detailed in Section 3.2.2, each actuation cycle is decomposed into 32 distinct phases in the dataset. All simulations, in the test dataset, feature actuation ($W^+ > 0$), exhibiting periodic behaviour that justifies phase-based analysis. The dependence on the phase of the oscillation cycle becomes more prominent with higher values of W^+. The separability analysis protocol is applied in Figure 3.15 to evaluate phase separability across models.

For the autoencoder approaches, similar trends are exhibited by both models, with the CAE consistently achieving substantially enhanced separability. A pattern analogous to the W^+ case is observed, where poorer separability is demonstrated by

temporally adjacent phases compared to distant phases. Notably, the highest intra-group D_B scores are exhibited by phases in the ranges $\varphi \in \left[\frac{T}{4} + \delta_\varphi, \frac{T}{2} - \delta_\varphi\right]$ and their symmetric counterparts $\varphi \in \left[\frac{3T}{4} + \delta_\varphi, T - \delta_\varphi\right]$, where δ_φ represents a small phase offset. These ranges, corresponding to the ascending drag regime (transitioning from $C_{f,\mathrm{min}}$ to $C_{f,\mathrm{max}}$), are distinguished by steeper temporal gradients in wall friction compared to other phases of the cycle, as identified in Section 2.5. Additionally, the phases $\varphi = \frac{T}{4}$ and $\varphi = \frac{3T}{4}$ correspond to the extremal values of the control signal $w_\mathrm{w}(t)$, where $w_\mathrm{w}\left(\frac{T}{4}\right) = W$ and $w_\mathrm{w}\left(\frac{3T}{4}\right) = -W$, which may explain the poor separability scores observed around these phases.

While the symmetric nature of friction variation arising from the sinusoidal wall oscillation (see Figure 2.6) might suggest difficulty in distinguishing phases separated by $\frac{T}{2}$, efficient separation is actually demonstrated in the autoencoder latent spaces for these cases. This suggests that other physical distinctions, such as Stokes strain sign reversal or streak orientation, are effectively captured by the autoencoders.

Distinct behaviour is exhibited by the extended POD approach, as evidenced by elevated D_B scores for phases φ near $\frac{T}{8}$ with their symmetric counterparts at $\varphi \approx \frac{5T}{8}$. This observation suggests that difficulties are encountered by the extended POD latent space in distinguishing these phase pairs, which correspond to similar drag and heat transfer characteristics but opposing streak orientations. More broadly, the presence of the corresponding non-diagonal line in Figure 3.15(c) reveals that phases $\varphi_1 \in \left[0, \frac{T}{4}\right]$ are not well separated from phases $\varphi_2 \approx \frac{3T}{4} - \varphi_1$. This implies that the extended POD methodology struggles to differentiate between friction-descending phase pairs φ_1/φ_2 originating from opposite halves of the actuation cycle and displaying nearly identical distances from $C_{f,\mathrm{min}}/C_{f,\mathrm{max}}$ phases. These phase pairs are characterised by substantially different streak orientations and opposite Stokes strain directions in the actuated case, which confirms the difficulty of extended POD to effectively capture related flow features within its latent space.

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

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

1. Superior performance in separating different W^+ amplitudes is consistently achieved by autoencoder architectures compared to extended POD, with the CAE achieving optimal capability. This suggests that amplitude-dependent flow features are better captured in the latent representations of autoencoders.

2. Similar phase separation characteristics are exhibited by both CAE and β-VAE, showing:

   - Reduced separability in ascending friction phases.
   - Superior overall performance achieved by the CAE.

- Effective discrimination of phases with different streak orientation or opposing Stokes strain direction.

3. The best global $D_B(\varphi)$ score is achieved by extended POD (see Table 3.6) and excellent separability is shown for certain phase combinations. However, weak separation of part of the descending friction phase-pairs is provided by this method, which differ significantly in terms of streak patterns and Stokes strain.

While extended POD is not matched by autoencoders in terms of latent orthogonality, more physically meaningful characterisations of the generative factors of the flow are provided by their learned representations, offering enhanced interpretability. The best performance overall was achieved by the CAE, especially with regard to oscillation phases. The approach to investigate these separabilities remains basic, considering one cluster for each phase or amplitude. Deeper insights into the separability of SWO amplitudes and phases, as well as mode correlations, may be provided by employing Cluster-based Reduced-Order Modelling (CROM) as proposed by Kaiser et al. (2014). Additionally, new information about the interactions between C_f and Nu, particularly concerning the breakdown of the Reynolds analogy, may be revealed. This extended cluster-based analysis remains a topic for future investigations.

3.7 Friction augmented CAE for enhanced latent physical coherence

Two fundamental quantities are of particular interest to the flow physics: the friction coefficient, C_f, and the Nusselt number, Nu. Within the prescribed oscillation period, SWO demonstrate notable capacity to enhance thermal transport efficiency without substantial drag penalties at appropriate amplitude settings. The inherent abstraction of neural networks traditionally obscures direct physical interpretation of their internal representations. The Augmented Convolutional AutoEncoder (A-CAE) model, inspired by the Lift Augmented Autoencoder of Fukami and Taira (2023), addresses this limitation through a sophisticated dual-objective framework. The system augments traditional autoencoding with a dedicated Multi-Layer Perceptron (MLP) network that derives C_f and Nu predictions from the latent representation, as illustrated in Figure 3.16. Optimizing both reconstruction accuracy and physical parameter prediction simultaneously introduces additional constraints on the organization of the latent space. This encourages the development of a better structured latent space with respect to C_f and Nu while preserving the computational benefits of dimensionality reduction. This integrated approach also establishes explicit connections between the compressed representation and two physical quantities of interest.

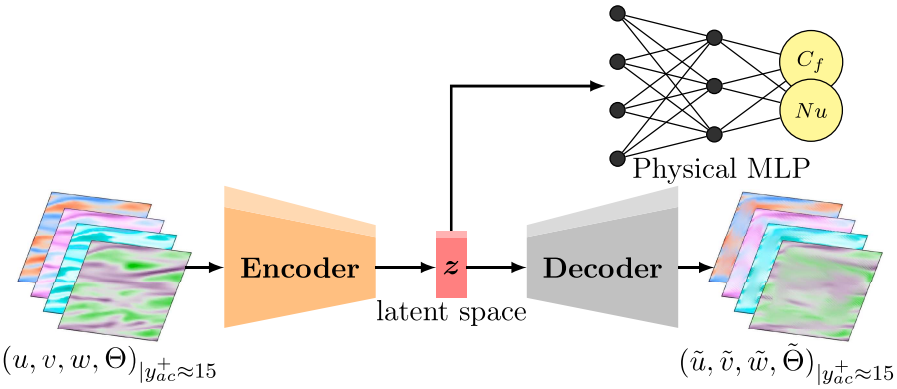

Figure 3.16: Friction-augmented-CAE representation: A physical MLP is incorporated into the CAE to predict the values of C_f and Nu from the latent vector z.

The field measurements, acquired at the height of maximum streak intensity and mixing phenomena, should provide relevant data for both Nu and C_f predictions. The architecture maintains the base convolutional structure described in Section 3.3.2, supplemented by an MLP comprising:

- Input layer: 500 neurons (matched to latent dimension).
- Hidden layers: 16, 8, and 4 neurons.
- Output layer: 2 neurons (Nu and C_f prediction).
- ELU activation functions throughout.

To address the scale disparity between Nu ($\mathcal{O}(10^1)$) and C_f ($\mathcal{O}(10^{-3})$), independent MinMax normalisation is applied to each parameter. The combined loss function incorporates both objectives:

$$\text{Loss}_{\text{A-CAE}}(\chi,\ C_f,\ \text{Nu},\ \alpha) = \text{MSE}(\chi) + \alpha \times \langle (C_{f\,j,k} - \widetilde{C_f}_{j,k})^2 + (\text{Nu}_{j,k} - \widetilde{\text{Nu}}_{j,k})^2 \rangle_{j,k}, \quad (3.19)$$

where $\widetilde{C_f}_{j,k}$ and $\widetilde{\text{Nu}}_{j,k}$ represent predicted values for sample k and simulation j, MSE(χ) denotes field reconstruction error (see Eq. (3.5)), and α functions as a weighting parameter balancing field reconstruction fidelity against prediction accuracy of physical quantities. This formulation enables the emergence of physically meaningful latent representations that capture both field dynamics and integrated flow characteristics.

Similar to the flow fields reconstruction, the evaluation of the C_f and Nu prediction follows a comparable principle of relative error. In this context, the coefficient of determination, R^2, is employed as a metric, as it is widely utilized in regression tasks to facilitate consistent comparisons. Considering $\xi \in \{C_f, \text{Nu}\}$, the indices j

and k denoting the simulations and samples as described in Section 3.2.4, and $\widetilde{\xi}_{j,k}$ representing the MLP prediction of $\xi_{j,k}$, R^2 is defined by the following equation:

$$R^2(\xi) = 100 \times \left(1 - \frac{\left\langle \xi_{j,k} - \widetilde{\xi}_{j,k} \right\rangle_{j,k}^2}{\langle \xi_{j,k} - \langle \xi_{j,k} \rangle_{j,k} \rangle_{j,k}^2} \right). \tag{3.20}$$

3.7.1 Improvements in Nusselt number and friction coefficient predictions

This section examines the ability of the A-CAE to reconstruct C_f and Nu. This provides insights into the potential for the latent space to be structured in a manner that more effectively distinguishes variations in these quantities. The primary objective is thus to ascertain whether the A-CAE is capable of embedding friction and heat transfer within the latent space. For this purpose, comparisons will be made with a standalone MLP trained to reconstruct C_f and Nu from the regular CAE latent samples. It is of particular interest to determine whether the integration of frictional information can be achieved without significantly compromising the reconstruction accuracy of the four fields, in a manner analogous to the β-VAE case. Additionally, it is essential to determine whether the latent samples are suitably separated based on these physical quantities.

	$R^2(C_f)$	R^2(Nu)
MLP test/training	49/66	64/78
A-CAE test/training	40/86	55/92

Table 3.7: Coefficient of determination (R^2) for the drag coefficient (C_f) and the Nusselt number (Nu), by a standalone MLP trained on the CAE latent space, and the A-CAE, both with $L_d = 500$. Results for the A-CAE are presented for an α value of 10^{-1}.

Table 3.7 compares the Nu and C_f reconstruction of the A-CAE with an α value of 0.1 (which yielded the best trade-off between Nu/C_f prediction and field reconstruction) and an MLP predicting Nu and C_f directly from the CAE with $L_d = 500$ latent space. The corresponding A-CAE field reconstruction score on the test dataset is $E_r(\chi) = [89, 72, 88, 85]$, demonstrating slighly reduced performance compared to the CAE (see Table 3.4). As in the β-VAE case, the v and Θ field are the most impacted by the regularization.

Comparison between both approaches in Table 3.7 demonstrates that the A-CAE manages to improve the C_f and Nu reconstruction (by 20% and 14% respectively) on the training dataset compared to the standalone MLP. This implies that the A-CAE approach successfully modified the latent space structure to contain more relevant information with regards to these physical quantities compared to the CAE baseline. However these gains do not translate to an improvement in friction prediction on

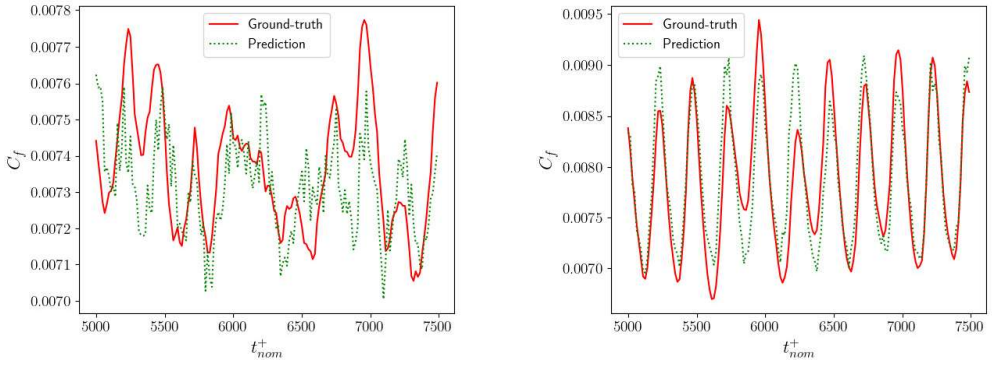

Figure 3.17: Temporal evolution of C_f, for the A-CAE with $\alpha = 0.1$, on the test dataset. (a) simulation with $W^+ = 6$ and (b) with $W^+ = 24$. Red solid lines represent the LES data and green dotted lines the A-CAE predicted values.

the test dataset, demonstrating the A-CAE struggles with generalization in terms of Nu and C_f predictions.

Both approaches actually exhibit a pronounced disparity between training and testing performance for C_f and Nu reconstructions, despite careful balancing of the training and testing datasets with respect to the distribution of W^+ values. Multiple regularisation strategies were implemented to address this generalisation challenge, including neural network dropout mechanisms, batch-normalization, and α variation for the A-CAE. However, these approaches proved ineffective in enhancing the model's generalisation capabilities. Therefore, while the results on the training dataset demonstrate the relevance of the A-CAE to enhance latent physical coherence with regards to Nu and C_f, this approach fails to effectively generalize this improvement on the test dataset.

Two principal hypotheses emerge to explain this behaviour. The first relates to the spatial decomposition methodology detailed in Section 3.2.4, where the division into six distinct subdomains introduces additional variability in the Nu and C_f distributions, potentially impeding the development of robust mapping functions. The second hypothesis suggests that the selected wall-normal position for data acquisition ($y^+_{\text{ac}} \approx 15$) might provide insufficient flow information for precise reconstruction of these quantities. Consequently, the model appears to develop dependencies on training patterns that fail to translate effectively across the complete dataset distribution.

As shown in Figure 3.17(b), C_f exhibits regular variations at higher amplitudes, where spanwise forcing predominates. This is coherent with the results of Chapter 2, as the temporal variation of C_f and Nu is less stochastic and more periodic for the actuated case due to their variations being driven by wall oscillations. This effect is further amplified with higher values of W^+. The A-CAE is able to capture these C_f dynamics fairly well, with errors arising primarily at friction peaks, which tend to be overestimated in the A-CAE predictions.

On the other hand, at lower wall oscillation amplitudes, streak generation remains largely unaffected due to the reduced effectiveness of the control. Consequently, the C_f and Nu fluctuations, driven by streak dynamics, become more stochastic and, thus, more challenging to predict accurately (Figure 3.17(a)). Despite the irregular patterns, the A-CAE still demonstrates an ability to capture some trends, as rises and falls in the ground truth generally align with corresponding rises and falls in the predictions.

An intriguing observation, is that Nu consistently achieves better prediction scores than C_f, despite their similar underlying dynamics. This discrepancy could be explained by a model bias, with the larger variance of Nu within the scaled dataset making it a priority for the model to capture its patterns more effectively.

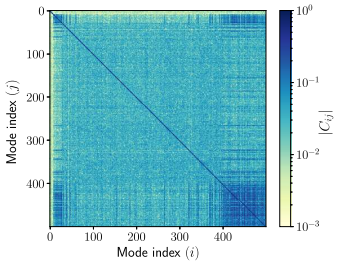

Figure 3.18: Correlation matrix obtained on the test dataset, with $L_d = 500$, for the A-CAE using $\alpha = 0.1$. Absolute values are applied to all coefficients. Colors range from light green (uncorrelated) to dark blue (highly correlated). Mode ranking algorithm of Kang et al. (2022) is applied.

The A-CAE regularization adversely affects latent space orthogonality, as evidenced by scores of $\text{Mean}_c = 4.21$ and $\text{MeanMax}_c = 22.08$. These metrics reveal substantial inter-mode correlations, particularly within modes indices ≈10-30 and ≈400-500, as clearly visualized in the corresponding correlation matrix (Figure 3.18). While the initial modes (1-10) maintain relative independence, the intermediate range (50-400) exhibits correlations of the same magnitude as those obtained in the baseline CAE case (Figure 3.13(a)). Overall, this demonstrates that optimizing Nu and C_f prediction accuracy deeply modifies the latent space structure, resulting in significant entanglement for near 25% of the modes.

3.7.2 Latent separability analysis

The A-CAE demonstrated significant limitations, as its improvements in predicting C_f and Nu did not generalize to the test dataset. Applying the separability analysis protocol described in Section 3.6 to clusters based on C_f and Nu values is thus a relevant follow-up, as it may demonstrate enhanced separability for the A-CAE on the test dataset.

Specifically, eight clusters $\{C_{f_i}\}_{i\in\{1,\dots,8\}}$ are created based on a regular subdivision from the minimal to maximal C_f within the dataset:

$$z_{j,k} \in C_{f\,i} \iff C_{f\,j,k} \in \left[C_f^{\min} + (i-1)\frac{C_f^{\max} - C_f^{\min}}{8}, C_f^{\min} + i\frac{C_f^{\max} - C_f^{\min}}{8} \right], \quad (3.21)$$

where $C_f^{\min} = \min_{j,k}(C_{f\,j,k})$, $C_f^{\max} = \max_{j,k}(C_{f\,j,k})$, and $C_{f\,j,k}$ is the ground truth friction coefficient value of latent sample $z_{j,k}$. The same procedure is applied to the Nusselt number. This number of clusters is chosen arbitrarily to strike a balance: it must be sufficiently large to effectively differentiate between C_f or Nu values, but not so excessive as to complicate cluster distinction.

Each latent sample is thus assigned to a specific cluster based on its C_f or Nu values, allowing for a fair evaluation of the separability of latent samples with respect to these quantities (see Section 3.6).

	CAE	A-CAE
$D_B(C_f)$	31.92	16.13
$D_B(\text{Nu})$	33.45	17.31

Table 3.8: Global Davies-Bouldin indicator as given by Eq. (3.16) applied to clusters of friction coefficient ($D_B(C_f)$) and Nusselt number ($D_B(\text{Nu})$), with $L_d = 500$, on the CAE and A-CAE with $\alpha = 0.1$.

Table 3.8 presents the Global Davies-Bouldin indicator, as defined in Eq. (3.16), for clusters of the friction coefficient and Nusselt number (Eq. (3.21)) on the A-CAE and CAE with $L_d = 500$, evaluated on the test dataset. For both physical quantities, the A-CAE achieved a significantly better separability score, nearly half that of the CAE. This suggests that, despite the lack of improvement in predicting Nu and C_f on the test dataset, the A-CAE regularization results in a more coherently organized latent space with respect to these quantities. Interestingly, the separability for Nu is slightly weaker than that for C_f, even though Nu is better predicted by both models. This observation further supports the notion that more accurate predictions do not necessarily lead to improved separability. This slight discrepancy may simply result from the larger variations in Nu compared to C_f within the dataset.

The latent separation of C_f and Nu clusters is further investigated in Figures 3.19 and 3.20, respectively, using the Davies-Bouldin indicator applied to each cluster combinations, as given by Eq. (3.17). This procedure is applied to both the CAE and A-CAE approaches, on the test dataset, with $L_d = 500$. Overall, the separation exhibits analogous tendencies for Nu and C_f. As expected, clusters corresponding to similar values are characterized by a worse separability compared to clusters with distant values. Indeed, the D_B score consistently decreases with increasing gaps in cluster ranges. Interestingly, the 8th cluster, corresponding to maximum friction cases, is easily separated from all other clusters, suggesting a better distinction of these specific cases where the highest dissimilarity occurs (see Section 2.5). Once again, the A-CAE largely outperforms the CAE in terms of latent separability with respect to these quantities. This superior performance is clearly visible at two levels:

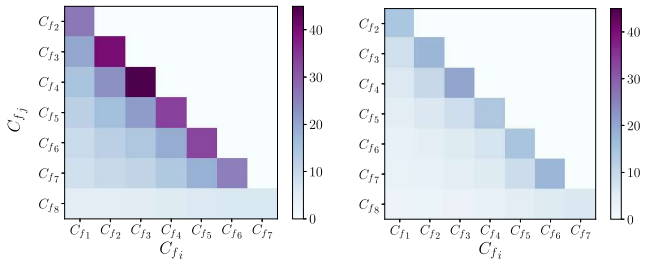

(a) (b)

Figure 3.19: Davies-Bouldin indicator applied to friction coefficient clusters ($D_B(C_{f_i}, C_{f_j})$), with $L_d = 500$, for models: (a) CAE, (b) A-CAE with $\alpha = 0.1$. Only the lower triangular part of the matrix is shown due to the symmetry of the operator.

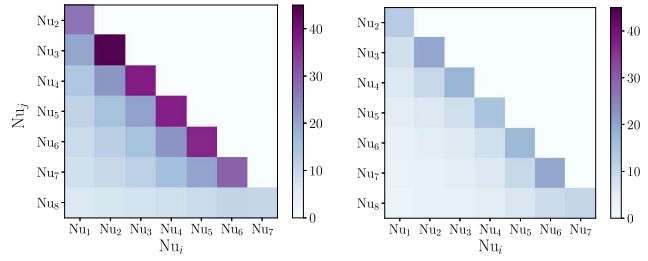

(a) (b)

Figure 3.20: Davies-Bouldin indicator applied to Nusselt number clusters ($D_B(\mathrm{Nu}_i, \mathrm{Nu}_j)$), with $L_d = 500$, for models: (a) CAE, (b) A-CAE with $\alpha = 0.1$. Only the lower triangular part of the matrix is shown due to the symmetry of the operator.

(i) adjacent clusters, corresponding to the diagonal line, exhibit scores near twice lower in the A-CAE case, and (ii) the A-CAE very efficiently separates all non-adjacent clusters, whereas the CAE struggles even with these easier cases.

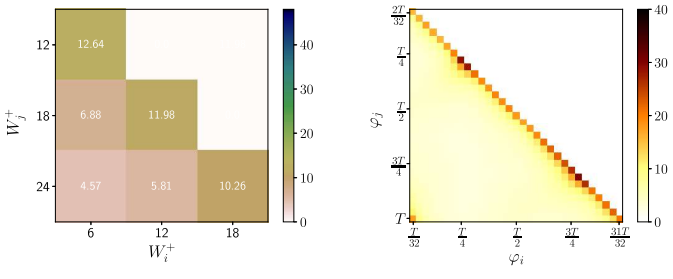

(a) (b)

Figure 3.21: Davies-Bouldin indicator applied to (a): amplitudes W^+ ($D_B(W_i^+, W_j^+)$), and (b): phases φ ($D_B(\varphi_i, \varphi_j)$), on the A-CAE with $\alpha = 0.1$ and $L_d = 500$ latent space. Only the lower triangular part of the matrix is shown due to the symmetry of the operator.

The superior separability of Nu and C_f representations in the A-CAE latent space compared to the baseline CAE suggests this improvement may naturally extend to oscillation phases and amplitudes, given their inherent relationship with friction dynamics. Consequently, the generative factor separability analysis from Section 3.6 is applied to the A-CAE for systematic comparison.

The results reveal a global $D_B(W^+)$ score of 11.88, with consistently low D_B values across all amplitude cluster combinations (Figure 3.21(a)), demonstrating the clear advantage of the A-CAE in distinguishing oscillation amplitudes within its latent space over all other methods examined in Figure 3.14. Similar to the CAE results (Figure 3.14(a)), higher amplitudes exhibit slightly better separability.

Further analysis of phase latent separation shows significant improvements, evidenced by Figure 3.21(b) and a global $D_B(\varphi)$ score of 21.44, outperforming the CAE, β-VAE, and extended POD cases from Figure 3.15. Non-adjacent phases are now distinctly separable, including in previously identified problematic regions corresponding to ascending friction phases. However, adjacent phases in this friction increase region still present the highest D_B scores amongst all phase pairs, indicating remaining challenges for these specific cases.

The enhancements in Nu and C_f separability have consequently produced a latent space organization that more effectively distinguishes different oscillation phases and amplitudes. This improvement likely stems from the fundamental physical connection between φ, W and the resulting variations in both Nusselt number and friction coefficient in the actuated case.

3.8 Conclusions and perspectives

This study comprehensively evaluated the effectiveness of proper orthogonal decomposition and various autoencoder architectures in reducing the dimensionality of turbulent channel flow data at $\mathrm{Re}_\tau = 200$ and $\mathrm{Pr} = 1$. The dataset, comprising velocity and temperature fields, was generated using large eddy simulations under spanwise wall oscillations with a fixed oscillation period and varying amplitudes. Specific attention was set towards small turbulent scales reconstruction, latent orthogonality, and latent separability of generative factors. The main results from the investigation are summarised as follows:

- CAEs consistently outperformed POD in terms of reconstruction accuracy across all flow fields, particularly at lower latent dimensions. The most significant improvement was observed in the wall-normal velocity component (v), highlighting the advantage of non-linearity offered by AE architectures in capturing complex turbulent dynamics.

- Compressing all four fields (u, v, w, Θ) simultaneously resulted in superior reconstruction performance compared to individual field compression. This

enhancement leverages the inherent physical correlations between the variables, especially between the streamwise velocity (u) and temperature (Θ) fields.

- Both POD and AE methodologies encountered difficulties in accurately reconstructing small-scale turbulent structures, particularly within the v-field. POD, constrained to 500 modes, was limited to capturing only large-scale flow characteristics, and exhibited a tendency towards mean-value approximation.

- β-VAEs achieved a balance between the reconstruction performance of CAEs and the latent space orthogonality of extended POD. Although there was a decrease in accuracy for the v and Θ fields, correlation metrics indicated improved latent space structure compared to CAEs.

- The β-VAE maintains optimal performance with only 424 latent variables ($L_d = 424$), preserving an effective balance between reconstruction fidelity and latent space orthogonality. Notably, in this pruned architecture, latent variables exhibiting greater KL divergence tend to show reduced inter-mode correlations.

- Both CAEs and β-VAEs demonstrate superior capability in distinguishing amplitude-related features (W^+), compared to extended POD. Autoencoders additionally incorporate key physical insights by differentiating phases (φ) with substantially different Stokes strain or streak orientation. Despite similar trends, the CAE consistently outperformed the β-VAE in terms of latent separability.

A fourth model, a friction-augmented CAE, was then introduced for further investigation. This model resulted in a trade-off between the predictive accuracy of the flow fields, mode entanglement, and latent space separability.

- While the A-CAE incurred a modest decrease in flow-field accuracy and enhanced mode entanglement, it achieved a notable improvement in reconstructing C_f and Nu from latent samples on the training data. However, this enhanced predictive accuracy failed to generalize to the test dataset.

- Nevertheless, the A-CAE regularization successfully induced a more coherently structured latent space, significantly improving the separability of samples based on Nu, C_f, oscillation phase, and amplitude, improvements that generalized well to unseen data.

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

- Address the limited predictability of the A-CAE for C_f and Nu by exploring sequence-based predictions using time-adjacent (or subdomain-adjacent) samples, and replacing the MLP with more suitable architectures such as recurrent neural networks (Sherstinsky, 2020).

- Employ CROM techniques (Kaiser et al., 2014) to elucidate the dynamics encoded within the latent spaces of each model. This approach, which is more complex than that of the present study, may reveal more meaningful information regarding how each model structures its latent space with regards to physical parameters.

- Investigate the application of three-dimensional convolutional layers to compress the entire 3D channel, as opposed to the current two-dimensional slice approach. This holistic compression strategy may capture spatial correlations more effectively, improving reconstruction accuracy and model generalisation. However, compressing the 3D space, characterized by a much larger dimension, may prove a challenging task.

- Contextualize the present results within the broader landscape of modern dimensionality reduction, which includes physics-informed autoencoders that enforce physical laws (such as divergence constraints) (Mohan et al., 2023), Koopman operator-based methods which aim to identify linear representations of nonlinear dynamics (Bevanda et al., 2021), sparse regression techniques like the Sparse Identification of Nonlinear Dynamics (SINDy) (Brunton et al., 2016) and its hybrid architecture SINDy-SHRED (Williams et al., 2024) that discover interpretable, parsimonious governing equations from data.

- Combine AE-based dimensionality reduction with long-short-term-memory networks or transformers (Solera-Rico et al., 2024) to develop an efficient predictive model of the flow. This integration could facilitate real-time flow control and optimisation in engineering applications.

Chapter summary

This chapter investigated the efficiency of AE variants and extended POD as dimensionality reduction techniques and serves as a first step towards building a ROM. The dataset is built upon the dissimilar heat transfer results of Chapter 2, composed of 2D $y^+_{\text{ac}} \approx 15$ slices of the velocity and temperature fields, from 21 LES at $\text{Re}_\tau = 200$ and $\text{Pr} = 1$, at a fixed period of $T^+ = 500$, with varying oscillation amplitudes in the $W^+ \in [0, 30]$ range.

After presenting the methods and evaluation metrics, CAEs are first compared to extended POD, mainly in terms of reconstruction accuracy. An attempt to find a balance between the latent orthogonality of extended POD and the reconstruction accuracy of CAEs is then explored through β-VAEs. A mode pruning algorithm is

applied to the β-VAE to eliminate over-regularized variables. All three models are then compared regarding separability of oscillation amplitudes W^+ and phases φ within their latent spaces, using traditional clustering metrics. Finally, a friction augmented CAE is considered, attempting to improve the structure of the latent space with regards to C_f and Nu.

The CAE outperforms extended POD in reconstruction accuracy across all tested latent dimensions, particularly for lower dimensions. It significantly better captures small flow scales, resulting in the most notable improvements appearing in the wall-normal velocity reconstruction. The pruned β-VAE version uses only 424 latent variables while maintaining accuracy identical to the full model. When compared to a CAE with $L_d = 400$, the pruned β-VAE achieves a good compromise between orthogonality and reconstruction quality. Separability analysis shows that AE models better distinguish oscillation amplitudes in their latent spaces and, unlike extended POD, can differentiate phases with similar values of C_f and Nu but opposite Stokes strain directions or different streak orientations. The CAE demonstrates the best performance in separating both parameters, suggesting β-VAE regularization offers no advantage in this regard.

The A-CAE showed improved prediction of C_f and Nu from latent samples without compromising the accuracy of the primary fields. Separability analysis revealed that while C_f and Nu prediction accuracy improvements did not generalize, the gains of the A-CAE in latent space separability did, resulting in a more coherently structured latent space regarding both physical quantities. The model regularization trades mode entanglement for improved phase/amplitude latent separability, which is likely a consequence of the enhanced Nu/C_f discrimination.

While not yet a complete ROM, as the latent space dynamical model remains to be developed, these results provide important groundwork. Possible approaches for this model include time series predictions using long-short-term-memory networks or transformers, or CROMs which may offer insights into the physical mechanisms behind dissimilar heat transfer. In any case, the large number of dimensions required to describe the flow suggests significant challenges in constructing such a dynamical model. The current results are therefore not yet directly applicable for parametric studies in subsequent chapters, but the LES configuration used to create the dataset is promising, as it maintains good accuracy while being computationally efficient.

Version française

Ce chapitre étudie plusieurs techniques de réduction de dimensionnalité: les réseaux de neurones autoencodeurs et la décomposition orthogonale propre (POD). Cette étude représente une première étape vers la construction d'un modèle à dimension réduite. Le jeu de données utilisé s'appuie sur les résultats de transfert thermique dissymétrique du Chapitre 2, composé de coupes 2D à $y^+_{\mathrm{ac}} \approx 15$ des champs de vitesse et température, issues de 21 simulations implicites aux grandes échelles à $\mathrm{Re}_\tau = 200$ et $\mathrm{Pr} = 1$. La période d'oscillations est fixée à $T^+ = 500$, avec des

variations d'amplitudes dans $W^+ \in [0, 30]$.

Après avoir présenté les méthodes de réduction de dimensionnalité et les métriques d'évaluation associées, les autoencodeurs convolutionnels (CAE) sont d'abord comparés à la POD étendue, principalement en termes de précision de reconstruction. Une approche visant à concilier l'orthogonalité latente de la POD et la reconstruction des CAEs est ensuite explorée via les autoencoders variationnels β (β-VAE). Un algorithme d'élagage des variables latentes est appliqué à ce modèle pour éliminer les variables sur-régularisées. Les trois modèles sont ensuite comparés quant à leur capacité à séparer les amplitudes W^+ et phases φ d'oscillations dans leurs espaces latents, à l'aide de métriques de clustering. Enfin, un autoencoder augmenté (A-CAE) est considéré, visant à améliorer la structure de l'espace latent par rapport aux quantités physiques C_f et Nu.

Le CAE surpasse la POD étendu en terme de précision de reconstruction pour toutes les dimensions latentes évaluées, particulièrement pour les plus basses. Il capture plus efficacement les petites échelles de l'écoulement, l'amélioration étant la plus marquée pour la reconstruction de la vitesse normale à la paroi. La version élaguée du β-VAE n'utilise que 424 variables latentes tout en maintenant une reconstruction identique à celle du modèle complet. Comparé à un CAE avec $L_d = 400$, le β-VAE élagué offre un bon compromis entre orthogonalité de l'espace latent et qualité de reconstruction. L'analyse de séparabilité des variables latentes révèle que les autoencodeurs distinguent mieux les amplitudes d'oscillations que la POD étendue. De plus, ils parviennent à différencier des phases présentant des valeurs similaires de C_f et Nu, mais avec des directions de couche de Stokes opposées ou des orientations de structures différentes. Le CAE obtient les meilleures performances pour la séparation des deux paramètres, ce qui suggère que la régularisation du β-VAE ne présente aucun avantage à cet égard.

La régularisation appliquée à l'A-CAE améliore la prédiction de Nu et C_f à partir des variables latentes, sans dégrader significativement la reconstruction des champs physiques primaires. L'analyse de séparabilité révèle que, bien que l' amélioration des prédictions de Nu et C_f ne se généralisent pas aux données de test, les avancées en matière d'organisation de l'espace latent conservent leur caractère généralisable. Plus précisément, une séparabilité accrue de l'espace latent vis-à-vis de ces deux grandeurs physiques est obtenue par rapport aux autres modèles. La régularisation compromet l'orthogonalité des variables latentes, mais elle améliore aussi la séparation des phases et amplitudes d'oscillations, vraisemblablement en raison de l'amélioration de la discrimination de Nu et C_f.

Bien que la création d'un modèle à dimension réduite complet reste à finaliser, en particulier par le développement d'un modèle dynamique dans l'espace latent, ces résultats établissent un fondement essentiel pour cette construction. Plusieurs approches pourraient être envisagées pour construire ce modèle, telles que des prédictions de séries temporelles basées sur des réseaux de neurones de type transformer, ou des modèles à dimension réduite fondés sur une analyse par clusters. Ces

méthodes pourraient contribuer à élucider les mécanismes physiques régissant le transfert thermique dissymétrique. Cependant, la dimensionalité élevée nécessaire à la description de l'écoulement représente un défi de taille pour la construction d'un tel modèle dynamique. Si les résultats actuels ne sont pas encore directement applicables aux études paramétriques des chapitres suivants, la configuration de simulation aux grandes échelles utilisée pour générer les données s'avère prometteuse pour cet objectif, offrant un compromis satisfaisant entre précision et efficacité computationnelle.

Chapter 4

Policy-based optimization of spanwise wall oscillations

Contents

4.1 Introduction . . . 112
4.2 Flow conditions and numerical simulation details . . . 115
4.2.1 Details of numerical simulations . . . 116
4.2.2 Validation of numerical simulations . . . 117
4.3 Control-law optimization using reinforcement Learning . . 118
4.3.1 Reinforcement Learning . . . 118
4.3.2 Policy-based optimization . . . 121
4.3.3 PBO configuration . . . 127
4.4 Optimal drag reduction . . . 128
4.5 Optimal net energetic balance . . . 132
4.6 Conclusions & Perspectives . . . 136
Extended Results . . . 139
4.7 Application to heat-transfer enhancement . . . 139
4.7.1 PBO-LES configuration . . . 139
4.7.2 Results . . . 140
4.7.3 DNS validation and comparative analysis . . . 143
4.7.4 Conclusion . . . 145
4.8 Improving Heat transfer with minimum drag penality . . . 146
4.8.1 Results . . . 147
4.8.2 Investigating optimal DHT period at fixed amplitude . . . 149
4.8.3 DNS validation and comparative analysis . . . 150
4.8.4 Conclusion . . . 154
Chapter summary . . . 155

This chapter presents a novel framework combining single-step reinforcement learning with implicit large eddy simulations. Specifically, the policy-based optimization algorithm is utilized. It is a state-of-the-art technique enhancing simplified reinforcement learning through evolutionary algorithm-inspired mechanisms. The approach capitalizes on both the algorithm's search efficiency (through its focus on

near-optimal parameter regions) and the computational advantages of LES (notably its parallelization capabilities, with which PBO is compatible).

Following validation of the implicit LES accuracy, the PBO-LES framework is first applied to drag reduction and net energetic gain optimization. While secondary to the primary research objectives, this initial application serves as a critical benchmark, leveraging well-established optimal parameters from prior studies. Subsequently, the framework is applied to: (i) heat-transfer enhancement maximisation, and (ii) dissimilar heat transfer optimization, which is the core objective of this thesis.

The chapter is structured as follows:

- A published paper (from Section 4.1 to Section 4.6) details the PBO-LES methodology and its validation through drag reduction applications.
- Section 4.7 extends the method to heat-transfer enhancement, identifying previously undetermined optimal SWO parameters.
- Section 4.8 applies PBO-LES to dissimilar heat transfer optimization, specifically targeting performance improvements over the reference case of Chapter 2.

4.1 Introduction

Reinforcement learning was introduced earlier in Section 1.4, with its connection to optimal control discussed by Bucci et al. (2019).

In Sonoda et al. (2023), drag reduction is performed in a turbulent channel-flow at $\mathrm{Re}_\tau = 150$ using wall blowing and suction for the action, and streamwise and wall normal velocity fluctuations at $y^+ = 15$ as a state. More precisely, Deep Deterministic Policy Gradient (DDPG) is used with a non-linear actor network in order to find the optimal policy, where the aim is to beat opposition control. The algorithm is applied on a minimal-channel and the converged control strategy is verified on the full-channel, featuring a larger domain and finer resolution, which is a conventional procedure in this field. A 37% drag reduction is achieved with the non-linear approach compared to 23% using opposition control, with similar results on the full-channel. Guastoni et al. (2023) used a similar approach combined with an innovative multi-agent system. The obtained control strategy reduces the drag by 20% more than opposition control, although at a more substantial energetic cost.

The effectiveness of DRL methods in discovering optimal control strategies depends on the speed at which the coupled environment can progress in time. Maintaining fair accuracy while minimizing the duration of each training iteration is therefore essential. Possible solutions include using coarse simulations on a small domain, as in Sonoda et al. (2023), or the implementation of a data-driven ROM, which aims to reproduce the simulation at a lower cost. A common approach in model-based RL is to frequently verify the policy with DNS accuracy, in order to

update the ROM, while learning the optimal policy. Linot and Graham (2023) introduce such ROM, using hybrid auto-encoders and the Dormand-Prince method for time integration. This model is then coupled with the Soft Actor-Critic (SAC) RL algorithm in Linot et al. (2023), in order to reduce the drag with two opposed streamwise slot jets. The agent discovered an innovative strategy, forming two unsustainable low-speed streaks that result in re-laminarization.

Recent advancements have seen the effective application of DRL to open-loop control in the form of Simplified (or single-step) Reinforcement Learning (SRL), which is the approach employed in the current study. The first algorithm introduced by Ghraieb et al. (2021); Lai et al. (2021) was a simplified variant of PPO, referred to as PPO-1, as outlined in Section 1.4.1. The effectiveness of PPO-1 has been demonstrated in various flow control applications:

- PPO-1 was coupled with a CFD environment to control conjugate heat transfer by Hachem et al. (2021).
- A similar approach was applied for shape optimization by Viquerat et al. (2020), generating the shape of a wing that maximises the lift-to-drag ratio. The shape was parameterised by a Bézier curve, using the point coordinates as input, with up to 12 degrees of freedom.
- The algorithm was used by Wang et al. (2023) to optimise both the spatial location and length of a downstream splitter plate to suppress vortex shedding behind a cylinder in a 2D flow. Their study compared PPO-1 with grid search and Bayesian Optimization (BO). While PPO-1 was more time-consuming than grid search, it explored the area around the optimum more precisely and avoided unnecessary exploration of unproductive regions. Conversely, BO required more wall-time but less CPU time than PPO-1 because BO is not compatible with parallelization. In practice, this means PPO-1 converges faster, as it can leverage parallel processing to speed up the search for the optimal solution, despite the higher CPU time usage.

Building on these advances, Viquerat et al. (2023) proposed Policy-Based Optimization (PBO), an improved and more robust extension of PPO-1, detailed in Section 1.4.1. PBO performance was found to be comparable to CMA-ES and significantly surpassed PPO-1 when minimising complex analytical functions and establishing parametric control laws for the Lorenz attractor. Increasing the number of individuals per generation led to PBO slightly outperforming CMA-ES, particularly in terms of convergence speed. PBO is also recently utilized in Mohammadikalakoo et al. (2024). The authors revealed PBO's exceptional performance in Tollmien-Schlichting wave suppression, using a single-frequency linear wall-blowing and suction control strategy. Indeed, PBO achieves threefold faster convergence compared to Particle Swarm Optimization (PSO), while identifying superior control parameters. However, this advantage diminished in strongly non-linear control scenarios,

highlighting the method's current limitations. These recent studies establish PBO as an efficient parameter optimization method, competitive with state-of-the-art optimization techniques.

It is worth noting that unlike evolutionary algorithms, which have seen widespread application across diverse fields (Slowik and Kwasnicka, 2020), SRL has, to the best of the authors' knowledge, not yet been utilized outside flow control. As the efficacy of this method is further demonstrated through additional cases in flow control, its expansion could extend to other domains.

This study presents a novel computational framework that integrates PBO, a simplified reinforcement learning approach, with implicit LES to optimise oscillating spanwise wall motions for drag reduction in turbulent channel flow. Conventionally, extensive parameter studies using DNS have been employed to identify optimal control parameters (Quadrio and Ricco, 2004; Gatti and Quadrio, 2013). However, this approach is computationally expensive and restricts the exploration of the parameter space. In contrast, the current work introduces a more efficient and cost-effective methodology that leverages recent advances in machine learning to significantly accelerate the discovery of optimal control strategies. The SWO drag reduction case serves as a benchmark to test the efficiency of the ML approach in discovering control strategies, as the optimal set of parameters has been established by previous studies. While other ML methods may potentially outperform PBO, the present study focuses on showcasing the ability of SRL to discover effective control strategies using fewer computational resources compared to the conventional grid search approach.

Section 4.2 describes the flow conditions, numerical procedures, and validation of the computational setup employed in this study. The optimization process utilizing SRL is detailed in Section 4.3.

The results, presented in Section 4.4, initially focus on identifying the SWO parameters that yield the maximum drag reduction. Based on the results of Quadrio and Ricco (2004), the optimal oscillation period and amplitude are expected to be approximately $T^+ \approx 100$ and $W^+ \approx 30$, respectively, within the investigated parameter range. Subsequently, Section 4.5 shifts the focus to optimizing the net energy balance, with the expectation of identifying more precise optimal parameters compared to those reported by Quadrio and Ricco (2004). While elucidating the physical mechanisms underlying drag reduction is beyond the scope of the present study, the primary objective is to showcase the efficiency of the PBO method in optimizing SWO parameters for both drag reduction and net energy balance.

The integration of machine learning techniques with computationally-efficient numerical simulations represents a promising approach for accelerating the discovery of optimal flow control strategies. By demonstrating the effectiveness of PBO in a well-established test case with known optimal parameters, this study aims to pave the way for the application of SRL methods to more complex flow control prob-

lems, ultimately contributing to the development of energy-efficient and sustainable technologies in various industrial sectors.

4.2 Flow conditions and numerical simulation details

Implicit LES are performed with the same configuration as described in Section 3.2.2, excluding the simulation of the scalar field which is unnecessary for drag reduction studies.

As stated in Section 1.2.4, SWO have been extensively used to reduce the drag (Agostini et al., 2015, 2014; Quadrio and Ricco, 2004). A relevant metric to quantify the drag reduction is provided in the following equation:

$$P_{\text{sav}} = 1 - \frac{\langle C_{f_{\text{ac}}} \rangle_t}{\langle C_{f_{\text{unac}}} \rangle_t}. \tag{4.1}$$

The spanwise wall velocity imposed by the control mechanism inherently incurs an energy cost. To fairly evaluate the efficiency of drag reduction, it is essential to account for this control cost. Neglecting potential mechanical losses, inevitable in real-world applications, the energy cost corresponds to the spanwise frictional power given in the numerator of Eq. (4.2). This power is defined as the time-integrated product of the spanwise shear and the spanwise wall velocity. For a meaningful comparison with the energy savings, which are expressed as a ratio relative to the unactuated case in Eq. (4.1), the control cost must also be expressed relative to an appropriate quantity from the unactuated flow. The unactuated streamwise frictional power serves this purpose, as indicated by the denominator of Eq. (4.2). Following Quadrio and Ricco (2004), the control cost is therefore represented as the actuated spanwise frictional power expressed as a percentage of the frictional power in the unactuated case:

$$P_{\text{req}} = \frac{\left[\int_{t_i}^{t_f} \left(\tau_{w,z}^{l} + \tau_{w,z}^{u} \right) w_{\text{w}} \, \mathrm{d}t \right]_{\text{ac}}}{\left[U_b \int_{t_i}^{t_f} \left(\tau_{w,x}^{l} + \tau_{w,x}^{u} \right) \, \mathrm{d}t \right]_{\text{unac}}}. \tag{4.2}$$

In this expression, t_i, t_f, $\tau_{w,x} = \nu \frac{\partial \langle u \rangle_{x,z}}{\partial y}$ and $\tau_{w,z} = \nu \frac{\partial \langle w \rangle_{x,z}}{\partial y}$ designate the initial and final time step for statistical collection and the streamwise and spanwise shear, respectively. The l and u superscript represent the lower and upper walls, respectively. The numerator is calculated at each iteration and then averaged over time for each actuated simulation, whereas the denominator can be directly obtained using $C_{f,\text{unac}}$ to determine the unactuated streamwise shear, along with the imposed value of $U_b = \frac{2}{3}$.

Therefore, the difference between the energy saved and the energy required to apply wall motions yields the net energetic balance:

$$P_{\text{net}} = P_{\text{sav}} - P_{\text{req}}. \tag{4.3}$$

Optimizing P_{sav} is therefore a simplified version of the problem, focusing solely on reducing drag, whereas incorporating P_{req} into the net energetic balance offers a more realistic metric for assessing control performance by accounting for the control cost.

4.2.1 Details of numerical simulations

The present optimization employs the implicit-LES S_1 configuration from Table 3.1, excluding the scalar field simulation, since the initial objective is drag reduction.

Each simulation is initiated using the same input snapshots taken from the steady-state of an unactuated simulation on the S_1 configuration. The warm-up period is fixed at approximately 3 200 (in t^+), an extended duration over a considerable time scale. It is selected in order to guarantee that the low drag state has been reached before initiating statistical collection. The duration of the statistical collection varies with the oscillation period to ensure it spans an appropriate number of cycles, as described in the subsequent equation:

$$t^+_{\text{collect}} = T^+ \lfloor \frac{1630}{T^+} \rfloor, \tag{4.4}$$

where $\lfloor \bullet \rfloor$ represents the floor function. As the speed of the simulations is important to reduce the total computational time of the optimization process, the only output of the simulations during training is set to P_{sav} and P_{req}, neglecting snapshots and additional statistics of the flow. These quantities will be used to construct the reward functions for policy-based optimization. The inherent stochasticity of the problem, coupled with the finite averaging time over short intervals, results in variations in the measured P_{req} and P_{sav}, thereby influencing the rewards obtained throughout the optimization process.

Configuration	Wall Time (hours)	CPU-hours
S_1	1.54	24.64
S_2	14.86	17 833.85

Table 4.1: Comparison of wall time and CPU-hours between S_1 and S_2.

Table 4.1 compares the computational costs of a DNS at $\text{Re}_\tau = 180$ and a LES at $\text{Re}_\tau = 200$ for a turbulent channel flow. The DNS utilizes the S_2 configuration of Table 2.1, while the LES employs the S_1 configuration. Both simulations are conducted over a total wall time of 4830 nominal wall units. The LES requires less than one-ninth of the wall-time and approximately 724 times fewer CPU hours compared to the DNS. This substantial reduction in computational resources demonstrates the significant savings achieved by using the S_1 configuration in the LES. Furthermore,

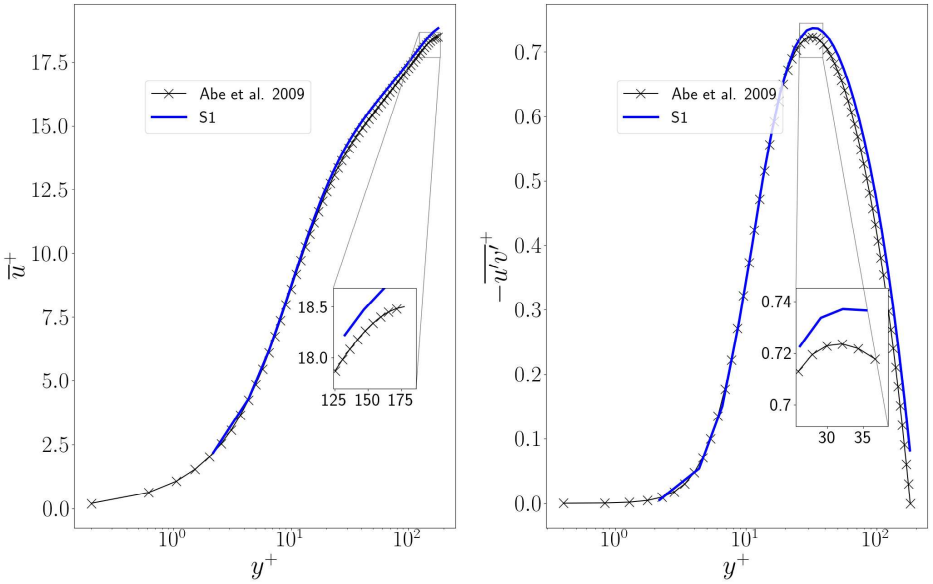

Figure 4.1: Wall-normal distribution of (a) mean streamwise velocity $\overline{u}$ and (b) turbulent shear stress $-\overline{u'v'}$. Statistics obtained from simulations on S_1 are represented by plain blue line. Reference results (Abe et al., 2009) are shown by the black line with crosses.

the dramatic decrease in CPU time makes the LES approach particularly well-suited for parallelization, as the computational load can be efficiently distributed across multiple processors. It is important to note that this advantage is of particular relevance in the context of parametric studies.

4.2.2 Validation of numerical simulations

The S_2 DNS configuration (see Table 2.1) is extensively validated in Section 2.3 through comparisons with reference simulations. The S_1 implicit LES configuration, on the other hand, is validated in Section 3.2.3 for temperature statistics, as well as for the actuation-induced variations in Nu and C_f. However, its accuracy for velocity statistics and for the recently introduced metrics P_{req}, and P_{net} has yet to be demonstrated.

In this section, the results of the current numerical simulations, with and without control, are compared to two reference results. The first reference simulation (Abe et al., 2009) is utilized to validate the unactuated relevant statistics, while the second reference simulation (Quadrio and Ricco, 2004) verifies the drag reduction results on the controlled flow. Their DNS resolution is much finer than the S_1 configuration, with a mesh 1.5, 3.1, 10 and 2.6 times more resolved in Δx, Δz, Δy_{min} and Δy_{max}, respectively, for Quadrio and Ricco (2004).

The wall-normal distributions of the mean streamwise velocity and turbulent

shear stress are conveyed by Figure 4.1. Simulations utilizing the S_1 mesh are denoted by the solid blue lines, while the black solid lines with crosses represent the reference data of Abe et al. (2009). For both quantities, the S_1 results exhibit a fair agreement with the reference data, with small discrepancies which could be attributed to the minor gap in Re_τ. The value of the time averaged friction coefficient $\langle C_f \rangle_t$ obtained on S_1 (7.62×10^{-3}) has a 4% discrepancy compared to the unactuated value reported in Quadrio and Ricco (2004). This is reasonable considering the differences in resolution. Overall, the unactuated statistics on the S_1 mesh display satisfactory precision, consistent with the reference data.

T^+	W^+	Configuration	$\%P_{\mathrm{sav}}$	$\%P_{\mathrm{req}}$	$\%P_{\mathrm{net}}$
125	18	Quadrio and Ricco (2004)	39.3	161.2	-121.9
125	18	S_1	41.3	162.1	-120.8
200	4.5	Quadrio and Ricco (2004)	12.7	7.7	5.0
200	4.5	S_1	14.3	7.7	6.6

Table 4.2: Comparison in P_{sav}, P_{req}, and P_{net}, between S_1 and Quadrio and Ricco (2004).

Two sets of parameters are investigated when applying spanwise wall oscillation: (i) $T^+ = 125, W^+ = 18$ and (ii) $T^+ = 200, W^+ = 4.5$. The results obtained from the current simulations are compared to those of Quadrio and Ricco (2004), as presented in Table 4.2, in terms of P_{sav}, P_{req}, and P_{net}. For the former set of parameters, a negative energetic balance ($P_{\mathrm{sav}} < P_{\mathrm{req}}$) is obtained, with only a 1% discrepancy observed between the present results and those of Quadrio and Ricco (2004). In contrast, the latter set of parameters ($T^+ = 200, W^+ = 4.5$) produces a positive, albeit relatively modest, energetic balance. For these flow parameters, a slightly more substantial contrast in P_{net} is exhibited, but it remains within an acceptable range. In both cases, excellent agreement is obtained for P_{req}, with the minor discrepancies in P_{net} arising from differences in P_{sav}. The relevant statistics of the controlled flow are adequately predicted, as will be further demonstrated in Sections 4.4 and 4.5.

4.3 Control-law optimization using reinforcement Learning

4.3.1 Reinforcement Learning

Control strategies are either active or passive depending on whether or not they introduce an energy cost. Closed-loop control is a type of control in which an output observable is used at each time step to adjust the inputs to the dynamic change.

On the contrary, open-loop control assumes that the action does not depend on the state of the system. In this case, the actuation is often steady or periodic. Figure

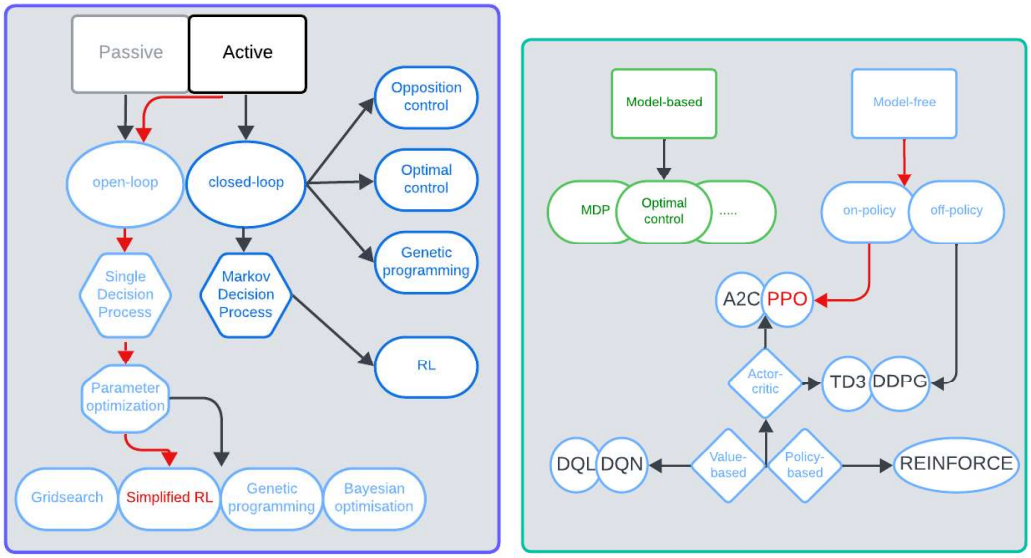

(a) (b)

Figure 4.2: Non-exhaustive diagrams describing: (a) The main flow control strategies. (b) Reinforcement Learning algorithms. Red arrows indicate the strategies or algorithms relevant to this work.

4.2(a) overviews the different types of flow control strategies, as well as related optimization techniques. Reinforcement Learning, a subset of machine learning, arises in the case of closed-loop control (see Section 1.4). RL is composed of many branches and sub-branches, as illustrated in Figure 4.2(b). Model-based approaches involve creating a representation of the environment, which can be updated as the agent learns the policy (see Section 1.4). In contrast, model-free methods learn the optimal policy directly through interaction with the environment, without relying on an explicit model. Model-free methods are often preferred in scenarios like flow control due to their reduced overhead, as they eliminate the need for building and maintaining an accurate model of the environment. In this context, the environment is typically provided by simulations or experimental setups. Additionally, model-free methods can be more robust in complex or not fully understood systems, where constructing an accurate model is challenging or impractical.

On-policy algorithms use the same policy for learning and choosing actions, while off-policy algorithms utilize separate policies for these tasks. On-policy algorithms therefore usually favour exploitation, while off-policy methods encourage exploration. The on-policy algorithm Proximal Policy Optimization (PPO) (Schulman et al., 2017) and the off-policy algorithm Twin Delayed Deep Deterministic Policy Gradient (TD3) (Dankwa and Zheng, 2019) are among the most commonly used in the literature, known for their efficiency in both convergence speed and stability.

The value function V of a state $\boldsymbol{s_t}$ at time step t is recursively defined by the

Bellman equation, which estimates the cumulative reward over time:

$$V(\boldsymbol{s_t}) = r_t + \gamma(1-d)V(\boldsymbol{s_{t+1}}), \tag{4.5}$$

where r_t is the reward at time step t and $\gamma \in [0,1]$ is the discount rate, introduced to favour short-term gains. Furthermore, $d = 1$ if $\boldsymbol{s_t}$ is a terminal state, while $d = 0$ if it is not. Value-based methods focus on estimating this value function, while policy-based methods concentrate on directly optimizing the policy. An alternative to the value function is the Q-function, which represents the cumulative reward from a given state $\boldsymbol{s_t}$ by taking a specific action $\boldsymbol{a_t}$ and then continuing to follow the optimal policy. It essentially estimates the future rewards an agent can expect from each action in a given state, helping guide decision-making by comparing different options. Mathematically, the Q-function is defined as:

$$Q(\boldsymbol{s_t}, \boldsymbol{a_t}) = r_t + \gamma(1-d)Q(\boldsymbol{s_{t+1}}, \boldsymbol{a_{t+1}}), \tag{4.6}$$

where $\boldsymbol{a}$, r, $\boldsymbol{s}$ indicate actions, rewards and states respectively, and the indices denote the iteration steps.

Practically, the Q-function helps agents make optimal decisions by evaluating and comparing the potential rewards of taking different actions. The action $\boldsymbol{a_{t+1}}$ is either chosen by the current policy (in the case of SARSA) or is the action that maximizes the Q-function for state $\boldsymbol{s_{t+1}}$ (in the case of Q-learning, see Vignon et al. (2023)). Q-learning has found widespread application in various fields, such as robotics (Sumanas et al., 2022) and video-games (Shao et al., 2019), where the Q-function helps identify the optimal actions at each state, ultimately enhancing long-term performance.

Actor-critic methods leverage the strengths of both value-based and policy-based approaches by learning the value function (critic) and the policy (actor) simultaneously. The actor can learn both from its own experience and from the feedback provided by the critic. This hybrid nature reduces the variance of the policy gradient by using the critic's value function as a reference, resulting in fewer samples needed to converge. Additionally, different types of critics, such as state value functions or advantage functions, can be used to address various problems and objectives.

The best-known Actor-Critic method is by far PPO. This on-policy method incorporates the following components:

- Stochastic Policy (Π_θ): The actor is represented by a neural network, which samples from its output distribution to predict future actions.
- Critic (V_ϕ): A separate neural network that computes the value function and is trained to verify the Bellman equation (Eq. (4.5)).
- Advantage Function $A^{\Pi_\theta}(\boldsymbol{s}, \boldsymbol{a})$: Measures the quality of an action compared to the average behavior of the policy in that state. It is essential for predicting

the loss function of the policy network. It is often computed from the value function using Generalized Advantage Estimation (GAE) (Schulman et al., 2015).

The clipped Loss function of the PPO actor network is given by:

$$L_{Clip}^{\theta,\theta_{old}}(\mathcal{D}) = \mathbb{E}_{(\boldsymbol{s},\boldsymbol{a})\sim\mathcal{D}}\left[\min\left(\rho_{\theta_{old}}^{\theta}A^{\Pi_\theta}(\boldsymbol{s},\boldsymbol{a}),\ \text{Clip}\left(\rho_{\theta_{old}}^{\theta},1-\varepsilon,1+\varepsilon\right)A^{\Pi_\theta}(\boldsymbol{s},\boldsymbol{a})\right)\right],$$
$$\text{where } \varepsilon\in[0,1],\ \rho_{\theta_{old}}^{\theta}=\frac{\Pi_\theta(\boldsymbol{a}|\boldsymbol{s})}{\Pi_{\theta_{\text{old}}}(\boldsymbol{a}|\boldsymbol{s})},\ \text{and Clip}(x,c,d)=\begin{cases}c & \text{if } x<c\\ x & \text{if } c\le x\le d\\ d & \text{if } x>d\end{cases}. \tag{4.7}$$

In this equation:

- $\mathcal{D}$ is the buffer containing all trajectories obtained with the current policy, cleared after each iteration.
- The expectation (or average) over the buffer $\mathcal{D}$ is formally written as: $\mathbb{E}_{(\boldsymbol{s},\boldsymbol{a})\sim\mathcal{D}}\left[f(\boldsymbol{s},\boldsymbol{a})\right]=\frac{1}{|\mathcal{D}|}\sum_{(\boldsymbol{s},\boldsymbol{a})\in\mathcal{D}}f(\boldsymbol{s},\boldsymbol{a})$.
- The policy probability $\Pi_\theta(\boldsymbol{a}|\boldsymbol{s})$ represents the likelihood of selecting action $\boldsymbol{a}$ given the state $\boldsymbol{s}$.
- The clipping function ensures the new policy remains close to the policy from the previous iteration $\Pi_{\theta_{old}}(\boldsymbol{a}|\boldsymbol{s})$, promoting stability.
- The clipping hyper-parameter ε balances stability and convergence speed.

Another common practice to promote stability is the use of a target policy, whose weights are updated every other iteration using Polyak Averaging (Dankwa and Zheng, 2019). A variant of PPO, proposed in Hämäläinen et al. (2018), implements Covariance Matrix Adaptation (CMA) to speed up optimization and avoid local optima.

4.3.2 Policy-based optimization

Spanwise wall oscillation is an active open-loop control method, which can be associated with single decision processes (see Figure 4.2(a)), as described in Lai et al. (2021). Many options are available in order to find the optimal control policy, including Grid Search (Wang et al., 2023), Bayesian Optimization (Wang et al., 2023), Genetic Programming (Slowik and Kwasnicka, 2020), and SRL (Ghraieb et al., 2021). SRL is a branch of DRL, specifically designed to solve open-loop problems (see Section 1.4.1).

Simplified-PPO or PPO-1 is introduced in Ghraieb et al. (2021); Lai et al. (2021). It is a version of the PPO algorithm, re-defined for SRL. The state $\boldsymbol{s_t}$ is fixed to a state $\boldsymbol{s_0}$, usually a zero vector and the discount rate γ is fixed to 1. This implies that the policy does not learn a state/action relationship, as the input state never varies. Essentially, this turns the problem into parameter optimization, where a transformation from a constant input into actions is learned. This leads to a substantial reduction in the number of parameters needed in the actor neural network. Consequently, the running of the environments accounts for nearly all the computational resources. Following any action, the environment reaches a terminal state, effectively adapting reinforcement learning to open-loop control. As a consequence of these changes, the advantage A^{Π_θ} reduces into the reward r_t in the actor loss function given by Eq. (4.7), as explained in Ghraieb et al. (2021). Because there is no longer need of the advantage function, the use of a critic network becomes obsolete in simplified-PPO, meaning it actually operates as a policy-based method. In the initial versions (Ghraieb et al., 2021; Lai et al., 2021), the action distribution is generated through mean and standard deviation, both produced by a unique neural network, from which the chosen actions are sampled. Initial implementations coupled a CFD environment with the PPO algorithm of a traditional RL library that is tuned to fit the PPO-1 requirements (as in Viquerat et al. (2020); Hachem et al. (2021); Ghraieb et al. (2021)). Coupling of the CFD and RL codes is either performed via an input/output IO stream (Ghraieb et al., 2021), MPI processes (Guastoni et al., 2023), or using code coupling frameworks such as Precice (Shams and Elsheikh, 2023).

In Viquerat et al. (2023), an updated version of PPO-1 is introduced, named Policy-Based Optimization (PBO). It takes inspiration from CMA-ES (Hansen, 2023), which is a leading method for evolutionary stochastic optimization of non-linear and non-convex functions. Indeed, PBO introduces a covariance matrix to compute the action distribution, in a similar fashion to the PPO version of Hämäläinen et al. (2018).

Figure 4.3 illustrates the PBO-LES framework, showcasing the interactions between the SRL, LES, and optimization process. In PBO, the single actor network of PPO-1 is replaced by three separate networks. These networks account for the mean $\boldsymbol{\mu}$, the standard deviation $\boldsymbol{\sigma}$, and the coefficients $\boldsymbol{c}$ used to compute the non-diagonal elements of the covariance matrix, which parameterize the action distribution as a multivariate Gaussian $\pi_\theta(\boldsymbol{a})$. LeCun normal weights are used for initialization, with zero bias for each layer. The $\boldsymbol{\mu}$ network employs ReLU activation functions for its hidden layers and hyperbolic tangent (tanh) for its output layer, while the $\boldsymbol{\sigma}$ and $\boldsymbol{c}$ networks utilize tanh and sigmoid activations, respectively, for their hidden and output layers. These activation functions, combined with the weight initialization, lead to the initial distribution parameters depicted in Figure 4.3. Figure 4.5 will further illustrate how these parameters ensure an efficient initial distribution for the early stages of exploration.

The covariance matrix $\boldsymbol{\Sigma}$ utilized in the policy distribution is computed using

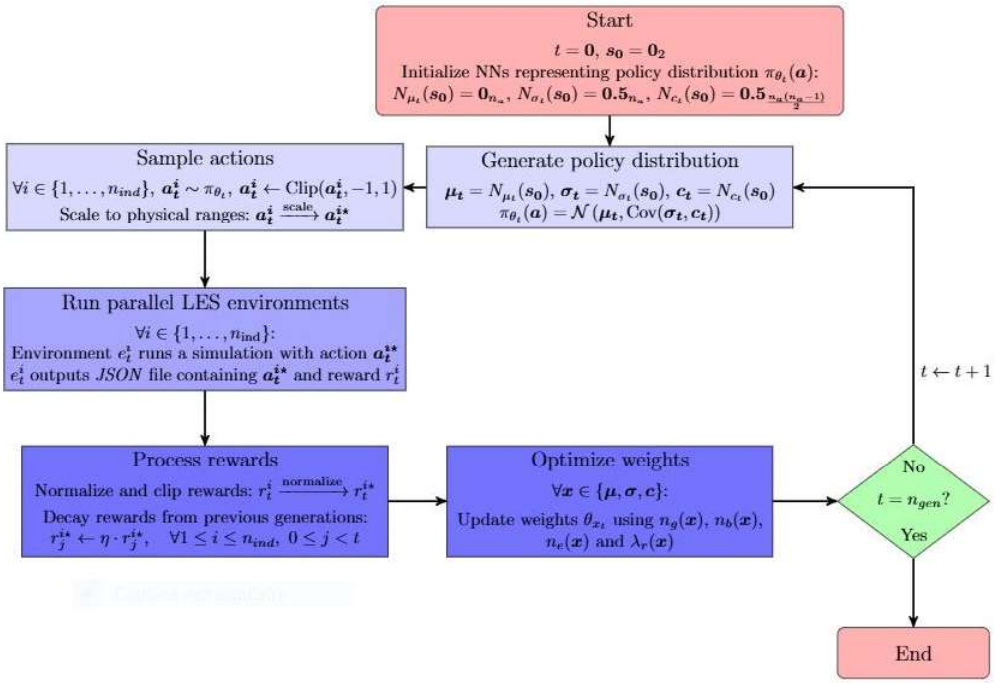

Figure 4.3: Flow chart of the Policy-Based Optimization cycle. Red rounded rectangles represent Start/End, blue rectangles represent Processes, green diamonds represent Decisions, and arrows indicate the flow of control between steps. Light blue represents the reinforcement learning processes, medium blue corresponds to the LES environments, and dark blue highlights the optimization steps. N_{x_t} represents the neural network corresponding to parameter $\boldsymbol{x}$, with θ_{x_t}, $n_g(\boldsymbol{x})$, $n_e(\boldsymbol{x})$, $n_b(\boldsymbol{x})$, and $\lambda_r(\boldsymbol{x})$ denoting its weights for generation t, number of previous generations included, number of epochs, number of batches, and learning rate, respectively. $\boldsymbol{a_i^t}$, e_i^t, and r_i^t denote the action, environment, and reward of individual i from generation t, respectively. n_a indicates the length of the action vector (or number of parameters) per individual, while n_{ind} and n_{gen} represent the number of individuals per generation and the number of generations, respectively. Normalized quantities are indicated by the notation $\star$ and $\sim \pi_{\theta_t}$ indicates the sampling operation from probability distribution $\pi_{\theta_t}(\boldsymbol{a})$.

parameters $\boldsymbol{\sigma}$ and $\boldsymbol{c}$, where $\boldsymbol{\sigma}$ corresponds to the diagonal coefficients. In order to guarantee that the final matrix is symmetric definite positive, the non diagonal coefficients are computed using $\boldsymbol{c}$ with hypersphere decomposition (more details in Viquerat et al. (2023); Numpacharoen and Atsawarungruangkit (2012)). More precisely, a lower triangular matrix $\boldsymbol{B}$ is parameterized by $\boldsymbol{c} \in [0,1]^{\frac{n_a(n_a-1)}{2}}$, where each component of $\boldsymbol{c}$, denoted c_{ij}, corresponds to a unique pair $(i,j) \in \{1,\dots,n_a\}^2$ with $i > j$. The matrix $\boldsymbol{B}$ is given by:

$$[\boldsymbol{B}(\boldsymbol{c})]_{ij} = \begin{cases} 1 & \text{if } i = j = 1, \\ \cos(\pi c_{ij}) & \text{if } i > j,\ j = 1, \\ \cos(\pi c_{ij}) \prod_{k=1}^{j-1} \sin(\pi c_{ik}) & \text{if } i > j > 1, \\ \prod_{k=1}^{j-1} \sin(\pi c_{ik}) & \text{if } i > 1,\ i = j \\ 0 & \text{otherwise.} \end{cases} \tag{4.8}$$

For example, with $n_a = 2$, which is relevant to the present work, the following matrix $\boldsymbol{B}$ is obtained, parameterized by only one parameter:

$$\boldsymbol{B}(c_{21}) = \begin{bmatrix} 1 & 0 \\ \cos(\pi c_{2,1}) & \sin(\pi c_{2,1}) \end{bmatrix}. \tag{4.9}$$

The product $\boldsymbol{B}(\boldsymbol{c})\boldsymbol{B}(\boldsymbol{c})^\top$ yields a valid correlation matrix, which is symmetric and positive semi-definite, with entries in $[-1,1]$.

The final covariance matrix of the policy distribution, at iteration (or generation) t, incorporating variance $\boldsymbol{\sigma_t}$ is then constructed as:

$$\boldsymbol{\Sigma_t} = \mathrm{Cov}(\boldsymbol{\sigma_t}, \boldsymbol{c_t}) = 0.85\ \mathrm{diag}(\boldsymbol{\sigma_t})\ (\boldsymbol{B}(\boldsymbol{c_t})\boldsymbol{B}(\boldsymbol{c_t})^\top)\ \mathrm{diag}(\boldsymbol{\sigma_t}), \tag{4.10}$$

where $\mathrm{diag}(\boldsymbol{\sigma_t})$ produces a diagonal matrix with the elements of $\boldsymbol{\sigma_t}$ as its diagonal entries. The coefficients are scaled down by a factor of 0.85. This results in a more expressive representation of the action distribution compared to PPO-1, which used a diagonal covariance matrix. In PBO, the covariance updates align the density of the sampling distribution with the direction of steepest descent.

Similar to evolutionary methods, an environment running with action $\boldsymbol{a}$ is referred to as an individual. The collection of environments that run simultaneously before the policy is updated is termed a generation. Therefore, n_{ind} and n_{gen} represent the number of individuals per generation and the number of generations, respectively. As illustrated in Figure 4.3, for each generation, the policy samples n_{ind} actions from its distribution. These actions are first clipped to the range $[-1,1]$ and then mapped to their relevant physical ranges using MinMax scaling.

For each of these n_{ind} scaled actions, a LES environment using the S_1 configuration (as described in Section 4.2.1) is executed. The exchange of data between the Fortran implicit-LES environment and the SRL Python code is facilitated through a standardised input/output (IO) stream. At each time step, the SRL agent provides

the scaled control actions, i.e. the amplitude and period of wall oscillations, which are then written into the LES input files. The LES simulations are then initiated and run either sequentially or in parallel, depending on the system configuration. Upon completion of each simulation, the relevant flow state variables, i.e. the corresponding control parameters (T^+ and W^+) and resulting rewards (P_{sav} or P_{net}), are encapsulated in a *JSON* file. This file acts as a container for the most important information required to assess the effectiveness of the control strategy in achieving the desired flow characteristics, such as drag reduction or net energy balance. The SRL algorithm then extracts the reward values from these *JSON* files and uses them to update the control policy, iteratively refining the strategy based on accumulated experience. This iterative feedback loop between the LES environment and the SRL agent enables the discovery of optimal control policies tailored to the specific flow configuration and objectives.

Aside from the introduction of the covariance matrix, PBO also introduces additional improvements over PPO-1. After collecting the rewards, they are normalized within their respective generations, with negative rewards being clipped to zero for enhanced stability:

$$\forall i \in \{1, \dots, n_{\text{ind}}\}, \quad r_i^{t\star} = \max\left(\frac{r_i^t - \mu_t^r}{\sigma_t^r}, 0\right), \tag{4.11}$$

where μ_t^r and σ_t^r are the mean and standard deviation of the rewards r_i^t for generation t, defined as $\mu_t^r = \frac{1}{n_{\text{ind}}}\sum_{i=1}^{n_{\text{ind}}} r_i^t$, $\sigma_t^r = \sqrt{\frac{1}{n_{\text{ind}}}\sum_{i=1}^{n_{\text{ind}}}(r_i^t - \mu_t^r)^2}$.

Furthermore, rewards from the last $n_g(\boldsymbol{x})$ generations, where $\boldsymbol{x} \in \{\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{c}\}$, are incorporated into the training process, with the inclusion of an exponential decay parameter: $\eta = 1 - e^{-\alpha n_a}$, with $\alpha \in \mathbb{R}^{+*}$. The hyper-parameter n_g may differ for the neural networks corresponding to $\boldsymbol{\mu}$, $\boldsymbol{\sigma}$, and $\boldsymbol{c}$. This inclusion of decayed data from previous policies enhances the reliability of the loss evaluation. The decay rate is slower for larger n_a, as more samples are needed to form a coherent covariance matrix. Consequently, PBO effectively becomes an off-policy algorithm.

The off-policy version of PBO is considered more robust because the inclusion of rewards from previous generations significantly increases the number of training samples. The larger dataset used for computing the loss function helps reduce the variance in policy updates, leading to more stable and efficient learning, whereas on-policy methods tend to suffer from lower sample efficiency. This is especially true for open-loop control, where the number of samples per generation is typically much lower than the number of samples within an episode in traditional reinforcement learning.

The loss function for a parameter $\boldsymbol{x} \in \{\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{c}\}$ at generation t, implemented in the code, is given by Eq. (4.12). In this equation using similar notation to Eq. (4.7), $\mathcal{D}_b(\boldsymbol{x})$ represents the buffer containing all the action/normalized reward pairs $(\boldsymbol{a}, r^\star)$ within a batch b, which is sampled from the last $n_g(\boldsymbol{x})$ generations:

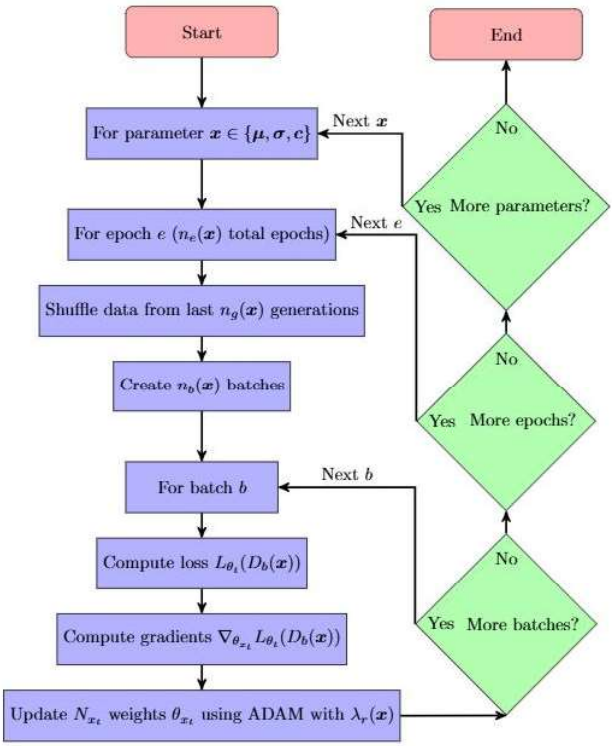

Figure 4.4: Flow chart of the optimization process. Red rounded rectangles represent Start/End, blue rectangles represent Processes, green diamonds represent Decisions, and arrows indicate the flow of control between steps. N_{x_t} represents the neural network corresponding to parameter $\boldsymbol{x}$, with θ_{x_t}, $n_g(\boldsymbol{x})$, $n_e(\boldsymbol{x})$, $n_b(\boldsymbol{x})$, and $\lambda_r(\boldsymbol{x})$ denoting its weights for generation t, number of previous generations included, number of epochs, number of batches, and learning rate, respectively. L_{θ_t} is the loss function at generation t and $\mathcal{D}_b(\boldsymbol{x})$ is the buffer containing actions and normalized rewards from batch b and parameter x.

$$L_{\theta_t}(\mathcal{D}_b(\boldsymbol{x})) = \mathbb{E}_{(\boldsymbol{a},r^\star)\sim\mathcal{D}_b(\boldsymbol{x})}\left[\log(\pi_{\theta_t}(\boldsymbol{a}))r^\star\right]. \tag{4.12}$$

The clipping component and probability ratio in the PPO loss function (Eq. (4.7)) are removed here due to a tendency to cause unstable updates in the final stages of the optimization process, as attested in Viquerat et al. (2023). The probability density function $\pi_{\theta_t}(\boldsymbol{a})$ can be efficiently computed using the Cholesky decomposition of the covariance matrix $\boldsymbol{\Sigma_t}$, and the mean $\boldsymbol{\mu_t}$, from which its log-probability can also be derived.

The optimization process for a fixed generation t is described in detail in Figure 4.4. For all $\boldsymbol{x} \in \{\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{c}\}$, the neural networks N_{x_t} are optimized separately, each with its own learning rate $\lambda_r(\boldsymbol{x})$, number of previous generations $n_g(\boldsymbol{x})$, number of epochs $n_e(\boldsymbol{x})$, and number of batches $n_b(\boldsymbol{x})$. The gradients of each loss function (Eq. (4.12)) are computed efficiently using the log-probability density function of $\pi_{\theta_t}(\boldsymbol{a})$. These gradients are then used to optimize the weights θ_{x_t} of each neural network N_{x_t} via the ADAM algorithm (Kingma and Ba, 2014). This entire process,

from action sampling to neural network optimization, is repeated iteratively over n_{gen} generations.

Ultimately, PBO can be considered an evolution strategy that uses gradient descent on the weights of the actor neural networks as an update rule. As discussed in Section 4.1, this method has proven to be competitive with state-of-the-art parameter optimization methods, which justifies its utilization in the present study.

4.3.3 PBO configuration

The open-source code of Viquerat et al. (2023) is utilized for the PBO implementation. The method comes with many hyper-parameters due to the use of 3 different neural networks to generate the action distribution. Table 4.3 summarizes such hyper-parameters, where the number of individuals per generation is set to 8 (n_{ind}), with 4 simulations running simultaneously in parallel (n_{para}), on 16 CPUs each (n_{CPUs}). This sets the time of one generation to approximately 384 CPU hours, corresponding to a wall-time of 6 hours, due to parallelization. This relatively high number of individuals per generation was selected due to the importance of this parameter on the quality of convergence, portrayed in Viquerat et al. (2023). As stated previously, each neural network includes its own hyper-parameters, such as the learning rate (step size at each iteration), the number of epochs (number of complete passes through the entire training dataset) and the number of batches (number of subsets the training data is divided into during one epoch). The values of these hyper-parameters (see Table 4.3) correspond to the default values available in Viquerat et al. (2023), from which they obtained excellent results. The decay parameter α is set to a relatively low value of 0.5, while the neural networks use their default architecture of 3 hidden layers with 2 neurons each. Consequently, the architectures were not specifically optimized for this problem. The models are not pre-trained and use the weight initialization previously described.

n_{ind}	n_{para}	n_{CPUs}	α	$\lambda_r(\boldsymbol{\mu},\boldsymbol{\sigma},\boldsymbol{c})$	n_e $(\boldsymbol{\mu},\boldsymbol{\sigma},\boldsymbol{c})$	n_b $(\boldsymbol{\mu},\boldsymbol{\sigma},\boldsymbol{c})$	n_g $(\boldsymbol{\mu},\boldsymbol{\sigma},\boldsymbol{c})$
8	4	16	0.5	$(5,5,1)\times 10^{-3}$	(128,16,16)	(1,4,8)	(1,8,16)

Table 4.3: PBO training configuration. λ_r is the learning rate, n_e is the number of epochs, n_b is the number of batches and n_g is the number of most recent generations included in each training iteration.

As reported in Section 4.1, this work considers two case studies, utilizing PBO: optimal drag reduction and optimal net energy balance. Predictably, the first case uses P_{sav} as a reward function, while the latter uses P_{net}. For both cases, the action vector (or parameters) corresponds to the oscillation period and amplitude, described in Section 1.2.4. The initial action distribution is illustrated by Figure 4.5. The joint PDF is normalized, as is customary, to ensure a unity integral over the action domain. It is centered within the imposed action boundaries, with a diagonal covariance matrix and a significantly high variance, imposed to favour ini-

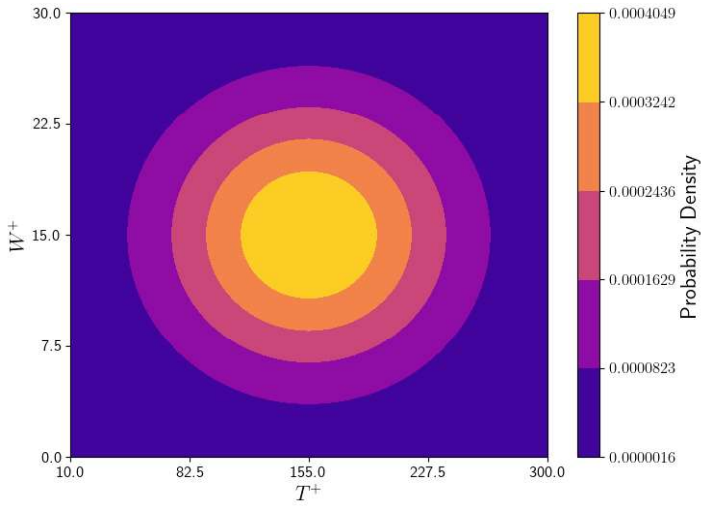

Figure 4.5: Initial joint probability density function (PDF) for control parameters T^+ and W^+. This centered distribution with diagonal covariance matrix arises from the neural networks initialization.

tial exploration. The oscillation parameters are sampled from a continuous action space, with $T^+ \in [10, 300]$, and $W^+ \in [0, 30]$. These boundaries verify the constraints $T^+_{\text{min}} > 125\Delta t$, $W^+_{\text{max}}\Delta t/\Delta z < 0.2$, and $t^+_{\text{collect}} \geq 5T^+$, to ensure stability and convergence. The parameter ranges were selected to closely align with those investigated by Quadrio and Ricco (2004) for comparison, where they demonstrate that the optimal parameters lie within these ranges.

4.4 Optimal drag reduction

This section aims to utilize the Policy-Based Optimization algorithm within an implicit Large Eddy Simulation environment to identify the optimal parameters for maximising drag reduction in a turbulent channel flow at $\text{Re}_\tau \approx 200$. This method is expected to convey a precise observation of the parameter space towards the optimum while maintaining a low computational cost due to the implicit LES environment. Based on the literature (Quadrio and Ricco, 2004; Gatti and Quadrio, 2013), it is anticipated that the optimal oscillation period T^+_{opti} lies within the range $[100, 125]$, while the optimal amplitude W^+_{opti} is approximately 30 within the investigated range. Higher values of W^+ are not considered for scrutiny at this Reynolds number, as previous studies have demonstrated that this amplitude is sufficient to reach re-laminarization.

The Policy-Based Optimization algorithm converged to a precise region of the parameter space within 50 generations, as illustrated in Figure 4.6(b), which shows the generation of each individual within the parameter space. The high initial variance of the policy distribution promotes initial exploration. The rapid convergence of the method towards the optimal region enables a thorough examination of the

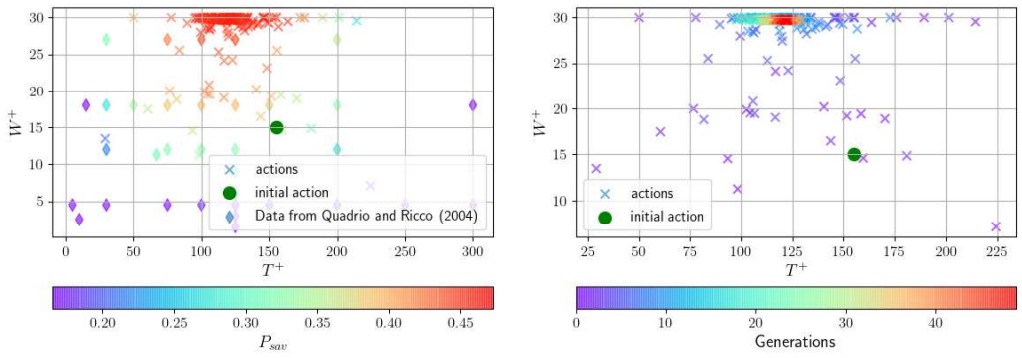

(a) (b)

Figure 4.6: Results for optimal drag reduction in the action space (T^+, W^+). Colormap corresponds to: (a) P_{sav}, (b) corresponding generation. Crosses indicate results obtained in the current optimization process and diamonds indicate reference data from Quadrio and Ricco (2004). Green round indicates the mean of the initial action probability distribution.

parameter space in this area, as anticipated, while bypassing the less productive regions. This corresponds to a total of 400 simulations, resulting in a total wall-time of approximately 300 hours, corresponding to approximately 20 000 CPU hours. The cost of only one DNS simulation, with the S_2 configuration requires 0.89 times fewer CPU hours, while being around 20 times faster in wall-time as visible in Table 4.1. This demonstrates the significant resource savings achieved through the use of implicit LES, while keeping a remarkably satisfying precision with regards to the obtained drag reduction, as demonstrated in Figure 4.6(a). Indeed, this figure compares the obtained drag reduction of each individual with the data of Quadrio and Ricco (2004). An excellent agreement is reached, further validating the accuracy of the proposed method.

The convergence of the reward to the maximum drag reduction is significantly faster than the convergence of the population as a whole. This is evident from Figure 4.7(a), which illustrates the time evolution of the best and averaged P_{sav} individual for each generation. The reward converges within 10 generations, reaching $P_{\text{sav}} \approx 0.47$ and $P_{\text{sav}} \approx 0.46$ for the best and averaged values, respectively, beyond which no additional gains are obtained in the subsequent generations. The variations in reward observed from the 10th generation are due to the stochastic nature of the problem, coupled with the short finite time intervals used for statistical collection, as stated in Section 4.2.1. One possible explanation for the faster convergence towards optimal drag reduction is the wide range of periods that allow for near-optimal parameters, as visible in Figure 4.6(a). The provided colormap indicates similar results for $T^+ \in [100, 150]$, while the optimal W^+ is clearly near the maximum value of 30, corresponding to the maximum Stokes layer thickness at the optimal period.

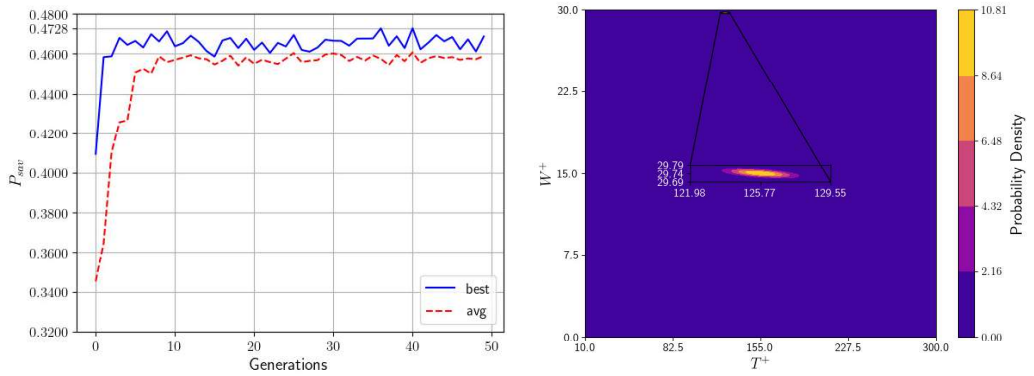

(a) (b)

Figure 4.7: PBO convergence for drag reduction, (a): P_{sav} evolution through generations. Blue line indicates the best P_{sav} individual within the generation, while the red dashed line indicates the average P_{sav} individual of the generation. (b): T^+-W^+ joint probability density function of the converged policy distribution.

This rapid convergence to the optimal drag reduction suggests that the drag is less sensitive to variations in the oscillation period within a certain range, provided that the amplitude is sufficiently high. This result is consistent with the observations of Quadrio and Ricco (2004), who reported that the drag reduction is almost independent of the oscillation period near the optimum when the maximum wall displacement is fixed.

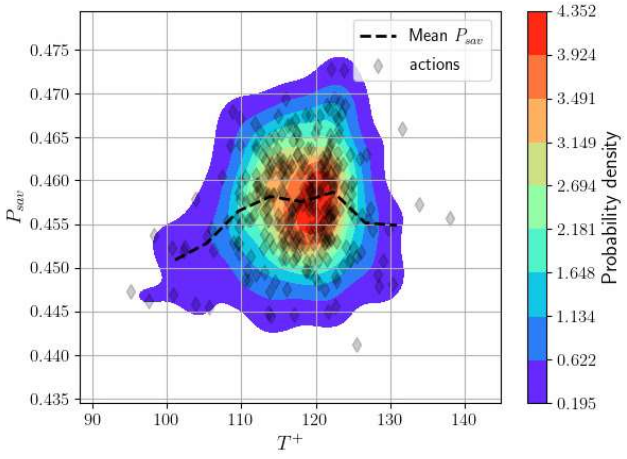

Figure 4.8: Joint probability density function of T^+ and P_{sav} around the optimal region, for actions with $W^+ > 29.5$ and $P_{\text{sav}} > 0.435$. The values of the colormap are obtained from Kernel Density Estimation (KDE). Black dashed line indicates the mean P_{sav} values varying in T^+, while grey diamonds indicate corresponding actions.

A closer examination is performed in Figure 4.8 to identify the precise range of periods resulting in optimal drag reduction. This figure displays the actions corresponding to $W^+ > 29.5$ and $P_{\text{sav}} > 0.435$, along with their joint probability density

function in terms of T^+ and P_{sav}, estimated using Kernel Density Estimation (KDE). A Gaussian kernel is applied to each data point, with the contributions summed to produce a smooth, continuous density estimate. As a non-parametric method, KDE does not assume a specific underlying distribution, providing a flexible, data-driven approach for estimating the T^+-P_{sav} distribution. The black dashed line represents the mean P_{sav} values for each T^+. The shape of the distribution clearly delineates the optimal region, characterised by superior mean and maximal values. The mean line emphasises that any period within the narrow range of $[110, 125]$ can be considered optimal. The period obtained after full convergence, $T^+_{\text{conv}} \approx 125$, falls within this range, as visible in Figure 4.7(b), and corresponds to the highest P_{sav} achieved during the optimization process. This figure displays the T^+-W^+ joint probability density function of the converged policy distribution. The shape of the density emphasizes the importance of the covariance matrix, not only in accelerating convergence but also in accurately representing the final distribution. These results align with the conclusions of Quadrio and Ricco (2004); Gatti and Quadrio (2013), where the best drag reduction is obtained for $T^+ \approx 100$.

Increasing W^+ appears to reliably enhance drag reduction up to the point of relaminarization, at least within the optimal period range. This presents challenges for practical applications, as higher wall-motion speeds could further complicate real-world implementations. Considering the control cost, however, provides a more realistic perspective and is likely to result in different optimal amplitude ranges.

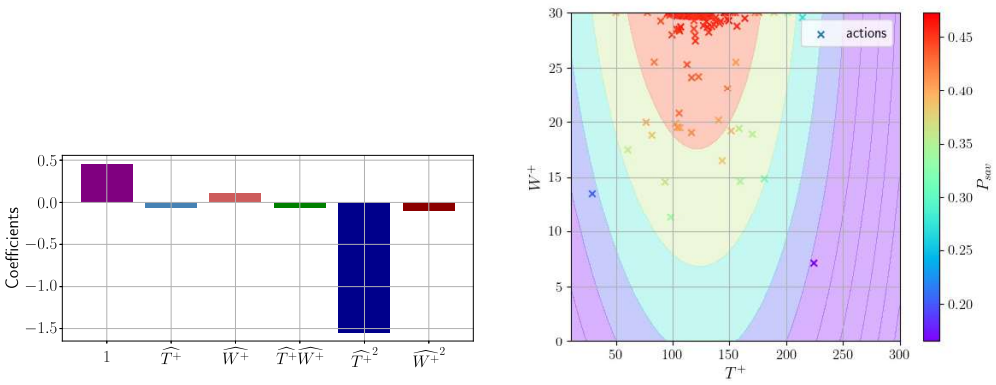

(a) (b)

Figure 4.9: Linear regression results with a 2D quadratic polynomial, using parameters $(\widehat{T^+}, \widehat{W^+})$ as inputs. (a) Magnitude of the obtained coefficients, (b) P_{sav} results in the action space. Contours represent the linear regression estimation, while crosses denote the actions explored during PBO.

To investigate the sensitivity of the system to the parameters T^+ and W^+, a linear regression analysis is performed using a 2D quadratic polynomial. This approach aims to fit the P_{sav} results within the action space by applying linear regression, using a function of the form:

$$P_{\text{sav}} \approx a_1 + a_2\widehat{T^+} + a_3\widehat{W^+} + a_4\widehat{T^+}\widehat{W^+} + a_5\widehat{T^+}^2 + a_6\widehat{W^+}^2. \tag{4.13}$$

$\widehat{\chi}$ indicates a variable χ scaled as $\widehat{\chi} = \frac{\chi - \chi_{\text{conv}}}{\chi_{\text{max}}}$, where χ_{max} is the maximum bound of χ and χ_{conv} is the PBO converged optimal value. This regression formulation effectively evaluates the sensitivity in P_{sav} to parameters T^+ and W^+, when straying from the optimal. All individuals obtained during the optimization process are used for training the regression model. Figure 4.9(a) highlights the magnitude of the coefficients a_i obtained for each term. Regarding the influence of W^+, the scaled action $\widehat{W^+}$ is almost always negative, and the terms $\widehat{W^+}$ and $\widehat{W^+}^2$ act jointly to reduce P_{sav} when W^+ is below W^+_{conv}. The quadratic term $\widehat{T^+}^2$ is characterized by a significant amplitude compared to the other terms. This higher sensitivity with regards to T^+ compared to W^+ is visible in the contour plot of Figure 4.9(b), which depicts the linear regression results in the action space. Additionally, it is worth noting that the quadratic approximation performs exceptionally well in predicting the values of P_{sav} across the available actions, with the average error in $\%P_{\text{sav}}$ remaining below 1%. Interestingly, the shape of the ellipse in Figure 4.9(b), along with the remarkably small coefficient for the interaction term $\widehat{T^+}\widehat{W^+}$, suggests an extremely low correlation between T^+ and W^+ for estimating P_{sav}. This implies that the two parameters could almost be optimized independently, further supporting their distinct roles in the optimization process.

The presented analysis offers a comprehensive examination of the drag reduction optimization. The method enabled a detailed exploration of the optimal region while circumventing the investigation of under-performing areas. This approach incurred a significantly lower computational cost compared to conducting a parametric study using DNS, with minimal loss in accuracy. Moreover, a thorough inquiry demonstrated that any period within the range $[110, 125]$ may be regarded as optimal. Therefore, for this particular case, it may not be necessary to converge to a specific point in the parameter space. Parameter sensitivity analysis revealed that the period T^+ plays a dominant role in influencing P_{sav} through its quadratic term. The interaction term has a negligible effect, suggesting both parameters may be optimized separately.

4.5 Optimal net energetic balance

The previous section determined optimal parameter ranges for drag reduction. However, the associated control cost was not considered, making this approach unrealistic. In this section, the energy cost is accounted for through the P_{req} term (see Eq. (4.2)), using P_{net} as the reward. Apart from this penalization in the reward function, the configuration remains the same as in the drag reduction case, with identical parameter ranges and LES at $\text{Re}_\tau \approx 200$. Increasing the period or decreasing the amplitude lowers the control cost, suggesting that the high-amplitude region which is optimal for drag reduction may no longer be ideal. Indeed, in Quadrio and

Ricco (2004), only low amplitudes ($W^+ \leq 6$), yielded a positive P_{net}, for periods in the $[75, 300]$ range, where the best net energetic gain $P_{\text{net}} = 7.7\%$ is reached for $T^+ = 125$ and $W^+ = 4.5$. Comparatively to the drag reduction case, simulations in the positive P_{net} region are fewer in Quadrio and Ricco (2004). The current method is therefore anticipated to offer deeper insights into the parameters that maximize the net energy balance, potentially leading to a higher P_{net}.

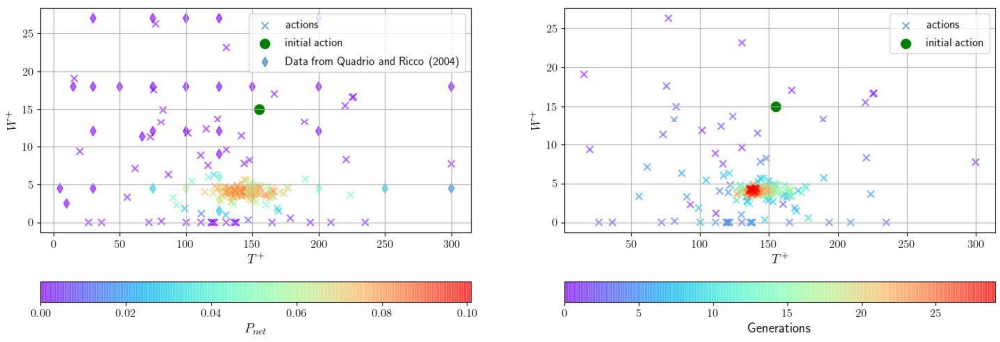

(a) (b)

Figure 4.10: Results for optimal net energetic balance in the action space (T^+, W^+). Colormap corresponds to: (a) P_{net} values, (b) corresponding generation. The minimal value of the P_{net} colormap is fixed to 0, highlighting the positive P_{net} region. Crosses indicate results obtained in the current optimization process and diamonds indicate reference data from Quadrio and Ricco (2004). Green round indicates the mean of the initial action probability distribution.

The optimization process converged to a precise location in the parameter space within 30 generations, as illustrated in Figure 4.10(b). Similar to the previous case, the initial generations broadly explored the parameter space, while the later iterations focused on refining the search in the optimal region. The optimal parameters derived from the converged mean action distribution (see Figure 4.11(b)) are $T^+_{\text{conv}} \approx 143$ and $W^+_{\text{conv}} \approx 4.22$, which are consistent with the previous results (Quadrio and Ricco, 2004). The converged density function reveals a more pronounced variance for the period than for the amplitude. The optimization process involved 240 LES runs, totalling approximately 180 wall-time hours and about 12 000 CPU hours. In comparison, a single Direct Numerical Simulation using the S_2 DNS configuration required 1.5 times more CPU hours (see Table 4.1), emphasizing the significant computational savings achieved with LES. Despite these substantial savings, the accuracy of the results remains intact, as demonstrated in Figure 4.10(a), which presents the P_{net} obtained for each individual alongside the results from Quadrio and Ricco (2004), confirming the strong consistency between the two studies. This underscores the robustness of the LES configuration in accurately predicting both drag reduction and the associated cost.

Figure 4.11(a) illustrates the evolution of the reward through generations. The best reward reaches $P_{\text{net}} = 9.4\%$, outperforming the previous best result reported by

Quadrio and Ricco (2004) by 1.8%, while the average reward of each generation stabilizes at $P_{\text{net}} \approx 8\%$. No significant gain in the best or averaged reward is observed after the 10th generation, highlighting that obtaining precise optimal parameters requires additional iterations due to the presence of a large near-optimal region. This is particularly relevant since the optimal W^+ range is broader than in the drag reduction case, where the highest value is clearly optimal. The reward curve is notably smoother compared to the drag reduction case, indicating greater stability in the results when accounting for P_{req}. This likely explains why convergence was achieved after 30 generations, compared to 50 in the drag reduction scenario.

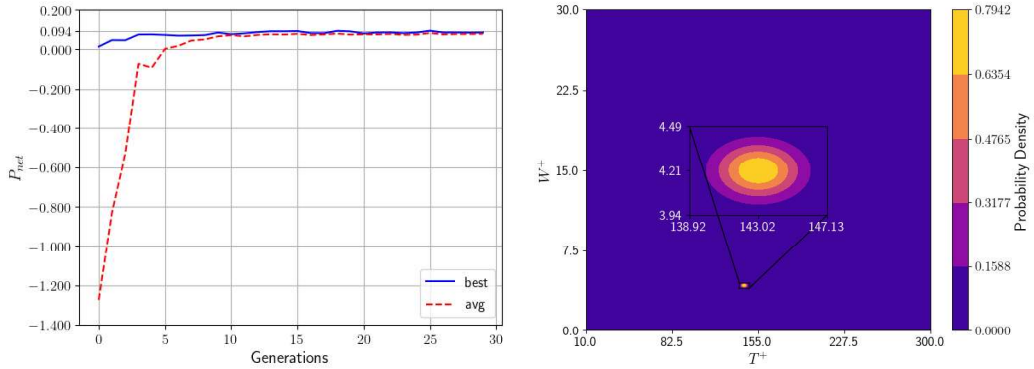

(a) (b)

Figure 4.11: PBO convergence for net energetic gain, (a): P_{net} evolution through generations. Blue line indicates the best P_{net} individual within the generation, while the red dashed line indicates the average P_{net} individual of the generation. (b): T^+-W^+ joint probability density function of the converged policy distribution.

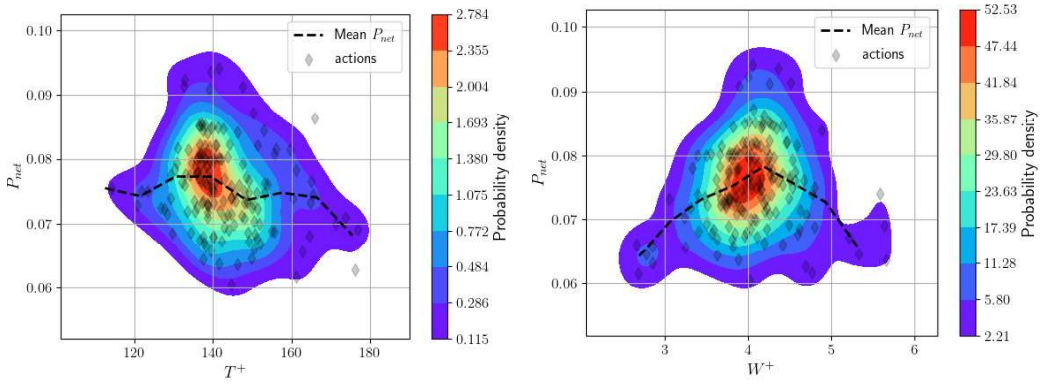

(a) (b)

Figure 4.12: Joint probability density function around the optimal region of (a) T^+ and P_{net}, for actions with $W^+ \in [3, 6]$ and $P_{\text{net}} > 0.06$, (b) W^+ and P_{net}, for actions with $T^+ \in [75, 250]$ and $P_{\text{net}} > 0.06$. Colormap is obtained from Kernel Density Estimation (KDE). Black dashed line indicates the mean P_{net}, while grey diamonds indicate corresponding actions.

From Figure 4.10(a), it is evident that the optimal region corresponds to low amplitudes ($W^+ \leq 5$) and a relatively wide range of periods, with $T^+ \in [115, 170]$. The low amplitude was expected due to its direct impact on control cost. The period range is slightly shifted to the right compared to the drag reduction case, which is coherent since the control cost decreases with an increasing period. These results align with expectations from Quadrio and Ricco (2004), while providing a more detailed exploration of the parameter space. Figure 4.12(a) presents the T^+-P_{net} joint probability density function in the near-optimal region. The distribution is narrower than in the drag reduction case (see Figure 4.8). From the mean P_{net} line, the optimal period range, $T^+_{\text{opti}} = [125, 155]$, can be deduced, confirming the period shift induced by the control cost. The converged period, $T^+_{\text{conv}} \approx 143$, falls within this range. Regarding the W^+-P_{net} distribution in the near-optimal region exhibited by Figure 4.12(b), the mean line indicates an even narrower optimal range. The optimal amplitude is estimated to lie between $[3.5, 5]$, which is consistent with $W^+_{\text{conv}} \approx 4.22$, aligning well with the highest mean P_{net} value. The averaged P_{net} at fixed parameters, estimated from the mean line, reaches a best value close to the 7.7% found by Quadrio and Ricco (2004), indicating that the current optimal parameters do not consistently exceed this benchmark. Therefore, while the best simulation yielded an improvement of 1.8% compared to Quadrio and Ricco (2004), this appears to be an outlier for the given oscillation parameters. The mean value is a more reliable indicator due to the uncertainty of the statistical convergence of a single simulation.

The lower range of optimal SWO amplitudes facilitates practical applications compared to the drag reduction scenario. Similarly, the higher range of optimal periods, when considering control costs, aligns with this trend, as lower frequencies are easier to manage in real-world implementations. These differences between the two considered cases are intuitive, as increasing W^+ or decreasing T^+ inherently raises the control cost.

As in Section 4.4, linear regression is performed on a 2D quadratic function of $(\widehat{T^+}, \widehat{W^+})$, estimating P_{net}, as shown in Figure 4.13. In practice, a model predicting P_{sav} is combined with an additional model predicting P_{req}. The coefficients in Figure 4.13(a) reveal once more a significant magnitude for the quadratic terms, especially for the $\widehat{W^+}^2$ coefficient. This result aligns with the strong influence of the oscillation amplitude on control cost and differ from the P_{sav} case, in which sensitivity towards T^+ is stronger. Another notable difference is the non-negligible interaction term, which is characterized by the highest amplitude, and corresponds to the ellipse orientation in Figure 4.13(b). This indicates that, when considering P_{net}, the two parameters exhibit highly correlated behavior, thus justifying the need for joint optimization. Unlike the P_{sav} linear regression, the predictions for control cost are less accurate, as seen in Figure 4.13(b). This lower accuracy is expected, given that control cost is related to the product of wall spanwise velocity w_{w} and spanwise friction $\tau_{w,z}$, which is not precisely estimated by the chosen 2D polynomials.

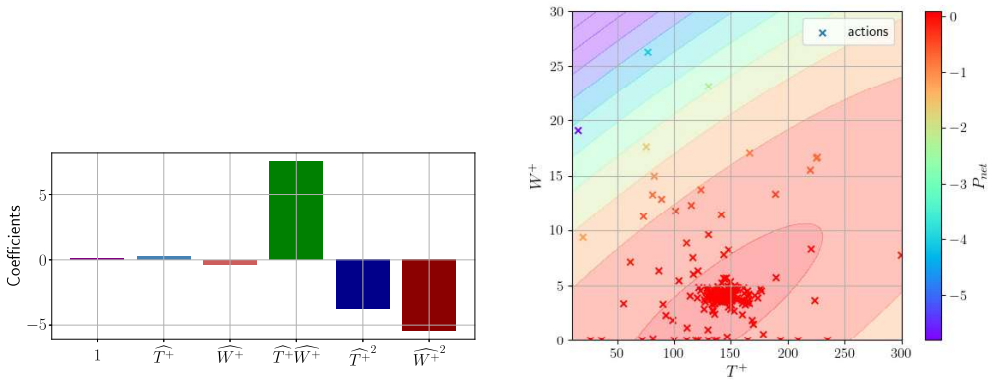

(a) (b)

Figure 4.13: Linear regression results with a 2D quadratic polynomial, using parameters $(\widehat{T^+}, \widehat{W^+})$ as inputs. (a) Magnitude of the obtained coefficients, (b) P_{net} results in the action space. Contours represent the linear regression estimation, while crosses denote the actions explored during PBO.

A detailed analysis of the optimization results was conducted, this time incorporating control cost in energy via spanwise friction power, with the goal of maximizing the net energy balance. As expected, this method provided a much more refined exploration of the parameter space compared to Quadrio and Ricco (2004), particularly in the near-optimal region. Additionally, it achieved a significantly lower computational cost than the drag reduction case. Despite the substantial savings in computational resources, the loss in precision remains minor. A thorough investigation identified the optimal parameter ranges as $T^+_{\text{opti}} \in [125, 155]$ and $W^+_{\text{opti}} \in [3.5, 5]$, with the best averaged P_{net} at fixed parameters matching the best case of Quadrio and Ricco (2004). Linear regression demonstrated the optimization of P_{net} requires joint consideration of both parameters due to the significant contribution of the interaction term. Furthermore, contrary to the drag reduction scenario, the impact of W^+ on the control cost is slightly more prominent than the one of T^+, highlighting the significantly different behavior when accounting for P_{req}.

4.6 Conclusions & Perspectives

This study investigates the control of turbulent channel flows at $\text{Re}_\tau = 200$ using spanwise wall oscillations. The primary objective was to demonstrate the effectiveness of the PBO algorithm, integrated with an Implicit LES approach, in identifying optimal control parameters for energy-efficient drag reduction. The PBO-LES framework successfully navigates the parameter space, achieving high precision in convergence while maintaining low computational costs. The well-established case of spanwise wall oscillations for drag reduction serves as a benchmark to validate

the effectiveness of this ML-enhanced optimization framework. The results not only corroborate previous studies but also provide new insights into the optimal conditions for achieving net energetic gains from drag reduction in turbulent channel flows.

The analysis was carried out in two stages: first, maximising drag reduction, and second, optimising the net energy balance by considering the cost of controlling spanwise friction. The main results of the study are as follows:

- The optimal oscillation parameters identified by the PBO-LES framework are in close agreement with the expectations from the existing literature, thus confirming the reliability of the framework for both the drag reduction and the net energy balance.

- The framework uses an initial high variance in the policy distribution to facilitate extensive exploration of the policy space. This is followed by rapid convergence, which allows detailed exploration of the optimal solution while avoiding less productive regions. The covariance matrix adaptation technique further accelerates convergence and improves the accuracy of the final policy distribution.

- The LES configuration requires fewer processors per environment compared to DNS, making it a more computationally efficient and highly parallelizable alternative. When combined with environment parallelization, the Implicit LES configuration significantly reduces computational cost compared to DNS while maintaining high accuracy for the quantities of interest.

- To maximize drag reduction, the optimal oscillation period is within the range $T^+_{\text{opti}} \in [110, 125]$. Increasing the amplitude of the applied sinusoidal signal enhances drag reduction up to the threshold where flow re-laminarization occurs. However, higher amplitudes result in a rapid increase in the energy required for control, which can negatively impact the overall energetic balance.

- The optimization of net energy balance is observed to converge rapidly towards precise parameters when the incorporation of control costs is implemented. This integration has been demonstrated to effectively mitigate resultant uncertainty. The optimal parameter ranges for net energetic balance are $T^+_{\text{opti}} \in [125, 155]$ and $W^+_{\text{opti}} \in [3.5, 5]$, providing greater precision than previously reported studies.

- Linear regression analysis indicates that power saved (P_{sav}) is primarily influenced by the T^+ quadratic term. The negligible interaction term allows for near-independent optimization of both parameters. In contrast, net power (P_{net}) exhibits significant contributions from quadratic and interaction terms, with a slightly higher sensitivity to W^+ compared to T^+. In this case, joint

optimization is required to account for the correlated effects of T^+ and W^+ on control costs.

- Beyond reproducing known optimal control parameters, this study demonstrates the potential of machine learning algorithms to accelerate the development of novel control strategies for real-world applications. It establishes a new paradigm for the use of machine learning in fluid dynamics, highlighting its transformative potential in the design of efficient flow control techniques.

This work highlights the potential of Policy-Based Optimization combined with computationally efficient CFD environments to thoroughly explore parameter spaces and deepen the understanding of flow responses to oscillations. However, the study has limitations, including the reliance on LES, where the development of more efficient models could be beneficial. Moreover, the short time integration in LES necessitated additional generations to achieve convergence. Exploring alternative optimization algorithms for comparison with PBO could also yield further valuable insights.

Extended Results

4.7 Application to heat-transfer enhancement

In this section, the PBO-LES framework described in Section 4.3.3 is applied to heat-transfer enhancement, aiming to find the SWO parameters that maximize this objective at $\mathrm{Re}_\tau = 200$ and $\mathrm{Pr} = 1$. Contrary to the drag reduction case, studies on heat transfer control with SWO are fewer, and to the author's knowledge, no rigorous parametric investigation has yet been performed in this case. It is therefore expected that the following results may reveal meaningful information about optimal parameter ranges and sensitivities for heat-transfer enhancement, as well as potential insights into the physical mechanisms strengthening the heat transfer.

4.7.1 PBO-LES configuration

The LES configuration and mesh S_1 of Section 4.2.1 are utilized with the additional temperature equation with MBC at $\mathrm{Pr} = 1$, described in Section 2.2. This configuration has been extensively validated in terms of drag and Reynolds stresses in Section 4.2.2. The temperature second-order stresses and Nusselt number variations are also validated in Section 3.2.3. Therefore, the same configuration as in Section 3.2.2 is utilized, though with numerical dissipation increased by a factor of 2.25 through the spectral vanishing viscosity scheme of Lamballais et al. (2011). This modification is required to address small instabilities that were observed in some of the drag/heat transfer increase cases. This slightly different configuration is referred to as S_1' to avoid confusion with S_1.

Configuration	$\%\Delta\mathrm{Nu}$	$\%\Delta C_f$	$\overline{A_n}$
S_1'	13.27	5.39	1.074
S_2 (Guérin et al., 2024)	14.93	7.72	1.064

Table 4.4: Comparison in heat transfer and drag SWO induced variations, for oscillation parameters $T^+ = 500$, $W^+ = 30$, for S_1' and S_2 (Guérin et al., 2024).

The augmentation in drag and heat transfer for the parameters $T^+ = 500$, $W^+ = 30$ is validated on S_1' in Table 4.4. Strong agreement with the reference DNS data (see Table 2.1 for details on S_2) is demonstrated, albeit with slightly reduced variations in both Nu and C_f, along with an acceptable discrepancy of 0.01 in the analogy factor. The S_1' LES configuration is thus considered sufficiently accurate for the present study. Further DNS validation will be conducted in Sections 4.7.3 and 4.8.3, within the identified optimal parameter ranges to verify the validity of the obtained solution.

The PBO setup, including hyper-parameters and neural network architecture, is exactly the same as in Section 4.3.3, aiming to find the parameters T^+, W^+ that

maximize the heat-transfer actuation-induced augmentation. A reward measuring this heat-transfer enhancement can easily be defined as:

$$\mathrm{Nu}_r = \frac{\langle \mathrm{Nu}_{\mathrm{ac}} \rangle_t}{\langle \mathrm{Nu}_{\mathrm{unac}} \rangle_t}, \tag{4.14}$$

expressed as the ratio between actuated and unactuated time-averaged Nusselt numbers. Previous parametric investigations (Quadrio and Ricco, 2004; Guérin et al., 2025) suggest that periods in the range $[10, 300]$ are not promising for heat-transfer enhancement, as they do not result in an augmentation of the drag for the considered amplitudes. The upper bound of $T^+ = 1\,600$ is fixed to guarantee at least one actuation period for statistical collection, resulting in the range $T^+ \in [300, 1\,600]$ for this investigation. This period range is wide compared to Sections 4.4 and 4.5, suggesting the algorithm may need more generations to converge. Regarding the oscillation amplitude, the range $W^+ \in [10, 40]$ is selected, with the lower bound guaranteeing sufficient actuation effect and upper bound above the previous limit of 30. The computational cost per simulation/generation remains very close to the one described in Section 4.3.3, despite the resolution of an additional scalar value.

Based on Section 3.2, in which the dataset contains simulations at various amplitudes and $T^+ = 500$, it is expected that the highest amplitude $W^+ = 40$ should be optimal, as was the case for drag reduction. However, the literature does not provide good intuition about the optimal period, where it is expected that it should be sufficiently high to not interrupt the regeneration cycle of the streaks while yielding a thinning of the thermal sublayer.

4.7.2 Results

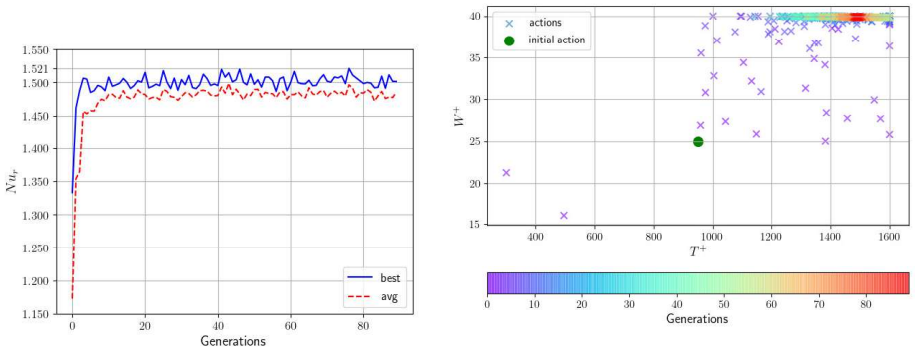

(a) (b)

Figure 4.14: PBO convergence for heat-transfer enhancement, (a): Nu_r evolution through generations. Blue line indicates the best Nu_r individual within the generation, while the red dashed line indicates the average Nu_r individual of the generation. (b): Generation of each individual within the action space T^+-W^+.

As shown in Figure 4.14, a large number of iterations are required to converge to a precise solution (nearly twice as many as in the drag reduction case) resulting in 720 simulations, corresponding to approximately 35 000 CPU-hours and a wall time of nearly 540 hours. The best-performing individual achieves a heat-transfer enhancement of 52%, while the average generation settles around 48%. The slow convergence is attributed to the broader range of oscillation periods considered, as well as the longer periods themselves, which reduce the number of actuation cycles available for statistical sampling, thereby increasing uncertainty. These results also suggest the existence of a wide range of near-optimal parameters.

Figure 4.14(b) illustrates the distribution of individuals within the parameter space. Notably, only two individuals are located in the region $T^+ < 950 \cap W^+ < 25$, demonstrating that the PBO algorithm concentrates almost exclusively on higher periods and amplitudes, which are more promising. Convergence is observed near $T^+ = 1\,480$, with particular emphasis on the range $[1\,200, 1\,600]$ during optimization, especially at the maximal amplitude $W^+ = 40$, which is confirmed as optimal.

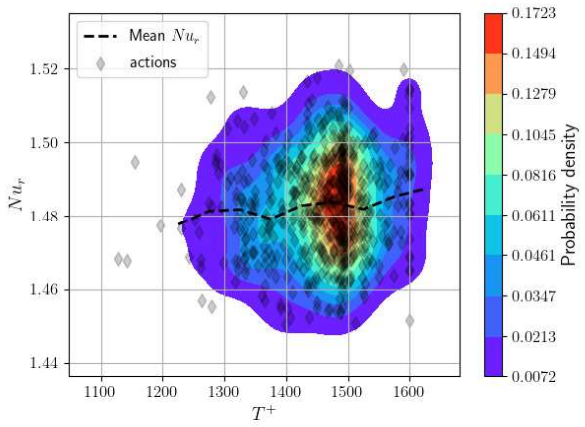

Figure 4.15: Joint probability density function of T^+ and Nu_r around the optimal region, for actions with $W^+ > 39.5$ and $\mathrm{Nu}_r > 1.45$. The values of the colormap are obtained from KDE. Black dashed line indicates the mean Nu_r values varying in T^+, while grey diamonds indicate corresponding actions.

By applying the same KDE approach as in Section 4.4, the range of periods that maximizes heat-transfer enhancement can be examined more closely. Figure 4.15 reveals a broad optimal period range ($T^+_{\mathrm{opti}} \in [1\,300, 1\,600]$), which partially explains the extended convergence time. Additionally, for similar parameters, variations in heat-transfer enhancement of up to 6.5% are observed, further supporting the hypothesis of increased statistical uncertainty at longer periods. This analysis justifies the extended number of generations required for convergence and suggests that periods $T^+ > 1\,600$ may yield further improvements.

Due to the broader range of oscillation cases considered here compared to Rouhi et al. (2025), the Stokes layer protrusion height $\ell_{0.01}$ is here defined as the wall-

normal position y where $\overline{\widehat{w}\widehat{w}}(y) = 0.01 \max_y(\overline{w''w''}(y))$ (see Eq. (2.9) for the definition of periodic fluctuations). This definition aligns with Rouhi et al. (2023) but omits their assumption about $\overline{w''w''}$ magnitude, which does not hold across all parameter ranges in the current study. The observed optimal period and amplitude within the investigated bounds suggest a strong correlation between heat transfer/drag enhancements and Stokes layer protrusion height. These enhancements appear to result from significant Stokes layer protrusion, which injects substantial energy into the smaller scales of the flow (see Rouhi et al. 2025). Consequently, for $T^+ > 1\,600$ where the Stokes layer becomes more protrusive, one would anticipate greater heat-transfer enhancement. Longer periods are also characterized by a prolonged lingering Stokes strain (or near zero Stokes strain rate) which is associated with turbulence recovery (Agostini et al., 2014; Knoop et al., 2025), thus potentially explaining the observed heat-transfer enhancement.

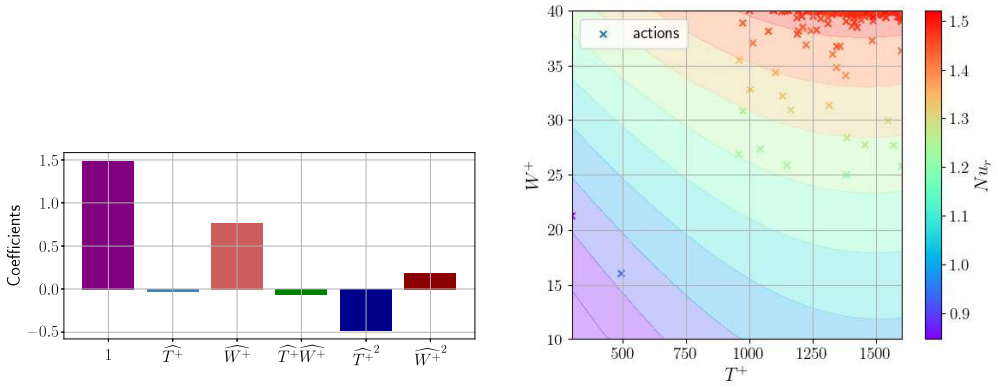

(a) (b)

Figure 4.16: Linear regression results with a 2D quadratic polynomial, using parameters $(\widehat{T^+}, \widehat{W^+})$ as inputs. (a) Magnitude of the obtained coefficients, (b) Nu_r results in the action space. Contours represent the linear regression estimation, while crosses denote the actions explored during PBO.

The linear regression analysis from Section 4.4 is applied to the heat-transfer enhancement case in Figure 4.16 to assess sensitivity to T^+ and W^+ near the optimum. Similar trends to the drag reduction case are observed, though with greater sensitivity to W^+. Indeed, while the quadratic term of the period remains dominant, the first-order term $\widehat{W^+}$ is the largest among non-constant terms. The cross-correlation term exhibits a negligible magnitude, suggesting that the optimization of period and amplitude could be decoupled. More importantly, this means that the optimal period remains stable across amplitude variations. However, the regression model's accuracy is likely confined to the region $T^+ > 950 \cap W^+ > 25$, where most evaluations are concentrated.

The optimal SWO parameters for heat-transfer enhancement have been thoroughly investigated. However, it remains unclear whether this optimum also maxi-

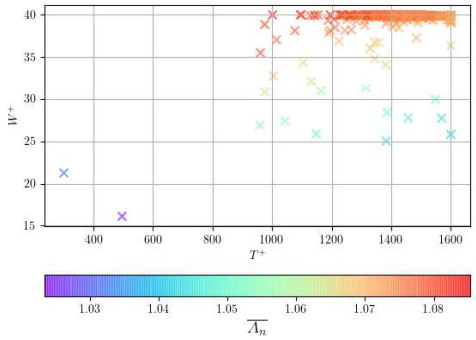

Figure 4.17: Results for optimal heat-transfer enhancement in the action space (T^+, W^+). Colormap corresponds to: $\overline{A_n}$.

mizes dissimilar heat transfer, effectively limiting the associated drag increase. Figure 4.17 displays the averaged analogy factor (Eq. (2.14)) for all individuals obtained during optimization. Lower periods consistently yield higher dissimilarity, with higher amplitudes again proving more effective. This suggests that contrary to heat-transfer enhancement, the mechanisms behind dissimilar heat transfer are more complex than trivial Stokes layer protrusion height maximisation. Thus, a separate optimization for dissimilar heat transfer is warranted, potentially focusing on lower periods compared to the currently studied range.

4.7.3 DNS validation and comparative analysis

This section aims to validate the results of the previous parametric investigation using DNS-level accuracy. Additionally, it explores the relationship between heat-transfer enhancement and Stokes layer protrusion height, while providing a comparative analysis with the results of Section 2.5.

To validate the optimal ranges identified via LES, a DNS simulation is performed at $T^+ = 1\ 400$, $W^+ = 40$ using the S_2 DNS configuration (Table 2.1) and an extended statistical collection time of $\approx 2\ 800\ t^+_{\text{nom}}$ on the S_1' configuration.

Configuration	%ΔNu	%ΔC_f	$\overline{A_n}$
S_1'	47.95	37.5	1.076
S_2	51.64	44.46	1.045

Table 4.5: Comparison in heat transfer and drag SWO induced variations, for oscillation parameters $T^+ = 1\ 400$, $W^+ = 40$, on S_1' and S_2.

The results for ΔNu, ΔC_f, and $\overline{A_n}$ are compared in Table 4.5. Enhancements are more pronounced in the DNS case, particularly for drag, where a 7% discrepancy in ΔC_f is observed. These differences remain acceptable considering the overall augmentation, supporting the reliability of the LES-derived optimal range. However,

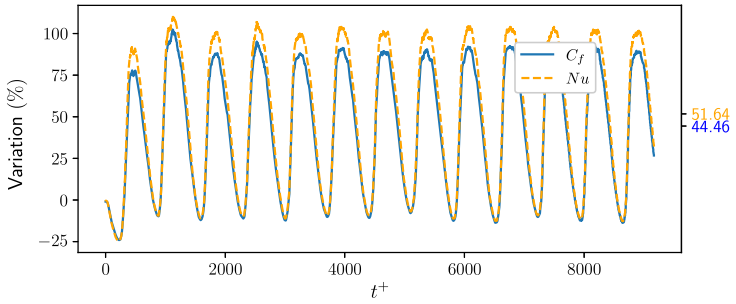

Figure 4.18: Temporal evolution of drag and heat-transfer augmentations obtained by spanwise oscillatory wall actuation for control parameters $T^+_{\text{nom}} = 1\ 400$ and $W^+_{\text{nom}} = 40$, on the S_2 DNS configuration.

the larger drag discrepancy reduces the analogy factor by 0.031 in the DNS, likely due to the S_1' mesh's limitations at elevated $\text{Re}_{\tau\text{ac}}$. This does not undermine the credibility of the results, as the focus was solely on Nusselt number optimization, which is accurately predicted by the LES. However, LES accuracy regarding the analogy factor should be carefully evaluated if dissimilar heat transfer is targeted, though this may be less critical for lower periods, where drag increases are expected to be smaller.

Temporal and Phase-Wise Dynamics of Nu and C_f

Figure 4.18 presents the temporal evolution of Nu and C_f for the S_2 DNS simulation. A transient similar to the $T^+ = 500$, $W^+ = 30$ case (Figure 2.6) is observed, followed by fluctuations between -10% and 100%, yielding average enhancements of 52% and 44% for heat transfer and drag, respectively. Dissimilarity arises primarily during peak friction phases, though less markedly than in the $T^+ = 500$ case, which is consistent with the lower analogy factor. The higher analogy factor achieved at $T^+ = 500$, $W^+ = 30$ confirms, with DNS accuracy, that the optimal period for dissimilar heat transfer must be significantly lower than for Nusselt maximization.

Protrusion height and sublayers thinning

The PBO algorithm converged to $T^+_{\text{conv}} \approx 1400$. Statistical analysis, however, showed that performance remained consistent across the T^+ range $[1\,300, 1\,600]$. This naturally leads to the question of whether increasing the period beyond the $T^+ = 1600$ threshold would improve heat-transfer performance, or if a plateau has already been reached where further increases in period no longer enhance performance.

At $\text{Re}_{\tau,\text{ac}} = 213$, the DNS results in a protrusion height of $\ell^+_{0.01,\text{ac}} = 177.2$. As this value approaches the channel half-height, it suggests that the protrusion height is near its maximum. Further increases in the oscillation period are then unlikely

to produce significant growth. Given the correlation between protrusion height and heat-transfer enhancement, this indicates that longer oscillation periods may not further improve heat-transfer performance. Therefore, at this specific Reynolds number, extending the oscillation period beyond $T^+ = 1600$ does not ensure any additional enhancement in heat transfer performance.

Another explanation behind the greater heat-transfer enhancement at larger periods is that longer oscillation periods extend the duration of lingering Stokes strain phases, thereby extending the time available for streak regeneration. This extended regeneration phase likely strengthens turbulence regeneration, ultimately enhancing heat transfer. Conversely, shorter oscillation periods reduce the time for streak development, resulting in weaker streaks, thus lower heat-transfer enhancement (or even heat-transfer decrease).

Considering the viscous sublayer thickness as δ_ν and the thermal sublayer thickness as δ_α, one can define the ratio of their respective thinning as:

$$R_{\alpha,\nu} = \frac{\delta_{\alpha,\mathrm{ac}}}{\delta_{\alpha,\mathrm{unac}}} / \frac{\delta_{\nu,\mathrm{ac}}}{\delta_{\nu,\mathrm{unac}}}. \tag{4.15}$$

In this case, $R_{\alpha,\nu} = 0.975$, indicating a slightly greater reduction in the thermal sublayer compared to the viscous sublayer. This behavior is also observed in the $T^+ = 500$, $W^+ = 30$ case from Chapter 2, suggesting that a more pronounced thermal sublayer reduction likely contributes to the observed dissimilarity. This result aligns with the observation of Rouhi et al. (2025) where a more pronounced dissimilarity is obtained when increasing Pr, which reduces the thermal sublayer thickness. Further investigation of optimal dissimilar heat transfer cases will enable verification of this hypothesis.

4.7.4 Conclusion

This study employed the PBO-LES framework to systematically identify optimal SWO parameters for enhanced heat transfer performance at $\mathrm{Re}_\tau = 200$ and $\mathrm{Pr} = 1$. Notably, this work constitutes the first exhaustive parametric study of SWO actuation for this specific purpose. The methodology is validated through excellent agreement between LES and DNS Nusselt numbers at optimal operating conditions. The principal research outcomes may be summarized as follows:

- A significantly larger number of generations was required for convergence compared to the drag reduction case. This extended convergence period is attributed to three factors: (i) the broader range of oscillation periods examined, (ii) the existence of an extended near-optimal parameter region, and (iii) increased statistical uncertainty associated with longer actuation periods and significant Re_τ augmentation.

- The optimization procedure identified an optimal parameter range of $T^+_{\mathrm{opti}} \in$

[1 300, 1 600], with maximum efficiency achieved at the upper amplitude limit of $W^+ = 40$. The observed correlation between improved thermal transport performance and the upper bounds of both oscillation period and amplitude suggests a direct dependence on Stokes layer protrusion height. Consequently, while periods exceeding $T^+ > 1\,600$ may further enhance both heat transfer and drag, such improvements might not manifest, as the Stokes layer protrusion height could already have reached a plateau.

- The parameters optimal for heat-transfer augmentation were found to be suboptimal for achieving dissimilar heat transfer, with lower oscillation periods showing greater promise for this specific objective. The optimal dissimilar heat transfer case will therefore correspond to lower heat-transfer enhancement compared to longer periods. This is likely because longer oscillation periods extend the duration of streak regeneration, which enhances heat transfer. Conversely, shorter periods limit streak development, reducing heat-transfer enhancement.

- Linear regression analysis revealed only weak correlation between T^+ and W^+, indicating that, as in the drag reduction case, the optimal period remains relatively insensitive to amplitude variations.

4.8 Improving Heat transfer with minimum drag penality

The preceding analysis demonstrated the potential to identify parameters yielding greater dissimilar heat transfer than the case of Section 2.5. Accordingly, the PBO-LES optimization framework is now employed to determine the oscillation parameters that maximize dissimilar heat transfer at $\mathrm{Re}_\tau = 200$ and $\mathrm{Pr} = 1$. The LES implementation utilizes the numerically validated S_1' configuration and mesh, as documented in Section 4.7.1.

The PBO hyperparameters and neural network architectures remain identical to previous investigations, with modifications limited to the parameter search ranges and the reward function definition. The time-averaged analogy factor (Eq. (2.14)) serves as an effective dissimilarity metric and consequently as the reward function. As established in Section 4.7, periods below $T^+ = 1\,000$ exhibit enhanced potential for maximizing heat-transfer dissimilarity. Therefore, while maintaining the same amplitude range as in Section 4.7, the investigated period range is reduced to $T^+ \in [10, 750]$. This optimization aims to surpass the performance of the Section 2.5 benchmark, with the expectation that maximal amplitude will again prove optimal within the specified range.

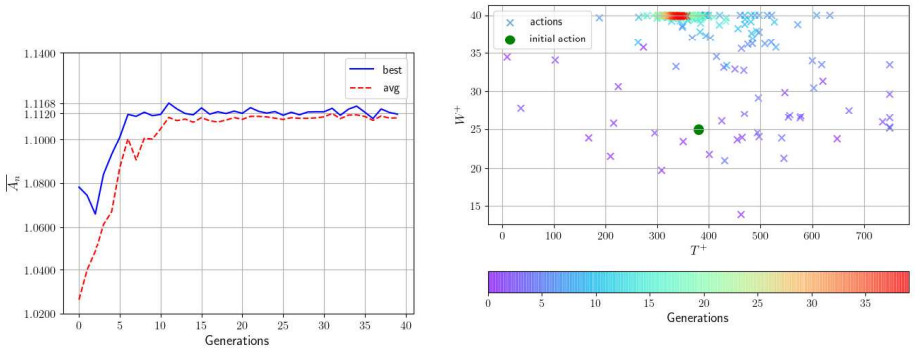

(a) (b)

Figure 4.19: PBO convergence for dissimilar heat transfer, (a): $\overline{A_n}$ evolution through generations. Blue line indicates the best $\overline{A_n}$ individual within the generation, while the red dashed line indicates the average $\overline{A_n}$ individual of the generation. (b): Generation of each individual within the action space T^+-W^+.

4.8.1 Results

The convergence to an optimal parameter region is achieved more than twice as rapidly as in the heat-transfer enhancement case (see Figure 4.19), resulting in significantly reduced computational costs. While this accelerated convergence is partially attributed to improved statistical convergence at shorter periods, it may also indicate a narrower range of optimal periods compared to the heat-transfer scenario. The enhanced statistical reliability is further supported by the close agreement between average and best-performing generations after the 15th generation, which stabilizes at $\overline{A_n} \approx 1.11$. Although this value may be reduced when verified with DNS accuracy (as the S_1' configuration tends to overestimate this quantity, see Tables 4.4 and 4.5), the results suggest that the obtained parameters should outperform the previous best configuration at $T^+ = 500$, $W^+ = 30$.

As shown in Figure 4.19(b), the algorithm initially explores the entire parameter space broadly before focusing more intensively on the near-optimal region. Convergence is observed at $T^+_{\text{conv}} \approx 347$ with the maximum amplitude of 40, as expected. The superior performance observed at maximum amplitude suggests a positive correlation between dissimilar heat transfer and Stokes layer protrusion height, when the oscillation period is fixed. This relationship implies that deeper momentum and thermal transport modulation may be instrumental in achieving greater dissimilarity. The significant deviation from the period $T^+ = 500$ further supports the hypothesis that the new parameters outperform the previous benchmark, which will be verified through DNS validation.

The optimal period range is examined using KDE analysis, focusing on top-performing individuals with amplitudes $W^+ > 39.5$. The observed uncertainty (below 0.01 maximum difference) is markedly lower than in the heat transfer case,

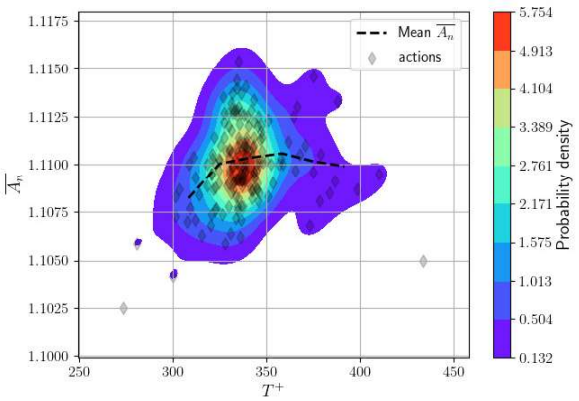

Figure 4.20: Joint probability density function of T^+ and $\overline{A_n}$ around the optimal region, for actions with $W^+ > 39.5$ and $\overline{A_n} > 1.1$. The values of the colormap are obtained from KDE. Black dashed line indicates the mean $\overline{A_n}$ values varying in T^+, while grey diamonds indicate corresponding actions.

consistent with the approximately four times smaller optimal periods obtained here. The KDE mean line indicates an optimal period range of $T^+_{\text{opti}} \in [320, 375]$, substantially narrower than that identified in Section 4.7, which partially explains the faster convergence in this case.

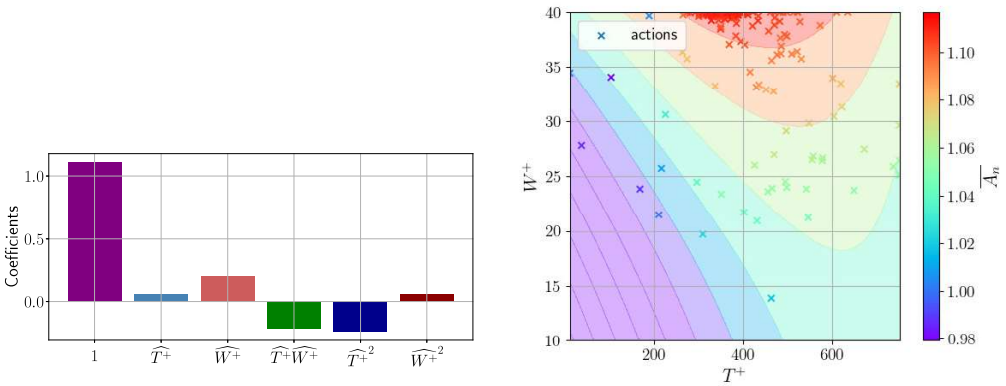

(a) (b)

Figure 4.21: Linear regression results with a 2D quadratic polynomial, using parameters $(\widehat{T^+}, \widehat{W^+})$ as inputs. (a) Magnitude of the obtained coefficients, (b) $\overline{A_n}$ results in the action space. Contours represent the linear regression estimation, while crosses denote the actions explored during PBO.

Parameter sensitivity analysis through linear regression (Figure 4.21) yields important insights. Similar to the heat-transfer enhancement case, the quadratic T^+ term dominates second-order effects while the W^+ term dominates first-order effects, indicating different sensitivity to each parameter in the near-optimal region.

However, the $\widehat{T^+W^+}$ coefficient is significantly larger in this case, revealing non-negligible correlation between period and amplitude. This result is corroborated by Figure 4.21(b), which clearly demonstrates the complex relationship between these parameters with respect to the analogy factor. The linear regression model shows excellent agreement with the sampled data points, confirming the reliability of the resulting mapping. Notably, periods below 200 combined with amplitudes below 25 produce inverse dissimilarity ($\overline{A_n} < 1$). Figure 4.21(b) clearly shows that the optimal period varies with amplitude: lower amplitudes require longer periods for optimal performance. For instance, at $W^+ = 30$, the regression estimates the optimum near the previously analyzed case $T^+ = 500$.

This significant result, namely that the optimal period for dissimilar heat transfer increases as amplitude decreases, will be validated in the following section through an additional optimization process at fixed amplitude ($W^+ = 30$). This validation will determine whether the $T^+ = 500$ case was indeed optimal.

4.8.2 Investigating optimal DHT period at fixed amplitude

This section employs the PBO-LES framework to identify the optimal period for dissimilar heat transfer at a fixed amplitude of $W^+ = 30$. Figure 4.21 suggests that the optimal period for this case should exceed that obtained at $W^+ = 40$.

Since only a single parameter requires optimization, the number of individuals per generation is reduced to $n_{\text{ind}} = 4$. All other hyperparameters in the PBO-LES framework remain unchanged from the previous dissimilar heat transfer investigation.

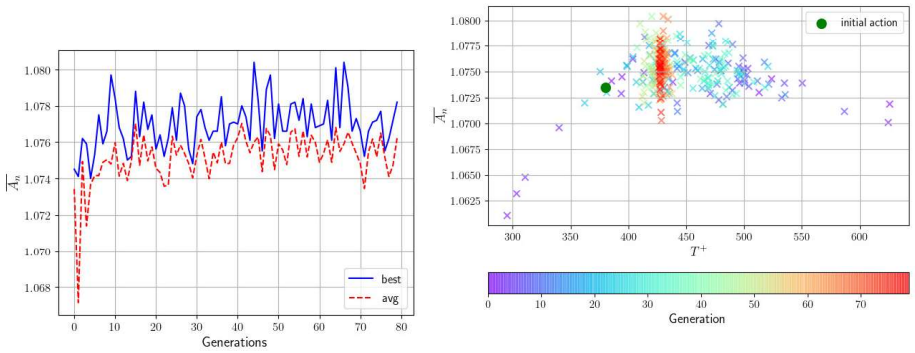

(a) (b)

Figure 4.22: PBO convergence for dissimilar heat transfer at fixed amplitude $W^+ = 30$: (a) Evolution of $\overline{A_n}$ across generations, where the blue line represents the best individual and the red dashed line shows the generation average, (b) Distribution of individuals in the T^+-$\overline{A_n}$ parameter space.

The PBO convergence results, presented in Figure 4.22, demonstrate an aver-

age reward function value of $\overline{A_n} = 1.075$ with a maximum of $\overline{A_n} = 1.08$. While this exceeds the analogy factor achieved at $T^+ = 500$ for the same amplitude in Section 2.5, DNS validation may alter this comparative performance.

The optimization requires twice as many generations as the previous case but compensates with half the individuals per generation, resulting in comparable computational costs. The algorithm converges to a well-defined optimal period of $T^+_{\text{opti}} \approx 430$, which indeed exceeds the optimal period at $W^+ = 40$, confirming that the optimal period increases as amplitude decreases. Figure 4.22(b) reveals that all periods within $[410, 490]$ yield similar $\overline{A_n}$ values, establishing this as the optimal range. This result aligns with the linear regression predictions from Figure 4.21, which estimated an optimum near $T^+ = 500$.

The final validation of these results requires DNS verification. For the $W^+ = 30$ case, only minimal performance gains over the baseline are anticipated when implementing $T^+ = 450$, given the breadth of the identified optimal period range.

4.8.3 DNS validation and comparative analysis

This section aims to validate the optimality of the DHT cases using DNS-level accuracy, ensuring that their performance remains analogous under high-fidelity conditions. Additionally, newly performed DNS simulations enable a comparative analysis of dissimilar heat transfer mechanisms, with results benchmarked against those presented in Chapter 2.

Configuration	(T^+, W^+)	%ΔNu	$\%\Delta C_f$	$\overline{A_n}$
S'_1	(325, 40)	19.94	7.95	1.108
S_2	(325, 40)	16.40	6.51	1.091
S'_1	(450, 30)	10.43	2.89	1.073
S_2	(450, 30)	12.42	5.31	1.068
S_2	(450, 40)	29.06	18.72	1.082

Table 4.6: Comparison in heat transfer and drag SWO induced variations, on S'_1 and S_2.

Table 4.6 presents the drag/heat-transfer enhancement and dissimilarity results for both S'_1 and S_2 configurations, enabling evaluation of both the optimal parameter ranges and the hypothesized period-amplitude relationship for dissimilar heat transfer.

For both optimal cases (at $W^+ = 30$ and $W^+ = 40$), LES and DNS results show good agreement, though the S'_1 configuration consistently overestimates the averaged analogy factor. However, compared to the heat-transfer enhancement case (Table 4.5), this discrepancy remains relatively small, with a maximum $\overline{A_n}$ deviation of 0.017 observed for the $T^+ = 325$, $W^+ = 40$ case.

The optimal SWO configuration at $W^+ = 40$ achieves $\overline{A_n} = 1.091$ in DNS, demonstrating significant improvement in dissimilarity and establishing SWO as an

effective technique for this purpose. For comparison, Uchino et al. (2017) reported $\overline{A_n} = 1.13$ in their DNS study of streamwise traveling wave-like deformations at $\mathrm{Re}_\tau = 180$ and $\mathrm{Pr} = 1$ using CTD boundary conditions, which are known to favour dissimilarity compared to MBC, as discussed in Kasagi et al. (2011). This is achieved with a drag increase of 51%, which is significantly higher than the 6% of the present study.

At $T^+ = 450$ and $W^+ = 30$, the DNS simulation yields $\overline{A_n} = 1.068$, representing a modest improvement over previous results at this amplitude (Table 4.4). This result supports the validity of the optimal period range determined with the PBO-LES methodology. Additional DNS validation at $T^+ = 450$ and $W^+ = 40$ produces an analogy factor 0.09 lower than the optimal $T^+ = 325$ case. This result not only confirms the optimality of the identified period range but also validates the inverse relationship between optimal period and amplitude.

Collectively, the DNS results are consistent with the conclusions drawn from the PBO-LES investigations.

Temporal and Phase-Wise Dynamics of Nu and C_f

The results of the parametric investigations for optimal dissimilar heat transfer have been validated in terms of averaged Nu and C_f variations. There remains to effectively analyse the temporal and phase-wise variations of these quantities, as a comparison with the results of Section 2.5 may yield valuable insights into the physical mechanisms behind dissimilar heat transfer.

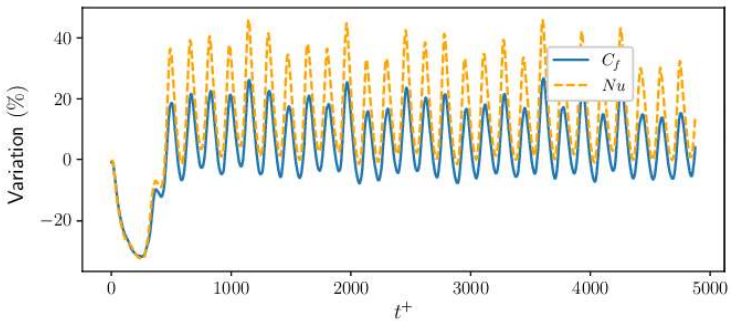

Figure 4.23: Temporal evolution of drag and heat transfer obtained by spanwise oscillatory wall actuation for control parameters $T^+_{\mathrm{nom}} = 325$ and $W^+_{\mathrm{nom}} = 40$, on the S_2 configuration.

Figure 4.23 presents the temporal variations of C_f and Nu for the simulation employing the optimal SWO parameters identified in this study for dissimilar heat transfer ($T^+ = 325$, $W^+ = 40$). The transient and equilibrium states exhibit similar patterns to the $T^+ = 500$, $W^+ = 30$ case (Figure 2.6), though with markedly greater disparity between drag and heat-transfer enhancements. This effect is particularly pronounced during maximum friction phases, which correspond to peak dissimilarity values, mirroring the behavior observed in Figure 2.6. This results in a reduced drag

enhancement compared to the $T^+ = 500$ case, coupled with greater Nusselt number augmentation (see Table 4.8), collectively yielding the improved analogy factor.

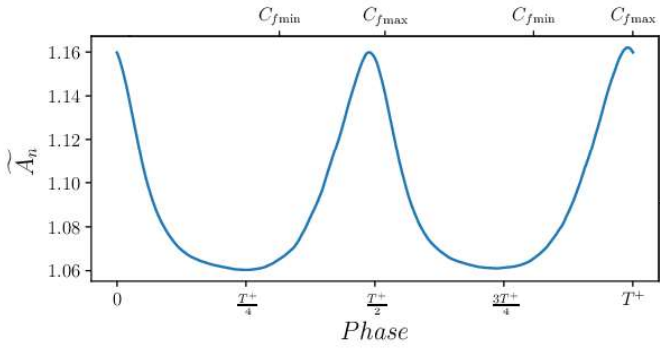

Figure 4.24: Phase-wise averaged evolution of the analogy factor $\widetilde{A_n}$ obtained by spanwise oscillatory wall actuation for control parameters $T^+_{\text{nom}} = 325$ and $W^+_{\text{nom}} = 40$, on the S_2 configuration.

Figure 4.24 displays the phase-averaged analogy factor for this optimal configuration. The analogy factor oscillates between 1.06 and 1.16, showing consistent improvement over the $T^+ = 500$ reference case (Figure 2.7) across all phases, with particularly enhanced values during $\widetilde{A_n}$ maxima. The profile exhibits a more extended plateau during minimum A_n phases, accompanied by sharper transitions between extreme values.

The phase analysis reveals perfect alignment between extremum analogy factor occurrences and quarter-period intervals. However, the correspondence with extremum friction phases appears less pronounced than in the $T^+ = 500$ case (Figure 2.7), particularly during minimum friction phases. This is due to an even more asymmetric distribution of drag variation phases, with ascending phases occurring approximately half as frequently as descending phases. While this asymmetry was previously observed in the $T^+ = 500$ case, its amplified manifestation in the optimal configuration suggests a link between drag variation asymmetry and the achievement of significant dissimilar heat transfer.

Protrusion height, sublayers thinning, and their negligible impact on dissimilarity

In the previous section, it was demonstrated that the maximal heat transfer case aligned with the maximum Stokes layer protrusion height for the selected parameter ranges. Additionally, a preferential thinning of the thermal sublayer relative to the viscous sublayer was observed, reinforcing the hypothesis (consistent with results in Section 2.5) that this mechanism could be a driver of dissimilarity.

To further explore this hypothesis, the protrusion height and sublayers thinning are analyzed in detail for the optimal DHT cases at $W^+ = 30$ and $W^+ = 40$.

The 20% greater protrusion height at $T^+ = 450$, $W^+ = 30$ compared to $T^+ =$

(T^+, W^+)	$R_{\alpha,\nu}$	$\ell_{0.01_{\text{ac}}}^+$
(450, 30)	0.989	65.05
(325, 40)	1.04	54.18

Table 4.7: Ratio between thermal and viscous sublayer thinning $R_{\alpha,\nu}$ and Stokes layer protrusion height $\ell_{0.01_{\text{ac}}}^+$ for the optimal SWO cases for dissimilar heat transfer, on the S_2 configuration.

325, $W^+ = 40$ (Table 4.7) occurs despite the analogy factor gap, confirming that optimal dissimilarity requires more than a simple protrusion height maximization. While the first case ($W^+ = 30$) causes stronger thermal sublayer thinning, an even stronger dissimilarity emerges in the second case ($W^+ = 40$), despite a greater reduction of the viscous sublayer thickness ($R_{\alpha,\nu} > 1$). This directly challenges the hypothesis made in Sections 4.7 and 2.5 that $\overline{A_n} > 1$ is linked to a larger thinning of the thermal sublayer compared to the viscous one.

Influence of Prandtl and Reynolds numbers on actuation performance

Section 2.7 analyzed how Prandtl (Pr) and Reynolds (Re) numbers influence the actuation's performance. It showed that increasing Re reduces the enhancement in both heat transfer (Nu) and drag (C_f) but leaves the dissimilarity (A_n) relatively unchanged. Conversely, increasing Pr enhances the dissimilarity. However, these conclusions were drawn from a single DNS simulation using fixed oscillation parameters ($T^+ = 500$ and $W^+ = 30$). This section aims to confirm if these trends hold true for a different set of parameters ($T^+ = 325$ and $W^+ = 40$).

	Pr			
	0.5	0.71	1	2
$\%\Delta Nu$ on S_2	12.14	14.15	16.40	21.90
$\%\Delta Nu$ on S_3	-0.63	2.16	5.21	12.44
$\overline{A_n}$ on S_2	1.052	1.071	1.091	1.142
$\overline{A_n}$ on S_3	1.032	1.060	1.091	1.166
$\%\Delta C_f$ on S_2	6.51			
$\%\Delta C_f$ on S_3	-3.81			

Table 4.8: ΔNu, ΔC_f and $\overline{A_n}$, for oscillation parameters $T^+ = 325$, $W^+ = 40$, on S_2 ($\text{Re}_\tau = 180$, see Table 2.1) and S_3 ($\text{Re}_\tau = 1020$, see Table 2.4), and for Pr = 0.5, 0.71, 1 and 2.

The variations of C_f, Nu, and A_n across Reynolds and Prandtl numbers are found to follow trends consistent with those shown in Figures 4.23 and 4.24 at $\text{Re}_\tau = 180$ and Pr = 1, which has led to the presentation of averaged values in Table 4.8. The DNS results at $\text{Re}_\tau \in \{180, 1020\}$ and $\text{Pr} \in \{0.5, 0.71, 1, 2\}$, for oscillation parameters $T^+ = 325$ and $W^+ = 40$, confirm the Prandtl number effects

previously observed in Section 2.7, with ΔNu and consequently $\overline{A_n}$ being increased at higher Pr values. The maximum analogy factor of $\overline{A_n} = 1.166$ is achieved at $\mathrm{Pr} = 2$ and $\mathrm{Re}_\tau = 1020$, representing the highest dissimilarity achieved throughout this thesis.

Attenuation of ΔNu and ΔC_f with increasing Reynolds number is observed, with discrepancies near 10%, matching earlier results. For the $\mathrm{Re}_\tau = 1020$ case, these effects manifest as a 3.81% reduction in drag coupled with heat-transfer enhancements ranging from 2.16% at $\mathrm{Pr} = 0.71$ to 12.44% at $\mathrm{Pr} = 2$. It is noted that a marginal reduction in Nu occurs at $\mathrm{Pr} = 0.5$, indicating that greater energy input is required to achieve heat-transfer enhancement at higher Reynolds numbers.

The combined effect of the Prandtl (Pr) and Reynolds (Re) numbers reveals a clear pattern. As the Reynolds number increases, its impact on dissimilarity ($\overline{A_n}$) depends on the Prandtl number:

- For fluids with $Pr > 1$, the dissimilarity is moderately amplified (it increases).
- For fluids with $Pr < 1$, the dissimilarity is slightly attenuated (it decreases).
- For fluids with $Pr = 1$, the dissimilarity does not change.

This identifies the Prandtl number of 1 as a transitional point between these two behaviors.

These results collectively confirm the Prandtl and Reynolds number effects on the SWO control, as established in Section 2.7.

4.8.4 Conclusion

This section employed the PBO-LES framework to identify optimal SWO parameters for maximizing heat transfer increase while minimizing drag penalties, at $\mathrm{Re}_\tau = 200$ and $\mathrm{Pr} = 1$. This investigation represents a critical objective for the present work and the first systematic optimization of SWO actuation for dissimilar heat transfer. Key results include:

- Accelerated PBO convergence relative to heat-transfer enhancement optimization, resulting from shorter and narrower optimal period ranges ($T^+_{\mathrm{opti}} \in [410, 490]$ at $W^+ = 30$; $T^+_{\mathrm{opti}} \in [320, 375]$ at $W^+ = 40$). DNS validation confirmed these optimal ranges, reaching $\overline{A_n} = 1.091$. The optimality of the maximum amplitude suggests a positive correlation (at fixed period) between the depth of momentum/thermal transport modulation and the resulting dissimilarity magnitude. However, the results clearly confirm that a more protrusive Stokes layer does not systematically yield greater dissimilarity. The optimal cases also establish that dissimilarity does not necessarily arise from a stronger thinning of the thermal sublayer relative to the viscous one.

- A unique amplitude-period correlation was observed, contrasting with drag reduction or heat-transfer enhancement scenarios. Interestingly, within the range of parameters investigated, the optimal period was seen to consistently decrease with increasing amplitude, a result also validated with DNS accuracy.
- The optimal configuration demonstrated enhanced dissimilarity compared to the reference case, as well as significant drag variation asymmetry. This phase-dependent behavior amplifies trends observed in the reference case, suggesting this asymmetry contributes to the thermal and momentum transport decoupling mechanism. Furthermore, the dissimilarity gains are most pronounced during high-friction phases, a result that also coincides with the reference case.

The Pr-Re effects on actuation performance established in Section 2.7 were confirmed, though the enhancement of the Pr impact on dissimilarity at higher Re is here precluded for Pr = 1. Recent studies (Rouhi et al., 2025) suggest these effects may hold for higher Pr values, making the Prandtl number an interesting additional parameter for future investigation. However, Rouhi et al. (2025) demonstrates that for plane oscillations, optimal drag reduction parameters remain stable with Pr because the time scale of near-wall energetic temperature scales shows little sensitivity to Prandtl number variations. Additionally, the observed trends suggest the optimal Pr for DHT will likely correspond to the maximum value within the defined action space. Thus, while Pr effects are significant, including Pr in the optimization process may prove unnecessary. However, performing parametric studies at various fixed Pr values remains interesting for investigating the Pr-independence of optimal DHT control strategies.

The streamwise travelling wave control strategy presents an interesting direction, as it introduces the streamwise frequency k_x as an additional parameter. The ω-k_x map presented in Knoop et al. (2025) indeed highlights an area of positive $\omega = \frac{1}{T}$ and k_x that results in drag enhancement and thus would be interesting to investigate with regards to dissimilar heat transfer.

Chapter summary

In this chapter, the optimization of SWO control parameters is performed through single-step reinforcement learning, with the ultimate objective being the optimization of dissimilar heat transfer. The optimization framework, composed of PBO and implicit LES at $\mathrm{Re}_\tau = 200$ and Pr = 1, is thoroughly described, and the accuracy of the simulations is validated against DNS data.

The framework is first applied to the optimization of drag reduction and net energetic gain, where optimal parameters are already established in the literature. The efficacy of the PBO-LES setup is validated, where the parameter space is broadly investigated by the PBO algorithm in initial iterations before being focused on near-optimal regions. Additional iterations are required for convergence due to statistical

uncertainty caused by the short integration times of the simulations. Parallelization is favored by the implicit-LES approach, and computational costs are drastically reduced compared to DNS parametric studies, while maintaining sufficient accuracy. Furthermore, this study provides additional precision regarding optimal parameters for net energetic gain, along with relevant insights into oscillation parameter sensitivities.

The optimization process is then applied to two less-explored SWO objectives: (i) heat-transfer enhancement and (ii) dissimilar heat transfer. The LES configuration employed in this investigation therefore incorporates an additional temperature equation with $\mathrm{Pr} = 1$, solved under mixed boundary conditions. The optimal period for heat-transfer enhancement is found to be significantly higher ($T^+_{\mathrm{opti}} \in [1\,300, 1\,600]$, yielding 52% Nusselt number increase) than in the drag reduction case, while following an analogous relationship to oscillation amplitude. These longer periods are associated with more pronounced statistical uncertainty and a prolonged lingering Stokes strain, which favours turbulence recovery. The optimal parameter ranges demonstrate a strong correlation between Stokes layer protrusion and heat-transfer enhancement. This suggests that (i) periods exceeding $T^+ > 1\,600$ may further improve thermal performance, provided no saturation threshold has been reached, and (ii) optimal periods could lengthen with an increasing Reynolds number.

However, these same parameters prove suboptimal for dissimilar heat transfer applications, indicating the involvement of more complex physical mechanisms than Stokes layer protrusion height maximization. For the dissimilar heat transfer case, the optimized parameters at $W^+ = 40$ achieve superior performance with $\overline{A_n} = 1.091$, while simultaneously producing a reduced Stokes layer protrusion height compared to the optimal $W^+ = 30$ configuration. An inverse relationship between optimal period and amplitude is observed. Through DNS validation and comparison with the reference $T^+ = 500$, $W^+ = 30$ case, pronounced asymmetry in the phase distribution of drag variations is revealed to be likely linked to the decoupling of thermal and momentum transport. However, this comparison also demonstrates that preferential thinning of the thermal sublayer relative to the viscous sublayer is not essential for achieving dissimilar heat transfer.

The previously established influence of Prandtl and Reynolds numbers on dissimilarity was confirmed. While omitting Pr from direct optimization appears justified, performing individual optimizations at specific Prandtl numbers remains an interesting avenue to verify the universality of optimal oscillation parameters for DHT applications. Additionally, the introduction of the streamwise frequency as an additional parameter constitutes a compelling direction for future parametric investigations. The validation of the simulations and optimization techniques provided in this chapter paves the way for waveform optimization of the spanwise wall signal, the subject of the next chapter.

Beyond drag reduction and heat transfer, integrating machine learning algo-

rithms with numerical simulations offers a promising pathway for discovering innovative flow control strategies tailored to diverse real-world challenges. Future studies could explore adapting different reinforcement learning algorithms to open-loop control systems, investigate alternative optimization techniques such as genetic algorithms or particle swarm optimization, and implement reduced-order models to improve computational efficiency. By demonstrating the effectiveness of PBO in a well-established test case with known optimal parameters, this study not only validates existing results but also paves the way for applying SRL methods to more complex flow control problems. Additionally, the coupling of open-source PBO and LES implementations provides a versatile and accessible toolset, facilitating broader adoption of PBO within CFD environments. This establishes a new paradigm for harnessing machine learning in fluid dynamics, underscoring its potential to revolutionize the design of efficient flow control techniques and contribute to the development of sustainable, energy-efficient technologies across various industrial sectors.

Version française

Ce chapitre présente l'optimisation des paramètres de contrôle à l'aide de méthodes d'apprentissage par renforcement simplifié, avec pour objectif final l'optimisation du transfert thermique dissymétrique. Le cadre d'optimisation combine l'algorithme d'optimisation par politique avec des simulations implicites aux grandes échelles à $\mathrm{Re}_\tau = 200$ et $\mathrm{Pr} = 1$. Après une description de ces méthodes, la précision des simulations est validée par comparaison avec des données issues de simulations numériques directes.

Ce dispositif d'optimisation est initialement appliqué à la réduction de traînée et au gain énergétique net, des objectifs pour lesquels les paramètres optimaux sont bien documentés dans la littérature. L'efficacité de l'approche est démontrée : l'algorithme explore d'abord de manière exhaustive l'espace des paramètres lors des premières itérations, avant de se concentrer progressivement sur les régions quasi-optimales. La convergence vers des paramètres précis nécessite des itérations supplémentaires, en raison de l'incertitude statistique introduite par les durées d'intégration courtes des simulations. Les simulations implicites aux grandes échelles permettent une parallélisation efficace et réduisent significativement les coûts computationnels comparé aux études paramétriques utilisant des simulations numériques directes, tout en maintenant une précision satisfaisante. Cette étude permet également de déterminer un intervalle plus précis de paramètres d'oscillations optimaux pour le gain énergétique net et révèle de nouveaux résultats concernant la sensibilité aux paramètres d'oscillations.

Le processus d'optimisation est ensuite étendu à deux objectifs moins documentés des oscillations transverses pariétales : (i) l'augmentation du transfert thermique et (ii) le transfert thermique dissymétrique. Le modèle numérique intègre par conséquent une équation supplémentaire pour la température, à $\mathrm{Pr} = 1$, et résolue avec des conditions aux limites mixtes. La période optimale pour l'augmentation du

transfert thermique s'avère significativement plus élevée que dans le cas de la réduction de traînée ($T^+_{\text{opt}} \in [1\,300, 1\,600]$, entraînant une augmentation de 52% du nombre de Nusselt), tout en suivant une relation analogue avec l'amplitude d'oscillations. Ces périodes plus longues sont associées à une incertitude statistique plus marquée et à des phases prolongées de contrainte de Stokes quasi-constante, favorisant la régénération turbulente. L'analyse des plages de paramètres optimaux révèle que la saillie de la couche de Stokes influence significativement l'augmentation du transfert thermique. Ces résultats suggèrent que (i) des périodes dépassant $T^+ > 1600$ pourraient encore améliorer les performances thermiques, sous réserve qu'aucun seuil de saturation ne soit atteint, et que (ii) les périodes optimales devraient s'allonger avec l'augmentation du nombre de Reynolds.

Cependant, ces mêmes paramètres s'avèrent sous-optimaux pour le transfert thermique dissymétrique, indiquant l'intervention de mécanismes physiques plus complexes qu'une simple maximisation de la saillie de la couche de Stokes. Dans ce cas, la période optimale à $W^+ = 40$ permet d'atteindre une performance d'$\overline{A_n} = 1.091$, tout en produisant une saillie de couche de Stokes réduite par rapport à la configuration optimale à $W^+ = 30$. Une relation inverse entre période et amplitude optimales est observée. Une validation par simulations numériques directes et comparaison avec le cas de référence à $T^+ = 500$, $W^+ = 30$ révèlent qu'une asymétrie prononcée dans la distribution de phase de la traînée est vraisemblablement liée au découplage entre les transports thermique et dynamique. Cependant, cette analyse montre également que l'amincissement préférentiel de la sous-couche thermique par rapport à la sous-couche visqueuse n'est pas essentiel pour obtenir un transfert thermique dissymétrique.

Bien que l'omission de Pr dans le processus d'optimisation semble justifiée, la stabilité des paramètres optimaux face aux variations du nombre de Prandtl reste à confirmer. La réalisation d'optimisations individuelles à des nombres de Prandtl spécifiques constituerait une piste pertinente pour vérifier ce point. De plus, l'introduction de la fréquence longitudinale des oscillations comme paramètre supplémentaire représente une direction pertinente pour de futures études. La validation des méthodes de simulation et d'optimisation employées dans ce chapitre ouvre la voie à l'optimisation de la forme du signal transverse pariétale, qui fera l'objet du prochain chapitre.

Au-delà de la réduction de traînée et du transfert thermique, l'intégration d'algorithmes d'apprentissage automatique avec des simulations numériques ouvre une voie prometteuse pour découvrir des stratégies innovantes de contrôle des écoulements, adaptées à des défis réels variés. Les études futures pourraient explorer l'adaptation de différents algorithmes d'apprentissage par renforcement aux systèmes de contrôle en boucle ouverte, étudier des techniques d'optimisation alternatives comme les algorithmes génétiques ou l'optimisation par essaim particulaire, et mettre en œuvre des modèles d'ordre réduit pour améliorer l'efficacité computationnelle par rapport aux LES.

Chapter 5

Optimizing the waveform of the spanwise wall signal

Contents

5.1 Introduction 160
5.2 Methods 161
5.2.1 Waveform parameterization 161
5.2.2 PBO-LES configuration 164
5.2.3 Validation of waveform parameterization 165
5.3 Optimal signal for drag reduction and net energetic gain . 166
5.3.1 Drag reduction 167
5.3.2 Net energetic gain 171
5.3.3 Conclusion 174
5.4 Optimal signal for dissimilar heat transfer 175
5.4.1 DNS validation and analysis 178
5.4.2 Conclusion 179
Chapter summary 181

In this chapter, the methodology developed in Chapter 4 is utilized, with the extented objective of optimizing the waveform of the wall-signals for both drag reduction and dissimilar heat transfer.

Closed-loop control offers unmatched control input flexibility by exploring a high-dimensional space of possible inputs. In the context of SWO, it achieves this through real-time adjustments of the spanwise wall velocity signal via feedback mechanisms, which deep reinforcement learning algorithms like PPO can optimize. However, this flexibility comes at the cost of the inherent stability found in open-loop control, which results from the steadiness or periodicity of the actuation. Indeed, with SWO, the periodicity and symmetry of sinusoidal signals ensure that the flow converges to an equilibrium state, a requirement for reliable statistical performance evaluation. Closed-loop systems lack these stability constraints, often producing transient performance that complicates consistent evaluation.

To enhance the diversity of the control signal in the investigated parameter space while guaranteeing consistent statistical evaluation, the authors maintain an open-

loop framework while enhancing signal diversity through parametric waveform generation. This approach preserves the computational advantages of open-loop control by avoiding complex state-action relationship learning while enabling exploration of varied signal waveforms via polynomial interpolation. The parameterization explicitly enforces periodicity and symmetry constraints, ensuring consistent statistical evaluation. This strategy will be coupled with the PBO-LES framework established in Chapter 4.

The chapter is structured as follows:

- Section 5.1 reviews relevant literature and contextualizes the waveform optimization objective.
- Section 5.2 details the optimization methodology, including PBO hyperparameters, waveform parameterization, and implicit LES validation.
- Section 5.3 presents waveform optimization results for drag reduction and net energetic gain, building on encouraging prior waveform studies.
- Section 5.4 achieves the final objective of this thesis by identifying optimal waveforms for dissimilar heat transfer.

5.1 Introduction

The fundamental concepts employed in this section (DRL, SRL, and SWO) have been comprehensively reviewed in Chapter 1. Consequently, this introduction focuses specifically on: (i) prior non-parametric investigations of SWO signal waveforms for drag reduction, and (ii) the advantages and limitations of DRL implementations in closed-loop control scenarios.

Cimarelli et al. (2013) investigate the performance of 9 non-sinuoidal spanwise wall signals with regards to drag reduction and net energetic gain, using DNS of a turbulent channel flow at $\mathrm{Re}_\tau = 200$. The authors provide a signal waveform (which is discontinuous) that outperforms the baseline sinusoidal signal with regards to drag reduction. This signal is constant at maximal and minimal amplitudes and results in an additional $\approx 5\%$ reduction in drag by using $T^+ = 125$ and $W^+ = 9$. Regarding net energetic gain, the sinusoidal signal performs the best within the investigated range of parameters: $T^+ \in \{62.5,\ 125,\ 250\}$ and $W^+ \in \{2.25,\ 4.5,\ 9\}$. However, the sinusoid is not the best waveform at every point in the parameter space, suggesting there may exist a non-sinusoidal signal that yields optimal results in a different region of the parameter space. Cimarelli et al. (2013) also provide an expression of the Stokes layer for any waveform, using Fourier series. This expression yields an accurate estimation of the Stokes layer penetration depth, for any wall-signal. According to Rouhi et al. (2025), the Stokes layer protrusion height ($\ell_{0.01}$) serves as a more reliable indicator than the Stokes layer penetration depth, owing to its near-proportionality to the maximum Stokes layer production. The expression of

$\ell_{0.01}$ also has the advantage of not being restricted to the sinusoidal waveform. Cimarelli et al. (2013) demonstrate that waveforms where the first Fourier mode does not dominate lead to substantially increased power consumption. Through consideration of just two dominant modes, they derive a waveform that achieves a net energetic gain improvement of $\Delta P_{\text{net}} = 0.01$ compared to a sinusoidal signal operating at its optimal period but sub-optimal amplitude of $\frac{W^+_{\text{opti}}}{2}$.

Several studies investigated closed-loop control strategies for drag reduction using DRL (Sonoda et al., 2023; Guastoni et al., 2023; Cavallazzi et al., 2024). In a recent study, Jeon et al. (2024) employ a multi-agent system integrated with group-invariant networks to enforce local symmetry invariance, alongside positional encoding to provide spatial information to the agents. The group-invariant networks enhance both learning speed and stability, whereas positional encoding contributes to achieving improved policies. This underscores the importance of carefully designing neural network architectures tailored to the specific constraints and objectives of the problem at hand.

Recently, Cavallazzi et al. (2024) applied closed-loop control to the STW method using the PPO algorithm in a turbulent channel flow ($\text{Re}_\tau = 200$). Their goal was to maximize drag reduction by continuously varying the signal's frequency $\omega(t)$. Although the learning agent found a strategy that temporarily surpassed the optimal fixed-frequency STW, it failed to sustain this performance, with the benefits decaying midway through the simulation. This demonstrates a core trade-off: while closed-loop control offers the flexibility to adapt the signal in real-time, it loses the stability provided by the strict periodicity and symmetry of open-loop control, making it difficult to achieve sustained, long-term drag reduction.

5.2 Methods

5.2.1 Waveform parameterization

In Chapter 4, the context of PBO and its current configuration was outlined. However, an efficient method to parameterize the signal waveform remains to be determined. This method should in particular minimize the number of parameters to simplify the optimization process and reduce the computational cost.

Unlike some of the signal waveforms explored in Cimarelli et al. (2013), this work aims to include only time-continuous wall-signals, which are more practical for real-world implementations. The periodicity of the signal, characterized by a period T, is preserved. Additionally, it is essential that the second half of the actuation cycle is the symmetric counterpart of the first half, as seen in the sinusoidal signal. Specifically, the second half is a time-reversed reflection of the first half, with an opposite amplitude. This inherent periodicity and symmetry enable the system to eventually achieve an equilibrium state (Ricco et al., 2021), which closed-loop strategies may find challenging to attain (Cavallazzi et al., 2024). Given this, the

signal only needs to be parameterized over the time interval $[0, T/2]$, with the signal on $[T/2, T]$ being inferred from the symmetry assumption.

Polynomial interpolation is an effective method for generating signals that are parameterized by control points. By specifying design points through which the signal must pass, a smooth polynomial is constructed to connect them, creating a continuous representation of the signal. Polynomials are advantageous since they are computationally efficient, differentiable, and capable of modeling complex functions with a simple mathematical form. In this study, five control points $\{(t_i^1, y_i^1)\}_{i=1,...5}$ are used to parameterize the signal over its first half period. Their coordinates in time t are given by the following equation :

$$\forall i \in \{1, \ldots, 5\}, \quad t_i^1 = \frac{(i-1)T/2}{4}. \tag{5.1}$$

The corresponding coordinates in space $\{y_i^1\}_{i=1,...,5}$ represent the ordinates (or weights) of the control points, which act as design parameters that enable the effective modulation of the signal. In order to simplify the waveform of the signal and to facilitate the application of the periodicity, symmetry and continuity constraints, y_1 and y_5 are fixed to 0 for any generated signal. The remaining ordinates $\{y_i^1\}_{i=2,...,4}$ are the control points that parameterize the signal. These ordinates are assumed to be positive to prevent the creation of waveforms with large and frequent variations. Considering the imposed symmetry, the control points for the second half period $\{(t_i^2, y_i^2)\}_{i=1,...,5}$ can be defined as:

$$\forall i \in \{1, \ldots, 5\}, \quad t_i^2 = t_i^1 + \frac{T}{2}, \; y_i^2 = -y_{6-i}^1. \tag{5.2}$$

Lagrangian interpolation was chosen as the method for generating the control signal due to its suitability for a small number of evenly spaced control points. The method is computationally efficient, as a polynomial that passes exactly through the design points is directly constructed, thereby avoiding the need to solve complex systems of equations often required by other techniques like splines. Additionally, by using only five control points, the primary drawbacks of high-degree Lagrangian interpolation, such as limited local control or Runge's phenomenon (excessive oscillations between points), are avoided. Furthermore, because the time coordinates ($\{t_i^1\}$) are fixed, the corresponding Lagrangian polynomials form a mathematical basis for the space of all 4th-order polynomials ($\mathbb{R}_4[X]$). It is therefore guaranteed that any 4th-order polynomial can be uniquely and completely described by adjusting the vertical coordinates ($\{y_i^1\}$), which provides the flexibility to generate a wide variety of signal waveforms.

The Lagrangian basis functions are defined separately for the first half ($k = 1$) and the second half ($k = 2$) of the signal's cycle. They are expressed as follows:

$$\forall i \in \{1,\dots,5\},\ \forall k \in \{1,2\}, \quad L_i^k(t) = \prod_{\substack{j=1 \\ j\neq i}}^{5} \frac{(t - t_j^k)}{(t_i^k - t_j^k)}. \tag{5.3}$$

The corresponding order-4 polynomials are computed as:

$$\forall k \in \{1,2\}, \quad P_k(t) = \sum_{i=1}^{5} y_i^k L_i^k(t), \tag{5.4}$$

using the ordinate points for each half of the signal's cycle.

Considering $\mathbb{1}$ as the indicator function, the spanwise velocity at the wall is thus described by the following signal, in which the mod operator is utilized to enforce periodicity:

$$w_\mathrm{w}(t) = \mathbb{1}_{t \bmod T \leq T/2}(t) P_1(t \bmod T) + \mathbb{1}_{t \bmod T > T/2}(t) P_2(t \bmod T). \tag{5.5}$$

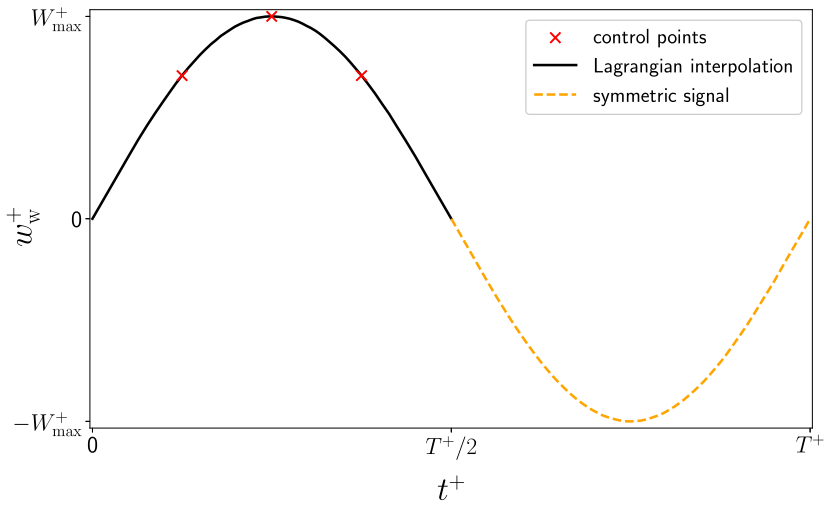

Figure 5.1: Approximation of the sinusoidal signal obtained through Lagrangian interpolation. T^+ indicates the period of the signal and $W^+_{\max}$ its maximal amplitude.

Figure 5.1 illustrates such a signal, approximating a sinusoid, and obtained through Lagrangian interpolation, using the control points (red crosses), with signal values constrained to zero at 0, T^+, and $T^+/2$. The second half period of the signal (dashed orange) is deduced using the reversed negative of the control points.

The signal's maximum amplitude is established through a scaling operation, with W_max serving as a parameter to adjust and regulate the amplitude. Consequently, the scaled spanwise wall signal, $\widetilde{w_\mathrm{w}(t)}$, ultimately applied for actuation, is expressed as:

$$\begin{aligned} &\forall i \in \{1,\dots,5\}, \quad \widetilde{y}_i^1 = y_i^1 \times \mathrm{sc}, \quad \widetilde{y}_i^2 = y_i^2 \times \mathrm{sc}, \quad \text{where } \mathrm{sc} = \frac{W_{\max}}{\max\limits_t(w_{\mathrm{w}}(t))} \\ &\implies \widetilde{P_1}(t) = P_1(t) \times \mathrm{sc}, \quad \widetilde{P_2}(t) = P_2(t) \times \mathrm{sc} \iff \widetilde{w_{\mathrm{w}}}(t) = w_{\mathrm{w}}(t) \times \mathrm{sc}. \end{aligned} \tag{5.6}$$

Due to the scaling approach, a constraint must be applied on the y_i^1 values before scaling, in order to guarantee the obtained polynomials are generated by a unique parameter combination. This uniqueness condition is essential for ensuring the convergence stability of the optimization process. The unactuated case, which corresponds to any parameter combination with $W_{\max} = 0$, does not satisfy the uniqueness property. However, this is considered acceptable, as the unactuated case is known to be suboptimal for all scenarios presently considered. Given this, it is considered sufficient to establish the uniqueness condition while excluding the $W_{\max} = 0$ value.

Therefore, the application $\mathcal{L}$ (see Eq. (5.8)) linking the signal parameters to the scaled polynomial $\widetilde{P_1}$ (or equivalently to $\widetilde{w_{\mathrm{w}}}$) must be injective. Surjectivity of this application, which is related to completeness (i.e., the ability of $\mathcal{L}$ to generate all polynomials in $\mathbb{R}_4[X]$), is not achieved here due to the constraints imposed on the control points. Let $\{a_i\}_{i=1,2}$ be parameters in $[0,1]$, the pre-scaling weights are defined by:

$$\forall i \in \{1,\dots,5\}, \quad y_i^1 = \begin{cases} 0 & \text{if } i \in \{1,5\}, \\ a_1 & \text{if } i = 2, \\ a_2(1-a_1) & \text{if } i = 3, \\ (1-a_1)(1-a_2) & \text{if } i = 4. \end{cases} \tag{5.7}$$

This definition verifies $\sum_{i=1}^{5} y_i^1 = 1$, which is a sufficient condition to guarantee the injectivity of application $\mathcal{L}$ (see proof in Appendix B). Thus, the signal waveforms are fully parameterized (excluding period and amplitude) using only two parameters: $\{a_i\}_{i=1,2}$.

$$\begin{aligned} \mathcal{L} : [0,1] \times [0,1] \times [10,300] \times]0,30] &\to \mathbb{R}_4[X] \\ (a_1, a_2, T^+, W_{\max}^+) &\mapsto \widetilde{P_1}(t) \end{aligned} \tag{5.8}$$

5.2.2 PBO-LES configuration

In the drag reduction case, the amplitude $W_{\max}^+$ is fixed at 15 to investigate whether the sinusoidal signal can be surpassed at this specified amplitude, where relaminarization is not yet achieved. As in the drag reduction case, the optimization of the amplitude is proved unnecessary for dissimilar heat transfer optimization, as established in Section 4.8. Thus, for this specific optimization task, the amplitude is set to $W_{\max}^+ = 30$ to facilitate a direct comparison with the optimal sinusoidal

	a_1	a_2	T^+	$W^+_{\max}$	rewards
Drag reduction	$[0,1]$	$[0,1]$	$[10,300]$	15	P_{sav} (Eq. (4.1))
Net energetic gain	$[0,1]$	$[0,1]$	$[10,300]$	$[0,30]$	P_{net} (Eq. (4.3))
Dissimilar heat transfer	$[0,1]$	$[0,1]$	$[10,750]$	30	$\overline{A_n}$ (Eq. (2.14))

Table 5.1: PBO parameter ranges and rewards for each case study.

case presented in Section 4.8. Consequently, PBO is conducted on three variables: T^+ and $\{a_i\}_{i=1,2}$.

However, as highlighted in Section 4.5, for a determined range of periods, the amplitude plays an important role in optimizing the signal for net energetic gain, as control requires more energy as amplitude increases. In this case, the optimization is thus expanded to four variables, including the additional parameter $W^+_{\max}$. The parameter ranges and rewards for each optimization case are summarized in Table 5.1, where the rewards and the ranges for T^+ and $W^+_{\max}$ align with those used in Chapter 4 for comparison, while $(a_1, a_2) \in [0,1]^2$.

n_{ind}	n_{para}	n_{CPUs}	α	$\lambda_r(\mu,\sigma,\Sigma)$	$n_e\ (\mu,\sigma,\Sigma)$	$n_b\ (\mu,\sigma,\Sigma)$	$n_g\ (\mu,\sigma,\Sigma)$
10	1	256	0.5	$(5,5,1)\times10^{-3}$	(128,16,16)	(1,4,8)	(1,8,16)

Table 5.2: PBO training configuration for waveform optimization. λ_r is the learning rate, n_e is the number of epochs, n_b is the number of mini-batches and n_g is the number of most recent generations included in each training iteration.

Table 5.2 summarizes the PBO hyperparameters for waveform optimization studies. The action space dimensionality increase necessitates raising the number of individuals per generation to 10. While hardware constraints enforce serial execution of the simulations ($n_{\text{para}} = 1$) due to migration to a different computational platform, the fourfold increase in available CPUs (n_{CPUs}) partially compensates this limitation. These adjustments lead to a net reduction in wall-time per generation by above 50% (to 2.8 hours) at the expense of nearly doubled CPU-time consumption (720 hours), reflecting a trade-off between optimization speed and computational resource allocation. All other PBO hyperparameters remain consistent with Chapter 4. The implicit LES configurations also maintain continuity with Chapter 4, employing the S_1 configuration for drag reduction studies, and the $S_1{}'$ configuration at $\text{Pr} = 1$ for dissimilar heat transfer optimization. The modifications with regards to the LES are limited to spanwise velocity boundary conditions as dictated by the control point parameterization.

5.2.3 Validation of waveform parameterization

While the accuracy of the S_1 and $S_1{}'$ simulation configurations for key performance metrics was established in Chapter 4, the Lagrangian interpolation method used for generating the signal waveforms must also be validated. To this end, the method

is tested on a known case wherein a standard sinusoidal signal is approximated (see Figure 5.1). The performance of this approximation is then compared against two benchmarks: a true sinusoidal signal in the same LES setup (S_1) and high-fidelity DNS data from Quadrio and Ricco (2004). For this specific validation, the S_1 configuration is considered sufficient.

Configuration	$\%P_{\text{sav}}$	$\%P_{\text{req}}$	$\%P_{\text{net}}$
Quadrio and Ricco (2004)	39.3	-161.2	-121.9
S_1	41.3	-162.1	-120.8
$S_{1,\text{lag}}$	39.27	-158.1	-118.83

Table 5.3: Comparison in P_{sav}, P_{req} and P_{net}, between S_1, Quadrio and Ricco (2004), and $S_{1,\text{lag}}$, with oscillation parameters $T^+ = 125$, $W^+ = 18$.

As shown in Table 5.3, strong consistency is demonstrated by the results from the Lagrangian-approximated signal ($S_{1,\text{lag}}$) when compared to both benchmarks. Notably, the power savings (P_{sav}) for the $S_{1,\text{lag}}$ case, at 39.27%, is found to align almost perfectly with the DNS reference, for which 39.3% was reported by Quadrio and Ricco (2004). Minor deviations observed in the required power (P_{req}) and net power (P_{net}) are attributed to small statistical variations and the inherent approximation error of fitting a polynomial to a sinusoid. This close agreement serves as confirmation that the Lagrangian parameterization is correctly implemented and functions as a reliable method for generating the control signals. With the parameterization method validated, further performance comparisons against the findings of Cimarelli et al. (2013) and new DNS simulations will be presented in subsequent sections.

5.3 Optimal signal for drag reduction and net energetic gain

This Section investigates the optimal waveform of SWO control signals for drag reduction in a turbulent channel flow at $\text{Re}_\tau = 200$. The optimization framework combines:

- PBO, an efficient single-step reinforcement learning algorithm inspired by evolutionary methods.
- Implicit LES environments, significantly reducing computational costs while maintaining high accuracy for key quantities.
- Wall-signal parameterization using Lagrangian polynomials, enabling diverse signal generation while enforcing periodicity and symmetry constraints.

The optimization is first applied to drag reduction and then extended to include

control energy costs, seeking the spanwise wall signal that maximizes net energetic gain.

5.3.1 Drag reduction

As established in Section 5.2, the signal amplitude $W^+_{\max}$ is fixed at 15 to prevent re-laminarization in the baseline sinusoidal case, enabling meaningful comparison of drag reduction performance. DNS validation will consequently evaluate gains relative to the optimal sinusoidal signal at this amplitude. While Cimarelli et al. (2013) demonstrated superior drag reduction with discontinuous signals, such performance is not guaranteed for continuous waveforms.

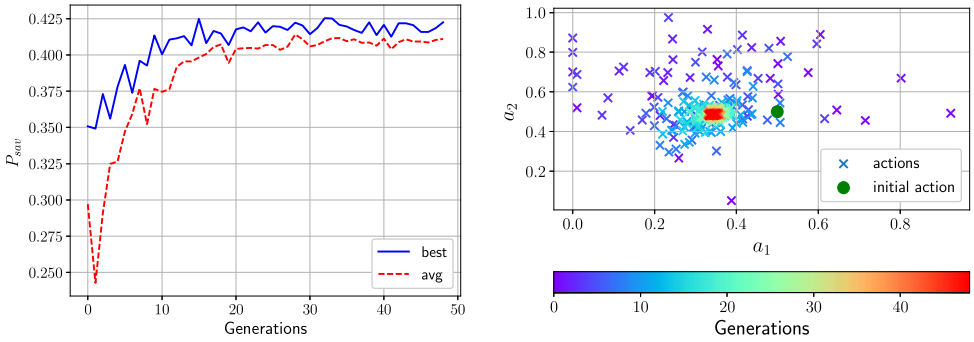

(a) (b)

Figure 5.2: PBO convergence for waveform optimization applied to drag reduction. (a): P_{sav} evolution through generations. Blue line indicates the best P_{sav} individual within the generation, while the red dashed line indicates the average P_{sav} individual of the generation. (b): PBO convergence in the a_1-a_2 space, with colormap corresponding to generations. Green round indicates the mean of the initial action probability distribution.

Figure 5.2(a) demonstrates PBO convergence to a precise signal within 48 generations (480 simulations), requiring 135 wall-time hours and 34 500 CPU hours. Despite the increased parameter space dimensionality, simulation counts remain comparable to the baseline drag reduction study (Section 4.4). The revised computational configuration reduces wall-time while increasing CPU-time but still leverages the advantages of LES for parametric exploration, as established in Chapter 4. The average reward stabilizes near $P_{\text{sav}} = 0.41$ and the best individual reaches $P_{\text{sav}} = 0.425$, which are encouraging results, as will be confirmed by the DNS validation.

Figure 5.2(b) presents the evolution of waveform parameters a_1 and a_2 throughout the optimization process, with the colormap indicating generation progression. The characteristic exploration pattern of the PBO algorithm is clearly visible: initial broad investigation of the parameter space during early generations (blue crosses) followed by progressively more focused exploration of the near-optimal region in later

generations (green to red crosses). The near-optimal parameter region is identified as $a_1 \in [0.25, 0.4]$ and $a_2 \in [0.4, 0.55]$, with the algorithm ultimately converging to $a_{1\mathrm{conv}} \approx 0.34$ and $a_{2\mathrm{conv}} \approx 0.49$. The corresponding optimal signal waveform is illustrated in Figure 5.3.

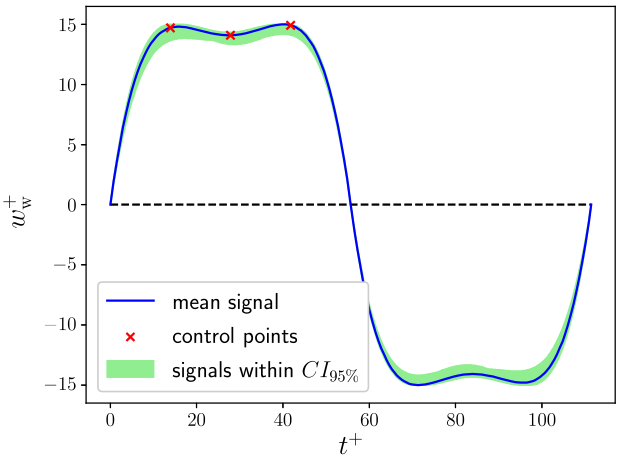

Figure 5.3: Optimal signal for drag reduction. The blue line represents the mean policy signal, while the green area indicates the waveforms within 95% confidence interval. Red crosses indicate control points.

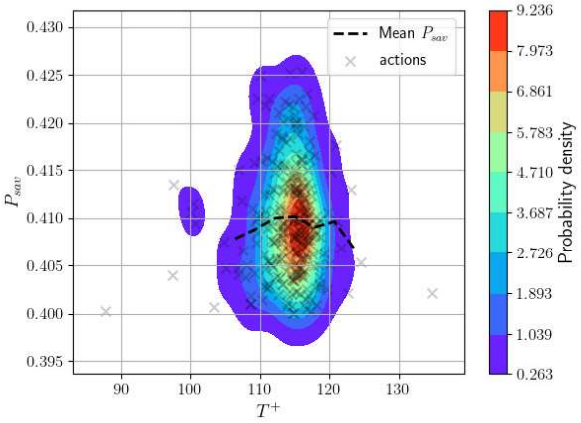

Figure 5.4: Joint probability density function of T^+ and P_{sav} around the optimal region, for actions with $P_{\mathrm{sav}} > 0.4$. The values of the colormap are obtained from KDE. Black dashed line indicates the mean P_{sav} values varying in T^+, while grey crosses indicate corresponding actions.

A remaining question is whether the optimal period remains invariant under changes in waveform. Following the methodology established in Section 4.4, KDE is applied to the periods of top-performing individuals in Figure 5.4. The resulting optimal period range $T^+_{\mathrm{opti}} \in [110, 120]$ closely aligns with the sinusoidal optimization results from Section 4.4, demonstrating remarkable consistency.

The algorithm converges to $T^+_{\mathrm{conv}} \approx 111$ (despite this value not being the most

extensively explored in the parameter space), further confirming that the optimal period remains largely unaffected by significant modifications to the signal waveform. This robustness suggests that the optimal period represents a fundamental characteristic of the flow physics that may be independent to the waveform.

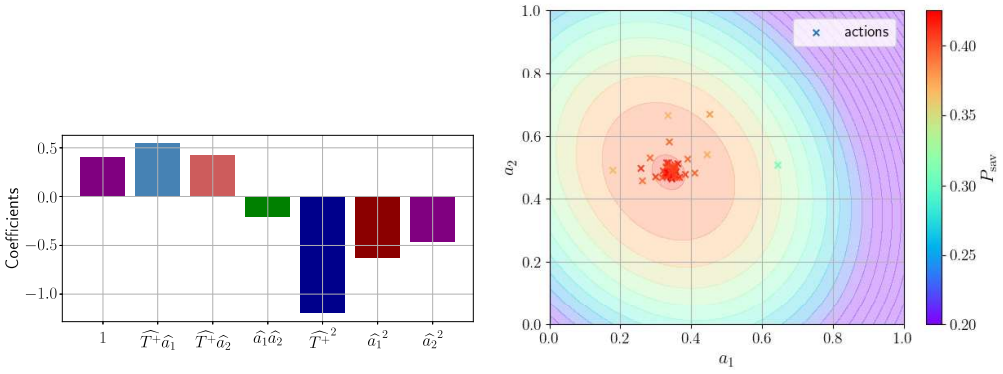

(a) (b)

Figure 5.5: Linear regression results with a 2D quadratic polynomial, using parameters $(\widehat{T^+}, \widehat{a_1}, \widehat{a_2})$ as inputs. (a) Magnitude of the obtained coefficients, (b) P_{sav} results in the a_1-a_2 space for $T^+_{\text{conv}} = 111$. Contours represent the linear regression estimation at $T^+_{\text{conv}} = 111$, while crosses denote the PBO actions where $T^+ \in [109, 113]$.

To verify parameter independence and evaluate sensitivities, the linear regression analysis described in Section 4.4 is here applied to a second-order polynomial of the form:

$$P_{\text{sav}} \approx c_1 + c_2\widehat{T^+}\widehat{a_1} + c_3\widehat{T^+}\widehat{a_2} + c_4\widehat{a_1}\widehat{a_2} + c_5\widehat{T^+}^2 + c_6\widehat{a_1}^2 + c_7\widehat{a_2}^2. \tag{5.9}$$

The coefficient magnitudes corresponding to the best fit on the PBO data are presented in Figure 5.5, along with the linear regression prediction in the a_1-a_2 parameter space for a fixed period of $T^+_{\text{conv}} = 111$. Positive correlations are observed between the period and the waveform parameters, with the strongest correlation being identified for the $\widehat{T^+}\widehat{a_1}$ term. This indicates that while the converged waveform shares the same optimal period as the sinusoidal case, this period is not necessarily optimal for all waveform configurations. This behavior thus differs from the period-amplitude independence observed in Section 4.4.

The waveform parameters $\widehat{a_1}$ and $\widehat{a_2}$ are shown to be negatively correlated, as confirmed by the ellipse orientation in Figure 5.5(b), in which the accuracy of the linear regression model is confirmed. Among the squared terms, the period is found to dominate in magnitude, followed by the a_1 term. These results suggest that P_{sav} exhibits greater sensitivity to variations in T^+ near the optimum, compared to waveform parameters.

The converged signal is presented in Figure 5.3, with the blue curve representing the mean optimized waveform and red markers indicating control points. The green

shaded region represents the 95% confidence interval derived from the waveform parameters distributions:

$$CI_{95\%} = \left\{ a_1, a_2 \in [0,1]^2 \mid \forall i \in \{1,2\},\ a_i \in [\mu_{a_i} - 1.96\sigma_{a_i}, \mu_{a_i} + 1.96\sigma_{a_i}] \right\}. \quad (5.10)$$

The narrow confidence band confirms satisfactory convergence. The optimized waveform exhibits several notable features:

- A distinctive plateau region at maximum amplitude.
- Gradual velocity decrease during the first half of the plateau phase.
- Symmetric velocity increase during the second half.
- Smooth transitions between extremal points.

While this waveform shares qualitative similarities with the discontinuous waveform shown to be effective by Cimarelli et al. (2013), several differences emerge. Most notably, the present waveform maintains continuity, and its plateau region exhibits a moderate modulation rather than a constant velocity. Although this modulation means a perfectly flat plateau is not achieved, the resulting waveform remains qualitatively similar.

Interestingly, the obtained waveform is also very close to the bi-modal one presented in Cimarelli et al. (2013), that yields P_{net} enhancement at sub-optimal low amplitudes. At $W^+ = 15$, however, it is expected that this waveform no longer outperforms the sinusoidal signal in terms of P_{net}.

DNS verification

In this section, the efficacy of the newly determined optimal signal will be evaluated, by performing DNS of this case and the sinusoidal baseline, at $T^+ = 111$, and $W^+ = 15$. It is expected that the Lagrangian waveform will outperform the sinusoidal in terms of drag reduction. However this may come at the expense of net energetic gain, where the newly found signal may require a larger control cost at this amplitude.

Waveform	$\%P_{\text{sav}}$	$\%P_{\text{req}}$	$\%P_{\text{net}}$
Optimal	40.64	184.92	-144.28
Sinusoidal	38.12	118.39	-80.27

Table 5.4: DNS comparison between sinusoidal and optimal waveform for drag reduction at $T^+ = 111$, $W^+ = 15$, on S_2.

The results of those DNS are presented in Table 5.4, in terms of P_{sav}, P_{req} and P_{net}. The DNS value for P_{sav} is in close agreement with the late generations averaged

rewards of Figure 5.2(a), further confirming the accuracy of the LES implementation. Furthermore, the Lagrangian signal obtains a gain of 2.5% in P_{sav} confirming its superiority in terms of drag reduction. This corresponds to a protrusion height 4.1% deeper than the baseline, meaning the agent found a waveform that results in deeper momentum transport modulation. However, this gain in performance comes at a significant cost in energy, with a 66.53% increase in P_{req}, resulting in the signal being sub-optimal with regards to P_{net}. This result is consistent with the findings of Cimarelli et al. (2013), in which a similar waveform demonstrates superior net energetic performance (P_{net}) compared to the sinusoidal baseline exclusively at sub-optimal oscillation amplitudes $W^+ < W^+_{\text{opti}}$. A second investigation is thus warranted to find out whether the optimal sinusoidal signal can be outperformed in terms of net energetic gain.

5.3.2 Net energetic gain

The preceding section established the optimal signal waveform for drag reduction without accounting for the associated control cost, rendering the approach somewhat idealized. In this section, the energy expenditure is incorporated into the analysis via the P_{req} term, with P_{net} serving as the performance metric. Furthermore, since Section 4.5 identified the signal's amplitude as a critical parameter influencing net energetic gain, the maximum amplitude W^+_{max} is introduced as a fourth optimization parameter (refer to Table 5.1). The remaining parameter ranges and PBO configuration remain consistent with the previous investigation.

While Cimarelli et al. (2013) demonstrated that sinusoidal signals achieve the best P_{net} performance among the tested waveforms, their results also indicated that non-sinusoidal waveforms can outperform sinusoidal signals in sub-optimal regions of the T^+-W^+ parameter space. This observation suggests the potential existence of alternative waveforms that might achieve peak net energetic gain in different regions of the parameter space. However, this hypothesis remains unverified, and the subsequent analysis will systematically evaluate whether performance improvements beyond the baseline can be realized.

The optimization process required 72 generations to achieve parameter convergence, corresponding to 720 simulations, approximately 200 hours of wall time, and nearly 50 000 CPU-hours. This represents a substantially greater computational investment compared to the drag reduction case, attributable in part to the optimization of an additional parameter. Figure 5.6 presents the evolution of both average and best rewards across generations. The algorithm exhibits stagnation beyond the 20th iteration, indicative of an extensive near-optimal region within the parameter space. Notably, the best and average individuals achieved $\%P_{\text{net}}$ values of 9.3% and 8% respectively, closely matching the performance metrics obtained during the sinusoidal waveform investigation (see Figure 4.11(a)). This observation suggests two possible interpretations: (i) the PBO converged to a signal closely resembling the sinusoidal waveform, thereby explaining the comparable performance,

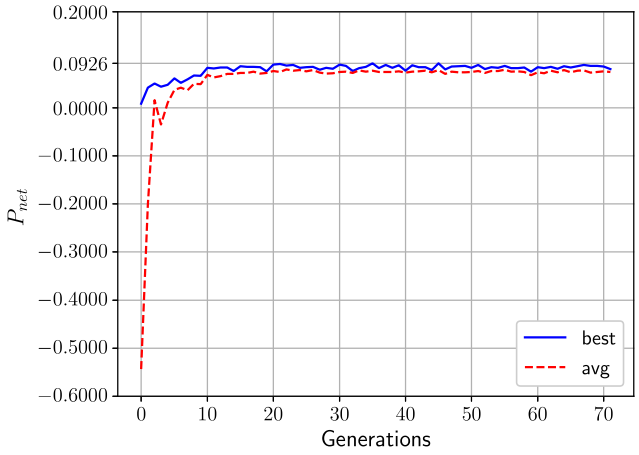

Figure 5.6: P_{net} evolution through generations. Blue line indicates the best P_{net} individual within the generation, while the red dashed line indicates the average P_{net} individual of the generation.

or (ii) the algorithm identified an alternative waveform yielding similar results, potentially located in a distinct region of the T^+-W^+ parameter space.

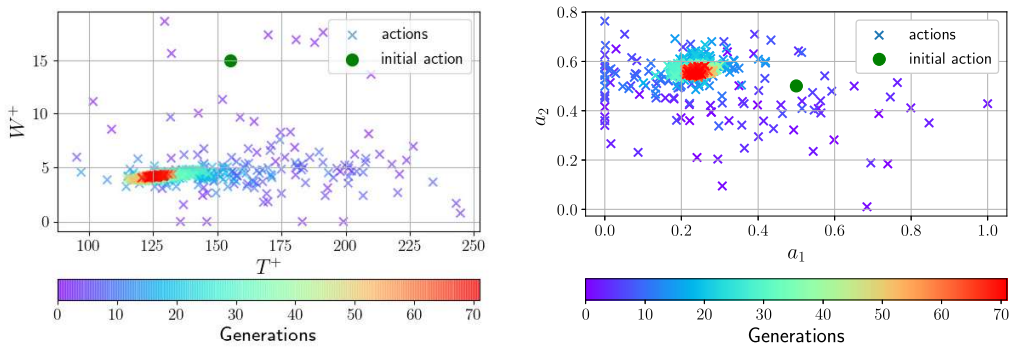

(a) (b)

Figure 5.7: PBO convergence in (a) the T^+-W^+, and (b) the a_1-a_2 space, with colormap corresponding to generations. Green round indicates the mean of the initial action probability distribution.

Figure 5.7 illustrates the distribution of individuals within the parameter space, with the left panel displaying the T^+-W^+ distribution and the right panel presenting the waveform parameter distribution (a_1-a_2). Concerning period and amplitude parameters, the optimization process preferentially explored lower amplitude regions, as higher amplitudes were found to substantially increase control costs. The PBO algorithm did not converge to a singular optimal point but rather identified an extended optimal region. Remarkably, the algorithm converged to T^+ and W^+ ranges analogous to those identified in the sinusoidal investigation, strongly suggesting that the optimal waveform closely approximates a sinusoidal signal. Should

an alternative waveform have been identified, its performance characteristics would necessarily be comparable to the baseline, precluding any significant performance enhancement.

Figure 5.7(b) reveals that the algorithm converged to a distinct region of the parameter space compared to the drag reduction case, yielding a different optimal waveform configuration. The increased parameter uncertainty, as evidenced by the dispersion of individuals in later generations, suggests that the extended convergence time stems not only from the additional parameter but also from a broader spectrum of potentially optimal wall signals.

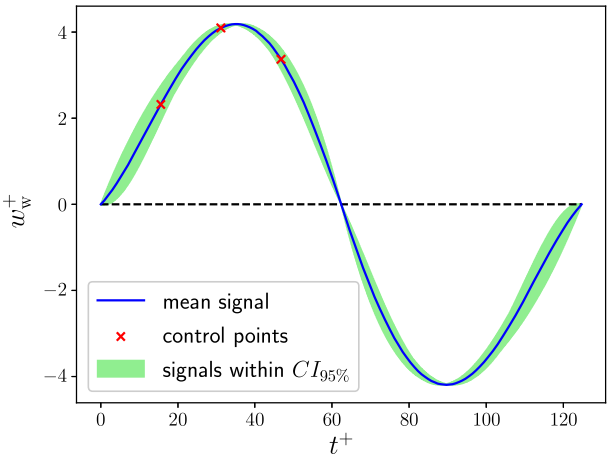

Figure 5.8: Optimal signal for net energetic gain. The blue line represents the mean policy signal, while the green area indicates the signals within 95% confidence interval. Red crosses indicate control points.

The optimal waveform corresponding to these parameters is depicted in Figure 5.8, with the 95% confidence interval indicated by green shading. This interval demonstrates greater waveform variability compared to the drag reduction case, consistent with the observations from Figure 5.7(b). All waveforms within the confidence interval exhibit near-sinusoidal characteristics. The mean optimal signal displays a slight temporal asymmetry, featuring a marginally prolonged ascent phase relative to the descent phase. However, this asymmetry remains subtle, and while the Lagrangian approximation of the sinusoidal signal (Figure 5.1) falls outside the confidence interval, it remains in close proximity to the identified optimal waveform. Consequently, PBO converged to a solution nearly identical to the optimal sinusoidal configuration, without achieving measurable performance improvement.

Synthesizing these findings leads to the conclusion that sinusoidal signals (and closely related waveforms) represent optimal configurations for net energetic gain, at least within the investigated parameter ranges. More sophisticated parameterizations might yield marginally enhanced performance, for example using Splines and additional control points. This study therefore refutes the possibility arising from Cimarelli et al. (2013) that non-sinusoidal waveforms could outperform the optimal

sinusoidal signal in alternative regions of the parameter space.

5.3.3 Conclusion

This study implemented the PBO-LES framework established in Chapter 4, augmented with a polynomial interpolation-based parameterization of the spanwise wall velocity signal. The optimization procedure focused on both drag reduction and net energetic gain objectives at $\mathrm{Re}_\tau = 200$. The investigation builds upon the foundational work of Cimarelli et al. (2013) while evaluating their hypotheses regarding the existence of potentially superior waveform configurations, compared to the optimal sinusoidal cases identified in Sections 4.4, 4.5.

The principal findings of this investigation can be summarized as follows:

- The combined PBO-LES methodology with Lagrangian polynomial interpolation has demonstrated effectiveness in identifying optimal signal waveforms for both drag reduction and net energetic gain. The implicit LES approach provided substantial computational efficiency gains while maintaining consistency with DNS validation results for optimal signal configurations, confirming its reliability for this class of optimization problems. The framework maintains extensibility to more sophisticated parameterization approaches, including Spline or Bézier curve implementations (as in Viquerat et al. (2020)) with additional control points, should greater waveform complexity be desired.

- Optimal period for drag reduction remains unchanged from baseline values. However, its stability across all waveform configurations is not necessarily preserved, as sensitivity analysis demonstrates significant correlations between T^+ and the waveform parameters. Furthermore, this analysis reveals higher sensitivity to period compared to waveform variations in the near-optimal region.

- The obtained drag reduction waveform exhibits a characteristic non-constant plateau phase at extremum amplitudes, consistent with the findings of Cimarelli et al. (2013). While achieving a 2.5% enhancement in drag reduction, which corresponds to a 4.1% increase in Stokes layer protrusion height, this configuration requires 66.53% additional energy expenditure, rendering it suboptimal for net energetic gain considerations.

- For net energetic gain optimization, the parameter space features substantially broader near-optimal regions, resulting in extended convergence durations. The algorithm consistently identified wall-signals closely resembling the optimal sinusoidal configuration, with comparable performance metrics. This finding suggests the baseline sinusoidal waveform resides within the optimal performance region, thereby challenging hypotheses derived from Cimarelli et al. (2013).

Building upon these promising results in the well-established domain of drag reduction, the subsequent investigation will focus on determining optimal waveforms for dissimilar heat transfer, a research direction of paramount importance to this thesis.

5.4 Optimal signal for dissimilar heat transfer

This section presents the application of the combined PBO-LES framework with Lagrangian interpolation waveform parameterization to the optimization of dissimilar heat transfer, with the dual objectives of maximizing thermal enhancement while minimizing drag penalty. A primary focus involves achieving more pronounced decoupling between momentum and thermal transport compared to the sinusoidal waveform optimal results at equivalent amplitude ($W^+ = 30$), as previously investigated in Section 4.8. The parameter ranges and reward function specifications for this dissimilar heat transfer optimization are detailed in Table 5.1.

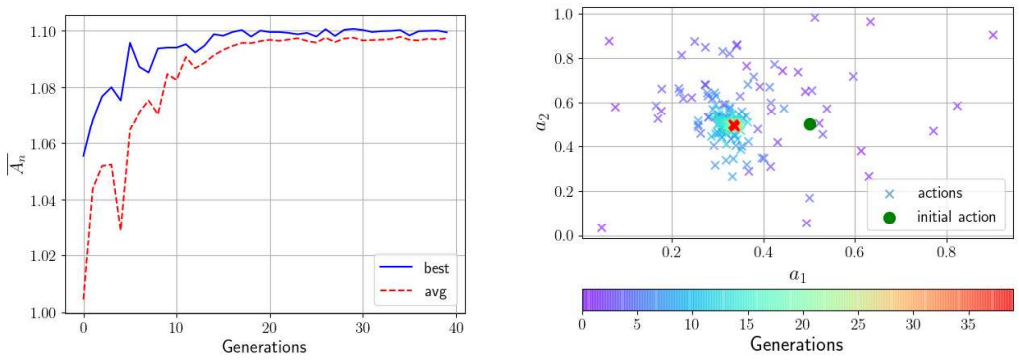

(a) (b)

Figure 5.9: PBO convergence for waveform optimization applied to dissimilar heat transfer. (a): $\overline{A_n}$ evolution through generations. Blue line indicates the best $\overline{A_n}$ individual within the generation, while the red dashed line indicates the average $\overline{A_n}$ individual of the generation. (b): PBO convergence in the a_1-a_2 space, with colormap corresponding to generations. Green round indicates the mean of the initial action probability distribution.

Figure 5.9 illustrates the convergence characteristics of the PBO algorithm for this application. The optimization process required 40 generations (400 simulations), representing only marginally greater computational expense than the sinusoidal optimization in Section 4.8, despite the additional parameter. The optimal individual achieved $\overline{A_n} = 1.10$, corresponding to a significant performance improvement of 0.02 over the sinusoidal case. This enhancement strongly suggests the identification of a non-sinusoidal optimal waveform, as will be subsequently demonstrated in Figure 5.12.

The optimization exhibits remarkably low uncertainty regarding the optimal waveform configuration, evidenced by both the rapid convergence and the tight

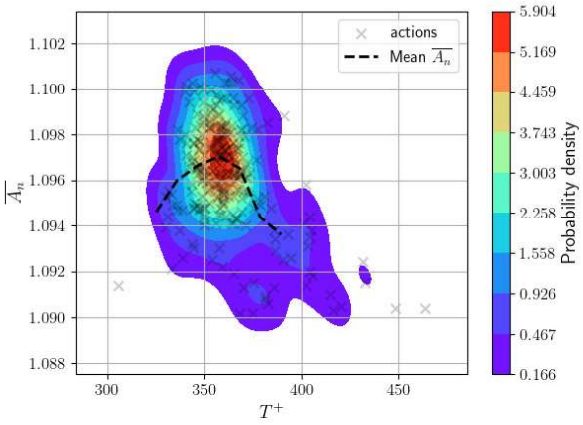

Figure 5.10: Joint probability density function of T^+ and $\overline{A_n}$ around the optimal region, for actions with $\overline{A_n} > 1.09$. The values of the colormap are obtained from KDE. Black dashed line indicates the mean $\overline{A_n}$ values varying in T^+, while grey crosses indicate corresponding actions.

representation of the latest generations in the a_1-a_2 action space (Figure 5.9(b)). This level of precision represents a notable improvement over all previous PBO investigations, implying reduced statistical uncertainty in the waveform's influence on A_n compared to its effects on C_f. This phenomenon may arise from the partial cancellation of statistical uncertainties in Nu and C_f through their ratio in the A_n formulation. The optimization process successfully balanced thorough exploration of the parameter space with effective exploitation of acquired knowledge in later generations.

As established in Section 4.8, the relationship between period and amplitude for dissimilar heat transfer proves to be more complex than for drag reduction, with optimal periods varying inversely with amplitude. This raises a critical question regarding the robustness of the optimal sinusoidal period range ($T^+_{\text{opti}} \in [410, 490]$) to waveform variations.

Figure 5.10 addresses this question through kernel density estimation performed on top-performing individuals ($\overline{A_n} > 1.09$) in the T^+-$\overline{A_n}$ space. Both the density distribution and mean line reveal an exceptionally narrow optimal period range, contrasting sharply with all previous case studies. This precision aligns with the observed convergence to well-defined waveform parameters, resulting in minimal statistical uncertainty across all optimized variables. From the mean line, an optimal period range of $T^+_{\text{opti}} \in [340, 365]$ can be deduced, which notably does not overlap with the optimal range for sinusoidal signals at $W^+ = 30$. This is a fundamental distinction from the drag reduction scenario and demonstrates that the optimal period depends not only on amplitude but also on waveform characteristics.

This observation is confirmed in Figure 5.11 through linear regression analysis (following the same methodology as in Eq. (5.9), but with $\overline{A_n}$ as the estimated quantity instead of P_{sav}). While $\widehat{T^+}$ is found to exhibit negligible correlation with

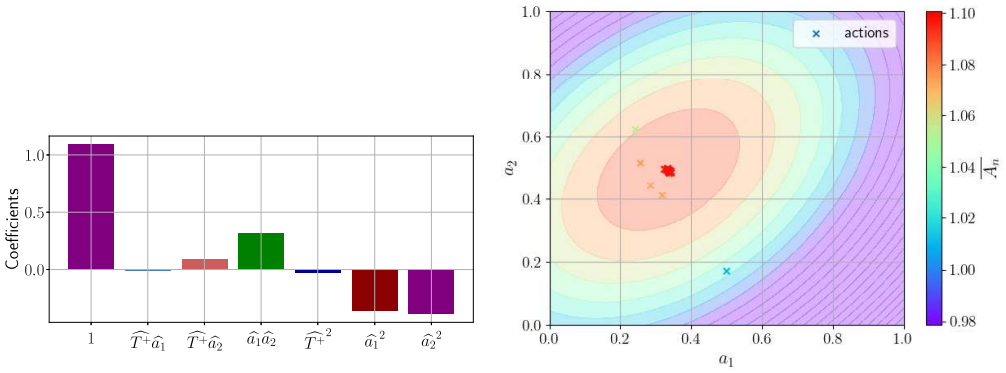

(a) (b)

Figure 5.11: Linear regression results with a 2D quadratic polynomial, using parameters $(\widehat{T^+}, \widehat{a_1}, \widehat{a_2})$ as inputs. (a) Magnitude of the obtained coefficients, (b) $\overline{A_n}$ results in the a_1-a_2 space for $T^+_{\text{conv}} = 350$. Contours represent the linear regression estimation at $T^+_{\text{conv}} = 350$, while crosses denote the PBO actions where $T^+ \in [345, 355]$.

$\widehat{a_1}$, the presence of a positive $\widehat{T^+}\widehat{a_2}$ correlation term provides confirmation that the optimal period is not invariant to waveform variations.

In contrast to the drag reduction case, the correlation between $\widehat{a_1}$ and $\widehat{a_2}$ is positive, suggesting fundamental differences in how waveform parameters influence
5265 dissimilar heat transfer compared to drag reduction. Further distinction is noted when observing the squared terms, as the analogy factor is shown to be more sensitive to waveform variations than to period changes.

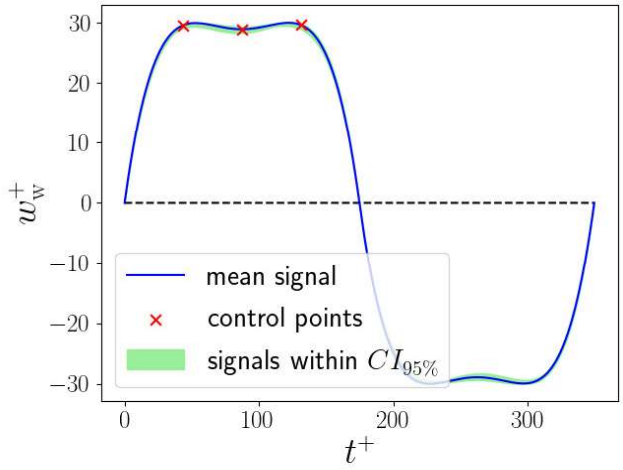

Figure 5.12: Optimal signal for dissimilar heat transfer. The blue line represents the mean policy signal, while the green area indicates the signals within 95% confidence interval. Red crosses indicate control points.

This finding is consistent with the observed convergence to a highly precise wave-

form configuration, as demonstrated in Figure 5.12, where the mean converged waveform and its associated 95% confidence intervals are presented. Indeed, the optimal waveform exhibits extremely low uncertainty, evidenced by the narrow confidence intervals. The identified waveform diverges significantly from sinusoidal characteristics, instead bearing strong resemblance to the optimal drag reduction waveform (Figure 5.3) and the bi-modal waveform of Cimarelli et al. (2013), particularly in its sustained extremum-amplitude plateau phases. The current waveform however displays slightly smoother transitions between control points compared to Figure 5.3. This waveform produces a more protrusive Stokes layer than the sinusoidal case (at identical amplitude and period), as demonstrated in Section 5.3. However, convergence to a lower period compared to the baseline in the present case may imply a reduced Stokes layer protrusion height, despite the use of a more protrusive waveform.

5.4.1 DNS validation and analysis

The final evaluation requires the obtained dissimilarity gains to be assessed with DNS-level accuracy. A relevant comparison can then be made with the DNS performance of the optimal sinusoidal signal at both $W^+ = 30$ and $W^+ = 40$, as documented in Table 4.6.

Configuration	ΔNu	ΔC_f	$\overline{A_n}$
$S_1{}'$	19.78	9.47	1.093
S_2	22.31	12.37	1.087

Table 5.5: Comparison in heat transfer and drag SWO induced variations, using the Lagrangian signal of Figure 5.12 with oscillation parameters $T^+ = 350$, $W^+ = 30$, on $S_1{}'$ and S_2.

The LES results for the $S_1{}'$ configuration and DNS results for S_2 are presented in Table 5.5 . As typically observed, the enhancement of both Nu and C_f is slightly underestimated by the LES. However, the average analogy factor shows only a 0.06 discrepancy with the DNS results, demonstrating strong agreement for this quantity. This close correspondence validates the reliability of the parametric investigation results.

The optimal waveform is found to provide a 0.019 improvement in analogy factor compared to the sinusoidal case at the same amplitude, with this result being well corroborated between LES and DNS. The performance in dissimilarity approaches that of the optimal case at $W^+ = 40$, indicating that the optimized waveform achieves similar performance to this amplitude increase when each case is considered at its respective optimal period. Despite the optimal waveform producing a more protrusive Stokes layer than the sinusoidal signal in the drag reduction case (Section 5.3.1), its reduced optimal period yields a lower $\ell_{0.01_{\mathrm{ac}}}^{+} = 62.45$ compared to the sinusoidal $T^+ = 450$, $W^+ = 30$ reference (see Table 4.7 for comparison). This result further demonstrates that dissimilar heat transfer remains uncorrelated

with Stokes layer protrusion when accounting for period variations. The optimal waveform case shows dominant viscous sublayer thinning over thermal sublayer reduction ($\%R_{\alpha,\nu} = 1.043$ (Eq. (4.15))), further excluding preferential thinning of the thermal sublayer as a primary driver of dissimilarity mechanisms.

Given the comparable performances between the two best dissimilar heat transfer cases, their relative energy costs must be examined. Of particular interest would be the determination of whether an appropriately optimized waveform could achieve dissimilarity enhancement equivalent to an amplitude increase, but with reduced control energy requirements. The optimal waveform at $W^+ = 30$ is characterized by $P_{\text{req}} = 462$, $P_{\text{net}} = -474.65$, $\ell_{0.01_{\text{ac}}}^{+} = 62.45$ and $\text{Nu}_r = 22.31$, whereas the optimal sinusoidal case at $W^+ = 40$ yields $P_{\text{req}} = 499.14$, $P_{\text{net}} = -505.43$, , $\ell_{0.01_{\text{ac}}}^{+} = 54.18$ and $\text{Nu}_r = 16.16$. Consequently, it is demonstrated that the optimal waveform at lower amplitude results in reduced control costs and improved net energetic gain, while simultaneously achieving a 15% larger Stokes layer protrusion height and 6.15% additional gains in heat transfer. This indicates superior performance across all considered metrics when control costs are taken into account, outperforming the sinusoidal signal at higher amplitude. Therefore, not only is enhanced dissimilarity achieved by the optimal waveform compared to the optimal sinusoidal case at equivalent amplitude, but superior performance is also obtained relative to amplitude increase scenarios when control costs are considered.

The bi-modal waveform partly resembles the squared waveform utilized in Knoop et al. (2025) with steady spanwise forcing. Their waveform is associated with a more prolonged lingering Stokes strain compared to the sinusoidal case, which Knoop et al. (2025) link to turbulence recovery in post-optimal (for drag reduction) conditions. This favoured turbulence regeneration may thus explain the enhanced drag and heat transfer compared to the baseline case (see Table 4.6). This comparison between temporal and spatial forcing is further supported by Viotti et al. (2009), who demonstrated fundamental similarities between these actuation approaches.

The phase-wise analysis in Figure 5.13 supports the idea of a fundamentally different physical interaction. In contrast to the sinusoidal case, the drag variation is nearly symmetric over the cycle, and the $\widetilde{A_n}$ profile exhibits distinct plateaus around $C_{f\text{min}}$ phases. These observations suggest that the PBO algorithm has identified a waveform that leverages a different, and more efficient, physical mechanism to achieve the decoupling of thermal and momentum transport. It is thus shown that the PBO approach is capable of identifying waveforms that produce markedly different physical interactions, leading to significant improvements in both dissimilarity performance and control cost efficiency.

5.4.2 Conclusion

Following drag reduction optimization, the optimal waveform and period of the spanwise wall velocity for maximizing dissimilar heat transfer were examined at $\text{Re}_\tau = 200$ and $\text{Pr} = 1$. The optimization employs the PBO-LES framework com-

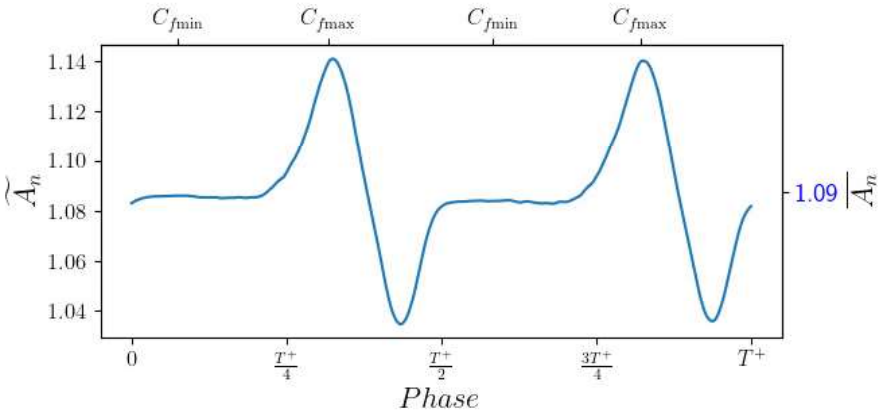

Figure 5.13: Phase-wise averaged evolution of the analogy factor $\widetilde{A_n}$ obtained by spanwise oscillatory wall actuation, using the waveform of Figure 5.12 with control parameters $T^+_{\text{nom}} = 350$ and $W^+_{\text{nom}} = 30$, on the S_2 configuration.

bined with Lagrangian interpolation for signal parameterization (see Section 5.2).

The principal findings are:

- The algorithm demonstrates rapid convergence, reaching precise parameters within just 40 generations. This accelerated convergence suggests potential cancellation of statistical uncertainties between Nu and C_f in the A_n term.

- The optimized waveform at $W^+ = 30$ closely resembles the drag-reduction bi-modal optimal waveform, improving the analogy factor by 0.019. DNS validation also demonstrates the superiority of this waveform compared to the optimal sinusoidal case at $W^+ = 40$ after accounting for control costs, highlighting the advantage of waveform modification over amplitude increase. The prolonged lingering Stokes strain associated to similar waveforms likely enhances heat transfer and drag beyond baseline levels by promoting a stronger turbulence regeneration.

- The identified optimal period differs from baseline values, indicating period dependence on both amplitude and waveform characteristics. Parameter sensitivity analysis reveals greater sensitivity to waveform compared to period variations, contrasting with drag reduction behavior. While converged signal waveforms are similar in both cases, their relationship to waveform parameters differ significantly.

- Phase-wise analogy factor variations vary considerably from baseline patterns, indicating that PBO identifies waveforms achieving enhanced dissimilarity through modified physical mechanisms.

Chapter summary

This chapter investigated optimal SWO waveforms for both drag reduction and dissimilar heat transfer applications at $\mathrm{Re}_\tau = 200$ and $\mathrm{Pr} = 1$, extending the PBO-LES framework established in Chapter 4 through polynomial signal parameterization. The analysis is conducted within an open-loop control scenario, where achieving a statistically steady-state is ensured through imposed periodicity and symmetry constraints. Following the presentation and validation of the Lagrangian interpolation-based signal parameterization, optimization is first performed for drag reduction and net energetic gain, building upon previous work by Cimarelli et al. (2013), before being applied to dissimilar heat transfer scenarios.

For drag reduction, a novel optimal waveform is identified, characterized by a non-constant plateau phase at extremum amplitudes. This configuration yields a 2.5% enhancement in drag reduction accompanied by a 4.1% increase in Stokes layer protrusion height. However, the associated increase in control costs renders this waveform suboptimal with respect to net energetic gain. While the optimal period matches the baseline value, parameter sensitivity analysis reveals that this period stability does not necessarily extend across all waveform variations. Greater sensitivity to period changes compared to waveform modifications in near-optimal regions is also observed. For net energetic gain, the identified optimal signals closely resemble the sinusoidal baseline case, suggesting no superior waveforms exist within the investigated parameter space.

In the dissimilar heat transfer case, the algorithm exhibits rapid convergence to precise parameters, potentially indicating cancellation of statistical uncertainties between Nu and C_f measurements. The resulting optimal waveform shows strong similarity to the drag-reducing configuration, demonstrating improved analogy factor performance. The prolonged lingering Stokes strain associated to plateau waveforms likely results in the enhanced heat transfer and drag performance through favoured turbulence regeneration. Analysis of the A_n and C_f phase distributions reveals different physical mechanisms governing dissimilar heat transfer compared to the sinusoidal case. Comparative evaluation with the optimal sinusoidal waveform at $W^+ = 40$ demonstrates the advantages of waveform optimization over amplitude increases when control costs are considered. The optimal period is shown to vary with waveform parameters, as evidenced by its deviation from baseline values. The relationship between waveform parameters and performance differs from the drag reduction scenario, particularly in the greater sensitivity to a_1 and a_2 compared to T^+.

The proposed methodology has demonstrated effectiveness in identifying optimal SWO waveforms across multiple flow control scenarios. Potential improvements could be achieved through implementation of additional control points and more sophisticated polynomial interpolation techniques. Furthermore, optimization of dissimilar heat transfer performance with explicit consideration of control costs presents an interesting avenue for future research (see Appendix A for preliminary

results).

Version française

Ce chapitre étudie les formes de signaux transverses pariétaux optimales pour la réduction de traînée et le transfert thermique dissymétrique, à $\mathrm{Re}_\tau = 200$ et $\mathrm{Pr} = 1$. Le cadre d'optimisation établi au Chapitre 4 est étendu via une paramétrisation polynomiale du signal pariétal. L'étude est conduite dans un scénario de contrôle en boucle ouverte, où l'atteinte de conditions statistiquement stationnaires est garantie par des contraintes de périodicité et de symétrie. Après la présentation et la validation de la paramétrisation du signal basée sur l'interpolation lagrangienne, l'optimisation est d'abord appliquée à la réduction de traînée et au gain énergétique net, s'appuyant sur les travaux antérieurs de Cimarelli et al. (2013). Cette approche est ensuite étendue au cas du transfert thermique dissymétrique.

Pour la réduction de traînée, une nouvelle forme optimale est identifiée, caractérisée par une phase de plateau non constant aux amplitudes extrêmes. Ce signal procure une amélioration de 2,5% de la réduction de traînée accompagnée d'une augmentation de 4,1% de la profondeur de saillie de la couche de Stokes. Cependant, l'augmentation des coûts de contrôle associée rend cette forme sous-optimale en termes de gain énergétique net. Bien que la période optimale corresponde à la valeur de référence, une analyse de sensibilité des paramètres révèle que cette stabilité ne s'étend pas nécessairement à toutes les variations de forme du signal. Une sensibilité accrue aux changements de période par rapport aux modifications de forme est également observée dans les régions quasi-optimales. Pour le gain énergétique net, les signaux optimaux identifiés sont analogue au cas sinusoïdal de référence, suggérant l'absence de formes plus performantes dans l'espace paramétrique étudié.

Pour le transfert thermique dissymétrique, l'algorithme converge rapidement vers des paramètres précis, suggérant une compensation mutuelle des incertitudes statistiques entre les mesures de Nu et C_f. La forme de signal optimale obtenue présente une similarité marquée avec celle de la réduction de traînée, tout en permettant une augmentation du facteur d'analogie. Les phases prolongées de contrainte de Stokes quasi-constante, caractéristiques des signaux à plateau, semblent responsables de l'amélioration du transfert thermique, en favorisant la régénération turbulente. L'analyse de la distribution de phase de A_n et C_f révèle des mécanismes physiques distincts de ceux observés dans le cas sinusoïdal. Une comparaison avec la forme sinusoïdale optimale à $W^+ = 40$ démontre l'avantage de l'optimisation de forme par rapport à une augmentation d'amplitude lorsque les coûts de contrôle sont pris en compte. La période optimale varie avec les paramètres de forme, comme en témoigne son écart par rapport aux valeurs de référence. La relation entre les paramètres de forme et la performance diffère du scénario de réduction de traînée, avec notamment une sensibilité accrue aux coefficients a_1 et a_2 comparée à celle observée pour T^+.

La méthodologie proposée a démontré son efficacité pour identifier des

formes optimales de signaux transverses pariétaux dans divers scénarios de contrôle d'écoulement. Des signaux plus complexes pourraient être obtenus par l'implémentation de points de contrôle supplémentaires et de techniques d'interpolation polynomiale plus sophistiquées. De plus, l'optimisation du transfert thermique dissymétrique avec une considération explicite des coûts de contrôle représente une voie prometteuse pour de futures recherches (voir l'annexe A pour des résultats préliminaires.).

Chapter 6

Conclusions and perspectives

Contents

6.1 Conclusions 185
6.2 Perspectives 188

6.1 Conclusions

The investigation of wall-bounded turbulence holds broad relevance, as the majority of turbulent flows encountered in engineering applications are constrained by solid boundaries. This thesis specifically examined the enhancement of heat transfer while minimizing associated friction increase in a turbulent channel flow through spanwise wall oscillations, a control strategy predominantly employed for drag reduction thus far. The planned method leveraged numerical simulations and machine learning methodologies, particularly deep reinforcement learning techniques, to identify optimal control strategies, potentially integrated with data-driven ROMs. This research carries practical implications for CSP plants, where improved heat transfer is critical. Transitioning from molten salts to pressurized air as the working fluid could simultaneously boost thermal performance and eliminate risks associated with fluid leakage. While the numerical configuration adopted presently differs from actual power plant conditions in several key aspects (including passive scalar temperature and simplified boundary conditions) this idealized setup offers distinct advantages: (i) it establishes a strong initial Reynolds analogy, making DHT achievement more challenging and therefore more meaningful, (ii) it simplifies the fundamental study of thermal-momentum transport decoupling mechanisms, and (iii) it provides clearer insights into wall-bounded turbulence physics. Crucially, demonstrating DHT under these stringent conditions may imply an even greater dissimilarity in real-world configurations where thermal-momentum decoupling occurs more readily.

An initial simulation is conducted at $\mathrm{Re}_\tau = 180$ and $\mathrm{Pr} = 1$, with oscillation parameters $T^+ = 500$ and $W^+ = 30$, guided by insights from drag control parametric studies in the literature. This control approach achieves an almost twofold preferential enhancement of heat transfer over drag, aligning perfectly with the objectives sought in power plant applications. Although the physical mechanisms underlying this deviation from the Reynolds analogy are not fully understood, they are

thoroughly examined in Chapter 2. It is demonstrated that dissimilarity primarily originates from the pressure term in the streamwise velocity equation, rather than from the forcing term difference. The actuation-induced rotation of the streaks is more pronounced for the temperature field, which could be associated with the decoupling process. Ultimately, the increase of drag and heat transfer occurs despite a reduction in the streak's intensity, due to an amplification of the sweep and ejection events in the wall-normal direction. This suggests that thermal-momentum decoupling relies on more nuanced flow modifications than mere variations in streak strength. Additionally, the work of Rouhi et al. (2025) suggests that the increase in friction is likely connected to a rise in energy within the flow scales at $\lambda_x^+ \approx 100$ due to a large Stokes layer production. Subsequent DNS conducted at $\mathrm{Re}_\tau = 1020$ across $\mathrm{Pr} \in \{0.5, 0.71, 1, 2\}$ demonstrate that higher Prandtl numbers amplify heat transfer enhancement and consequently dissimilarity. This behavior aligns with the mechanism identified in Rouhi et al. (2025), wherein thermal sublayer thinning exposes temperature fluctuations to greater Stokes layer production, thereby more effectively energizing small scales and augmenting heat transfer. The robustness of dissimilarity to Reynolds number variations is maintained across the studied parameter space, with the exception of a slight amplification of Prandtl number effects at elevated Re. However, both drag and heat transfer enhancements are observed to attenuate with increasing Re, indicating that greater energy input becomes necessary to achieve comparable performance at higher Reynolds numbers. These results indicate potential superior response of molten salts to this control strategy compared to pressurized air, with Re-robustness in terms of dissimilarity supporting industrial use. However, idealized numerical conditions and practical challenges, such as mechanical constraints and energy costs at relevant scales may limit real-world applicability.

Based on this initial simulation, a dataset is generated through implicit LES at $\mathrm{Re}_\tau = 200$ and $\mathrm{Pr} = 1$, with fixed $T^+ = 500$ and amplitude variations to evaluate performance on unknown actuation parameters. This dataset, consisting of temperature and velocity fields at $y_{\mathrm{ac}}^+ \approx 15$, is created to develop a data-driven ROM. Such a model has the potential to be integrated with the chosen optimization framework to reduce computational costs. In Chapter 3, the focus is set on a detailed investigation of various dimensionality reduction techniques rather than the creation of a reduced-order dynamical model. More precisely, this examination compares CAEs, extended POD, β-VAEs, and A-CAEs, in terms of compression performance and physical interpretability. The CAE surpasses extended-POD in compression ability, consistently achieving more accurate reconstructions, particularly for lower latent dimensions. The discrepancy in reconstruction primarily affects the smaller scales of the flow, where POD provides only a large-scale approximation. However, extended-POD yields nearly orthogonal modes, unlike the CAE. β-VAEs, when combined with pruning of their latent variables, offer a fine balance between orthogonality and reconstruction accuracy. Autoencoders provide a more effective separation of latent samples with different oscillation amplitudes compared to extended-POD, as well

as a more physically coherent separation of oscillation phases. The β-VAE regularization fails to provide any advantage in terms of separability of the generative factors. The A-CAE demonstrated enhanced prediction of C_f and Nu from latent samples while maintaining primary field accuracy. While the prediction accuracy improvements failed to generalize, the A-CAE nevertheless produced a more coherently structured latent space. This manifests through the best latent separability of physical quantities (C_f, Nu) and generative flow factors (W^+, φ) across all models considered, despite significant entanglement of latent dimensions.

The implementation of a dynamical model in the latent space, necessary for coupling with the optimization process, has not yet been performed. However, the implicit-LES configuration used to create the dataset presented an interesting prospect for this coupling, as they have proven sufficiently accurate while significantly reducing computational costs compared to DNS.

Given the challenges of closed-loop optimization, including extended training times, complex implementation, and difficulties in achieving a statistically steady state (Cavallazzi et al., 2024), it was chosen to operate in an open-loop control scenario, focusing directly on optimizing the control parameters. Despite the initial goal of employing DRL in a closed-loop setup, the chosen approach remains within the realm of reinforcement learning through SRL, an adaptation of DRL suitable for open-loop control. Specifically, the PBO algorithm is employed in Chapter 4, enhancing the simplified PPO algorithm with insights drawn from CMA-ES. By coupling PBO with implicit LES environments, computational costs are significantly reduced, although this introduces non-negligible statistical uncertainty, necessitating additional iterations for convergence. Initially, the optimization targets only the oscillation period and amplitude, focusing first on drag reduction and net energetic gain, areas where optimal parameters are well established. This initial phase validates the accuracy and effectiveness of the PBO-LES framework, and provides insights into parameter sensitivities and more precise optimal ranges for net energetic gain. Subsequently, the optimization framework is applied to heat transfer enhancement and dissimilar heat transfer optimization at $\mathrm{Pr} = 1$, which aligns closely with the initial objectives and received considerably less attention in the literature. Optimal parameter ranges are identified for both scenarios. Maximizing heat transfer enhancement is found to be strongly correlated with maximizing Stokes layer protrusion height, likely up to a certain threshold, suggesting that the optimal period could increase for higher Reynolds numbers. However, this is not the case for dissimilar heat transfer, where a higher dissimilarity sometimes corresponds to lower Stokes layer protrusion height. Contrary to the drag reduction and heat transfer augmentation scenarios, the optimal period for dissimilar heat transfer is not stable with amplitude variations, instead exhibiting an inverse relationship between the two parameters.

Finally, to introduce additional complexity while maintaining an open-loop control framework, the waveform of the signal is parameterized through polynomial interpolation. The PBO-LES waveform optimization approach is applied to both drag

reduction and dissimilar heat transfer scenarios in Chapter 5. A new bi-modal waveform is obtained for optimal drag reduction, resulting in a more protrusive Stokes layer and enhanced drag reduction compared to the sinusoidal case. Although the waveform parameters exhibit non-negligible correlations with the oscillation period, the algorithm converges towards the same optimal T^+ as the sinusoidal case. However, the algorithm fails to identify a superior waveform to the sinusoidal one for net energetic gain, even when considering different periods and amplitudes. The optimal waveform for dissimilar heat transfer at $W^+ = 30$ is highly similar to the optimal drag reduction waveform. In this case, however, the optimal period is lower than that of the optimal sinusoidal case, resulting in a less protrusive Stokes layer yet an increase in $\overline{A_n}$ nearly equivalent to a W^+ increase of 10 and at a reduced control cost. This suggests that optimizing dissimilar heat transfer, while considering control costs, should incorporate waveform parameters to enhance performance. The prolonged lingering Stokes strain characteristic of plateau waveforms likely enhances drag and heat transfer augmentations by promoting more efficient turbulence regeneration. The physical mechanisms associated with dissimilar heat transfer are also shown to be significantly different from those of the baseline case.

6.2 Perspectives

Throughout this thesis, numerous choices were made, each presenting alternative options and potential avenues for follow-up studies.

Regarding the numerical setup, investigations were conducted using a simple configuration to focus on the fundamental characterization of Reynolds analogy breaking. Exploring a configuration closer to practical applications, such as those used in concentrated solar power plants, thus holds great potential. Extending the current results to such a setup could directly impact the efficiency of these power plants, potentially offering significant advancements in renewable energy and contributing to climate change mitigation.

Although comprehensively investigated, the physical mechanisms behind the breaking of Reynolds analogy at a Prandtl number of 1 are not yet completely comprehended. Further investigations into the vorticity components or turbulent budgets could help elucidate how the pressure term in the momentum equation results in the decoupling of thermal and momentum transport.

The implementation of a dynamical model within the latent space of the CAE, A-CAE, or β-VAE offers a promising outlook, as it would complete the dimensionality reduction investigation. A possible approach involves using time series forecasting, for example, utilizing LSTMs or transformers (Solera-Rico et al., 2024). While such a model may yield accurate predictions on the present dataset, this dataset only accounts for amplitude variations. Because the flow's response to amplitude changes is straightforward and predictable, this model would not be well-suited for parametric studies. As demonstrated in Chapter 5, such studies require analyzing

how the flow responds to changes in multiple parameters (such as frequency and waveform) rather than amplitude alone. Therefore, to be relevant for parametric studies, a ROM must also be robust to variations in frequency and waveform of the input signal, which presents an additional challenge. Furthermore, other methods to estimate drag and heat transfer from the latent samples could be considered, potentially involving established empirical laws (Kendall and Koochesfahani, 2006; Kader, 1981) using the reconstructed snapshots. The creation of a CROM (Kaiser et al., 2014) with the current dataset may yield additional insights into the physical mechanisms behind dissimilar heat transfer.

In this thesis, various parametric studies considered the optimization of amplitude, period, and waveform. Although excluding Pr from the optimization process seems warranted, conducting targeted optimizations at fixed Pr values offers a valuable opportunity to assess the generalizability of optimal oscillating conditions in DHT scenarios. The streamwise frequency (k_x) as an additional parameter is also an interesting possibility, as it demonstrated a strong potential for drag enhancement (Knoop et al., 2025). Additionally, incorporating control costs into the waveform optimization for DHT presents an exciting prospect. Indeed, while the optimal waveform for DHT has shown promising results in this regard, a preliminary study (Appendix A) highlights the significant challenge of achieving heat transfer enhancement with limited energetic cost when using sinusoidal waveforms. In terms of waveform optimization, more complex parameterizations could be considered, using harmonic functions or Splines with additional control points for greater flexibility and local control.

An extension to closed-loop flow control applications using classical deep reinforcement learning is also a promising direction for future studies, as learning a state-action relationship could significantly enhance performances. However, this approach may lead to temporarily improved outcomes, as achieving a statistically steady state is not guaranteed (Cavallazzi et al., 2024).

Appendices

Appendix A

Maximizing heat-transfer enhancement with control cost penalty

The ability of SWO to achieve significant dissimilar heat transfer has been demonstrated in Chapter 4 and Chapter 5. However, these optimizations did not incorporate the energetic cost of control within their reward functions, a critical notion for practical applications.

This appendix presents preliminary investigations aimed at achieving heat-transfer enhancement ($\mathrm{Nu}_r > 1$) while ensuring a positive global relative gain in energy. The relative gain in heat transfer is expressed as $\mathrm{Nu}_r - 1$, while the energetic gains due to drag variations are represented by P_{net}, accounting for control costs through P_{req} as defined in Eq. (4.3).

An initial objective is to satisfy:

$$(\mathrm{Nu}_r - 1) + P_{\mathrm{net}} > 0, \tag{A.1}$$

which ensures that the control strategy yields a net energetic benefit. However, this approach has limitations, as it directly compares thermal energy gains with momentum energy losses due to friction and control costs. In some applications, the importance of thermal energy gains may outweigh that of momentum losses, leading to overall efficiency improvements even when Eq. (A.1) is not achieved. To avoid ambiguity, an ideal control strategy would result in:

$$\mathrm{Nu}_r > 1 \quad \text{and} \quad P_{\mathrm{net}} > 0, \tag{A.2}$$

ensuring simultaneous improvements in both thermal and momentum efficiency.

Two PBO investigations were conducted to explore these objectives, focusing on achieving at least the goal defined in Eq. (A.1). The optimizations were performed on a sinusoidal signal, over 10 generations with 8 individuals each, which was sufficient to reach a plateau in the reward function and assess the feasibility of achieving efficient dissimilar heat transfer with control cost constraints. The parameter space was constrained to $T^+ \in [10, 750]$ and $W^+ \in [3, 30]$, using the PBO configuration detailed in Section 4.3.3. The lower bound $W^+ > 3$ is imposed to prevent premature convergence to the unactuated case, particularly in the case where Eq. (A.1) should prove challenging to achieve.

The first optimization process aimed to maximize $(\mathrm{Nu}_r - 1) + P_{\mathrm{net}}$ by adopting it as the reward function. The best rewards plateaued at -0.019, indicating that a net energetic gain is not achievable under the current framework, considering both temperature and momentum equally. This also implies that the objective defined in Eq. (A.2) remains unattainable, as no individual simultaneously satisfies $P_{\mathrm{net}} > 0$ and $\mathrm{Nu}_r > 1$. The best-performing individuals exhibited $P_{\mathrm{net}} \in [0, 0.04]$ and $\mathrm{Nu}_r \in [0.935, 0.98]$, suggesting that the optimization prioritizes favorable P_{net} values at the expense of heat-transfer reduction (see Figure A.2). This highlights the significant energetic penalty associated with heat-transfer enhancement. The best-performing individuals also exhibit $\overline{A_n}$ values within the range $[0.996, 1.004]$, reinforcing the observation that the optimization prioritizes P_{net} over dissimilarity. Furthermore, the agent consistently favors longer periods ($T^+ > 400$) and converges toward the lower boundary of $W^+ = 3$, strongly suggesting that the unactuated case may indeed be optimal under the given constraints.

The second optimization focused on achieving Nusselt number enhancement while still incorporating a control cost penalty, which may be particularly relevant in certain industrial contexts. The reward function for this investigation is defined as:

$$(\mathrm{Nu}_r - 1) + P_{\mathrm{net}} + \frac{1}{1 + e^{-100(\mathrm{Nu}_r - 1)}}. \tag{A.3}$$

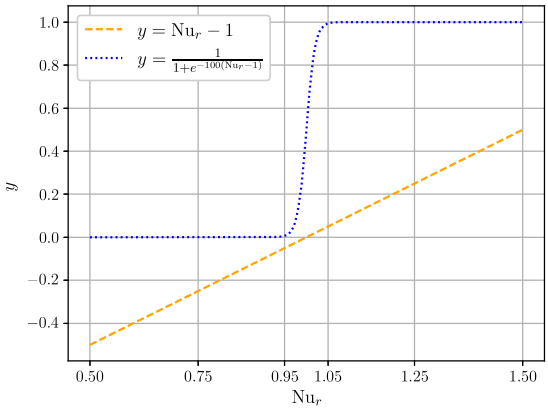

Figure A.1: Illustration of the Nu_r terms in the Eq. (A.3) reward function.

The sigmoid term acts as a soft constraint, creating a strong incentive for the agent to push Nu_r above a threshold. As illustrated in Figure A.1, this term features a steep slope for Nu_r values between 0.95 and 1.05. This sharp transition encourages the agent to achieve a Nusselt number enhancement of at least 5%, ensuring this objective is prioritized over the maximization of P_{net}.

As anticipated, none of the individuals achieved the objective defined in Eq. (A.1) (see Figure A.2). However, the best-performing individuals did attain a moderate heat-transfer enhancement ($\mathrm{Nu}_r \in [1.014, 1.047]$ and $\overline{A_n} \in [1.022, 1.036]$), albeit

with corresponding P_{net} values in the range $[-0.83, -0.46]$. While this approach yields lower combined values of $P_{\text{net}} + (\text{Nu}_r - 1)$ compared to the first case (with the best score being -0.45 among the top-performing individuals and -0.22 globally), it ensures that the control strategy successfully achieves the targeted heat-transfer enhancement, along with moderate dissimilar heat transfer. Notably, the agent converges toward the upper boundary of $T^+ = 750$, indicating that longer periods may improve outcomes, and stabilizes around $W^+ \approx 17$, reflecting a reduction in the P_{net} penalty due to the reward function design.

The results, illustrated in Figure A.2, confirm that neither Eq. (A.1) nor Eq. (A.2) was achieved in either optimization process, underscoring the limitations of SWO in delivering effective heat-transfer enhancement when control costs and drag penalties are considered. The figure also reveals that even modest heat-transfer enhancements incur substantial energetic costs, with P_{net} dropping below -0.5. Larger enhancements of above 5% further reduce momentum net energetic gains, reaching as low as $P_{\text{net}} \approx -1.1$.

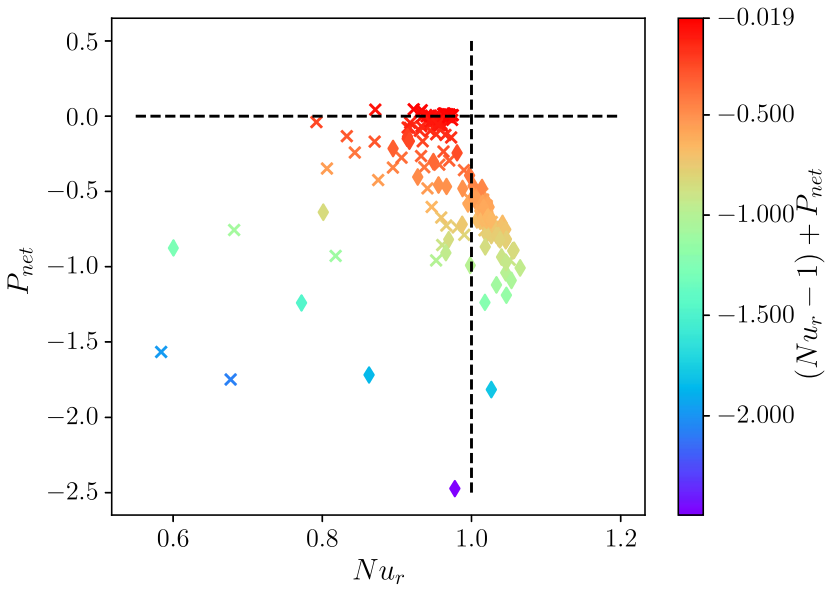

Figure A.2: Summary of all individuals obtained during the first (crosses) and second (diamonds) PBO investigations, focusing on optimal dissimilar heat transfer with a control cost penalty. The colorbar represents the combined relative gains in thermal and momentum energy, including control costs, while the black dashed lines indicate the thresholds for positive gains in P_{net} and Nu_r.

In summary, this appendix has explored the feasibility of achieving dissimilar heat-transfer enhancement while accounting for the energetic cost of control, a critical consideration for practical applications. The investigations revealed that, under the current optimization framework, simultaneously satisfying both heat-transfer enhancement ($\text{Nu}_r > 1$) and a positive net energetic gain ($P_{\text{net}} > 0$) remains challenging. Additionally achieving the easier objective $(\text{Nu}_r - 1) + P_{\text{net}} > 0$ also proved

unattainable. The results underscore the inherent trade-off between thermal and momentum energy gains, with even modest improvements in heat transfer often accompanied by significant energy costs.

However, this study remains preliminary and the inclusion of waveform parameter optimization or the consideration of larger oscillation periods may pave the way for more favorable outcomes in future studies (see Section 5.4). Therefore, future work could incorporate these changes or even consider a closed-loop control implementation of SWO, optimized with DRL methods. While the energetic penalty associated with enhanced heat transfer can be substantial, such enhancements may still hold considerable value in specific applications where thermal performance is prioritized. This highlights the importance of tailoring control strategies to the unique demands of each application through reward shaping, balancing the relative significance of thermal and momentum considerations.

Appendix B

Injectivity of $\mathcal{L}$

The purpose of this appendix is to provide a formal mathematical proof for the uniqueness of the signal generation method described in Chapter 5.2. Specifically, it is demonstrated that the function $\mathcal{L}$, defined in Eq. (5.8), which maps the parameters $(a_1, a_2, T^+, W^+_{max})$ to the resulting signal waveform (represented by the polynomial $\widetilde{P_1}(t)$), is injective. Injectivity, in this context, guarantees that distinct combinations of input parameters consistently produce distinct output signals, ensuring that no two different parameter sets lead to the exact same waveform. Establishing this property is mandatory for the optimization process conducted in Chapter 5. It ensures a well-defined search space by confirming that each point explored by the optimization algorithm corresponds to a unique wall signal. This uniqueness eliminates ambiguity and supports the reliable convergence of the algorithm towards an optimal parameter set.

Proposition: The application $\mathcal{L}$ defined in Eq. (5.8):

$$\begin{aligned} \mathcal{L} : [0,1] \times [0,1] \times [10,300] \times]0,30] &\to \mathbb{R}_4[X] \\ (a, b, T, W) &\mapsto \widetilde{P_1}(t) \end{aligned}$$

is injective.

Proof. It is assumed that $\widetilde{Q_1}(t)$ is the output of two different sets of parameters, i.e.

$$\widetilde{Q_1}(t) = \mathcal{L}(a_1, b_1, T_1, W_1) = \mathcal{L}(a_2, b_2, T_2, W_2).$$

Given that $W_1 > 0$ and $W_2 > 0$, the equalities $T_1 = T_2$ and $W_1 = W_2$ follow directly, as the period and maximal amplitude of $\widetilde{Q_1}(t)$ are uniquely determined by these parameters.

Let $\{y_{1i}\}_{i=1,\dots,5}$ represent the pre-scaling control point weights associated with parameters a_1 and b_1, and $\{y_{2i}\}_{i=1,\dots,5}$ those arising from a_2 and b_2 (see Eq. (5.7) for the mapping). Since the same scaled polynomial $\widetilde{Q_1}(t)$ is produced by these control weights, they must be equal to each other after scaling (see Eq. (5.6)) is applied, i.e., they must be proportional before this operation. In other words, there exists $k \in \mathbb{R}$ such that

$$y_{1i} = k\, y_{2i} \quad \forall\, i \in \{1, \dots, 5\}.$$

Now, as is visible in Eq. (5.7), the sum of the control weights in both cases is equal to 1. That is, $\sum_{i=1}^{5} y_{1i} = k \sum_{i=1}^{5} y_{2i} = 1$.

Since $\sum_{i=1}^{5} y_{2i} = 1$, it follows that $k = 1$. Thus, it is concluded that:

$$y_{1i} = y_{2i} \quad \forall\, i \in \{1, \ldots, 5\}.$$

Given that the mapping of Eq. (5.7) is evidently injective, it can now be deduced that $a_1 = a_2$ and $b_1 = b_2$. Thus, it has been proven that:

$$\mathcal{L}(a_1, b_1, T_1, W_1) = \mathcal{L}(a_2, b_2, T_2, W_2) \quad \Longrightarrow \quad (a_1, b_1, T_1, W_1) = (a_2, b_2, T_2, W_2),$$

which means that the application $\mathcal{L}$ is injective.

□

Therefore, it has been formally demonstrated that the signal parameterization method employed, based on constrained Lagrangian interpolation points, leads to a unique mapping from the parameter space (a_1, a_2, T^+, W^+_{max}) to the resulting polynomial $\widetilde{P}_1(t)$ (or equivalently to the spanwise wall velocity signal $\widetilde{w}_{\text{w}}(t)$). This proven injectivity validates the robustness of the parameter space definition and underpins the reliability of the optimization process detailed in Chapter 5.

Bibliography

R. Abadía-Heredia, M. López-Martín, B. Carro, J. Arribas, J. Pérez, and S. Le Clainche. A predictive hybrid reduced order model based on proper orthogonal decomposition combined with deep learning architectures. Expert Systems with Applications, 187:115910, Jan. 2022. ISSN 09574174. doi: 10.1016/j.eswa.2021.115910. URL https://linkinghub.elsevier.com/retrieve/pii/S0957417421012653. (Cited on page 27.)

H. A. Abdulbari, R. Zulkifli, M. Z. Salleh, and M. Z. Yusoff. Skin-Friction Drag Reduction Using Riblets, Dimples, and Oscillating Wall. Journal of Mechanical Engineering and Sciences, 4:452–464, 2013. (Cited on pages 27 and 32.)

H. Abe, H. Kawamura, and H. Choi. Very Large-Scale Structures and Their Effects on the Wall Shear-Stress Fluctuations in a Turbulent Channel Flow up to Reτ=640. Journal of Fluids Engineering, 126(5):835–843, Sept. 2004a. ISSN 0098-2202, 1528-901X. doi: 10.1115/1.1789528. URL https://asmedigitalcollection.asme.org/fluidsengineering/article/126/5/835/464663/Very-LargeScale-Structures-and-Their-Effects-on. (Cited on page 41.)

H. Abe, H. Kawamura, and Y. Matsuo. Surface heat-flux fluctuations in a turbulent channel flow up to Reτ=1020 with Pr=0.025 and 0.71. International Journal of Heat and Fluid Flow, 25(3):404–419, June 2004b. ISSN 0142727X. doi: 10.1016/j.ijheatfluidflow.2004.02.010. URL https://linkinghub.elsevier.com/retrieve/pii/S0142727X04000153. (Cited on page 59.)

H. Abe, R. A. Antonia, and H. Kawamura. Correlation between small-scale velocity and scalar fluctuations in a turbulent channel flow. Journal of Fluid Mechanics, 627:1–32, May 2009. ISSN 0022-1120, 1469-7645. doi: 10.1017/S0022112008005569. URL https://www.cambridge.org/core/product/identifier/S0022112008005569/type/journal_article. (Cited on pages 35, 37, 38, 39, 40, 70, 117 and 118.)

L. Agostini. Exploration and prediction of fluid dynamical systems using auto-encoder technology. Physics of Fluids, 32(6):067103, June 2020. ISSN 1070-6631. doi: 10.1063/5.0012906. URL https://doi.org/10.1063/5.0012906. (Cited on pages 21 and 25.)

L. Agostini and M. Leschziner. On the departure of near-wall turbulence from the quasi-steady state. Journal of Fluid Mechanics, 871, 2019. Publisher: Cambridge University Press. (Cited on page 9.)

L. Agostini and M. Leschziner. Auto-encoder-assisted analysis of amplitude and wavelength modulation of near-wall turbulence by outer large-scale structures in

channel flow at friction Reynolds number of 5200. Physics of Fluids, 34(11): 115142, Nov. 2022. ISSN 1070-6631. doi: 10.1063/5.0123119. URL https://doi.org/10.1063/5.0123119. (Cited on pages 4 and 68.)

L. Agostini, E. Touber, and M. Leschziner. Spanwise oscillatory wall motion in channel flow: drag-reduction mechanisms inferred from DNS-predicted phase-wise property variations at. Journal of Fluid Mechanics, 743:606 –635, 2014. (Cited on pages 17, 18, 27, 45, 50, 115 and 142.)

L. Agostini, E. Touber, and M. Leschziner. The turbulence vorticity as a window to the physics of friction-drag reduction by oscillatory wall motion. International Journal of Heat and Fluid Flow, 51:3–15, 2015. (Cited on pages 16, 18, 46 and 115.)

R. Akhavan, W. Jung, and N. Mangiavacchi. Control of wall turbulence by high frequency spanwise oscillations. In 3rd Shear Flow Conference, Orlando,FL,U.S.A., July 1993. American Institute of Aeronautics and Astronautics. doi: 10.2514/6.1993-3282. URL https://arc.aiaa.org/doi/10.2514/6.1993-3282. (Cited on page 17.)

F. Alcántara Ávila. Study of the thermal field of turbulent channel flows via Direct Numerical Simulations. PhD thesis, Universitat Politècnica de València, Valencia (Spain), Dec. 2021. URL https://riunet.upv.es/handle/10251/180122. (Cited on pages 13, 14 and 15.)

S. Asidin, W. Kuntjoro, and M. Z. Salleh. A Review on Flow Control Methods in Turbine Internal Cooling. Applied Sciences, 9(16):3233, 2019. (Cited on page 32.)

S. Banerjee, R. Krahl, F. Durst, and C. Zenger. Presentation of anisotropy properties of turbulence, invariants versus eigenvalue approaches. Journal of Turbulence, 8: N32, Jan. 2007. ISSN 1468-5248. doi: 10.1080/14685240701506896. URL https://www.tandfonline.com/doi/full/10.1080/14685240701506896. (Cited on page 5.)

D. Bank, N. Koenigstein, and R. Giryes. Autoencoders, Apr. 2021. URL http://arxiv.org/abs/2003.05991. arXiv:2003.05991 [cs]. (Cited on pages 24 and 25.)

P. Bartholomew, G. Deskos, R. A. Frantz, F. N. Schuch, E. Lamballais, and S. Laizet. Xcompact3D: An open-source framework for solving turbulence problems on a Cartesian mesh. SoftwareX, 12:100550, July 2020. ISSN 23527110. doi: 10.1016/j.softx.2020.100550. URL https://linkinghub.elsevier.com/retrieve/pii/S2352711019303620. (Cited on pages 15, 37 and 70.)

G. Berkooz, P. Holmes, and J. L. Lumley. The proper orthogonal decomposition in the analysis of turbulent flows. Annual review of fluid mechanics, 25(1):539–575, 1993. (Cited on pages 21, 22 and 23.)

P. Bevanda, S. Sosnowski, and S. Hirche. Koopman operator dynamical models: Learning, analysis and control. Annual Reviews in Control, 52:197–212, 2021. ISSN 13675788. doi: 10.1016/j.arcontrol.2021.09.002. URL https://linkinghub.elsevier.com/retrieve/pii/S1367578821000729. (Cited on page 107.)

J. Bird, M. Santer, and J. F. Morrison. Experimental Control of Turbulent Boundary Layers with In-plane Travelling Waves. Flow, Turbulence and Combustion, 100(4):1015–1035, June 2018. ISSN 1386-6184, 1573-1987. doi: 10.1007/s10494-018-9926-2. URL http://link.springer.com/10.1007/s10494-018-9926-2. (Cited on page 17.)

J. Borée. Extended proper orthogonal decomposition: a tool to analyse correlated events in turbulent flows. Experiments in Fluids, 35(2):188–192, Aug. 2003. ISSN 1432-1114. doi: 10.1007/s00348-003-0656-3. URL https://link.springer.com/article/10.1007/s00348-003-0656-3. Company: Springer Distributor: Springer Institution: Springer Label: Springer Number: 2 Publisher: Springer-Verlag. (Cited on pages 68, 76 and 83.)

P. Bradshaw and N. S. Pontikos. Measurements in the turbulent boundary layer on an 'infinite' swept wing. Journal of Fluid Mechanics, 159(-1):105, Oct. 1985. ISSN 0022-1120, 1469-7645. doi: 10.1017/S0022112085003123. URL http://www.journals.cambridge.org/abstract_S0022112085003123. (Cited on page 43.)

S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. Proceedings of the National Academy of Sciences, 113(15):3932–3937, Apr. 2016. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1517384113. URL https://pnas.org/doi/full/10.1073/pnas.1517384113. (Cited on page 107.)

M. A. Bucci, O. Semeraro, A. Allauzen, G. Wisniewski, L. Cordier, and L. Mathelin. Control of chaotic systems by deep reinforcement learning. Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences, 475(2231):20190351, 2019. doi: 10.1098/rspa.2019.0351. URL https://royalsocietypublishing.org/doi/abs/10.1098/rspa.2019.0351. _eprint: https://royalsocietypublishing.org/doi/pdf/10.1098/rspa.2019.0351. (Cited on page 112.)

G. M. Cavallazzi, L. Guastoni, R. Vinuesa, and A. Pinelli. Deep Reinforcement Learning for the Management of the Wall Regeneration Cycle in Wall-Bounded Turbulent Flows. Flow, Turbulence and Combustion, Nov. 2024. ISSN 1386-6184, 1573-1987. doi: 10.1007/s10494-024-00609-4. URL https://link.springer.com/10.1007/s10494-024-00609-4. (Cited on pages 161, 187 and 189.)

D. Chandran, A. Zampiron, A. Rouhi, M. Fu, D. Wine, B. Holloway, A. Smits, and I. Marusic. Turbulent drag reduction by spanwise wall

forcing. Part 2. High-Reynolds-number experiments. Journal of Fluid Mechanics, 968:A7, Aug. 2023. ISSN 0022-1120, 1469-7645. doi: 10.1017/jfm.2023.498. URL https://www.cambridge.org/core/product/identifier/S0022112023004986/type/journal_article. (Cited on page 17.)

K.-S. Choi and B. R. Clayton. The mechanism of turbulent drag reduction with wall oscillation. International Journal of Heat and Fluid Flow, 22(1):1–9, Feb. 2001. ISSN 0142727X. doi: 10.1016/S0142-727X(00)00070-9. URL https://linkinghub.elsevier.com/retrieve/pii/S0142727X00000709. (Cited on page 16.)

A. Cimarelli, B. Frohnapfel, Y. Hasegawa, E. De Angelis, and M. Quadrio. Prediction of turbulence control for arbitrary periodic spanwise wall movement. Physics of Fluids, 25(7), July 2013. ISSN 1070-6631. doi: 10.1063/1.4813807. URL https://pubs.aip.org/aip/pof/article/25/7/075102/314579/Prediction-of-turbulence-control-for-arbitrary. Publisher: AIP Publishing. (Cited on pages 18, 160, 161, 166, 167, 170, 171, 173, 174, 178, 181 and 182.)

E. Climate. Global electricity review 2024. Report, Ember, 2025. URL https://ember-climate.org/insights/research/global-electricity-review-2024/. (Cited on page 1.)

L. Cordier and M. Bergmann. Proper Orthogonal Decomposition: an overview. In Lecture series 2002-04, 2003-03 and 2008-01 on post-processing of experimental and numerical data, pages 1–46. Von Kármán Institute for Fluid Dynamics, 2008. ISBN 978-2-930389-80-X. (Cited on page 75.)

R. V. Cruz. High-fidelity simulation of conjugate heat transfer between a turbulent flow and a duct geometry. phdthesis, Université de Poitiers, Dec. 2021. URL https://theses.hal.science/tel-03605404. (Cited on pages 13 and 15.)

H. Csala, S. T. M. Dawson, and A. Arzani. Comparing different nonlinear dimensionality reduction techniques for data-driven unsteady fluid flow modeling. Physics of Fluids, 34(11), Nov. 2022. ISSN 1070-6631. doi: 10.1063/5.0127284. URL https://pubs.aip.org/aip/pof/article/34/11/117119/2848789/Comparing-different-nonlinear-dimensionality. Publisher: AIP Publishing. (Cited on pages 21 and 77.)

S. Dankwa and W. Zheng. Twin-Delayed DDPG: A Deep Reinforcement Learning Technique to Model a Continuous Movement of an Intelligent Robot Agent. In Proceedings of the 3rd International Conference on Vision, Image and Signal Processing, pages 1–5, Vancouver BC Canada, Aug. 2019. ACM. ISBN 978-1-4503-7625-9. doi: 10.1145/3387168.3387199. URL https://dl.acm.org/doi/10.1145/3387168.3387199. (Cited on pages 119 and 121.)

M. David, A. Toutant, and F. Bataille. Study of asymmetrically heated flows passing through gas-pressurized solar receivers using Direct Numerical Simulations. International Journal of Heat and Mass Transfer, 201:123577, Feb. 2023. ISSN 00179310. doi: 10.1016/j.ijheatmasstransfer.2022.123577. URL https://linkinghub.elsevier.com/retrieve/pii/S0017931022010468. (Cited on pages 3 and 11.)

W. Ding and T. Bauer. Progress in Research and Development of Molten Chloride Salt Technology for Next Generation Concentrated Solar Power Plants. Engineering, 7(3):334–347, Mar. 2021. ISSN 20958099. doi: 10.1016/j.eng.2020.06.027. URL https://linkinghub.elsevier.com/retrieve/pii/S2095809921000473. (Cited on pages 2 and 3.)

H. Dong, Y. Zhang, F. Deng, and W. Liu. The Research on Fluid Properties with Spanwise-Wall Oscillation by LES WALE Model. In 2019 IEEE 4th International Conference on Advanced Robotics and Mechatronics (ICARM), pages 966–969, Toyonaka, Japan, July 2019. IEEE. ISBN 978-1-72810-064-7. doi: 10.1109/ICARM.2019.8834155. URL https://ieeexplore.ieee.org/document/8834155/. (Cited on page 43.)

P. Drawdown. Concentrated solar power: Technology brief. Report, Project Drawdown, 2025. URL https://drawdown.org/solutions/concentrated-solar-power. (Cited on page 2.)

H. Eivazi, S. Le Clainche, S. Hoyas, and R. Vinuesa. Towards extraction of orthogonal and parsimonious non-linear modes from turbulent flows. Expert Systems with Applications, 202:117038, 2022. URL https://www.sciencedirect.com/science/article/pii/S0957417422004535. Publisher: Elsevier. (Cited on pages 25, 68, 82 and 83.)

Enerdata. Global energy statistical yearbook 2024. Technical report, Enerdata, 2024. URL https://yearbook.enerdata.net/. (Cited on page 1.)

J. Fang and L. Lu. Large eddy simulation of compressible turbulent channel flow with active spanwise wall fluctuations. Modern Physics Letters B, 24(13):1457–1460, May 2010. ISSN 0217-9849, 1793-6640. doi: 10.1142/S0217984910023864. URL https://www.worldscientific.com/doi/abs/10.1142/S0217984910023864. (Cited on pages 19, 33 and 45.)

J. Fang, L. Lu, and L. Shao. Large eddy simulation of compressible turbulent channel flow with spanwise wall oscillation. Science in China Series G: Physics, Mechanics and Astronomy, 52(8):1233–1243, Aug. 2009. ISSN 1862-2844. doi: 10.1007/s11433-009-0165-3. URL https://link.springer.com/article/10.1007/s11433-009-0165-3. Company: Springer Distributor: Springer Institution: Springer Label: Springer Number: 8 Publisher: SP Science in China Press. (Cited on pages 19 and 33.)

J. Fang, L.-P. Lu, and L. Shao. Heat transport mechanisms of low Mach number turbulent channel flow with spanwise wall oscillation. Acta Mechanica Sinica, 26(3): 391–399, June 2010. ISSN 0567-7718, 1614-3116. doi: 10.1007/s10409-010-0343-6. URL http://link.springer.com/10.1007/s10409-010-0343-6. (Cited on pages 19 and 33.)

C. Flageul, S. Benhamadouche, E. Lamballais, and D. Laurence. DNS of turbulent channel flow with conjugate heat transfer: Effect of thermal boundary conditions on the second moments and budgets. International Journal of Heat and Fluid Flow, 55:34–44, Oct. 2015. ISSN 0142727X. doi: 10.1016/j.ijheatfluidflow.2015.07.009. URL https://linkinghub.elsevier.com/retrieve/pii/S0142727X15000910. (Cited on pages 13, 38 and 70.)

P. I. Frazier. A Tutorial on Bayesian Optimization, 2018. URL https://arxiv.org/abs/1807.02811. Version Number: 1. (Cited on page 29.)

K. Fukagata, K. Iwamoto, and N. Kasagi. Contribution of Reynolds stress distribution to the skin friction in wall-bounded flows. Physics of Fluids, 14 (11):L73–L76, Oct. 2002. ISSN 1070-6631. doi: 10.1063/1.1516779. URL https://doi.org/10.1063/1.1516779. (Cited on pages 33 and 50.)

K. Fukami and K. Taira. Grasping extreme aerodynamics on a low-dimensional manifold. Nature Communications, 14(1):6480, Oct. 2023. ISSN 2041-1723. doi: 10.1038/s41467-023-42213-6. URL https://www.nature.com/articles/s41467-023-42213-6. (Cited on pages 26 and 98.)

K. Fukami, T. Nakamura, and K. Fukagata. Convolutional neural network based hierarchical autoencoder for nonlinear mode decomposition of fluid field data. Physics of Fluids, 32(9), Sept. 2020. ISSN 1070-6631. doi: 10.1063/5.0020721. URL https://pubs.aip.org/aip/pof/article/32/9/095110/1061028/Convolutional-neural-network-based-hierarchical. Publisher: AIP Publishing. (Cited on pages 26 and 77.)

P. Garnier, J. Viquerat, J. Rabault, A. Larcher, A. Kuhnle, and E. Hachem. A review on Deep Reinforcement Learning for Fluid Mechanics, 2019. URL https://arxiv.org/abs/1908.04127. Publisher: arXiv Version Number: 2. (Cited on pages 4 and 28.)

D. Gatti and M. Quadrio. Performance losses of drag-reducing spanwise forcing at moderate values of the Reynolds number. Physics of Fluids, 25(12):125109, Dec. 2013. ISSN 1070-6631, 1089-7666. (Cited on pages 17, 114, 128 and 131.)

H. Ghraieb, J. Viquerat, A. Larcher, P. Meliga, and E. Hachem. Single-step deep reinforcement learning for open-loop control of laminar and turbulent flows. Physical Review Fluids, 6(5):053902, May 2021. ISSN 2469-990X. doi: 10.1103/PhysRevFluids.6.053902. URL https://link.aps.org/doi/10.1103/PhysRevFluids.6.053902. (Cited on pages 29, 30, 113, 121 and 122.)

T. Gomez, V. Flutet, and P. Sagaut. Contribution of Reynolds stress distribution to the skin friction in compressible turbulent channel flows. Physical Review E, 79(3):035301, Mar. 2009. doi: 10.1103/PhysRevE.79.035301. URL https://journals.aps.org/pre/abstract/10.1103/PhysRevE.79.035301. Publisher: American Physical Society. (Cited on page 54.)

A. Gruber, M. Gunzburger, L. Ju, and Z. Wang. A comparison of neural network architectures for data-driven reduced-order modeling. Computer Methods in Applied Mechanics and Engineering, 393:114764, Apr. 2022. ISSN 00457825. doi: 10.1016/j.cma.2022.114764. URL https://linkinghub.elsevier.com/retrieve/pii/S004578252200113X. (Cited on pages 26 and 68.)

L. Guastoni, J. Rabault, P. Schlatter, H. Azizpour, and R. Vinuesa. Deep reinforcement learning for turbulent drag reduction in channel flows. The European Physical Journal E, 46(4):1–12, Apr. 2023. ISSN 1292-895X. doi: 10.1140/epje/s10189-023-00285-8. URL https://link.springer.com/article/10.1140/epje/s10189-023-00285-8. Company: Springer Distributor: Springer Institution: Springer Label: Springer Number: 4 Publisher: Springer Berlin Heidelberg. (Cited on pages 112, 122 and 161.)

L. Guérin, C. Flageul, L. Cordier, S. Grieu, and L. Agostini. Preferential enhancement of convective heat transfer over drag via near-wall turbulence manipulation using spanwise wall oscillations. International Journal of Heat and Fluid Flow, 110:109564, Dec. 2024. ISSN 0142727X. doi: 10.1016/j.ijheatfluidflow.2024.109564. URL https://linkinghub.elsevier.com/retrieve/pii/S0142727X24002893. (Cited on pages 70, 72 and 139.)

L. Guérin, L. Cordier, C. Flageul, S. Grieu, and L. Agostini. Policy-based optimization for drag reduction via spanwise wall oscillations. Neural Computing and Applications, Feb. 2025. ISSN 0941-0643, 1433-3058. doi: 10.1007/s00521-025-11067-y. URL https://link.springer.com/10.1007/s00521-025-11067-y. (Cited on page 140.)

E. Hachem, H. Ghraieb, J. Viquerat, A. Larcher, and P. Meliga. Deep reinforcement learning for the control of conjugate heat transfer with application to workpiece cooling, Mar. 2021. URL http://arxiv.org/abs/2011.15035. arXiv:2011.15035 [physics]. (Cited on pages 113 and 122.)

N. Hansen. The CMA Evolution Strategy: A Tutorial, Mar. 2023. URL http://arxiv.org/abs/1604.00772. arXiv:1604.00772 [cs]. (Cited on pages 29 and 122.)

Y. Hasegawa and N. Kasagi. Dissimilar control of momentum and heat transfer in a fully developed turbulent channel flow. Journal of Fluid Mechanics, 683:57–93, Sept. 2011. ISSN 1469-7645, 0022-1120. doi: 10.1017/jfm.2011.248. Publisher: Cambridge University Press. (Cited on page 54.)

D. Hilberg, W. Lazik, and H. E. Fiedler. The application of classical POD and snapshot POD in a turbulent shear layer with periodic structures. Applied Scientific Research, 53(3-4):283–290, Dec. 1994. ISSN 0003-6994, 1573-1987. doi: 10.1007/BF00849105. URL http://link.springer.com/10.1007/BF00849105. (Cited on page 23.)

M. D. Hoffman and M. J. Johnson. Elbo surgery: yet another way to carve up the variational evidence lower bound. In Workshop in advances in approximate bayesian inference, NIPS, volume 1, 2016. (Cited on page 81.)

P. Holmes, J. L. Lumley, G. Berkooz, and C. W. Rowley. Turbulence, Coherent Structures, Dynamical Systems and Symmetry. Cambridge University Press, 2 edition, Feb. 2012. ISBN 978-0-511-91970-1 978-1-107-00825-0. doi: 10.1017/CBO9780511919701. URL https://www.cambridge.org/core/product/identifier/9780511919701/type/book. (Cited on page 26.)

P. Hämäläinen, A. Babadi, X. Ma, and J. Lehtinen. PPO-CMA: Proximal Policy Optimization with Covariance Matrix Adaptation, 2018. URL https://arxiv.org/abs/1810.02541. Version Number: 9. (Cited on pages 121 and 122.)

R. Jahanbakhshi and T. A. Zaki. Nonlinearly most dangerous disturbance for high-speed boundary-layer transition. Journal of Fluid Mechanics, 876:87–121, Oct. 2019. ISSN 0022-1120, 1469-7645. doi: 10.1017/jfm.2019.527. URL https://www.cambridge.org/core/product/identifier/S0022112019005275/type/journal_article. (Cited on page 29.)

J. Jeon, J. Rabault, J. Vasanth, F. Alcántara-Ávila, S. Baral, and R. Vinuesa. Advanced deep-reinforcement-learning methods for flow control: group-invariant and positional-encoding networks improve learning speed and quality, 2024. URL https://arxiv.org/abs/2407.17822. Version Number: 2. (Cited on page 161.)

J. Jimenez. Coherent structures in wall-bounded turbulence, 2017. URL https://arxiv.org/abs/1710.07493. Publisher: arXiv Version Number: 3. (Cited on pages 5 and 8.)

J. Jiménez. Cascades in Wall-Bounded Turbulence. Annual Review of Fluid Mechanics, 44(1):27–45, Jan. 2012. ISSN 0066-4189, 1545-4479. doi: 10.1146/annurev-fluid-120710-101039. URL https://www.annualreviews.org/doi/10.1146/annurev-fluid-120710-101039. (Cited on page 8.)

W. Jung, N. Mangiavacchi, and R. Akhavan. Suppression of turbulence in wall-bounded flows by high-frequency spanwise oscillations. Physics of Fluids A: Fluid Dynamics (1989-1993), 4(8):1605–1607, 1992. (Cited on pages 16, 19, 32, 33 and 45.)

B. Kader. Temperature and concentration profiles in fully turbulent boundary layers. International Journal of Heat and Mass Transfer, 24(9):1541–1544, Sept. 1981.

ISSN 00179310. doi: 10.1016/0017-9310(81)90220-9. URL https://linkinghub.elsevier.com/retrieve/pii/0017931081902209. (Cited on page 189.)

E. Kaiser, B. R. Noack, L. Cordier, A. Spohn, M. Segond, M. Abel, G. Daviller, J. Östh, S. Krajnović, and R. K. Niven. Cluster-based reduced-order modelling of a mixing layer. Journal of Fluid Mechanics, 754:365–414, 2014. Publisher: Cambridge University Press. (Cited on pages 26, 98, 107 and 189.)

Y.-E. Kang, S. Yang, and K. Yee. Physics-aware Reduced-order Modeling of Transonic Flow via $β$-Variational Autoencoder, 2022. URL https://arxiv.org/abs/2205.00608. Publisher: arXiv Version Number: 2. (Cited on pages 25, 82, 91, 92, 93, 94 and 102.)

N. Kasagi, Y. Tomita, and A. Kuroda. Direct Numerical Simulation of Passive Scalar Field in a Turbulent Channel Flow. Journal of Heat Transfer, 114(3): 598–606, Aug. 1992. ISSN 0022-1481, 1528-8943. doi: 10.1115/1.2911323. URL https://asmedigitalcollection.asme.org/heattransfer/article/114/3/598/383011/Direct-Numerical-Simulation-of-Passive-Scalar. (Cited on pages 13, 34 and 70.)

N. Kasagi, Y. Hasegawa, K. Fukagata, and K. Iwamoto. Control of Turbulent Transport: Less Friction and More Heat Transfer, Mar. 2011. URL https://dx.doi.org/10.1115/IHTC14-23344. (Cited on pages 13, 20, 47 and 151.)

A. Kendall and M. Koochesfahani. A Method for Estimating Wall Friction in Turbulent Boundary Layers. In 25th AIAA Aerodynamic Measurement Technology and Ground Testing Conference, San Francisco, California, June 2006. American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-029-1. doi: 10.2514/6.2006-3834. URL https://arc.aiaa.org/doi/10.2514/6.2006-3834. (Cited on page 189.)

J. Kennedy and R. Eberhart. Particle swarm optimization. In Proceedings of ICNN'95 - International Conference on Neural Networks, volume 4, pages 1942–1948, Perth, WA, Australia, 1995. IEEE. ISBN 978-0-7803-2768-9. doi: 10.1109/ICNN.1995.488968. URL http://ieeexplore.ieee.org/document/488968/. (Cited on page 29.)

J. Kim. Progress in pipe and channel flow turbulence, 1961–2011. Journal of Turbulence, 13:N45, Jan. 2012. ISSN 1468-5248. doi: 10.1080/14685248.2012.726358. URL https://www.tandfonline.com/doi/full/10.1080/14685248.2012.726358. (Cited on pages 7 and 8.)

D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization, 2014. URL https://arxiv.org/abs/1412.6980. Version Number: 9. (Cited on pages 24 and 126.)

D. P. Kingma and M. Welling. An Introduction to Variational Autoencoders. Foundations and Trends® in Machine Learning, 12(4):307–392, 2019. ISSN 1935-8237, 1935-8245. doi: 10.1561/2200000056. URL http://www.nowpublishers.com/article/Details/MAL-056. (Cited on page 80.)

M. W. Knoop, F. H. Hartog, F. F. J. Schrijer, O. W. G. Van Campenhout, M. Van Nesselrooij, and B. W. Van Oudheusden. Experimental assessment of square-wave spatial spanwise forcing of a turbulent boundary layer. Experiments in Fluids, 65(5):65, May 2024. ISSN 0723-4864, 1432-1114. doi: 10.1007/s00348-024-03799-9. URL https://link.springer.com/10.1007/s00348-024-03799-9. (Cited on page 17.)

M. W. Knoop, R. Deshpande, F. F. J. Schrijer, and B. W. Van Oudheusden. Response of a turbulent boundary layer to steady, square-wave-type transverse wall-forcing. Physical Review Fluids, 10(6):064607, June 2025. ISSN 2469-990X. doi: 10.1103/PhysRevFluids.10.064607. URL https://link.aps.org/doi/10.1103/PhysRevFluids.10.064607. (Cited on pages 16, 17, 19, 142, 155, 179 and 189.)

Y. Kuwata. Dissimilar turbulent heat transfer enhancement by Kelvin–Helmholtz rollers over high-aspect-ratio longitudinal ribs. Journal of Fluid Mechanics, 952:A21, Dec. 2022. ISSN 0022-1120, 1469-7645. doi: 10.1017/jfm.2022.915. URL https://www.cambridge.org/core/product/identifier/S0022112022009156/type/journal_article. (Cited on pages 33 and 58.)

P. Lai, R. Wang, W. Zhang, and H. Xu. Parameter optimization of open-loop control of a circular cylinder by simplified reinforcement learning. Physics of Fluids, 33(10):107110, Oct. 2021. ISSN 1070-6631, 1089-7666. doi: 10.1063/5.0068454. URL https://pubs.aip.org/pof/article/33/10/107110/1065279/Parameter-optimization-of-open-loop-control-of-a. (Cited on pages 29, 113, 121 and 122.)

S. Laizet and E. Lamballais. High-order compact schemes for incompressible flows: A simple and efficient method with quasi-spectral accuracy. Journal of Computational Physics, 228(16):5989–6015, Sept. 2009. ISSN 00219991. doi: 10.1016/j.jcp.2009.05.010. URL https://linkinghub.elsevier.com/retrieve/pii/S0021999109002587. (Cited on pages 37 and 70.)

S. Laizet and N. Li. Incompact3d: A powerful tool to tackle turbulence problems with up to O (10^5) computational cores. International Journal for Numerical Methods in Fluids, 67(11):1735–1757, Dec. 2011. ISSN 0271-2091, 1097-0363. doi: 10.1002/fld.2480. URL https://onlinelibrary.wiley.com/doi/10.1002/fld.2480. (Cited on pages 37 and 70.)

E. Lamballais, V. Fortuné, and S. Laizet. Straightforward high-order numerical dissipation via the viscous term for direct and large eddy simulation. Journal of Computational Physics, 230(9):3270–3275, May 2011. ISSN 00219991. doi:

10.1016/j.jcp.2011.01.040. URL https://linkinghub.elsevier.com/retrieve/pii/S0021999111000659. (Cited on pages 38, 70 and 139.)

H. Lee and J. Romero. IPCC, 2023: Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, H. Lee and J. Romero (eds.)]. IPCC, Geneva, Switzerland. Technical report, Intergovernmental Panel on Climate Change (IPCC), July 2023. URL https://www.ipcc.ch/report/ar6/syr/. Edition: First. (Cited on page 1.)

A. J. Linot and M. D. Graham. Dynamics of a data-driven low-dimensional model of turbulent minimal Couette flow. Journal of Fluid Mechanics, 973:A42, Oct. 2023. ISSN 0022-1120, 1469-7645. doi: 10.1017/jfm.2023.720. URL https://www.cambridge.org/core/product/identifier/S0022112023007206/type/journal_article. (Cited on pages 27, 68 and 113.)

A. J. Linot, K. Zeng, and M. D. Graham. Turbulence control in plane Couette flow using low-dimensional neural ODE-based models and deep reinforcement learning. International Journal of Heat and Fluid Flow, 101:109139, June 2023. ISSN 0142727X. doi: 10.1016/j.ijheatfluidflow.2023.109139. URL https://linkinghub.elsevier.com/retrieve/pii/S0142727X23000383. (Cited on pages 68 and 113.)

L. Liu, H. Jiang, P. He, W. Chen, X. Liu, J. Gao, and J. Han. On the variance of the adaptive learning rate and beyond. arXiv preprint arXiv:1908.03265, 2019. (Cited on page 79.)

S. P. Mahulikar and H. Herwig. Fluid friction in incompressible laminar convection: Reynolds' analogy revisited for variable fluid properties. The European Physical Journal B, 62(1):77–86, Mar. 2008. ISSN 1434-6028, 1434-6036. doi: 10.1140/epjb/e2008-00115-0. URL http://link.springer.com/10.1140/epjb/e2008-00115-0. (Cited on page 14.)

I. Marusic, D. Chandran, A. Rouhi, M. K. Fu, D. Wine, B. Holloway, D. Chung, and A. J. Smits. An energy-efficient pathway to turbulent drag reduction. Nature Communications, 12(1):5805, Oct. 2021. ISSN 2041-1723. doi: 10.1038/s41467-021-26128-8. URL https://www.nature.com/articles/s41467-021-26128-8. Number: 1 Publisher: Nature Publishing Group. (Cited on pages 17 and 60.)

R. Mathis, I. Marusic, N. Hutchins, and K. Sreenivasan. The relationship between the velocity skewness and the amplitude modulation of the small scale by the large scale in turbulent boundary layers. Physics of Fluids, 23(12):121702, 2011. (Cited on page 10.)

B. Mohammadikalakoo, M. Kotsonis, and N. A. K. Doan. Optimization of Tollmien-Schlichting waves control: comparison between a deep reinforcement learning and particle swarm optimization approach, Jan. 2024. (Cited on pages 30 and 113.)

A. T. Mohan, N. Lubbers, M. Chertkov, and D. Livescu. Embedding hard physical constraints in neural network coarse-graining of three-dimensional turbulence. Physical Review Fluids, 8(1):014604, Jan. 2023. doi: 10.1103/PhysRevFluids.8.014604. URL https://link.aps.org/doi/10.1103/PhysRevFluids.8.014604. Publisher: American Physical Society. (Cited on pages 26 and 107.)

A. Mosavi, Y. Faghan, P. Ghamisi, P. Duan, S. F. Ardabili, E. Salwana, and S. S. Band. Comprehensive Review of Deep Reinforcement Learning Methods and Applications in Economics. Mathematics, 8(10):1640, Sept. 2020. ISSN 2227-7390. doi: 10.3390/math8101640. URL https://www.mdpi.com/2227-7390/8/10/1640. (Cited on page 28.)

W. Ni, L. Lu, C. Ribault, and J. Fang. Direct Numerical Simulation of Supersonic Turbulent Boundary Layer with Spanwise Wall Oscillation. Energies, 9(3):154, Mar. 2016. ISSN 1996-1073. doi: 10.3390/en9030154. URL https://www.mdpi.com/1996-1073/9/3/154. (Cited on pages 19 and 33.)

R. Nian, J. Liu, and B. Huang. A review On reinforcement learning: Introduction and applications in industrial process control. Computers & Chemical Engineering, 139:106886, 2020. ISSN 0098-1354. doi: https://doi.org/10.1016/j.compchemeng.2020.106886. URL https://www.sciencedirect.com/science/article/pii/S0098135420300557. (Cited on page 28.)

K. Numpacharoen and A. Atsawarungruangkit. Generating Correlation Matrices Based on the Boundaries of Their Coefficients. PLoS ONE, 7(11):e48902, Nov. 2012. ISSN 1932-6203. doi: 10.1371/journal.pone.0048902. URL https://dx.plos.org/10.1371/journal.pone.0048902. (Cited on page 124.)

M. Quadrio and P. Ricco. Initial response of a turbulent channel flow tospanwise oscillation of the walls. Journal of Turbulence, 4(1):007, Mar. 2003. ISSN 1468-5248. doi: 10.1088/1468-5248/4/1/007. URL https://iopscience.iop.org/article/10.1088/1468-5248/4/1/007/meta. Publisher: IOP Publishing. (Cited on page 44.)

M. Quadrio and P. Ricco. Critical assessment of turbulent drag reduction through spanwise wall oscillations. Journal of Fluid Mechanics, 521:251–271, 2004. (Cited on pages 16, 17, 20, 32, 33, 43, 44, 68, 114, 115, 117, 118, 128, 129, 130, 131, 132, 133, 134, 135, 136, 140 and 166.)

M. Quadrio, P. Ricco, and C. Viotti. Streamwise-travelling waves of spanwise wall velocity for turbulent drag reduction. Journal of Fluid Mechanics, 627:161–178, 2009. (Cited on page 16.)

M. Quadrio, A. Chiarini, J. Banchetti, D. Gatti, A. Memmolo, and S. Pirozzoli. Drag reduction on a transonic airfoil. Journal of Fluid Mechanics, 942:R2, July 2022. ISSN 0022-1120, 1469-7645. doi: 10.1017/

jfm.2022.369. URL https://www.cambridge.org/core/product/identifier/S002211202200369X/type/journal_article. (Cited on page 20.)

A. D. Rasamoelina, F. Adjailia, and P. Sinčák. A review of activation function for artificial neural network. In 2020 IEEE 18th World Symposium on Applied Machine Intelligence and Informatics (SAMI), pages 281–286. IEEE, 2020. (Cited on pages 24 and 79.)

REN21. Renewables 2024 global status report. Report, REN21 Secretariat, 2024. URL https://www.ren21.net/gsr-2024/. (Cited on page 3.)

P. Ricco, M. Skote, and M. A. Leschziner. A review of turbulent skin-friction drag reduction by near-wall transverse forcing. Progress in Aerospace Sciences, 123: 100713, 2021. Publisher: Elsevier. (Cited on pages 7, 16, 32 and 161.)

S. Rodat and R. Thonig. Status of Concentrated Solar Power Plants Installed Worldwide: Past and Present Data. Clean Technologies, 6(1):365–378, Mar. 2024. ISSN 2571-8797. doi: 10.3390/cleantechnol6010018. URL https://www.mdpi.com/2571-8797/6/1/18. (Cited on pages 2 and 3.)

A. Rouhi, S. Endrikat, D. Modesti, R. D. Sandberg, T. Oda, K. Tanimoto, N. Hutchins, and D. Chung. Riblet-generated flow mechanisms that lead to local breaking of Reynolds analogy. Journal of Fluid Mechanics, 951:A45, Nov. 2022. ISSN 0022-1120, 1469-7645. doi: 10.1017/jfm.2022.880. URL https://www.cambridge.org/core/product/identifier/S0022112022008801/type/journal_article. (Cited on pages 33 and 58.)

A. Rouhi, M. Fu, D. Chandran, A. Zampiron, A. Smits, and I. Marusic. Turbulent drag reduction by spanwise wall forcing. Part 1. Large-eddy simulations. Journal of Fluid Mechanics, 968:A6, Aug. 2023. ISSN 0022-1120, 1469-7645. doi: 10.1017/jfm.2023.499. URL https://www.cambridge.org/core/product/identifier/S0022112023004998/type/journal_article. (Cited on pages 18, 33 and 142.)

A. Rouhi, M. Hultmark, and A. J. Smits. Spanwise wall forcing can reduce turbulent heat transfer more than drag. Journal of Fluid Mechanics, 1010:A59, May 2025. ISSN 0022-1120, 1469-7645. doi: 10.1017/jfm.2025.310. URL https://www.cambridge.org/core/product/identifier/S0022112025003106/type/journal_article. (Cited on pages 19, 60, 141, 142, 145, 155, 160 and 186.)

D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning representations by back-propagating errors. Nature, 323(6088):533–536, Oct. 1986. ISSN 0028-0836, 1476-4687. doi: 10.1038/323533a0. URL https://www.nature.com/articles/323533a0. (Cited on page 24.)

J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel. High-Dimensional Continuous Control Using Generalized Advantage Estimation, 2015. URL https://arxiv.org/abs/1506.02438. Version Number: 6. (Cited on page 121.)

J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal Policy Optimization Algorithms, Aug. 2017. URL http://arxiv.org/abs/1707.06347. arXiv:1707.06347 [cs]. (Cited on pages 29 and 119.)

F. Sebilleau, R. Issa, S. Lardeau, and S. P. Walker. Direct Numerical Simulation of an air-filled differentially heated square cavity with Rayleigh numbers up to 10 11. International Journal of Heat and Mass Transfer, 123:297–319, Aug. 2018. ISSN 00179310. doi: 10.1016/j.ijheatmasstransfer.2018.02.042. URL https://linkinghub.elsevier.com/retrieve/pii/S0017931018302588. (Cited on page 12.)

Y. Seki, K. Iwamoto, and H. Kawamura. Prandtl number Effect on Turbulence Quantities through High Spatial Resolution DNS of Turbulent Heat Transfer in a Channel Flow. In Proceedings of the International Symposium on Turbulence, Heat and Mass Transfer, pages 301–304, Dubrovnik, Croatia, 2006. Begellhouse. ISBN 978-1-56700-229-4. doi: 10.1615/ICHMT.2006.TurbulHeatMassTransf.560. URL http://www.dl.begellhouse.com/references/1bb331655c289a0a,4a93e7ca5fe696c8,7a3638b515acd70e.html. (Cited on pages 35, 37, 38, 39, 40, 41, 70 and 71.)

M. Shams and A. H. Elsheikh. Gym-preCICE: Reinforcement learning environments for active flow control. SoftwareX, 23:101446, July 2023. ISSN 23527110. doi: 10.1016/j.softx.2023.101446. URL https://linkinghub.elsevier.com/retrieve/pii/S2352711023001425. (Cited on page 122.)

K. Shao, Z. Tang, Y. Zhu, N. Li, and D. Zhao. A Survey of Deep Reinforcement Learning in Video Games, Dec. 2019. URL http://arxiv.org/abs/1912.10944. arXiv:1912.10944 [cs]. (Cited on pages 28 and 120.)

A. Sherstinsky. Fundamentals of Recurrent Neural Network (RNN) and Long Short-Term Memory (LSTM) network. Physica D: Nonlinear Phenomena, 404:132306, Mar. 2020. ISSN 01672789. doi: 10.1016/j.physd.2019.132306. URL https://linkinghub.elsevier.com/retrieve/pii/S0167278919305974. (Cited on page 107.)

A. Slowik and H. Kwasnicka. Evolutionary algorithms and their applications to engineering problems. Neural Computing and Applications, 32(16):12363–12379, Aug. 2020. ISSN 0941-0643, 1433-3058. doi: 10.1007/s00521-020-04832-8. URL http://link.springer.com/10.1007/s00521-020-04832-8. (Cited on pages 114 and 121.)

A. Solera-Rico, C. Sanmiguel Vila, M. Gómez-López, Y. Wang, A. Almashjary, S. T. M. Dawson, and R. Vinuesa. β-Variational autoencoders and transformers for reduced-order modelling of fluid flows. Nature Communications, 15(1):1–15, Feb. 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-45578-4. URL https://www.nature.com/articles/s41467-024-45578-4. Number: 1 Publisher: Nature Publishing Group. (Cited on pages 25, 27, 68, 82, 83, 107 and 188.)

T. Sonoda, Z. Liu, T. Itoh, and Y. Hasegawa. Reinforcement learning of control strategies for reducing skin friction drag in a fully developed turbulent channel flow. Journal of Fluid Mechanics, 960:A30, Apr. 2023. ISSN 0022-1120, 1469-7645. doi: 10.1017/jfm.2023.147. URL https://www.cambridge.org/core/product/identifier/S0022112023001477/type/journal_article. (Cited on pages 112 and 161.)

M. Sumanas, A. Petronis, V. Bucinskas, A. Dzedzickis, D. Virzonis, and I. Morkvenaite-Vilkonciene. Deep Q-Learning in Robotics: Improvement of Accuracy and Repeatability. Sensors, 22(10):3911, May 2022. ISSN 1424-8220. doi: 10.3390/s22103911. URL https://www.mdpi.com/1424-8220/22/10/3911. (Cited on pages 28 and 120.)

S. Tang and R. Antonia. Scaling of small-scale wall turbulence. Journal of Fluid Mechanics, 948:A25, Oct. 2022. ISSN 0022-1120, 1469-7645. doi: 10.1017/jfm.2022.707. URL https://www.cambridge.org/core/product/identifier/S0022112022007078/type/journal_article. (Cited on page 7.)

S. Tang and R. Antonia. Similarity for dissipation-scaled wall turbulence. Journal of Fluid Mechanics, 960:A18, Apr. 2023. ISSN 0022-1120, 1469-7645. doi: 10.1017/jfm.2023.197. URL https://www.cambridge.org/core/product/identifier/S0022112023001970/type/journal_article. (Cited on page 7.)

I. Tiselj. Tracking of large-scale structures in turbulent channel with direct numerical simulation of low Prandtl number passive scalar. Physics of Fluids, 26(12):125111, Dec. 2014. ISSN 1070-6631, 1089-7666. doi: 10.1063/1.4905018. URL https://pubs.aip.org/pof/article/26/12/125111/1023020/Tracking-of-large-scale-structures-in-turbulent. (Cited on page 41.)

E. Touber and M. Leschziner. Near-wall streak modification by spanwise oscillatory wall motion and drag-reduction mechanisms. Journal of Fluid Mechanics, 693: 150–200, 2012. (Cited on pages 16, 44 and 45.)

K. Uchino, H. Mamori, and K. Fukagata. Heat transfer in fully developed turbulent channel flow with streamwise traveling wave-like wall deformation. Journal of Thermal Science and Technology, 12(1):JTST0003–JTST0003, 2017. ISSN 1880-5566. doi: 10.1299/jtst.2017jtst0003. URL https://www.jstage.jst.go.jp/article/jtst/12/1/12_2017jtst0003/_article/-char/ja/. Publisher: The Japan Society of Mechanical Engineers and The Heat Transfer Society of Japan. (Cited on pages 19, 47 and 151.)

C. Vignon, J. Rabault, and R. Vinuesa. Recent advances in applying deep reinforcement learning for flow control: Perspectives and future directions. Physics of Fluids, 35(3):031301, Mar. 2023. ISSN 1070-6631, 1089-7666. doi: 10.1063/5.0143913. URL https://pubs.aip.org/pof/article/35/3/031301/2881609/Recent-advances-in-applying-deep-reinforcement. (Cited on pages 4, 27, 28 and 120.)

C. Viotti, M. Quadrio, and P. Luchini. Streamwise oscillation of spanwise velocity at the wall of a channel for turbulent drag reduction. Physics of fluids, 21(11): 115109, 2009. (Cited on pages 16 and 179.)

J. Viquerat, J. Rabault, A. Kuhnle, H. Ghraieb, A. Larcher, and E. Hachem. Direct shape optimization through deep reinforcement learning, Dec. 2020. URL http://arxiv.org/abs/1908.09885. arXiv:1908.09885 [cs]. (Cited on pages 113, 122 and 174.)

J. Viquerat, R. Duvigneau, P. Meliga, A. Kuhnle, and E. Hachem. Policy-based optimization: single-step policy gradient method seen as an evolution strategy. Neural Computing and Applications, 35(1):449–467, Jan. 2023. ISSN 1433-3058. doi: 10.1007/s00521-022-07779-0. URL https://link.springer.com/article/10.1007/s00521-022-07779-0. Company: Springer Distributor: Springer Institution: Springer Label: Springer Number: 1 Publisher: Springer London. (Cited on pages 29, 30, 113, 122, 124, 126 and 127.)

A. W. Vreman and J. G. M. Kuerten. Comparison of direct numerical simulation databases of turbulent channel flow at $Re\ \tau = 180$. Physics of Fluids, 26(1):015102, Jan. 2014. ISSN 1070-6631, 1089-7666. doi: 10.1063/1.4861064. URL https://pubs.aip.org/pof/article/26/1/015102/314284/Comparison-of-direct-numerical-simulation. (Cited on pages 38 and 40.)

F. Waleffe. On a self-sustaining process in shear flows. Physics of Fluids, 9(4): 883–900, 1997. Publisher: AIP. (Cited on page 19.)

J. M. Wallace. Quadrant Analysis in Turbulence Research: History and Evolution. Annual Review of Fluid Mechanics, 48(1):131–158, Jan. 2016. ISSN 0066-4189, 1545-4479. doi: 10.1146/annurev-fluid-122414-034550. URL https://www.annualreviews.org/doi/10.1146/annurev-fluid-122414-034550. (Cited on pages 10 and 11.)

C. Wang, P. Yu, and H. Huang. Reinforcement-learning-based parameter optimization of a splitter plate downstream in cylinder wake with stability analyses. Physical Review Fluids, 8(8):083904, Aug. 2023. ISSN 2469-990X. doi: 10.1103/PhysRevFluids.8.083904. URL https://link.aps.org/doi/10.1103/PhysRevFluids.8.083904. (Cited on pages 30, 113 and 121.)

J. P. Williams, O. Zahn, and J. N. Kutz. Sensing with shallow recurrent decoder networks. Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences, 480(2298):20240054, Sept. 2024. ISSN 1364-5021, 1471-2946. doi: 10.1098/rspa.2024.0054. URL https://royalsocietypublishing.org/doi/10.1098/rspa.2024.0054. (Cited on page 107.)

D. Xu and Y. Tian. A Comprehensive Survey of Clustering Algorithms. Annals of Data Science, 2(2):165–193, June 2015. ISSN 2198-5804, 2198-5812.

doi: 10.1007/s40745-015-0040-1. URL http://link.springer.com/10.1007/s40745-015-0040-1. (Cited on pages 27 and 84.)

A. Yamamoto, Y. Hasegawa, and N. Kasagi. Optimal control of dissimilar heat and momentum transfer in a fully developed turbulent channel flow. Journal of Fluid Mechanics, 733:189–220, Oct. 2013. ISSN 0022-1120, 1469-7645. doi: 10.1017/jfm.2013.436. Publisher: Cambridge University Press. (Cited on page 20.)

X. Zhang, T. Li, and N. Jiang. Flow Control Strategies for Improving Performance of Turbomachinery: A Review. Progress in Aerospace Sciences, 115:100595, 2020. (Cited on page 32.)

Improving solar-receiver efficiency by manipulating turbulence in heat-collecting, fluid-conveying channels: *From physical understanding to Machine-Learning-based optimization*

Abstract

This study investigates the enhancement of Dissimilar Heat Transfer (DHT) in wall-bounded turbulent flows through spanwise wall oscillations, combining numerical simulations with machine learning approaches. The research develops progressively through three interconnected phases. First, baseline Direct Numerical Simulations (DNS) of an idealized channel flow configuration with strong unactuated Reynolds analogy establish the fundamental potential for preferential heat transfer enhancement over drag increase, with particular emphasis on revealing the physical mechanisms behind thermal-momentum decoupling. These mechanisms are shown to involve complex flow modifications extending beyond simple streak intensity variations. Second, the study conducts a systematic comparison of the following dimensionality reduction techniques for developing data-driven Reduced-Order Models (ROMs): Convolutional Autoencoders (CAEs), Proper Orthogonal Decomposition (POD), and β-Variational Autoencoders (β-VAEs). Third, the application of the Policy Based Optimization (PBO) algorithm, which is adapted from deep reinforcement learning for open-loop control and coupled with implicit Large Eddy Simulation (LES) environments, enables comprehensive parametric studies of drag reduction, heat transfer enhancement, and DHT scenarios. This approach is subsequently extended to include waveform optimization of the input signal through polynomial parameterization. The collective insights from these numerous investigations significantly advance the understanding of spanwise oscillations as a control strategy for thermal transport optimization. The developed methodology and analysis provide a foundation for extending these results to more realistic scenarios, such as concentrated solar power plant applications, where such control strategies could contribute meaningfully to energy efficiency improvements, thus to climate change mitigation efforts.

Keywords: Numerical simulations, wall-bounded turbulent flows, near wall turbulence, dissimilar heat transfer, spanwise wall oscillations, simplified reinforcement learning, policy based optimisation, autoencoders, reduced-order modeling.

Amélioration de l'efficacité des récepteurs solaires par manipulation de la turbulence dans les canaux caloporteurs : *De la compréhension physique à l'optimisation par apprentissage automatique*

Résumé

Cette thèse étudie l'efficacité des oscillations transverses pariétales pour induire un transfert thermique dissymétrique dans les écoulements turbulents pariétaux. L'approche méthodologique repose sur une combinaison de simulations numériques et de techniques d'apprentissage automatique. Le travail se développe selon trois axes interconnectés. Premièrement, des simulations numériques directes d'un canal plan oscillant sont réalisées pour établir des cas de référence. Cette maquette numérique présente une forte analogie de Reynolds dans le cas non contrôlé, ce qui permet d'évaluer le potentiel d'une amélioration préférentielle du transfert thermique par rapport à la traînée. Une attention particulière est portée aux mécanismes physiques à l'origine du découplage thermique-dynamique. Ces mécanismes impliquent des modifications complexes de l'écoulement pariétal. Deuxièmement, une comparaison systématique de techniques de réduction de dimensionnalité est menée dans le but de développer un modèle à dimension réduite. Cette comparaison inclut plusieurs méthodes d'apprentissage automatique, notamment les autoencodeurs convolutionnels, la décomposition orthogonale propre et les autoencodeurs variationnels β. Troisièmement, l'algorithme d'optimisation par politique, adapté de l'apprentissage par renforcement profond et couplé à des simulations implicites aux grandes échelles, permet des études paramétriques exhaustives de l'amplitude et de la période des oscillations. Ces études sont appliquées à la réduction de traînée, à l'augmentation du transfert thermique et au transfert thermique dissymétrique. Cette approche est ensuite étendue à l'optimisation de la forme du signal de contrôle via une paramétrisation polynomiale. Les résultats obtenus améliorent significativement la compréhension des oscillations transverses pariétales en tant que stratégie de contrôle pour l'optimisation du transfert thermique. La méthodologie développée et les analyses menées fournissent des bases solides pour étendre ces travaux à des scénarios plus réalistes. Ces avancées sont particulièrement pertinentes dans le cadre des centrales solaires à concentration, où de telles stratégies de contrôle pourraient significativement améliorer l'efficacité énergétique et contribuer aux efforts d'atténuation du changement climatique.

Mots-clés: Simulations numériques, écoulements turbulents pariétaux, turbulence proche paroi, transfert thermique dissymétrique, oscillations transverses pariétales, apprentissage par renforcement simplifié, optimisation par politique, autoencodeurs, modélisation d'ordre réduit.