



Multi-fidelity Co-Kriging Optimization using Hybrid Injected RANS and LES

Javier Fatou Gómez

4-12-2018



Multi-fidelity Co-Kriging Optimization using Hybrid Injected RANS and LES

MASTER OF SCIENCE THESIS

For obtaining the degree of Master of Science in Aerospace
Engineering at Delft University of Technology

Javier Fatou Gómez

4-12-2018



Copyright © Javier Fatou Gómez
All rights reserved.

DELFT UNIVERSITY OF TECHNOLOGY
DEPARTMENT OF AERODYNAMICS

The undersigned hereby certify that they have read and recommend to the Faculty of Aerospace Engineering for acceptance a thesis entitled “**Multi-fidelity Co-Kriging Optimization using Hybrid Injected RANS and LES**” by **Javier Fatou Gómez** in partial fulfillment of the requirements for the degree of **Master of Science**.

Dated: 4-12-2018

Prof.dr. S. Hickel

Dr. R.P. Dwight

Dr. S.J. Hulshoff

Dr.ir. M. Pini

Abstract

The computation of complex turbulent flows design optimization processes is currently limited by the lack of accuracy of Reynolds-Averaged Navier-Stokes (RANS) in massively separated flows and the infeasible cost of multiple Large-Eddy Simulation (LES) evaluations. A novel method is presented, injecting data from LES or other high-fidelity source such as DNS into the RANS equations, forming a Hybrid Injected RANS (HIRANS) model. The aim is to construct a multi-fidelity design optimization framework that outperforms single-fidelity RANS and LES variants. Two different formulations, injecting a scaled version of the non-dimensional anisotropic part of the Reynolds stress tensor and both isotropic and anisotropic components, are tested in the periodic hill case. A cost-effective LES configuration is assessed, and the agreement of the RANS and HIRANS results with respect to the LES reference is investigated. The original geometry is parametrized using a hill width multiplier, computing several LES evaluations. The injection of LES information from the same geometry into HIRANS and the prediction capabilities when using interpolated data from different geometries are tested. A global design optimization process is computed, using single-fidelity RANS, HIRANS and LES Kriging and multi-fidelity RANS-LES and HIRANS-LES Co-Kriging surrogates. The objective function is based on a combination of turbulent mixing and total pressure losses.

The correction of the mean velocity components required of the injection of both isotropic and anisotropic components for the test case. The LES setup analysis yielded similar results to the reference data with one tenth of grid points and forty percent of its averaging period. The locally corrected HIRANS model successfully reduced the L^2 norms of the Reynolds stresses with respect to LES to a third part of the original RANS values in the fifty-nine LES samples tested, with a modest improvement in the mean velocity components. The non-local corrections yielded irregular results for the mean velocity components, with successful corrections of the Reynolds stresses despite the long distances in the parameter space and different flow features of neighbouring LES cases to interpolate from. In the optimization process, the Co-Kriging LES-HIRANS was not able to outperform the Co-Kriging LES-HIRANS and Kriging LES methods. It improved the initial prediction of the underlying function, but the surrogate yielded artificially low predicted errors far away from the LES samples, leading to an overly exploitative method. An error correction formulation combining two HIRANS fidelity levels was simulated using a modified Kriging believer criterion, outperforming the original formulation and achieving similar results as Kriging LES. The computational efficiency improvements for future research of the Co-Kriging HIRANS are suggested to be linked to an adequate error estimation integration into the surrogate model.

Preface and acknowledgements

When in August 2016 I came back to Delft, after my experience in the UK, it was one of those rare sunny occasions when the Dutch weather forms beautiful rays of light in the canals. Back to the old colleagues, who were already graduating. It felt like...home. “I won’t take much longer”, I said to myself. Formula 1 efficiency at the service of the individual. The result is here, slightly later than expected, 192 pages of work.

I would like to start thanking my two supervisors, who have been together with me in this long journey. It all started when we were ambitiously talking about the periodic hill as an introductory step to test our new ideas, which would lead to new geometries after the concept validation. Never could I have imagined that so much could be extracted from a periodic hill. Professor Stefan Hickel was breathtaking in his broad understanding, where any idea that asked to myself could be easily presented and discussed. The feedback received was consistently tackling the challenges that we were facing at each step, including the increasingly long INCA meetings which ended in succeeding in our goals. From Dr. Richard Dwight, I can only be thankful for his detailed review of the different parts of this document. His willing to explain all the necessary topics that I knew so little about were inspiring, and greatly improved the quality of this work. Furthermore, I would like to thank Martin Schmelzer. Without his initial code, which he openly shared, and his contributions to my doubts, this work would not be the same. In addition, I would also like to thank the discussions in the “Periodic Hill club”, specially to Max, whose contributions in the OpenFOAM/Python integration and honesty were very much appreciated.

I cannot sign this document without acknowledging all the people at Mercedes who were a part of who I am now, and who gave me the tools to perform at the level (whatever that one is) that I can give. I would like to start with Dan. Your passive-aggressive insistence in inserting Python in my head completely changed my contributions level there, and allowed me to learn a beautiful programming language (my acknowledgements to Guido van Rossum). I cannot simply imagine how much time would have all the couplings between different methods taken otherwise. Tom and Felix, thank you for your patience and your human contribution. The feeling of being with incredibly talented people in the technical and the human side made my understand what was the feeling of being in a Team. And thank you, Dani, for your paradigm as a peer, the late nights together and the rides in the Beast.

49. 49 people which I could count in the panel. I did not know all of them. But I do know something: I was in the Basement before that panel was bought. And soon will be my turn. I am sure I will be forgetting people, and I will be only mentioning some, but thanks to Martin, Mathijs, Matthias, Arun, Blanca, Lluís, Edo, Gigi, Mitro, Corrado, Liu, Mart, Dorian, Lucas, David, Neils, Carlos and all the people who have been listening to my complaints and my “I am gonna graduate” whenever a break was over. If you are not here, and we shared nice moments, thank you! My memory tends to not work especially well after 24 hours of not sleeping. I shall not leave out Derek, whose surprise vegetarian pizza are worth the Library effort; Jaime, whom we suffered every Monday, and Fancy, who honoured his name as my artistic consultant. His ideas helped to create and improve many of the nicest plots in this document (Figure 4.9 as an example).

My very first group in Delft started this path. Mario, Edu, Rajat, Amrit, Vishal, and many others. Roberto, our experiences with Hypersonic and with certain Rotterdam place cannot be forgotten. Gonzalo, you have been greatly missed in this time, in the beauty of maths discussion or with Kobi and his merengue. Queena, our honest and almost endless talks and experiences will be long cherished. Andrea, you have been fantastic organizing barbecues and forgetting about how to be late to successfully escape from The Morge. Jordi, (Casacuberta et al., 2018), we finally made it. My gratitude is extended to my second group in Delft. People like María, Yogurín and Ramsés. I cannot forget Ángel. Football, crypto and so many things in between, your support during those months was incredible. Marta, Chufi (DEP), Víctor, Rocío. Madrid, a town surrounded by highways and its never-ending vibe.

Ricky, Nando. You have been with me since I am myself. And you know what that means when the years together surpass an order of magnitude. Esther, we evolve as time passes, and still our cycle keeps running. Paco, there will hopefully be a time when I play with your Swiss robots. Mora, your overwhelming randomness shares injection/implanted levels similar to this work. Alex (del Estal Herrero et al., 2018a,b). You unilaterally decided to be my friend and came to Delft for a year to be together. This time, I am the second one. My words are short in this language to express what I feel.

My family, all of you who have been with me at the different stages, my grandmas, aunts, uncles and cousins, thank you. To my parents and sister, thank you for your support in the uncertainty of these years, in the frustrations, in the excitement, in the moments we shared when time was not looking at us. Thank Yuun. For your love and support, for my cute Latte and my right partner, for Tongariro and Aachen, for the ultimate comfiness levels and for our increasingly adopted us. I cannot wait for the time we provide and celebrate it accordingly in Japan again.

Contents

Abstract	v
Preface and acknowledgements	vii
List of Figures	xviii
List of Tables	xix
Nomenclature	xxi
1 Introduction	1
1.1 Motivation: design optimization in fluid mechanics	1
1.2 Research objectives	4
1.3 Contribution of this work	5
1.4 Outline	7
2 Numerical schemes and computational setup	9
2.1 Introduction	9
2.2 RANS methodology	9
2.2.1 Derivation of RANS equations	9
2.2.2 OpenFOAM $k - \omega$	10
2.2.3 Missing term in OpenFOAM	11
2.2.4 Numerical discretization	12
2.3 LES methodology	13
2.3.1 Derivation of LES equations	13
2.3.2 LES solver: INCA	14
2.4 Kriging and Co-Kriging	14
2.4.1 Surrogate model: Kriging	14
2.4.2 Balance exploration and exploitation: Expected Improvement . . .	17
2.4.3 Extension to multi-fidelity analysis: Co-Kriging	17
2.5 Summary	18

3	Local RANS corrections for reference geometry	19
3.1	Hybrid Injected RANS models formulation	19
3.1.1	HIRANS _a : injection of anisotropy information	19
3.1.2	HIRANS _b : injection of full LES Reynolds stress tensor	21
3.2	Numerical setup	22
3.2.1	Flow parameters and mesh quality	22
3.2.2	Mass flow control	23
3.2.3	Injection strategy and convergence control	23
3.3	Results	24
3.3.1	Baseline RANS model	24
3.3.2	HIRANS _a : injection of anisotropy information	29
3.3.3	HIRANS _b : injection of full LES Reynolds stress tensor	35
3.3.4	Differences between anisotropy and mean flow corrections	38
3.3.5	Sensitivity of the solution to the injection factors ξ_i	40
3.4	Summary	45
4	LES setup and analysis	47
4.1	Introduction and LES setup	47
4.2	Computational performance	47
4.3	Effect of spanwise extent, mesh resolution and averaging time	49
4.3.1	Spanwise extent	49
4.3.2	Grid resolution	52
4.3.3	Averaging period: start and extent	56
4.4	Configuration choice and summary	62
5	HIRANS for different periodic hill geometries	63
5.1	Introduction	63
5.2	Local injection of LES data	63
5.2.1	Geometry definition and study cases	64
5.2.2	Results	65
5.3	HIRANS non-local prediction with sparse LES data	67
5.3.1	Introduction	67
5.3.2	Interpolation techniques	67
5.3.3	Results	69
5.4	Summary	73

6	Kriging and Co-Kriging optimization	75
6.1	Introduction	75
6.2	Setup	76
6.2.1	Parameter space and objective function	76
6.2.2	Routines description	76
6.2.3	Hyperparameter estimation	77
6.2.4	Additional samples acquisition	78
6.2.5	Flowchart	79
6.3	Initial sampling plans	80
6.3.1	Case A: $\beta_A = (-5.0/2)\alpha$, $n_{e_{init}} = 6$	80
6.3.2	Case B: $\beta_B = (-5.83/2)\alpha$, $n_{e_{init}} = 7$	83
6.4	Results	86
6.4.1	Objective function improvements	86
6.4.2	Interpolation and regression	90
6.4.3	Optimum geometries compared to baseline	92
6.5	Influence of HIRANS and Co-Kriging parameters	94
6.5.1	Influence of the exploration-exploitation factor ξ_{EI}	94
6.5.2	Influence of the number of initial samples	96
6.5.3	Influence of the HIRANS injection factors ξ_m, ξ_t	98
6.6	Predicted error correction using multi-level HIRANS	100
6.6.1	Simulation using modified Kriging believer	103
6.7	RANS/HIRANS with respect to LES along the parameter space	107
6.7.1	Local injection	107
6.7.2	Non-local prediction accuracy	109
6.8	Summary	112
7	Conclusions and recommendations	115
7.1	Conclusions	115
7.1.1	Local RANS corrections	115
7.1.2	LES setup and analysis	116
7.1.3	Non-locally injected HIRANS prediction	117
7.1.4	Kriging and Co-Kriging optimization	117
7.1.5	Main research question	119
7.2	Recommendations	120
7.2.1	Local RANS corrections	120
7.2.2	HIRANS interpolation and prediction	121
7.2.3	Kriging and Co-Kriging optimization	121
	References	123
	Appendices	131

A	Injection factor sweep results	133
B	Convergence and results for higher HIRANS injection factors	137
C	Barycentric map for turbulence anisotropy	141
C.1	Invariants and eigenvalues of the anisotropy tensor	141
C.2	Barycentric map	142
D	Periodic hill parametrization for different geometries	145
E	SIP formulation	149
F	Additional LES cases	151
F.1	Extended D, RANS and HIRANS cases	151
F.2	Baseline, Extended A and B cases	154
F.3	Extended B, C and D cases	156
G	Kriging/Co-Kriging setup parameters	159
G.1	Bounds for hyperparameters	159
G.2	Exploration-exploitation bias ξ_{EI} per iteration	160
H	Snapping integration in Kriging/Co-Kriging predictor	163

List of Figures

1.1	Periodic hill geometry.	7
3.1	Left: grid used for the RANS case. Right: y^+ values for the RANS bottom (- - - -) and top (- - - -) walls and LES bottom (———) and top (- - - -) walls.	22
3.2	Streamwise velocity contours for RANS (left) and LES (right) solutions.	25
3.3	Differences of RANS streamwise (left) and vertical (right) velocity components with respect to the LES solution.	25
3.4	Left: Reattachment locations. Right: Wall shear stress distribution for RANS (- - - -) and LES (———).	26
3.5	Top: $\langle u'u' \rangle$ for RANS (left) and LES (right). Rest: Reynolds stresses differences with respect to LES values.	27
3.6	Contours of turbulent kinetic energy for RANS (left) and LES (right) solutions.	28
3.7	Baseline RANS case ($\xi_i = 0$). Left: filtered residuals of $\langle p \rangle$ (———), $\langle U_x \rangle$ (- - - -), $\langle U_y \rangle$ (- - - -), k (- - - -) and ω (- - - -). Right: Non-dimensional filtered lift (———) and drag (- - - -) coefficients gradient per iteration.	29
3.8	HIRANS cases for $\xi_i = 0.8$ (top), $\xi_i = 0.9$ (center) and $\xi_i = 1$ (bottom). Left: filtered residuals of $\langle p \rangle$ (———), $\langle U_x \rangle$ (- - - -), $\langle U_y \rangle$ (- - - -), k (- - - -) and ω (- - - -). Right: Non-dimensional filtered lift (———) and drag (- - - -) coefficients gradients per iteration.	31
3.9	Differences of HIRANS streamwise (left) and vertical (right) velocity components with respect to the LES solution.	32
3.10	Reynolds stresses differences with respect to LES values	33
3.11	Contours of turbulent kinetic energy for HIRANS (left) and LES (right) solutions.	34
3.12	Differences of HIRANS streamwise (left) and vertical (right) mean velocity components with respect to the LES solution.	35

3.13	Left: Reattachment locations. Right: Wall shear stress distribution for the bottom wall for LES (—), RANS (- - - - -), HIRANS with anisotropy injection (- · - · -) and HIRANS with full τ_{LES} injection (- - - - -).	36
3.14	Reynolds stresses differences with respect to LES values	37
3.15	Contours of turbulent kinetic energy for HIRANS (left) and LES (right) solutions.	38
3.16	Barycentric map for the anisotropy tensor.	39
3.17	Barycentric maps for $x/h = 3.97$ (left) and $x/h = 8.01$ (right). Eigenvalues for LES (\circ), RANS (\odot), HIRANS with LES anisotropy injection (\diamond) and HIRANS with τ_{LES} injection (\triangle).	39
3.18	Separation (left) and reattachment (right) locations for HIRANS with LES anisotropy injection (\triangle) and HIRANS with τ_{LES} injection (\diamond).	40
3.19	L^1 (······), L^2 (- - - - -) and L^∞ (- - - - -) norms of $\langle U_x \rangle$ (\circ), $\langle U_y \rangle$ (\odot), $\langle u'u' \rangle$ (\diamond), $\langle v'v' \rangle$ (\triangle), $\langle w'w' \rangle$ (\square) and $\langle u'v' \rangle$ (∇) of HIRANS with LES anisotropy injection (left) and HIRANS with τ_{LES} injection (right).	42
3.20	L^2 (left) and L^∞ (right) norms of $\langle U_x \rangle$ (top) and $\langle U_y \rangle$ (centre) and τ_w (bottom) for HIRANS with τ_{LES} injection. Contours linearly interpolated from 81 sample points (\times).	44
4.1	Mean velocity components at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), full (- - - - -), semi (- · - · -), and narrow (- · - · -) INCA cases.	50
4.2	Turbulent stresses at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), full (- - - - -), semi (- · - · -), and narrow (- · - · -) INCA cases (experimental data not available for $\langle w'w' \rangle$).	51
4.3	Mesh for fine (top) and coarse (bottom) cases.	52
4.4	Mean velocity components at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), fine (- - - - -), and coarse (- · - · -) INCA cases.	53
4.5	Turbulent stresses at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), fine (- - - - -), and coarse (- · - · -) INCA cases (experimental data not available for $\langle w'w' \rangle$).	55
4.6	Volumetric forcing (top) and root-mean-square velocity (bottom) for the Extended D simulation (left) and averaging period (right). The mean quantities over the averaging period are represented by (- - - - -).	57
4.7	FFT for the volumetric forcing (left) and root-mean-square velocity (right) for Extended D, flow-through-times 89-144 (—), Extended D, flow-through-times 22-55 (- - - - -), fine mesh, flow-through-times 22-55 (- · - · -). Reference slopes (······).	58
4.8	Mean and standard deviation of L^2 (- - - - -) and L^∞ (- - - - -) norms computed against PIV data over the different x/h stations of $\langle U_x \rangle$ (\circ), $\langle U_y \rangle$ (\odot), $\langle u'u' \rangle$ (\diamond), $\langle v'v' \rangle$ (\triangle) and $\langle u'v' \rangle$ (∇) for the different cases.	59
4.9	Mean velocity components at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), baseline (- - - - -), Extended A (- · - · -) and Extended D (- · - · -) INCA cases.	60

4.10	Turbulent stresses at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), baseline (- - - -), Extended A (- - - -) and Extended D (- - - -) INCA cases.	61
5.1	Geometry for different hill width values: baseline, $\psi = 1.0$ (—), steeper cases, $\psi = (0.25, 0.5, 0.75)$, (- - - -), and less steep cases, $\psi = (2.0, 3.0, 4.0)$, (- - - -).	64
5.2	Streamlines with streamwise velocity contours of LES cases for the steepest ($\psi = 0.25$, left) and least steep ($\psi = 4.0$, right) cases.	65
5.3	L^1 (.), L^2 (- - - -) and L^∞ (- - - -) norms for different geometries of $\langle U_x \rangle$ (\circ), $\langle U_y \rangle$ (\odot), $\langle u'u' \rangle$ (\diamond), $\langle v'v' \rangle$ (\triangle), $\langle w'w' \rangle$ (\square) and $\langle u'v' \rangle$ (∇) of RANS (left) and HIRANS with τ_{LES} injection (right).	66
5.4	Mean and standard deviation of L^2 (top) and L^∞ (bottom) norms for different geometries of $\langle U_x \rangle$ (\circ), $\langle U_y \rangle$ (\odot), $\langle u'u' \rangle$ (\diamond), $\langle v'v' \rangle$ (\triangle), $\langle w'w' \rangle$ (\square) and $\langle u'v' \rangle$ (∇) of interpolation with neighbouring (left) and all (right) cases.	70
5.5	Turbulent kinetic energy contours (top) and streamlines (bottom) for $\psi = 1.0$. Left: all cases, parameter interpolation. Right: LES solution.	71
5.6	L^2 (- - - -) norms for different geometries of $\langle U_x \rangle$ (\circ), $\langle U_y \rangle$ (\odot), $\langle u'u' \rangle$ (\diamond), $\langle v'v' \rangle$ (\triangle), $\langle w'w' \rangle$ (\square) and $\langle u'v' \rangle$ (∇) with respect to local injection (left) and pure RANS (right) values for neighbouring parameter interpolation.	72
6.1	Kriging and Co-Kriging flowchart.	79
6.2	Initial expensive (\circ) samples, Kriging $\hat{y} \pm 2\hat{s}$ (- - - -), Expected Improvement (—) and maximum EI location (\star) for the single fidelity methods for case A. Kriging LES (top), RANS (centre) and HIRANS (bottom).	81
6.3	Initial expensive (\bullet) and cheap (\odot) samples in expensive locations. Rest of cheap samples (\circ). Co-Kriging $\hat{y} \pm 2\hat{s}$ (- - - -), Expected Improvement (—) and maximum EI location (\star) for the multi-fidelity methods for Case A. Co-Kriging RANS (top) and HIRANS (bottom).	82
6.4	Initial expensive (\circ) samples, Kriging $\hat{y} \pm 2\hat{s}$ (- - - -), Expected Improvement (—) and maximum EI location (\star) for the single fidelity methods for Case B. Kriging LES (top), RANS (centre) and HIRANS (bottom).	83
6.5	Initial expensive (\bullet) and cheap (\odot) samples in expensive locations. Rest of cheap samples (\circ). Co-Kriging $\hat{y} \pm 2\hat{s}$ (- - - -), Expected Improvement (—) and maximum EI location (\star) for the multi-fidelity methods for case B. Co-Kriging RANS (top) and HIRANS (bottom).	85
6.6	LES samples for cases A (top) and B (bottom). Entire parameter space (left) and detail of the minima vicinities (right). Kriging LES (\circ), Kriging RANS (\odot), Kriging HIRANS (\diamond), Co-Kriging RANS (\triangle) and Co-Kriging HIRANS (\square). LES initial (\bullet) and additional samples (\circ). Visual guide using regressive Kriging LES regressive surrogate with $n_e = 59$ (- - - -).	87
6.7	Objective function improvement per LES evaluation for cases A (left) and B (right). Kriging LES (\circ), Kriging RANS (\odot), Kriging HIRANS (\diamond), Co-Kriging RANS (\triangle) and Co-Kriging HIRANS (\square). Exploring (- - - -) and exploiting (—) samples.	89
6.8	L^2 norms of surrogates with respect to regressive Kriging LES surrogate with $n_e = 59$ for cases A (left) and B (right). Initial sampling plan (\blacksquare) and end of optimization process (\boxtimes) values.	90

- 6.9 Top: LES k (\diamond) and f_p (\circ) samples with estimated 3σ regions. Bottom: Objective function for case A (\triangle) with estimated 3σ regions. Entire parameter space (left) and detail of the minimum vicinities (right). 92
- 6.10 Mean streamwise velocities (left) and turbulent kinetic energy (right) for baseline geometry ($\psi = 1.00$, top) and optimum geometries for cases A ($\psi = 0.759$, center) and B ($\psi = 1.69$, bottom). 93
- 6.11 Top: Initial expensive (\circ) and cheap (\circ) samples for case A, Co-Kriging HIRANS $\hat{y} \pm 2\hat{s}$ (----). Bottom: Expected Improvement and maximum EI locations with $\xi_{EI} = -0.5$ (———, \star), $\xi_{EI} = 0$ (———, \star) and $\xi_{EI} = 0.5$ (———, \star). 94
- 6.12 Top: Co-Kriging HIRANS $\hat{y} \pm 2\hat{s}$, expensive and cheap samples for 11 (———, \bullet , \diamond), 9 (———, \bullet , \circ) and 6 (———, \bullet , ∇) initial LES points. Bottom: Expected Improvement and maximum EI locations with $\xi_{EI} = -0.5$ (———, \star), $\xi_{EI} = 0$ (———, \star) and $\xi_{EI} = 0.5$ (———, \star). 96
- 6.13 Top: Co-Kriging HIRANS $\hat{y} \pm 2\hat{s}$, expensive and cheap samples for 48 (———, \diamond), 24 (———, \circ) and 12 (———, ∇) interior HIRANS points. LES samples (\bullet). Bottom: Expected Improvement and maximum EI locations with 48 (———, \star), 24 (———, \star) and 12 (———, \star) interior HIRANS points. 97
- 6.14 Top: Co-Kriging HIRANS $\hat{y} \pm 2\hat{s}$ for $\alpha = 1$ (\diamond), $\alpha = 0.75$ (\triangle), $\alpha = 0.5$ (\circ), $\alpha = 0.25$ (\square) and $\alpha = 0$ (∇). Bottom: Expected Improvement and maximum EI locations with $\alpha = 1$ (———, \star), $\alpha = 0.75$ (———, \star), $\alpha = 0.5$ (———, \star), $\alpha = 0.25$ (———, \star) and $\alpha = 0$ (———, \star). 99
- 6.15 Top: Initial expensive (\bullet) and cheap (\bullet) samples, Co-Kriging HIRANS $\hat{y} \pm 2\hat{s}$ for locally injected (———), non-locally injected (———) and corrected (———) models. Visual guide using regressive Kriging LES surrogate with $n_e = 59$ (- - - - -). Bottom: Expected Improvement and maximum EI locations for locally injected (———, \star), non-locally injected (———, \star) and corrected (———, \star) models. 102
- 6.16 Top: Kriging believer simulation for case A. Entire parameter space (left) and detail of the minima vicinities (right). Kriging believer with Co-Kriging HIRANS (\diamond) and Kriging LES (\blacksquare) sampling strategies. Kriging LES (\circ), Kriging RANS (\circ), Kriging HIRANS (\diamond), Co-Kriging RANS (\triangle) and Co-Kriging HIRANS (\square). Initial samples (\bullet) and additional samples (\circ). Visual guide using regressive Kriging LES surrogate with $n_e = 59$ (- - - - -). Bottom: Objective function improvement per LES evaluation for Kriging believer Co-Kriging HIRANS (left) and Kriging LES (right) sampling strategies. Exploring (- - - - -) and exploiting (———) samples. 104
- 6.17 Top: Kriging believer simulation for case B. Entire parameter space (left) and detail of the minima vicinities (right). Kriging believer with Co-Kriging HIRANS (\diamond) and Kriging LES (\blacksquare) sampling strategies. Kriging LES (\circ), Kriging RANS (\circ), Kriging HIRANS (\diamond), Co-Kriging RANS (\triangle) and Co-Kriging HIRANS (\square). Initial samples (\bullet) and additional samples (\circ). Visual guide using regressive Kriging LES surrogate with $n_e = 59$ (- - - - -). Bottom: Objective function improvement per LES evaluation for Kriging believer Co-Kriging HIRANS (left) and Kriging LES (right) sampling strategies. Exploring (- - - - -) and exploiting (———) samples. 105

6.18	L^2 norms of surrogates with respect to Kriging LES surrogate with $n_e = 59$ for cases A (left) and B (right). Initial sampling plan (■) and end of optimization process (⊠) values. The indices 1 and 2 are for Co-Kriging HIRANS and Kriging LES sampling strategies, respectively.	106
6.19	L^2 (-----) and L^∞ (-·-·-·-) norms for different geometries of $\langle U_x \rangle$ (○), $\langle U_y \rangle$ (⊙), $\langle u'u' \rangle$ (△), $\langle v'v' \rangle$ (◇), $\langle w'w' \rangle$ (□) and $\langle u'v' \rangle$ (▽) of RANS (left) and HIRANS with τ_{LES} injection (right).	107
6.20	Mean and standard deviation of flow variables norms with respect to LES values for RANS (■) and HIRANS (⊠).	108
6.21	Mean and standard deviation of objective function variables norms with respect to LES values for RANS (■) and HIRANS (⊠).	109
6.22	L^2 (.....) norms non-dimensionalized with RANS values for increasing distance to neighbouring LES samples of $\langle U_x \rangle$ (○), $\langle U_y \rangle$ (⊙), $\langle u'u' \rangle$ (△), $\langle v'v' \rangle$ (◇), $\langle w'w' \rangle$ (□) and $\langle u'v' \rangle$ (▽). $x_1 = 0.75/3.75$ (top-left), $x_1 = 1.0/3.75$ (top-right), $x_1 = 2.0/3.75$ (bottom-left) and $x_1 = 3.0/3.75$ (bottom-right).	110
6.23	L^∞ (.....) norms non-dimensionalized with RANS values for increasing distance to neighbouring LES samples of $\langle U_x \rangle$ (○), $\langle U_y \rangle$ (⊙), $\langle u'u' \rangle$ (△), $\langle v'v' \rangle$ (◇), $\langle w'w' \rangle$ (□) and $\langle u'v' \rangle$ (▽). $x_1 = 0.75/3.75$ (top-left), $x_1 = 1.0/3.75$ (top-right), $x_1 = 2.0/3.75$ (bottom-left) and $x_1 = 3.0/3.75$ (bottom-right).	111
A.1	L^2 (left) and L^∞ (right) norms of τ_w for the bottom wall for HIRANS with τ_{LES} injection. Contours linearly interpolated from 81 sample points (×).	133
A.2	L^1 (top), L^2 (center) and L^∞ (bottom) norms of $\langle U_x \rangle$ (left) and $\langle U_y \rangle$ (right) for HIRANS with τ_{LES} injection. Contours linearly interpolated from 81 sample points (×).	134
A.3	L^1 (top), L^2 (center) and L^∞ (bottom) norms of $\langle u'u' \rangle$ (left) and $\langle v'v' \rangle$ (right) for HIRANS with τ_{LES} injection. Contours linearly interpolated from 81 sample points (×).	135
A.4	L^1 (top), L^2 (center) and L^∞ (bottom) norms of $\langle w'w' \rangle$ (left) and $\langle u'v' \rangle$ (right) for HIRANS with τ_{LES} injection. Contours linearly interpolated from 81 sample points (×).	136
B.1	Filtered residuals for $t^* = 5000$ (left) and $t^* = 0$ (right) iterations, for $\xi_m, \xi_t = 0.85$ (top), $\xi_m, \xi_t = 0.90$ (centre) and $\xi_m, \xi_t = 0.95$ (bottom) of $\langle p \rangle$ (—), $\langle U_x \rangle$ (- - - - -), $\langle U_y \rangle$ (- · - · - · -), k (- · - · - · -) and ω (- - - - -).	138
B.2	L^1 (.....), L^2 (-----) (top) and L^∞ (-·-·-·-) (bottom) norms of $\langle U_x \rangle$ (○), $\langle U_y \rangle$ (⊙), $\langle u'u' \rangle$ (◇), $\langle v'v' \rangle$ (△), $\langle w'w' \rangle$ (□) and $\langle u'v' \rangle$ (▽) of HIRANS with $t^* = 5000$ (left) and $t^* = 0$ (right).	139
B.3	Differences of HIRANS streamwise (top) and vertical (centre) mean velocity components and turbulent kinetic energy (bottom) for $\{\xi_m, \xi_t\} = \{0.80, 0.80\}$, $t^* = 5000$ iterations (left) and $\{\xi_m, \xi_t\} = \{0.90, 0.85\}$, $t^* = 0$ iterations (right) with respect to the LES solution.	140
C.1	Barycentric map for the anisotropy tensor.	142

F.1	Mean velocity components at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), Extended D (- - - -), RANS (- · - · - ·), HIRANS with τ_{LES} injection from Breuer et al. (2009) (- · - · - ·) and INCA (- - - -) cases.	152
F.2	Turbulent stresses at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), Extended D (- - - -), RANS (- · - · - ·), HIRANS with τ_{LES} injection from Breuer et al. (2009) (- · - · - ·) and INCA (- - - -) cases.	153
F.3	Mean velocity components at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), INCA baseline (- - - -), Extended A (- · - · - ·) and Extended B (- · - · - ·) cases.	154
F.4	Turbulent stresses at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), INCA baseline (- - - -), Extended A (- · - · - ·) and Extended B (- · - · - ·) cases.	155
F.5	Mean velocity components at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), INCA Extended B (- - - -), Extended C (- · - · - ·) and Extended D (- · - · - ·) cases.	156
F.6	Turbulent stresses at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), INCA Extended B (- - - -), Extended C (- · - · - ·) and Extended D (- · - · - ·) cases.	157

List of Tables

3.1	Reattachment locations for RANS and LES	26
3.2	Reattachment locations for RANS, HIRANS, and LES	36
4.1	Grid properties of the different cases for the spanwise extent study.	49
4.2	Grid properties of the different cases for the grid resolution study.	52
4.3	Test matrix of averaging period cases. Both starting time and averaging period are in flow-through times ($t_{\text{flow-through}} = x_{\text{max}}/U_b$). Values truncated to integer numbers.	56
5.1	Values of the weighting coefficient for the “left” LES neighbour, α_{ψ_L} , for neighbouring interpolation. The right neighbour value is $\alpha_{\psi_R} = 1 - \alpha_{\psi_L}$	69
6.1	Number of acquired LES evaluations for the different methods with their corresponding exploring/exploiting decisions.	78
6.2	Improvements of the objective function for the different methods with respect to the baseline periodic hill geometry, J_{BSL} ($\psi = 1.0$) and minimum of the initial sampling plans, $\min(J)_{IS}$	88
D.1	Default values for the reference periodic hill geometry.	145
G.1	Acquisition strategies and ξ_{EI} values of for the first additionally acquired samples.	160
G.2	Acquisition strategies and ξ_{EI} values of for the second additionally acquired samples.	160
G.3	Acquisition strategies and ξ_{EI} values of for the third additionally acquired samples.	161
G.4	Acquisition strategies and ξ_{EI} values of for the fourth additionally acquired samples.	161
G.5	Acquisition strategies and ξ_{EI} values of for the fifth additionally acquired samples.	162
G.6	Acquisition strategies and ξ_{EI} values of for the sixth additionally acquired samples.	162

Nomenclature

Latin Symbols

a_{ij}	Reynolds-stress anisotropy tensor	$kg\ m^{-1}\ s^{-2}$
b_{ij}	Non-dimensional Reynolds-stress anisotropy tensor	-
\mathbf{C}	Augmented covariance matrix	-
d_z	Periodic hill spanwise extent	m
G	Filter function	-
f_p	Volumetric pressure-gradient forcing term	-
\hat{f}	Surrogate prediction	-
h	Periodic hill height	m
\mathbf{I}	Identity matrix	-
J	Objective function	-
k	Turbulent kinetic energy	m^2/s^2
Kn	Knudsen number	-
L^i	L- i norm	-
Ma	Mach number	-
p	Static pressure, Kriging hyperparameter	$kg\ m^{-1}\ s^{-2}$, -
P	Production of turbulent kinetic energy	$m^2\ s^{-3}$
Re	Reynolds number	-
Re_b	Reynolds number based on bulk velocity	-
\hat{s}	Predicted mean squared error	-
σ	Variance	-
S_{ij}	Rate of strain tensor	$kg\ m^{-1}\ s^{-2}$

t	Time	s
$t_{\text{flow-through}}$	Flow-through time	s
$\langle U_i \rangle$	Mean velocity component i	$m s^{-1}$
\bar{U}_i	Filtered velocity component i	$m s^{-1}$
$\langle u'_i u'_j \rangle$	Reynolds stress component ij	$m^2 s^{-2}$
u'	Residual velocity component	$m s^{-1}$
u_{RMS}	Root-mean-square velocity	$m s^{-1}$
w_i	Basis function weight i	i
x_i	Space,parameter space coordinate i	m,-
y	Objective function response	-
Z_c	Co-Kriging cheap Gaussian process	-
Z_d	Co-Kriging difference Gaussian process	-
Z_e	Co-Kriging expensive Gaussian process	-

Greek Symbols

α	Turbulence kinetic energy objective function multiplier	-
β	Pressure-gradient forcing term objective function multiplier	-
δ_{ij}	Kronecker delta	-
η	Non-dimensional distance	-
Γ	Massflow control geometric factor	-
λ	Regularization parameter	-
μ	Dynamic viscosity, mean	$kg m^{-1} s^{-1},-$
ν_t	Eddy viscosity	$m^2 s^{-1}$
ϕ	Probability density function	-
Φ	Cumulative distribution function	-
ψ	Periodic hill width parameter/basis function	-
Ψ	Gram matrix	-
τ_{ij}	Reynolds stress tensor	$kg m^{-1} s^{-2}$
ρ	Density/Co-Kriging scaling factor	-
τ^*	HIRANS hybrid Reynolds stress tensor	$kg m^{-1} s^{-2}$
τ_w	Wall shear stress	$kg m^{-1} s^{-2}$
θ	Kriging/Co-Kriging hyperparameter	-
ω	Specific rate of dissipation	-
ξ_{EI}	Exploration-exploitation bias	-
ξ_m	Momentum equation injection factor	-

ξ_t Turbulent transport equations injection factor -

Abbreviations

BSL	Baseline
DES	Detached Eddy Simulation
DMD	Dynamic Mode Decomposition
EI	Expected Improvement
FFT	Fast Fourier Transform
HIRANS	Hybrid Injected RANS
IS	Initial samples
LES	Large Eddy Simulation
PID	Proportional-integro-derivative
PIV	Particle Image Velocimetry
POD	Proper Orthogonal Decomposition
RANS	Reynolds-Averaged Navier-Stokes
SIP	Stone's Strongly Implicit Procedure
URANS	Unsteady Reynolds-Averaged Navier-Stokes

Other Symbols

∂	Partial derivative
D	Substantial derivative
$\nabla \cdot$	Divergence operation
∇	Gradient operation
\int	Integral
$\langle \rangle$	Mean operation

Chapter 1

Introduction

1.1 Motivation: design optimization in fluid mechanics

In the context of modelling complex turbulent flows, Large-Eddy Simulation (LES) has been widely successful at capturing features such as separation and recirculation in many scenarios where other models, such as Reynolds-Averaged Navier-Stokes (RANS), often find difficulties with. In industry nowadays, using a large number of LES simulations for a design optimization process is prohibitively expensive. On the other hand, RANS or even less computationally expensive methods show their limitations with tight optimization margins or in environments where their intrinsic assumptions introduce large errors.

To achieve the ultimately goal of a computationally efficient multi-fidelity design optimization process, several areas need to be considered, including the individual capabilities of each of the solvers used, how to calculate their sensitivity to design parameters and the techniques which allow to combine them.

To start with, the modelling capabilities of the different methods can greatly vary. Areas like massively separated flows or low aspect ratio flow features, away from the usually Thin Shear Flows trained regions of many turbulence models (Spalart et al., 1997), face great challenges to RANS, while LES has a broader range in terms of accuracy, which comes with its corresponding price tag.

In contrast, LES also has its own limitations: gradient-based optimization methods based on adjoint formulations can suffer from the butterfly effect, where starting from very similar initial conditions, the adjoint solution diverges when running backwards in time. Current successful results in overcoming this problem, such as (Talnikar et al., 2015b) often rely on stabilization techniques such as introducing artificial dissipation in some regions. Other approaches have been implemented, such as the Least Squares Shadowing, which “approximates the “shadow” trajectory in the phase space, avoiding the high sensitivity of the initial value problem” (Blonigan et al., 2016), but with a cost orders of magnitude higher than the actual flow simulation (Blonigan et al., 2016). This shadowing technique has also been used with the single step one-shot optimization method,

which couples solver, adjoint and optimization equations (Günther et al., 2017), previously developed in (Özkaya and Gauger, 2009) and applied to Unsteady RANS (URANS) in (Günther et al., 2016).

RANS can also be corrected in different ways. In the last years, several data processing models, such as active learning with Machine Learning to produce augmented RANS models (Duraisamy, 2015; Duraisamy et al., 2017), have been developed. This method constructs a new field variable, which can be introduced either as a source term or as a multiplier of some physical term, such as the dissipation part of the turbulent transport equations. This variable is learnt by Machine Learning formulating an inverse problem based on high fidelity data. Some of its caveats are exposed in Duraisamy et al. (2017), where it is strongly suggested that one important factor to obtain accurate predictions is not the singular values of the different terms, but obtaining the correct balance between them. Poroseva et al. (2016) commented this issue, yielding large discrepancies even with very small errors from second moment closure terms extracted from DNS data. Other methods include nonlinear autoregressive models with exogenous input (Pham et al., 2010) or RANS calibration and uncertainty analysis by Bayesian estimation (Cheung et al., 2011; Edeling et al., 2014a,b; Guillas et al., 2014). These techniques often have two main drawbacks: the computational resources required to correct the RANS method can be of the order of magnitude of actually solving additional LES cases, and their prediction capabilities are usually clustered in regions close to the local corrections, which are often far away from a multi-parameter optimization process requirements.

Several attempts to combine RANS and LES by splitting the flowfield have also been performed, such as in Detached-Eddy Simulation (DES) (Spalart et al., 1997; Spalart, 2009), where a criterion decides which regions are solved with each model. These are not exempt of challenges too, specially concerning the transition regions between the two models, as well as a consistent treatment of the two regions during the whole optimization envelope. An alternative method to combine both solvers was proposed in (Marsden et al., 2007). In the context of airfoil trailing edge aeroacoustics with load constraints, the capabilities of RANS were used to firstly calculate the lift and drag, so only potential candidates which meet these constraints were promoted to the aeroacoustics calculations. These were performed splitting the domain into a LES region around the trailing edge, and RANS for the rest of the domain.

The previous comments have been made about the accuracy of the solver in a *local context*, but during a possibly non-convex optimization routine with multiple local minima, the acquisition of new points also needs to be considered. Surrogate models based on physical characteristics can be constructed, such as Proper Orthogonal Decomposition, POD (Berkooz et al., 1993), which uses a kinetic energy criterion to identify the key modes relevant to the flow. Extensions to it have been formulated, such as hyper reduction (Amsallem et al., 2015, 2012; Carlberg et al., 2011). This technique constructs approximations that reduce the computational cost of the surrogate construction. Dynamic Mode Decomposition, DMD (Schmid, 2010) aims to find the modes not necessarily with the greatest amount of kinetic energy, but most relevant in the dynamics of the flow instead. Interpolation schemes can then be used to cover the whole studied envelope, nevertheless these models have a similar drawback than calibration techniques, which is their accuracy and response to non-linearities in the flow dynamics far away from the already sampled points. Apart from well-known methods such as radial basis functions,

alternative formulations can be found in literature, such as nonlinear manifold learning, Isomap (Tenenbaum et al., 2000). Isomap aims to “preserve the *geodesic* distances between all pairs of points,” (Franz et al., 2014) which can be significantly different from Euclidean distances, allowing for a more appropriate choice of neighbours than in traditional interpolation techniques.

Furthermore, a related branch to improve this dichotomy resources-accuracy involves the data-driven surrogates based on Gaussian processes (O’Hagan and Kingman, 1978), such as Bayesian optimization (Kushner, 1964) and Kriging methods (Krige, 1951; Matheron, 1962; Sasena, 2002), which have been utilized to construct response surfaces based on flow or objective function quantities. Moreover, they have been also combined with POD-based variants to predict behaviours outside of the snapshots points used for the POD model (Xiao et al., 2010; Braconnier et al., 2011). These methods rely exclusively on the available data without taking consideration of the underlying physics of the problem, and can be cheaply updated every time a new sample point is acquired to increase their accuracy. Their main advantages are their ability to capture non-linearities in the objective function response outside of the known locations and the cheap cost of choosing new acquisition points. Gradient information can be introduced, as in Gradient-Enhanced Kriging (Fischer and Grandhi, 2015), they can be parallelized to acquire multiple samples (Talnikar et al., 2015a; Snoek et al., 2012) and tolerance to noise and the addition of inequality constraints can be added (Gardner et al., 2014). These surrogates are often found in global optimization processes, combined with criteria that combine exploration and exploitation (searching in the region of the best current point or exploring new ones), such as Expected Improvement, EI (Talnikar et al., 2015a).

Bayesian optimization has been successfully applied to fluid mechanics problems, as in turbulent channel drag reduction (Talnikar et al., 2015a) and transonic turbine blade shape design (Talnikar et al., 2014), where downstream pressure loss was computed for 35 LES samples. The advantages and drawbacks of this method are reflected in this reference: while Bayesian optimization, by itself, might be able to efficiently acquire the region of the maximum of the objective function, it still required pure LES data. Extending the problem to a larger amount of design parameters would likely make it infeasible in terms of computational resources.

A key extension of Kriging for the present document was presented in Forrester et al. (2007, 2008): *Co-Kriging*, which is able to compute different layers of fidelity, constructing independent surrogate methods for each of them and incorporating the cheap dataset information into the expensive one, for both translation and scaling. It has been already applied for different levels of fidelity of RANS in Brooks et al. (2011) for a transonic compressor rotor design.

The current scenario manifests that the gap between RANS and LES in global design optimization for problems such as massively separated flows has not been bridged yet. This is harnessing the potential of multi-fidelity frameworks such as Co-Kriging, where the large differences between the two methods limits the usefulness of the information that can be extracted from the cheap dataset. Acknowledging this situation, there seems to be room for the formulation of new models focused in the local context of increasing the computational efficiency within a design envelope. The aim is a qualitatively similar at its best, not misleading at its worst, accuracy with respect to LES, while keeping a much more contained cost.

1.2 Research objectives

The main goal of this work can be stated in the proposed research question:

Does the combination of a hybrid LES-data-injected RANS solver with pure LES in a Co-Kriging multi-fidelity global design optimization process outperform a single fidelity method for a fixed computational budget in the context of turbulent incompressible fluid flows for the specified periodic hill test cases?

Based on it, secondary questions arise:

- (a) *How reliably and consistently does a hybrid RANS method match the local LES model in the mean flow quantities?*
- (b) *What is the range in the parameter space (width of the periodic hill) where locally corrected RANS methods still resemble the time-averaged LES dynamics?*
- (c) *Which interpolation techniques provide a more robust set of tools within the computational budget? How well can they resemble the higher fidelity features in the hybrid model within the whole design envelope?*
- (d) *What is the optimum amount of resources to be spent in higher and lower fidelity computations to outperform a single-fidelity method? How is this balance linked with the accuracy of the corrections of the hybrid model?*
- (e) *Could this method be generalized to other problems, such as multi-parameter optimization for the periodic hill test case?*

These five sub-questions are related to the five main parts of the project, and provide a framework to work on the proposal, testing and application of the hybrid model in the optimization process (first, third and fourth question) and on the understanding of limitation of boundaries of the studied formulation (second and fifth ones) on the different stages of the project.

The objectives to be met based on these questions are:

- (a) Introduce local corrections injecting LES information into the base RANS model for the periodic hill case, and analyse their effect in terms of the main flow features, such as velocity components, turbulent stresses and wall shear stress. Achieve a model that is able to follow a similar trend as the time-averaged LES for the different sample points within a specified range.
- (b) Analyse the effect of LES corrections in new sample points for RANS, understanding if the proposed model is suitable for the different conditions and correcting it if necessary.
- (c) Introduce interpolation schemes to introduce the sparse LES data into multiple RANS cases, and combine this with the information of local corrections to determine the minimum amount of LES cases needed.

- (d) Compute a Co-Kriging multi-fidelity design optimization routine based on the points stated above, analysing different levels of the RANS-LES budget balance to outperform a single-fidelity global optimization method.

1.3 Contribution of this work

As it has been exposed, there are several limitations in the current methods for design global optimization in the case of computationally expensive fluid flows. There is always a compromise between the accuracy of the solution, the spent resources, and how well can the model be generalized to other problems.

This document proposes a novel framework, which injects turbulent stresses from a high-fidelity solution, such as LES, into a lower-fidelity model (RANS). The aim is to correct it, constructing in practice a hybrid model, HIRANS (Hybrid Injected Reynolds-Averaged Navier-Stokes). This allows to provide the necessary corrections to improve its accuracy and its predictions capabilities beyond its original boundaries, while keeping the same order of magnitude of cost as the additional RANS model. To the knowledge of the author, no similar approaches can be found in literature in a similar context at the time of the publication.

These corrections are based on the physical limitations of the RANS turbulence model used, the OpenFOAM implementation of the 2-equation $k - \omega$ (Wilcox, 1998), which is further described in Section 3.1.2. Due to the usage of Boussinesq's hypothesis, where the anisotropic part of the Reynolds stress tensor is modelled in an analogous way as the viscous stress of a Newtonian fluid,

$$\langle u'_i u'_j \rangle - \frac{2}{3} k \delta_{ij} = -\nu_t \left(\frac{\partial \langle U_i \rangle}{\partial x_j} + \frac{\partial \langle U_j \rangle}{\partial x_i} \right), \quad (1.1)$$

the limitation of assuming a linear relation between the anisotropy and the mean velocity gradients with the scaling of the turbulent viscosity highly constrains the flow in several situations.

Understanding the importance of the anisotropy in the calculation of the Reynolds stresses, information of the non-dimensional anisotropy tensor, as well as the full Reynolds stress tensor from LES, can be injected, in the momentum and turbulent model equations. In this document, different injection methods are proposed and evaluated to provide the required numerical stability and accuracy level. In fact, the corrections where the Reynolds stress tensor from LES is introduced in the momentum equation and the non-dimensional anisotropy tensor b_{ij} is injected in the production term for the turbulent kinetic energy transport equation provides much closer results to the time-averaged LES results than the baseline RANS model.

The corrections over the whole envelope of the optimization routine are performed using sparse LES cases and interpolating between them, using different methods based in parameter, geometry and flow quantities.

The availability of a hybrid model whose predictions are much closer to its higher-fidelity counterpart is a core factor in the usage of a multi-fidelity global optimization routine.

The technique used in this document is Co-Kriging (Forrester et al., 2008), which allows for a seamless introduction of a higher and a lower-fidelity model, using the information of the lower one, with the assumption that the auto-regressive model which includes both solutions has the form

$$Z_e(\mathbf{x}) = \rho Z_c(\mathbf{x}) + Z_d(\mathbf{x}), \quad (1.2)$$

where the expensive results are approximated using a constant scaling factor ρ and a difference Z_d using a Gaussian process between the scaled cheap code ρZ_c and the expensive one Z_e .

To validate this novel framework, the test case chosen is the widely known periodic hill (Figure 1.1), following the geometry definition from Mellen et al. (2000), which retains the original hill shape of Almeida et al. (1993) while doubling the distance between hills, halving the height of the channel and removing the side walls and assuming periodic boundary conditions in the spanwise and streamwise directions. Moreover, the Reynolds number at the crest of the hill is $Re = 10595$, based on the hill height and the bulk velocity U_b at the crest of the hill. The channel is comprised of a bottom wall where the flow separates, creating a recirculation region, and a top wall where a boundary layer is formed. Reference LES (Breuer et al., 2009) and Particle Image Velocimetry, PIV (Rapp, 2009) data is available for this case at $Re = 10595$, and additional LES solutions can be computed at a reasonable cost. The design parameter corresponds to a hill width multiplier in the streamwise direction. The suitability of this case is due to being a 2D geometry simple to parametrize, where RANS solvers usually struggle with, and also hybrid solvers, such as DES. In previous studies, as in Šarić et al. (2007), the limitations of a DES model with the same grid as a LES one are argued, observing that at the Reynolds number considered the differences are still notable with respect to a pure LES case. In Davidson and Peng (2003), a hybrid model where RANS is used in the near-wall region, with the switch to LES in the logarithmic part of the boundary layer, struggles in the feeding of RANS data into the LES region due to the poor spectral properties of this information, but achieves acceptable mean velocity profiles. Moreover, the hybrid URANS-LES of Temmerman et al. (2005), at a Reynolds number based on the bulk velocity $Re_b = 21560$ predicts the reattachment point around 0.7-1.1 hill heights downstream of the reference data.




Figure 1.1: Periodic hill geometry.

The previous discussion about the contributions of this work can be summarized in two main arguments. On the one hand, a multi-fidelity Co-Kriging with LES and HIRANS information is proposed and analysed. The goal is *not* to obtain the most accurate representation of the local flow features, which could be achieved by other techniques, such as Bayesian estimation of the coefficients of the turbulence model or diverse Machine Learning processes. Instead, the aim is to construct a method that *qualitatively* captures these features and remains robust and consistent in its predictions within a parameter space, comprised of a significantly larger amount of cheap samples compared to the expensive ones to inject data from. The second of the contributions of this document is a valuable database of several costly LES evaluations for different geometries. This information can be used for future research of any novel or existing method that requires a construction or validation set of this fidelity level.

1.4 Outline

This document is structured in four main chapters, which are directly linked with the four research sub-questions. Firstly, a general overview of the numerical schemes and computational setup is given in Chapter 2. After that, Chapter 3 tackles the first challenge of this thesis: how to implement corrections over the baseline $k - \omega$ model using injected LES information, constructing the novel scheme Hybrid Injected RANS. The tests in the periodic hill case are shown, and the two different HIRANS formulations are analysed. The next part of the multi-fidelity framework is portrayed in Chapter 4, where a reliable and computationally feasible LES setup is explored.

With HIRANS and LES configurations already established, Chapter 5 explores the effect of the HIRANS *prediction* when using interpolated LES injection from non-local geometries. This information is used to construct for single and multi-fidelity optimization methods in Chapter 6. A final discussion of the results and recommendations for future work in Chapter 7.

Numerical schemes and computational setup

2.1 Introduction

In this chapter, the main equations and numerical setup for the RANS, LES and Kriging/Co-Kriging models are derived and presented. Additional formulations, such as the different HIRANS models, are presented in subsequent chapters as they are employed along the document.

2.2 RANS methodology

2.2.1 Derivation of RANS equations

In order to derive the RANS equations, the different variables are decomposed into mean and fluctuating components (Reynolds decomposition). For example, for the velocity,

$$U(x, t) = \langle U(x, t) \rangle + u'(x, t). \quad (2.1)$$

Starting for the continuity equation for incompressible flow,

$$\frac{\partial U_i}{\partial x_i} = \frac{\partial}{\partial x_i} (\langle U_i \rangle + u'_i) = 0. \quad (2.2)$$

It is then observed that both mean and fluctuating component are solenoidal, as taking the mean of (2.2) leads to $\partial \langle U_i \rangle / \partial x_i = 0$, and subtracting this component the result $\partial u'_i / \partial x_i = 0$ is obtained.

The momentum equation includes a nonlinear convective term, which is slightly more elaborate to treat. Starting from the conservative form for the substantial derivative,

$$\frac{DU_j}{Dt} = \frac{\partial U_j}{\partial t} + \frac{\partial}{\partial x_i}(U_i U_j). \quad (2.3)$$

Taking the mean of this equation,

$$\left\langle \frac{DU_j}{Dt} \right\rangle = \frac{\partial \langle U_j \rangle}{\partial t} + \frac{\partial}{\partial x_i} \langle U_i U_j \rangle. \quad (2.4)$$

If the Reynolds decomposition is applied for the nonlinear component, and applying the fact that the mean of a single fluctuation component is null,

$$\begin{aligned} \langle U_i U_j \rangle &= \langle (\langle U_i \rangle + u'_i)(\langle U_j \rangle + u'_j) \rangle \\ &= \langle U_i U_j \rangle + \langle u'_i u'_j \rangle. \end{aligned} \quad (2.5)$$

The last term, which corresponds to the velocity covariances $\langle u'_i u'_j \rangle$, is called *Reynolds stresses*.

Applying this result to the incompressible momentum equation results in

$$\rho \langle U_j \rangle \frac{\partial \langle U_i \rangle}{\partial x_j} = \frac{\partial}{\partial x_j} \left[- \langle p \rangle \delta_{ij} + \mu \left(\frac{\partial \langle U_i \rangle}{\partial x_j} + \frac{\partial \langle U_j \rangle}{\partial x_i} \right) - \rho \langle u'_i u'_j \rangle \right], \quad (2.6)$$

which is similar to the general form of the momentum conservation equation, but with an additional apparent stress, which comes from the fluctuations of the velocity field in the form of the Reynolds stresses, which transfers momentum using the mechanism of this fluctuating velocity field (Pope, 2000). The modelling of this tensor is crucial in capturing the physics underlying it, and several models have been developed along the years. Simple mixing length assumptions for the eddy viscosity ν_t from Boussinesq's hypothesis (1.1) have been formulated, such as Smagorinsky (Smagorinsky, 1963), Cebeci-Smith (Smith and Cebeci, 1967) and Baldwin-Lomax (Baldwin and Lomax, 1978). An alternative definition of the eddy viscosity in the popular Spalart-Allmaras model (Spalart and Allmaras, 1992), and two-equation turbulent-kinetic-energy related models have been widely adopted, such as $k-\varepsilon$ (Jones and Launder, 1972; Launder and Sharma, 1974) $k-\omega$ (Wilcox, 1988) and SST $k-\omega$ (Menter, 1993). More complex and costly Reynolds stress models, with individual equations for each of the stress components, can be found in Launder et al. (1975) and Speziale et al. (1991).

2.2.2 OpenFOAM $k-\omega$

The open source software OpenFOAM (Weller et al., 1998, 2018) is used for the implementation of the turbulence model.

Starting from the incompressible RANS momentum equation

$$\rho \langle U_j \rangle \frac{\partial \langle U_i \rangle}{\partial x_j} = \frac{\partial}{\partial x_j} \left[- \langle p \rangle \delta_{ij} + \mu \left(\frac{\partial \langle U_i \rangle}{\partial x_j} + \frac{\partial \langle U_j \rangle}{\partial x_i} \right) - \rho \langle u'_i u'_j \rangle \right]. \quad (2.7)$$

The turbulence model used is the OpenFOAM implementation of the $k - \omega$ model, which is based on the 1988 Wilcox version (Wilcox, 1998), whose transport equations for k and ω are

$$\rho \frac{\partial k}{\partial t} + \rho \langle U_j \rangle \frac{\partial k}{\partial x_j} = \tau_{ij} \frac{\partial \langle U_i \rangle}{\partial x_j} - \beta^* \rho k \omega + \frac{\partial}{\partial x_j} \left[(\mu + \sigma^* \mu_T) \frac{\partial k}{\partial x_j} \right], \quad (2.8a)$$

$$\rho \frac{\partial \omega}{\partial t} + \rho \langle U_j \rangle \frac{\partial \omega}{\partial x_j} = \alpha \frac{\omega}{k} \tau_{ij} \frac{\partial \langle U_i \rangle}{\partial x_j} - \beta \omega^2 + \frac{\partial}{\partial x_j} \left[(\mu + \sigma \mu_T) \frac{\partial \omega}{\partial x_j} \right], \quad (2.8b)$$

with the closure coefficients

$$\alpha = 0.52, \quad \beta = 0.072, \quad \beta^* = 9/100, \quad \sigma = 1/2, \quad \sigma^* = 1/2, \quad (2.9)$$

and the eddy viscosity computed as $\mu_T = \rho k / \omega$.

2.2.3 Missing term in OpenFOAM

Apart from this, a remark must be made about the OpenFOAM model: in this eddy viscosity model, Boussinesq's hypothesis is used in order to provide the anisotropic part of the Reynolds stresses:

$$\langle u'_i u'_j \rangle - \frac{2}{3} k \delta_{ij} = -\nu_t \left(\frac{\partial \langle U_i \rangle}{\partial x_j} + \frac{\partial \langle U_j \rangle}{\partial x_i} \right) \quad (2.10)$$

Nevertheless, the implementation in OpenFOAM lacks in both, momentum and k transport equation, the second term in the left hand side of (2.10), which effectively means that there is not a linear relation between the anisotropic stresses and the mean rate of strain, but with the *full* Reynolds stress tensor instead. In the momentum equation, the term corresponding to k can be absorbed as a modified pressure, as $p^* = p + 2/3k$. Nevertheless, in the k transport equation, this term, which appears in the production as $P = \tau_{ij} (\partial \langle U_i \rangle / \partial x_j)$, cannot be absorbed in such a trivial way, and the difference between the standard Boussinesq's hypothesis and the OpenFOAM implementation is:

$$P_{Boussinesq} = \left(\langle u'_i u'_j \rangle - \frac{2}{3} k \delta_{ij} \right) \frac{\partial \langle U_i \rangle}{\partial x_j} \quad (2.11a)$$

$$P_{OpenFOAM} = \langle u'_i u'_j \rangle \frac{\partial \langle U_i \rangle}{\partial x_j} \quad (2.11b)$$

For this document, it has been decided to add the isotropic turbulent contribution in the k and ω equations on top of the OpenFOAM standard $k - \omega$ model, while leaving the momentum equation untouched, which means that the pressure value obtained will effectively be $p^* = p + 2/3k$.

The effect of this missing term is further assessed in Section 2.2.3, where it is shown that its influence is not negligible, specially in regions close to walls.

2.2.4 Numerical discretization

In order to discretize the continuity, momentum and $k - \omega$ transport equations, the open source software OpenFOAM has been used. This framework uses a finite volume method (FVM) to evaluate the set of partial differential equations in the discrete plane on a given mesh. The domain is divided into an arbitrary number of control volumes (cells), where the different variables to be computed are located at the centroids of these cells. The equations to be discretized are used in the integral formulation, applying it to each control volume. To describe variations between different cell centroids (like gradients), interpolation techniques are used. Volume integrals with divergence terms are converted to surface integrals using the Gauss' divergence theorem:

$$\int_V (\nabla \cdot \mathbf{F}) dV = \int_S (\mathbf{F} \cdot \mathbf{n}) dS \quad (2.12)$$

The integration of these quantities in the discrete space is performed in the numerical setup OpenFOAM, in general, using the Gauss quadrature rule, which, for a function g and a weighting function W , uses the approximation (Moroney (2006)):

$$\int_a^b g(u)W(u)du \approx \sum_{j=1}^p w_j g(u_j) \quad (2.13)$$

The sum of these quantities is performed with the values of the face centres, which need to be interpolated from the calculated values at the cell centres. From the variety of interpolation schemes available, linear interpolation has been applied by default, using upwind schemes for the discretisation of the velocity gradients in the terms $\nabla \cdot (\phi U)$, $\nabla \cdot (\phi k)$, $\nabla \cdot (\phi \omega)$, where ϕ represents the velocity flux along a cell face.

After applying the discretization for the different quantities, the different equations need to be solved. The continuity and momentum equation can be used to solve velocity components and pressure. OpenFOAM stores all the primitive variables in the cell centres, but as explained before,

Inherent to this formulation is the conservation of these numerical fluxes, which leads to the conservation of quantities (mass, momentum, etc.), which is exactly satisfied for an arbitrary number of control volumes, for any grid and for any part of the domain. The flexibility of FVM approaches to create different structured and unstructured grid shapes make them also suitable for complex geometries.

2.3 LES methodology

2.3.1 Derivation of LES equations

LES (Smagorinsky, 1963; Lilly, 1966; Deardorff, 1974) was first developed for meteorological applications, and it has been widely used to perform high fidelity simulations of fluid flows since then. In this formulation, the larger 3D unsteady turbulent motions are solved, while the smaller scales use different models to approach reality. Because of it, its computational cost is between RANS and direct numerical simulation (DNS). Using Pope (2000) structure, the four steps in LES can be stated as follows:

- (a) The velocity field $U(x, t)$ is decomposed into a filtered component $\bar{U}(x, t)$ and a residual, or subgrid-scale (SGS) component, $u'(x, t)$. By filtering this field, it represents the features of the large scales, or eddies.
- (b) The equations for the filtered velocity field can be derived from the Navier-Stokes, showing a standard form with an additional residual-stress tensor (SGS stress tensor) from the residual motions.
- (c) The modellisation of this tensor leads to closure of the equations.
- (d) The equations can be solved for the filtered velocity, providing an approximation in one realization to the turbulent large scale motions of the flow.

The low-pass filtering (Leonard, 1975) can be defined as

$$\bar{U}(x, t) = \int G(r, x)U(x - r, t)dr, \quad (2.14)$$

with an integration over the whole domain, and a filter function G which satisfies

$$\int G(r, x)dr = 1. \quad (2.15)$$

Many different filter functions have been tested in literature, with different properties, such as box, Gaussian or sharp spectral filters. The filtered components will differ depending on a filter width Δ . For example, for a one-dimensional problem, the sharp spectral filter function can be defined as $G = \sin(\pi r/\Delta)/\pi r$, which gives a cut-off wavenumber $\kappa_c = \pi/\Delta$. This means that in wavenumber space, this filter will cut all Fourier modes which have a wavenumber $|\kappa|$ greater than the cut-off value, producing a filtered energy spectrum.

The velocity field, using the filtered and residual components, is finally decomposed as

$$U(x, t) = \bar{U}(x, t) + u'(x, t). \quad (2.16)$$

While it looks very similar to the Reynolds decomposition applied for the RANS equations, in this case, $\bar{U}(x, t)$ is a random field, and the filtered residual is non-zero in general,

$$\overline{u'} \neq 0. \quad (2.17)$$

2.3.2 LES solver: INCA

The LES simulations are run using the academic solver INCA. This finite volume approach for the filtered Navier-Stokes equation uses a third order Runge-Kutta scheme for time integration, with a second order central difference scheme for the discretization of the convective and diffusive fluxes. The second-order Poisson equation for the pressure can be solved employing a wide variety of techniques, including Krylov subspace and multigrid methods. An immersed boundaries method (Peskin, 1972) with first-order boundary reconstruction is used.

The discretization is performed using implicit LES, where the subgrid-scale model and numerical discretization are fully merged. An adaptive local deconvolution method (ALDM, Hickel et al. (2006)) is implemented in INCA, providing “a truncation error which acts as a subgrid-scale model, consistent with asymptotic turbulence theory” (Hickel and Adams, 2007). The wall modelling of the solver can be found in this last reference, comprising a van Driest damping function “to compensate for increased SGS dissipation by reducing the dissipative weight of the numerical flux function” (Hickel and Adams, 2007).

2.4 Kriging and Co-Kriging

As introduced in Section 1.1, efficient gradient-based methods such as adjoint have not been widely developed for LES. Moreover, there is no guarantee that the global optimum will be found if the objective function is non-convex, or that the optimizer will not get trapped in large “valley” regions in this space. The lack of knowledge about the behaviour of the objective function in this study adds another reason to consider global optimization instead. Nevertheless, global optimization usually comes with a large number of samples needed to explore the parameter space. Techniques such as Genetic Algorithms (Goldberg, 1989) require several generations for selection, crossover and mutation operations, which rapidly becomes too costly for the problem considered in this document.

On the other hand, a surrogate model for the objective function can be constructed. Using this information, it can be *predicted* how it will behave across the parameter space, and apply a selection criterion to explore and exploit this surrogate model. The formulation chosen in this work is the Kriging/Co-Kriging models, with the Expected/Co-Expected Improvement criterion. A comprehensive discussion of the different equations can be found in Forrester et al. (2008).

2.4.1 Surrogate model: Kriging

To start the optimization process, the design variables vector $\mathbf{x} = \{x_1, x_2, \dots, x_k\}^T$ can be constructed. This vector leads to the objective function values y , forming the data pairs $\{(\mathbf{x}^{(1)}, y^{(1)}), (\mathbf{x}^{(2)}, y^{(2)}), \dots, (\mathbf{x}^{(n)}, y^{(n)})\}$ for n data samples. To construct a surrogate model, from the wide range of functions available in literature, interpolating Radial Basis Functions (RBF) are considered.

Starting from the sampling plan $\mathbf{X} = \{x^{(1)}, x^{(2)}, \dots, x^{(n)}\}$, with responses $\mathbf{y} = \{y^{(1)}, y^{(2)}, \dots, y^{(n)}\}$, the radial basis function approximation $\hat{f}(\mathbf{x})$ has the form

$$\hat{f}(\mathbf{x}) = \mathbf{w}^T \boldsymbol{\psi} = \sum_{i=1}^{n_c} w_i \psi(\|\mathbf{x} - \mathbf{c}^{(i)}\|) \quad (2.18)$$

being $c^{(i)}$ the centre of the i th of the n_c basis function centres, and $\boldsymbol{\psi}$ the n_c -vector with the values of the basis functions ψ , using Euclidean distances between the prediction location \mathbf{x} and the centres $\mathbf{c}^{(i)}$ of the basis functions. There are several choices of the basis functions themselves, such as linear, $\psi(r) = r$; cubic, $\psi(r) = r^3$; or Gaussian, $\psi(r) = e^{-r^2/(2\sigma^2)}$.

Whichever the choice of the basis functions, \mathbf{x} can be estimated easily, using the interpolation condition

$$\hat{f}(\mathbf{x}^{(j)}) = \sum_{i=1}^{n_c} w_i \psi(\|\mathbf{x} - \mathbf{c}^{(i)}\|) = y^{(j)}, \quad j = 1, \dots, n. \quad (2.19)$$

The linear equation 2.19 allows to express nonlinear responses. With $n_c = n$, this leads to

$$\boldsymbol{\Psi} \mathbf{w} = \mathbf{y}, \quad (2.20)$$

being $\boldsymbol{\Psi}$ the Gram matrix, which can be established as $\Psi_{i,j} = \psi(\|\mathbf{x}^{(i)} - \mathbf{x}^{(j)}\|)$, $i, j = 1, \dots, n$. Solving this equation for \mathbf{w} allows to determine the different weighting parameters.

If noisy data is expected, this model may be endangered by *overfitting*. In that case, a *regularization parameter* λ (Hoerl and Kennard, 1970; Tikhonov and Arsenin, 1977; Poggio and Girosi, 1990) can be added to the diagonal of the Gram matrix, as

$$\mathbf{w} = (\boldsymbol{\Psi} + \lambda \mathbf{I})^{-1} \mathbf{y}, \quad (2.21)$$

where \mathbf{I} is the identity matrix. This regularization parameter is also labeled as *nugget* in other references, such as in Sasena (2002), where it is applied to all elements of the correlation vector of a new point with the existing samples. It is observed that the model starts drift away from being purely *interpolating* and becomes more *regressive* as this parameter λ is increased. Very close sample points can produce ill-conditioning of the Gram matrix in the computations, which also encourages the addition of this parameter in a very small quantity ε in several cases.

The main choice for the basis function in this work is the *Kriging* method (Krige, 1951; Sacks et al., 1989), which is defined as

$$\psi^{(i)} = \exp \left(- \sum_{j=1}^k \theta_j |x_j^{(i)} - x_j|^{p_j} \right). \quad (2.22)$$

This defines a vector of hyperparameters $\boldsymbol{\theta} = \{\theta_1, \theta_2, \dots, \theta_k\}^T$, with an additional varying exponent $\mathbf{p}_j = \{p_1, p_2, \dots, p_k\}^T$ for each parameter dimension, with typical values of $p_j \in [1, 2]$. If \mathbf{p} is fixed to 2 and a constant θ_j is used for all dimensions, the Gaussian basis function is obtained.

Using the assumption that the already sampled responses are coming from a stochastic process, and using the set of random vectors

$$\mathbf{Y} = \begin{pmatrix} Y(\mathbf{x}^{(1)}) \\ \vdots \\ Y(\mathbf{x}^{(n)}) \end{pmatrix} \quad (2.23)$$

the random field has a mean of $\mathbf{1}\mu$, being $\mathbf{1}$ a column vector of n ones, and the random variables are correlated as

$$\text{cor}[Y(\mathbf{x}^{(i)}), Y(\mathbf{x}^{(l)})] = \exp\left(-\sum_{j=1}^k \theta_j |x_j^{(i)} - x_j^{(l)}|^{p_j}\right), \quad (2.24)$$

which allows to construct the $n \times n$ correlation matrix of the observed data

$$\boldsymbol{\Psi} = \begin{pmatrix} \text{cor}[Y(\mathbf{x}^{(1)}), Y(\mathbf{x}^{(1)})] & \cdots & \text{cor}[Y(\mathbf{x}^{(1)}), Y(\mathbf{x}^{(n)})] \\ \vdots & \ddots & \vdots \\ \text{cor}[Y(\mathbf{x}^{(n)}), Y(\mathbf{x}^{(1)})] & \cdots & \text{cor}[Y(\mathbf{x}^{(n)}), Y(\mathbf{x}^{(n)})] \end{pmatrix} \quad (2.25)$$

and a covariance matrix

$$\text{Cov}(\mathbf{Y}, \mathbf{Y}) = \sigma^2 \boldsymbol{\Psi}. \quad (2.26)$$

This assumes that the function is smooth and continuous. A different value of a hyperparameter θ_j will affect the correlation, essentially the “width” of its influence away from a sample point. Also, p_j allows for a smooth correlation for $p_j = 2$ near a sample point, which becomes almost discontinuous as this parameter decreases. The hyperparameters can be estimated using Maximum Likelihood Estimation (MLE, see [Forrester et al. \(2008\)](#)) and, in order to predict new points, a vector of correlations between the observed data and the new prediction is defined as

$$\boldsymbol{\psi} = \begin{pmatrix} \text{cor}[Y(\mathbf{x}^{(1)}), Y(\mathbf{x})] \\ \vdots \\ \text{cor}[Y(\mathbf{x}^{(n)}), Y(\mathbf{x})] \end{pmatrix} = \begin{pmatrix} \psi^{(1)} \\ \vdots \\ \psi^{(n)} \end{pmatrix}. \quad (2.27)$$

The MLE gives the final result for the predictions $\hat{y}(\mathbf{x})$

$$\hat{y}(\mathbf{x}) = \hat{\mu} + \boldsymbol{\psi}^T \boldsymbol{\Psi}^{-1} (\mathbf{y} - \mathbf{1}\hat{\mu}). \quad (2.28)$$

The mean squared error (MSE) is then

$$\hat{s}^2(\mathbf{x}) = \sigma^2 (1 - \boldsymbol{\psi}^T \boldsymbol{\Psi}^{-1} \boldsymbol{\psi}), \quad (2.29)$$

omitting the last term from [Sacks et al. \(1989\)](#), as it is usually very small and does not appear if the derivation is performed from a Bayesian approach ([Forrester et al., 2008](#)).

2.4.2 Balance exploration and exploitation: Expected Improvement

Once the surrogate model is constructed, a criterion which allows for a balance of local minimization and local exploration needs to be assessed. The Expected Improvement is used, which has been extensively applied in literature for Bayesian optimization (Wang et al., 2016) and LES problems (Talnikar et al., 2015a). Its expression is

$$E[I(\mathbf{x})] = \begin{cases} (y_{\min} - \widehat{y}(\mathbf{x}) - \xi_{\text{EI}})\Phi(Z) + s\phi(Z) & \text{if } s > 0, \\ 0 & \text{if } s = 0. \end{cases} \quad (2.30a)$$

$$Z = \frac{y_{\min} - \widehat{y}(\mathbf{x}) - \xi_{\text{EI}}}{\widehat{s}(\mathbf{x})} \quad (2.30b)$$

In this expression, $\Phi(\cdot)$ and $\phi(\cdot)$ are the cumulative distribution function and the probability density function, respectively. The parameter ξ_{EI} controls a balance between exploration and exploitation. A value $\xi_{\text{EI}} > 0$ promotes exploration, with $\xi_{\text{EI}} < 0$ doing it for exploitation.

This criterion balances in each point the probability of finding an improvement over the current minimum with the amount of additional information that can be gathered to improve the quality of the surrogate model. This aims to avoid the optimizer getting stuck in regions with marginal probabilities of improvement and also to simply improve the accuracy of the predictions while not optimizing significantly the objective function.

2.4.3 Extension to multi-fidelity analysis: Co-Kriging

While a Kriging surrogate model allows the prediction over a small quantity of expensive data, there exists the possibility of employing a greater amount of lower fidelity, cheaper data corrected in some way to be coupled in the process and guide the surrogate model. This essence lead to the Co-Kriging formulation (Forrester et al., 2007, 2008; Brooks et al., 2011). This Kriging version correlates multiple datasets using the c and e sub-indices for the cheap and expensive data, the sample locations are

$$\mathbf{X} = \begin{pmatrix} \mathbf{X}_c \\ \mathbf{X}_e \end{pmatrix} = \begin{pmatrix} \mathbf{x}_c^{(1)} \\ \vdots \\ \mathbf{x}_c^{(n_c)} \\ \mathbf{x}_e^{(1)} \\ \vdots \\ \mathbf{x}_e^{(n_e)} \end{pmatrix}. \quad (2.31)$$

Using the auto-regressive model (Kennedy and O'Hagan, 2000), assuming that

$$\text{cov}\{Y_e(\mathbf{x}^{(i)}), Y_c(\mathbf{x}) | Y_c(\mathbf{x}^{(i)})\} = 0,$$

which is translated to the hypothesis that nothing else can be learnt about $Y_e(\mathbf{x}^{(i)})$ from the cheaper solution if this value is already known for the expensive code (Markov property which states that the expensive simulation is treated as the ground truth).

The local features of the cheap and expensive solutions are represented by Gaussian processes $Z_c(\cdot)$ and $Z_e(\cdot)$. If the expensive code is approximated by the expensive one multiplied by a constant scaling ρ and a difference Gaussian process $Z_d(\cdot)$:

$$Z_e(\mathbf{x}) = \rho Z_c(\mathbf{x}) + Z_d(\mathbf{x}). \quad (2.32)$$

The covariance matrix now needs to include the correlations between the cheap and expensive models, as

$$\mathbf{C} = \begin{pmatrix} \sigma_c^2 \Psi_c(\mathbf{X}_c \mathbf{X}_c) & \rho \sigma_c^2 \Psi_c(\mathbf{X}_c \mathbf{X}_e) \\ \rho \sigma_c^2 \Psi_c(\mathbf{X}_e \mathbf{X}_c) & \rho^2 \sigma_c^2 \Psi_c(\mathbf{X}_e \mathbf{X}_e) + \sigma_d^2 \Psi_d(\mathbf{X}_e \mathbf{X}_e) \end{pmatrix}. \quad (2.33)$$

The hyperparameters to estimate are now $\boldsymbol{\theta}_c$, $\boldsymbol{\theta}_d$, \mathbf{p}_c and \mathbf{p}_d . As with Kriging, MLE can be used with an augmented dataset, which leads to a prediction for the expensive points

$$\hat{y}_e(\mathbf{x}) = \hat{\boldsymbol{\mu}} + \mathbf{c}^T \mathbf{C}^{-1} (\mathbf{y} - \mathbf{1} \hat{\boldsymbol{\mu}}), \quad (2.34)$$

begin \mathbf{c} the column vector of the covariance of \mathbf{X} and a new predicted location \mathbf{x} . The estimated mean squared error is then

$$\hat{s}^2(\mathbf{x}) \approx \rho^2 \hat{\sigma}_c^2 + \hat{\sigma}_d^2 - \mathbf{c}^T \mathbf{C}^{-1} \mathbf{c} + \frac{1 - \mathbf{1}^T \mathbf{C}^{-1} \mathbf{c}}{\mathbf{1}^T \mathbf{C}^{-1} \mathbf{1}}. \quad (2.35)$$

This prediction can be used with a Co-Expected Improvement criterion to explore and exploit the parameter space, in an analogous way as in 2.30a.

2.5 Summary

In this chapter, the general formulation of the different methods has been presented. Starting from the derivation of the RANS and LES equations, a discussion about the surrogate models, Kriging and Co-Kriging, with the acquisition criterion used, Expected Improvement, followed.

The details regarding each specific model for each part of this work are further discussed in each chapter.

Local RANS corrections for reference geometry

3.1 Hybrid Injected RANS models formulation

To efficiently perform a multi-fidelity global optimization routine, the two models, higher and lower fidelity, ideally should have a high degree of correlation. In the present document, RANS is considered the lower fidelity model, and LES the ground truth. Nevertheless, it should be noted that, for the current correction framework, the higher fidelity method could well be any other (such as DNS, for example), as long as it contains the necessary inputs for the RANS model.

RANS models have been proven moderately successful in predicting mean loads in a variety of fields, including aerospace applications, but easily find limitations when challenged in certain situations. The contribution of the anisotropic components which can be found in the Reynolds stresses is often mispredicted, depending on the assumptions made in the specific turbulent model. Due to this, there is a need for correcting the RANS model before introducing it in a multi-fidelity optimization process, in order to get it closer to its higher fidelity counterpart.

3.1.1 HIRANS_a: injection of anisotropy information

From the hypothesis that the anisotropy components can be highly relevant in the prediction of the flow in near-wall regions, it is natural to aim to introduce this information into the RANS model. To do this, a hybrid non-dimensional anisotropic part of the Reynolds stresses can be formulated. Starting from the conventional normalized anisotropy tensor

$$b_{ij} = \frac{a_{ij}}{2k} = \frac{\langle u'_i u'_j \rangle}{\langle u'_l u'_l \rangle} - \frac{1}{3} \delta_{ij}, \quad (3.1)$$

the HIRANS hybrid Reynolds stress tensor is defined using the non-dimensionalized anisotropy information b_{ij} from LES. Nevertheless, introducing the stress tensor information directly from LES can result in ill-conditioning of the equations (Wu et al., 2018), which can lead to large errors in the mean velocity components even if the stresses are represented with great accuracy, as it is the case of the injection of information from very accurate DNS simulations. In addition to this study, (Poroseva et al., 2016) shows that this effect becomes more pronounced as the Reynolds number increases. To overcome this, which does not only affect the accuracy of the mean velocity predictions but also the numerical stability of the solver, the hybrid tensor is scaled with the value of k obtained in the RANS transport equation, giving

$$\tau_a^* = 2k_{RANS}(b_{ijLES} + \frac{1}{3}\delta_{ij}). \quad (3.2)$$

Furthermore, injection parameters ξ_i can be used, in a similar way as the blending functions in spatial hybrid RANS/LES approaches (Fan et al., 2001), with the difference that, in the current case, the simulation is an entirely RANS model, and this parameter remains constant spatially. These injection parameters control the amount of RANS and LES information which is introduced in the HIRANS model. The stress tensor is then expressed as

$$\tau_{HI} = \xi_i \tau^* + (1 - \xi_i) \tau_{k-\omega}, \quad (3.3)$$

where $\tau_{k-\omega}$ is the part corresponding to the RANS stress tensor obtained using the standard $k-\omega$ equations. This hybrid tensor acts in both the momentum and the production term of the k and ω transport equations, introducing a percentage of the LES anisotropy in both of them, but consistently scaling it with the RANS isotropic component. Both equations are then

$$\rho \langle U_j \rangle \frac{\partial \langle U_i \rangle}{\partial x_j} = \frac{\partial}{\partial x_j} \left[- \langle p \rangle \delta_{ij} + 2\mu \langle S_{ij} \rangle - \rho \left(\underbrace{\xi_m \tau^*}_{\text{Hybrid form}} + \underbrace{(1 - \xi_m) \tau_{ij_{k-\omega}}}_{\text{Standard } k-\omega} \right) \right], \quad (3.4a)$$

$$\rho \frac{\partial k}{\partial t} + \rho \langle U_j \rangle \frac{\partial k}{\partial x_j} = \left(\underbrace{\xi_t \tau^*}_{\text{Hybrid form}} + \underbrace{(1 - \xi_t) \tau_{ij_{k-\omega}}}_{\text{Standard } k-\omega} \right) \frac{\partial \langle U_i \rangle}{\partial x_j} - \beta^* \rho k \omega + \frac{\partial}{\partial x_j} \left[(\mu + \sigma^* \mu_T) \frac{\partial k}{\partial x_j} \right], \quad (3.4b)$$

$$\rho \frac{\partial \omega}{\partial t} + \rho \langle U_j \rangle \frac{\partial \omega}{\partial x_j} = \alpha \frac{\omega}{k} \left(\underbrace{\xi_t \tau^*}_{\text{Hybrid form}} + \underbrace{(1 - \xi_t) \tau_{ij_{k-\omega}}}_{\text{Standard } k-\omega} \right) \frac{\partial \langle U_i \rangle}{\partial x_j} - \beta \omega^2 + \frac{\partial}{\partial x_j} \left[(\mu + \sigma \mu_T) \frac{\partial \omega}{\partial x_j} \right], \quad (3.4c)$$

where $\langle S_{ij} \rangle = 1/2(\partial \langle U_i \rangle / \partial x_j + \partial \langle U_j \rangle / \partial x_i)$ is the mean strain rate and the subscripts m, t correspond to independent values for the momentum and k transport equations, respectively. The transport equation for ω is modified similarly as the one for k .

3.1.2 HIRANS_b: injection of full LES Reynolds stress tensor

The correction of the anisotropic contribution of the Reynolds stresses does only partially take into account the overall mismatch between RANS and LES stress components. A larger amount of information can be injected, using the full isotropic and anisotropic components from LES. Nonetheless, major difficulties appeared in the convergence process of the solver in the course of this document in this case, even using ξ_i injection parameters.

A slightly less greedy approach can be taken, if the turbulent transport equations remain as in the previous model, but the direct contribution in the momentum equation is modelled using the full τ_{LES} . This introduces two different hybrid tensors,

$$\tau_{HI_m} = \xi_m \tau_{LES} + (1 - \xi_m) \tau_{k-\omega}, \quad (3.5a)$$

$$\tau_{HI_t} = \xi_t \tau^* + (1 - \xi_t) \tau_{k-\omega}, \quad (3.5b)$$

where τ^* is the same hybrid tensor formulated in the previous HIRANS model in (3.2).

If the isotropic part of the Reynolds stress tensor is absorbed into a modified pressure, $p^* = \langle p \rangle + 2/3k$, the effective anisotropic part of the Reynolds stress tensor in the momentum equation is

$$a_{ij_m} = \begin{cases} 2k_{RANS} \left[\xi_m b_{ij_{LES}} + (1 - \xi_m) b_{ij_{RANS}} \right] & \text{for HIRANS}_a, \\ 2k_{RANS} \left[\xi_m b_{ij_{LES}} \frac{k_{LES}}{k_{RANS}} + (1 - \xi_m) b_{ij_{RANS}} \right] & \text{for HIRANS}_b. \end{cases} \quad (3.6)$$

The difference between both terms, non-dimensionalizing it with k_{LES} , is

$$\frac{a_{ij_b} - a_{ij_a}}{2k_{LES}} = \xi_m b_{ij_{LES}} \left(1 - \frac{k_{RANS}}{k_{LES}} \right), \quad (3.7a)$$

$$\frac{\tau_{ij_b} - \tau_{ij_a}}{2k_{LES}} = \xi_m \left(b_{ij_{LES}} + \frac{1}{3} \delta_{ij} \right) \left(1 - \frac{k_{RANS}}{k_{LES}} \right). \quad (3.7b)$$

The expression (3.7b) draws the following conclusion: the relative levels of anisotropy between both models depend on only two factors, for the same value of the injection parameter ξ_m : the ratio between the turbulent kinetic energy of RANS and LES and the local non-dimensional anisotropy of the LES solution. That is, a non-accurate prediction of the isotropic part of the turbulence in the RANS model means also that HIRANS_a will compute anisotropy levels further away from HIRANS_b, which uses more LES information. Both models would be identical if $k_{RANS} \rightarrow k_{LES}$. For the case of purely isotropic turbulence, if $b_{ij_{LES}} \rightarrow 0$, assuming finite values of k for RANS and LES, both solutions would also collapse into the same anisotropy, but, as (3.7a) shows, the Reynolds stress tensor will only be the same if $k_{RANS} = k_{LES}$. As it has already been mentioned, the isotropic part can be absorbed as part of a modified pressure p^* , so this does not affect the computations in the current model.

The momentum equation is then

$$\rho \langle U_j \rangle \frac{\partial \langle U_i \rangle}{\partial x_j} = \frac{\partial}{\partial x_j} \left[- \langle p \rangle \delta_{ij} + 2\mu \langle S_{ij} \rangle - \rho \left(\underbrace{\xi_m \tau_{LES}}_{\text{Pure LES}} + \underbrace{(1 - \xi_m) \tau_{ij_{k-\omega}}}_{\text{Standard } k-\omega} \right) \right]. \quad (3.8)$$

As previously stated, the k and ω transport equations are the same as in the previous HIRANS model, (3.4b, 3.4c).

3.2 Numerical setup

The HIRANS framework has been developed by a Python 2.7 framework (van Rossum, 1995; Team, 2018), which, given the input parameters, is able to automatically mesh, setup, run and postprocess multiple simulations.

3.2.1 Flow parameters and mesh quality




Figure 3.1: Left: grid used for the RANS case.
Right: y^+ values for the RANS bottom (- - -) and top (- . - .) walls and LES bottom (—) and top (- . . -) walls.

To ensure that the quality of the solution is at the required standards, the grid has been properly assessed. As observed in Figure 3.1, the values of y^+ have been kept below 1 in the majority of the bottom wall, with the exceptions of the first flat part after the leeward hill and on the high peak of shear stress in the windward hill. The resolution of the top wall is less relevant for the different features of interest to be corrected and has been kept around 2 for computational cost purposes and to avoid high aspect ratio cells near the wall. A sensitivity study has been performed with the HIRANS solution with respect to near-wall resolution. The results of it do not provide any overall benefits with respect to the mesh shown. The reference LES data used for the injection of information provides very low values of y^+ for the bottom wall and a DNS-like resolution for the top one, not achievable without a large cost penalty in the current study. The number of cells

in the RANS case is 120×130 in streamwise and normal direction, respectively, while the reference LES uses around $298 \times 220 \times 200$ cells in streamwise, normal and spanwise directions.

The simulations were performed at $Re = 10595$ using the previously shown SIMPLE OpenFOAM algorithm for pressure-velocity coupling, with second order Gauss integration for gradient and divergence terms as a default, using upwind options for the fluxes of U , k and ω . The turbulence model used was the OpenFOAM implementation of $k - \omega$ with the addition of the isotropic missing term explained in Section 2.2.3. The pressure was solved using geometric-algebraic multigrid with a Gauss-Seidel iterative solver for the smoother step. Velocity components, k and ω were solved using a Diagonal-Incomplete-LU-preconditioned BiCG method.

3.2.2 Mass flow control

In order to keep the mass flow in the required level to achieve a Reynolds number at the crest of the hill of $Re_S = 10595$, a forcing term needs to be added. The relation between the Reynolds number at the hill crest and the one using the bulk velocity along the whole domain, Re_B , can be calculated via a geometric factor Γ (Balakumar et al., 2014), as

$$Re_B = \Gamma Re_S. \quad (3.9)$$

This geometric factor can be calculated as

$$\Gamma = \frac{L_{y_{x=0}}}{1/L_x \int_0^{L_x} L_y(x) dx}. \quad (3.10)$$

This geometric factor has an approximate value of $\Gamma \approx 0.72$ for the reference periodic hill geometry. This factor is used to ensure that there is the proper momentum in the global field. A source term f_p in the form of a pressure-gradient volumetric force is added to the right hand side of the momentum equation, which corrects the field momentum to achieve the correct value of U_B at every iteration.

3.2.3 Injection strategy and convergence control

The process of injecting an exogenous input to the RANS equation may introduce numerical instabilities, potentially unsurpassable for the method. Due to this, an injection strategy for the parameters ξ_i is proposed,

$$\xi_i = \begin{cases} \xi_{i_0} + (\xi_{i_f} - \xi_{i_0}) \left(\frac{t}{t_f - t^*} \right) & \text{for } t \leq t_f - t^*, \\ \xi_{i_f} & \text{for } t > t_f - t^*, \end{cases} \quad (3.11)$$

where the parameters ξ_i are not fully introduced from the first iteration, but rather start from values ξ_{i_0} and then ramp up linearly until achieving their maximum values ξ_{i_f} , with

t^* iterations left until the simulation stops at its final t_f . This method fulfils two goals: on the one hand, it allows the solver to gradually treat the external information, and on the other hand, the t^* iterations allow it to converge (if convergence exists) to the solution once the injection factor is at the maximum desired value. Unless explicitly mentioned, the parameters used in the different simulations are $\xi_{i_0} = \xi_{i_f}/2$ and $t^* = 5000$ iterations, with $t_f = 10000$ iterations.

To monitor the convergence process, and also to formulate criteria with respect to the numerical stability of the solution, the residuals from the different main quantities (pressure, velocity components, k and ω) are taken, as well as the force coefficients from the flowfield. The assessment of them is performed after a filtering process, as

$$\bar{R}^{n+1} = \frac{\bar{R}^n \min(n, \alpha - 1) + R^{n+1}}{\min(n + 1, \alpha)}, \quad (3.12)$$

where the quantity R is any residual or load, the filtered quantity at the first iteration is $\bar{R}_0 = R_0$ and the filtering constant is fixed to $\alpha = 1000$ iterations. For the force coefficients, a similar filtering is applied to a non-dimensional numerical gradient with respect to the previous iteration \bar{c}_i^n , as

$$\bar{c}_i^n = \frac{\bar{c}_i^n - \bar{c}_i^{n-1}}{\bar{c}_{i_{BSL}}}. \quad (3.13)$$

In this expression, the subscript i refers to the difference force coefficients (lift and drag), and the superscripts refer to the correspondent iteration number. The filtered force coefficients $\bar{c}_{i_{BSL}}$ are computed from the baseline RANS iteration (that is, pure $k - \omega$). This numerical gradient, which evaluates the influence of the current iteration on the filtering quantity, allows to show the growth of the perturbations for the unstable cases.

3.3 Results

3.3.1 Baseline RANS model

As expressed before, the periodic hill test case has been chosen based on previous results from literature which show the limitations of several RANS models to quantitatively predict features such as the recirculation region and the reattachment location. While qualitatively these phenomena are greatly influenced by the geometric constraints, as seen in Figure 3.2, it is expected that the baseline $k - \omega$ model is a significant departure from the reference LES case in both mean and turbulence quantities. The differences between RANS and LES cases are evaluated as

$$\Delta\{u_i\} = u_{i_{RANS}} - u_{i_{LES}}, \quad (3.14)$$

which are then appropriately non-dimensionalized using the bulk velocity at the crest of the hill U_b . The results for the velocity components are illustrated in Figure 3.3. There is




Figure 3.2: Streamwise velocity contours for RANS (left) and LES (right) solutions.

a significant difference with respect to the LES solution in both streamwise and vertical velocities, which is pronounced near the separation region of the first hill and becomes highly relevant close to the reattachment point. For this baseline case, the $k - \omega$ model predicts a much larger separation region, with a reattachment point much further than the reference LES.




Figure 3.3: Differences of RANS streamwise (left) and vertical (right) velocity components with respect to the LES solution.

Figure 3.4 shows this last point, where it is observed how the shear stress distribution at the bottom wall is significantly different from the one of the reference LES case. Starting from the separation location, this one occurs roughly at the same location, as it could be expected based on the strong local pressure gradient, which is the dominant factor at this location. Nevertheless, a much higher peak of shear is observed close to the top of the first hill in RANS. This peak also corresponds to a locally higher Re close to the wall,

Method	Reattachment location $[x/h]$
RANS	5.74
LES	4.71




Figure 3.4 & Table 3.1: Left: Reattachment locations.
 Right: Wall shear stress distribution for RANS (- - - -) and
 LES (———).

as can be observed if looking at the distribution of streamwise velocity, which somewhat prevents its separation for some time, until the pressure gradient becomes too strong. The explanation of this difference in the local Re close to the wall is also related to the physics close to the second hill. While RANS predicts the increment in wall shear stress due to the slope of the geometry fairly well in location, it fails to capture its magnitude, having lesser stresses during all the second section, but also not going close to zero as the LES solution on the cyclic boundary. Due to this, these conditions are translated to the first cyclic boundary, and a higher shear at the beginning prevents the flow to separate at the same location. It can also be observed that this shift is almost irrelevant at the next location downstream close to the wall, as RANS is able to capture quite well the separated region downhill. Nevertheless, the difference in the intensity of the fluctuations in the shear layer shifts the reattachment point to around one x/h unit more downstream, which explains the streamwise velocity deficit observed between the two hills in Figure 3.3.

With respect to the hypothesis of the lack of accuracy of the baseline RANS in the prediction of the anisotropy of the flow, Figure 3.5 shows that these deviations do indeed exist, but they are not the most dominant ones for this case. The most important stress component which is underpredicted is $\langle u'u' \rangle$, which presents large deviations in the shear layer region on the separation after the first hill. This contribution is dominant in the underprediction of the turbulent kinetic energy in this region. On the other hand, it is observed how RANS overpredicts the $\langle v'v' \rangle$ stress close to the top of the second hill. This is the region of the highest wall shear stress of the flow (which, as it has been shown, is slightly shifted downstream in the RANS solution). Moreover, the spanwise component $\langle w'w' \rangle$ is underpredicted as well, as it could be expected due to the lack of three-dimensional vortical features in the RANS solution.

To explain this behaviour, it is relevant to look at the different structures which can be found in the LES reference data (Breuer et al., 2009). In this reference, for the case of low Re , that is, $Re = 200$, there are highly elongated structures close to the second hill, which

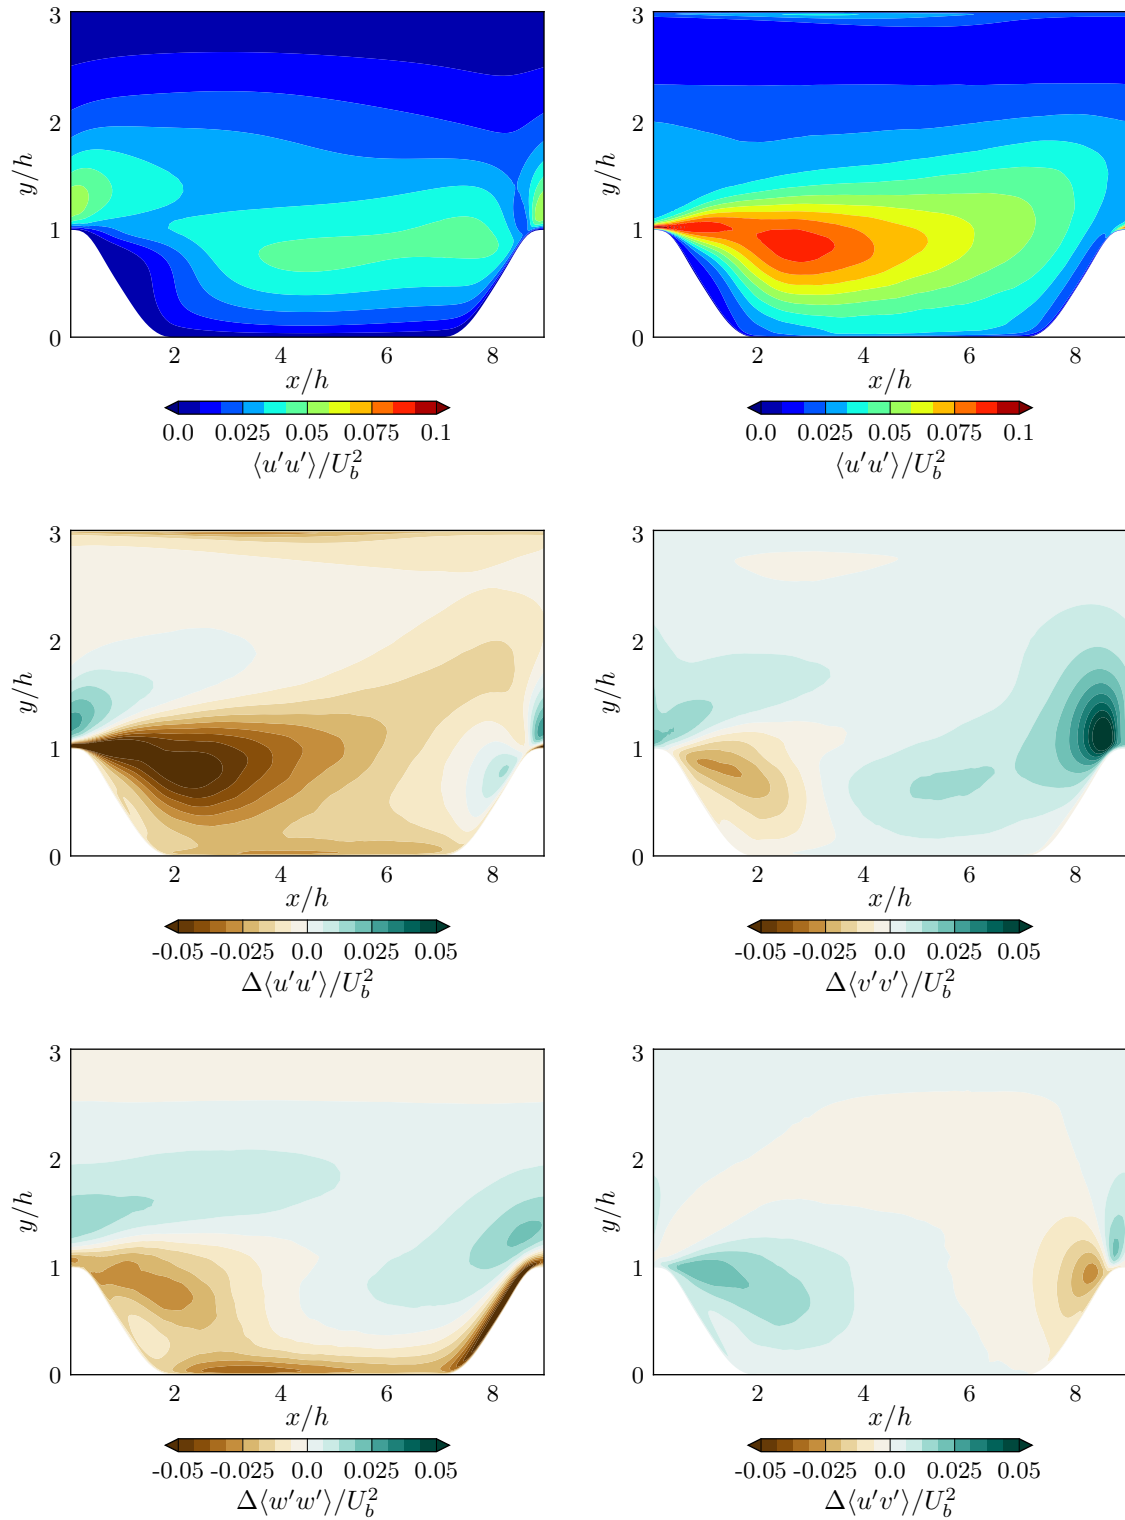


Figure 3.5: Top: $\langle u'u' \rangle$ for RANS (left) and LES (right).
Rest: Reynolds stresses differences with respect to LES values.




Figure 3.6: Contours of turbulent kinetic energy for RANS (left) and LES (right) solutions.

are argued to be created by centrifugal forces, as shown by (Görtler, 1940), or by Craik-Leibovich type-2 instabilities (Phillips and Wu, 1994). These structures are located along a large region of the second hill, highly affecting the vertical and spanwise components. At the higher Reynolds number considered in this document, those structures would be expected to be substantially smaller. This is supported by (Fröhlich et al., 2005), which shows that for $Re = 10595$, these structures are found in different spanwise locations, introducing flow rotation in the spanwise component. This simply cannot be predicted by the 2D RANS analysis, where $\langle w'w' \rangle \approx 2/3k$.

Furthermore, observing Figure 3.6, not only is the anisotropy of the flow poorly predicted by the $k - \omega$ model, but also the isotropic turbulent kinetic energy. There is a tendency of the RANS model to overpredict the value of k on the windward side of the hill, close to the top. However, the Kelvin-Helmholtz-type vortices observed for the LES solution, in the high shear region where separation occurs, do not seem to be predicted quantitatively in the steady RANS simulation.

Using these observations, it seems appropriate to hypothesize that, to increase the accuracy of the RANS model, the anisotropy components, while relevant, are not the only factor to consider. If structures such as the fluctuating shear layer in this flow (or any other kind of instability or large-scale fluctuating component) are not able to be captured appropriately by the steady RANS, injecting uniquely information about the anisotropy will likely still give stresses that deviate from the higher fidelity solution and, with that, misleading mean components. This hypothesis motivates the two formulations studied in the current document. In the first model, scaled anisotropy information is injected, and in the second one, the full Reynolds stress tensor is injected into the momentum equation, while keeping the anisotropy correction of the first formulation for the turbulent transport equations.

3.3.2 HIRANS_a: injection of anisotropy information

The first hybrid model is constructed following the previously shown expression

$$\tau^* = 2k_{RANS}(b_{ijLES} + \frac{1}{3}\delta_{ij}). \quad (3.15)$$

This hybrid Reynolds stress tensor utilizes the anisotropy of the LES solution and is injected using ξ_i factors in the momentum equation and the production term for k . It is expected, a priori, that the mean flow quantities will be affected by the modification in the anisotropy. Nevertheless, in the regions where the mismatch of the RANS isotropic contribution with respect to LES is important quantitatively, such as in the shear layer after the separation region, the improvements are expected to be modest.




Figure 3.7: Baseline RANS case ($\xi_i = 0$).
 Left: filtered residuals of $\langle p \rangle$ (—), $\langle U_x \rangle$ (- - -), $\langle U_y \rangle$ (- · - · -),
 k (· · · · ·) and ω (- - -).
 Right: Non-dimensional filtered lift (—) and drag (- - -) coefficients
 gradient per iteration.

Convergence

Before presenting the results, it is relevant to mention the strong effect which the injection factors ξ_i have in the convergence and stability of the numerical solution. It is desirable to increase these factors as close to 1 as possible, to maximize the LES anisotropy injection. Nevertheless, the tests performed show numerical stability only for injection factors $\xi_i \leq 0.8$, with $\xi_{i_0} = 0$ and $t^* = t_f/2 = 5000$ iterations, as shown in Figure 3.7 for the baseline RANS model and Figure 3.8 for the HIRANS tests. There is a clear trend in the different cases: they are numerically stable as the injection factor is increased, until it reaches a certain point where the solver cannot withstand it. In all the tests performed, a small kink appears when $\xi_{i_t} = \xi_{i_{\max}}$. For the stable cases, the solution keeps converging after a certain number of iterations after this kink (around 1000), but for the unstable ones, the instabilities grow in amplitude as the iteration number increases with

an exponential trend and an increasing slope as the injection factor grows, as observed for the cases regarding $\xi_i = 0.9$ and $\xi_i = 1$.

The spectral radius change on the matrices for the iterating procedure when injecting the LES information is significant, as the replacement of the stabilizing term of the Reynolds stresses is changed by an exogenous source, τ^* . The importance of this term is highlighted in the fact that more greedy approaches in the production term definition, as introducing mean velocity components from LES, failed in the convergence progress. It is specially relevant in the case of the solving process of the velocity components, which suffer a sharp change of trend around values of $\xi_m \approx 0.9$. This observation could potentially allow to perform simulations close to this value, injecting linearly ξ_i with a cut-off of the solving process when $\xi_m = 0.9$, without allowing any additional iterations to relax. This method has the advantage of maximizing the amount of LES information injected, going beyond $\xi_i = 0.8$, but the clear downside is that the simulation stops in a region where the discrete system for calculating the velocity components is already unstable. It could be argued that a more refined analysis could be performed, evaluating where is the instability threshold, and the exact mechanism which promotes it, while seeking for solutions for its damping.

For the current study, injection factors of $\xi_i = 0.8$ have been used, as they provide a stable solution in terms of the filtered residuals and load coefficients, both for local and non-local analysis for small departures of the baseline geometry, which allows to evaluate the performance of the injection techniques in those cases in a robust way. Higher injection factors closer to the instability threshold would result in different values of ξ_i for different geometries, which would introduce an additional variable in the analysis clouding the results. For a further discussion about the convergence process, the reader can refer to Appendix A, where higher injection factors are tested for the second HIRANS model.

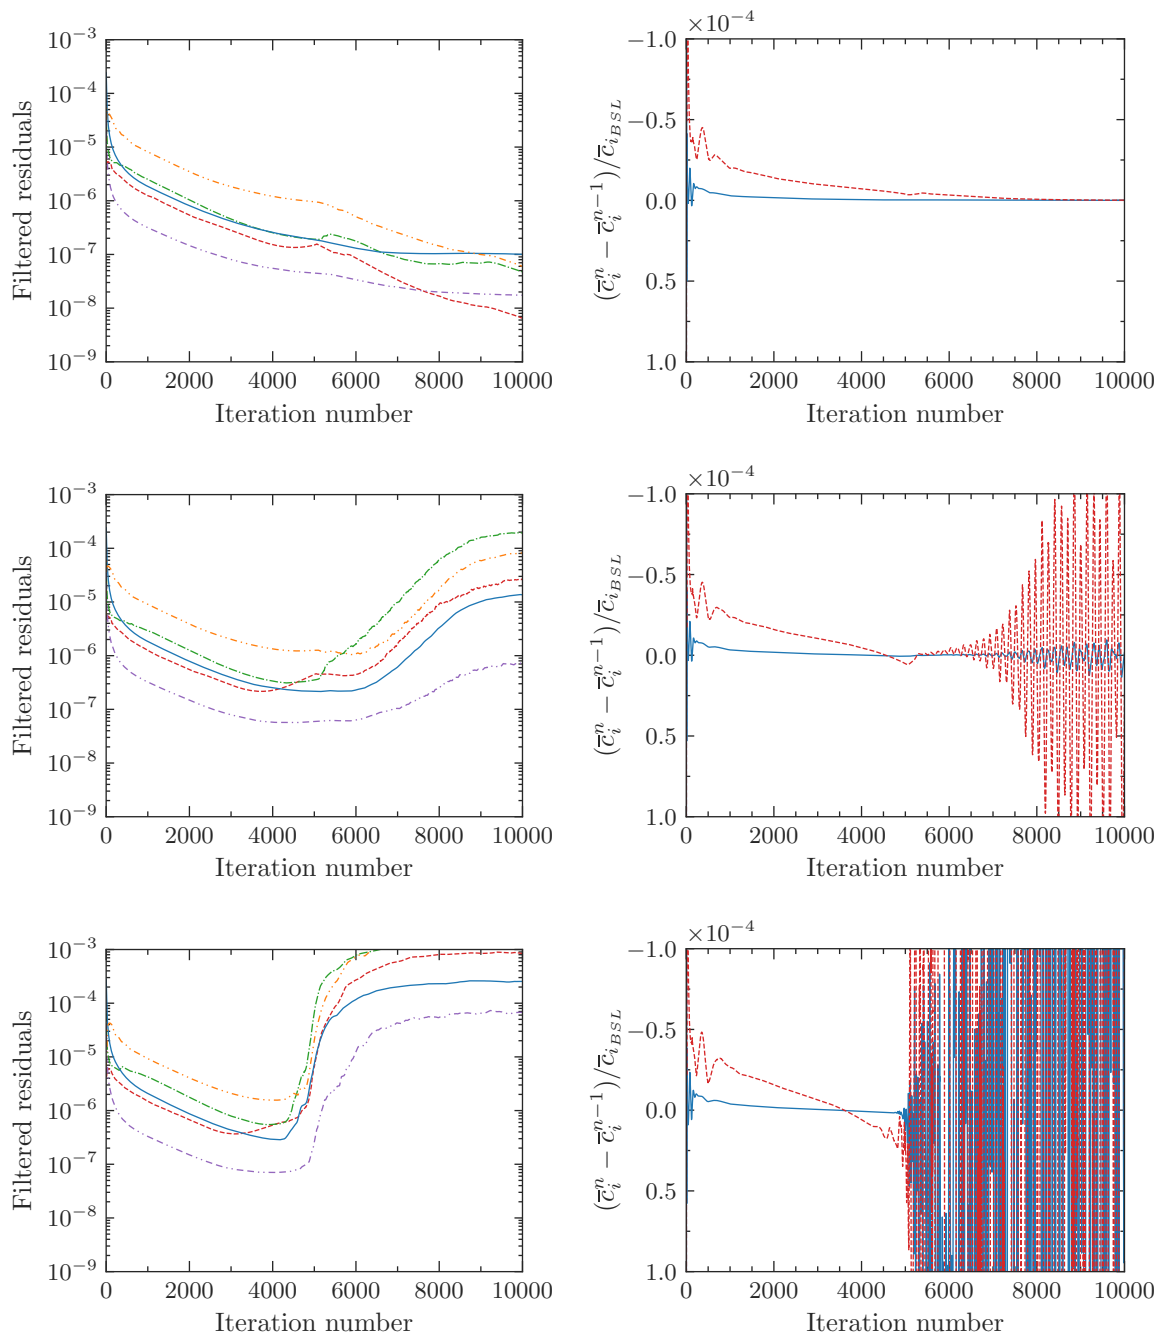


Figure 3.8: HIRANS cases for $\xi_i = 0.8$ (top), $\xi_i = 0.9$ (center) and $\xi_i = 1$ (bottom).
 Left: filtered residuals of $\langle p \rangle$ (—), $\langle U_x \rangle$ (- - -), $\langle U_y \rangle$ (- · - · -),
 k (· - - · -) and ω (- - - - -).
 Right: Non-dimensional filtered lift (—) and drag (- - -) coefficients
 gradients per iteration.

Results




Figure 3.9: Differences of HIRANS streamwise (left) and vertical (right) velocity components with respect to the LES solution.

The results for the mean velocity components with respect to the LES solution, which can be observed in Figure 3.9, are not specially encouraging for this case: there is hardly any difference compared to the baseline RANS result on the shear layer and recirculation region. However, the correction on the anisotropy leads to a much better agreement near the top wall than the baseline RANS results.

The corrections for the different components of the Reynolds stresses show (Figure 3.10) an improvement over the baseline results. It should be noted that the contours presented correspond to the effective stress tensor which enters in the momentum equation, that is, $\tau_{HI_m} = \xi_m \tau^* + (1 - \xi_m) \tau_{k-\omega}$. With respect to the streamwise normal stress $\langle u'u' \rangle$, the differences in the shear layer are reduced to around half the magnitude with respect to the baseline RANS case. Even though, it can be argued that they are still very noticeable, and that there is a deficit of these fluctuations in the recirculating regions. For the $\langle v'v' \rangle$ component, the differences also become much smaller than on the previous case, improving specially in the high shear region close to the windward part of the hill. The $\langle w'w' \rangle$ component seems less affected by the results, still showing an underprediction in the accelerating flow part of the second hill. Finally, for the shear component $\langle u'v' \rangle$, the best correction has been observed, with a close agreement with respect to the LES results.

Looking at the individual components of the Reynolds stresses and their modest improvement in the first HIRANS model, the question of why the mean velocity components do not seem to follow the same path arises. Firstly, while along the whole flowfield there is indeed a closer correlation between the HIRANS and LES solutions, in the places where the individual components are dominant quantitatively in the latter (separated region and shear layer region for $\langle u'u' \rangle$ and $\langle v'v' \rangle$, windward side of the hill for $\langle w'w' \rangle$) are still far off from the LES values. Especially relevant are the large fluctuations of streamwise and normal velocity components on the leeward part of the hill, which are not being able to be captured accurately by this HIRANS method.




Figure 3.10: Reynolds stresses differences with respect to LES values

Moreover, Figure 3.11 shows that while there has been an improvement of the individual components, the trace of the Reynolds stress tensor is still far off from the LES values. The turbulent kinetic energy is largely underpredicted after the separation point, and the picture of it does not look much different from the baseline RANS case. This is not surprising, as if the expression used for the production of k in its transport equation (2.8a) is recalled, only the anisotropy has been modified. This corrected anisotropy is scaled by the mean velocity gradients, which are not very different from the standard $k - \omega$ case. In the end, it results in a very “mild” correction, which does not lead to a great difference in effect in the momentum equation. If the anisotropic part is following Boussinesq’s hypothesis and the isotropic part is hardly changed, it is then expected that the mean velocity profiles should also be quite similar, as it is the case, showing the interlinking between the momentum and k transport equations.




Figure 3.11: Contours of turbulent kinetic energy for HIRANS (left) and LES (right) solutions.

Using the exact expression for the production term

$$\mathcal{P} = -\langle u'_i u'_j \rangle \frac{\partial \langle U_i \rangle}{\partial x_j}, \quad (3.16)$$

and knowing that only the symmetric part of the velocity-gradient tensor and the anisotropic part of the Reynolds stresses affect the production, (3.16) can also be written in terms of the anisotropic stresses and the mean rate of strain as

$$\mathcal{P} = -a_{ij} \langle S_{ij} \rangle. \quad (3.17)$$

Examining (3.17), it is clear that if only the anisotropic parts of the stresses are “implicitly” injected (the “implicit” nomenclature arises from the similarity of this treatment to the one found in Wu et al. (2018), as they are non-dimensionalized with k_{RANS}), the effect is going to likely be only limited.

The hybrid stress tensor constructed, τ^* , is damped by the local value of k_{RANS} , and this damped tensor is introduced in the momentum equation, which eventually leads to similar velocity components, and with that, velocity gradients, as the baseline RANS. While this is a “virtuous” circle in terms of the stability of the solver, it is a “vicious” one in terms of the actual departure from the $k - \omega$ solution, where even high values of ξ_i are eventually dominated by the local values of k computed by the $k - \omega$ transport equation.

3.3.3 HIRANS_b: injection of full LES Reynolds stress tensor

The second HIRANS model introduces an injection of the full Reynolds stress tensor τ_{LES} from a LES (or any other higher fidelity) solution, as formulated in (3.18). The aim is to directly act on the contribution of the mean velocity components in the momentum equation, but still retaining the hybrid stress tensor τ^* in the k transport equation which provides numerical stability to the method, resulting in

$$\tau_{HI_m} = \xi_m \tau_{LES} + (1 - \xi_m) \tau_{k-\omega}, \quad (3.18a)$$

$$\tau_{HI_t} = \xi_t \tau^* + (1 - \xi_t) \tau_{k-\omega}. \quad (3.18b)$$

If the mean velocity gradients have been explicitly modified by the contribution of the LES stresses in (3.8), the production of the turbulent kinetic energy would also be affected, as formulated in (3.4b), modifying its quantitative value towards the LES solution but still retaining a certain amount of damping which would help the solution to not diverge after introducing an exogenous input.

In terms of the numerical stability of this approach, the results were qualitatively similar to the ones shown for the anisotropy injection, with values of $\xi_i \leq 0.8$ providing stable solutions. A more detailed report of the convergence for different cases can be found in Appendix B, where the convergence and results for $\xi_i \geq 0.8$ and two different values of t^* are analysed. The results presented, unless explicitly specified, correspond to injection factors of $\xi_m = 0.8$, $\xi_t = 0.3$. The motivation behind this choice is discussed in Section 3.3.5.

Results




Figure 3.12: Differences of HIRANS streamwise (left) and vertical (right) mean velocity components with respect to the LES solution.

The results for the mean velocity components were much better for this case, being in a close agreement with the reference LES solution in the majority of the flowfield, as




Figure 3.13 & Table 3.2: Left: Reattachment locations.
 Right: Wall shear stress distribution for the bottom wall for
 LES (—), RANS (- - -),
 HIRANS with anisotropy injection (- · - ·) and
 HIRANS with full τ_{LES} injection (- - -).

shown in Figure 3.12. The differences in the mean streamwise component were reduced to less than $\pm 5\%$ in the whole domain, excepting a small region close to the separation point, where the value of $\langle U_x \rangle$ in the boundary layer was overpredicted in the HIRANS solution. It should be noted that, apart from the differences on the models used in the different solvers, the mesh resolution in the HIRANS case was not as fine as the reference LES case near the peak of the first hill, and that the upstream conditions (that is, the accelerated part of the second hill, which is introduced as a periodic boundary condition) showed already mean velocity differences within the boundary layer. The value of y^+ for the HIRANS case achieved a peak of around 1.25 in that region, albeit still slightly failed to predict the velocity magnitude there, which is propagated downstream.

The wall shear stress distribution shown in Figure 3.13 for the different models highlights the improvement achieved with this method: from a reattachment location further away one x/h in the $k - \omega$ model, the HIRANS case offered an almost perfect match between $x/h = 4$ and $x/h = 7$, and a virtually identical reattachment point as the LES solution used for the mapping. In addition, this new HIRANS framework achieved a much closer behaviour on the leeward hill, but still underestimated the negative streamwise velocity component in the recirculating part close to the wall.

There is one common feature of all the RANS and HIRANS models: the inability to accurately represent the peak of wall shear stress in the windward hill. This shift in location and magnitude affected also the separation point, which was in practically the same location in all the non-LES models. A careful inspection of the velocity differences with respect to the LES cases (Figures 3.9 and 3.12 for HIRANS_a and HIRANS_b, respectively) shows that the boundary layer profiles looked quite different in this region, thinner in the HIRANS cases, but also with lower gradients close to the wall, which can be related to the turbulence model not being able to capture the gradients from the acceleration region. It is also worth mentioning that the RANS mesh was coarser in the streamwise component,

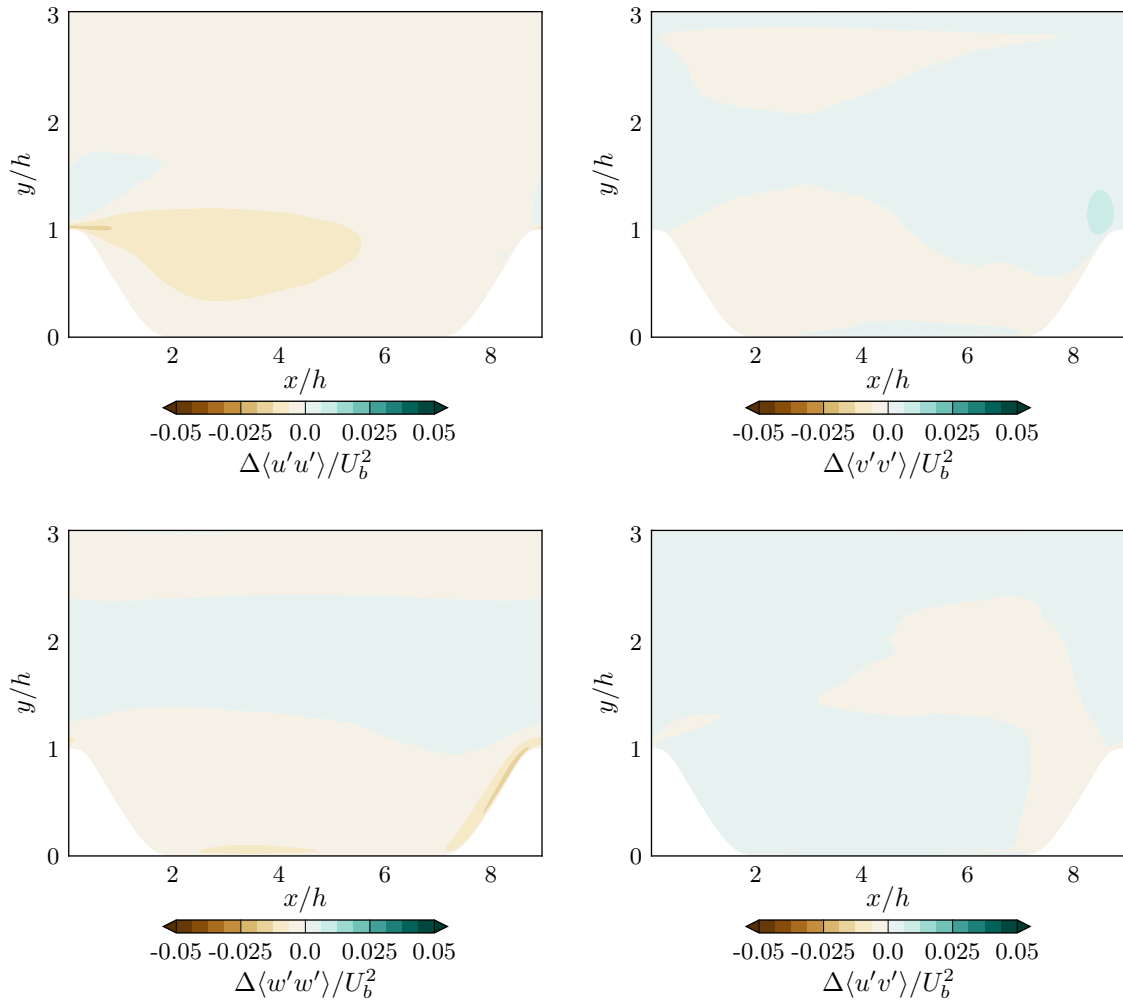


Figure 3.14: Reynolds stresses differences with respect to LES values

with approximately half of the cells of the LES solution, so it is a potentially problematic area for the data injection due to the high velocity gradients found in this streamwise direction. A change of interpolation procedure from linear to cubic of the LES data did not significantly improve the results overall.

The limitations of the RANS part of the model were revealed in the strongest geometry-related dynamics: the flow acceleration on the windward hill and the separation on the leeward one. The LES injection was able to work successfully in the rest of the flowfield, but exhibited difficulties in sections with large fluctuating components and strong pressure gradients.

A conclusion can be drawn looking at the individual stresses components in Figure 3.14. The regions with the higher degree of discrepancy in the wall shear stress for the bottom wall were also the ones with greater differences in these contours. In particular, the separation region still exhibited an underprediction of $\langle u'u' \rangle$, and the recirculating region




Figure 3.15: Contours of turbulent kinetic energy for HIRANS (left) and LES (right) solutions.

and the highly accelerated windward hill showed the largest differences in the spanwise component $\langle w'w' \rangle$. Both regions were likely to be very difficult to be corrected further: for the separation region, the oscillating solution in the unsteady region which has been averaged in the LES solution contains a richness of different structures that the RANS model cannot capture. For the $\langle w'w' \rangle$ discrepancies, the absence of a 3D solution, and hence, a null anisotropic contribution following Boussinesq's hypothesis for this component, greatly limits the RANS part of the model.

Lastly, it is observed from the turbulent kinetic energy contours (Figure 3.15) very close distributions with respect to the LES reference. All the regions had a qualitatively proper distribution of k , with the peak seen on the previous RANS results on the windward hill almost gone and the shear layer fluctuations after the separation very close in size, while slightly underpredicted in magnitude.

3.3.4 Differences between anisotropy and mean flow corrections

The tests performed before for the two HIRANS models have highlighted the fact that a correction of the anisotropy levels of the Reynolds stresses using a higher fidelity solution does not need to necessarily correlate with an improvement of the mean dynamics of the flow prediction. To show this, the barycentric map (Banerjee et al., 2007) for the anisotropy tensor b_{ij} (Figure 3.16) is used as a metric. This map uses the eigenvalues of this non-dimensional anisotropy tensor to describe the state of the turbulent anisotropy at each location. For further details of this metric, please refer to Appendix C.

Figure 3.17 shows the barycentric maps for two different x/h locations. The $k-\omega$ solution gave turbulence components which did not deviate from the linear plane strain line, as expected from the Boussinesq's hypothesis of the model. This was clearly not the case for the LES solution, with significant departures from it, with a variety of eigenvalues closer to the linear contraction region and closer to the two-component limit.




Figure 3.16: Barycentric map for the anisotropy tensor.

The first HIRANS model, with the anisotropy injection, certainly provided a much closer eigenvalue distribution of the Reynolds stress tensor, but as it was stated in Section 3.3.2, the mean dynamics of the flow were not accurately represented, and they even provided a further departure from the LES solution in certain regions. When looking at the second HIRANS model, with full τ_{LES} injection, the eigenvalues distribution was even closer to the LES one, but not significantly different from the first HIRANS case. However, as observed in Section 3.3.3, this model provided a much more accurate solution of the LES mean flow dynamics.




Figure 3.17: Barycentric maps for $x/h = 3.97$ (left) and $x/h = 8.01$ (right). Eigenvalues for LES (\circ), RANS (\odot), HIRANS with LES anisotropy injection (\diamond) and HIRANS with τ_{LES} injection (\triangle).

This observation allows to draw an important conclusion: there could be, for some test cases as the one observed in this periodic hill, a dichotomy between uniquely correcting the anisotropy of the flow and actually accurately represent the rest of the mean flow quantities. The eigenvalues of the Reynolds stress tensor certainly give an idea about the turbulent state of the flow, but both corrections are not always correlated, and this should be taken into account if this metric is used, as it may need to be complemented with others (for instance, the deviation of other flow quantities as velocity components with respect to the reference one, or other test case-based quantities such as the reattachment point for the periodic hill) to fulfil the original goal.

3.3.5 Sensitivity of the solution to the injection factors ξ_i




Figure 3.18: Separation (left) and reattachment (right) locations for HIRANS with LES anisotropy injection (\triangle) and HIRANS with τ_{LES} injection (\diamond).

The previous sections have shown the performance of the two HIRANS models with respect to LES. While the majority of the analysis has been centred in offering qualitative results about the improvements of the different variables on the flowfield, it is also necessary to show the magnitude of the improvements achieved. In this section, it is exposed how the variation of the two different injection factors, ξ_m and ξ_t , affects the solution.

Some obvious metrics for this case are the separation and reattachment locations at the lower wall. Figure 3.18 shows the results for a variation of the two injection factors ξ_i at the same time. Starting from the separation location, there was only a clear trend which can be observed in the first HIRANS case, where only the non-dimensional turbulent anisotropy is injected. Fairly surprisingly, and despite the not particularly good results in capturing the mean velocity components, the separation location seemed to get closer to the LES solution as the injection factors ξ_i were increased. This is consistent with Figure 3.13, which showed that in the small region near the separation, this is the model which behaves closest to the LES reference. Moreover, this first HIRANS model exhibited a good match to the LES reference in the mean streamwise velocity component in that area (Figure 3.9), despite the poor agreement in the fluctuating components (Figure 3.10), specially in $\langle u'u' \rangle$. The explanation for this could be motivated in the non-dimensional formulation for the hybrid Reynolds stress tensor τ^* : as the anisotropy components are non-dimensionalized with k_{RANS} from the modified k transport equation, the numerator and denominator could be minimizing these discrepancies. Nevertheless, it is not well understood why the tendency was so consistent as more LES information is injected, or why the second HIRANS model, which showed very small discrepancies of the different turbulent stresses in this area, was not able to achieve a similar behaviour. This metric should be taken with care, as the cell size in streamwise direction was of approximately $d_x \approx 0.08$, so the maximum variations achieved were of around a 5 cells distance.

With respect to the reattachment point, the second HIRANS model achieved a remarkable result, getting towards the LES location as the ξ_i factors are increased. For the first HIRANS model, the discrepancy in the different mean velocity components was also reproduced in the reattachment location. The jump observed between $\xi_i = 0.4$ and

$\xi_i = 0.5$ was due to the presence of a secondary recirculation region on the windward hill, which merged with the main one as the injection factors were increased, creating a large recirculation region between the two hills.

The global field data discrepancy can also be quantified. In order to perform that, the different norms can be defined,

$$L^1\{y_{RANS}\} = \frac{\sum_{i=1}^N |y_{RANS} - y_{LES}| V_i}{\sum_{i=1}^N |y_{LES}| V_i}, \quad (3.19a)$$

$$L^2\{y_{RANS}\} = \sqrt{\frac{\sum_{i=1}^N (y_{RANS} - y_{LES})^2 V_i}{\sum_{i=1}^N (y_{LES})^2 V_i}}, \quad (3.19b)$$

$$L^\infty\{y_{RANS}\} = \frac{\sup |y_{RANS} - y_{LES}|}{\sup |y_{LES}|}. \quad (3.19c)$$

In these expressions, y_{RANS} and y_{LES} are any flow quantity, and V_i is the local cell volume. The norms are non-dimensionalized with the LES results and integrated for each cell, to minimize the effect of the mesh and the scaling of the variable to study, excepting variables along a line, as the wall shear stress, which are integrated along the line segment and not the cell volume, to avoid a weighting based on the scaling on the normal direction.

The results for the different field variables are presented in Figure 3.19. It is observed how the first HIRANS model achieved dubious results: while there were certain norms that are indeed reduced (especially the ones related with $\langle v'v' \rangle$), not all of them followed the same trajectory. The L^1 and L^2 norms for $\langle U_x \rangle$ and $\langle U_y \rangle$ increased from the baseline RANS model. This reinforces the stated hypothesis that, for this case, an improvement of the non-dimensional turbulent anisotropy does not necessarily need to be correlated with an improvement of the mean velocity components.

The second HIRANS model behaved as expected from the previous results, significantly lowering the majority of the norms throughout the ξ_i parameter space as these factors were increased. One exception escaped from this trend: the L^∞ norm for $\langle U_y \rangle$. A further scrutinisation of it reveals that the location of it was approximately the same for the different cases: the second cell close to the wall on the windward hill, around $x = 8.51 - 8.57$ for the different cases, shifting its location between two adjacent cells. As it has been already mentioned, the accelerating part on this second hill has proven to be challenging for this HIRANS model, not being able to significantly correct the wall shear stress and the velocity components very close to the wall there.

To further visualize the independent contribution of the two ξ_i parameters in the HIRANS model with τ_{LES} injection, Figure 3.20 shows the trend over different variables. It was observed that, in general, for the L^2 norms, the quantitatively dominant factor was the one corresponding to the injection of LES information in the momentum equation, ξ_m . For every variable which has been studied in this document, this data injection has introduced a significant reduction in its L^2 norm, achieving a value of around 20-25% of the baseline

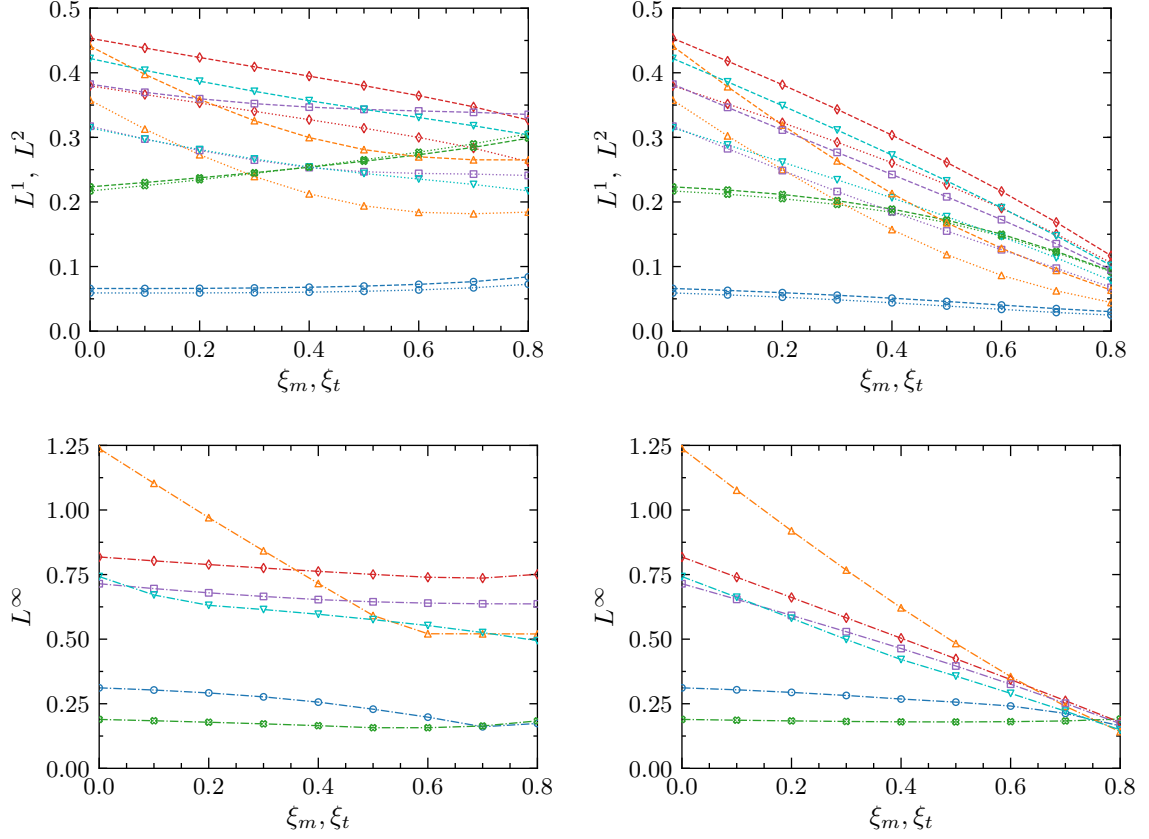


Figure 3.19: L^1 (.....), L^2 (- - - -) and L^∞ (- · - · -) norms of $\langle U_x \rangle$ (\circ), $\langle U_y \rangle$ (\diamond), $\langle u'u' \rangle$ (\odot), $\langle v'v' \rangle$ (\triangle), $\langle w'w' \rangle$ (\square) and $\langle u'v' \rangle$ (∇) of HIRANS with LES anisotropy injection (left) and HIRANS with τ_{LES} injection (right).

RANS values for the velocity components and the different shear stresses. The wall shear stress effect was less significant, as while there were certain regions, such as the recirculating part, which were greatly improved, the discrepancies in the accelerating part of the windward hill remained similar to the RANS case.

The effect of the hybrid tensor τ^* for the k transport equation was less straightforward. Depending on the field variable and norm, its effect became beneficial or detrimental. Globally, there was a worsening of the L^2 norms for $\langle U_x \rangle$ and $\langle U_y \rangle$, which mainly came from the cells close to the top wall. Nevertheless, its effect in the L^∞ norms was always positive, as well as in the wall shear stress norms. This suggested that introducing this factor improved the behaviour of the RANS model close to the walls, where the majority of the L^∞ norms are located (see Appendix A), but did not necessarily correlate with an improvement along the whole flowfield. This idea can be linked with the observation of the positive behaviour of the first HIRANS model in predicting a separation location closer to the LES model, using the same τ^* tensor.

Based on this, which injection factor should be chosen? The answer is, as many times in science and engineering: “depends”. Concerning the momentum equation factor ξ_m , keeping it as high as possible will likely result in the best agreement with the higher fidelity results, as the information is directly affecting the mean velocity components solving process. The turbulent transport equations factor ξ_t does not have such a straightforward influence, and it has been shown that, depending on the quantity of study, the correlation with the LES results can be better or worse when increasing it. It is also reasonable to expect this trend to be case-dependent. A compromise between all the different metrics concluded in utilizing $\xi_m = 0.8$, $\xi_t = 0.3$ for the rest of this work.

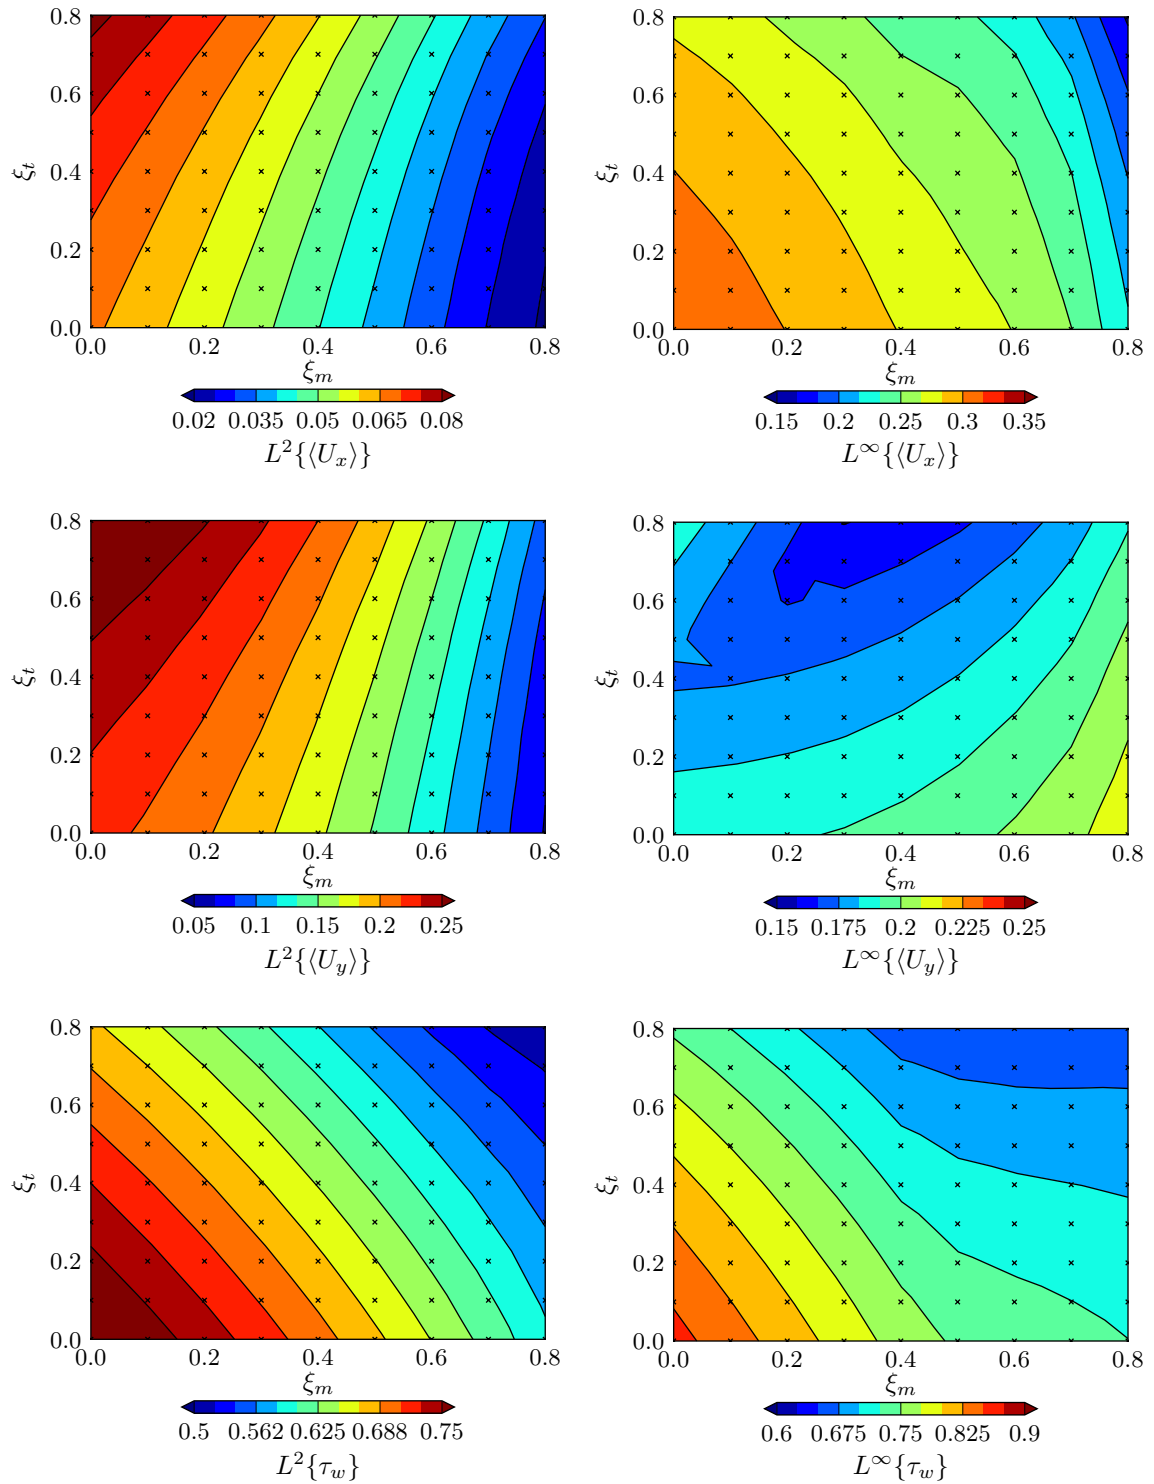


Figure 3.20: L^2 (left) and L^∞ (right) norms of $\langle U_x \rangle$ (top) and $\langle U_y \rangle$ (centre) and τ_w (bottom) for HIRANS with τ_{LES} injection. Contours linearly interpolated from 81 sample points (\times).

3.4 Summary

In this chapter, the formulation and results from the baseline $k - \omega$ RANS model and the two different HIRANS models, with injected anisotropy information and the full LES Reynolds stress tensor, have been presented.

For the baseline RANS model, the discrepancies with respect to the reference LES case were manifested in both mean velocity components and stresses. The Boussinesq's hypothesis applied for the turbulence model tended to underpredict the magnitude of the different stresses, both isotropic and anisotropic contributions, and created a very large recirculating region.

The correction of the anisotropy by the first HIRANS model was successfully achieved, but this was not translated to an overall better distribution of the mean velocity components in the flowfield. The isotropic component k was qualitatively unaffected, and as the constructed hybrid tensor τ^* was scaled by this value of k_{RANS} , there was not a great departure from the baseline RANS model.

On the other hand, the injection of the full τ_{LES} fulfilled the objective: the differences on the different stresses components with respect to the LES values became almost negligible quantitatively for the purpose of this study in the majority of the flowfield (within $\pm 1\%$ of the scaled quantities with U_b^2). There were some small region exceptions, located at the separation region for $\langle u'u' \rangle$, the top of the windward hill for $\langle v'v' \rangle$ and the foot of it for $\langle w'w' \rangle$. In all those cases, though, the differences were still smaller than the baseline RANS case. Furthermore, the wall shear stress distribution on the lower wall was much better represented, capturing a virtually identical reattachment point as the LES solution, but still showing an underprediction of the high shear region on the top of the second hill and of the negative shear on the flat recirculating part.

Furthermore, from the results of the barycentric map for the eigenvalues of the Reynolds stress tensor for the different cases, it was extracted that a correction metric based on the anisotropy of the flow and its resemblance to a higher fidelity solution may not always be correlated with a better prediction of other quantities within the flowfield.

The positive results from the second HIRANS model justify its application for the next part of this document, where sparse LES cases are interpolated into non-locally corrected (in terms of geometry) RANS cases, and compared with LES references and locally corrected ones, as well as the baseline $k - \omega$ model for those geometries.

LES setup and analysis

4.1 Introduction and LES setup

In order to validate the HIRANS framework for different geometries, a higher fidelity LES code must be used. In this chapter, a brief description of the different configurations which have been studied is given, to introduce a reliable, yet computationally feasible setup for the comparison with the HIRANS results and to produce several samples for the Kriging and Co-Kriging optimization routines.

While the LES solver INCA has already been tested for the reference periodic hill geometry (Hickel et al., 2008), several modifications have been implemented during the course of this work to increase its computational efficiency. The simulations have been performed in the TU Delft HPC, with 20 cores Intel(R) Xeon(R) CPU E5-2670v2 at 2.5 GHz and 128 GB of RAM.

The results of these simulations are compared with the LES results from Breuer et al. (2009), used as a reference to validate the HIRANS framework, and with PIV experimental data from Rapp (2009).

4.2 Computational performance

Despite the moderately low Reynolds number of the simulations ($Re = 10595$), there is a rich variety of structures which make the LES simulations particularly challenging in terms of the computational effort needed to achieve accurate results. Investigations from Fröhlich et al. (2005) and Breuer et al. (2009) show this fact, where the former remarks the high level of turbulent intensity in the spanwise direction for the post-reattachment region, confirmed by the features also captured by the latter author in a wide range of Reynolds numbers. Thus, a high streamwise and normal resolution is in principle needed to capture the shear layer at the separation location, and a relatively large domain of around $L_z = 4.5h$ is recommended by Fröhlich et al. (2005), with an averaging period of 55

flow-through times. Breuer et al. (2009) employed an even larger averaging period of 140 flow-through times. The summation of all of these effects introduces several constraints for an accurate LES simulation, and a practically infeasible configuration for the available resources and the multiple samples needed for the study.

There are two main ways to tackle this problem: increasing the computational performance of the solver and/or accepting a trade-off between accuracy and simulation time.

With respect to the first argument, the main bottleneck found during the course of this study has been the Poisson solver for the pressure, spending well more than 50% of the computational time in the TU Delft cluster arrangement. Using this knowledge, different state-of-the-art iterative solvers were tested: from V-cycle and full multigrid configurations to Fast Fourier Transform and Krylov subspace methods such as BiCGStab and GMRES, with further preconditioning given by Gauss-Seidel and multigrid. While some of the methods were slightly faster than others, being BiCGStab with an algebraic multigrid preconditioner generally the superior one, the order of magnitude of their speed were qualitatively the same. It is important to note that, due to memory requirements, the GMRES computations were always restarted, damaging its convergence properties. For the finer grids, it even failed to converge if applied for the first time steps. Starting the simulation with BiCGStab and continuing with GMRES later on still did not deliver a significant boost of performance. Some other methods were implemented, such as GCR (Vuik and Lahaye, 2012) and Recursive GMRES, GMRESR, (Van der Vorst and Vuik, 1994). However, for the test case, they resulted in a lack of a speed improvement and convergence, respectively.

After all these unsuccessful attempts, one of the methods tested finally achieved its goal: the Stone's Strongly Implicit Procedure, SIP, (Stone, 1968). Its formulation can be found in Appendix E. In the present document, this method was used for preconditioning the Poisson equation system, which was then solved with BiCGStab. This allowed to accelerate the code around five to six times in the test case compared to full and V-cycles multigrid, Jacobi and Gauss-Seidel preconditioners, offering a consistently similar performance on further computations, allowing a greater variety of mesh resolution and averaging time tests. Incomplete LU factorization was unavailable at the time of the computations for the INCA solver, which did not allow a direct comparison with it. It is believed that the periodic boundary conditions combined with the immersed boundaries first-order reconstruction were the main reason of the poor conditioning of the system. Indeed, the latter part of it was significant in the stability of the system, which was not able to converge when a second-order reconstruction of the hill shape was attempted. In terms of the conditioning of the preconditioned system, while the multigrid methods were able to provide a significant reduction of Poisson iterations with respect to Jacobi or Gauss-Seidel, their cost per iteration was significantly higher, which overall only slightly increased the computational performance. The SIP method was comparable in terms of overall Poisson iterations, but the large reduction of cost per iteration was the main responsible of the improvement in CPU time.

4.3 Effect of spanwise extent, mesh resolution and averaging time

The INCA LES implementation is compared with reference LES and experimental data. The two main LES works of Fröhlich et al. (2005) and Breuer et al. (2009) are both with curvilinear grids, with a number of interior cells of $N_x \times N_y \times N_z = 196 \times 128 \times 186$ for the first case, and an increase of the resolution up to around $298 \times 220 \times 200$ for the second one, allowing for a DNS-like representation of the upper wall in this case. Due to the readily available field data, the comparison in this and the next sections is performed with the latter simulation as the reference data. Three main parameters are investigated: spanwise extent, mesh resolution and averaging time. Unless explicitly specified, the averaging time considered is 33 flow-through times, starting the samples at 22 flow-through times, with the flow-through time defined as $t_{\text{flow-through}} = x_{\text{max}}/U_b$.

4.3.1 Spanwise extent

Case	L_z	N_x	N_y	N_z	N
Full	4.5	140	280	144	$5.6 \cdot 10^6$
Semi	1.5	140	280	48	$1.2 \cdot 10^6$
Narrow	0.5	140	280	14	$3.5 \cdot 10^5$
Breuer et al. (2009)	4.5	298	220	200	$1.3 \cdot 10^7$

Table 4.1: Grid properties of the different cases for the spanwise extent study.

Using the recommendation of Fröhlich et al. (2005), for both LES and hybrid RANS-LES results, while the correlation coefficients as a function of the separation d_z for the turbulent fluctuations are non-zero for $L_z = 4.5$, the error introduced is estimated to be small. Longer spanwise extents would provide little quantitative benefits for the extra cost needed for their simulation. As there exist wide vortical structures, reducing the spanwise domain length would hurt the averaging representativeness.

To confirm it, and to evaluate the errors with respect to the baseline extent of $4.5h$, in the aid of bringing down the computational cost, three different configurations are studied, as shown in Table 4.1, with noticeable reductions in computational requirements.

Nevertheless, Figure 4.1 represents a clear picture: reducing the spanwise extent rapidly reduced the accuracy of the LES computations. The recirculation region was wrongly calculated even in the “semi” case, becoming larger as the domain gets narrower. This was in agreement with the discussion in Fröhlich et al. (2005), following the data of Klostermeier (2001) and the dissertation of Fröhlich (1990). There, it is stated that only the very largest structures get “squeezed” when narrowing the domain size, which could be partially the case of the shear layer region, damaging the spanwise average if the domain is not large enough. The 2-point correlations in Klostermeier (2001) suggest that structures close to the upper wall and the separation region have a considerable size in the spanwise direction, the latter with an approximate size of 3 hill heights. Therefore, not even this “semi” case could properly represent some of the flow features. Figure 4.2

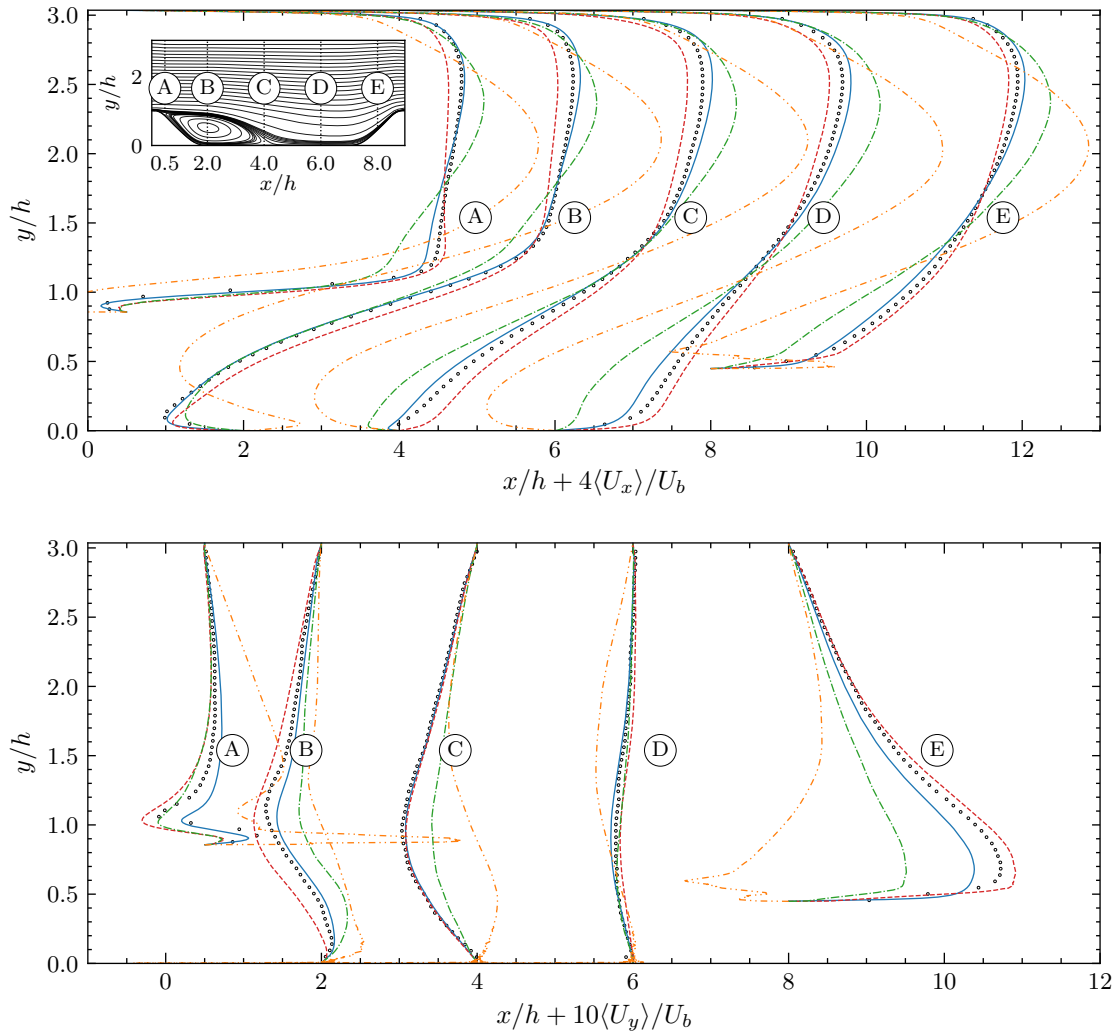


Figure 4.1: Mean velocity components at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), full (- - -), semi (- · - ·), and narrow (·····) INCA cases.

shows the large underprediction of the different turbulent stress components. The narrow case, in every single variable and location, showed a completely wrong pattern, with very high turbulent stresses probably coming from non-fully represented structures.

The simulation performed for the “full” case showed a much closer agreement with respect to the reference LES and PIV data, despite its considerably shorter averaging time, 33 flow-through times for the INCA simulations versus 140 for the reference data. The recirculation region was underpredicted, which, due to the conservation of momentum within the periodic domain, is reflected in a lower momentum close to the upper wall, and the turbulent stresses show a similar shape as the reference data, being qualitatively similar to both LES and experimental cases.

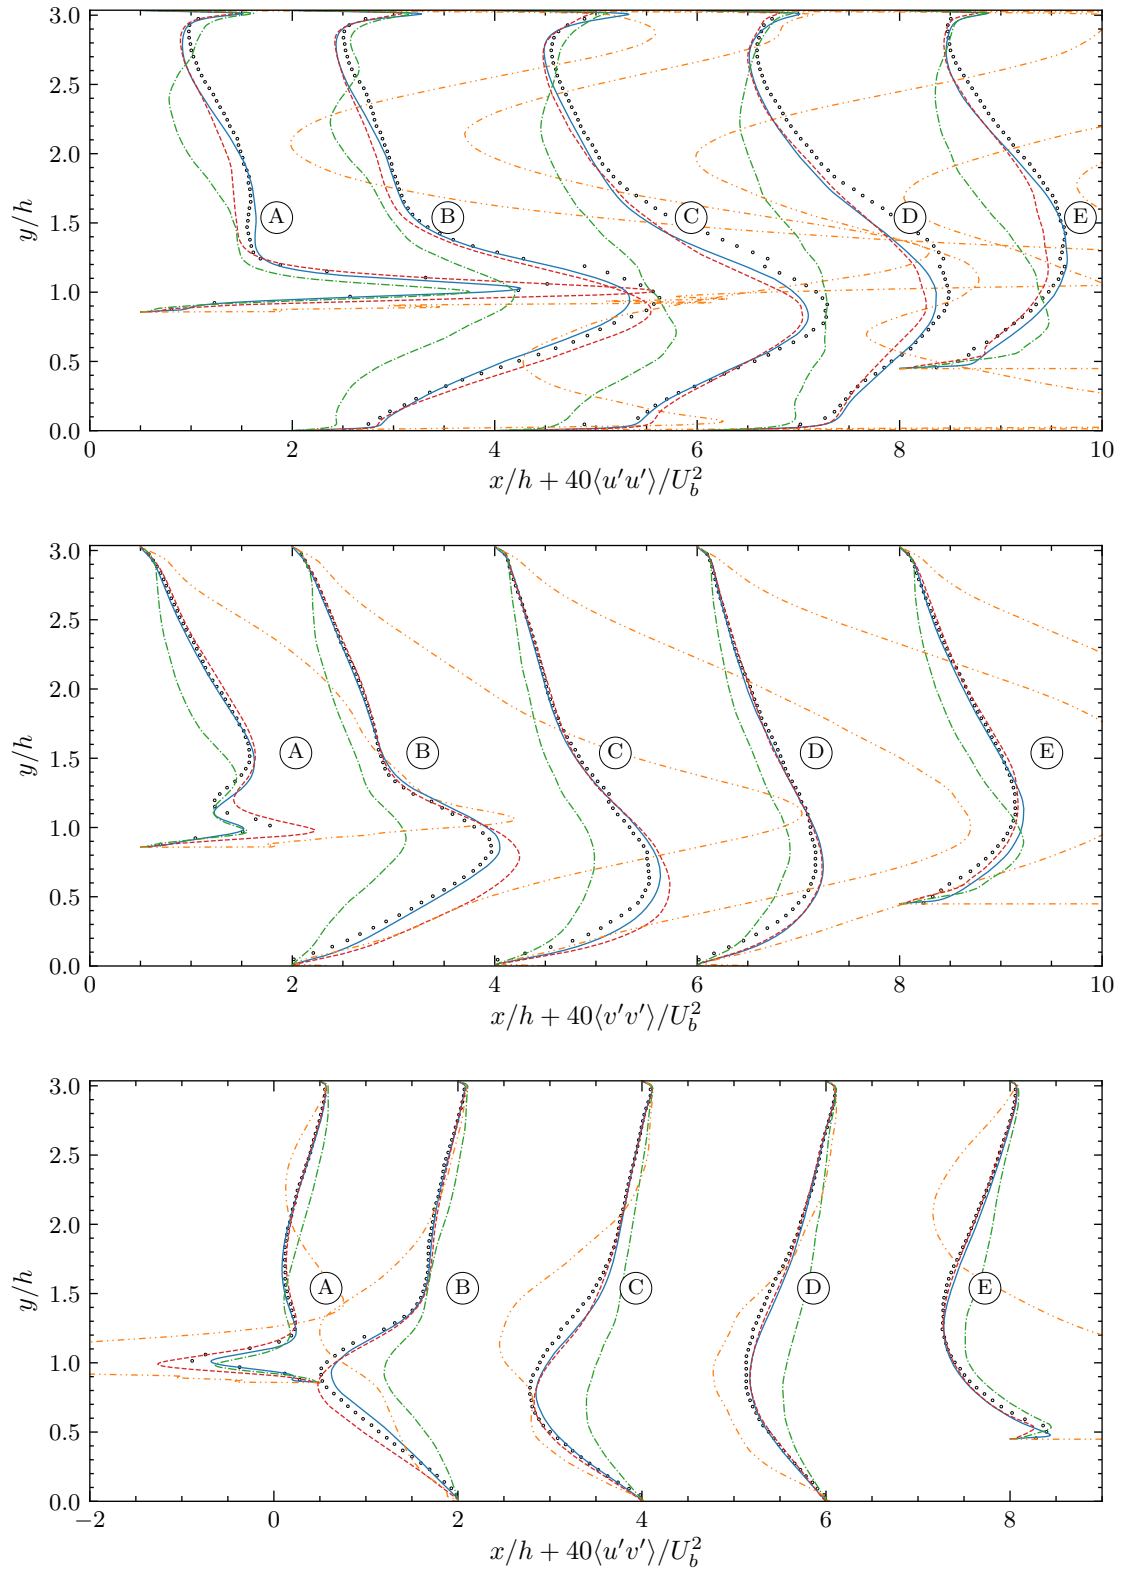


Figure 4.2: Turbulent stresses at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), full (- - -), semi (- · - ·) INCA cases (experimental data not available for $\langle w'w' \rangle$).




Figure 4.3: Mesh for fine (top) and coarse (bottom) cases.

4.3.2 Grid resolution

The previous section showed the need to perform a simulation over the recommended spanwise extent of $L_z = 4.5h$. However, the INCA case with the fine grid resolution considered is computationally too expensive to produce several samples for an optimization routine in this study, with an approximate cost of around 20 CPU days per case. An alternative grid with a coarser resolution principally in the normal and spanwise directions is also tested, with almost a fifth of the nodes of the finer mesh, as shown in Table 4.2.

Method	n_x	n_y	n_z	N	$y_{bottom\ wall}^+$
Fine mesh	140	280	144	$5.6 \cdot 10^6$	$< 1 - 14$
Coarse mesh	128	100	96	$1.2 \cdot 10^6$	$< 1 - 20$
Breuer et al. (2009)	298	220	200	$1.3 \cdot 10^7$	$< 1 - 1.2$
RANS/HIRANS	120	130	-	15600	$< 1 - 1.5$

Table 4.2: Grid properties of the different cases for the grid resolution study.

Figure 4.3 shows the key differences between the two grids. While the finer one has a much better resolution overall, it is particularly refined in the separation region and both bottom and top walls. As observed in Table 4.2, the Cartesian grids for the INCA

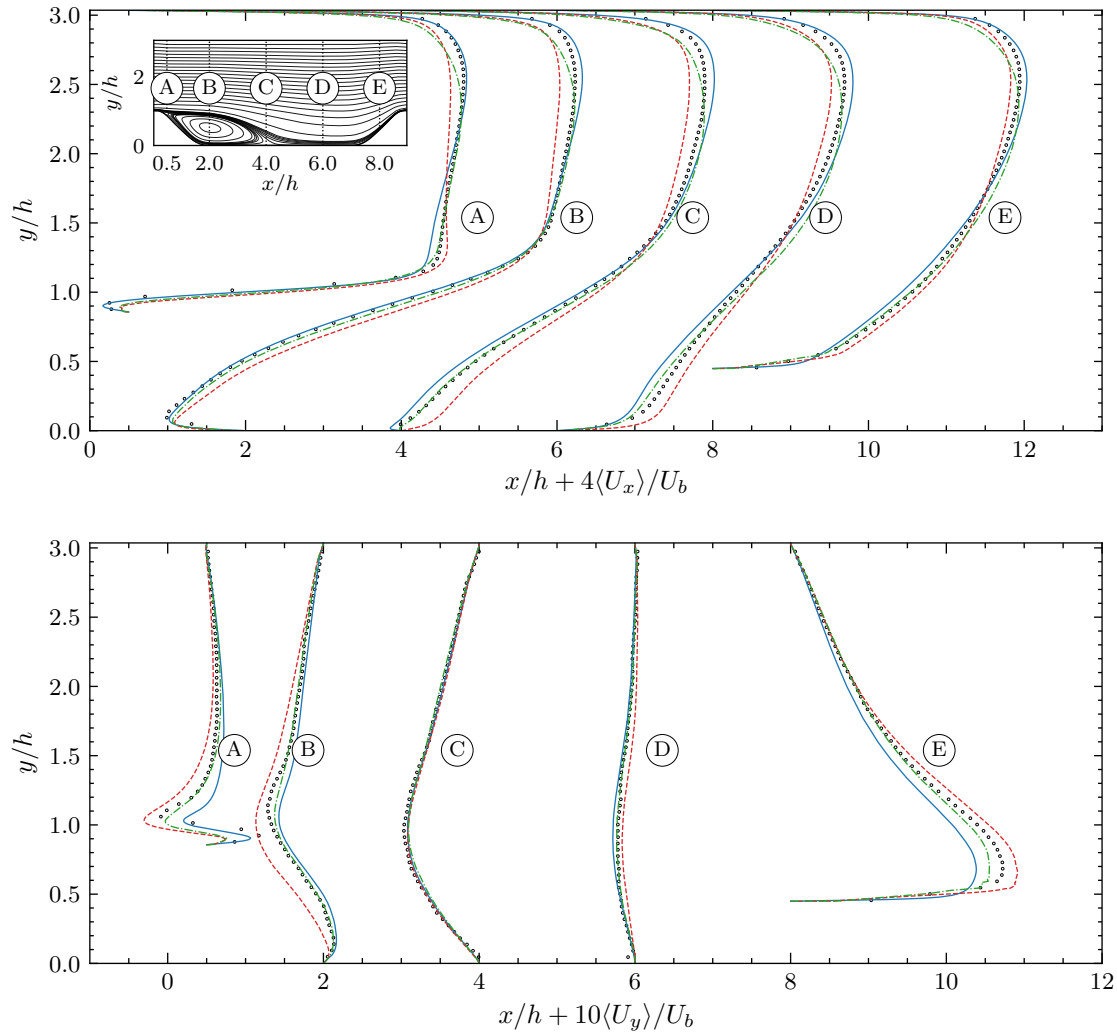


Figure 4.4: Mean velocity components at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), fine (- - -), and coarse (- · - ·) INCA cases.

cases struggle with the near-wall resolution, with relatively high peak y^+ values close to the highly curved, not streamwise aligned regions of the bottom wall, compared to the body-fitted grids of Breuer et al. (2009) and RANS. Nevertheless, the values are still within the range for adequate wall function treatment. The coarser mesh keeps a good resolution near the separation region, but sacrifices it in the rest of the domain. Despite this coarsening, the aspect ratio of the cells is improved, as in the finer case they become very elongated in the streamwise direction due to the very high resolution in the normal one.

The results of the grid coarsening for the mean velocity components can be observed in Figure 4.4. Starting from the streamwise velocity results, there seemed to be a much better prediction of the recirculation extent with the fine mesh. Indeed, both lines at $x/h = 4.0$ and $x/h = 6.0$ had a very close agreement with the experimental data, being

virtually identical on the first case. Not only that, but the prediction on the lower part of the domain showed excellent agreement, with a slight deficit of momentum at around $y/h = 1.0 - 2.0$, which resulted in a worse prediction of the top wall. The discrepancy with Breuer et al. (2009) in this upper region can be attributed to the fact that this reference had a DNS-like representation there, while, as it was observed in Figure 4.3, the resolution in the normal direction between $y/h = 2.0 - 3.0$ in the INCA coarse case was not so high, with around an order of magnitude of difference in y^+ . The results for the vertical component of the mean velocity were also remarkably good, achieving an arguably similar or better result in the majority of the domain, excepting the near-wall separation region.

While the mean velocity components were very satisfactory, the turbulent stresses were, in general, too. As Figure 4.5 shows, the overprediction of the streamwise velocity fluctuations in the shear layer region after separation became slightly smaller and closer to the experimental data, following this trend with increasing x/h . The region near the top wall did not follow the trends of the rest of the simulations. It is suggested that this part of the domain had not completely developed for this case, and its effect is further analysed in Section 4.3.3. Despite this result, the rest of the domain seemed to follow the expected shape of the finer grid. With respect to $\langle v'v' \rangle$, it was also observed how the exaggerated peak in the beginning of the shear layer followed much closer the reference LES data, and a similar result occurred for $\langle u'v' \rangle$. Nevertheless, it must be noted that, for these last two components, the results from Breuer et al. (2009) showed a closer agreement with the PIV data.

Overall, the mesh coarsening does not seem to be a step back, but even a step forward. It is somewhat surprising, as it would be expected that the finer grid achieved more accurate results, especially taking into account the evident lack of full development of the solution on the top wall for the coarser case. Nevertheless, the next part of this document, Section 4.3.3, shows that the averaging time for these simulations suggest that the case is not fully developed until well beyond the calculated period here. Hence, the error due to the lack of development and an excessively short averaging period could be greater than the discretization error in this case. The lack of accuracy of the fine mesh could be also affected by the high aspect ratio present in the elongated cells in streamwise direction, but conclusive results about this hypothesis would need additional research.

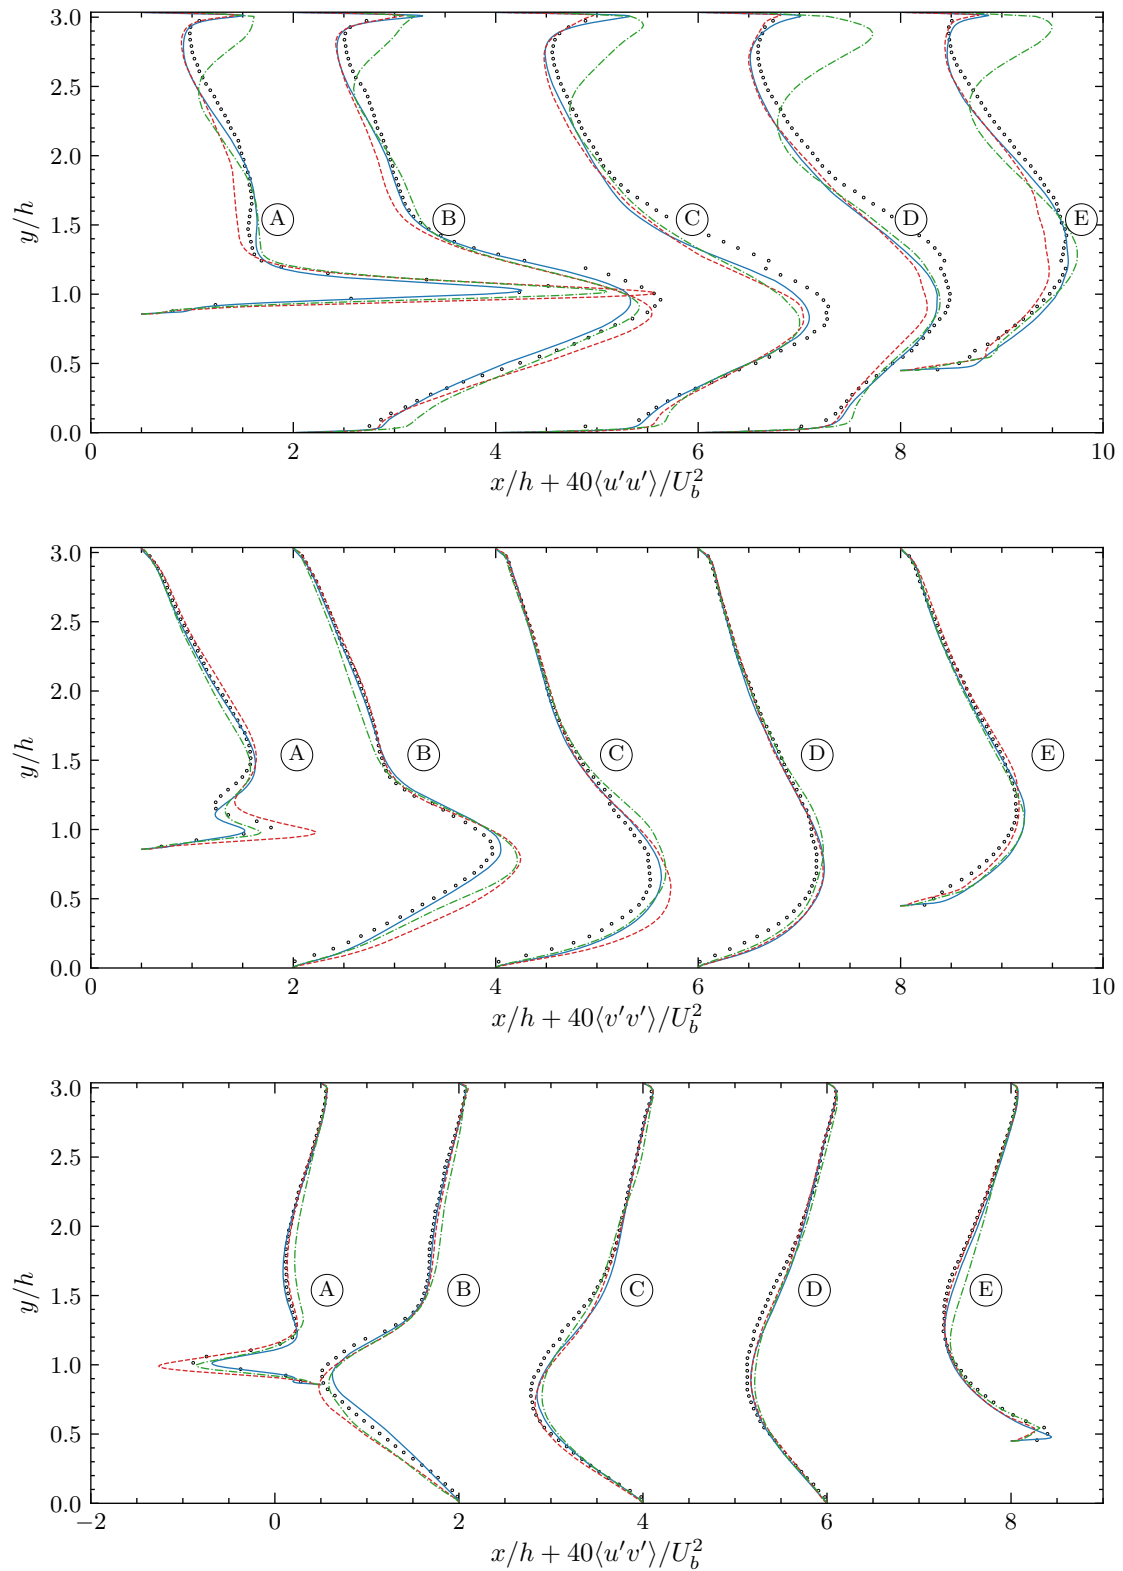


Figure 4.5: Turbulent stresses at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), fine (- - -), and coarse (- · - ·) INCA cases (experimental data not available for $\langle w'w' \rangle$).

4.3.3 Averaging period: start and extent

The results from the mesh resolution study show a potentially important issue with the simulations performed: having a properly developed flow over a sufficiently long averaging time. This section addresses both questions, in order to be able to produce reliable computations for the different geometries considered along the document.

Method	Starting averaging time	Averaging period
Baseline	22	33
Extended A	55	33
Extended B	55	55
Extended C	89	33
Extended D	89	55
Fröhlich et al. (2005)	23	55
Breuer et al. (2009)	-	140

Table 4.3: Test matrix of averaging period cases. Both starting time and averaging period are in flow-through times ($t_{\text{flow-through}} = x_{\text{max}}/U_b$). Values truncated to integer numbers.

Table 5.1 shows the different starting times and averaging periods tested. Due to the successful results of the coarse mesh and resource constraints, all the simulations were performed with it. As it is observed, both effects, averaging period and starting averaging time are considered. For the remaining part of this section, only results for the Extended A and Extended D simulations are shown. The rest of the cases are included in Appendix F.

To study the adequacy of the selected values for starting the averaging process and the length of the average period, two different variables are considered. The first one is the additional forcing term which is added to drive the channel at constant Re , f_p . This is added as a proportional-integral-derivative (PID) controller each time step, as

$$f_p = K_{static} \left(K_p e(t) + K_i \int_0^t e(t') dt' + K_d \frac{de(t)}{dt} \right), \quad (4.1)$$

where the $e(t)$ function refers to the bulk velocity over the domain at the instant i with respect to the target (solution) one, as $e(t) = (U_{b_{sol}} - U_{b_i})/U_{b_{sol}}$, and K_i are the different constants.

The second variable to study is a measure of the fluctuations of the velocity components in a root mean squared sense, u_{RMS} , as

$$u_{RMS} = \sqrt{\sum_{i=1}^3 (u_i - \bar{u}_i)^2}, \quad (4.2)$$

where \bar{u}_i is the averaged velocity component over the entire simulation.




Figure 4.6: Volumetric forcing (top) and root-mean-square velocity (bottom) for the Extended D simulation (left) and averaging period (right). The mean quantities over the averaging period are represented by (- - - -).

Figure 4.6 shows the development of these quantities over the course of the Extended D simulation, the longest one from the tests. In these figures, the initial transient state can be clearly observed in the u_{RMS} variable. The fluctuations in the first flow-through times started to decay and seem to achieve a value which oscillates around $u_{RMS} \approx 0.25$. Looking at the f_p term, it could be argued that both metrics seemed valid to start the averaging period as soon as 25 flow-through times. Nevertheless, the overall picture is not so simple: around 75 flow-through times, there was an increase of the u_{RMS} over the previously predicted value, which could lead to a misleading average if this one was performed between 75 and 100 flow-through times. Based on this observation, low frequency oscillations may play an important role in the dynamics of the system, and, looking at u_{RMS} , it seems that after around 85 flow-through times, the system started to decrease the magnitude of these peaks.

A further analysis applying a Fast Fourier Transform (FFT) to f_p and u_{RMS} is shown in Figure 4.7. It is observed how, for both variables, for earlier times, $t_{\text{flow-through}} = [22, 55]$, the Extended D simulation showed a very low frequency oscillation at around 6×10^{-2} flow-through times. This component was dampened as the simulation progresses and it




Figure 4.7: FFT for the volumetric forcing (left) and root-mean-square velocity (right) for Extended D, flow-through-times 89-144 (—), Extended D, flow-through-times 22-55 (- - -), fine mesh, flow-through-times 22-55 (- · - ·). Reference slopes (······).

did not appear in the converged solution. Furthermore, the fine mesh did not show this behaviour in the averaging period. This oscillation could be the cause for the upper wall non-converged behaviour in $\langle u'u' \rangle$ in the coarse mesh, not present in the finer one. Looking at the spectrum, while the fine mesh starts to resemble the converged solution, it still needed further development. The f_p plot showed a peak which progressed in the direction towards a single flow-through period for the converged Extended D case, while the u_{RMS} showed it at approximately one order of magnitude lower in frequency. This oscillation was likely related with the Kelvin-Helmholtz instability present in the shear layer after separation. This information motivated the reasoning to run the simulation well beyond the original value of 55 flow-through times. Not only that, but it also suggested that, with the current forcing in the initial time step of perturbations of 0.32% of U_b in the streamwise velocity component, a transient state significantly longer than the 23 flow-through times present in Fröhlich et al. (2005) is preferred to be calculated before starting the averaging period. Based on the frequency of the different oscillations present, it is also inferred that this longer averaging period could be beneficial in the calculation of a robust solution for different geometries. It is not trivial to understand a priori how the oscillations of the transient system for the different geometries will be until the simulations have already been run, which is not practical within the optimization routine framework.

Figure 4.9 shows the mean velocity components behaviour for the baseline simulation, the effect of starting averaging later (Extended A) and both the effects of a delayed start and a longer averaging period (Extended D). Starting with $\langle U_x \rangle$, the greatest differences were located in the central and upper parts of the domain for every x/h station considered. Despite the relatively coarse mesh, which was around 10% of the size of Breuer et al. (2009), the results for the extended simulations, especially for Extended D, were strikingly similar to the PIV reference in a large part of the domain and even closer than the LES reference. The only exception was found near the separation region. This could be related with the first-order reconstruction in the immersed boundaries setting for the INCA LES cases and the relatively high y^+ values in the curved regions of the bottom wall. This




Figure 4.8: Mean and standard deviation of L^2 (- - - -) and L^∞ (- - - - -) norms computed against PIV data over the different x/h stations of $\langle U_x \rangle$ (\circ), $\langle U_y \rangle$ (\diamond), $\langle u'u' \rangle$ (\triangle), $\langle v'v' \rangle$ (∇) for the different cases.

discrepancy, as explained before, also affects the top wall results due to the boundary conditions imposed. The vertical velocity component seemed to be hardly affected in the extended simulations, only being noticeable in the windward hill, with an increase of $\langle U_y \rangle$ over the baseline simulation, showing the convergence of the results to the PIV solution as it is extended.

The turbulent stresses showed a much clearer trend in Figure 4.10. For $\langle u'u' \rangle$, the results significantly differed from the baseline solutions, and also seemed to converge when extending the simulation. The Extended A computation was a major alteration of the baseline results, correcting the wrong shape near the top wall, but not necessarily bringing them closer to the PIV reference in some parts of the domain. Notwithstanding, both extended simulations showed similar shapes qualitatively to the LES reference. A similar conclusion can be drawn of the rest of the components, $\langle v'v' \rangle$ and $\langle u'v' \rangle$: they seemed to converge at every station quantitatively, but for values that differed more to the PIV reference than for the mean velocity components. This was not an issue only related with these simulations, as Breuer et al. (2009) achieved similar results qualitatively.

To quantify the magnitude of the discrepancy with respect to the PIV results, L^2 and L^∞ norms were calculated between the different cases and the PIV data, using 3.19 and shown in Figure 4.8. The comparison was performed using all x/h lines available in the PIV data, calculating the pointwise discrepancy with respect to the LES solution (interpolated in the PIV points) and averaging these discrepancies over all the x/h stations. The conclusions from these figures are closely related to the ones drawn from the previous analysis. The baseline simulation was a significant departure of the extended cases, which seemed to be very similar between each other in the majority of the flow variables considered. The only exception was found in $\langle u'u' \rangle$, which was highly affected by the fluctuating shear layer and recirculation regions, suggesting that a fully representative average of this variable would

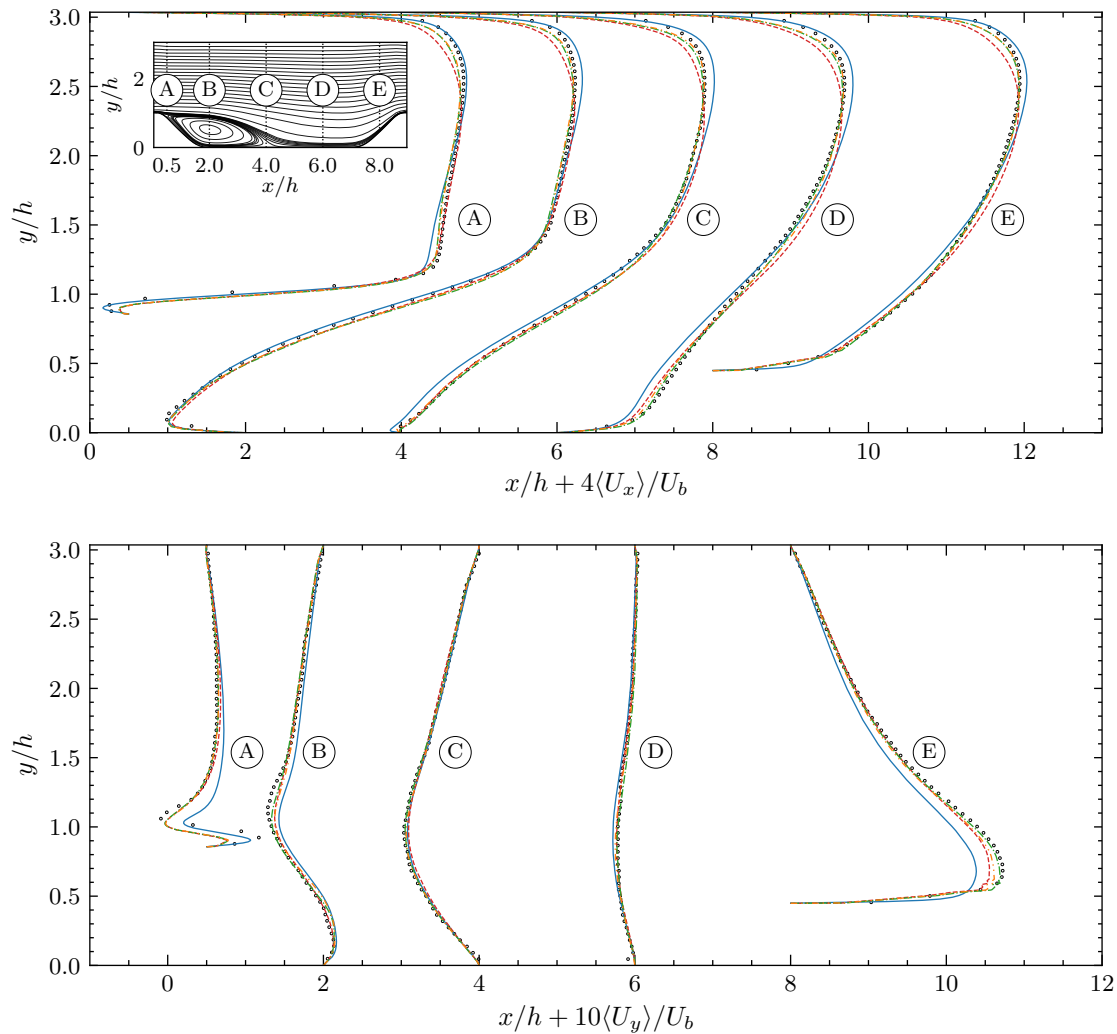


Figure 4.9: Mean velocity components at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), baseline (- - -), Extended A (- · - ·) and Extended D (- · · -) INCA cases.

need a longer averaging period. Nevertheless, it is noted that for every norm, the mean results for the Extended D simulation were comparatively roughly as good compared to PIV data as the ones achieved in Breuer et al. (2009), with a notable emphasis in the accurate characterization of the mean streamwise velocity component. With respect to the standard deviation, similar results can be drawn, with a particularly low value for the $\langle u'u' \rangle$ component in the Extended D case.

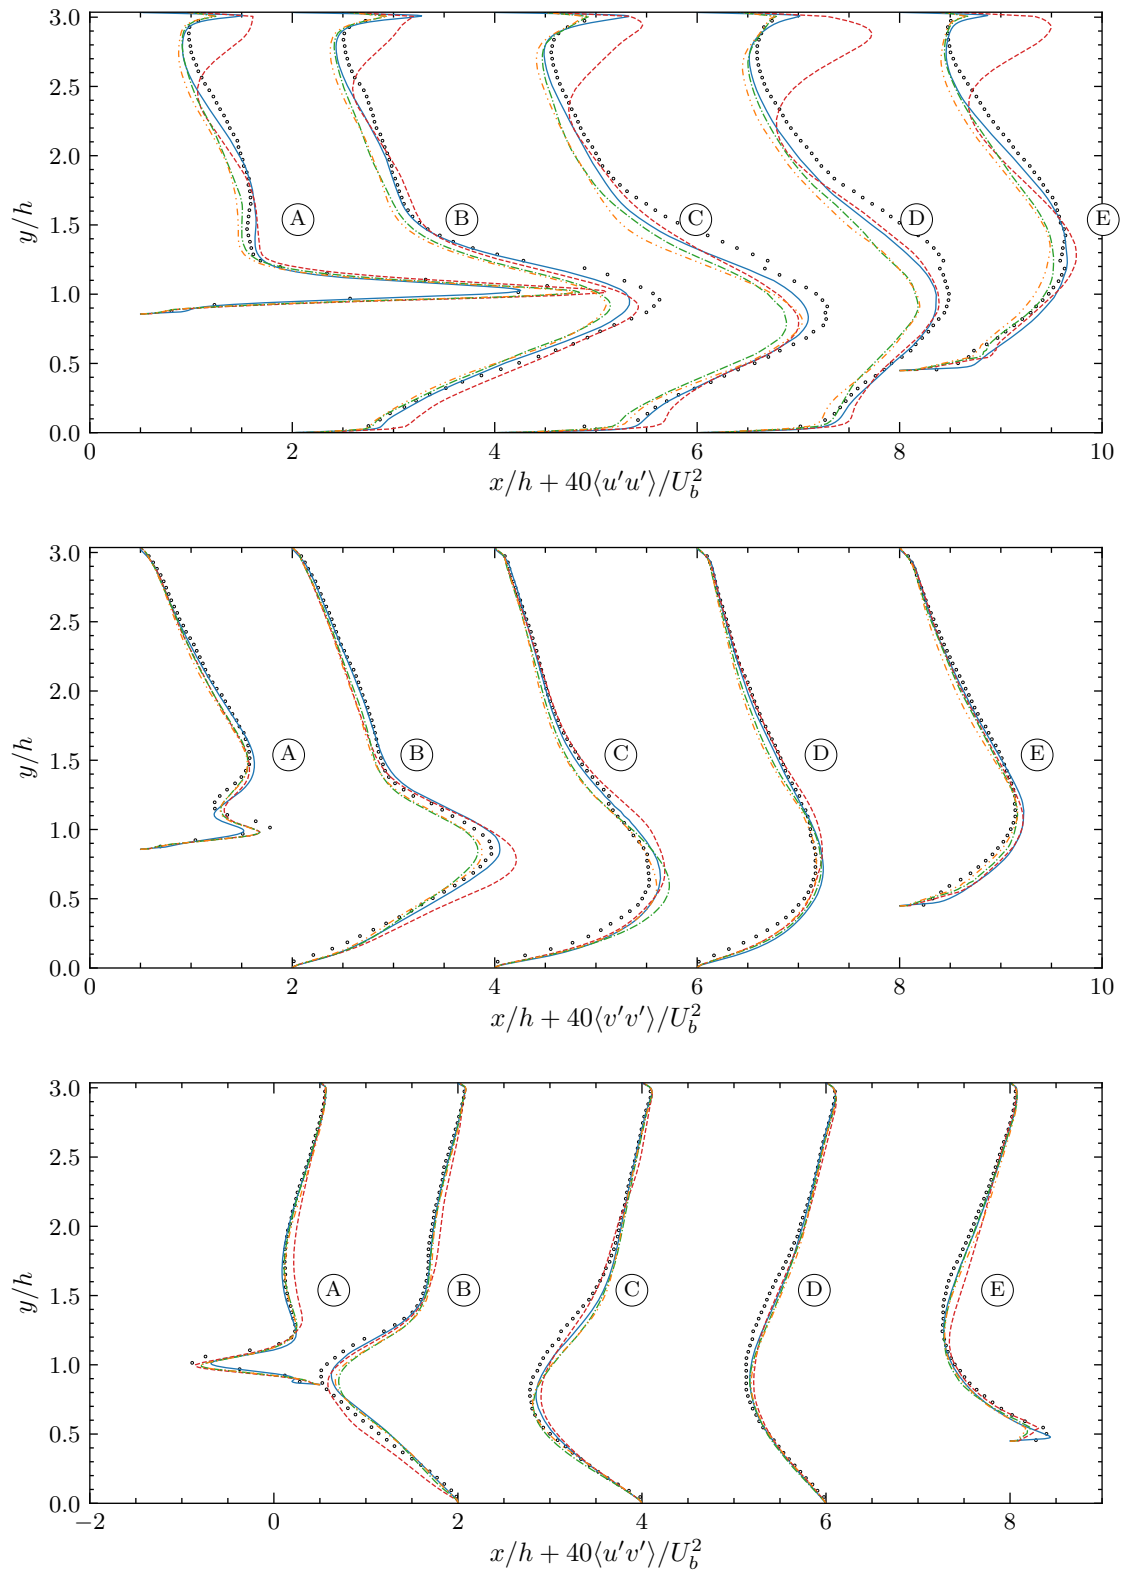


Figure 4.10: Turbulent stresses at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), baseline (- - -), Extended A (- · - ·) and Extended D (···) INCA cases.

4.4 Configuration choice and summary

The LES setup has established itself as a challenging part of the current study, even for the simple geometry and moderate Reynolds number considered. The constraints in computational resources available and the accuracy required to be used for the HIRANS model determined the testing of different strategies until finding the final configuration.

Starting from the computational performance, the main limitation was found in the Poisson solver, which was tackled using several methods, such as multigrid, Krylov subspace solvers or Fast Fourier Transform. The introduction of the Stone's Strongly Implicit Procedure allowed to accelerate the overall computational time around five to six times by decreasing the iteration time with respect to multigrid methods, giving also the possibility to test finer grids.

To reduce even further the computational cost, narrower domains in the spanwise direction were considered. Neither the narrower or even the "semi" domains were successful in any way. The large flow structures constrained the restriction of the spanwise domain, to at least an extent of $L_z = 4.5h$ between the three considered cases.

The grid resolution study revealed that even a coarser mesh of around one million cells was able to capture the main dynamics of the flow and be qualitatively representative in comparison with LES and PIV reference data. Coarsening the mesh resolution also affected greatly the recirculation length, which was unexpectedly better represented in the coarser case.

The most fruitful part was the averaging period study: it detected that, for the coarser mesh case and the tested forcing term in the initial time step, the simulation needed as long as at least 55 to 89 flow-through times to achieve better converged statistics. Increasing the averaging period could also become beneficial if low-frequency oscillations are present in different geometries, such as steeper ones with several recirculating regions. In addition, this effect could also explain why the finer grid achieved a worse performance compared to reference data, based on the balance between discretization and convergence errors. Extending the simulation showed an excellent agreement in the mean velocity components with respect to PIV data, being also positive in the turbulent stresses, with similar discrepancy levels as [Breuer et al. \(2009\)](#).

Due to the time constraints of this work, solving a simulation with the finer grid for a longer amount of time was not studied. With the current configuration, solving up to 55 flow-through times took around 15-20 CPU days in the TU Delft cluster arrangement, compared to the approximately 2 CPU days of the coarser mesh for the Extended D case, which ran until around 144 flow-through times. Triplicating the needed time for a single configuration in the fine mesh is out of the scope of this document, as it is focused on calculating several samples within the design optimization framework.

Thus, the final configuration, which is then translated to the different geometries, is the Extended D setup: spanwise extent of $L_z = 4.5h$, 1.2×10^6 grid points and a total running period of 144 flow-through times. It provides a robust, yet not prohibitively expensive framework which can be used to acquire several samples in the different geometries, interpolation and design optimization parts of this work.

HIRANS for different periodic hill geometries

5.1 Introduction

The HIRANS framework with τ_{LES} injection has proven successful in correcting the mean velocity components and turbulent stresses in the periodic hill test. However, in the context of a multi-fidelity design optimization routine, it would be desirable to be able to compute a large number of HIRANS samples with a significantly lower number of LES cases.

This chapter addresses how well does the hybrid injected RANS formulation adapt to variations of the original periodic hill geometry with varying flow features. Firstly, the periodic hill is parametrized using a single degree of freedom. After this, RANS, LES and HIRANS cases are run for them, to compare the capabilities of the model with local injection from the same geometry.

Furthermore, the *prediction* capabilities of the framework are analysed, in order to provide insight of how accurate the different flow magnitudes can be with respect to the LES information when no local data is available, and interpolation between different geometries is needed.

A comparison with the known LES information, locally injected HIRANS and RANS is given, discussing the shortcomings of the method as long as potential pitfalls of it.

5.2 Local injection of LES data

This section involves the parametrization of the original geometry, as well as the different results for the local injection of LES data for 7 different cases.

5.2.1 Geometry definition and study cases

The baseline periodic hill geometry is defined in Almeida et al. (1993). The hill is divided in six segments with seven control points, with each segment defined as a fourth-order polynomial. This allows for a simple way of controlling different sections. To introduce a parameter space of multiple dimensions, a linear system can be constructed, fixing the x/h positions of the different control points and imposing continuity in the y/h locations, first and second derivatives. This results in a 24×24 system, which can be easily solved to obtain the new coefficients for the piecewise continuous function. From the multiple possibilities to parametrize this, the following one has been implemented, with varying x_i and y_i positions for the bottom wall coordinates,

$$x_i = \psi \left(\frac{x_{i0} + x_{i_{mult}}}{h} \right), \quad (5.1a)$$

$$y_i = \left(\frac{y_{i0} + y_{i_{mult}}}{h} \right), \quad (5.1b)$$

where the subscripts 0, *mult* refer to the baseline and the new multiplier values, respectively. Fixing the extrema, there are 5 interior points which, with varying $x_{i_{mult}}$, $y_{i_{mult}}$ and the width of the hill ψ , allow for up to 11 degrees of freedom. The complete derivation for the calculation of the system can be found in Appendix D.




Figure 5.1: Geometry for different hill width values:
baseline, $\psi = 1.0$ (—),
steeper cases, $\psi = (0.25, 0.5, 0.75)$, (- - -),
and less steep cases, $\psi = (2.0, 3.0, 4.0)$, (- · - ·).

In the context of this study, only the hill width factor ψ has been modified. A total of 7 geometries are tested in this chapter, from configurations significantly steeper than the reference one, $\psi < 1$, to the less steep ones, $\psi > 1$. Figure 5.1 shows that, even with this single degree of freedom, the geometry can become significantly different from the reference one. Doing so, it also involves different kinds of flow features, with little to no




Figure 5.2: Streamlines with streamwise velocity contours of LES cases for the steepest ($\psi = 0.25$, left) and least steep ($\psi = 4.0$, right) cases.

separation at all expected in the least steep configurations and secondary recirculating regions in the steepest ones. For some configurations with high values of ψ , the two hills even start to touch, joining in a single one with a sharp junction. Figure 5.2 shows how distant the flow features are in the two limiting cases. On the steepest one, $\psi = 0.25$, the leeward hill shows a secondary recirculation region, and the main one is significantly bigger than in the baseline case. Moreover, the windward hill exhibits its own reversed flow region close to its foot. On the other hand, the $\psi = 4.0$ case has attached flow along the entirety of the bottom wall.

5.2.2 Results

In a similar way as it was presented in Section 3.3.5, L^1 , L^2 and L^∞ norms for different flow variables are calculated to determine the difference between the RANS/HIRANS and LES solutions. It must be noted that all the results shown in this and the following sections for the HIRANS model correspond to the injection factors $\xi_m = 0.8$ and $\xi_t = 0.3$.

Figure 5.3 shows the performance of the HIRANS model compared to the baseline $k - \omega$ results. Despite having a much coarser LES solution to be injected from, which includes sections near to the walls with y_{LES}^+ significantly higher than its RANS counterpart, the model performed well in the majority of the envelope considered. Starting with the mean velocity components, there was a notable improvement in many points, which was especially reflected in locations where the baseline RANS differed significantly from the LES solution, such as $\psi = 2.0$. This did not always happen, though, as there were exceptions to it, which were the points with the best performance in the baseline RANS model. In this cases, such as $\psi = 3.0$ and $\psi = 4.0$, the norms even slightly increased. So, while the HIRANS model was not able to always decrease the difference between RANS and LES solutions for these field variables, it did keep them below a certain threshold for the whole envelope, achieving a much more *consistent* difference. This is highly relevant for the design optimization part, as the Co-Kriging model would have a much easier task

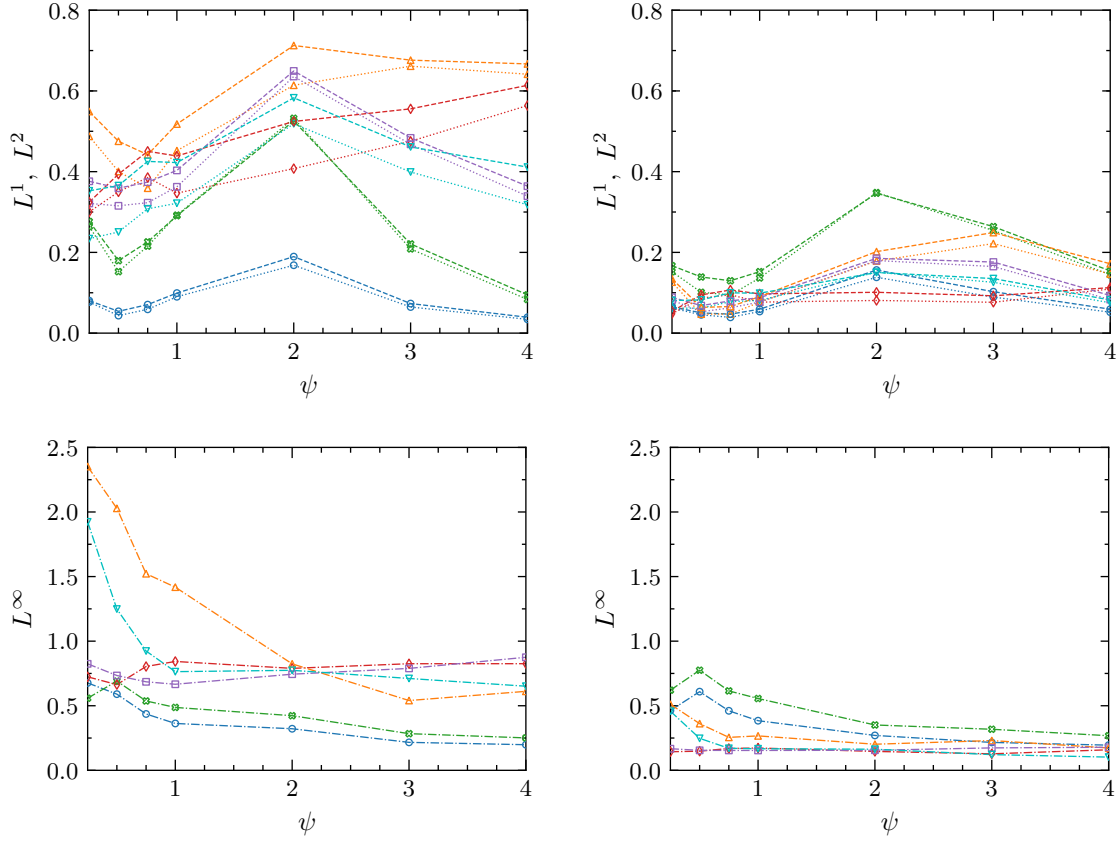


Figure 5.3: L^1 (.....), L^2 (- - - -) and L^∞ (- · - · - ·) norms for different geometries of $\langle U_x \rangle$ (\circ), $\langle U_y \rangle$ (\odot), $\langle u'u' \rangle$ (\diamond), $\langle v'v' \rangle$ (\triangle), $\langle w'w' \rangle$ (\square) and $\langle u'v' \rangle$ (∇) of RANS (left) and HIRANS with τ_{LES} injection (right).

of combining the information of low and high fidelity data if the differences between them are bounded more tightly.

With respect to the turbulent stresses, the results showed a vast improvement, for every single sample point considered. The injection of LES data allowed to severely correct the different components. As an example, the L^∞ norm for $\langle u'v' \rangle$ and $\psi = 0.25$ was reduced to around one fourth of its original magnitude, and the $\langle v'v' \rangle$, to one fifth in the same geometry.

For the mean velocity components, a great part of the L^∞ values were obtained in the highly accelerated part of the windward hill in a high curvature region, close to the top, and in the first 5 cells from the wall. The LES solution employed as an input for the injection process had, in some of these locations, a value of y^+ of around an order of magnitude higher of the HIRANS case. This created a large discrepancy between the true behaviour of the flow and the linear interpolation performed to inject the information from the LES cells. Alternative interpolation procedures close to the wall, following similar progressions as the physical development of the different quantities, could achieve a better match in these regions.

5.3 HIRANS non-local prediction with sparse LES data

5.3.1 Introduction

This section challenges the HIRANS model, using LES data from different geometries to predict the value at a new location. A description of the different interpolation techniques used to combine the LES cases is given, and their results are assessed. Additional comments on the potential non-physical features which can be predicted and a comparison with RANS and locally injected HIRANS follows.

In the previous chapter, the performance of the HIRANS model was validated for the periodic hill test cases when the injected information came from a LES case with the same geometry. Nevertheless, for a design optimization routine, sampling each point with the high fidelity data is not only prohibitive in terms of computational cost, but also does not bring any tangible benefit in the interpolating Co-Kriging framework, as one of the assumptions of this model, discussed in Section 2.4.3, is that *no more can be learnt using low fidelity data in a point already sampled with the high fidelity one*.

In consequence, it brings the necessity of being able to use the HIRANS model for geometries where no high fidelity data is available. This would lead to an envelope with numerous HIRANS point which only use *sparse* LES (or any other high fidelity) data. The choice of which high fidelity information should be used is certainly broad: when interpolating non-locally, in this highly non-linear flow, sophisticated techniques might be needed.

5.3.2 Interpolation techniques

The different techniques to interpolate τ_{LES} data into a HIRANS case are described in this section. While all of them are based in different quantities, they are also a linear combination of weightings based on the chosen metric. It must also be noted that the interpolation was performed in a cell-to-cell basis, irrespective of the position of the cell centres in each geometry. As the mesh grows in y direction, fixing the x positions, only a difference in height between corresponding nodes existed.

Parameter space interpolation

The first and simplest choice for interpolation is to perform it in the ψ parameter space. For each LES case i , and considering the new width of the hill as ψ_{new} , the Euclidean distances p_i can be computed as

$$p_i = |\psi_{new} - \psi_i|. \quad (5.2)$$

Using this, a linear weighting procedure is applied, giving higher importance to the cases closest to the new HIRANS sample, as

$$\alpha_i = \frac{1}{N_{LES} - 1} \left(1 - \frac{p_i}{\sum_{j=1}^{N_{LES}} p_j} \right). \quad (5.3)$$

Finally, any interpolated variable β from the LES solution is the linear combination

$$\beta_{new} = \sum_{i=1}^{N_{LES}} \alpha_i \beta_{LES_i}. \quad (5.4)$$

In this formulation, all the LES cases are taken into account into the interpolation procedure. However, due to the different flow features that can be found in the envelope of the parameter space, this could potentially create very noisy and/or non-physical solutions. Due to this, instead of taking all the LES cases, only two of them, the ones which are “left” and “right” of the current one ψ_{new} , can be used to interpolate from, as

$$\psi_L = \underset{\psi \in \psi_{LES}}{\operatorname{argmin}}(|\psi - \psi_{new}|, \psi < \psi_{new}), \quad (5.5a)$$

$$\psi_R = \underset{\psi \in \psi_{LES}}{\operatorname{argmin}}(|\psi - \psi_{new}|, \psi > \psi_{new}). \quad (5.5b)$$

Bottom wall distance interpolation

Interpolating on the parameter space could be misleading, as a linearity in this one might not be aligned with the non-linearities in the geometry definition and the flow solution. Recalling 5.1, the parameter ψ acts linearly over the x/h dimension, but not over y/h . This one (see Appendix D) is a fourth order piecewise continuous polynomial.

Using the bottom wall y coordinates, a geometry-based interpolation distance can be defined, as

$$p_i = \sum_{k=1}^{n_x} |y_{new_k} - y_{i_k}|. \quad (5.6)$$

In this expression, n_x is the number of nodes in the x direction, and the coordinates y are referred to the first cell centres.

After the computation of the distances, the weighting coefficients and the interpolated variables are computed in an analogous way to (5.3) and (5.4).

Kriging flow variable interpolation

Using a geometry-based interpolation allows to capture large differences in the geometry topology, which do not necessarily correlate with the *flow* topology. That is, a small change in the curvature of a region could be taken as a small difference in a geometry-based interpolation, but could also mean the limit between a separated and an attached flow, or the formation of a secondary recirculation region.

While the flow topology might not be known *a priori*, it can certainly be approximated. From the myriad of ways of doing this, a Kriging surrogate model can be constructed. Using this information, the value of a certain flow variable can be predicted and used as a metric for the interpolation weighting coefficients p_i .

With the Kriging framework developed in this work, 2.22, a prediction of a flow variable β_{new} is performed, as $\hat{\beta}_{new}$, using all the LES data available. With the predictor constructed for the chosen variable space, this leads to

$$p_i = |\hat{\beta}_{new} - \beta_i|. \quad (5.7)$$

While this formulation might have the advantage of better capturing the flow characteristics in the design envelope, it is not exempt of the discussion of which variable to choose and how well can it be predicted in a highly non-linear flow with sparse initial data. A Co-Kriging could be implemented to improve this introducing RANS or HIRANS as cheap datasets, albeit it is not clear how misleading could be the information if large discrepancies are found between RANS and LES. Also, employing HIRANS would be biased by the fact that every new sample point would be dependent on the weighting coefficients which were also sampled before using the initial high fidelity data.

5.3.3 Results

Choice of interpolation technique

ψ_{new}	Parameter	Bottom wall	Kriging k	Kriging f_p
0.50	0.50	0.086	0.60	0.47
0.75	0.50	0.092	0.50	0.45
1.0	0.80	0.40	0.78	0.61
2.0	0.50	0.35	0.24	0.97
3.0	0.50	0.30	0.26	0.31

Table 5.1: Values of the weighting coefficient for the “left” LES neighbour, α_{ψ_L} , for neighbouring interpolation. The right neighbour value is $\alpha_{\psi_R} = 1 - \alpha_{\psi_L}$.

In this part, the different interpolation techniques (parameter, bottom wall and Kriging flow variable) are compared. The main metrics to be used are the L^i norms with respect to the corresponding LES solution in the local geometry.

Figure 5.4 shows the average values of the norms of all the ψ cases considered. It is visually evident how interpolating using the neighbouring cases provided a much closer picture with respect to the LES solution in terms of the different discrepancies. This was especially noticeable in the different turbulent stress components, whose L^1 and L^2 norms were up to 3-5 times higher when injecting information of all the available LES data, whichever interpolation technique (parameter, bottom wall or Kriging) was utilized. This effect similarly happened for the L^∞ norms, where the turbulent stresses showed around twice the difference with respect to the LES case. This was not shockingly surprising, as, recalling Figure 5.5, the dynamics in the steepest and least steep cases are very different, and injecting turbulent stresses information of a non-separated case into a geometry where a single (or double) recirculation region is expected, even with low weighting values, introduces data which does not correspond with the reality of the flowfield.

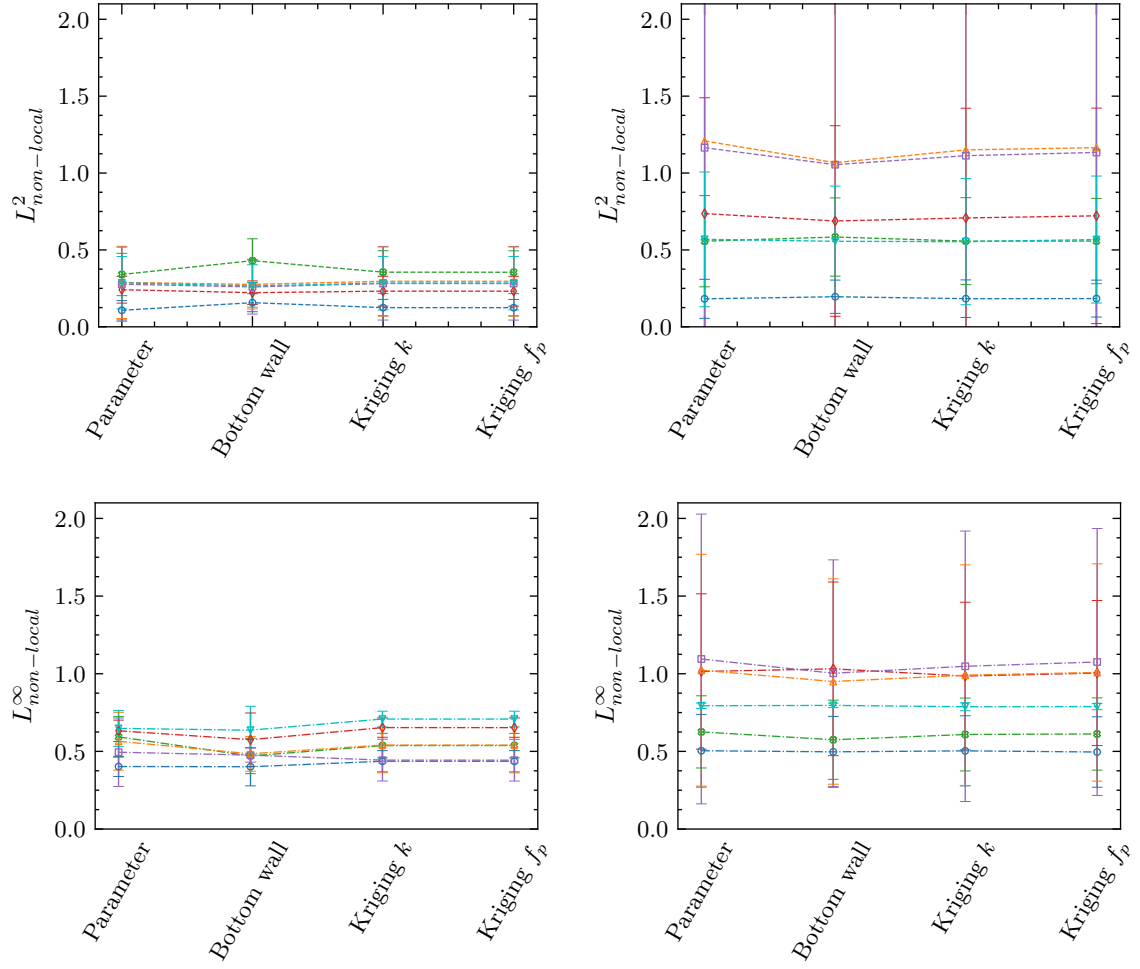


Figure 5.4: Mean and standard deviation of L^2 (top) and L^∞ (bottom) norms for different geometries of $\langle U_x \rangle$ (\circ), $\langle U_y \rangle$ (\diamond), $\langle u'u' \rangle$ (\diamond), $\langle v'v' \rangle$ (\triangle), $\langle w'w' \rangle$ (\square) and $\langle u'v' \rangle$ (∇) of interpolation with neighbouring (left) and all (right) cases.

Non-physical features discussion

This phenomenon can be observed in Figure 5.5. The distribution of the turbulent kinetic energy in this case, when using all LES data available, did not contribute to a distribution that was even physical with these boundary conditions. It is observed that there was some resemblance of the peak k around $x/h = 3.0$ and $y/h \approx 0.8$, corresponding to the Kelvin-Helmholtz instability produced after the separation region. However, tracing that back downstream, there was an exaggerated curvature in the contour levels, which probably were due to the different curvature around the body-fitted mesh in the $\psi \leq 0.75$ case which it had been interpolated from. This is an important point, as the linear weighting procedure gave a relatively high influence to the $\psi = 0.25$ and $\psi = 0.5$ cases, as they are close in the parameter space. Moreover, there was no clear increase of k along the instability, and there were very low values along the leeward hill. This was probably




Figure 5.5: Turbulent kinetic energy contours (top) and streamlines (bottom) for $\psi = 1.0$. Left: all cases, parameter interpolation. Right: LES solution.

related to the delay of the separation in the x/h direction for $\psi \geq 2.0$. As the grid locations were fixed in this axis, and the interpolation was performed in a cell-to-cell basis, the $(i, 0)$ cell, where the subindex i refers to the x direction and the 0 is the first cell starting from the bottom wall, received injected values of a flow where the separation did not occur yet.

Moreover, the streamlines show that there was still a separation region, but what happened downstream was highly non-physical: there was a secondary recirculation region with highly curved streamlines near the foot of the leeward hill, representing a much higher adverse pressure gradient dynamics than the real scenario. The main recirculation region was also significantly extended.

It is concluded that choosing neighbouring cases is the appropriate choice in all the metrics considered. Even in this case, if very different physical features were used for the injection process, non-physical results could still appear.

Comparison with local injection and RANS

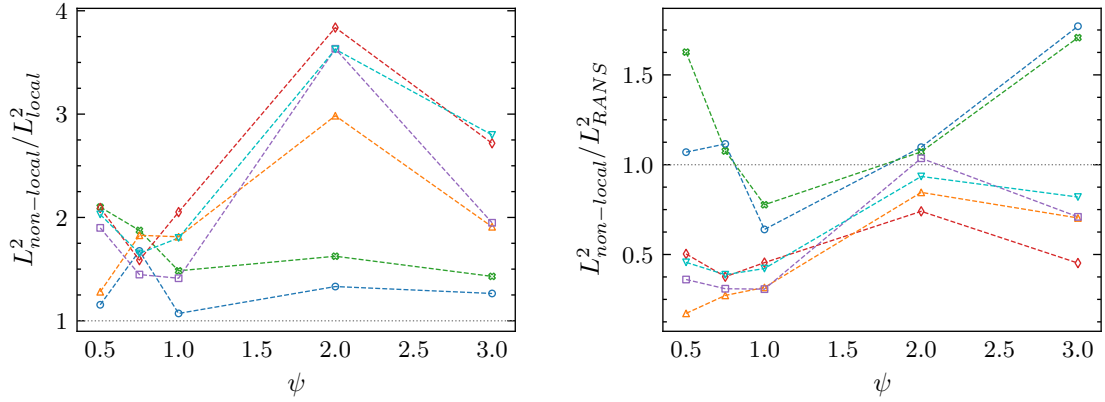


Figure 5.6: L^2 (- - - -) norms for different geometries of $\langle U_x \rangle$ (\circ), $\langle U_y \rangle$ (\odot), $\langle u'u' \rangle$ (\diamond), $\langle v'v' \rangle$ (\triangle), $\langle w'w' \rangle$ (\square) and $\langle u'v' \rangle$ (∇) with respect to local injection (left) and pure RANS (right) values for neighbouring parameter interpolation.

To quantitatively analyse how the non-local interpolation of the injected stresses performs, Figure 5.6 shows the norms for the different geometries non-dimensionalized with the values of local injection and pure RANS cases. As expected, the non-local injection with sparse LES data performed significantly worse to the local injection for virtually every value considered. While the mean velocity components seemed to be relatively similar, the turbulent stresses showed major variations, being three-four times worse than the local solution at $\psi = 2.0$. With respect to the RANS solution, the conclusion which can be derived is that non-local injection seemed to consistently work well for the different components of τ , with non-concluding results for the velocity components, which were significantly better ($\psi = 1.0$) or worse ($\psi = 0.5$, $\psi = 4.0$) depending on the specific case considered. Interestingly, the greatest variations in turbulent stresses did not correlate with the ones in the mean velocity components. Further analysis is performed in Section 6.7.2 with a larger amount of LES samples.

5.4 Summary

In this chapter, the HIRANS formulation with τ_{LES} injection has been tested for different periodic hill geometries, to understand its capabilities and its limitations.

On the first place, the periodic hill geometry was parametrized with respect to the width hill with a single factor ψ . This allowed to test a variety of geometries with different flow features, ranging from double recirculation regions to no flow separation at all. Using information from a local LES case, it was observed that the HIRANS model consistently achieved closer result to the reference data in comparison with its $k - \omega$ counterpart. This was especially notable in the turbulent stresses. The mean velocity components were corrected in a lesser degree.

With respect to the prediction cases when applying non-local injection, three main different interpolation techniques for the τ_{LES} term were tested: in the parameter, bottom wall and flow quantities space. The three of them were tested choosing only the neighbouring geometries and all available cases for interpolation. As expected, neighbouring cases were the most suitable choice. When all cases were injected, there was a risk of non-physical features developing, highlighting the importance of injecting similar flow features, as vortices or recirculating regions, as the expected in the new geometry. With respect to the interpolation methods, they were qualitatively similar, being the parameter and Kriging methods almost identical. This was attributed to the fact that the final interpolation step was based on a linear combination, and the different techniques only affected the values of the weightings of each left and right cases. The results for the non-local prediction resulted in a less reliable method than using local injection, but still offered a significant advantage over RANS in the turbulent stresses, whilst the mean velocity components were only improved in some particular cases, and significantly worsened in others.

This chapter reveals that the HIRANS method can robustly correct turbulent stresses in the test case considered, which makes it suitable to be introduced in a multi-fidelity optimization routine with LES as the expensive solver if the objective function is linked with turbulent quantities, such as turbulent mixing. Nevertheless, the model could also potentially not provide an increase of accuracy over RANS, or even underperform it, if the mean velocity components are quantitatively relevant in the objective function determination.

Kriging and Co-Kriging optimization

6.1 Introduction

The previous chapters have given an analysis of the potential of the HIRANS framework in both local injection and prediction. In this final chapter, all the different pieces which have been previously investigated in the context of global design optimization are combined.

Kriging (2.22) and its multi-fidelity extension Co-Kriging (2.32) are used for this purpose. As stated in the introduction, RANS and LES solutions might yield significantly different results, which can consequently lead to different regions of local and global minima when using both of them separately. The Co-Kriging framework allows to obtain a model which interpolates the expensive, assumed to be the ground truth, solver data points, while regressing over the cheaper ones. However, if the cheap solution is far away from the true shape of the objective function along the parameter space, this regression might not be able to reduce the true error compared to a single fidelity model and even introduce a misleading direction in the global search.

This chapter starts by presenting the ten optimization processes setup: Kriging with LES, RANS and HIRANS and Co-Kriging using LES-RANS and LES-HIRANS, each of them for two different objective functions. After this, the results of the different optimization strategies are presented, analysing the performance of the different methods and debating over the suitability of each of them for the current test cases.

This global optimization framework has a significant number of parameters, comprising the balance between exploration and exploitation ξ_{EI} , the injection parameters ξ_t, ξ_m for the HIRANS model and the number of initial samples. All of these are discussed, especially focusing in the shortcomings of the current formulation. Moreover, error corrections over the tested Co-Kriging with LES and HIRANS model are proposed.

To conclude, all the LES samples from every single optimization process are compared with RANS and HIRANS data, giving a quantitative figure of the improvement of this novel method compared to the general $k-\omega$ in both flow and the chosen objective function quantities.

6.2 Setup

A brief description of the tests performed is given, comprising the parameter space and objective function, the five different methods for the two cases tested, in addition to the surrogate construction and acquisition criteria. The algorithms used correspond to Forrester et al. (2008), whose original MATLAB[®] code has been adapted to Python.

6.2.1 Parameter space and objective function

The parameter space is a scaled version of the width hill ψ presented in (5.1), as

$$x_1 = \frac{\psi - 0.25}{3.75}, \quad x \in [0, 1], \quad \psi \in [0.25, 4.0]. \quad (6.1)$$

The objective function J to use for this project is a linear combination which balances turbulent mixing with a penalty on the total pressure losses, as

$$J = -\frac{\alpha k + \beta f_p}{|J_{LES_{BSL}}|}, \quad \alpha = 1, \quad \beta = \left\{ - (5/2)\alpha, - (5.83/2)\alpha \right\}, \quad (6.2)$$

where k is the mean turbulent kinetic energy over the flowfield and f_p is the needed pressure gradient volumetric forcing term to drive the periodic hill channel at the desired Reynolds number. The multipliers α and β control their balance, which is modified for the two different cases by changing the latter. For the LES samples, this quantity is the mean over the averaging period. This function is scaled with $J_{LES_{BSL}}$, which corresponds to the LES value of the original periodic hill geometry. Finally, the minus sign before the quotient is given as the Kriging framework used in this document is setup for minimization, hence the opposite of the desired quantity to maximise is formulated as J .

6.2.2 Routines description

The objective function J is evaluated by five different methods:

- (a) Kriging LES.
- (b) Kriging RANS.
- (c) Kriging HIRANS.
- (d) Co-Kriging LES-RANS.
- (e) Co-Kriging LES-HIRANS.

The only combination which has not been studied is the Co-Kriging framework combining HIRANS and RANS, as both of them have a comparable computational cost. If HIRANS is believed to offer a significant advantage in accuracy over RANS, employing it as an expensive model would not be particularly useful for this framework, and vice versa. In

addition, the HIRANS model would also be limited by the amount of LES information provided in the initial sampling plan. If the accuracy of both methods is unknown, there would exist a dilemma over which one to use as the expensive model. In situations where this could be predicted *a priori* with certain confidence, this could be a potential option. Based on this, a dynamic HIRANS model, with varying ξ_m and ξ_t injection factors is explored in Section 6.5.3.

All the methods have been tested in two different cases which differ in the number of initial LES sample points (with the trivial exception of the Kriging RANS method) and the balance between turbulent mixing and pressure losses, as:

- (a) Case A: $\beta = (-5.0/2)\alpha$, with 6 initial samples, corresponding to the ones studied in Section 5.2.1, leaving out $\psi = 0.75$.
- (b) Case B: $\beta = (-5.83/2)\alpha$, with all the 7 initial samples from Section 5.2.1.

The reasoning over this choice is to provide two different scenarios in terms of the magnitude and location of the minima, whilst also testing two different degrees of fidelity for the HIRANS prediction, as the left-out point ($\psi = 0.75$) concerns a region with significant changes in the flow topology in a relatively small distance in the parameter space.

6.2.3 Hyperparameter estimation

The determination of the different hyperparameters θ , θ_c , θ_d and ρ for the surrogate models has been performed for all the methods using a Maximum Likelihood Estimation (MLE). Their global search was computed using a Basin-Hopping algorithm (Wales and Doye, 1997), with a L-BFGS-B minimization criterion (Malouf, 2002). It is important to remark that this method has an important cost if a large amount of samples exist or are going to be predicted. This is due to the multiple matrix-inversion operations needed for the correlation matrix Ψ (2.25). This is executed by Cholesky factorization, with a penalty in the MLE applied to the situations where, if \mathbf{x} is densely sampled, this may fail due to encountering a non-positive-definite matrix. There are different ways to reduce this cost. The first one is reducing the amount of the iterations given to the MLE in the global estimation of hyperparameters, or reduce the amount of iterations needed until convergence. This could be applied starting from the second additional sample, where a reasonable estimation of the hyperparameters has already been computed, and it can be given as the the initial point. An alternative is to keep the same hyperparameters for a certain number of iterations, and only computing them after the sampled space has been considerably modified. The balance of improving the surrogate quality and its cost penalty depends on each particular application. In this study, the maximum cost of performing this procedure has been comparable to a RANS sample, significantly lower to an expensive evaluation, and hence it was decided to update the surrogate each iteration.

Moreover, a regularization parameter λ has been introduced in order to improve the conditioning of the covariance matrices, as described in 2.4.1. It is especially relevant when many samples are very close to each other, as it happens with a densely populated parameter space or in some of the locally exploiting evaluations, where the acquisition criterion decides to apply a very small departure from an existing sample point, usually the

current minimum. This can lead to ill-conditioning of these matrices (Micchelli, 1984). An increase of this regularization parameter would create a more regressing, but less interpolating model. A small value was chosen for all the methods, $\lambda = 10^{-9}$, preventing the ill-conditioning but yielding an interpolating model over the high fidelity samples.

6.2.4 Additional samples acquisition

Method	LES evaluations		
	Total number	Exploring	Exploiting
Kriging LES	6	1 st -3 rd , 5 th	4 th , 6 th
Kriging RANS	1	-	1 st
Kriging HIRANS	3	-	1 st -3 rd
Co-Kriging RANS	6	3 rd , 5 th -6 th	1 st -2 nd , 4 th
Co-Kriging HIRANS	6	3 rd , 5 th -6 th	1 st -2 nd , 4 th

Table 6.1: Number of acquired LES evaluations for the different methods with their corresponding exploring/exploiting decisions.

In all the methods, the predicted minimum was assessed by a LES evaluation. In the single-fidelity RANS and HIRANS methods, the x_1 location given by these solvers was then evaluated with LES. The total number of additional LES samples is presented in Table 6.1. For Kriging RANS, there was no possibility of updating the model after the predicted point, hence a single LES sample was computed. In Kriging HIRANS, the model could be resampled using the new LES information, but it was decided during the process to keep the number of acquired samples to 3. The method focused in the vicinities of local minima, and very slight departures of it where obtained regardless of the exploration-exploitation bias within the bounds tested.

As described in Section 2.4.2, the acquisition of new samples was performed using the Expected Improvement (EI) criterion, allowing to balance between exploration and exploitation along the optimization process. Nevertheless, the large cost for a single LES evaluation did not permit a large number of additional samples. With the computational resources available at the time of the first evaluations, the original budget for the Co-Kriging methods was 2 additional samples. Using the unmodified EI criterion, they would have had at least one of the two points in regions where intuitively there was a very small chance of improving the initial minima. Due to this, it was decided to intervene in the bias of the exploration/exploitation balance, using the ξ_{EI} parameter from (2.30a). A complete guideline of the parameters used for the different methods can be found in Appendix G. The first iterations of the Kriging LES methods did not necessarily need this bias, as the lower amount of samples provided larger estimated mean squared errors, which yielded sufficiently suitable locations. In these situations, they are labelled as “exploring” samples, but they were given an almost neutral value of $\xi_{EI} = 0.01$, as in Lizotte (2008). Additional computational budget to perform LES simulations was available after this, which allowed to increase the number of additional samples of these methods to 6.

6.2.5 Flowchart

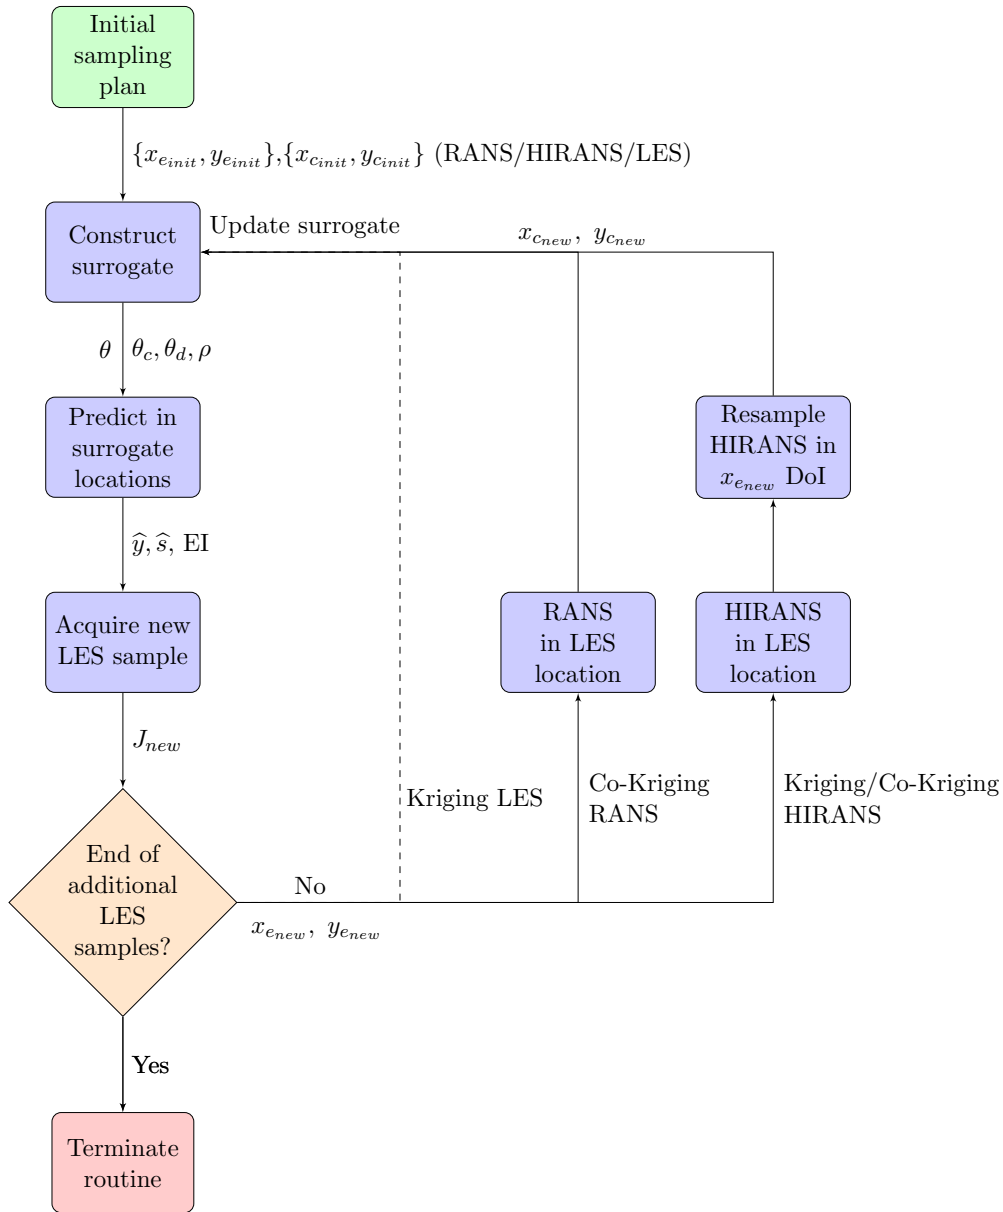


Figure 6.1: Kriging and Co-Kriging flowchart.

Figure 6.1 shows the process followed in the different methods for a fixed computational budget, where the termination condition could be trivially replaced for a target improvement. In this chart, the Kriging RANS scheme is not shown in the loop, as, inherently to its construction, it is only possible to run a single LES point in the minimum location of the Kriging surrogate, with no possibility of updating the model with the new LES information. In the HIRANS methods, an additional resampling was performed in the domain of influence (DoI) of the new LES location, that is, between its left and right LES neighbours, in the order of $n_{c_{resampled}} \approx n_c / (n_e - 1)$ locations, updating the injected information in the new region.

6.3 Initial sampling plans

In this section, the initial sampling plans for the two different cases are shown, commenting the particularities of each different method. In all the methods involving RANS and HIRANS, a linearly spaced division of the parameter space has been performed, with a number of cheap points of $n_c = 50$. In the Co-Kriging methods, the cheap data in the expensive locations is also added, as needed by the surrogate model. This augments the number of cheap samples to $n_c = 50 - 2 + n_e$, as the two extrema of the linear spacing also correspond to LES points.

6.3.1 Case A: $\beta_A = (-5.0/2)\alpha$, $n_{e_{init}} = 6$

The results of the initial sampling plan for the first case can be observed in Figures 6.2, 6.3. Starting with the single fidelity Kriging, the first feature which can be observed is the substantial difference between the RANS and LES surrogates. While it is uncertain how the LES solution might develop when acquiring new samples, the Kriging model predicted a possibly non-convex optimization problem. In contrast, and excepting the first two sample points, with x_1 close to zero, the RANS prediction was a very smooth function with a clear minimum. The Expected Improvement criterion, biased towards exploitation for this latter method, gave the next sample point in a location close to $x_1 = 0.5$. Nonetheless, the bias in this method was almost irrelevant, as the linear spacing between points still drove the acquisition criterion towards regions close to the predicted minimum. In contrast, the LES surrogate showed quite a different picture. There were several points which could be sampled within a $\hat{y} \pm 2\hat{s}$ region with a potentially lower value than the initial minimum. Thus, there were different distinct peaks in this Expected Improvement function, with a different magnitude as the parameter ξ_{EI} is modified. As the first sample was decided to be exploratory, the acquisition model promoted a new evaluation in a region with high uncertainty, to improve the quality of the surrogate model before exploiting.

The non-uniform division of the samples, with larger distance between them in the ones corresponding to the higher values of x_1 , seemed to be an apparent disadvantage in this aspect. Despite a better representation of the steeper configuration regions, where more variations are expected in J due to the greater flow features change in the vicinity of these points in the parameter space, a neutral or exploration-biased surrogate tended to provide at least one or two samples in the least steep regions, $\psi > 2.0$. If the number of additional expensive samples is sufficient, this does not provide major complications in the overall process, otherwise, it becomes a considerable limitation. An initial sampling plan with linearly spaced points could overcome this, as the predictor would be inclined to explore the segment in the regions with higher probability of improvement.

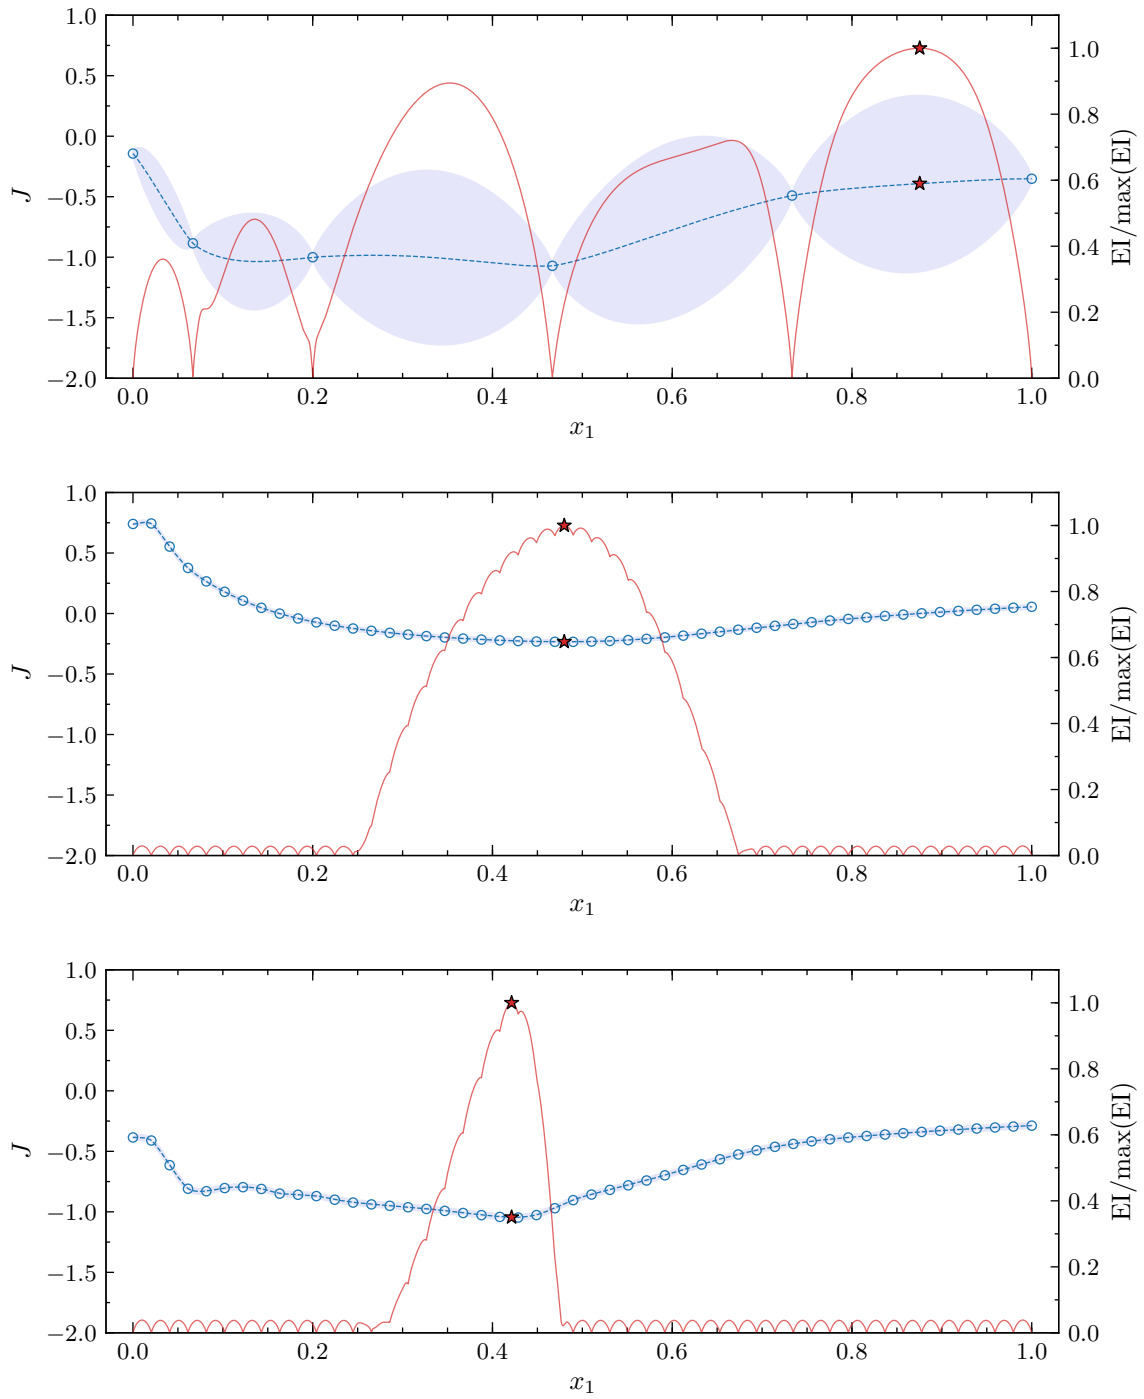


Figure 6.2: Initial expensive (\circ) samples, Kriging $\hat{y} \pm 2\hat{s}$ (---), Expected Improvement (—) and maximum EI location (\star) for the single fidelity methods for case A. Kriging LES (top), RANS (centre) and HIRANS (bottom).




Figure 6.3: Initial expensive (\bullet) and cheap (\circ) samples in expensive locations. Rest of cheap samples (\circ). Co-Kriging $\hat{y} \pm 2\hat{s}$ (---), Expected Improvement (—) and maximum EI location (\star) for the multi-fidelity methods for Case A. Co-Kriging RANS (top) and HIRANS (bottom).

With respect to the Co-Kriging methods, major differences between using RANS or HIRANS as the cheap dataset can be observed in Figure 6.3. For the former, the magnitude of the predicted error, while still lower than in regular Kriging, was still significantly high. RANS predicted a shape which resembled in some way the surrogate close to $x_1 \approx 0$ and $x_1 > 0.5$, whilst differing substantially in the rest of them. In particular, the second, third and fourth LES sample points, $x_1 \in [0.05, 0.50]$, had a similar value of the objective function J , but in RANS they were markedly decreasing.

Conversely, using the HIRANS model, the predicted error was notably reduced, and the Co-Kriging surrogate did closely regress over the cheap dataset. This allowed to see a clear minimum located at around $x_1 \approx 0.42$, with a steep gradient of J after it. The Expected Improvement criterion, biased towards exploitation, predicted very different locations. While in the RANS method the hyperparameter determination provided a smooth function over the second and LES sample, $x_1 \approx 0.12 - 0.20$, hence deciding that an improvement could be located in that region, the HIRANS predicted the opposite trend there.

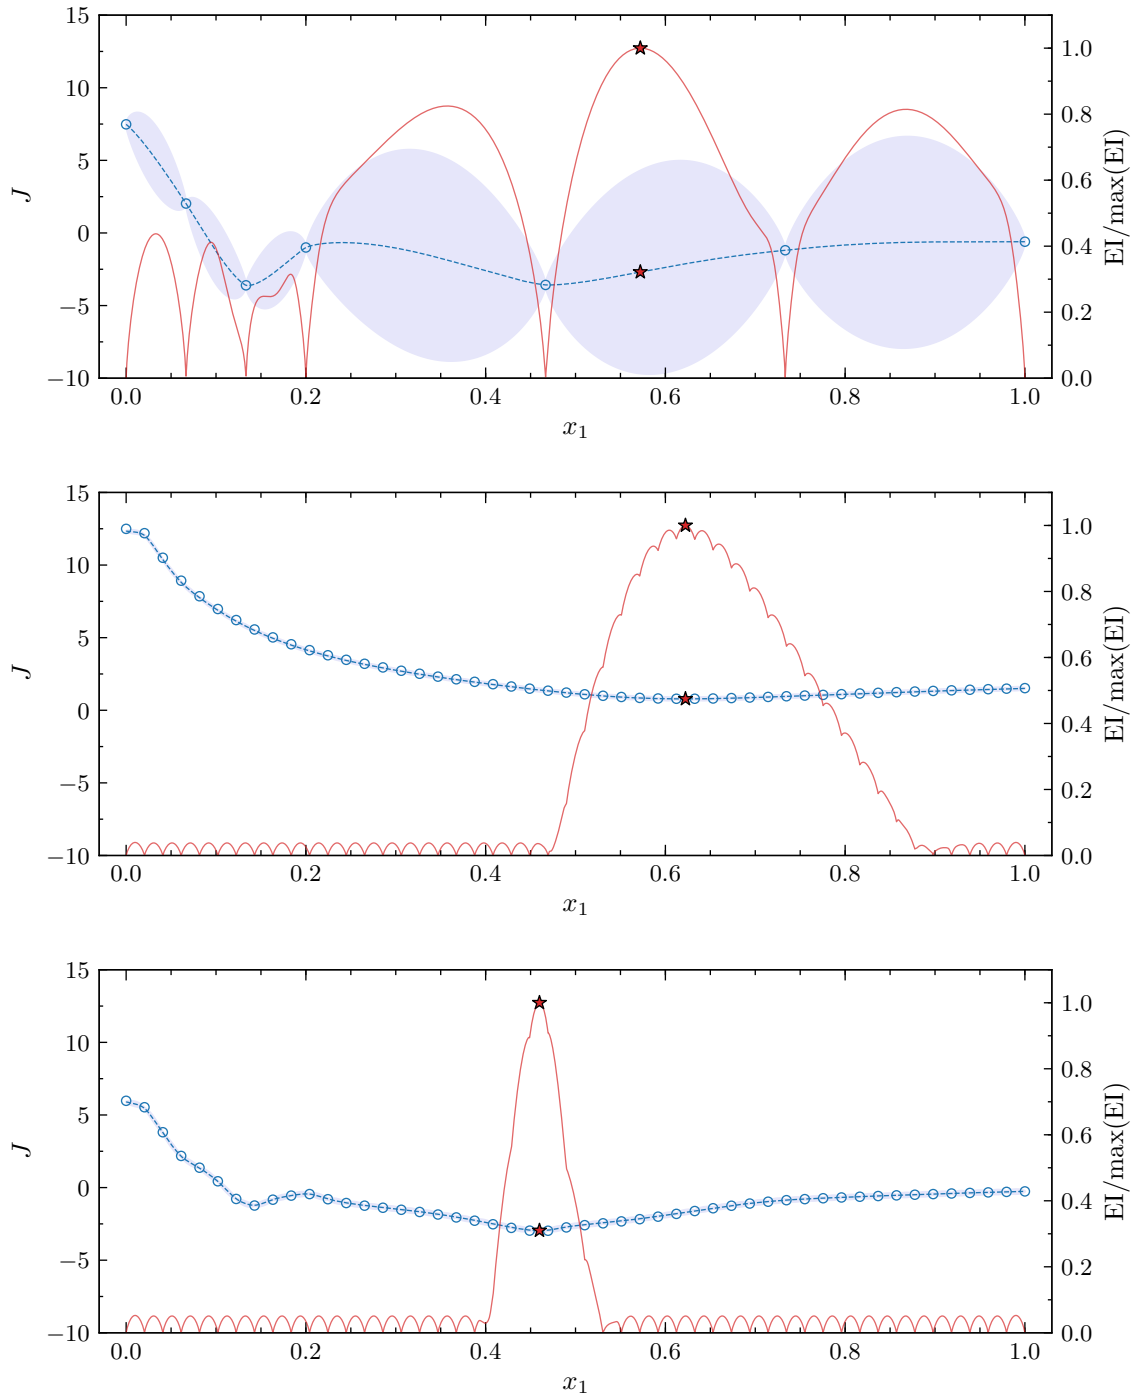
6.3.2 Case B: $\beta_B = (-5.83/2)\alpha$, $n_{e_{init}} = 7$ 

Figure 6.4: Initial expensive (\circ) samples, Kriging $\hat{y} \pm 2\hat{s}$ (---), Expected Improvement (—) and maximum EI location (\star) for the single fidelity methods for Case B. Kriging LES (top), RANS (centre) and HIRANS (bottom).

This second case included an additional sample point, the one corresponding to $\psi = 0.75$, as well as an additional weight over the total pressure losses with respect to the turbulent mixing, which produced two comparable minima, at $\psi = 0.75$ and $\psi = 2.0$, in the initial LES samples, challenging the surrogate model to exploit in the neighbourhood of these two locations and exploring the rest of the space.

Looking at Figure 6.4, the differences between RANS and LES were significant for the single fidelity methods. RANS showed a very similar pattern as the one observed in case A, with a predicted minimum around $x_1 \approx 0.62$ and no sign of the change in convexity of the LES surrogate. In contrast, HIRANS was able to qualitatively predict the locations of the change in curvature of the objective function, but the magnitude of the predicted mean differed from it, yielding a potential improvement of the objective function at around $x_1 \approx 0.45$, considering the one at $x_1 \approx 0.14$ as lower in comparison. This outlines the importance of a sufficiently space-filling initial sampling plan when using HIRANS methods: this method was not able to predict accurately the steep gradient in J found in the $\psi = 0.75$ sample, which shifted the surrogate away from this potential global minimum location.

Similarly as in the first configuration (case A), the Co-Kriging models responded in unrelated directions when introducing RANS and HIRANS information. The first one barely changed the shape of the regular LES Kriging, with relatively high predicted confidence intervals, while the second followed much closer the expensive samples and has a clear regression over the cheap ones. The bias towards exploitation in these two models was noticeable in the HIRANS method, with a very small departure from the vicinity of an existing sample, performing a highly localized exploitation.




Figure 6.5: Initial expensive (\bullet) and cheap (\odot) samples in expensive locations. Rest of cheap samples (\circ). Co-Kriging $\hat{y} \pm 2\hat{s}$ (---), Expected Improvement (—) and maximum EI location (\star) for the multi-fidelity methods for case B. Co-Kriging RANS (top) and HIRANS (bottom).

6.4 Results

In this section, the results from the ten different configurations are discussed. The performance of the different methods is assessed, in terms of the final improvements achieved, the explored regions and the progression per iteration. A separate discussion follows on the behaviour of the Co-Kriging with HIRANS data method, commenting its advantages and shortcomings. The balance between a purely interpolating or regressing surrogate is also investigated, with a new proposal for the regularization parameter λ for future research. Lastly, the final results are shown and compared with the original periodic hill geometry. A visual guide of the parameter space knowledge is given in the figures, computed by a regressive Kriging LES surrogate, with $n_e = 59$ and a regularization parameter of $\lambda = 10^{-6}$.

6.4.1 Objective function improvements

The results of the two different routines for the five methods can be observed in Figure 6.6. Each scatter symbol and colour corresponds to the acquired samples using one method (e.g., Kriging LES). In these plots, the initial and acquired samples from all the different configurations corresponding to each case (A and B) are introduced. Furthermore, the rest of the LES samples available, corresponding to the other case or additional space-filling evaluations are labelled as “additional samples”, for visualisation purposes of where the “true” minima lie. The right plots are detailed views of the samples within the most improved regions.

Starting from case A, two minima can be observed, one of them very close to the removed sample for this case, $\psi = 0.75$. The only surrogate model which was able to find it was the Co-Kriging with RANS information. The rest of the methods decided to explore other regions of the parameter space, especially the other minimum in the neighbourhood of $x_1 \approx 0.35 - 0.50$. The Kriging LES surrogate successfully exploited this region, followed closely by the Co-Kriging with HIRANS. In contrast, both single fidelity Kriging for RANS and HIRANS fixated in a limited location of the parameter space. The former did not achieve a significant improvement over the initial sampling data, while the latter, even if updated with new injected information, was slow in the hill descending process compared to its Co-Kriging version.

The results for case B showed a similar picture: Co-Kriging with RANS and Kriging with LES were similar in their improvement results, with Co-Kriging with HIRANS after them. In the last acquired sample for this method, it jumped to a different location of the parameter space, $x_1 \approx 0.31$, at the other side of the “true valley”, $x_1 \approx 0.39$. With respect to the single fidelity methods, Kriging with HIRANS sampled in an almost identical location as its Co-Kriging counterpart, but it was not able to progress much further.

An argument could be made about how physical are the features observed in these figures, especially the ones seen in the first valley, $x_1 \approx 0.15$. It is observed how the regressing Kriging surrogate does not pass through them, and instead gives a similar value for the two minima. A discussion about an estimation of the noise levels of the LES evaluations is given in Section 6.4.2.

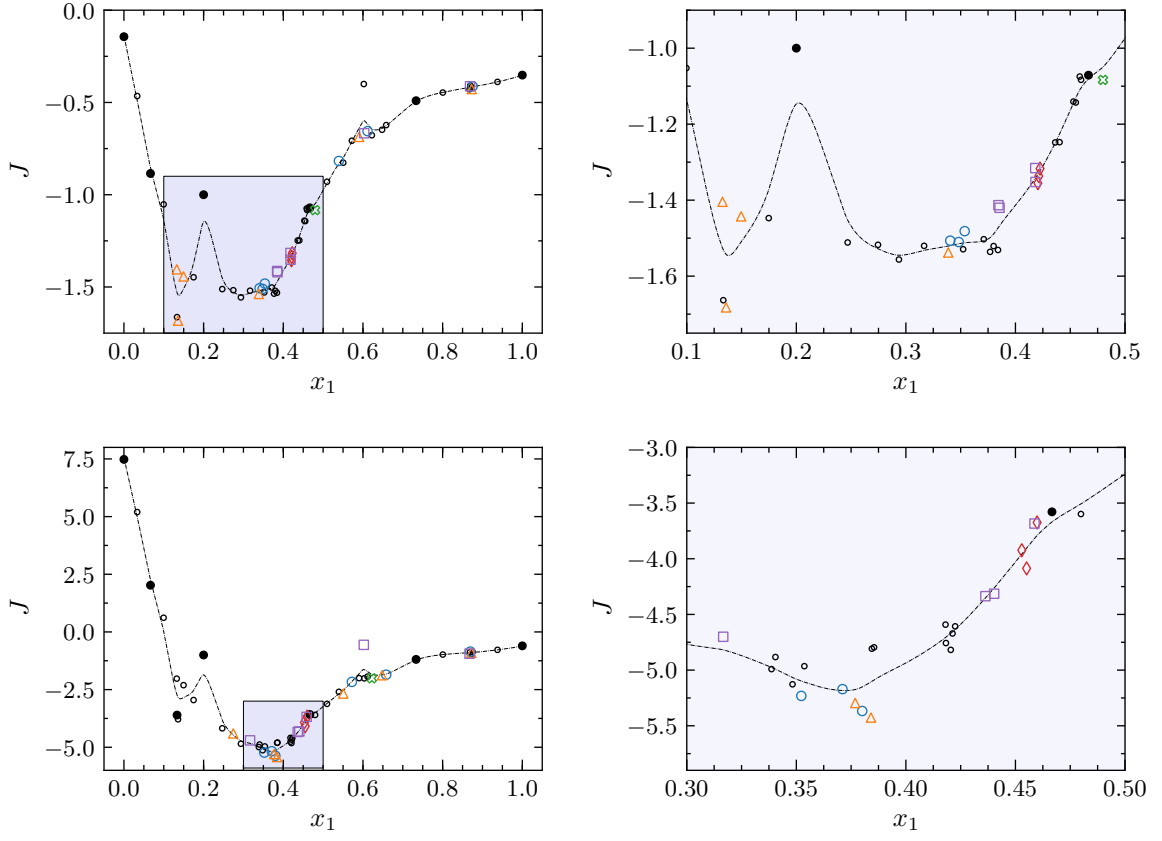


Figure 6.6: LES samples for cases A (top) and B (bottom). Entire parameter space (left) and detail of the minima vicinities (right). Kriging LES (\circ), Kriging RANS (\odot), Kriging HIRANS (\diamond), Co-Kriging RANS (\triangle) and Co-Kriging HIRANS (\square). LES initial (\bullet) and additional samples (\circ). Visual guide using regressive Kriging LES regressive surrogate with $n_e = 59$ (-----).

Quantitative metrics for the improvements can be calculated as

$$(\Delta J)_{BSL}[\%] = \frac{\min(J) - J_{BSL}}{J_{BSL}} \cdot 100, \quad (6.3a)$$

$$(\Delta J)_{IS}[\%] = \frac{\min(J) - \min(J)_{IS}}{\min(J)_{IS}} \cdot 100. \quad (6.3b)$$

The first equation is a reference to the baseline values J_{BSL} corresponding to the original periodic hill geometry, $\psi = 1.0$, while the second metric is the one which shows the improvements with respect to the best point of the corresponding initial sampling plan $\min(J)_{IS}$. Both metrics are calculated using the confirmation LES evaluations, including the single-fidelity RANS and HIRANS methods.

The improvements of the different methods are shown in Table 6.2. As previously outlined, Co-Kriging with RANS information outperformed the rest of the formulations for case

Method	Case A		Case B	
	$(\Delta J)_{BSL}[\%]$	$(\Delta J)_{IS}[\%]$	$(\Delta J)_{BSL}[\%]$	$(\Delta J)_{IS}[\%]$
Kriging LES	51.03	41.00	423.2	45.11
Kriging RANS	1.151	1.151	0.000	0.000
Kriging HIRANS	35.52	26.53	308.7	13.35
Co-Kriging RANS	68.23	57.06	429.5	46.87
Co-Kriging HIRANS	42.01	32.58	370.1	30.38

Table 6.2: Improvements of the objective function for the different methods with respect to the baseline periodic hill geometry, J_{BSL} ($\psi = 1.0$) and minimum of the initial sampling plans, $\min(J)_{IS}$.

A, but in Case B, its result is within the tolerance margin of the Kriging LES method. Both HIRANS methods were also able to provide significant improvements, with a better performance of its multi-fidelity version for both of them. One of the main assumptions of this work is that, in many situations, such as in massively separated flows, RANS and LES can have disproportionately large differences in their calculations of certain quantities, such as the turbulent kinetic energy. This has been corroborated in this test: the single fidelity RANS model achieved either a negligible improvement (case A) or no improvement at all (case B), as the predicted minimum location in this last case yielded a worse result than the original geometry.

Improvements per iteration

The progression of the different frameworks along the iterations is presented in Figure 6.7. A particular feature of this plot is that it does not only show the quantitative improvements, but *how* were they achieved, in terms of the bias between exploration and exploitation. This illustrates how three of the methods (Kriging HIRANS and the two Co-Kriging) achieved their highest improvements in exploitation-biased ($\xi_{EI} < 0$) samples for case A, while Kriging LES did it in an exploratory one. It should be borne in mind that this metric is related to the particular bias achieved with the ξ_{EI} parameter, which does not always correlate with local exploitation or global exploration, especially for a sparsely sampled parameter space. Indeed, dating back to Figure 6.3, the exploitation sample of the Co-Kriging with RANS was performed in a region with significant predicted mean squared error, as the RANS model did not adjust particularly well to the LES information.

The results corresponding to case B show that both HIRANS methods initially succeeded in exploiting the surrogate, being unmatched for the first two iterations. After that, the Co-Kriging HIRANS exploration allowed it to achieve a further improvement in the last acquired sample. A consistent picture can be constructed for these models using this knowledge. In both cases, they outperformed the rest of the methods (excepting the global minimum found in case A by the Co-Kriging HIRANS model) for the first iterations (albeit with different ξ_{EI}), while being harnessed in further ones by the lack of exploration. Even for the sparsely sampled parameter space at the beginning of the process, the HIRANS models were able to encounter regions of local minima. In addition, working with a multi-fidelity framework was especially successful as the routines progressed.




Figure 6.7: Objective function improvement per LES evaluation for cases A (left) and B (right). Kriging LES (\circ), Kriging RANS (\odot), Kriging HIRANS (\diamond), Co-Kriging RANS (\triangle) and Co-Kriging HIRANS (\square). Exploring (- - - -) and exploiting (———) samples.

Moreover, Figure 6.8 represents the L^2 norm of the different predictors with respect to the “true” LES values, computed with a Kriging surrogate with $n_e = 59$. It can be observed that for both cases, the Co-Kriging HIRANS is the model which better resembles the true data, doing it 12 and 22 percent better than the single-fidelity Kriging LES for cases A and B, respectively, and 13 and 14 percent in the vicinities of the minima, that is, $x_1 \in [0.1, 0.5]$ and $x_1 \in [0.3, 0.5]$ for A and B cases. Nevertheless, after the entire optimization process, the overly exploitative model only achieved a modest increase of its prediction capabilities, being inferior to Kriging LES by 23 and 40 percent. While global accuracy is not required *per se*, as a better mapping in the regions of the true minima is usually preferred for the optimization purpose, it also gives an estimation of how far the models are in comparison with the true values. Consequently, Co-Kriging RANS had the best resemblance with respect to the true values, while also being able to find the global minima. As it has been observed, Kriging RANS showed notable differences compared to the rest of the methods, being substantially more unreliable in shape and magnitude, while its multi-fidelity version achieved the best prediction of the objective function in terms of the L^2 error overall.

Co-Kriging HIRANS discussion

A relevant question to be asked is: if the Co-Kriging with HIRANS information was the model which most closely resembled the “true” shape of J , why was it not able to yield the best results in the optimization process? There are (at least) two main coupled factors which can contribute to this.

On the first place, the surrogate is oblivious of the loss of accuracy of the HIRANS points as they become further away from LES locations. The Co-Kriging model can only employ the information of the error at known LES locations (2.32), and in Section 5.3, it was highlighted how much this accuracy was worsened when using non-local injection. Therefore, the predictor cannot calculate how much should it regress over these additional cheap points, or how should the uncertainties grow away from the known high fidelity




Figure 6.8: L^2 norms of surrogates with respect to regressive Kriging LES surrogate with $n_e = 59$ for cases A (left) and B (right). Initial sampling plan (■) and end of optimization process (▨) values.

locations. This results in artificially low predicted error bands, which do not correspond with the true accuracy of the model and conduce to an excessively exploitative method.

The second added challenge is inherent to this particular tested configuration: the objective function contains k as one of its main variables, which is linearly interpolated from neighbouring LES samples and injected in the HIRANS solution. The result of this is that the HIRANS model tends to have a close-to-linear k distribution in many regions of the parameter space, which makes even more difficult to the surrogate model to find hidden non-linearities in the objective function. Generally speaking, if HIRANS is used with very high injection factors ξ_m, ξ_t , it is possible that this flow solution behaves somewhat similarly to a linear interpolation over the J space, due to the constraints imposed. If the k contribution is quantitatively relevant in the objective function, the additional layer of running a flow solver, which might also be able to predict non-linearities in other parameters such as f_p , is the added benefit which contributes to the increased accuracy of the surrogate model over a simple linear interpolation model.

6.4.2 Interpolation and regression

One additional remark must be made about the accuracy of the flow solution. In all the calculated methods, the values of the LES quantities have been treated as the ground truth. This is a strong assumption, which ignores the different uncertainties concerning the obtained results. An estimation of the expected deviations in k and f_p is given. Using subdomain branching (Algorithm 1), the averaging period is divided in progressively smaller domains. Distinct mean quantities are computed for each division. These results are treated as a distribution to calculate the standard deviation from. Up to 55 divisions ($N_{\max} = 10$ subdomains) were chosen in the calculations, combined with (H.2b) to obtain

Algorithm 1: Subdomain branching for error prediction of k and f_p .

Function subdomain_branching(t_{start_0} , t_{end_0} , u_{RMS} , f_p , N_{max}):

```

m = 1.
Δt = t_end0 - t_start0.
for j=1:1:N_max do
  # Number of subdomains.
  for i=1:1:j do
    # Current subdomain.
    t_start_i = t_start0 + Δt/j * i.
    t_end_i = t_start0 + Δt/j * (i + 1).
    u_RMS_m = Σ_{l=t_start_i}^{t_end_i} 1/Δt * u_RMS_l.
    f_p_m = Σ_{l=t_start_i}^{t_end_i} 1/Δt * f_p_l.
    m = m + 1.
  end for
end for
# Statistics over all subdomains.
ū_RMS = Σ_{l=1}^m 1/m * u_RMS_l.
f̄_p = Σ_{l=1}^m 1/m * f_p_l.
ŭ_RMS = √(1/(m-1) * (Σ_{l=1}^m (u_RMS_l - ū_RMS)^2)).
k̂ = 1/√2 * ŭ_RMS * ū_RMS.
f̂_p = √(1/(m-1) * (Σ_{l=1}^m (f_p_l - f̄_p)^2)).
return (k̂, f̂_p).

```

\hat{J} , particularized for the baseline geometry. While these calculations suffer from some undesirable properties (correlation between the subdomains, dependence on the arbitrary number of domains considered, small averaging time, calculated only for the original geometry), they serve as an estimation of how robust could be the values of the inputs for the objective function. Figure 6.9 illustrates this. The observed valley in the objective function at $x_1 \approx 0.14$ is produced by a local increase in k and a specially notable decrease in f_p . The exacerbated peak in J of around 20% is mainly a contribution of the forcing term which, despite not visually evident, drops around 15% of its value in a very localized region, to later recover the general trend. The estimated standard deviation does suggest that part of it could be attributed to the uncertainty over the accuracy of the averaging

process, as a 3σ region would include the general trend. These estimations suggest that further analysis would need to be performed to ensure that the valley corresponding to the global minimum for case A, $x_1 \approx 0.14$, physically exists.




Figure 6.9: Top: LES k (\diamond) and f_p (\circ) samples with estimated 3σ regions. Bottom: Objective function for case A (\triangle) with estimated 3σ regions. Entire parameter space (left) and detail of the minimum vicinities (right).

A proposal connected with the uncertainty analysis can be found in [Talnikar et al. \(2015a\)](#) as *snapping*. This technique was used in Bayesian optimization with LES samples, where, apart from exploration and exploitation, the additional option of running an existing sample for a longer averaging time was also introduced. It allowed to reduce the uncertainty of the design point, with a lower cost than a new sample, due to the savings of avoiding the computation of the transient period. In the simulations tested in this document, this idea was not considered, as the total number of samples for each particular method was rather low $n_{e_{\max}} = 13$ samples, but it would be something which could be taken into account once the parameter space fills and local exploitation appears. A discussion about possible integrations of snapping in the predictor can be found in [Appendix H](#).

6.4.3 Optimum geometries compared to baseline

To finish with the results section, the optimum geometries obtained are presented in [Figure 6.10](#). The difference in the balance between turbulent mixing and total pressure losses is clearly visualized, with the best case A geometry being focused mainly in the former and the case B one in the latter. This second geometry is a larger departure from the original geometry, with a much smaller recirculation region.

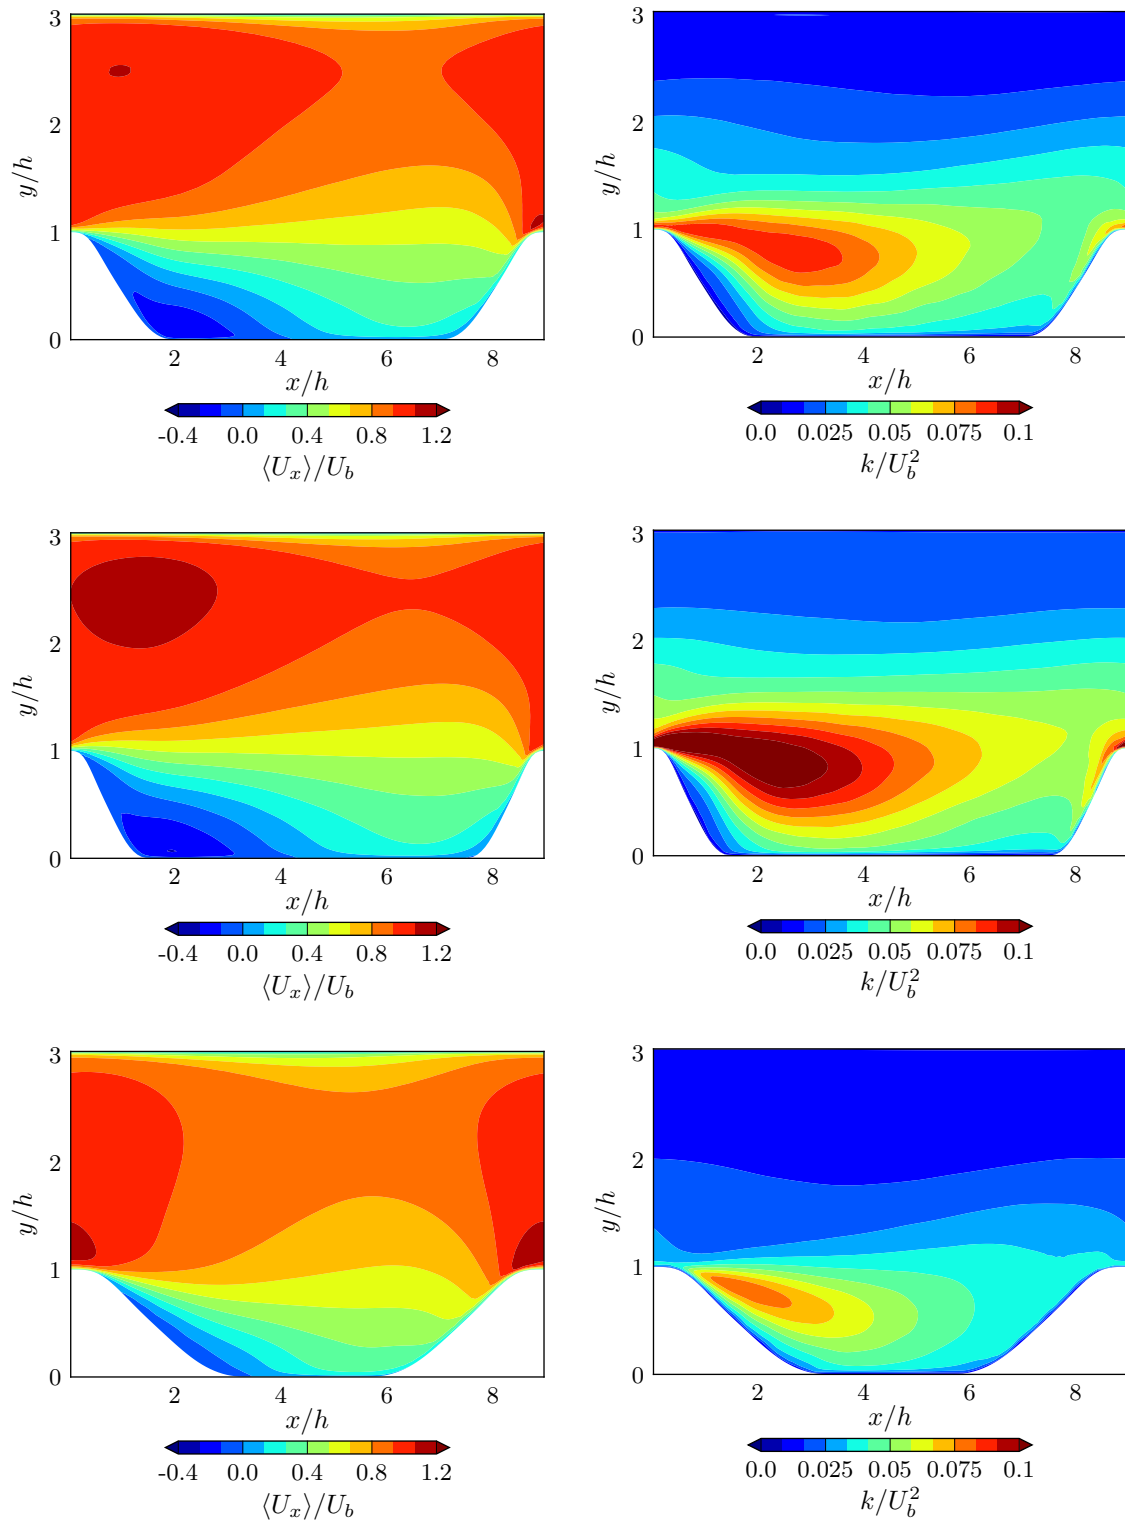


Figure 6.10: Mean streamwise velocities (left) and turbulent kinetic energy (right) for baseline geometry ($\psi = 1.00$, top) and optimum geometries for cases A ($\psi = 0.759$, center) and B ($\psi = 1.69$, bottom).

6.5 Influence of HIRANS and Co-Kriging parameters

In this section, the influence of different parameters regarding the HIRANS and surrogate models is studied. Unless explicitly specified, all these studies are performed using the initial sampling plan of case A and $\xi_{EI} = 0$.

6.5.1 Influence of the exploration-exploitation factor ξ_{EI}




Figure 6.11: Top: Initial expensive (\circ) and cheap (\circ) samples for case A, Co-Kriging HIRANS $\hat{y} \pm 2\hat{s}$ (---). Bottom: Expected Improvement and maximum EI locations with $\xi_{EI} = -0.5$ (—, \star), $\xi_{EI} = 0$ (—, \star) and $\xi_{EI} = 0.5$ (—, \star).

The exploration-exploitation balance can be tuned using the Expected Improvement parameter ξ_{EI} . Figure 6.11 illustrates the importance of choosing a right value for this parameter if the number of samples to acquire is small. In this occasion, choosing a neutral magnitude ($\xi_{EI} = 0$) would still give a fairly exploitative method for this specific sample, due to the reduced predicted error of the Co-Kriging method with respect to its single fidelity counterpart. Switching to a exploitative value ($\xi_{EI} = -0.5$) exaggerates the regions with lower predicted values, as it is observed in the small secondary peak in the EI for $x_1 \approx 0.1$. Recalling Figure 6.6, this was a region which actually corresponded to

the true global minimum, and which was not predicted accurately in the surrogate model with HIRANS information, as the cheap samples suggested a concave shape of J in that vicinity.

On the other hand, an exploratory value ($\xi_{\text{EI}} = 0.5$) can unduly promote regions with slightly larger uncertainty, but with remote possibilities of being an improvement over the current minimum at the initial stages of this routine. This observation needs to be put into the context of this specific situation, as progressing into the optimization process would narrow the predicted error, and the exploring bias could become relevant to query in regions outside the clustered predicted minima.

It is unclear which value of this parameter is optimum for each configuration. Several examples in literature do not consider this balance in their acquisition formulation, as in the popular *EGO* original algorithm (Jones et al., 1998), despite its somewhat related appearance as a parameter ϵ in the Probability of Improvement criterion (Kushner, 1964), being this the sought improvement over the current minimum. A generalised Expected Improvement with a *cool* criterion is described in Sasena (2002), with a modified *EGO* algorithm which starts a global search to focus in local exploitation as the number of iterations increase. Other methods such as the lower confidence bounds (Cox and John, 1992) inherently incorporate an arbitrary value which shapes the bias of the routine in this respect. Additional studies in Bayesian Optimization analysis, such as Lizotte (2008), propose an almost neutral value of $\xi_{\text{EI}} = 0.01$, after stating that it provided the most robust configuration of this study, including in the comparison a cool criterion. Nonetheless, it is noted that the configuration of each specific setup is crucial for this parameter. In this last reference, 500 samples were acquired in a single fidelity framework, a significantly larger value than the available high fidelity budget of the present document, $n_{e_{\text{max}}} = 13$.

Lastly, the objectives of the optimization process must also be considered. Is it purely based on achieving a marginally better design or is there a desire of increasing the quality of the knowledge of the behaviour of J along the parameter space? The latter could provide solutions that, in an industrial context and given several comparable minima, would be more suitable for subsequent processes, such as fabrication or cost constraints, which would lead to an increased bias towards exploration.

6.5.2 Influence of the number of initial samples

High fidelity samples

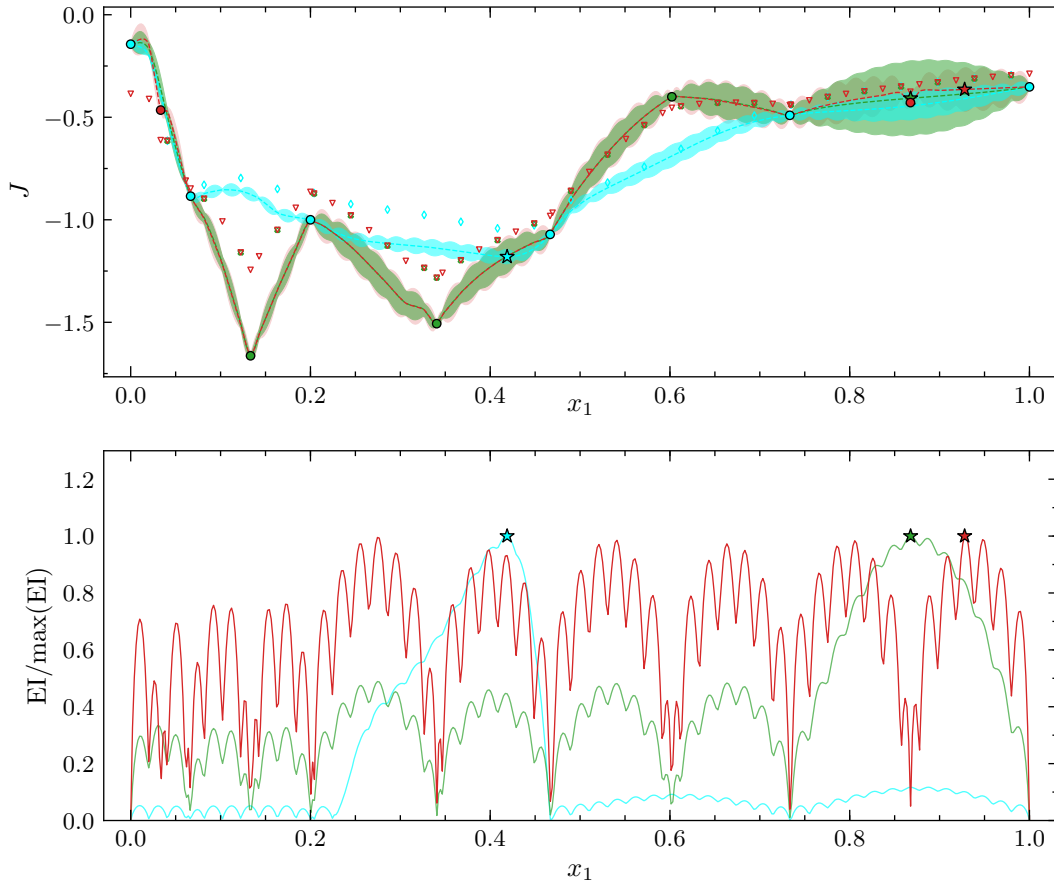


Figure 6.12: Top: Co-Kriging HIRANS $\hat{y} \pm 2\hat{s}$, expensive and cheap samples for 11 (—, ●, ◆), 9 (—, ●, ◊) and 6 (—, ●, ▽) initial LES points. Bottom: Expected Improvement and maximum EI locations with $\xi_{EI} = -0.5$ (—, ★), $\xi_{EI} = 0$ (—, ★) and $\xi_{EI} = 0.5$ (—, ★).

A good space-filling initial sampling plan is crucial for the purpose of global optimization, but if the computational resources are limited to a small amount of total sample points, it can also prevent from being able to exploit local minima. Figure 6.12 shows the result of the initial Co-Kriging sampling plan with 6, 9 and 11 samples, where the additional samples have been located at the regions which were more sparsely sample in the middle part of the domain and close to the extrema, respectively. It is observed that, for this particular case, the change from 6 to 9 initial expensive samples was highly relevant in the Co-Kriging predictor shape. In contrast, the 2 additional points for 11 initial samples did not provide relevant qualitative differences in any of the potential minima regions. The reasoning behind this is that in those regions, both k and f_p behave quasi-linearly (Figure 6.9), similarly to the prediction given by the cheap dataset utilised for regression. This is related to the flow topology in these regions, with distinct two counter-rotating structures

and no separation in the left and right extrema of the parameter space, respectively. A single recirculating region, as in the original geometry, $\psi = 1.0$, is encountered in more interior locations.

Low fidelity samples




Figure 6.13: Top: Co-Kriging HIRANS $\hat{y} \pm 2\hat{s}$, expensive and cheap samples for 48 (—, \diamond), 24 (—, \odot) and 12 (—, ∇) interior HIRANS points. LES samples (\bullet). Bottom: Expected Improvement and maximum EI locations with 48 (—, \star), 24 (—, \star) and 12 (—, \star) interior HIRANS points.

One of the main drawbacks of the current Co-Kriging HIRANS formulation, as discussed in Section 6.4.1, is its overexploiting tendency. Figure 6.13 illustrates this: decreasing the number of cheap evaluations did not significantly affect the quality of the Kriging predictor, as the HIRANS surrogate itself was remarkably smooth, without showing high noisy peaks in any location of the parameter space. Decreasing the cheap interior samples to 12 only seemed to slightly alter the Co-Kriging prediction around $x_1 \approx 0.1$. The Expected Improvement maximum location did not significantly change, but instead promoted regions with secondary peaks such as in $x_1 \approx 0.3$. This is a positive result, which shows that the number of interior HIRANS points computed, $n_c = 48 + n_e$, was probably not the most efficient choice, both in terms of resources and potential. A lower number of

cheap initial samples, such as $n_c \approx 3n_e$, would provide computational resources savings, whilst also preventing the surrogate model to contrive to become trapped in a possibly non-optimal local exploitation region.

This reasoning does not necessarily imply that the number of the cheap dataset points should be static with updates uniquely in the locations of high fidelity queries. An alternative strategy could be to keep a constant ratio between lower and higher fidelity samples, fixating it at a value such as $n_c \approx (2, 3)n_e$ in contrast with the value of $n_c \approx (5, 6)n_e$ tested, using internal equispaced locations inside each segment comprising two high fidelity points. This would enforce sufficient space-filling while avoiding to overconstrain the surrogate in regions without enough expensive data to ensure the reliability of the HIRANS solution.

6.5.3 Influence of the HIRANS injection factors ξ_m, ξ_t

To evaluate the influence of the injection factors in the multi-fidelity analysis, additional routines are computed with lower values of ξ_m and ξ_t .

The injection factors are scaled using a factor α , such as

$$\xi_m = \alpha \xi_{m_{\max}}, \quad (6.4a)$$

$$\xi_t = \alpha \xi_{t_{\max}}, \quad (6.4b)$$

with $\xi_{m_{\max}} = 0.8$ and $\xi_{t_{\max}} = 0.3$.

Figure 6.14 shows the progression followed when decreasing the injection factors. Similarly to the results obtained for the majority of the flow variables for the original periodic hill geometry (Figure 3.19), the largest differences were obtained in the range close to the maximum amount of LES information injected, that is, $\xi_m = [0.6, 0.8]$, $\xi_t = [0.225, 0.3]$. The regression shape of these two configurations was nearly identical, shifting their convexity along the different high fidelity segments when approximately half of the maximum LES data was injected. This was corroborated by the Expected Improvement criterion, which decided to sample in the vicinity of $x_1 \approx 0.4$ for the three methods with higher high fidelity data injected, and closer to $x_1 \approx 0.9$ as the predicted mean squared error was growing.

It can be concluded that, for the current cases, the highest injection factors tested are desirable. If a dynamic ξ_m, ξ_t approach were attempted, a non-physical sawtooth predictor would be obtained.




Figure 6.14: Top: Co-Kriging HIRANS $\hat{y} \pm 2\hat{s}$ for $\alpha = 1$ (\diamond), $\alpha = 0.75$ (\triangle), $\alpha = 0.5$ (\circ), $\alpha = 0.25$ (\square) and $\alpha = 0$ (∇).
 Bottom: Expected Improvement and maximum EI locations with $\alpha = 1$ (—, \star), $\alpha = 0.75$ (—, \star), $\alpha = 0.5$ (—, \star), $\alpha = 0.25$ (—, \star) and $\alpha = 0$ (—, \star).

6.6 Predicted error correction using multi-level HIRANS

While the idea of including dynamic injection factors may not be able to be naively introduced in this test case, there is an indirect way to include more than two fidelity levels in the surrogate model. This multi-level approach, combining locally and non-locally injected HIRANS surrogates in the LES locations, can be formulated to correct the predicted error in the rest of the linearly spaced, non-locally injected HIRANS points. The new predicted error can be computed using a combination of the two surrogates.

The process can be outlined as follows. Firstly, a regular Co-Kriging HIRANS surrogate model is constructed. A *secondary* surrogate is constructed by replacing the interior locally-injected HIRANS points by non-locally injected values using the neighbouring left and right LES samples, conserving the rest of the previous HIRANS values. The predicted errors from the two Co-Kriging surrogates are then combined.

A desirable property would be that the error converges to the regular Co-Kriging with locally injected values close to the LES locations and grows towards the predictor with non-locally injected samples when increasing the distance from them. A distance metric with respect to the neighbouring left and right LES points $\{x_L, x_R\}$ can be formed as

$$d(\mathbf{x}) = \left(\frac{(\mathbf{x} - x_L)^2 + (\mathbf{x} - x_R)^2}{(x_L - x_R)^2} \right)^{1/2}. \quad (6.5)$$

This distance is maximum at an LES location, and minimum when the new sample is equidistant to existing LES points. By scaling it to the range $[0, 1]$,

$$\eta(\mathbf{x}) = \frac{1}{1 - \sqrt{1/2}} [d(\mathbf{x}) - 1] + 1. \quad (6.6)$$

Finally, this parameter η is used to correct the estimated error as

$$\hat{s}^2(\mathbf{x}) = \hat{s}_{local}^2(\mathbf{x}) + \alpha(1 - \eta) |\hat{s}_{local}^2 - \hat{s}_{non-local}^2|. \quad (6.7)$$

where the subscripts *local* and *non-local* refer to the regular (locally injected HIRANS) and secondary (non-locally injected HIRANS) Co-Kriging surrogates. The factor α controls the fact that the error estimates for the non-local predictions have been computed with artificially long sample distances with respect to the ones present in the primary routine. As an example, for equispaced expensive samples $\mathbf{x}_1 = \{0.0, 0.2, 0.4, 0.6, 0.8, 1.0\}$, the HIRANS value for the second surrogate at $x_1 = 0.4$ would be calculated with $x_L = 0.2$ and $x_R = 0.6$. Thus, the distance with respect to its neighbouring points would be $|x - x_L|_{\max} = |x - x_R|_{\max} = 0.2$ for each of them, twice the maximum possible distance in the primary surrogate, as $|x - x_L|_{\max} = 0.1$ with local injection. A simple value of $\alpha = 1/4$ can be used, assuming a linear progression of the predicted error when walking away from LES locations.

The entire process is outlined in Algorithm 2. There are several advantages about this method. The first one is the computational cost: only $n_e - 2$ (as the extrema are excluded)

Algorithm 2: Multi-level HIRANS error correction and prediction.

Function `corrected_surrogate_multi_level`($x_e, y_e, x_c, y_c, \xi_{EI}$):

Construct	Co-Kriging surrogate $\rightarrow \hat{y}, \hat{s}_{local}^2$. <i>Use LES and locally injected HIRANS points.</i>
Run	additional HIRANS in LES interior points. <i>Use non-local information with left-right neighbours (5.3.2).</i>
Construct	secondary Co-Kriging surrogate $\rightarrow \hat{s}_{non-local}^2$. <i>Use LES and non-locally injected HIRANS points.</i>
Compute	distance of predictor points to LES neighbours (6.5).
Scale	distances to $[0, 1]$ range (6.6).
Correct	\hat{s}^2 with \hat{s}_{local}^2 and $\hat{s}_{non-local}^2$ (6.7).
Compute	Expected Improvement and choose new sample point. $\rightarrow x_{e_{new}}, \hat{y}_{e_{new}}$.
return ($x_{e_{new}}, \hat{y}_{e_{new}}$).	

additional cheap function evaluations are needed. Moreover, the model is roughly as robust as the underlying HIRANS solver, and it should increase this robustness as the optimization process progresses and the distances between sample points start to narrow. The possibility of breakdown exists if HIRANS fails to converge in non-locally injected samples, albeit it has not been the case in any of the tests performed in this document. If that occurred, the non-local injection could be performed globally with decreased ξ_m, ξ_t . It would still give an estimation (albeit arguably less accurate) which would more reasonably resemble the true estimates than the original method, unless the error growth when increasing the distance from an LES location is far from the linear assumption. As the most critical situation would be in the first iteration, using the initial sampling plan, this could be easily assessed by running locally and non-locally injected points and assuring that this assumption could hold before progressing with the optimization process. In the case of breakdown, using an arbitrary multiplier for the first iterations, as $\hat{s}_{non-local}^2 = 1/\alpha \max(\hat{s}_{local}^2)$, should provide a reasonable estimate, until HIRANS converged with non-locally injected information and could be introduced into the predictor.

Figure 6.15 shows the result of applying this correction. In the vicinities of the LES locations, the estimates closely followed the shape of the primary predictor, growing towards approximately half the distance between the locally and non-locally injected dataset predictions at the centre of each segment. Comparing the $\hat{y} \pm 2\hat{s}$ region with respect to the original one, it yielded a much more realistic balance, covering a larger part of the underlying “true” function but without the excessive error estimation in regions with less variations of J , $x_1 > 0.5$. It was also supported by the Expected Improvement: the next point to be acquired was located in a similar region with respect to the original Co-Kriging predictor, but the region with high locations of it was significantly wider, giving additional weight to the points with the increase predicted error. Moreover, it did not deceptively accept disproportionately large uncertainties in regions with low possibilities of improving the current minimum, as in the case of the non-locally injected surrogate.

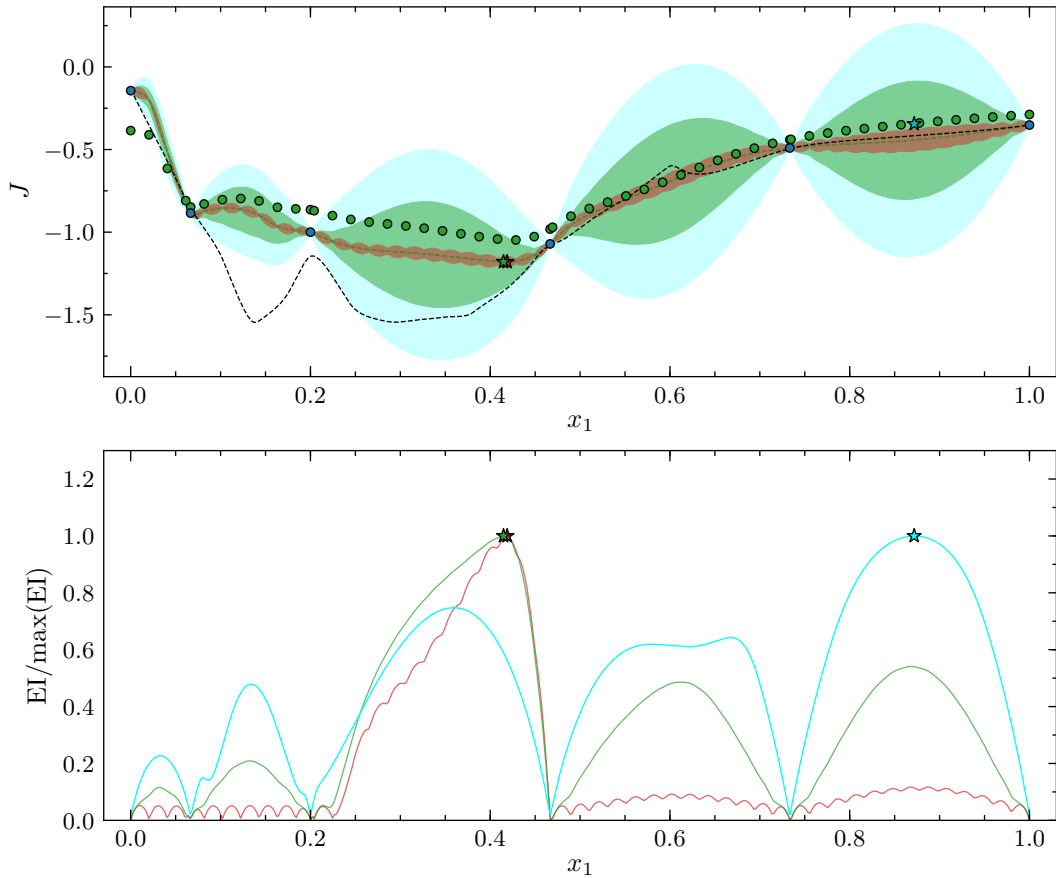


Figure 6.15: Top: Initial expensive (\bullet) and cheap (\bullet) samples, Co-Kriging HIRANS $\hat{y} \pm 2\hat{s}$ for locally injected (—), non-locally injected (—) and corrected (—) models. Visual guide using regressive Kriging LES surrogate with $n_e = 59$ (-----). Bottom: Expected Improvement and maximum EI locations for locally injected (—, \star), non-locally injected (—, \star) and corrected (—, \star) models.

While this simple correction achieved the purpose of addressing the main drawback of the current formulation, there are certainly additional questions which are raised. Choosing non-locally injected points to construct the multi-level surrogate, with unduly long distances in the parameter space, could lead to an excessively conservative model in the error estimates. It could also be unaware of the non-linearities present in the region between the left and right neighbours which it has been interpolated from, aliasing potential minimum candidates in the unexplored regions. In addition, it has not been thoroughly studied how does the error grow when going away from LES locations, and a simple linear distance metric could also not represent the reality. Lastly, the additional cost of the multi-level approach needs to be taken into account, as it needs to be updated in every iteration. In the current test cases, as the cost of RANS/HIRANS is orders of magnitude lower than a new LES sample, and the updates only need to take into account the new sampled region each iteration, it is not a limiting factor. Indeed, the cost is rather small if $n_c \gg n_e$, as each segment update requires the update of the points without a LES solution and the

non-locally injected cheap samples in the two extrema. This could be a relevant cost if the cheap and expensive function evaluations are similar in computational cost, and in occasions with a large number of expensive samples, where $n_c \approx n_e$.

6.6.1 Simulation using modified Kriging believer

Algorithm 3: Modified Kriging believer sampling strategy

Read initial HIRANS and LES samples.
Read master HIRANS and LES Kriging surrogates, $n_e = n_c = 59$.
Construct initial Co-Kriging surrogate.
Obtain first expensive location to sample, x_{add_0} (Algorithm 2).

for $j = 0, n_{add}$ **do**

Acquire	$j - th$ expensive sample, y_{add_j} <i>Use modified Kriging believer and LES master surrogate.</i>
Obtain	HIRANS response at new expensive location. <i>Use modified Kriging believer and HIRANS master surrogate.</i>
Resample	HIRANS in x_{add_j} domain of influence. <i>Use inner HIRANS loop with n_{init} and $j - th$ locations.</i>
Construct	new Co-Kriging surrogate.
Scale	$\hat{s}_{non-local_j}^2(\mathbf{x}_{new}), (6.9), \quad \mathbf{x}_{new} \in [x_{LES_L}, x_{LES_R}]$.
Correct	\hat{s}_j^2 (6.7). <i>Use $\hat{s}_{non-local_j}^2$ and new Co-Kriging surrogate.</i>
Compute	EI and choose new sample point.

end for

To study the potential performance improvement using this correction, a simulation based on a modified *Kriging believer* criterion can be performed. The classic Kriging believer (Ginsbourger et al., 2010; Talnikar et al., 2014) is a heuristic strategy which allows to perform multi-point acquisitions. This method is based on believing the value of the Kriging predictor and modifying the response of the model based on that information to apply the acquisition function for subsequent sample points. An alternative is the *constant liar* strategy, where the response of each acquisition point is given by a constant value L , such as $L = \max(J)$. This strategy deceives the Expected Improvement criterion from getting trapped in small regions.

The Kriging believer criterion can be slightly modified for the Co-Kriging multi-fidelity analysis and the gathered knowledge in this *a posteriori* analysis, as its main aim is to simulate how this model could have performed if these corrections had been applied in the first place. For every new sample point, the expensive and cheap responses can be approximated by the Kriging surrogates constructed by performing single fidelity Kriging

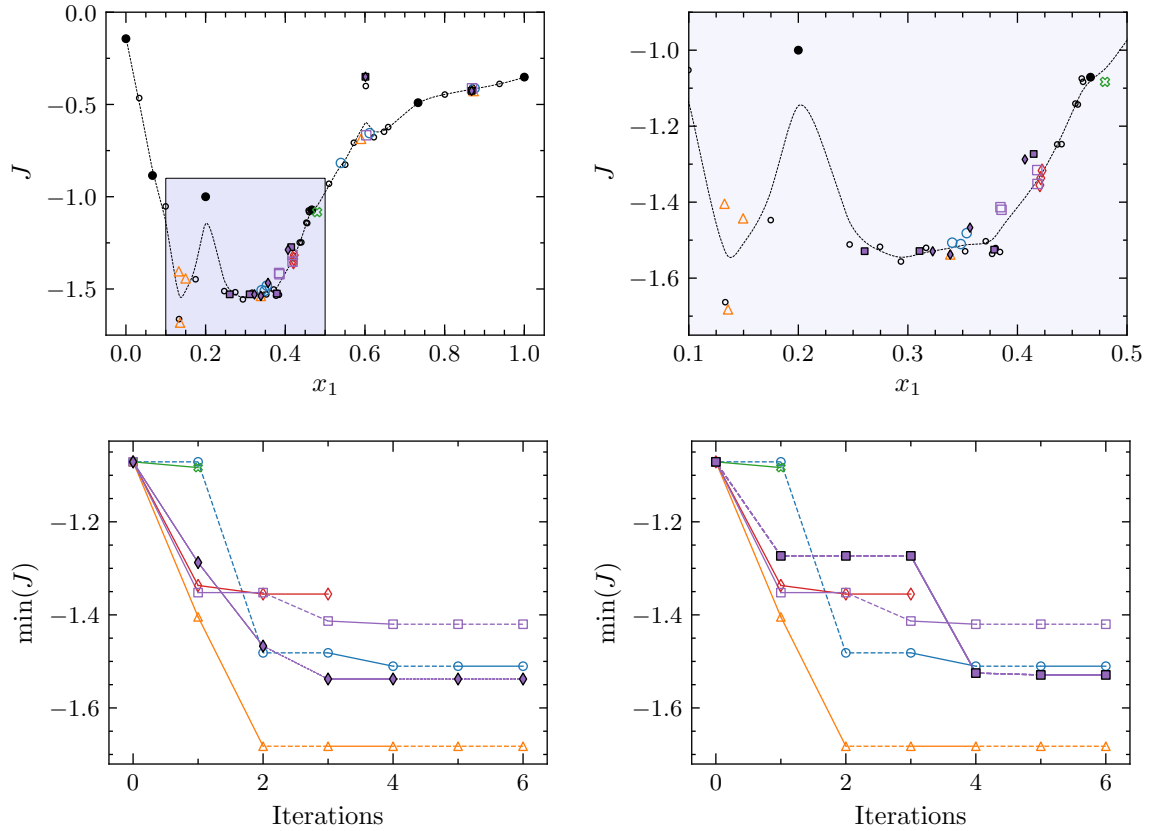


Figure 6.16: Top: Kriging believer simulation for case A. Entire parameter space (left) and detail of the minima vicinities (right). Kriging believer with Co-Kriging HIRANS (\blacklozenge) and Kriging LES (\blacksquare) sampling strategies. Kriging LES (\circ), Kriging RANS (\odot), Kriging HIRANS (\diamond), Co-Kriging RANS (\triangle) and Co-Kriging HIRANS (\square). Initial samples (\bullet) and additional samples (\circ). Visual guide using regressive Kriging LES surrogate with $n_e = 59$ (-----). Bottom: Objective function improvement per LES evaluation for Kriging believer Co-Kriging HIRANS (left) and Kriging LES (right) sampling strategies. Exploring (-----) and exploiting (——) samples.

predictors for all the available LES and HIRANS samples, $n_e = 59$. In this way, no additional flow solutions are needed to be computed, and the method can roughly resemble the conditions of the real parameter space.

As no HIRANS flow variables are computed, the non-locally injected values cannot be directly updated. Instead, the values from the first iteration are used, $\hat{s}_{non-local_0}^2$. In the new location, x_{add_j} , the domain between its left and right neighbours is split by it, and the error is scaled by the new length of each left and right segment. This prevents the solver from computing large errors in new locations close to a LES point.

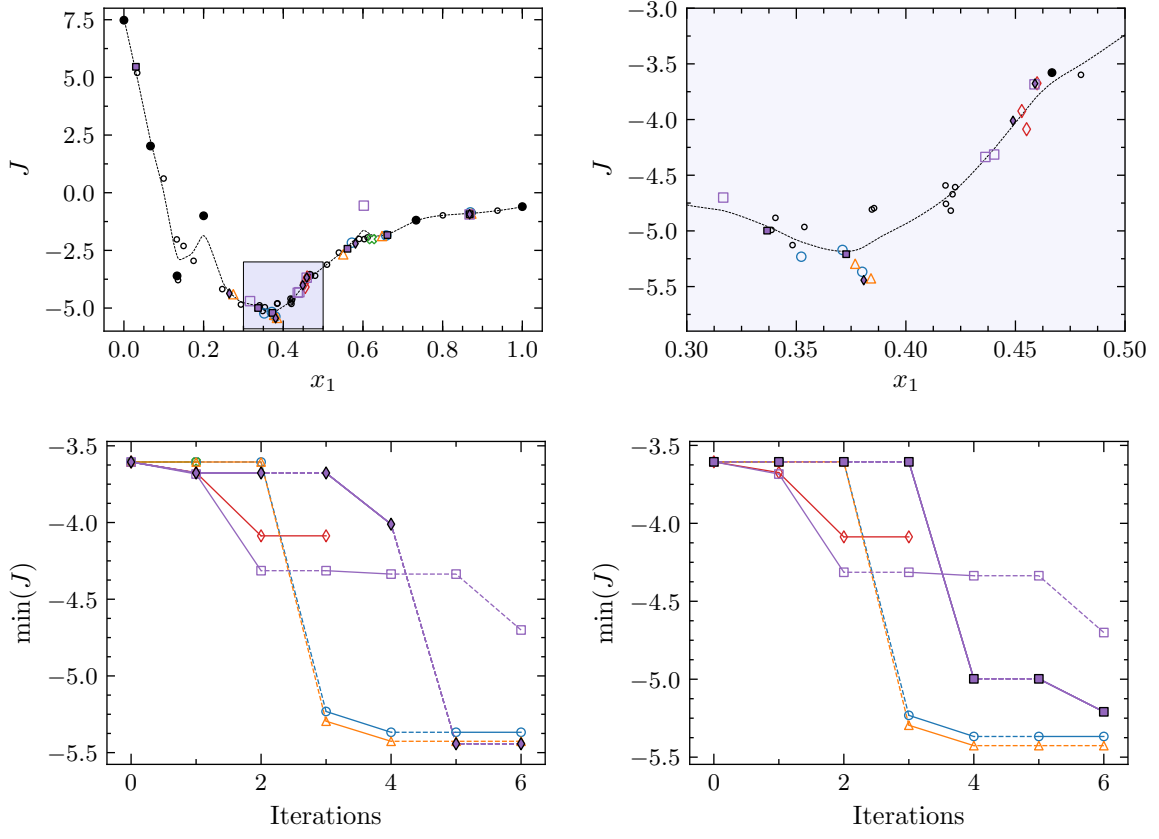


Figure 6.17: Top: Kriging believer simulation for case B. Entire parameter space (left) and detail of the minima vicinities (right). Kriging believer with Co-Kriging HIRANS (\diamond) and Kriging LES (\blacksquare) sampling strategies. Kriging LES (\circ), Kriging RANS (\odot), Kriging HIRANS (\diamond), Co-Kriging RANS (\triangle) and Co-Kriging HIRANS (\square). Initial samples (\bullet) and additional samples (\circ). Visual guide using regressive Kriging LES surrogate with $n_e = 59$ (- - - -). Bottom: Objective function improvement per LES evaluation for Kriging believer Co-Kriging HIRANS (left) and Kriging LES (right) sampling strategies. Exploring (- - - -) and exploiting (———) samples.

The scaling factors from the left and right segments are defined as

$$\rho_L = \frac{|x_L - x_{add_j}|}{|x_L - x_R|}, \quad (6.8a)$$

$$\rho_R = \frac{|x_{add_j} - x_R|}{|x_L - x_R|}. \quad (6.8b)$$

Using these, the non-local component of the error is computed as

$$\hat{s}_{non-local_j}^2(\mathbf{x}) = \begin{cases} \hat{s}_{non-local_{j-1}}^2(\mathbf{x})\rho_L^2 & \text{if } \mathbf{x} \in [x_L, x_{add_j}], \\ \hat{s}_{non-local_{j-1}}^2(\mathbf{x})\rho_R^2 & \text{if } \mathbf{x} \in [x_{add_j}, x_R]. \end{cases} \quad (6.9)$$




Figure 6.18: L^2 norms of surrogates with respect to Kriging LES surrogate with $n_e = 59$ for cases A (left) and B (right). Initial sampling plan (■) and end of optimization process (▨) values. The indices 1 and 2 are for Co-Kriging HIRANS and Kriging LES sampling strategies, respectively.

The main drawback of this method is that the lack of updating of the error distribution apart from the scaling could derive in non-realistic predicted errors as the routine progresses. Nevertheless, it allows for an estimation of how the real routine could respond. The entire process is detailed in Algorithm 3.

The modified Kriging believer strategy has been tested for the two different cases, using the same ξ_{EI} parameters as in the Kriging LES and Co-Kriging HIRANS methods. In case A, Figure 6.16 shows that, regardless of the bias given in the acquisition plan, the corrections outperform the original Co-Kriging HIRANS framework at the end of the process. The key difference is that, even in the exploitative samples, the Expected Improvement decided to test longer departures from the previous point. This created a more global search while constrained to the local minimum vicinities, sampling in similar places than the Kriging LES and the Co-Kriging RANS, despite not finding the global minimum, $x_1 \approx 0.15$

For case B, Figure 6.17, the corrected model achieved similar results, being able to find the global minimum region in both acquisition plans. In this case, it was slightly slower (1-2 iterations) in finding the best region with respect to Kriging LES. Moreover, Figure 6.18 shows that the quality of the surrogate model with respect to the underlying function also increased when introducing this method, as the more global search in regions with high gradients achieved a better mapping overall.

In all the configurations tested the model outperformed the original formulation, and while these results should be prudently considered, they suggest that thoroughly studying ways to inject more realistic error predictions in the surrogate model could potentially increase the efficiency of this method.

6.7 RANS/HIRANS with respect to LES along the parameter space

6.7.1 Local injection

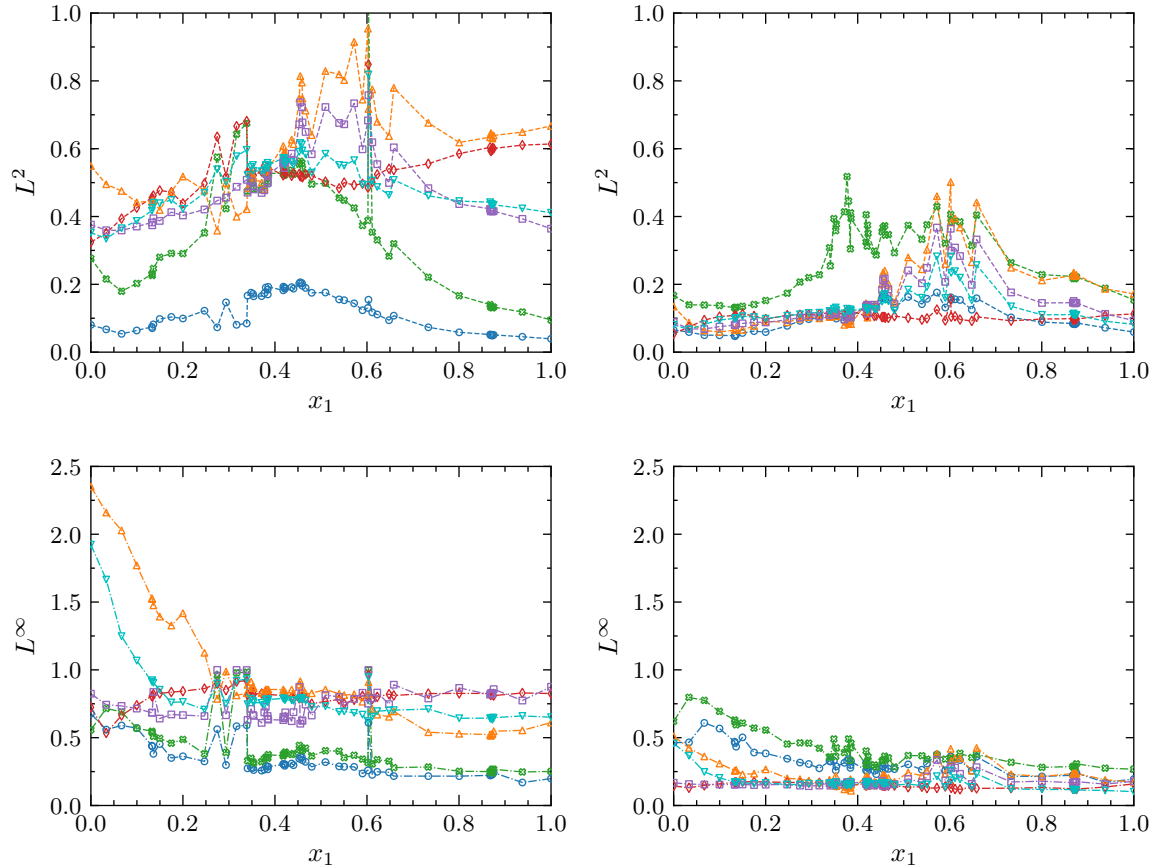


Figure 6.19: L^2 (- - - -) and L^∞ (- · - · - ·) norms for different geometries of $\langle U_x \rangle$ (\circ), $\langle U_y \rangle$ (\odot), $\langle u'u' \rangle$ (\triangle), $\langle v'v' \rangle$, (\diamond), $\langle w'w' \rangle$ (\square) and $\langle u'v' \rangle$ (∇) of RANS (left) and HIRANS with τ_{LES} injection (right).

In addition to the results of the different optimization methods, the total number of LES samples computed for all the different parts of this project can be analysed as a whole in a comparison with respect to RANS and HIRANS results. This allows to draw a similar comparison as the one established in 5.2.2.

The norms of RANS and HIRANS methods with respect to their corresponding LES simulations along the parameter space can be observed in Figure 6.19. As before, the HIRANS model performed much better in the majority of the parameter space in almost all the different norms considered with respect to the normal $k - \omega$ RANS. The effect was more pronounced in the fluctuating components, with L^2 norms which rarely surpassed the 0.4 region for HIRANS, where the baseline RANS model exceeded this value for several quantities and samples. In the L^∞ , a similar result was obtained. In the steepest




Figure 6.20: Mean and standard deviation of flow variables norms with respect to LES values for RANS (■) and HIRANS (▨).

periodic hill geometries ($x_1 \rightarrow 0$), there was a notable reduction of the $\langle u'u' \rangle$ and $\langle u'v' \rangle$ norms, to around a fifth and a fourth of their original values, respectively.

Moreover, these results can be averaged over the parameter space, to obtain some statistics of their distribution. This is shown in Figure 6.20. As visually captured before, the greatest differences were obtained in the Reynolds stresses components. The most notable reduction came from $\langle u'u' \rangle$, where not only the mean L^2 norm was around a fifth of its original value, but the standard deviation was also highly diminished. All the rest of the turbulent stresses followed the same trend, which were also supported by their L^∞ results.

With respect to the velocity components, the results were in general still positive in the L^2 case, but with smaller reductions than for the turbulent quantities. In the L^∞ norm, the HIRANS method yielded very slight reductions with respect to RANS.

These results show that the HIRANS model can be specially suitable for cases where the turbulent quantities from RANS and LES are far away from each other, as it has been able to consistently achieve accurate results over the different samples considered. In the context of a design optimization process, utilizing values related with it, such as the turbulent kinetic energy, this method could be an useful tool to increase the accuracy of the simulations. Indeed, Figure 6.21 shows the improvement of the different parts of the considered objective function prediction, both in the turbulent kinetic energy and in the pressure losses. These conclusions are very positive, as it can be seen that for objective functions, not only they became much closer, but their standard deviation greatly diminished, providing both accuracy and consistency with respect to the LES solution, a desired property in order to infer the shape of unknown regions of the parameter space.

The latter reflection, combined with the better representation of the underlying “true” function with respect to a single fidelity Kriging LES delved in 6.8, allows to prudently infer that the HIRANS method could be seamlessly generalized for multi-parameter optimization. To start with, the formulation itself is independent of any parameter assumption, as the injection of data is based on having neighbouring samples available.




Figure 6.21: Mean and standard deviation of objective function variables norms with respect to LES values for RANS (■) and HIRANS (▣).

Moreover, achieving a more accurate initial surface response, which has been achieved by the Co-Kriging HIRANS formulation, would be even more determinant in a multi-dimensional parameter space, where the acquisition criterion would work with multiple potential minima with different activity in each parameter in its objective function response to its change. A limitation of the current framework could be attributed to the necessary interpolation of τ_{LES} , which would oblige the initial sampling plan to either use extrapolation of this quantity or evaluate all the corners in each dimension, which is not often performed when popular methods such as a Latin Hypercube are used. This could be overcome by utilizing RANS in the corners and introducing a correction in the predicted error in those regions, similarly as in Section 6.6.

6.7.2 Non-local prediction accuracy

With the large database of LES samples available, the accuracy of HIRANS with respect to the distance to neighbouring high fidelity samples to inject data from can also be computed. This is performed using an average distance metric to the neighbouring sample points, as

$$d_{neighb_i} = \frac{1}{2} \left(|x_{1_i} - x_L| + |x_{1_i} - x_R| \right), \quad (6.10)$$

where x_{1_i} is the position in the parameter space to sample, and x_L and x_R are the positions of the left and right LES samples to use for the injection of Reynolds stresses. A value of 0 indicates local injection.

Using four of the interior initial LES samples, that is, $\psi = \{0.75, 1.0, 2.0, 3.0\}$, HIRANS solutions have been calculated jumping one position left and right each time to choose the high fidelity data from the available 59 LES simulations. That is, the first iteration corresponds to the direct neighbours and the next one to the left and right samples of

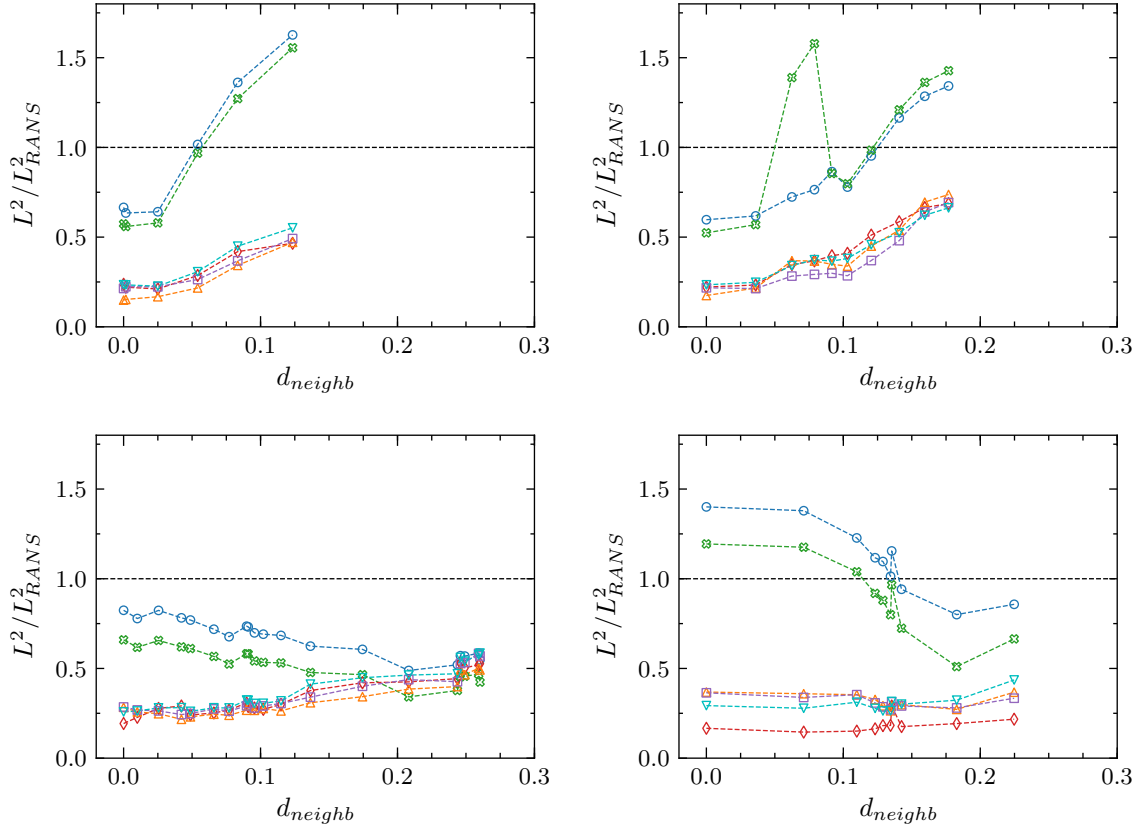


Figure 6.22: L^2 (.....) norms non-dimensionalized with RANS values for increasing distance to neighbouring LES samples of $\langle U_x \rangle$ (\circ), $\langle U_y \rangle$ (\odot), $\langle u'u' \rangle$ (\triangle), $\langle v'v' \rangle$, (\diamond), $\langle w'w' \rangle$ (\square) and $\langle u'v' \rangle$ (∇).
 $x_1 = 0.75/3.75$ (top-left), $x_1 = 1.0/3.75$ (top-right),
 $x_1 = 2.0/3.75$ (bottom-left) and $x_1 = 3.0/3.75$ (bottom-right).

these neighbours. The initial sample corresponding to $\psi = 0.5$ was not included due to the small amount of left neighbours in the parameter space.

Figure 6.22 shows the results of applying the injection with progressively further LES simulations. Even in the long distances tested in the available samples, which covered a quarter in the parameter space in average for $\psi = 2.0$ (that is, the distance between left and right neighbours were roughly separated by half of the parameter space), the Reynolds stresses L^2 norms outperformed the baseline RANS values. This became more evident in the least steep configurations quantitatively. Furthermore, it should be noted that for $\psi = 1.0$, the furthest left neighbour was the extreme sample, $\psi = 0.25$, with double recirculating regions, which posed a challenge to the original periodic hill geometry.

With respect to the L^∞ norms, illustrated in Figure 6.23, they also showed a positive trend. In practically every single flow quantity and sample point, the HIRANS samples with furthest LES points outperformed RANS in this metric. It was especially significant in the steepest configurations for $\langle u'u' \rangle$, due to improvements in the fluctuating shear layer ($\psi = 0.75$) and the vicinities of the windward hill crest ($\psi = 1.0$). The mean

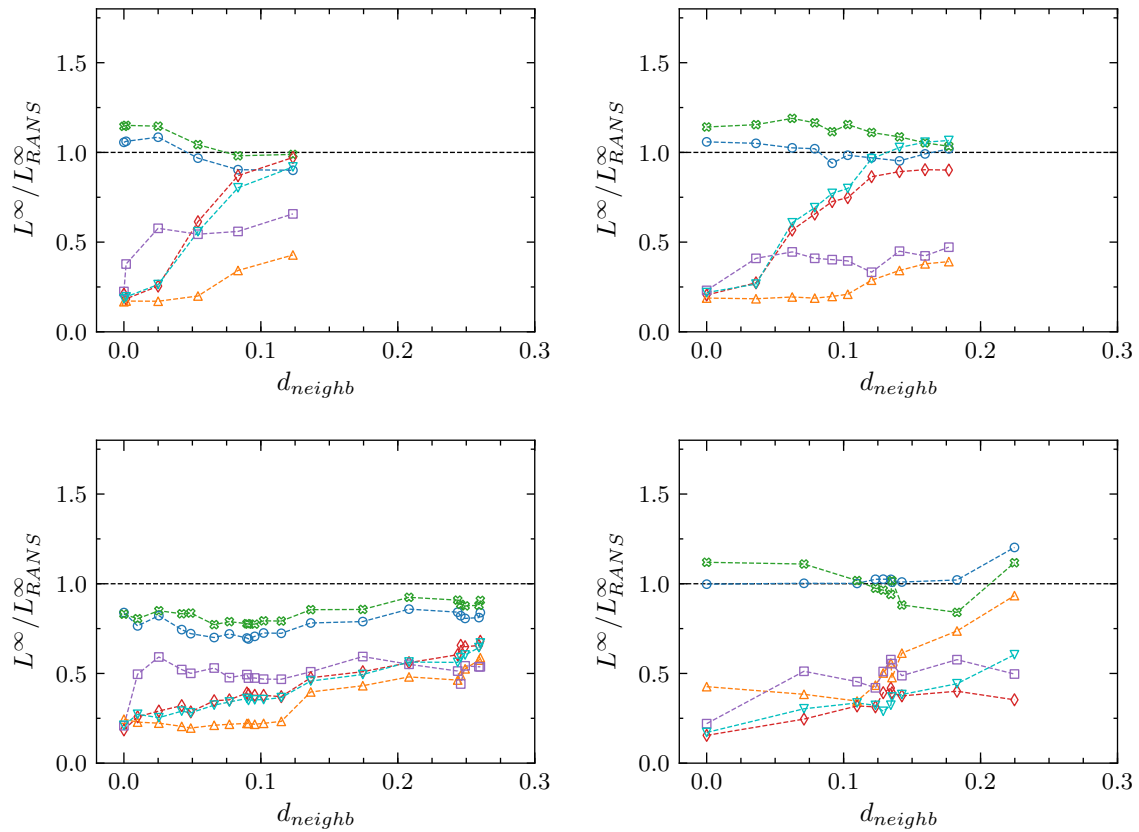


Figure 6.23: L^∞ (.....) norms non-dimensionalized with RANS values for increasing distance to neighbouring LES samples of $\langle U_x \rangle$ (\circ), $\langle U_y \rangle$ (\odot), $\langle u'u' \rangle$ (\triangle), $\langle v'v' \rangle$, (\diamond), $\langle w'w' \rangle$ (\square) and $\langle u'v' \rangle$ (∇).
 $x_1 = 0.75/3.75$ (top-left), $x_1 = 1.0/3.75$ (top-right),
 $x_1 = 2.0/3.75$ (bottom-left) and $x_1 = 3.0/3.75$ (bottom-right).

velocity components were quite irregular, as non-locally injected samples even improved the local HIRANS results for the steeper geometries, slightly worsening them for the less steep ones.

These results corroborate the conclusions from the rest of the document: HIRANS was robust in the prediction of the turbulent stresses throughout the parameter space, with local and non-local injection, despite the long distances tested in the parameter space, with very different flow features. With respect to the velocity components, their improvement with respect to RANS depended on the part of the parameter space considered. On the steeper geometries, the average distance threshold to outperform RANS was around $d_{neighb} \approx (0.05, 0.12)$ in the L^2 norms. In contrast, in the least steep configurations, no threshold was found, consistently outperforming or being qualitatively similar to RANS.

6.8 Summary

This chapter constitutes the piece which combines all the different parts of the HIRANS-LES jigsaw in a global design optimization context.

To start with, single-fidelity Kriging for RANS, HIRANS and LES, in addition to Co-Kriging with RANS and HIRANS in combination with LES, have been studied. Using as a parameter space a scaled version of the hill width presented in Chapter 5, two different cases were tested. The cases differed in the number of initial sample points and the balance between the two variables of the objective function, the turbulent kinetic energy k and the forcing term f_p . The surrogates were constructed by a Maximum Likelihood Estimation for the hyperparameters, with an Expected Improvement acquisition criterion. The computational resources available at the time of the study also constrained the bias of the EI towards exploitation of early samples in the Co-Kriging methods, with a maximum number of additional points of 6 per case and method.

The results of the different tests were partially surprising. The method which provided the best improvement of the objective function, validated with a LES evaluation for all the methods, was Co-Kriging with RANS information in both cases. For case A, it explored the region where the global minimum was located, which the rest of the methods did not query. In case B, it was comparable to Kriging LES. The HIRANS models were able to successfully exploit the objective function in the first iterations, but the single-fidelity model stalled in a local region, while the multi-fidelity version was able to explore further. Nevertheless, the results were far away from the best two methods. In addition, Kriging with RANS was qualitatively not able to improve the initial minimum at all.

Two main coupled restrictions were delved in the Co-Kriging HIRANS method. Inherently to this formulation, the surrogate could not quantify the true accuracy of the non-locally injected cheap sample points. In addition, and particular to this objective function, as τ_{LES} was injected using linear weightings, k_{HIRANS} was constrained by a quasi-linear distribution in certain regions where high non-linearities exist. The result of this, combined with the high number of cheap initial samples, $n_c \approx 8n_e$ samples, resulted in inordinately low error bands, conducting to an overly exploitative model.

A discussion over the trustworthiness of the LES results was given. Interpolating predictors were used throughout the document. Nonetheless, an estimation of the noise in calculating J concluded that the global minimum found by Co-Kriging RANS at $x_1 \approx 0.15$ might be influenced by noise, as a regressive surrogate within the 3σ regions did greatly diminish the magnitude of the peaks. A proposal for future research to incorporate noise information was formulated, modifying the regularization parameter to be individual to each sample, while combining it with the *snapping* technique of Talnikar et al. (2015a) to resample LES for a longer period to reduce the uncertainty over the point.

The influence of different parameters corresponding to the HIRANS formulation and the optimization side was also studied. Starting with the exploration-exploitation bias ξ_{EI} , its optimum value was highly dependant of the objectives of the particular optimization process. For a very low amount of additional samples, $n_{e_{add}} \approx 2$ samples, a neutral balance would result in an overly exploratory model. Conversely, an overly exploitative value disproportionately promoted very small departures from existing samples, harming the potential improvements in longer processes.

With respect to the number of initial samples, 7 space-filling expensive points, biased towards the steepest slopes region, were sufficient to provide a surrogate model which resembled the main features of the underlying function. For the multi-fidelity methods, the number of 48 interior cheap samples was considered excessive in terms of computational efficiency, as a lower number, of around $n_c \approx \{3, 4\}n_e$ was sufficient to provide a similar surrogate, with more realistic predicted errors.

In terms of the injection factors ξ_m, ξ_t , it was found that the closest results to the LES values were obtained when they were set as high as possible. Using dynamic ξ_m, ξ_t far away from LES locations would create a non-physical sawtooth shape.

The error correction using multi-level HIRANS was successful in its application, by introducing information from locally and non-locally injected HIRANS samples in the LES locations, computing a linear combination of them for the predicted error. Using a modified version of the Kriging believer technique, it was simulated and compared against the rest of the methods, outperforming its non-corrected Co-Kriging HIRANS version in the four computed configurations, and situated on par with the single-fidelity Kriging LES.

Using all the available high fidelity simulations, $n_e = 59$ samples, norms of the flow and objective function variables for RANS and HIRANS with respect to LES were computed, with the latter method achieving much closer results. The improvements in the mean velocity components were modest, while performing significantly better in the Reynolds stresses. In terms of the L^2 norms for J , HIRANS was around eight times closer to LES than its original $k - \omega$ model, with much lower standard deviation as well, with maximum deviations for J around four to ten times lower for cases A and B, respectively. The non-local prediction capabilities showed that HIRANS was also able to reach a better agreement to the LES results in the L^2 norms for the Reynolds stresses in every case, even with left and right neighbours were separated by a distance covering 50% of the parameter space. The mean velocity components were worsened significantly in the steeper geometries, albeit they even improved with more distant LES neighbours in the less steep ones.

Two main conclusions can be drawn. Starting with the quality of the predictor, the initial response of J was improved with respect to Kriging LES and Co-Kriging RANS. The HIRANS framework, only dependent of available neighbouring high-fidelity information, could extend its capabilities to multi-parameter design optimization. In terms of the objective function improvements, the uncorrected Co-Kriging HIRANS was already able to successfully exploit local minima in one-two iterations, and a glimpse of its potential has been given by the multi-level Kriging believer simulations. Nevertheless, a thorough study still needs to be performed to improve the predicted error corrections of the surrogate, which will be the next step in increasing the computational efficiency of the method.

Conclusions and recommendations

7.1 Conclusions

7.1.1 Local RANS corrections

Two different Hybrid Injected RANS models were formulated based on injecting non-dimensional turbulent anisotropy b_{ij} and both isotropic and anisotropic components of the Reynolds stress tensor τ from a higher fidelity source, such as LES or DNS. The models acted in both the momentum and the production term of the $k - \omega$ equations. For the periodic hill test case, an analysis of the eigenvalues of the anisotropy tensor using a barycentric map showed that injecting b_{ij} resulted in a successful correction of the turbulence anisotropy. Nonetheless, this was not translated to an improvement of the mean velocity components, as the isotropic component k was not predicted accurately, neither in the baseline RANS model or in the anisotropy-corrected HIRANS. Conversely, using a model with the injection of both isotropic and anisotropic contributions from τ_{LES} did achieve significantly closer results to the LES solution in the different flow variables, stating the importance of the isotropic component for the periodic hill case.

Linking this to the research question:

How reliably and consistently does a hybrid RANS method match the local LES model in the mean flow quantities?

The HIRANS model with τ_{LES} injection was successful in its corrections for the original periodic hill geometry when data from Breuer et al. (2009) was injected, with a virtually identical reattachment location to the LES solution and with the mean velocity components and turbulent stresses within $\pm 5\%$ and $\pm 1\%$ of the reference values when non-dimensionalized with the bulk velocity U_b in the majority of the flowfield, respectively.

What is the range in the parameter space (width of the periodic hill) where locally corrected RANS methods still resemble the time-averaged LES dynamics?

LES cases corresponding to 59 variations of the periodic hill geometry were compared to the primary and the new hybrid formulations. The HIRANS model successfully reduced the RANS L^2 norms with respect to LES in the Reynolds stresses to a third part of their original values in average with significantly smaller standard deviations, especially for $\langle u'u' \rangle$. A modest improvement in the mean velocity components was also obtained. These results were achieved with the added constraint of employing injected LES information with significantly higher y^+ values than the RANS solution. This resulted in smaller improvements in the L^∞ norms, which were located in certain cases in nodes close to the highly accelerated part of the windward hill. In that region, the LES information had been injected using a linear interpolation procedure from higher y^+ values, $y_{LES}^+ \sim \mathcal{O}(10)$, $y_{HIRANS}^+ \sim \mathcal{O}(1)$, not accordingly with the physical progression of the turbulent stresses close to the wall.

The role of the ξ_t injection factor, present in the production term of the $k - \omega$ equations, remains as a future research question. The results were found to be much less dependent on it than on the momentum equation corrections. A value of $\xi_t = 0.3$ was chosen, without a clear indication of performing significantly better than others overall. Regarding the limit in the stability of the corrections, a factor of $\xi_m = 0.8$ for the momentum equation was found to be sufficiently robust, without a single diverged case in the hundreds of simulations performed in this study.

7.1.2 LES setup and analysis

Constructing a INCA LES setup with an adequate balance between accuracy and computational resources available was a challenging part. Starting from the solver, several preconditioning and Krylov subspace methods were tested to increase the speed of the Poisson equation calculations. This resulted in the implementation of the Stone's Strongly Implicit Procedure as a preconditioner for a BiCGSTAB algorithm, which accelerated the procedure around 5-6 times with respect to multigrid preconditioners.

The study of spanwise extent, grid resolution and averaging time allowed to conclude that the large flow structures present in the geometry required an extent of around $4.5 L_z$ to be correctly represented, in accordance with the results from Fröhlich et al. (2005). A grid with around 10^6 nodes was sufficient to present good agreement with PIV (Rapp, 2009) and reference LES (Breuer et al., 2009) data. Furthermore, extending the simulation to a minimum of 55 flow-through times for the averaging period was considered necessary to compute reliable statistics, with 55 to 89 initial flow-through times for the development of the solution. The final configuration resulted in similar differences with respect to PIV in mean velocity components and turbulent stresses than the solution of Breuer et al. (2009), with around 40% of its averaging time and 10% of its grid points. This allowed to perform one LES simulation on the TU Delft HPC with a single node and 20 cores within around 35 to 75 CPU hours, depending on the geometry considered.

7.1.3 Non-locally injected HIRANS prediction

The prediction of flow quantities by using non-local LES information has demonstrated, despite its limitations, to provide consistent results over the envelope of the parameter space considered, the width of the hill ψ . The formulation was tested in the most limiting cases in 5.3, where information from neighbouring LES cases had very different flow features. Even in those configurations, the method succeeded to qualitatively improve the turbulent stresses in a L^2 metric. A study on the distance threshold to LES neighbours to outperform RANS yielded an improvement of turbulent stresses in every single test considered, even for LES samples separated by 50% of the parameter space. With respect to the velocity components, their improvement with respect to RANS depended on the part of the parameter space considered. On the steeper geometries, the average distance threshold to a LES neighbour to outperform RANS was around $d_{neighb} \approx (0.05, 0.12)$ in the L^2 norms. In contrast, in the less steep configurations, no threshold was found, consistently outperforming or being qualitatively similar to RANS.

In the initial sampling plan with $n_e = 6$ samples, the single fidelity Kriging HIRANS predictor decreased the L^2 norm of the objective function J with respect to the underlying function to less than a third part of the Kriging RANS surrogate, allowing to successfully exploit J from the first function evaluation. A remarkable property of the method was its lack of significant noise along the parameter space for the initial sampling plans, setting very smooth regression guides for the Co-Kriging surrogates.

In terms of the research question:

Which interpolation techniques provide a more robust set of tools within the computational budget? How well can they resemble the higher fidelity features in the hybrid model within the whole design envelope?

All the tested interpolation techniques to provide non-local predictions have a similar cost in their application. Expectedly, interpolating LES information from only neighbouring cases in the parameter space achieved better results, whilst using all the available cases could produce non-physical flow features. Linear weightings were computed using Euclidean distances in the hill width ψ parameter space, in the bottom wall geometry and in Kriging surrogates of k and f_p . Similar results were obtained with the parameter and Kriging methods, being the former slightly superior. The main limitation was found to be in the linear weightings procedure. Similar flow features, as the number of recirculating regions, were concluded to be highly desirable in order to obtain accurate representations.

7.1.4 Kriging and Co-Kriging optimization

Combining information from RANS and LES in single and multi-fidelity design optimization processes was the last step of this work. By using a scaled version of the hill width ψ as a parameter, with a combination of turbulent mixing and total pressure losses as the objective function J , five optimization methods for two different definitions of J were computed. One of the main premises of the present document was confirmed: a single-fidelity Kriging RANS surrogate was not able to qualitatively improve J , as it predicted a wrong minimum location. Conversely, its multi-fidelity version achieved the best improvements in both routines, partially based on a “lucky guess” which found the global minimum in

one of the routines, and its overly exploratory properties. In average, the RANS addition to the LES information was even slightly *globally deceiving* the multi-fidelity surrogate with respect to Kriging LES in terms of the L^2 norm with respect to the “true” underlying function. A regressing Kriging surrogate with $n_e = 59$ samples showed that the quantitative results might be significantly affected by noise in calculating J for the LES simulations.

The single-fidelity Kriging HIRANS was successful exploiting J in 1-2 expensive function evaluations, but it was not able to sufficiently explore outside a narrow region close to its predicted minimum. Co-Kriging with HIRANS information reduced the initial L^2 norm of the surrogate by the regression of the cheap samples, but yielded less improvements of J . This model was overly exploitative and was oblivious of the loss of HIRANS accuracy when walking away from the locally injected LES samples vicinities, producing artificially low predicted errors. A modified Kriging believer simulation was performed using a multi-level HIRANS model with a combination of local and non-locally injected data. This corrected model outperformed the original Co-Kriging HIRANS, both in terms of the surrogate error and the J exploitation, sampling in similar locations as in the single-fidelity Kriging LES method.

Answering the research questions:

What is the optimum amount of resources to be spent in higher and lower fidelity computations to outperform a single-fidelity method? How is this balance linked with the accuracy of the corrections of the hybrid model?

The influence of initial expensive and cheap evaluations was discussed. A sufficiently space-filling initial sampling plan for the expensive data was highly relevant in this test case, with a minimum of 9 data points needed to represent the shape of J without failing in the prediction of large valleys. In contrast, the value of $n_c = 48 + n_e$ was excessive in terms of computational efficiency. Lower values such as 24 and even 12 cheap points for the initial sampling, $n_e = 6$, resulted in very similar predicted responses \hat{y} . A value of around $n_c = \{2, 3\}n_e$ has been proposed, which could be augmented with increasing expensive samples and hence, improved non-locally injected HIRANS predictions. The balance is highly dependent on the local accuracy of the HIRANS evaluations. In case that a large number of samples is computed in regions far away from LES locations, where its accuracy is resented, it could deceive the predictor, both in terms of its response \hat{y} and the non-realistic reduced predicted error. One of the core future lines of research that would lead to an efficient balance of resources is the representation of the accuracy of the non-local HIRANS predictions along the parameter space, to be used as an input for the surrogate model.

Could this method be generalized to other problems, such as multi-parameter optimization for the periodic hill test case?

The knowledge gathered in this document suggests that the answer to this question is affirmative. The HIRANS formulation is not dependent on the dimension of the parameter space, only requiring neighbouring high fidelity information to use as an input. If these conditions are met, the formulation has shown to provide a more accurate initial response of the underlying function than a single fidelity LES predictor. It is not exempt of limitations, such as the need to sample each corner of the parameter space to provide interpolation of quantities or using other methods to substitute the data in these locations.

In an environment where more search directions exist, such as in a high-dimensional parameter space, this initial accuracy improvement could be an even more relevant feature of the model than in the tested case of this document.

Related to the last reflection, in many optimization problems there is a large number of design parameters, but not such a considerable change in the flow dynamics as in the scope of this document. In a case of an airfoil optimization, the location of separated flow or a shockwave in transonic/supersonic cases may be the main flow topology change, with a singular feature that varies in location and strength. Those situations could be potential test cases for the HIRANS model, given its positive results when similar flow features are used for interpolation. Also, it could be suitable for channel/constricted flows design, given its good agreement with LES in the computation of turbulent kinetic energy and total pressure losses.

7.1.5 Main research question

Does the combination of a hybrid LES-data-injected RANS solver with pure LES in a Co-Kriging multi-fidelity global design optimization process outperform a single fidelity method for a fixed computational budget in the context of turbulent incompressible fluid flows for the specified periodic hill test cases?

The response to this question is currently negative if constrained to the raw results of the tests performed. Co-Kriging HIRANS successfully found a new minimum in two expensive function evaluations, but was outperformed by Kriging LES and Co-Kriging RANS in the entire processes, albeit with different exploration-exploitation bias with respect to the single-fidelity model.

A more refined analysis outlines that the multi-fidelity surrogate was able to give more accurate initial *response* predictions, but failed in the characterization of the *predicted errors* with respect to the true underlying function. The error correction with the Kriging believer simulation showed that even a simple modification of the predicted mean square error \hat{s} by a multi-level approach could put it on par, qualitatively, with the single-fidelity surrogate, both in terms of results and computational efficiency. The current framework certainly offers potential advantages, and future research in a better representation of the trustworthiness of the predictions in different regions away from high fidelity locations could be able to unlock them.

7.2 Recommendations

7.2.1 Local RANS corrections

In terms of the HIRANS formulation, four main proposals can be formulated.

On the first place, the framework could be extended to modify additional terms of the closure equations for the turbulence model. Currently, only the production term is corrected, using a hybrid tensor τ^* which is scaled with the turbulent kinetic energy computed by HIRANS. Formulations which uniquely depend on the LES information injected, or which affect terms such as the diffusion or the dissipation rate could be tested, in order to provide a more accurate representation of the LES dynamics. Additionally, albeit $k - \omega$ has been successful in its local corrections, alternative turbulence models could be tested, such as Reynolds-stress models (Hallback et al., 1989), introducing the high-fidelity data in the equations for the individual stress components.

A second consideration must be made in cases where, as the INCA tested cases, the LES grid is coarser than HIRANS. Throughout this document, linear interpolation has been used in this respect, but this condition clearly does not represent the reality close to the walls, which increases the difficulty of providing an accurate representation of the quantities in that region, such as the wall shear stress. A modification to the current algorithm could be defined, where a wall function scaling would substitute the linear interpolation based on the original information calculated in a fine LES grid or from similar reference data. Dynamic injection levels close to the walls could also be attempted, albeit the results from the present document in that respect predict difficulties in that respect.

With respect to this, a study about the accuracy of the HIRANS solution with varying high fidelity sources, such as DNS and LES or LES with different grids would serve as a way to quantify the importance of the input data in the trustworthiness of the model. This has been outlined in this study, as the correction using the information from Breuer et al. (2009) achieved closer results to its source data than the one using the coarser INCA grid.

Finally, the HIRANS model has been computed for cases ranging from massively separated flows with counter-rotating structures to situations with very low values of turbulent kinetic energy. Testing the formulation with alternative flow features and phenomena such as laminar-turbulent transition and turbulence development (swept wings with crossflow instabilities and vortex generators) or sharp gradients in flow quantities (transonic compressible flows) would improve the knowledge about the boundaries in the applicability of the formulation.

7.2.2 HIRANS interpolation and prediction

In this respect, two main paths could be explored. The first one corresponds to a more sophisticated method to interpolate neighbouring cases with respect to the linear weightings procedure presented in this document, such as basis function or Kriging formulations.

A second topic, especially relevant in multi-parameter optimization, would be related to a suitable choice of neighbouring cases. Methods such as nonlinear manifold learning, Isomap (Tenenbaum et al., 2000; Franz et al., 2014) could be introduced, by the computation of *geodesic* instead of Euclidean distances in the parameter space.

7.2.3 Kriging and Co-Kriging optimization

In the global design optimization side, there are three main aspects to be discussed.

The first one corresponds to the Co-Kriging framework. To correct the estimations for the predicted mean squared error, approaches can be tested containing estimations of the HIRANS model with non-local information, such as the multi-level formulation with two HIRANS layers. This model could be updated every expensive iteration by the resampling of the HIRANS data, if this one is significantly cheaper to compute than the expensive one, or if it provides a key benefit to the quality of the surrogate. Besides, a model which takes into account the uncertainty of the input data, such as the individual regularization parameter per sample proposed in Appendix H, could provide a better balance between interpolation and regression.

The methods in this document included a maximum number of acquired samples of $n_{e_{add}} = 6$ samples. A longer process would answer the question of the advantages of the HIRANS formulation once the distances between LES cases become smaller. The next intuitive step would be to augment the parameter space to multiple dimensions. The current periodic hill parametrization allows a total of 11 degrees of freedom. A test with several dimensions would challenge the initial accuracy improvement of the Co-Kriging HIRANS surrogate and its ability to exploit with a small number of samples.

Finally, an alternative way of introducing a large number of dimensions in the parameter space would be an *adjoint* method. As exposed in Blonigan et al. (2016); Blonigan (2016), LES simulations in complex turbulent flows can suffer from diverging sensitivities when running the adjoint computations backwards in time. Instead of stabilizing the LES computations, the adjoint sensitivities could be calculated in a cheaper HIRANS model, either for initial search directions, using locally injected information, or globally, if the non-locally injected model is sufficiently accurate. Furthermore, this information could be combined with Gradient-Enhanced Kriging (Fischer and Grandhi, 2015) or any other global optimization method.

References

- Almeida, G., Durao, D., and Heitor, M. (1993). Wake flows behind two-dimensional model hills. *Experimental Thermal and Fluid Science*, 7(1):87–101.
- Amsallem, D., Zahr, M., Choi, Y., and Farhat, C. (2015). Design optimization using hyper-reduced-order models. *Structural and Multidisciplinary Optimization*, 51(4):919–940.
- Amsallem, D., Zahr, M. J., and Farhat, C. (2012). Nonlinear model order reduction based on local reduced-order bases. *International Journal for Numerical Methods in Engineering*, 92(10):891–916.
- Balakumar, P., Park, G., and Pierce, B. (2014). DNS, LES, and wall-modeled LES of separating flow over periodic hills. In *Proceedings of the Summer Program*, pages 407–415.
- Baldwin, B. and Lomax, H. (1978). Thin-layer approximation and algebraic model for separated turbulent flows. In *16th aerospace sciences meeting*, page 257.
- Banerjee, S., Krahl, R., Durst, F., and Zenger, C. (2007). Presentation of anisotropy properties of turbulence, invariants versus eigenvalue approaches. *Journal of Turbulence*, (8):N32.
- Berkooz, G., Holmes, P., and Lumley, J. L. (1993). The Proper Orthogonal Decomposition in the analysis of turbulent flows. *Annual review of fluid mechanics*, 25(1):539–575.
- Blonigan, P. J. (2016). *Least Squares Shadowing for sensitivity analysis of large chaotic systems and fluid flows*. PhD thesis, Massachusetts Institute of Technology.
- Blonigan, P. J., Wang, Q., Nielsen, E. J., and Diskin, B. (2016). Least Squares Shadowing sensitivity analysis of chaotic flow around a two-dimensional airfoil. In *54th AIAA Aerospace Sciences Meeting*, page 0296.
- Braconnier, T., Ferrier, M., Jouhaud, J.-C., Montagnac, M., and Sagaut, P. (2011). Towards an adaptive pod/svd surrogate model for aeronautic design. *Computers & Fluids*, 40(1):195 – 209.

- Breuer, M., Peller, N., Rapp, C., and Manhart, M. (2009). Flow over periodic hills—numerical and experimental study in a wide range of Reynolds numbers. *Computers & Fluids*, 38(2):433–457.
- Brooks, C. J., Forrester, A., Keane, A., and Shahpar, S. (2011). Multi-fidelity design optimisation of a transonic compressor rotor. *University of Southampton, Rolls-Royce Aerothermal Methods Group*.
- Carlberg, K., Bou-Mosleh, C., and Farhat, C. (2011). Efficient non-linear model reduction via a least-squares Petrov–Galerkin projection and compressive tensor approximations. *International Journal for Numerical Methods in Engineering*, 86(2):155–181.
- Casacuberta, J., Groot, K. J., Tol, H. J., and Hickel, S. (2018). Effectivity and efficiency of selective frequency damping for the computation of unstable steady-state solutions. *Journal of Computational Physics*, 375:481 – 497.
- Cheung, S. H., Oliver, T. A., Prudencio, E. E., Prudhomme, S., and Moser, R. D. (2011). Bayesian uncertainty analysis with applications to turbulence modeling. *Reliability Engineering & System Safety*, 96(9):1137–1149.
- Choi, K.-S. and Lumley, J. L. (2001). The return to isotropy of homogeneous turbulence. *Journal of Fluid Mechanics*, 436:59–84.
- Cox, D. D. and John, S. (1992). A statistical method for global optimization. In *[Proceedings] 1992 IEEE International Conference on Systems, Man, and Cybernetics*, pages 1241–1246. IEEE.
- Davidson, L. and Peng, S.-H. (2003). Hybrid LES-RANS modelling: a one-equation SGS model combined with a $k - \omega$ model for predicting recirculating flows. *International Journal for Numerical Methods in Fluids*, 43(9):1003–1018.
- Deardorff, J. W. (1974). Three-dimensional numerical study of the height and mean structure of a heated planetary boundary layer. *Boundary-Layer Meteorology*, 7(1):81–106.
- del Estal Herrero, A., Percin, M., Karasek, M., and van Oudheusden, B. (2018a). Flow visualization around a flapping-wing micro air vehicle in free flight.
- del Estal Herrero, A., Percin, M., Karasek, M., and van Oudheusden, B. (2018b). Flow visualization around a flapping-wing micro air vehicle in free flight using large-scale piv. *Aerospace*, 5(4).
- Duraisamy, K. (2015). A framework for turbulence modeling using Big Data. *NASA Aeronautics Research Mission Directorate (ARMD), 2015 LEARN/Seedling Technical Seminar*.
- Duraisamy, K., Singh, A. P., and Zhang, Z. J. (2017). Augmentation of turbulence models using field inversion and machine learning. In *55th AIAA Aerospace Sciences Meeting*, page 0993.
- Edeling, W., Cinnella, P., and Dwight, R. P. (2014a). Predictive RANS simulations via Bayesian model-scenario averaging. *Journal of Computational Physics*, 275:65–91.

- Edeling, W., Cinnella, P., Dwight, R. P., and Bijl, H. (2014b). Bayesian estimates of parameter variability in the $k - \varepsilon$ turbulence model. *Journal of Computational Physics*, 258:73–94.
- Fan, T., Tian, M., Edwards, J., Hassan, H., and Baurle, R. (2001). Validation of a hybrid Reynolds-averaged/large-eddy simulation method for simulating cavity flameholder configurations. In *15th AIAA Computational Fluid Dynamics Conference*, page 2929.
- Ferziger, J. H. and Peric, M. (2012). *Computational methods for fluid dynamics*. Springer Science & Business Media.
- Fischer, C. C. and Grandhi, R. V. (2015). A surrogate-based adjustment factor approach to multi-fidelity design optimization. In *17th AIAA Non-Deterministic Approaches Conference*, page 1375.
- Forrester, A. I., Sóbester, A., and Keane, A. J. (2007). Multi-fidelity optimization via surrogate modelling. In *Proceedings of the royal society of london a: mathematical, physical and engineering sciences*, volume 463, pages 3251–3269. The Royal Society.
- Forrester, A. I., Sóbester, A., and Keane, A. J. (2008). *Engineering design via surrogate modelling: a practical guide*. John Wiley & Sons.
- Franz, T., Zimmermann, R., Görtz, S., and Karcher, N. (2014). Interpolation-based reduced-order modelling for steady transonic flows via manifold learning. *International Journal of Computational Fluid Dynamics*, 28(3-4):106–121.
- Fröhlich, J. (1990). *Résolution numérique des équations de Navier-Stokes à faible nombre de Mach par méthode spectrale*. PhD thesis, Nice.
- Fröhlich, J., Mellen, C. P., Rodi, W., Temmerman, L., and Leschziner, M. A. (2005). Highly resolved large-eddy simulation of separated flow in a channel with streamwise periodic constrictions. *Journal of Fluid Mechanics*, 526:19–66.
- Gardner, J. R., Kusner, M. J., Xu, Z. E., Weinberger, K. Q., and Cunningham, J. P. (2014). Bayesian optimization with inequality constraints. In *ICML*, pages 937–945.
- Ginsbourger, D., Le Riche, R., and Carraro, L. (2010). Kriging is well-suited to parallelize optimization. In *Computational intelligence in expensive optimization problems*, pages 131–162. Springer.
- Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1st edition.
- Görtler, H. (1940). Über eine dreidimensionale instabilität laminaren grenzschichten am konkaven wander. *Naschr Wiss Gas, Gottingen Math Phys Klasse*, 2(1).
- Guillas, S., Glover, N., and Malki-Epshtein, L. (2014). Bayesian calibration of the constants of the $k - \varepsilon$ turbulence model for a CFD model of street canyon flow. *Computer Methods in Applied Mechanics and Engineering*, 279:536–553.

- Günther, S., Gauger, N. R., and Wang, Q. (2016). Simultaneous single-step one-shot optimization with unsteady PDEs. *Journal of Computational and Applied Mathematics*, 294:12–22.
- Günther, S., Gauger, N. R., and Wang, Q. (2017). A framework for simultaneous aerodynamic design optimization in the presence of chaos. *Journal of Computational Physics*, 328:387–398.
- Hallback, M., Groth, J., and Johansson, A. (1989). A Reynolds stress closure for the dissipation in anisotropic turbulent flows. In *7th Symposium on Turbulent Shear Flows, Volume 2*, volume 2, pages 17–2.
- Hickel, S. and Adams, N. (2007). On implicit subgrid-scale modeling in wall-bounded flows. *Physics of Fluids*, 19(10):105106.
- Hickel, S., Adams, N. A., and Domaradzki, J. A. (2006). An adaptive local deconvolution method for implicit les. *Journal of Computational Physics*, 213(1):413–436.
- Hickel, S., Kempe, T., and Adams, N. (2008). Implicit large-eddy simulation applied to turbulent channel flow with periodic constrictions. *Theoretical and Computational Fluid Dynamics*, 22(3-4):227–242.
- Hoerl, A. E. and Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67.
- Jones, D. R., Schonlau, M., and Welch, W. J. (1998). Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492.
- Jones, W. and Launder, B. E. (1972). The prediction of laminarization with a two-equation model of turbulence. *International journal of heat and mass transfer*, 15(2):301–314.
- Kennedy, M. C. and O’Hagan, A. (2000). Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87(1):1–13.
- Klostermeier, C. (2001). Correlation analysis of large structures on LES of periodic hill flow. Technical report, Internal report, Institut für Hydromechanik, Universität Karlsruhe.
- Krige, D. G. (1951). A statistical approach to some basic mine valuation problems on the witwatersrand. *Journal of the Southern African Institute of Mining and Metallurgy*, 52(6):119–139.
- Kushner, H. J. (1964). A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *Journal of Basic Engineering*, 86(1):97–106.
- Launder, B. E., Reece, G. J., and Rodi, W. (1975). Progress in the development of a reynolds-stress turbulence closure. *Journal of fluid mechanics*, 68(3):537–566.
- Launder, B. E. and Sharma, B. (1974). Application of the energy-dissipation model of turbulence to the calculation of flow near a spinning disc. *Letters in heat and mass transfer*, 1(2):131–137.

- Leonard, A. (1975). Energy cascade in large-eddy simulations of turbulent fluid flows. In *Advances in geophysics*, volume 18, pages 237–248. Elsevier.
- Lilly, K. (1966). The representation of small-scale turbulence in numerical simulation experiments. *National Center for Atmospheric Research*.
- Lizotte, D. J. (2008). *Practical Bayesian optimization*. PhD thesis, University of Alberta.
- Lumley, J. L. (1979). Computational modeling of turbulent flows. In *Advances in applied mechanics*, volume 18, pages 123–176. Elsevier.
- Lumley, J. L. and Newman, G. R. (1977). The return to isotropy of homogeneous turbulence. *Journal of Fluid Mechanics*, 82(1):161–178.
- Malouf, R. (2002). A comparison of algorithms for maximum entropy parameter estimation. In *proceedings of the 6th conference on Natural language learning-Volume 20*, pages 1–7. Association for Computational Linguistics.
- Marsden, A. L., Wang, M., Dennis, J., and Moin, P. (2007). Trailing-edge noise reduction using derivative-free optimization and large-eddy simulation. *Journal of Fluid Mechanics*, 572:13–36.
- Matheron, G. (1962). *Traité de géostatistique appliquée. 1 (1962)*, volume 1. Editions Technip.
- Mellen, C. P., Fröhlich, J., and Rodi, W. (2000). Large-eddy simulation of the flow over periodic hills. In: *16th IMACS World Congress. Lausanne, Switzerland*.
- Menter, F. R. (1993). Zonal two equation $k - \omega$ turbulence models for aerodynamic flows. In *23rd fluid dynamics, plasmadynamics, and lasers conference*, page 2906.
- Micchelli, C. A. (1984). Interpolation of scattered data: distance matrices and conditionally positive definite functions. In *Approximation theory and spline functions*, pages 143–145. Springer.
- Moroney, T. J. (2006). *An investigation of a finite volume method incorporating radial basis functions for simulating nonlinear transport*. PhD thesis, Queensland University of Technology.
- O’Hagan, A. and Kingman, J. (1978). Curve fitting and optimal design for prediction. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 1–42.
- Özkaya, E. and Gauger, N. R. (2009). Single-step one-shot aerodynamic shape optimization. In *Optimal control of coupled systems of partial differential equations*, pages 191–204. Springer.
- Peskin, C. S. (1972). Flow patterns around heart valves: A numerical method. *Journal of Computational Physics*, 10(2):252 – 271.
- Pham, H. T., Yang, B.-S., et al. (2010). A hybrid of nonlinear autoregressive model with exogenous input and autoregressive moving average model for long-term machine state forecasting. *Expert Systems with Applications*, 37(4):3310–3317.

- Phillips, W. and Wu, Z. (1994). On the instability of wave-catalysed longitudinal vortices in strong shear. *Journal of Fluid Mechanics*, 272:235–254.
- Poggio, T. and Girosi, F. (1990). Regularization algorithms for learning that are equivalent to multilayer networks. *Science*, 247(4945):978–982.
- Pope, S. B. (2000). *Turbulent Flows*. Cambridge University Press.
- Poroseva, S. V., Colmenares F, J. D., and Murman, S. M. (2016). On the accuracy of RANS simulations with DNS data. *Physics of Fluids*, 28(11):115102.
- Rapp, C. (2009). *Experimentelle Studie der turbulenten Strömung über periodische Hügel*. PhD thesis, Technische Universität München.
- Sacks, J., Welch, W. J., Mitchell, T. J., and Wynn, H. P. (1989). Design and analysis of computer experiments. *Statistical science*, pages 409–423.
- Šarić, S., Jakirlić, S., Breuer, M., Jaffrézic, B., Deng, G., Chikhaoui, O., Fröhlich, J., Von Terzi, D., Manhart, M., and Peller, N. (2007). Evaluation of detached eddy simulations for predicting the flow over periodic hills. In *ESAIM: proceedings*, volume 16, pages 133–145. EDP Sciences.
- Sasena, M. J. (2002). *Flexibility and efficiency enhancements for constrained global design optimization with Kriging approximations*. PhD thesis, General Motors.
- Schmid, P. J. (2010). Dynamic mode decomposition of numerical and experimental data. *Journal of fluid mechanics*, 656:5–28.
- Smagorinsky, J. (1963). General circulation experiments with the primitive equations: I. the basic experiment. *Monthly weather review*, 91(3):99–164.
- Smith, A. and Cebeci, T. (1967). Numerical solution of the turbulent-boundary-layer equations. Technical report, Douglas Aircraft Co, Long Beach CA, Aircraft Division.
- Snoek, J., Larochelle, H., and Adams, R. P. (2012). Practical Bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, pages 2951–2959.
- Spalart, P. and Allmaras, S. (1992). A one-equation turbulence model for aerodynamic flows. In *30th aerospace sciences meeting and exhibit*, page 439.
- Spalart, P., Jou, W., Strelets, M., and Allmaras, S. (1997). Comments on the feasibility of LES for wings, and on a hybrid RANS/LES approach. *Advances in DNS/LES*, 1:4–8.
- Spalart, P. R. (2009). Detached-eddy simulation. *Annual review of fluid mechanics*, 41:181–202.
- Speziale, C. G., Sarkar, S., and Gatski, T. B. (1991). Modelling the pressure–strain correlation of turbulence: an invariant dynamical systems approach. *Journal of fluid mechanics*, 227:245–272.
- Stone, H. L. (1968). Iterative solution of implicit approximations of multidimensional partial differential equations. *SIAM Journal on Numerical Analysis*, 5(3):530–558.

- Talnikar, C., Blonigan, P. J., Bodart, J., and Wang, Q. (2014). Parallel optimization for LES. In *Proceedings of the Summer Program*, page 315.
- Talnikar, C., Blonigan, P. J., Bodart, J., and Wang, Q. (2015a). Optimization with LES – algorithms for dealing with sampling error of turbulence statistics. In *53rd AIAA Aerospace Sciences Meeting*, page 1954.
- Talnikar, C., Wang, Q., and Laskowski, G. M. (2015b). Unsteady adjoint of pressure loss for a fundamental transonic turbine vane. *ArXiv Preprint*.
- Team, P. C. (2018). Python: A dynamic, open source programming language. *Python Software Foundation*. URL <https://www.python.org/>.
- Temmerman, L., Hadžiabdić, M., Leschziner, M., and Hanjalić, K. (2005). A hybrid two-layer URANS–LES approach for large eddy simulation at high Reynolds numbers. *International journal of heat and fluid flow*, 26(2):173–190.
- Tenenbaum, J. B., De Silva, V., and Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323.
- Tikhonov, A. N. and Arsenin, V. Y. (1977). *Solutions of Ill-posed problems*. W.H. Winston.
- Van der Vorst, H. A. and Vuik, C. (1994). GMRESR: a family of nested GMRES methods. *Numerical linear algebra with applications*, 1(4):369–386.
- van Rossum, G. (1995). Python tutorial. Technical Report CS-R9526, Centrum voor Wiskunde en Informatica (CWI), Amsterdam.
- Vuik, C. and Lahaye, D. (2012). Scientific computing (WI4201). *Lecture notes for WI4201, TU Delft*.
- Wales, D. J. and Doye, J. P. (1997). Global optimization by basin-hopping and the lowest energy structures of lennard-jones clusters containing up to 110 atoms. *The Journal of Physical Chemistry A*, 101(28):5111–5116.
- Wang, Z., Hutter, F., Zoghi, M., Matheson, D., and de Feitas, N. (2016). Bayesian optimization in a billion dimensions via random embeddings. *Journal of Artificial Intelligence Research*, 55:361–387.
- Weller, H., Greenshields, C., Santos, B., Bainbridge, W., Janssens, M., et al. (2018). OpenFOAM website. www.openfoam.org.
- Weller, H. G., Tabor, G., Jasak, H., and Fureby, C. (1998). A tensorial approach to computational continuum mechanics using object-oriented techniques. *Computers in Physics*, 12(6):620–631.
- Wilcox, D. C. (1988). Reassessment of the scale-determining equation for advanced turbulence models. *AIAA journal*, 26(11):1299–1310.
- Wilcox, D. C. (1998). *Turbulence modeling for CFD*, volume 2. DCW industries La Canada, CA, 2nd edition.

-
- Wu, J., Xiao, H., Sun, R., and Wang, Q. (2018). RANS equations with Reynolds stress closure can be ill-conditioned. *arXiv preprint arXiv:1803.05581*.
- Xiao, M., Breikopf, P., Coelho, R. F., Knopf-Lenoir, C., Sidorkiewicz, M., and Villon, P. (2010). Model reduction by CPOD and Kriging. *Structural and multidisciplinary optimization*, 41(4):555–574.

Appendices

Appendix A

Injection factor sweep results

In this appendix, a full sweep of the injection factors ξ_m , ξ_t is calculated for different variables, for the case of HIRANS_b, that is, HIRANS with τ_{LES} injection. Figures A.1, A.2, A.3 and A.4 portray this. The injection factor ξ_m seems to be a positive influence in all the norms, excepting the L^∞ for $\langle U_y \rangle$, while the effect of ξ_t depends on the variable considered.




Figure A.1: L^2 (left) and L^∞ (right) norms of τ_w for the bottom wall for HIRANS with τ_{LES} injection. Contours linearly interpolated from 81 sample points (\times).

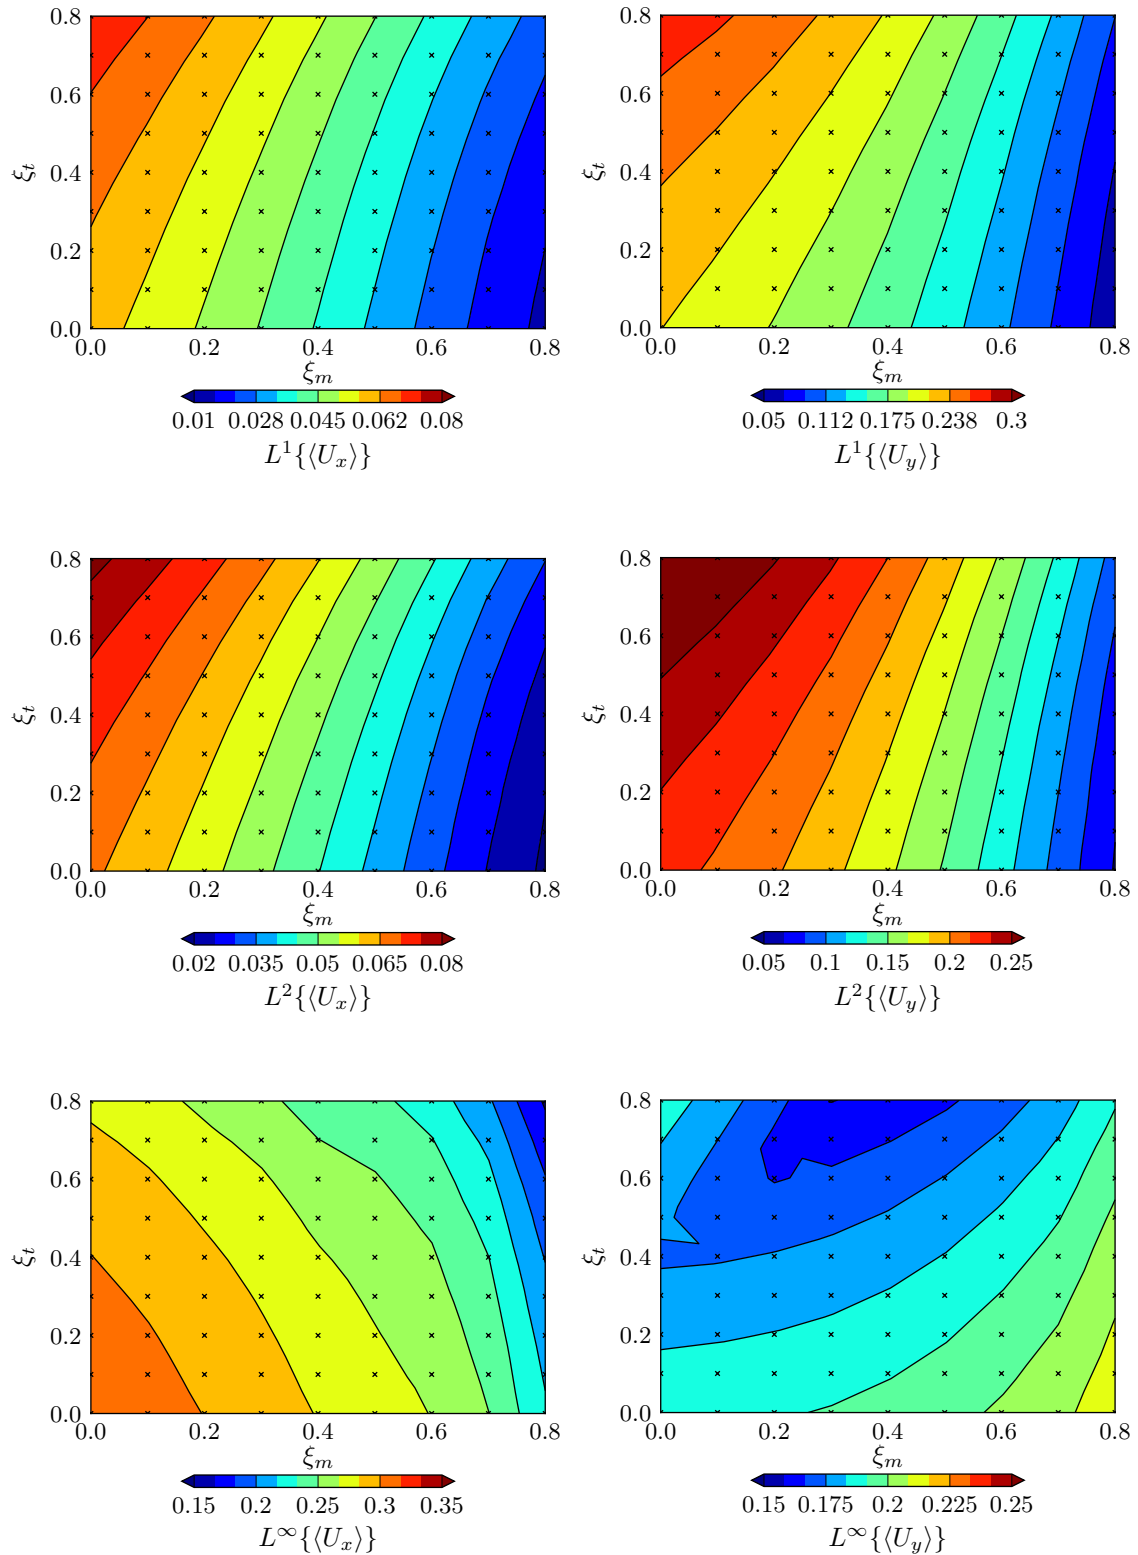


Figure A.2: L^1 (top), L^2 (center) and L^∞ (bottom) norms of $\langle U_x \rangle$ (left) and $\langle U_y \rangle$ (right) for HIRANS with τ_{LES} injection. Contours linearly interpolated from 81 sample points (\times).

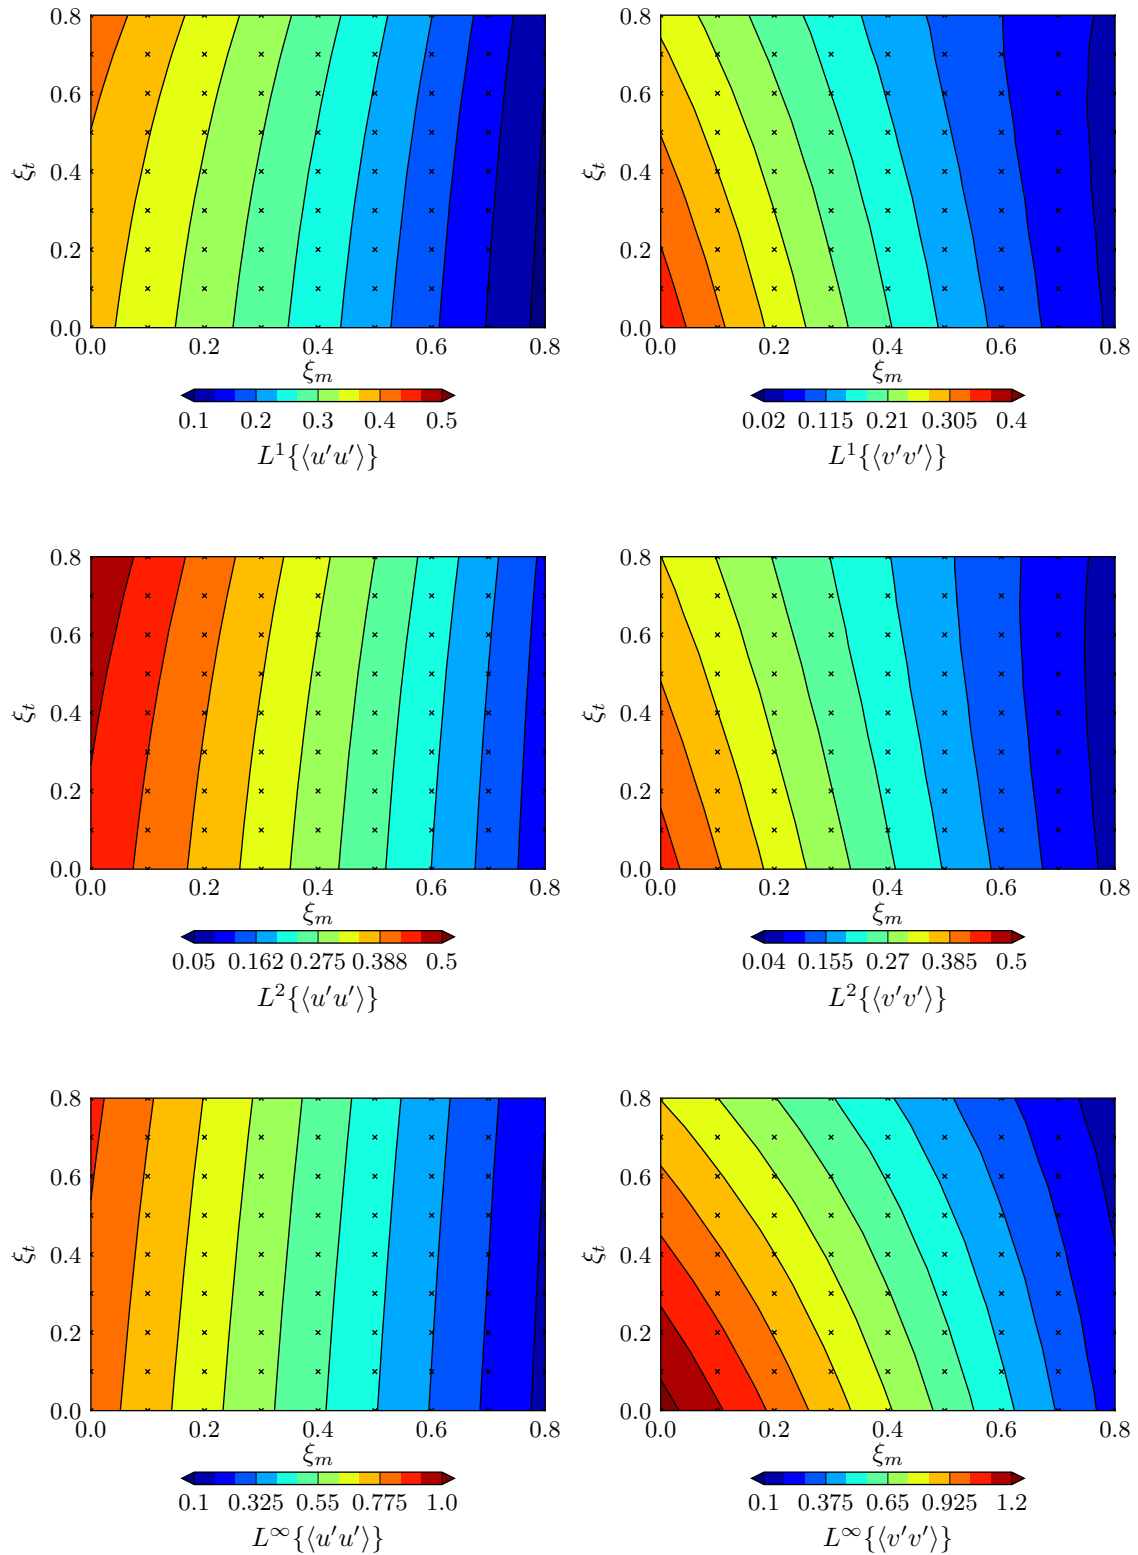


Figure A.3: L^1 (top), L^2 (center) and L^∞ (bottom) norms of $\langle u'u' \rangle$ (left) and $\langle v'v' \rangle$ (right) for HIRANS with τ_{LES} injection. Contours linearly interpolated from 81 sample points (\times).

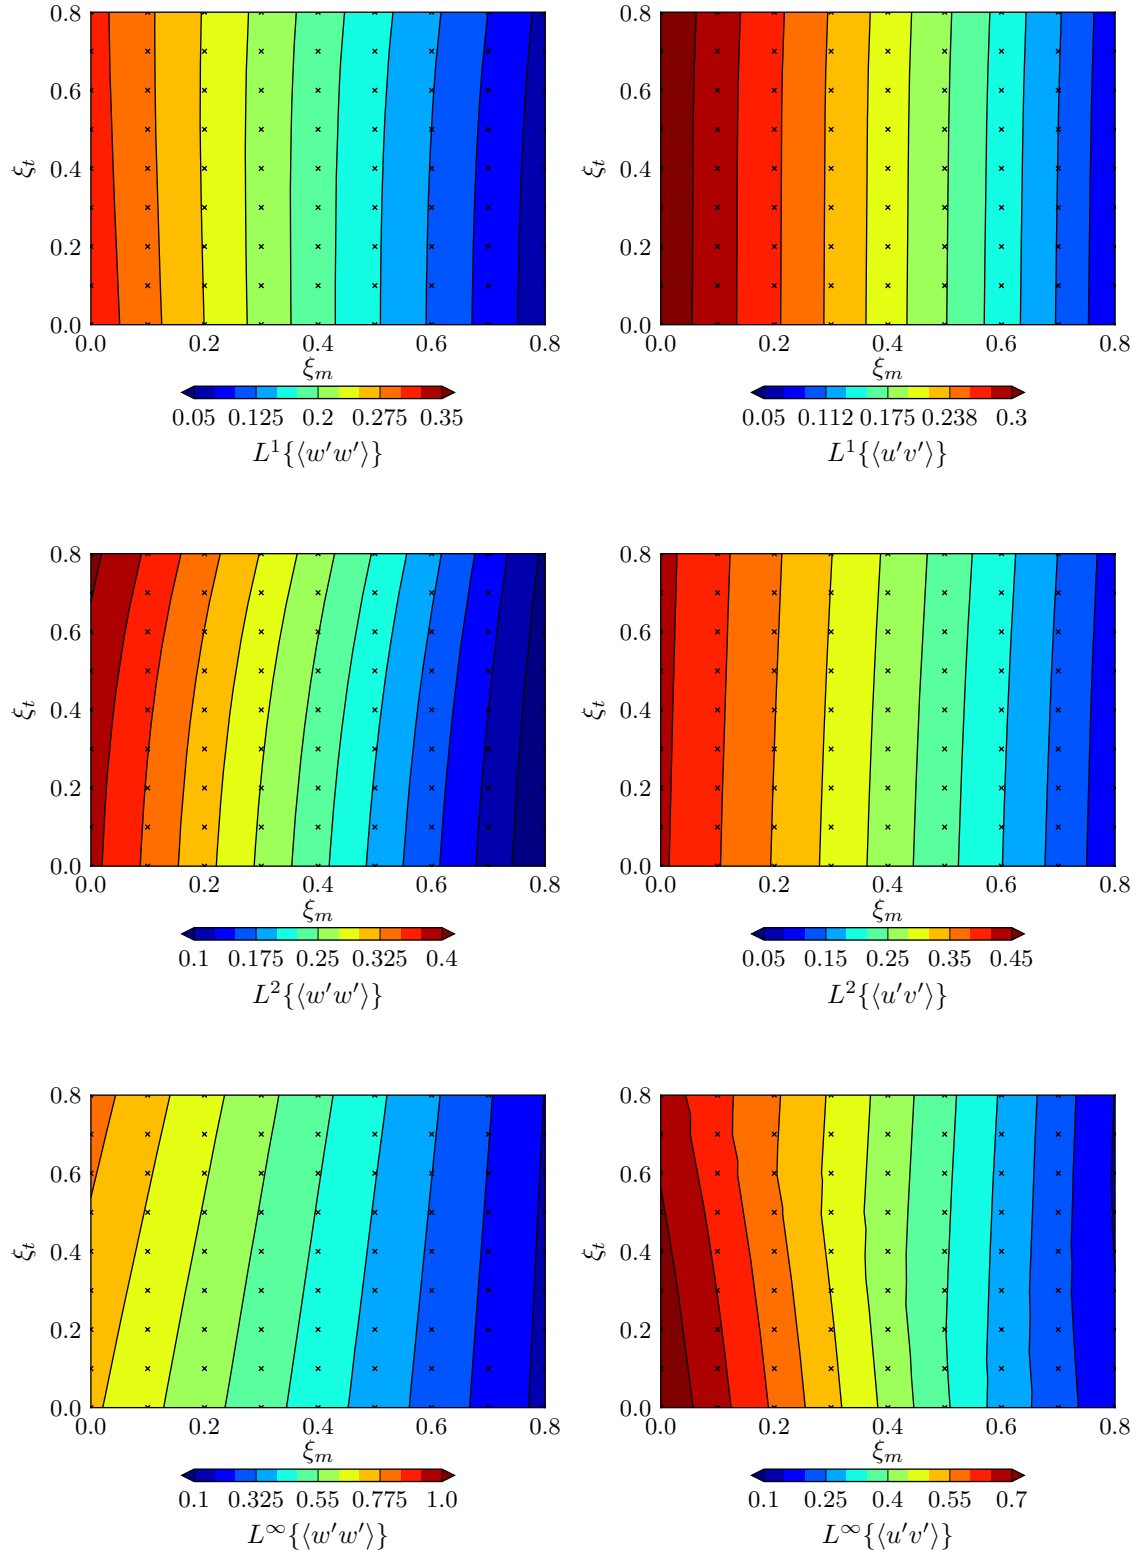


Figure A.4: L^1 (top), L^2 (center) and L^∞ (bottom) norms of $\langle w'w' \rangle$ (left) and $\langle u'v' \rangle$ (right) for HIRANS with τ_{LES} injection. Contours linearly interpolated from 81 sample points (x).

Appendix B

Convergence and results for higher HIRANS injection factors

The convergence and results of the HIRANS model with τ_{LES} injection from [Breuer et al. \(2009\)](#) with higher injection factors than $\xi_i = 0.8$ are shown in this appendix.

In the course of this document, the ramping up of the ξ_m , ξ_t factors has been performed using a value of $t^* = 5000$ iterations, that is, the maximum values of these factors would be achieved 5000 iterations before the simulation ended. If the method was not stable, it would start to diverge after these iterations, creating non-physical values. Nevertheless, it is also possible to not allow the solver any iterations at all, $t^* = 0$ iterations, after the maximum injection factors are reached. While this raises questions about the validity of a solution which would diverge if run for a longer time, it also achieved better results quantitatively with respect to the reference LES values.

Figure [B.1](#) illustrates this. For $t^* = 5000$ iterations, the limiting injection factors were already in the threshold of $\xi_m, \xi_t \approx (0.80, 0.85)$. In contrast, higher factors, such as 0.9 values, can be introduced if no additional iterations are allowed to the solver after reaching the maximum injection values.

With respect to the results, Figure [B.2](#) shows that this method outperformed the $t^* = 5000$ iterations in all the norms, calculated using [\(3.19\)](#), considered. The values for the turbulent stresses were greatly reduced, and the $\langle U_y \rangle$ also improved with respect to the $\xi_m, \xi_t = 0.8$ case. In addition, Figure [B.3](#) shows that the separation region was better represented, while the top wall differences became larger. The k distribution was almost identical to the LES solution. Nonetheless, the divergence in the mean velocity components was evident from the $\{\xi_m, \xi_t\} = \{0.90\}$ case.

Following these results, an upper threshold of 0.80 for the two factors was considered a safe choice for the course of this document. It was successfully applied with $t^* = 5000$ iterations without significantly diverging in any of the hundreds of HIRANS cases computed.

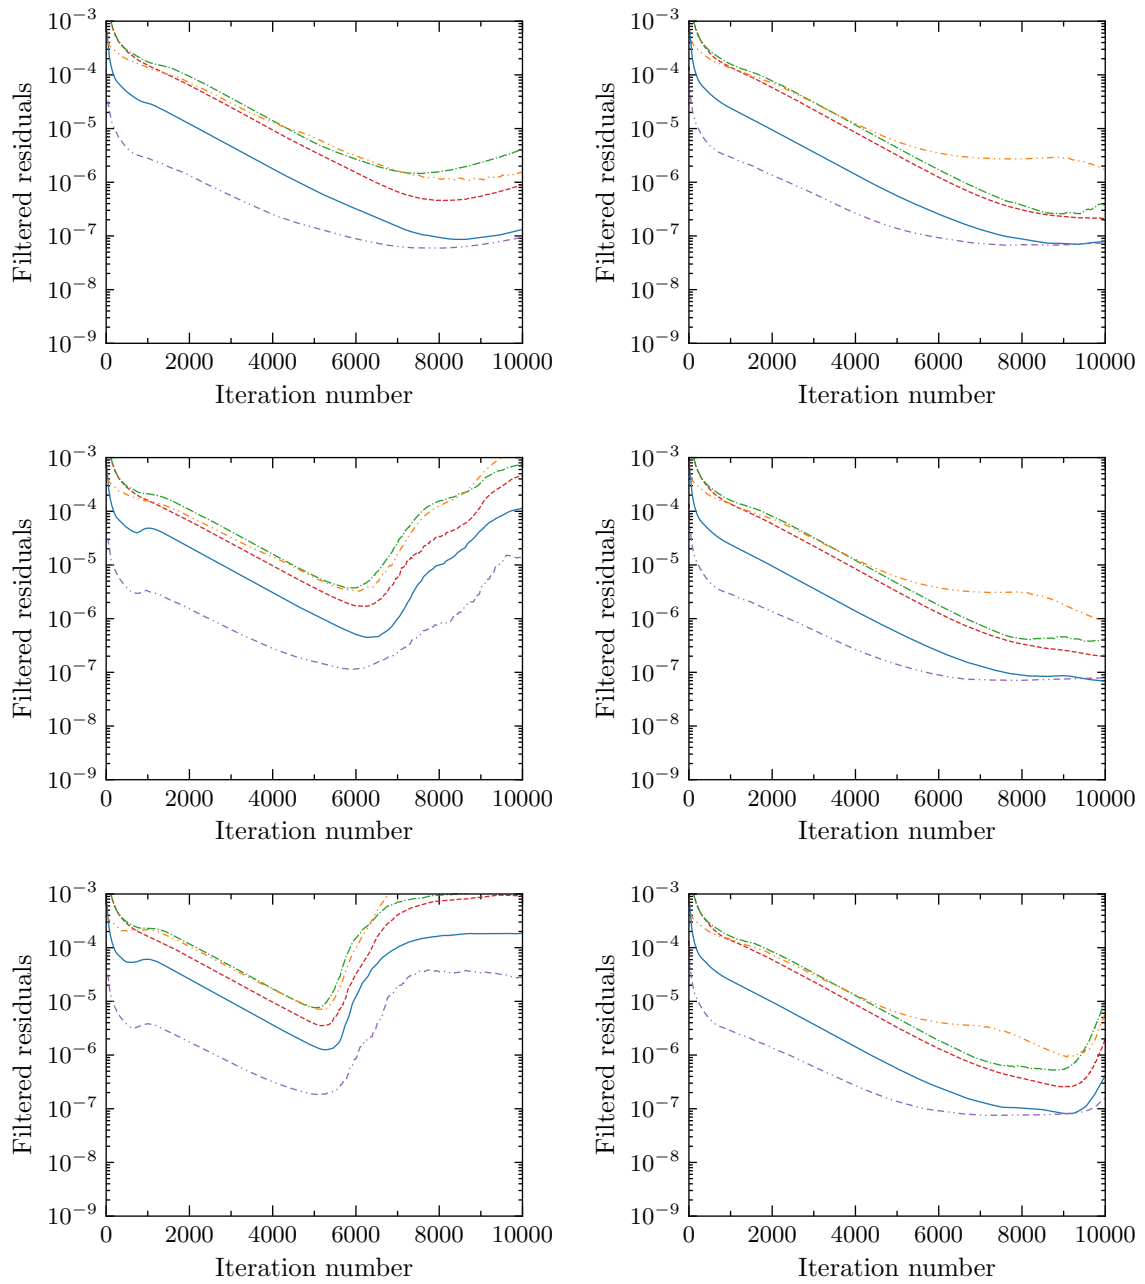


Figure B.1: Filtered residuals for $t^* = 5000$ (left) and $t^* = 0$ (right) iterations, for $\xi_m, \xi_t = 0.85$ (top), $\xi_m, \xi_t = 0.90$ (centre) and $\xi_m, \xi_t = 0.95$ (bottom) of $\langle p \rangle$ (—), $\langle U_x \rangle$ (---), $\langle U_y \rangle$ (-.-.), k (....) and ω (-.-.-).

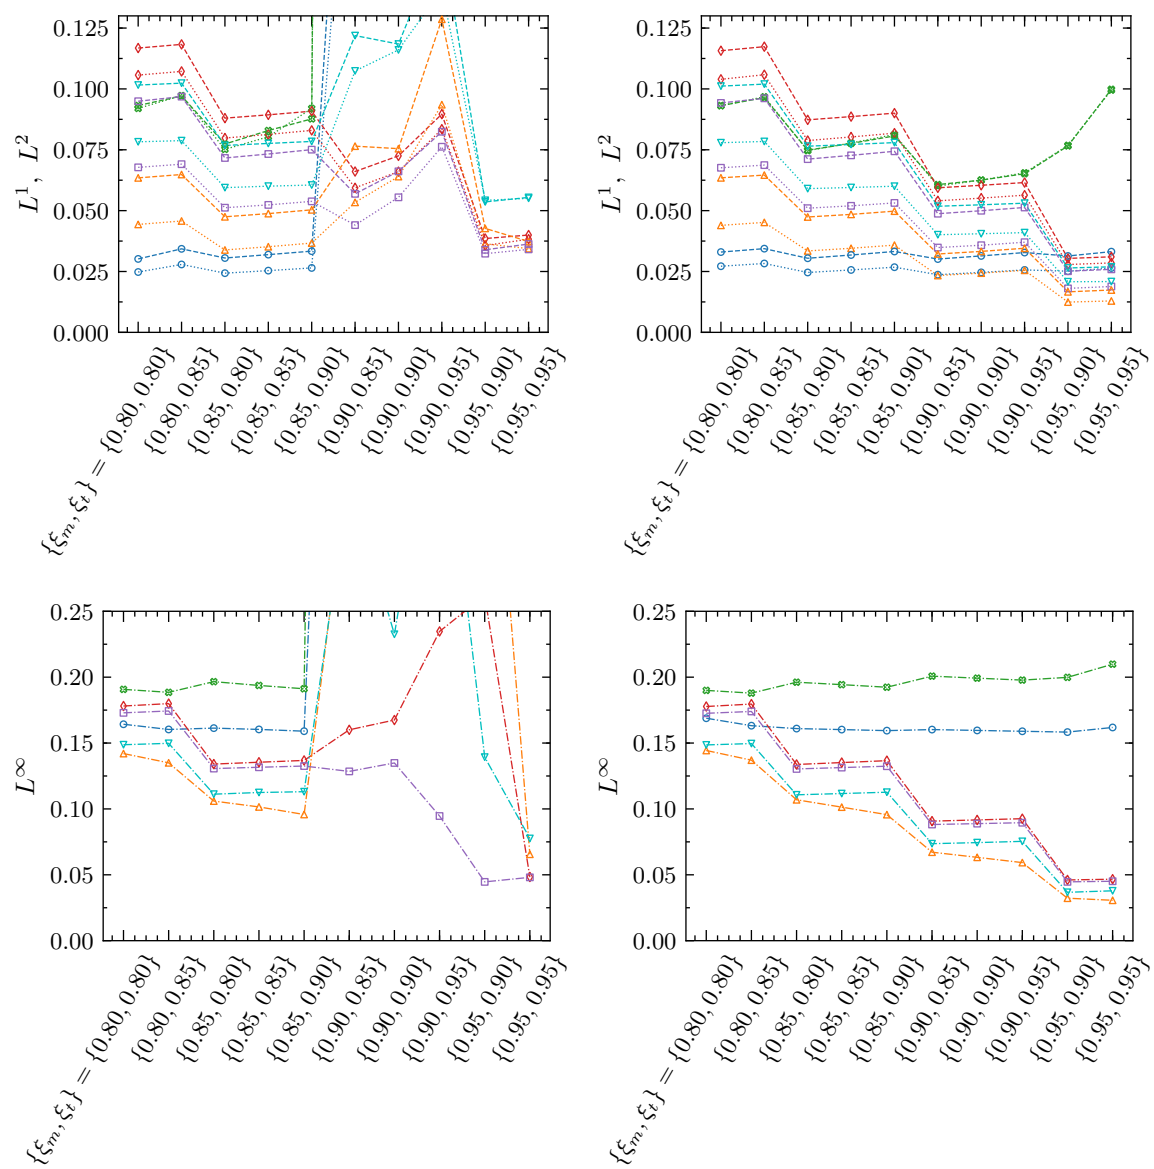


Figure B.2: L^1 (.....), L^2 (----) (top) and L^∞ (-----) (bottom) norms of $\langle U_x \rangle$ (○), $\langle U_y \rangle$ (⊙), $\langle u'u' \rangle$ (◇), $\langle v'v' \rangle$ (△), $\langle w'w' \rangle$ (□) and $\langle u'v' \rangle$ (▽) of HIRANS with $t^* = 5000$ (left) and $t^* = 0$ (right).

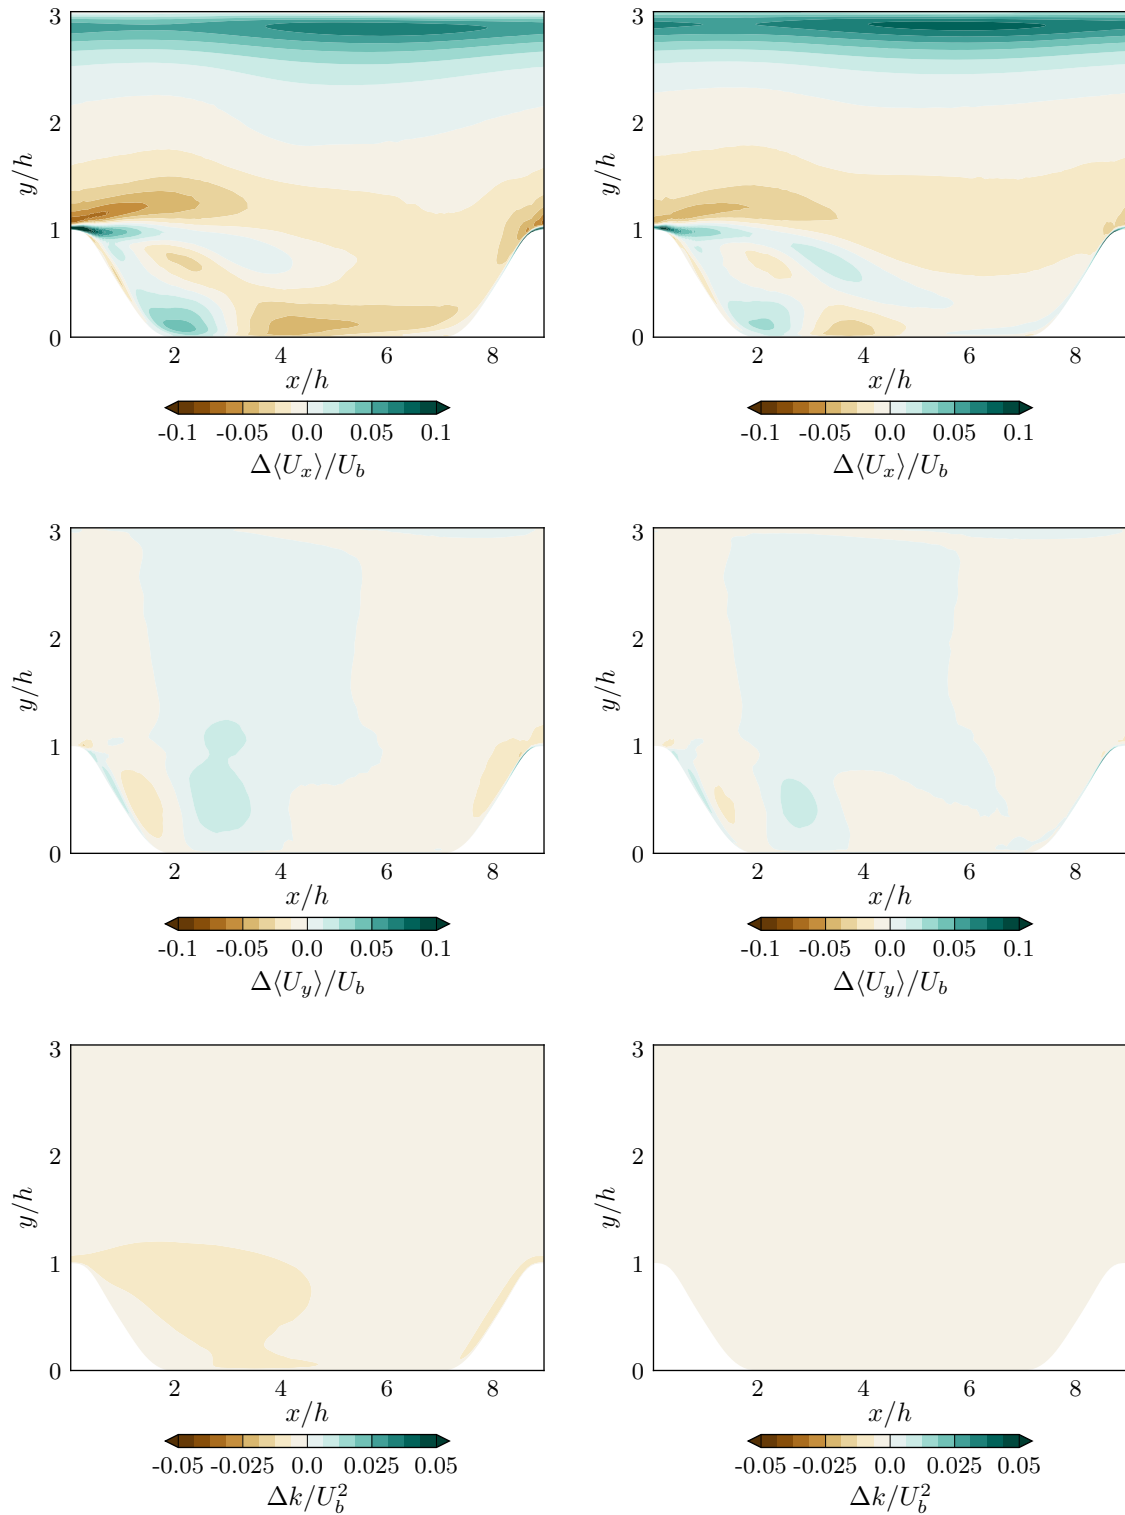


Figure B.3: Differences of HIRANS streamwise (top) and vertical (centre) mean velocity components and turbulent kinetic energy (bottom) for $\{\xi_m, \xi_t\} = \{0.80, 0.80\}$, $t^* = 5000$ iterations (left) and $\{\xi_m, \xi_t\} = \{0.90, 0.85\}$, $t^* = 0$ iterations (right) with respect to the LES solution.

Barycentric map for turbulence anisotropy

C.1 Invariants and eigenvalues of the anisotropy tensor

The barycentric map was introduced in (Banerjee et al., 2007) as an alternative way of visualising the anisotropy within the Reynolds stresses. This was motivated by the fact the previous techniques, based on the invariants of this stress tensor, could be visually misleading due to the non-linearity of its functional space.

Starting from the Reynolds stress tensor, the non-dimensional anisotropy tensor is defined as

$$b_{ij} = \frac{a_{ij}}{2k} = \frac{\langle u'_i u'_j \rangle}{\langle u'_i u'_i \rangle} - \frac{1}{3} \delta_{ij}. \quad (\text{C.1})$$

Its three principal invariants are

$$I = b_{ii}, \quad II = b_{ij} b_{ji}, \quad III = b_{ij} b_{in} b_{jn}. \quad (\text{C.2})$$

The study of the functional space of the (II , III) invariants allows to understand the dynamics of the turbulence anisotropy, as the invariant space is bounded. Based on this, two different representations have been proposed in literature. The first one, from (Lumley and Newman, 1977), is built with the following relations between II and III

$$II \geq \frac{3}{2} \left(\frac{4}{3} |III| \right)^{2/3}, \quad II \leq \frac{2}{9} + 2III. \quad (\text{C.3})$$

The second one (Choi and Lumley, 2001) introduces the variables (ξ , η) as

$$\xi^3 = III/2, \quad \eta^2 = -II/2. \quad (\text{C.4})$$

Alternatively, a representation based on the eigenvalues of the anisotropy tensor, proposed in (Lumley, 1979), can also be constructed, using

$$\lambda_1 \geq (3|\lambda_2| - \lambda_2)/2, \quad \lambda_1 \leq 1/3 - \lambda_2. \quad (\text{C.5})$$

The representations based on the eigenvalues and the invariants have the functional relationship

$$II = 2(\lambda_1^2 + \lambda_1\lambda_2 + \lambda_2^2), \quad III = -3\lambda_1\lambda_2(\lambda_1 + \lambda_2). \quad (\text{C.6})$$

One of the main advantages of this type of representations is that, are they are based on invariants of the second order tensor, they do not depend on the frame of reference.

C.2 Barycentric map




Figure C.1: Barycentric map for the anisotropy tensor.

From the previous relations, a barycentric map is defined using limiting (x, y) coordinates for the different states of turbulence, as shown in Figure C.1. This map allows to easily compare the anisotropy levels found on the RANS model with higher fidelity solutions, such as the LES cases. Due to the linear eddy viscosity assumption, the RANS eigenvalues of the anisotropy tensor are located in the line of plane strain, while this constraint is not found in LES or DNS cases, allowing to understand the relationship between the anisotropy correction and its effect on the rest of the parameters of the flowfield.

The linear functional space of the eigenvalues, in contrast with the non-linear space using invariants, allows for an easy visual representation of the different states of turbulence, which can be defined as:

Isotropic turbulence

No directional dependence appears, and thus, all eigenvalues of the anisotropy tensor are null,

$$\lambda_1 = 0, \quad \lambda_2 = 0. \quad (\text{C.7})$$

Axisymmetric turbulence

Two fluctuating components have the same statistical quantities, thus the anisotropy tensor has two multiple eigenvalues. Based on the traceless nature of this tensor, three different situations can arise:

Axisymmetric contraction

The first and second eigenvalues are the same, so

$$\begin{pmatrix} \lambda_2 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & -2\lambda_2 \end{pmatrix}. \quad (\text{C.8})$$

Axisymmetric expansion

The second and third eigenvalues are the same,

$$\begin{pmatrix} -2\lambda_2 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_2 \end{pmatrix}. \quad (\text{C.9})$$

Third case

The first and third eigenvalues are the same,

$$\begin{pmatrix} -\lambda_2/2 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & -\lambda_2/2 \end{pmatrix}. \quad (\text{C.10})$$

Two component turbulence

For this state, at least one zero eigenvalue has to appear in the Reynolds stress tensor. The anisotropy tensor in canonical form has then to contain one eigenvalue equal to $-1/3$. Because of the non-increasing-order condition, where $\lambda_1 \leq \lambda_2 \leq \lambda_3$, only two solutions are possible,

$$\lambda_1 = \frac{2}{3}, \quad \lambda_2 = -\frac{2}{3}, \quad (\text{C.11a})$$

$$\lambda_1 = \frac{1}{3} - \lambda_2. \quad (\text{C.11b})$$

Plane-strain turbulence

In this case, it is the anisotropy tensor the one which must have at least one zero eigenvalue. Three different cases can be found,

$$\begin{pmatrix} 0 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & -\lambda_2/2 \end{pmatrix}, \quad \lambda_1 = 0, \quad \lambda_2 = 0, \quad (\text{C.12a})$$

$$\begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -\lambda_1 \end{pmatrix}, \quad 0 \leq \lambda_1 \leq \frac{2}{3}, \quad \lambda_2 = 0, \quad (\text{C.12b})$$

$$\begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & -\lambda_1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad \lambda_1 = 0, \quad \lambda_2 = 0. \quad (\text{C.12c})$$

Where it is shown that the first and three cases correspond to the isotropic limiting state of turbulence, and the line corresponding to the variation of λ_1 of the second case corresponds to the plane strain in the barycentric map.

One, two and three-component limiting states

The rank of the tensor b_{ij} represents the different limiting states of the turbulence. The one, two and three-component limiting states of turbulence are found when one, two and three eigenvalues of the Reynolds stress tensor are non-zero, respectively.

Appendix D

Periodic hill parametrization for different geometries

In this appendix, the whole system for the parametrization of the periodic hill geometry is shown. Following [Almeida et al. \(1993\)](#) for the hill shape, and [Mellen et al. \(2000\)](#) for the distance between hills, the geometry for any generic configuration, if the hill height is $h = 28$, can be defined as

$$y(x) = \begin{cases} \min(28, \alpha_{01} + \alpha_{02}x + \alpha_{03}x^2 + \alpha_{04}x^3) & \text{if } x \geq 0 \text{ and } x < 9, \\ \alpha_{11} + \alpha_{12}x + \alpha_{13}x^2 + \alpha_{14}x^3 & \text{if } x \geq 9 \text{ and } x < 14, \\ \alpha_{21} + \alpha_{22}x + \alpha_{23}x^2 + \alpha_{24}x^3 & \text{if } x \geq 14 \text{ and } x < 20, \\ \alpha_{31} + \alpha_{32}x + \alpha_{33}x^2 + \alpha_{34}x^3 & \text{if } x \geq 20 \text{ and } x < 30, \\ \alpha_{41} + \alpha_{42}x + \alpha_{43}x^2 + \alpha_{44}x^3 & \text{if } x \geq 30 \text{ and } x < 40, \\ \max(0, \alpha_{51} + \alpha_{52}x + \alpha_{53}x^2 + \alpha_{54}x^3) & \text{if } x \geq 40 \text{ and } x < 54. \end{cases} \quad (\text{D.1})$$

The values for the default configuration can be found in table [D.1](#).

Point	x_i	y_i
x_0	0	28
x_1	9	27
x_2	14	24
x_3	20	19
x_4	30	11
x_5	40	4
x_6	54	0

Table D.1: Default values for the reference periodic hill geometry.

$$\text{RHS} = \begin{pmatrix} y_0 \\ y_1 \\ y_1 \\ y_2 \\ y_2 \\ y_3 \\ y_3 \\ y_4 \\ y_4 \\ y_5 \\ y_5 \\ y_6 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}. \quad (\text{D.3})$$

The first 12 rows correspond to the continuity condition. Rows 13-17 are related to the continuity of the first derivative between two segments. Rows 18-22 are second derivative continuity and finally, rows 23-24 are $f'(x_0) = 0$ and $f'(x_6) = 0$, respectively, being $f'(x_i)$ the first derivative of $y(x)$ at the point x_i . Finally, it must be remarked that the values in these calculations are normalised by y/h .

Appendix E

SIP formulation

A popular way of computing the solution of a sparse linear system $Ax = b$ is to apply LU factorization, with L being lower-triangular and U upper triangular. The system is solved as $Ly = b$, $Ux = y$. A disadvantage of this method for sparse systems is that L and U may suffer from fill-in diagonals, which can severely penalise the memory requirements to solve the system. One alternative is to perform an incomplete factorization such that $A \approx LU$, where the L and U matrices have the same sparsity pattern as the original matrix A , eliminating the fill-in bands. It can be generalised to the ILU(k) method, where the sparsity patterns of L and U are not the ones from A , but from A^{k+1} instead. This matrix $M = LU$, which does not yield by itself as the exact solution of the original system, can be introduced as a preconditioner in an iterative method such as a Krylov subspace solver, using the decomposition

$$A = M - N, \tag{E.1}$$

where N is typically small.

In [Ferziger and Peric \(2012\)](#), it is stated that this incomplete LU factorization used for certain problems may suffer from rather slow convergence. To overcome this, the *strongly implicit procedure* (SIP) was designed specifically for systems that are discretisations of partial differential equations. The derivation of the method is given for a 2D problem in orthogonal grids, with a matrix A composed of 5 non-zero diagonals, which can be extended to 3D problems. The LU factorization of the matrix A for this case results in

$$\begin{pmatrix} \cdot & & & & \\ \cdot & \cdot & & & \\ \cdot & \cdot & \cdot & & \\ \cdot & \cdot & \cdot & \cdot & \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{pmatrix} \begin{pmatrix} \cdot & & & & \\ \cdot & \cdot & & & \\ \cdot & \cdot & \cdot & & \\ \cdot & \cdot & \cdot & \cdot & \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{pmatrix} = \begin{pmatrix} \cdot & & & & \\ \cdot & \cdot & & & \\ \cdot & \cdot & \cdot & & \\ \cdot & \cdot & \cdot & \cdot & \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{pmatrix}, \tag{E.2}$$

L_w L_s L_p 1 U_N U_E M_W M_{NW} M_S M_P M_N M_{SE} M_E

where the subindices N , S , E , W , NW , SE indicate north, south, east, west, north-west and south-east nodes, respectively. It can be observed that at least two extra non-zero diagonals are obtained after performing the product LU . If these two diagonals, as in ILU (and also in SIP), are set to zero for the matrix M , then N must have at least non-zero elements in those two diagonals, where the rest of the diagonals of M are equated to the corresponding diagonals of A . This choice is the one of ILU(0).

Instead, N can be allowed to have non-zero elements if *all* the seven non-zero diagonals of LU . Using a vector ϕ at a row P ,

$$(M\phi)_P = M_P\phi_P + M_S\phi_S + M_N\phi_N + M_E\phi_E + M_W\phi_W + M_{NW}\phi_{NW} + M_{SE}\phi_{SE}, \quad (\text{E.3})$$

the two last terms are the ones corresponding to the extra filling.

If the matrix N needs to contain the two extra diagonals of M , and the rest of the elements are desired to be chosen such as $N\phi \approx 0$,

$$N_P\phi_P + N_N\phi_N + N_S\phi_S + N_E\phi_E + N_W\phi_W + N_{NW}\phi_{NW} + N_{SE}\phi_{SE} \approx 0, \quad (\text{E.4})$$

then the two extra terms need to nearly cancel the contribution of the remaining 5. That is, E.4 can be written as

$$M_{NW}(\phi_{NW} - \phi_{NW}^*) + M_{SE}(\phi_{SE} - \phi_{SE}^*) \approx 0. \quad (\text{E.5})$$

In this equation, ϕ_{NW}^* and ϕ_{SE}^* are approximations of ϕ_{NW} and ϕ_{SE} .

The idea proposed by Stone is based on the idea that the solution is expected to be smooth, being based on the approximation of an elliptic partial equation. With this, ϕ_{NW}^* and ϕ_{SE}^* can be set using the values of ϕ at nodes which correspond to the diagonals of A . The proposed approximation was

$$\phi_{NW}^* \approx \alpha(\phi_W + \phi_N - \phi_P), \quad (\text{E.6a})$$

$$\phi_{SE}^* \approx \alpha(\phi_S + \phi_E - \phi_P). \quad (\text{E.6b})$$

If the factor $\alpha = 1$, the interpolation is second-order accurate, but not stable, as Stone's work stated that stability requires $\alpha < 1$. In the present document, this method was used for preconditioning the system using $\alpha = 0.92$, solved with BiCGStab.

Appendix F

Additional LES cases

In this appendix, line plots at different x/h stations can be found for different INCA, RANS and HIRANS cases.

F.1 Extended D, RANS and HIRANS cases

This first section highlights the differences between the two LES cases used for injecting data in the original periodic hill geometry, $\psi = 1.0$, and the RANS and HIRANS cases. Figure F.1 shows that, for both INCA and Breuer et al. (2009) LES data, the HIRANS case achieves a much better agreement with the injected case than the regular $k-\omega$ model. The corrections on the bottom wall seem satisfactory, while the top wall values are more problematic in the streamwise component. With respect to the turbulent stresses, Figure F.2, they are greatly improved, providing a similar shape with respect to the LES and PIV solutions, while being slightly underpredicted in general in magnitude.

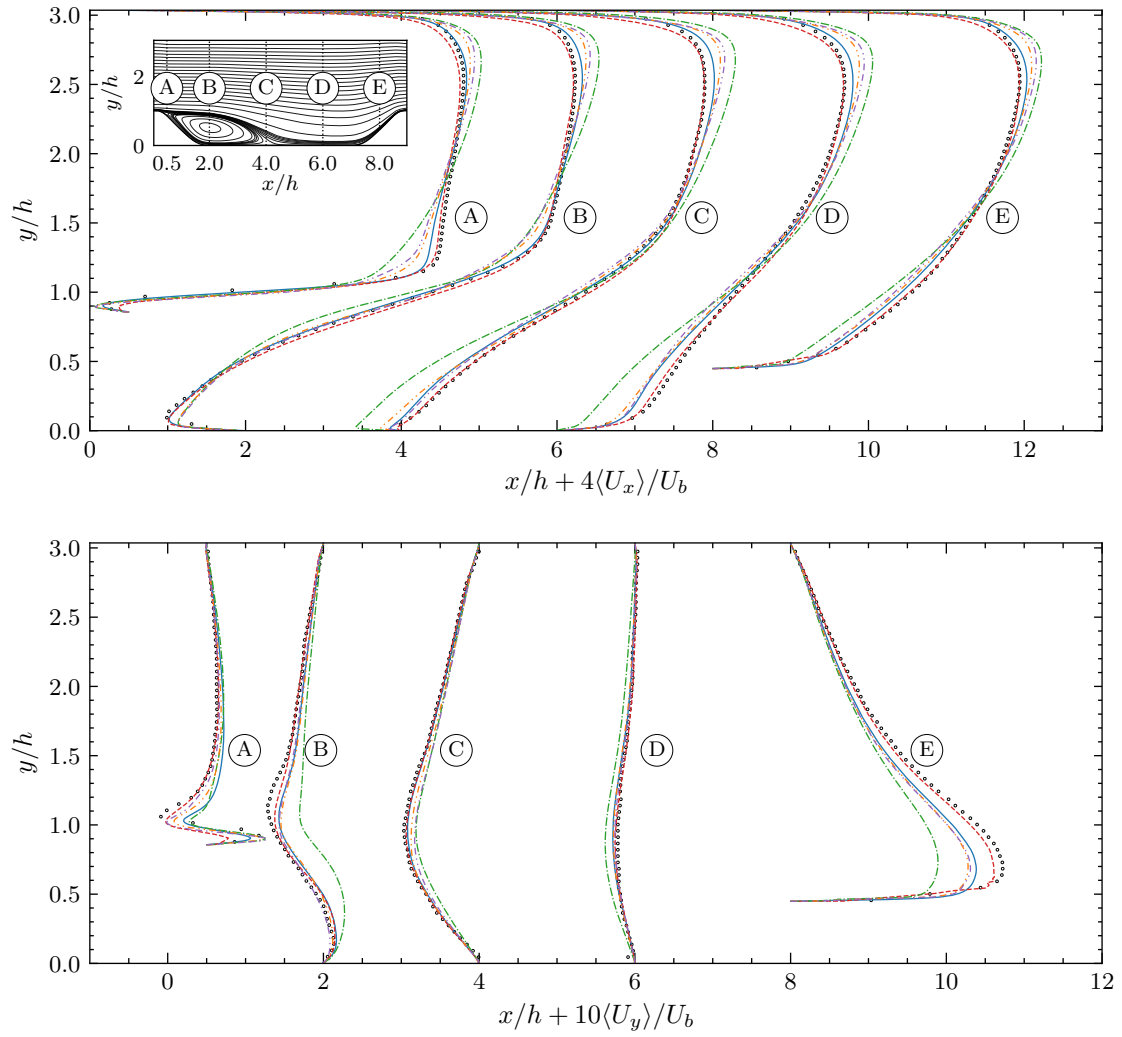


Figure F.1: Mean velocity components at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), Extended D (- - -), RANS (·····), HIRANS with τ_{LES} injection from Breuer et al. (2009) (- · - ·) and INCA (- · - ·) cases.

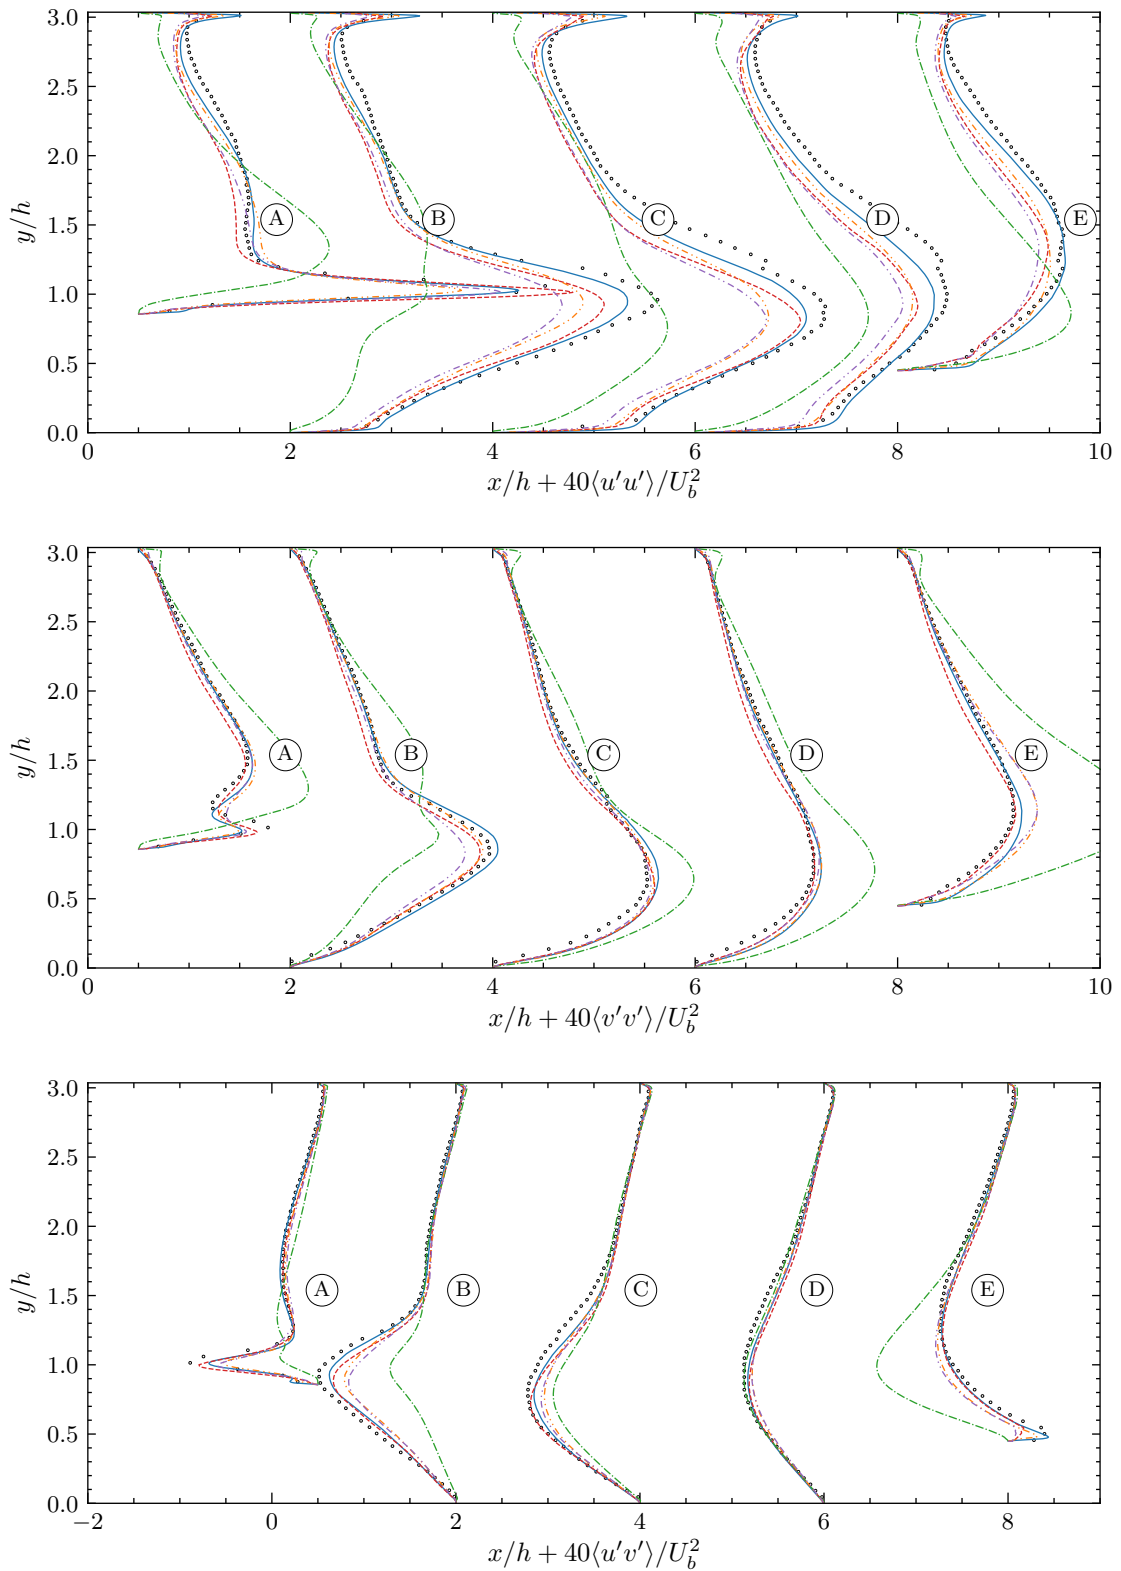


Figure F.2: Turbulent stresses at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), Extended D (- - -), RANS (····), HIRANS with τ_{LES} injection from Breuer et al. (2009) (- · - ·) and INCA (- · - · -) cases.

F.2 Baseline, Extended A and B cases

This section includes the Extended A and B cases in the comparison, analysing the effect of a delayed start of the averaging period of 55 flow-through times, averaging over 33 flow-through times (Extended A) and a delayed and extended averaging period of 55 flow-through times.

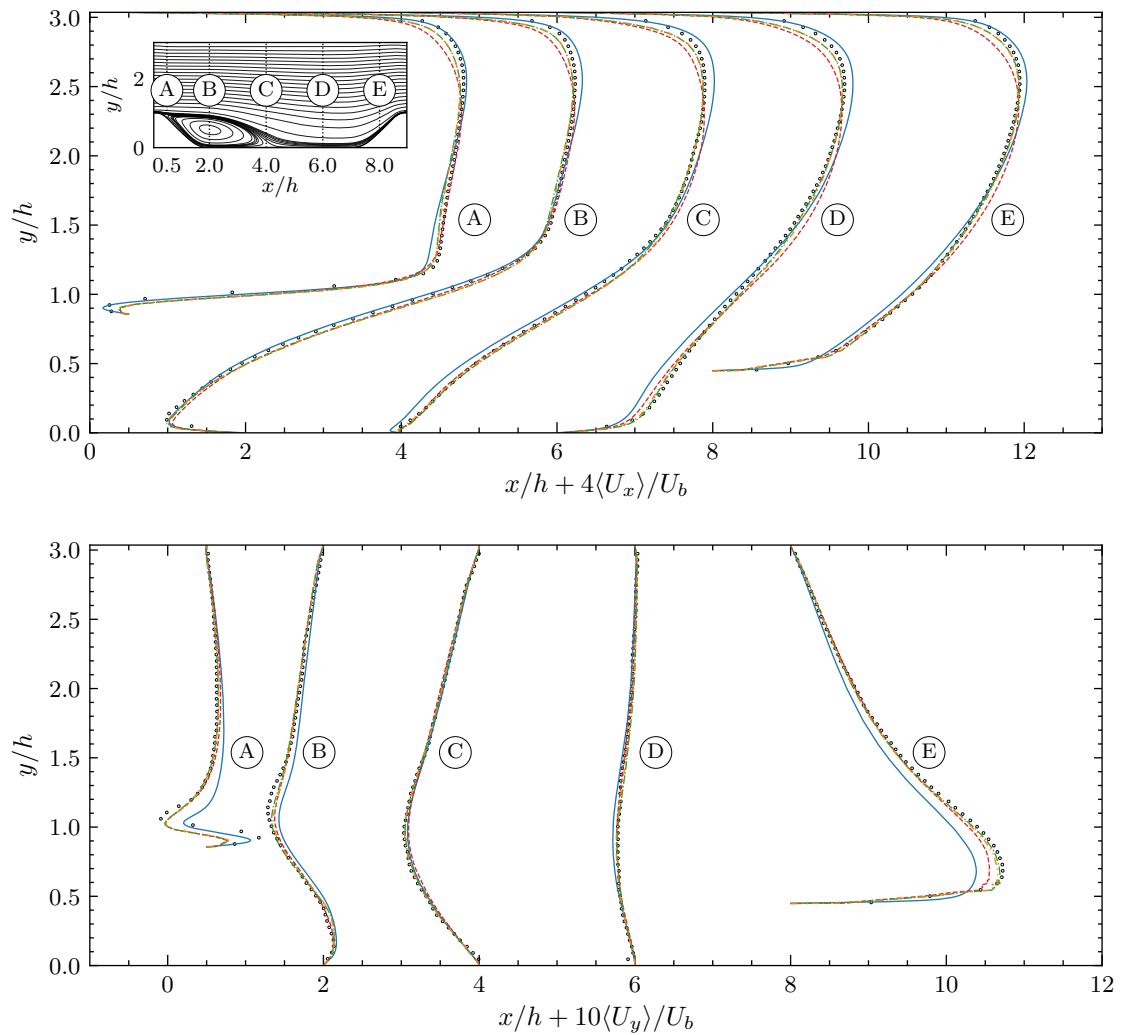


Figure F.3: Mean velocity components at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), INCA baseline (- - -), Extended A (- · - ·) and Extended B (- · · -) cases.

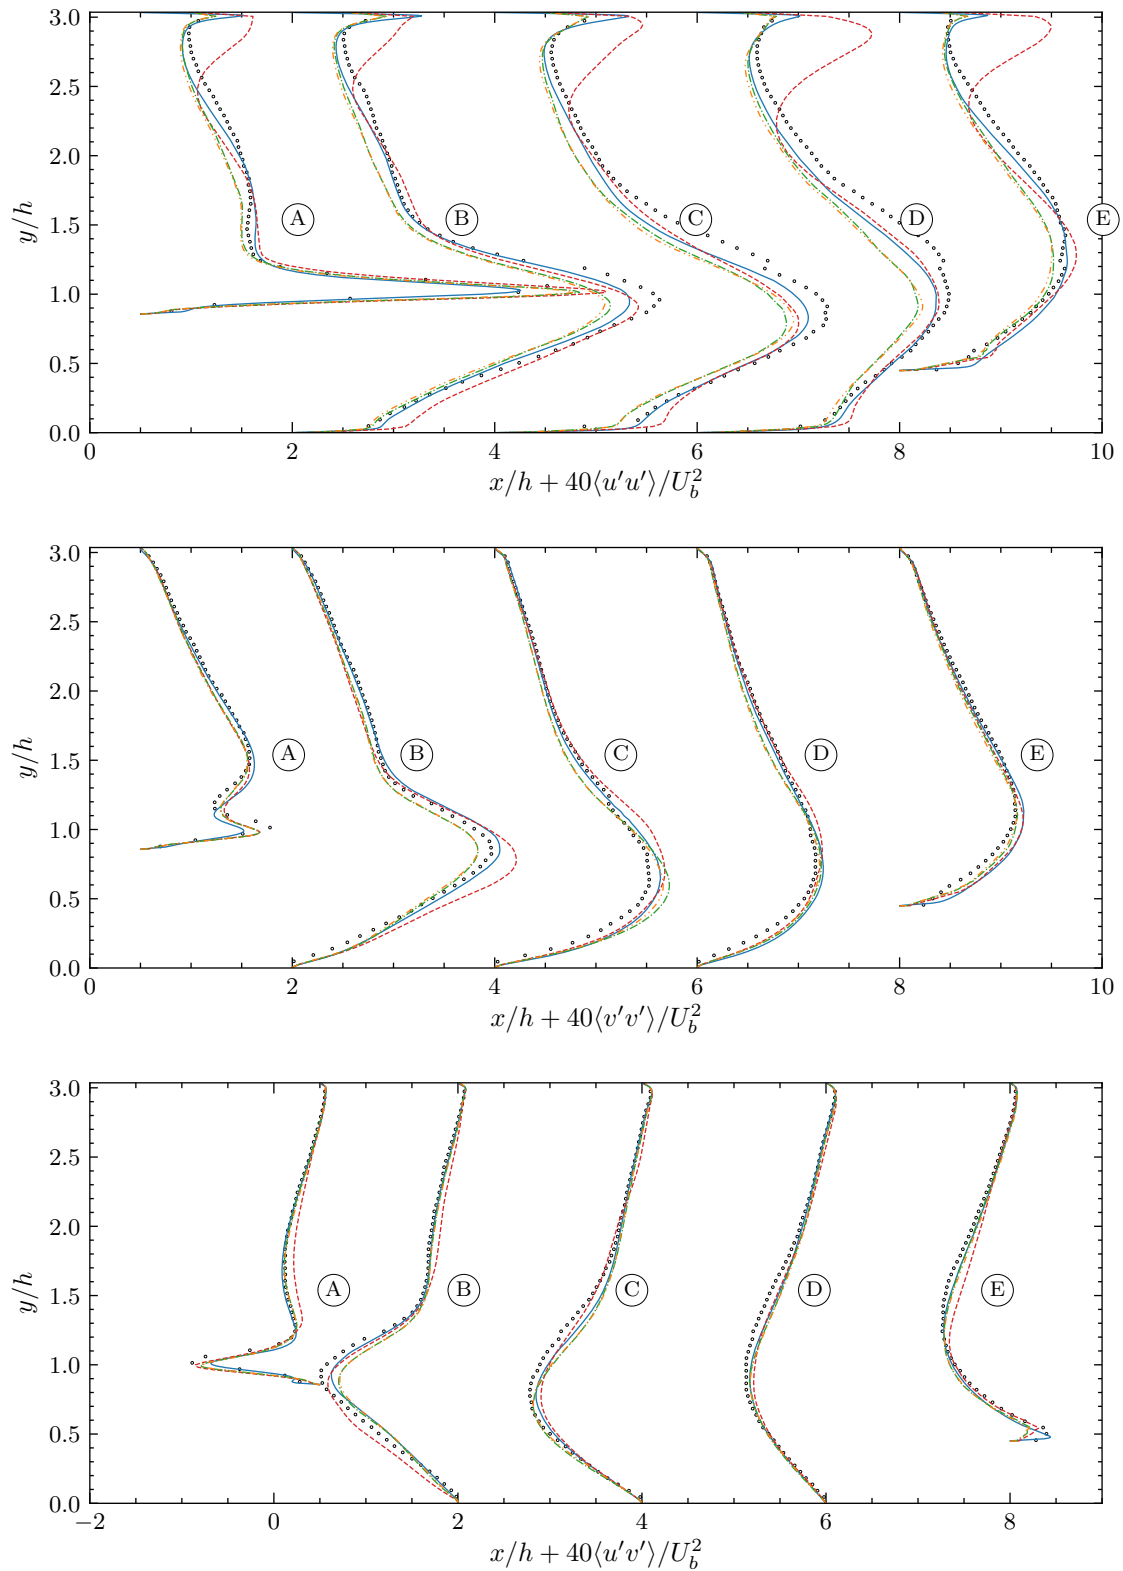


Figure F.4: Turbulent stresses at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), INCA baseline (- - -), Extended A (···) and Extended B (- · -) cases.

F.3 Extended B, C and D cases

The final section includes the effect of furthering delayed start of the averaging period, 89 flow-through times, averaging over 33 flow-through times (Extended C) and a delayed start and extended averaging period of 89 and 55 flow-through times, respectively.

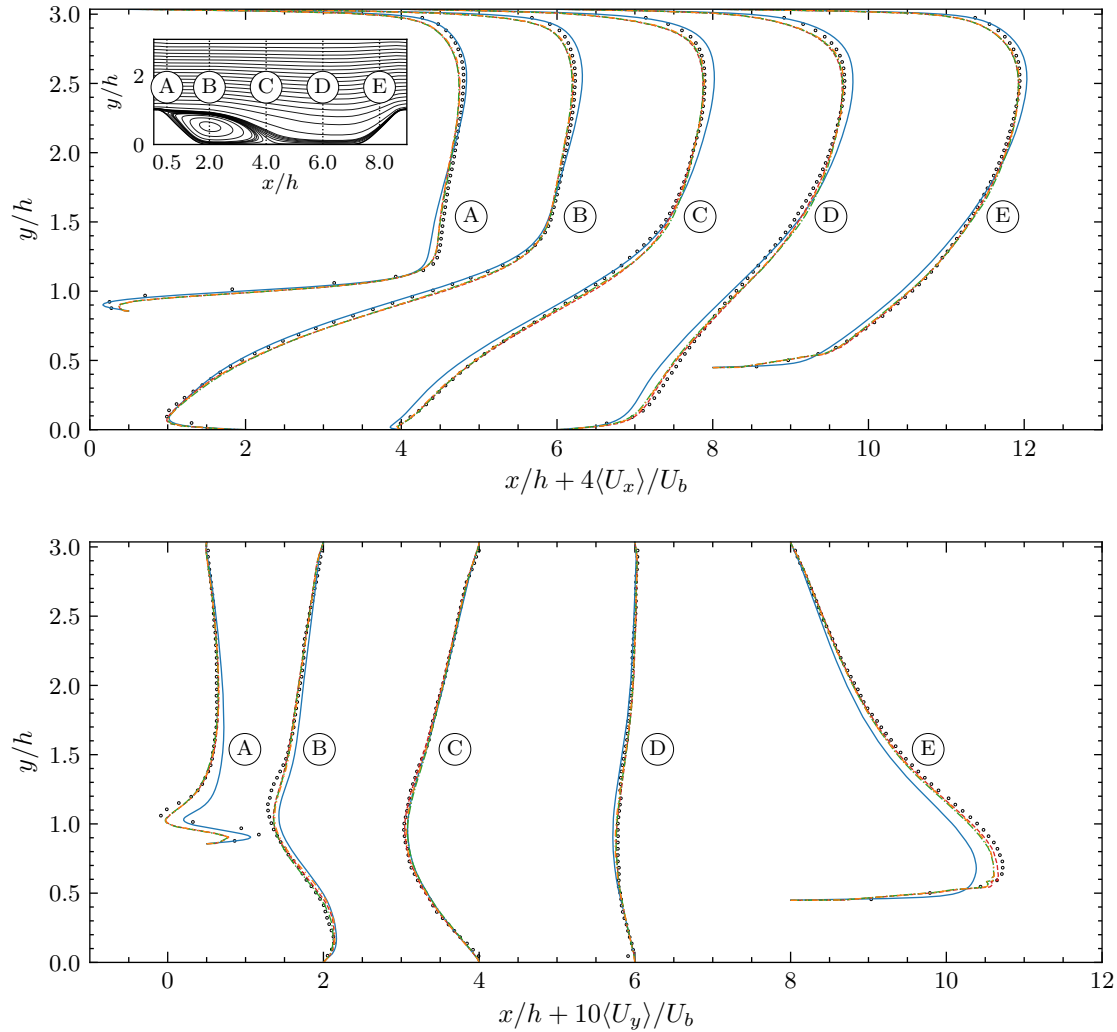


Figure F.5: Mean velocity components at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), INCA Extended B (- - -), Extended C (- · - ·) and Extended D (- · · - ·) cases.

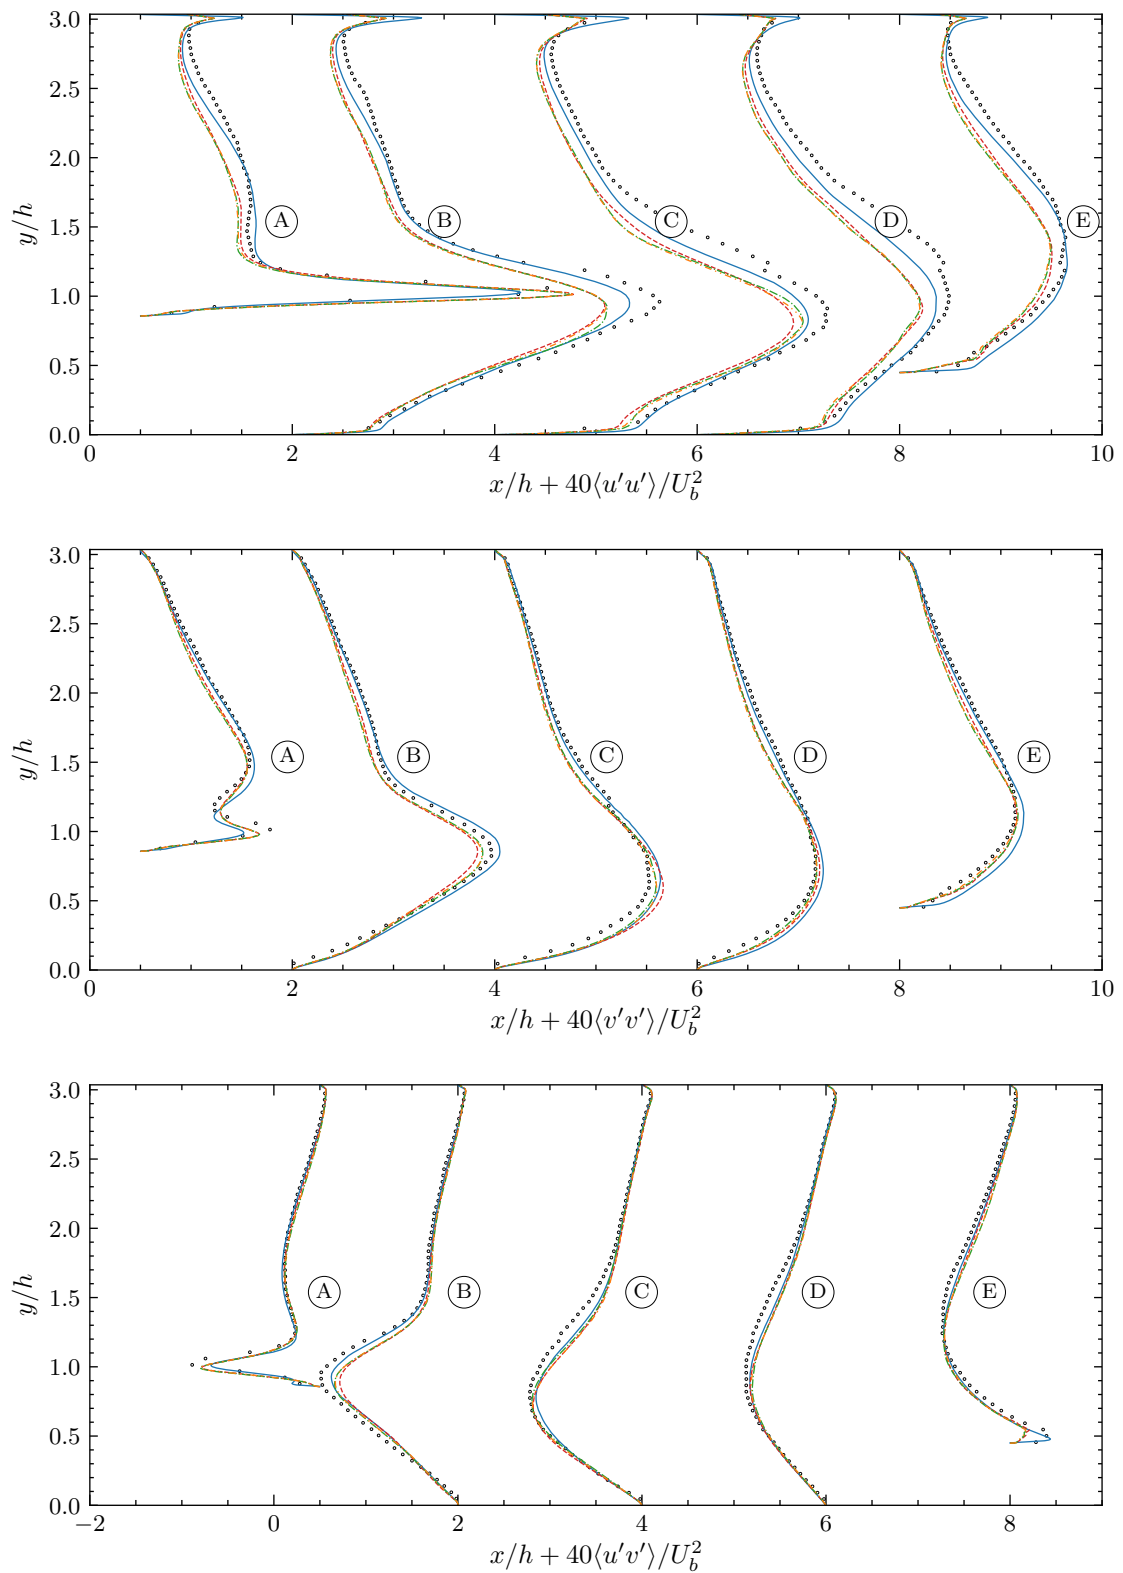


Figure F.6: Turbulent stresses at different x/h stations for Breuer et al. (2009) LES (—), Rapp (2009) PIV (\circ), INCA Extended B (- - -), Extended C (- · - ·) and Extended D (···) cases.

Kriging/Co-Kriging setup parameters

G.1 Bounds for hyperparameters

The bounds for the hyperparameters have been tested, to aim to keep them as wide as possible while not inordinately increasing the cost of the Maximum Likelihood Estimation procedure, which was run for 1000 iterations in the majority of the simulations, and reduced to 100 for the most expensive ones and the additional tests performed in Chapter 6. The bounds were specified as

$$\text{Kriging: } \begin{cases} \theta & = [-3.0, 1.0], \\ p & = [1.5, 2.0]. \end{cases} \quad (\text{G.1a})$$

$$\text{Co-Kriging: } \begin{cases} \theta_c & = [-3.0, 2.0], \\ \theta_d & = [-1.0, 1.0], \\ p_c & = [1.0, 2.0], \\ p_d & = [1.8, 2.0], \\ \rho & = [-5.0, 5.0]. \end{cases} \quad (\text{G.1b})$$

G.2 Exploration-exploitation bias ξ_{EI} per iteration

Similarly as the hyperparameter estimation, the Expected Improvement criterion was run for 5000 samples within the parameter space for the main routines in the first iterations, and reduced to 500 samples in later stages due to its increasing computational cost.

Acquisition strategies for $n_{e_{\text{add}}} = 1$

Method	Routine	ξ_{EI}	Explore/Exploit
Kriging LES	A	0.01	Explore
Kriging LES	B	0.01	Explore
Kriging RANS	A	-0.1	Exploit
Kriging RANS	B	-0.5	Exploit
Kriging HIRANS	A	-0.1	Exploit
Kriging HIRANS	B	-0.5	Exploit
Co-Kriging RANS	A	-0.1	Exploit
Co-Kriging RANS	B	-0.5	Exploit
Co-Kriging HIRANS	A	-0.1	Exploit
Co-Kriging HIRANS	B	-0.5	Exploit

Table G.1: Acquisition strategies and ξ_{EI} values of for the first additionally acquired samples.

Acquisition strategies for $n_{e_{\text{add}}} = 2$

Method	Routine	ξ_{EI}	Explore/Exploit
Kriging LES	A	0.01	Explore
Kriging LES	B	0.01	Explore
Kriging HIRANS	A	-0.2	Exploit
Kriging HIRANS	B	-0.5	Exploit
Co-Kriging RANS	A	-0.2	Exploit
Co-Kriging RANS	B	-0.5	Exploit
Co-Kriging HIRANS	A	-0.2	Exploit
Co-Kriging HIRANS	B	-0.5	Exploit

Table G.2: Acquisition strategies and ξ_{EI} values of for the second additionally acquired samples.

Acquisition strategies for $n_{e_{add}} = 3$

Method	Routine	ξ_{EI}	Explore/Exploit
Kriging LES	A	0.01	Explore
Kriging LES	B	0.01	Explore
Kriging HIRANS	A	-0.2	Exploit
Kriging HIRANS	B	-0.5	Exploit
Co-Kriging RANS	A	-0.1	Explore
Co-Kriging RANS	B	-0.25	Explore
Co-Kriging HIRANS	A	-0.1	Explore
Co-Kriging HIRANS	B	-0.25	Explore

Table G.3: Acquisition strategies and ξ_{EI} values of for the third additionally acquired samples.

Acquisition strategies for $n_{e_{add}} = 4$

Method	Routine	ξ_{EI}	Explore/Exploit
Kriging LES	A	-0.1	Exploit
Kriging LES	B	-0.5	Exploit
Co-Kriging RANS	A	-0.2	Exploit
Co-Kriging RANS	B	-0.5	Exploit
Co-Kriging HIRANS	A	-0.2	Exploit
Co-Kriging HIRANS	B	-0.5	Exploit

Table G.4: Acquisition strategies and ξ_{EI} values of for the fourth additionally acquired samples.

Acquisition strategies for $n_{e_{add}} = 5$

Method	Routine	ξ_{EI}	Explore/Exploit
Kriging LES	A	0.0	Explore
Kriging LES	B	0.0	Explore
Co-Kriging RANS	A	0.0	Explore
Co-Kriging RANS	B	0.0	Explore
Co-Kriging HIRANS	A	0.0	Explore
Co-Kriging HIRANS	B	0.0	Explore

Table G.5: Acquisition strategies and ξ_{EI} values of for the fifth additionally acquired samples.Acquisition strategies for $n_{e_{add}} = 6$

Method	Routine	ξ_{EI}	Explore/Exploit
Kriging LES	A	-0.1	Exploit
Kriging LES	B	-0.5	Exploit
Co-Kriging RANS	A	0.0	Explore
Co-Kriging RANS	B	0.0	Explore
Co-Kriging HIRANS	A	0.0	Explore
Co-Kriging HIRANS	B	0.0	Explore

Table G.6: Acquisition strategies and ξ_{EI} values of for the sixth additionally acquired samples.

Appendix H

Snapping integration in Kriging/Co-Kriging predictor

Snapping [Talnikar et al. \(2015a\)](#) allows to give the Expected Criterion a third option apart from exploration and exploitation, which is to recompute an existing sample for a longer averaging period to reduce its uncertainties. This technique could be integrated in the Kriging predictor in different ways. A proposal for future research is given by the modification of the regularization parameter λ . Recalling the Gram matrix equation, 2.4.1,

$$\mathbf{w} = (\Psi + \lambda \mathbf{I})^{-1} \mathbf{y}, \quad (\text{H.1})$$

with the weightings \mathbf{w} , Ψ contains the correlation matrix of the different existing samples, λ is a scalar and \mathbf{y} yields the existing responses.

A simple correction of this model could be given by actuating over the regression close to the existing points, by making the parameter λ a vector, $\lambda = \lambda(\mathbf{x})$, dependent on the individual uncertainty of each sample point. This uncertainty could be calculated by computing each individual estimation of the different variables used in the objective function J by the subdomain branching method, and then applying linear weightings based on the resultant statistical distribution over the parameter space. A simple estimation could be given, parting the particular choice of J of this document using error propagation of the individual components,

$$\sigma_{J_{aux}}^2 = \alpha^2 \sigma_k^2 + \beta^2 \sigma_{f_p}^2, \quad (\text{H.2a})$$

$$\sigma_J^2 = J^2 \left[\left(\frac{\sigma_{J_{aux}}}{J} \right)^2 + \left(\frac{\sigma_{J_{BSL}}}{J_{BSL}} \right)^2 \right]. \quad (\text{H.2b})$$

This estimated variance would then be introduced in the calculation of the regularization parameter as

$$\lambda(\mathbf{x}) = \lambda_{\max} \frac{\sigma_J^2(\mathbf{x})}{\sigma_{J_{\max}}^2}, \quad (\text{H.3})$$

where λ_{\max} would be set as the maximum regression allowed in the routine. This idea can be linked with the formulation present in [Sasena \(2002\)](#), where λ is not a scalar, but instead a vector with an entry per dimension of the parameter space, allowing different noise levels per design parameter. In the proposed formulation of this document, this idea would be extended to noise levels *per individual sample*, with a straightforward application for multi-dimensional optimization.

Another alternative, which is outlined in [Forrester et al. \(2008\)](#), is to include λ as a hyper-parameter to be determined by MLE. This has the advantage of providing a data-driven approach over the smoothness of the surrogate, with the drawbacks of being oblivious of the true uncertainty of a specific point and its neighbours and the increased cost of the MLE process.

