University of London
Imperial College of Science, Technology and Medicine
Department of Aeronautics

Scale-Resolving Simulations and Optimisation of Wind Farms in Complex Terrain

Christian Jané-Ippel

September 2025

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Aeronautics of the University of London and
the Diploma of Imperial College.

To my beloved family and friends.

A mi querida familia y amigos.

A la meva estimada família i amics.

Declaration of Originality

All the work presented in this thesis is that of the author. Any work or contributions made by others is referenced appropriately.

Copyright Declaration

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Abstract

Wind energy is a cornerstone of the global transition to net zero, yet wind farm efficiency is often constrained by wake interactions and complex terrain. Optimising turbine layouts to maximise power capture is therefore critical. Conventional methods, such as analytical wake models or steady RANS, are efficient in flat terrain but fail to resolve the unsteady turbulence and terrain–wake coupling dominating hilly landscapes. Large-eddy simulation (LES) provides the necessary fidelity, but computational costs have traditionally limited it to validation. This thesis addresses this gap by developing a framework coupling LES with Bayesian Optimisation (BO) for wind-farm layout optimisation in complex terrain.

Methodologically, the high-order solver `WInc3D` is extended using immersed boundary and wall-stress modelling to simulate realistic onshore farms. In parallel, BO is embedded within an HPC workflow, introducing physics-informed seeding strategies to efficiently initialise optimisation in high-dimensional design spaces.

Validation against wind-tunnel experiments and SCADA data confirms the solver reproduces boundary-layer inflows, terrain-induced flow, and turbine performance. A two-turbine optimisation around a Gaussian hill provides proof of concept, demonstrating that terrain features can be exploited through hub-height or wake steering. Finally, the framework is scaled to a 16-turbine farm in synthetic three-dimensional terrain. Here, physics-informed BO proves essential: while black-box BO struggles, seeding with flow-based layouts enables tractable optimisation, revealing consistent exploitation of ridges and avoidance of recirculation zones.

This thesis demonstrates that LES–BO is feasible for wind-farm design in complex terrain, with typical simulations taking one to four hours. It advances methodology by integrating high-fidelity simulation with optimisation, extending the understanding of mechanisms like ridge acceleration, recirculation, wake deflection, and valley sheltering. Ultimately, it establishes LES as a design-oriented framework for next-generation wind-farm optimisation.

Acknowledgements

I would like to express my deepest gratitude to those who have guided, encouraged, and accompanied me throughout this doctoral journey.

First and foremost, I am sincerely thankful to my supervisors, Prof. Sylvain Laizet and Prof. Rafael Palacios, for their mentorship and support. Working under their guidance has been both a privilege and an honour. Sylvain has been instrumental in shaping my scientific development, teaching me the importance of persistence, clarity of thought, and the need to balance exploration with concrete results. His encouragement and confidence in me have been of great value. Rafa has greatly enriched my research with his broader perspective on science and his thoughtful insights, always shared with generosity. I owe them both my deepest thanks.

I would also like to thank Imperial College London, for providing an intellectually stimulating environment that has inspired me both before and during my doctoral studies. The atmosphere of excellence, curiosity, and collaboration at Imperial has played a central role in my education and personal growth. I am particularly grateful to the Department of Aeronautics for its world-class research environment and the opportunities it has offered me.

My appreciation goes to Dr. Nikos Bempedelis, whose guidance, discussions, and friendship have been a constant support throughout this journey. His generosity with his time and knowledge, both in research and beyond, has been invaluable. I am also thankful to my collaborator, Dr. Kostas Steiros, for the chance to collaborate and share ideas. To Prof. Oliver Buxton for welcoming me into his football team and the time spent together. I also wish to acknowledge Dr. Chris Cantwell and Prof. Spencer Sherwin, whose support after my MSc allowed me to remain connected to Imperial during a pivotal stage, ultimately making this PhD possible.

I thank NERC and SSE Renewables Services (UK) Limited for funding this PhD studentship. I am grateful to SSE Renewables for their guidance and insights as an industrial partner, particularly Paul Housley, Eoin Fanning, and Jessica Lunn for their advice and support. I am also thankful to the Grantham Institute and the Science and Solutions for a Changing Planet DTP at Imperial College London for fostering an inspiring cohort of students united by the shared goal of addressing climate change. Beyond the technical training, these programmes offered opportunities to broaden my perspective on sustainability, including memorable gatherings such as our meetings at the Natural History Museum.

I gratefully acknowledge the use of the ARCHER2 UK National Supercomputing Service via a Pioneer project and the UK Turbulence Consortium. I would also like to thank the organisers of the International HPC Summer School (2022) and the von Kármán Institute Workshop on CFD for Atmospheric Flows and Wind Engineering (2023), both of which enriched my understanding and provided opportunities to connect with the wider research community.

I am especially thankful to my colleagues and friends at Imperial, whose companionship has made this journey both meaningful and enjoyable. In particular, I would like to thank Chi Hin, Malavika, Francisco, and Ignasi for their enduring friendship and support through the highs and lows of the PhD. I am also deeply grateful to the many colleagues who shared this experience with me — Kazuki, Steffi, Mario, Grace, Josephine, Henrik, Friedrich, Maria, Benita, Ali, Kelvin, Alex, Kaloyan, Parv, and many others — for the countless discussions, memories, and moments of joy we created together.

I also wish to express my gratitude to my closest friends outside Imperial, particularly those from my undergraduate studies — Gorka, Pep, Roger, Òscar, Oriol, Xavi, Alexandra, Marta, Guillem, and Ruth — with whom I shared the challenges of engineering studies. The bonds we forged during those years remain unbreakable and have shaped who I am today.

I owe everything to my family: to my parents, Luis and Ana, for their unconditional love, encouragement, and faith in me, and to my brother Luis, who has always understood me better than anyone else and whose presence I value most. I would also like to thank my sisters and cousins — Patri, Silvia, Laura, Óscar, Hans, and Marc — who are like siblings to me, and the wider family — Joan, Maribel, Rosa, Toni, HH, Betty, and Sergio — whose support and joy have always been present in my life. Finally, I wish to dedicate this thesis to my grandparents, Luis and Maria, and Gloria and Derek, whose lives and values continue to inspire me.

Lastly, I extend my thanks to all those I have met through conferences, seminars, summer schools, and travels. As David Eagleman has remarked, "All the experiences in your life shape the microscopic details of your brain. Neurally speaking, who you are depends on where you've been."[1] This thesis, and the person I have become, are the result of all those experiences and the people who have been part of them. To each of you, I am deeply thankful.

[1] Eagleman, D. (2015). *The Brain: The Story of You*. Canongate Books, Edinburgh.

Contents

Declaration of Originality **5**

Copyright declaration **7**

Abstract **9**

Acknowledgements **11**

List of Tables **19**

List of Figures **21**

1 Introduction **29**

1.1 Context and Research Motivation 29

1.2 Complex Terrain Flow Physics 31

1.3 Modelling and Optimisation Across Fidelities 34

1.3.1 Low-Order Methods 34

1.3.2 Mid-Fidelity RANS Approaches 36

1.3.3 High-Fidelity Large-Eddy Simulation 37

1.4 Research Gaps and Objectives 39

1.5 Research Innovation and Contributions 41

1.6 Thesis Structure 42

1.7 List of Publications 43

2 Computational Methods **45**

2.1 Large-Eddy Simulation Framework 46

2.1.1 Introduction to `WInc3D` 47

2.1.2 Governing Equations 49

2.1.3 Numerical Solver 50

2.1.4 Atmospheric Boundary Layer Modelling 62

2.1.5 Complex Terrain Simulations 65

2.1.6 Turbine Modelling 67

2.2 Bayesian Optimisation Framework 71

2.2.1 Design of Experiments 72

2.2.2 Gaussian Process Surrogate and Hyperparameter Strategy 75

2.2.3 Acquisition Functions and Batch Strategy 78

3 Validation of the LES Framework **81**

3.1 Idealised Benchmark 82

3.1.1 Precursor and 2D Hill Validation 83

3.1.2 Extension to 3D Hill 87

3.1.3 Turbine Performance around the 2D Hill . . . 89
3.1.4 Yawed Turbine Validation . . . 93
3.2 Operational Wind Farm Validation . . . 94
3.2.1 Computational Domain and Terrain Representation . . . 95
3.2.2 Grid Resolution and Computational Cost . . . 96
3.2.3 Inflow Profile . . . 98
3.2.4 Roughness Map . . . 99
3.2.5 ADM Configuration . . . 100
3.2.6 Power Comparison with SCADA . . . 102
3.2.7 Flow-Field Interpretation . . . 104
3.3 Conclusions . . . 106

4 Two-Turbine Optimisation over the 2-D Hill 108

4.1 Problem Setup . . . 109
4.1.1 Design Variables and Objective Function . . . 109
4.1.2 Optimisation Configuration . . . 112
4.2 Results . . . 113
4.2.1 Power Production . . . 115
4.2.2 Velocity and Turbulence Intensity Fields . . . 116
4.2.3 Influence of the upstream wind turbine . . . 119
4.2.4 Hill Recirculation Changes . . . 121
4.3 Conclusions . . . 122

5 Utility-Scale Wind-Farm Optimisation over Synthetic Terrain 125

5.1 Problem Setup . . . 126

5.1.1 Site and Terrain Model . . . 126

5.1.2 Atmospheric Boundary Layer Inflow . . . 127

5.1.3 Design Space, Constraints and Optimisation Cases . . . 128

5.1.4 Simulation Configuration . . . 131

5.1.5 Optimisation Configuration . . . 132

5.2 Results . . . 133

5.2.1 Physics-Informed BO: Motivation and Performance . . . 133

5.2.2 Optimal Layouts over the Synthetic Multi-Hill Topography . . . 139

5.2.3 Flow Physics over the Synthetic Multi-Hill Topography . . . 146

5.3 Conclusions . . . 150

6 Conclusions and Future Work 152

6.1 Conclusions . . . 152

6.1.1 Methodological Contributions . . . 153

6.1.2 Synthesis of Results Across Chapters . . . 153

6.1.3 Scientific and Practical Impact . . . 155

6.1.4 Limitations . . . 156

6.1.5 Closing Statement . . . 157

6.2 Future Work . . . 157

Bibliography **160**

A Copyright Permissions **175**

List of Tables

2.1 Sixth-order discrete operators. The column α/α_f gives the implicitness coefficient α for derivative and interpolation stencils, and the damping coefficient α_f for the filter. 54

4.1 Summary of optimum design variables and objective functions obtained from BO1 and BO2, alongside the Reference case. 113

5.1 Characteristics of the individual hills forming the composite terrain. 127

5.2 Normalised power outputs for the optimal layouts across WFLO cases. For WFLO-2D-MWD, individual inflow angles are also reported. 140

A.1 Permissions summary table for reproduced figures 175

List of Figures

1.1 Trend of onshore and offshore wind turbine size (rotor diameter and tip height), 1980–2030. Figure reproduced from Global Wind Energy Council (2024), with permission from GWEC. 31

1.2 A wind farm in mountainous terrain with vegetation. Figure licensed from Getty Images. 32

1.3 Wind turbines located near a steep valley. Figure reproduced with permission from Nordex SE. 32

2.1 Half-staggered pressure field, showing the location of velocity and pressure nodes. Figure reproduced from Laizet and Lamballais (2009), with permission from Elsevier. 53

2.2 Schematic of the polynomial reconstruction in the alternating-direction IBM. Along each sweep direction, a one-dimensional Lagrangian polynomial is fitted to fluid reference points (filled blue circles) and constrained to satisfy $u_i = 0$ at the interface (black circles), providing smooth extended values (red squares) for the compact derivative stencil inside the solid region. Figure reproduced from Gautier et al. (2014), with permission from Informa UK Limited (Taylor & Francis Group). 59

2.3 Two-dimensional pencil decomposition for $P_{\text{row}} = 4$ and $P_{\text{col}} = 4$, illustrating the three pencil orientations: (a) X-pencil, (b) Y-pencil, and (c) Z-pencil. Each colour represents a distinct MPI process. Figure reproduced from Laizet and Li (2011), with permission from John Wiley & Sons. 61

2.4 An example of BO on a toy one-dimensional problem. Figure reproduced from Tanaka and Iwata (2018), with permission from Springer Nature. 73

3.1 Computational domain representation of the validation set-up, featuring the constant section hill and an actuator disc at the hilltop, along with the boundary conditions. 84

3.2 Horizontally- and time-averaged streamwise velocity (a) and resolved stress profiles (b) of the TBL in the precursor simulation compared against analytical profiles and experimental data of Cao and Tamura (2006). 85

3.3 Comparison between `WInc3D` simulations of the 2D hill and the experimental data of Cao and Tamura (2006) for (a) normalised streamwise velocity and (b) turbulence intensity. The white line in (a) shows the edge of reverse flow. Vertical profiles are plotted with a horizontal offset corresponding to their streamwise location for visualisation. 86

3.4 Time-averaged streamwise velocity profiles from simulations over the 3D hill compared with the experimental data of Ishihara et al. (1999). Vertical profiles are shown at the mid-plane for various streamwise locations (left), plotted with a horizontal offset corresponding to their streamwise position, and a spanwise profile at $x/H = 3.75$, $y/H = 1$ (right). `WInc3D` simulations (H1: red, H2: black) are compared with Diebold et al. (2013) (green circles), Fang and Porté-Agel (2016) (blue circles), and Liu and Stevens (2020b) (orange circles). 88

3.5 Streamwise turbulence intensity profiles from simulations over the 3D hill compared with the experimental data of Ishihara et al. (1999). Vertical profiles are shown at the mid-plane for various streamwise locations (left), plotted with a horizontal offset corresponding to their streamwise position, and a spanwise profile at $x/H = 3.75$, $y/H = 1$ (right). `WInc3D` simulations (H1: red, H2: black) are compared with Diebold et al. (2013) (green circles), Fang and Porté-Agel (2016) (blue circles), and Liu and Stevens (2020b) (orange circles). 89

3.6 Normalised power output of a wind turbine placed at various streamwise positions relative to the 2D hill centred at $x/D = 0$, as computed by `WInc3D`, Liu and Stevens (2020a), and Zhang et al. (2022). 90

3.7 Velocity deficit fields and wake-centre trajectories (white dashed lines) for wind turbines placed at different streamwise positions relative to the 2D hill: upstream ($x/D = -5$), crest ($x/D = 0$), and near-lee ($x/D = 1.25$). Vertical profiles are plotted with a horizontal offset corresponding to their streamwise location for visualisation. 92

3.8 Normalised velocity deficit along the wake-centre trajectory for wind turbines placed at different streamwise positions relative to the 2D hill and in flat terrain. 93

3.9 Power output of a yawed wind turbine normalised by the power output of a turbine at zero yaw angle. Comparison between different resolutions with `WInc3D` and the LES of Heck et al. (2023). The predictions of the low-order $\cos(\gamma)$ and $\cos(\gamma)^3$ models are also shown for reference. 94

3.10 Top view of the terrain elevation map of the operational wind farm with indicative computational domains aligned with inflow directions at 200°, 210°, and 220° (left), and close view of the turbine layout with IDs at 210° (right). 95

3.11 Computational domain representation of the operational wind farm simulations, showing terrain elevation, turbine positions, and instantaneous velocity field at lateral and outflow planes. 97

3.12 Horizontally- and time-averaged streamwise velocity of the ABL in the operational wind farm precursor simulation compared with log-law fits (a), and resolved stress profiles compared against theoretical stress (b). 99

3.13 Surface roughness map prescribed for the operational wind farm: $y_0 = 0.032$ m for open terrain (beige) and $y_0 = 0.5$ m for forested regions (green). 100

3.14 Turbine thrust coefficient and power curve, based on vendor data and used in the ADM for the operational wind farm. 101

3.15 Normalised turbine power at 210° inflow for the operational wind farm. LES predictions compared with SCADA mean values and standard deviations. 103

3.16 Normalised turbine power averaged over the 200°–220° sector for the operational wind farm. LES predictions compared with SCADA mean values and standard deviations. 103

3.17 Time-averaged hub-height streamwise velocity field at 210° inflow over the operational wind farm. Terrain contours (dashed lines) are shown at 25-m intervals, and turbine IDs are indicated. 105

3.18 Time-averaged hub-height streamwise turbulence intensity field at 210° inflow over the operational wind farm. Terrain contours (dashed lines) are shown at 25-m intervals, and turbine IDs are indicated. 106

4.1 Schematic of the 2D hill with two wind turbines placed around it showing the design variables for optimisations BO1 (a) and BO2 (b). The example is for $x_1 = 0.0$, $h_1 = 1.0$, $\alpha_1 = 0°$, $x_2 = 3.75$, and $h_2 = 1.0$. 110

4.2 Power signals with total and cumulative averages and 99% CI for (a) wind turbine placed at the hilltop ($x/D = 0$) and (b) wind turbine immersed in the hill's wake ($x/D = 5.0$). The grey region indicates the initial transient discarded. 111

4.3 Deficit from the optimum power as a function of the cumulative simulation number for (a) BO1 and (b) BO2. The red dashed line indicates the -1% power deficit. 113

4.4 Instantaneous flow visualisation of the optimum cases BO1 (a) and BO2 (b), showing normalised velocity magnitude and vorticity contours (white). 114

4.5 Power series with cumulative averages and 99% CI of (a) the Reference case and the optimum cases of (b) BO1 and (c) BO2. 116

4.6 Normalised time-averaged streamwise velocity fields (left column) and streamwise turbulence intensity fields (right column) at the symmetry plane for the Reference case (top row) and both optimised scenarios, BO1 (middle row) and BO2 (bottom row). The black lines are streamlines of the velocity field, while the white lines indicate the contours of zero streamwise velocity. 118

4.7 Velocity deficit caused by the upstream turbine in (a) BO1 and (b) BO2 evaluated 0.25D upstream of the downstream turbine. Contours are shown with dashed lines and the grey area indicates the hill's section. 120

4.8 Mean and standard deficit power prediction of the GP surrogate model against the tilt angle by fixing x_1, x_2 and h_2 to the BO2 optimum values. The GP surrogate model is built with the 100 evaluations of BO2, plus the BO2 optimum without tilt. 121

4.9 Visualisation of the hill recirculation edge depicted by the contour of zero-velocity for the Reference (a), BO1 (b), and BO2 (c) cases. 122

5.1 Horizontally- and time-averaged streamwise velocity (left) and streamwise turbulence intensity profile (right) of the ABL in the precursor simulation of the utility-scale wind farm, compared against analytical profiles. 128

5.2 Synthetic complex terrain elevation map with example turbine layout showing the 32 design variables. Labels (A–F) correspond to major terrain features. The dashed lines represent 20-m contours of the terrain elevation. The black rectangle is the available land to place turbine centres. 129

5.3 Normalised power output versus cumulative simulation number for WFLO-2D-270°, comparing uninformed BO with physics-informed cases. The first 32 LHS cases are marked by the vertical grey dashed line. 134

5.4 Farm layouts from each of the four sampling strategies, overlaid on the time-averaged hub-height velocity field from the base flow (no turbines). Top left: Cluster-Based Sampling; top right: Ridge-Based Layouts; bottom left: Velocity-Ordered Zone Sampling; bottom right: Physics-Informed Custom Layouts. Contours indicate increasing velocity regions. 135

5.5 Deficit from the optimum power as a function of the cumulative simulation number for (a) WFLO-1D-270°, (b) WFLO-2D-270°, (c) WFLO-2D-210°, and (d) WFLO-2D-MWD. Blue points show the deficit of each evaluated layout relative to the best case; the black line tracks the best layout found so far. Vertical dashed lines mark the end of the initial sampling phase (64 layouts) and the transition from exploration to exploitation (192 layouts). The red dashed line indicates a 1% power deficit. 138

5.6 Optimal layouts for the four WFLO cases overlaid on terrain height. Each layout was obtained after 512 LES evaluations using physics-guided BO. 139

5.7 Turbine-by-turbine normalised power output ($\hat{P}_T$) for the three optimal layouts: WFLO-1D-270°, WFLO-2D-270°, and WFLO-2D-210°. Turbines are sorted by power from highest to lowest for each layout, with original turbine IDs annotated above the bars. 141

5.8 Time-averaged hub-height flow fields for the optimal single-direction layouts. Left column: streamwise velocity. Right column: turbulence intensity. Top to bottom: WFLO-1D-270°, WFLO-2D-270°, WFLO-2D-210°. Terrain contours (dashed lines) are shown at 20-m intervals. Flow fields are sampled at turbine hub height, defined locally as a constant vertical distance of D above the terrain. 141

5.9 Power output for each turbine across four wind directions of the optimal layout for WFLO-2D-MWD. Turbines are sorted by weighted average power, with error bars indicating directional sensitivity. 143

5.10 Time-averaged hub-height streamwise velocity fields for the optimal multi-directional layout under each wind direction. Top row: 300° (left) and 270° (right). Bottom row: 240° (left) and 210° (right). Terrain contours (dashed lines) are shown at 20-m intervals. 144

5.11 Time-averaged hub-height streamwise turbulence intensity fields for the optimal multi-directional layout under each wind direction. Terrain contours (dashed lines) are shown at 20-m intervals. 144

5.12 Instantaneous flow visualisation showing velocity magnitude contours at the boundaries and vorticity contours (white) for the optimal multi-directional layout under 270° inflow. 147

5.13 Cross-sectional view at z = 1865 m showing time-averaged velocity field (top) and velocity deficit (bottom) for turbines 10 and 11 in the optimal WFLO-1D-270° layout. 148

5.14 Hub-height velocity deficit fields for the optimal multi-directional layout across four wind directions. Red regions indicate velocity deficits (turbine wakes), with blue regions showing flow acceleration induced by turbine presence. Top row: 300° (left) and 270° (right). Bottom row: 240° (left) and 210° (right). Terrain contours (dashed lines) are shown at 20-m intervals. 149

Chapter 1

Introduction

1.1 Context and Research Motivation

Wind energy has become a central pillar of the renewable transition, delivering large-scale, low-carbon electricity at competitive costs. By the end of 2024, global installed capacity reached 1,136 GW, following a record 117 GW of new installations during the year, of which 109 GW came from onshore projects and 8 GW offshore (Global Wind Energy Council, 2025). This confirms both the maturity of the technology and its continuing expansion as a reliable contributor to decarbonisation.

The COP28 agreement set a global target to triple renewable energy capacity to at least 11,000 GW by 2030, a milestone considered essential for meeting the Paris Agreement and longer-term net-zero pathways. For the wind sector, this implies a step change in the pace of deployment. GWEC Market Intelligence projects that annual additions will rise to 139 GW in 2025, with about 981 GW of new capacity expected by 2030 under current policies. This corresponds to an average of 164 GW per year, which, while significant, remains well below the $\sim$320 GW per year required for wind energy to align with the COP28 ambition (Global Wind Energy Council, 2025). This gap underlines the urgency of accelerating wind deployment and motivates continued research into technologies and strategies that can deliver higher efficiency and greater reliability at scale.

The United Kingdom has positioned wind energy as a cornerstone of its long-term decarbonisation strategy. In line with the national commitment to achieve net zero greenhouse gas emissions by 2050 (Department for Business, Energy & Industrial Strategy, 2019), offshore and onshore wind are expected to play a central role in the transformation of the power system. As of mid-2024, Great Britain had approximately 14.8 GW of offshore wind and 14.2 GW of onshore wind capacity installed (Department for Energy Security and Net Zero, 2024). The UK Government's *Clean Power 2030 Action Plan* outlines national targets for 2030 of 43–50 GW of offshore wind and 27–29 GW of onshore wind (Department for Energy Security and Net Zero, 2024). While offshore capacity is set to dominate, onshore deployment remains a critical component of the UK strategy, particularly as available sites become scarcer and projects increasingly extend into more challenging environments.

A central factor behind the rapid cost reduction and competitiveness of wind energy has been the continuous increase in turbine size. Current commercial onshore machines now reach 8–10 MW with rotor diameters above 200 m, and estimates suggest onshore rotors could approach 270 m by 2030 (Figure 1.1). Offshore platforms have grown even more rapidly, moving from the Siemens 3.6 MW "workhorse" with a 120 m rotor diameter to machines in the 15–20 MW class with diameters exceeding 250 m. Recent developments by Dongfang (16 MW, 128 m blades), CSSC Haizhuang (18 MW, 260 m rotor), and Mingyang (18–20 MW, 292 m rotor, with plans for 22 MW and 310+ m) demonstrate the accelerating pace of offshore turbine scaling (Global Wind Energy Council, 2024).

The engineering implications of rotor scaling are profound. At the scale of modern utility turbines, accurate representation of the dynamic interactions between the atmospheric boundary layer (ABL), terrain, and wind turbine wakes becomes essential for predicting wind farm power output and turbine loading. As blade tips extend higher into the ABL, interactions with vertical shear, turbulence intensity, and stability-driven stratification become more pronounced. Existing rules for turbine spacing and farm operation, traditionally expressed as multiples of rotor diameter, require reassessment to ensure their validity at these scales. Reliable prediction of these effects demands computational approaches that can resolve the coupled dynamics of the atmosphere, terrain, and turbine wakes. These challenges emphasise the need to link scales

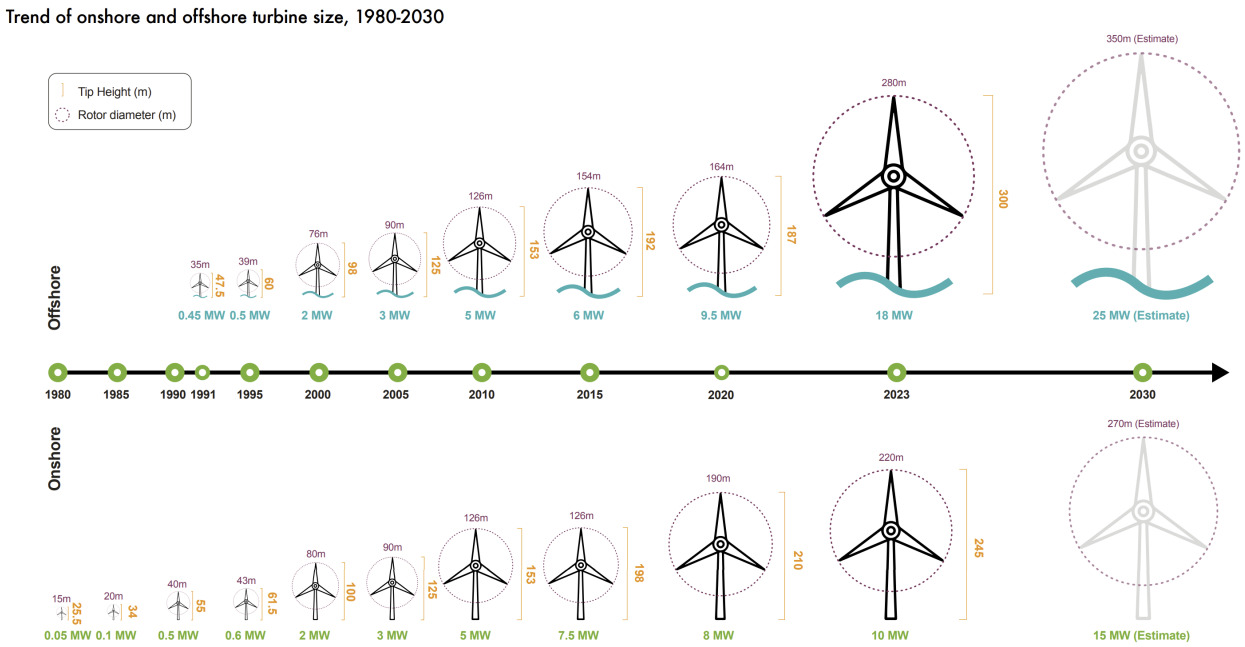


Figure 1.1: Trend of onshore and offshore wind turbine size (rotor diameter and tip height), 1980–2030. Figure reproduced from Global Wind Energy Council (2024), with permission from GWEC.

across the atmospheric system, from synoptic and mesoscale weather down to rotor-scale turbulence, and to develop models capable of integrating aerodynamics, structural dynamics, and energy system considerations for resilient design (Veers et al., 2019).

As prime flat terrain and offshore sites become increasingly occupied, and with limited space offering optimal conditions alongside social acceptance constraints near urban areas, wind farms are increasingly deployed in complex landscapes (Clifton et al., 2022). Examples of such installations are shown in Figures 1.2 and 1.3, where turbines are sited across mountainous regions with valleys, steep slopes, and vegetation. These environments amplify the challenges of rotor scaling, as terrain-induced variability interacts with boundary-layer dynamics and turbine wakes.

1.2 Complex Terrain Flow Physics

The interaction between atmospheric flow, terrain-induced features, and wind turbine wakes represents one of the key challenges for reliable performance prediction in complex terrain.



Figure 1.2: A wind farm in mountainous terrain with vegetation. Figure licensed from Getty Images.



Figure 1.3: Wind turbines located near a steep valley. Figure reproduced with permission from Nordex SE.

In these environments, suboptimal layouts can lead to substantial energy losses and increased fatigue loading (Elgendi et al., 2023). Terrain irregularities produce localised speed-up on windward slopes and deceleration with possible flow separation and recirculation in the lee, all of which alter turbine wake recovery, deflection, and power generation (Stevens and Meneveau, 2017; Porté-Agel et al., 2019).

The first studies on wind turbines in complex terrain were conducted via wind tunnel experiments and field measurements during the early 1990s (Taylor and Smith, 1991; Helmis et al., 1995). More recently, Tian et al. (2013) performed wind-tunnel flow measurements and recorded the power output and fatigue loads of an array of five wind turbines under the effect of a 2D Gaussian hill. It was found that the optimal placement of a wind turbine is on the windward side of the hill due to the higher velocity and lower turbulence intensity. Howard et al. (2015) and Howard et al. (2016) performed a wind tunnel investigation on the effects on a wind turbine placed in the wake induced by a 3D sinusoidal hill and found a drop in performance. Hyvärinen and Segalini (2017a), Hyvärinen and Segalini (2017b), and Hyvärinen et al. (2018) compared the performance of wind turbines located in a flat terrain with a hilly terrain both experimentally and numerically and found that a faster wake recovery is obtained for the hilly terrain case. Dar and Porté-Agel (2022) used particle-image velocimetry in a wind tunnel to study the wake of wind turbines located on different escarpments and found that the

high levels of turbulence intensity induced by the forward-facing step escarpment led to faster wake recovery compared with the ramp-shaped ones.

Field measurements are more limited, usually due to industry confidentiality, and surveys among experts identify an interest in more public validation campaigns (Wingerden et al., 2020). Subramanian et al. (2016) measured the wake evolution of multi-megawatt wind turbines located in flat (Altenbruch II) and complex terrain (Mont Crosin) with an instrumented drone. It was found that the near-wake extension in complex terrain was shorter than in flat terrain. Hansen et al. (2016) extracted supervisory control and data acquisition (SCADA) data from a wind farm in Shaanxi with complex terrain and indicated that the far wake is diurnal dependent. During the daytime and early evening, the wake was lifted, while it moved down and followed the terrain during the nighttime. Several studies have been made from the field measurement campaign of Perdigão, a double-ridge complex terrain with a wind turbine on top with an extensive database of the flow characteristics around it (Mann et al., 2017; Fernando et al., 2019; Menke et al., 2018; Barthelmie and Pryor, 2019). Nevertheless, such examples remain rare, and the limited availability of public datasets leaves uncertainty in validating computational models and in exposing their added value. This validation gap between academic simulations and operational wind-farm data undermines confidence in adopting high-cost methods for industrial practice.

The main coupling mechanisms between terrain and turbine wakes can be summarised as speed-up effects, recirculation zones, flow separation, and change in wake recovery or deflection. While canonical forms of these mechanisms are relatively well understood from wind-tunnel and field studies, their prediction in realistic configurations remains uncertain. The inherently unsteady nature of these interactions, combined with their sensitivity to atmospheric stability and turbine operating parameters, requires modelling frameworks that balance tractability with sufficient fidelity to capture the dominant physics. In recent years, numerical studies have gained prominence, particularly turbulence-resolving approaches such as large-eddy simulation (LES), which provide the necessary physical detail to explore terrain–ABL–wake interactions.

1.3 Modelling and Optimisation Across Fidelities

The modelling of wind farm aerodynamics can be organised into a hierarchy of methods that balance accuracy and computational cost. At one end of the spectrum, low-order analytical wake models provide rapid, time-averaged predictions of velocity deficits suitable for applications requiring thousands of evaluations. Mid-order methods, typically based on Reynolds-averaged Navier–Stokes (RANS) formulations, resolve additional flow physics at a moderate computational cost. At the high-fidelity end, LES and related approaches explicitly resolve a broad range of turbulent scales and capture dynamic terrain–wake coupling, but at a cost several orders of magnitude higher. Traditionally, wind-farm design and optimisation workflows have followed this hierarchy: low-order methods are used to explore large design spaces efficiently, while higher-fidelity models are applied selectively for detailed analysis and validation.

1.3.1 Low-Order Methods

Low-order methods are typically based on analytical wake models that represent the mean velocity deficit downstream of a turbine using simplified assumptions about wake expansion and entrainment. These models neglect turbulent fluctuations and meandering, but their negligible cost (typically $\mathcal{O}(10^{-3})$ CPU-h per evaluation; i.e. milliseconds on a laptop) makes them attractive for optimisation problems where tens of thousands of simulations are required. Such models underpin both academic and industrial wind farm design tools. In academia, they are commonly implemented within open-source frameworks such as *Flow Local Optimisation & Redirection via Induction in Steady-state* (FLORIS) (Mudafort et al., 2023), developed by NREL for wind farm simulation and control studies. In industry, similar engineering wake models are embedded within commercial, bankable software suites such as WindPRO or WindFarmer, which are widely used for wind farm design and assessment.

The Park model, introduced by Jensen (1983), is among the earliest and most widely used wake models. It assumes a top-hat velocity deficit with linear wake expansion, controlled by an empirical entrainment coefficient. The simplicity of this formulation allows for runtimes on the order

of milliseconds, but the assumption of a uniform top-hat deficit and lack of momentum conservation limit its predictive accuracy. Frandsen et al. (2006) introduced momentum-conserving corrections, while Bastankhah and Porté-Agel (2014) derived a Gaussian deficit model that more realistically represents the radial distribution of velocity. The Gaussian formulation conserves momentum and better matches experimental measurements of wake recovery (Porté-Agel et al., 2019), and is therefore widely adopted in both academic and industrial contexts.

For wind-farm layout optimisation (WFLO), these analytical formulations estimate the velocity deficit behind each rotor and reconstruct the farm-scale flow by superposition. However, the steady-state assumptions and simplified wake shape and lateral expansion underlying these models were derived for flat terrain and break down in topographically complex sites. Politis et al. (2011), Segalini (2017), and Porté-Agel et al. (2019) demonstrated that low-order models significantly underpredict terrain-driven flow separation and pressure-gradient effects. Moreover, linearised formulations assume small flow perturbations from turbine-induced forces, thus neglecting non-linear interactions and flow separation common in complex terrain (Segalini, 2017).

Recent efforts have attempted to extend analytical wake models to incorporate terrain effects whilst retaining computational efficiency. Zong and Porté-Agel (2020) and Du et al. (2024) derived momentum-conserving wake superposition formulae that remain valid under strong pressure gradients typical of mountainous terrain. Farrell et al. (2021) developed modifications to the FLORIS model, improving predictions on heterogeneous terrain without explicitly modelling the wind farm geometry. Although such advances reduce error, analytical wake models still exhibit large uncertainty in complex terrain. Nonetheless, they remain attractive in commercial software due to millisecond runtimes and have been incorporated into terrain-aware WFLO frameworks (Brogna et al., 2020; Liu et al., 2022).

Consequently, analytical wake models are best regarded as efficient but limited tools, to be complemented by higher-fidelity simulations when terrain and atmospheric complexity play a dominant role. This limitation motivates the development of mid-fidelity approaches that can capture additional flow physics at moderate computational cost.

1.3.2 Mid-Fidelity RANS Approaches

RANS methods provide a balance between computational tractability and physical fidelity. RANS simulations decompose the velocity field into mean and fluctuating components and solve for the time-averaged flow field, requiring turbulence models to close the system of equations. While this approach cannot capture unsteady wake meandering or turbulent fluctuations directly, it can resolve terrain-induced flow features such as speed-up, blockage, and flow curvature through direct modelling of turbine-terrain interactions at moderate cost (typically $\mathcal{O}(10\text{–}10^2)$ CPU-h per evaluation; i.e. minutes–hours on a workstation). RANS simulations typically employ actuator disc models (ADM) to represent turbines, allowing investigation of terrain effects that strongly influence turbine performance.

RANS validations have been demonstrated across diverse terrains: 25-turbine sites in China (Sessarego et al., 2018), the Alaiz experimental farm (Bretos et al., 2024), southern Italian installations (Astolfi et al., 2018), and mountainous Chinese locations (Cheng et al., 2024). Despite their efficiency and widespread use in engineering practice, RANS models are known to be overly diffusive in turbine wakes, limiting their ability to represent unsteady and anisotropic turbulence (Sanderse et al., 2011).

For WFLO, RANS offers significant advantages through adjoint-enabled frameworks that allow gradient-based optimisation at practical computational cost. King et al. (2017) pioneered the use of adjoint methods with steady-state RANS to optimise turbine placement in utility-scale farms, enabling optimisation across full wind roses with realistic inflow distributions. Allen et al. (2020) extended NREL's WindSE framework to include terrain-aware adjoints, demonstrating a reduction in mean error relative to SCADA data and significant annual energy production (AEP) gains by exploiting terrain-induced flow features. Antonini et al. (2020) generalised this adjoint CFD methodology to arbitrary topographies and wind resources, achieving a 6.6% AEP increase in real terrain whilst maintaining computational feasibility on modern high-performance computing (HPC) systems. The development of CFD-based surrogate models has further extended RANS applicability. Wang et al. (2024) demonstrated that Kriging surrogates built from limited CFD training sets can significantly reduce computational requirements and

maintain optimisation fidelity through adaptive sampling strategies, including optimisations on an 8-turbine case with Gaussian hill terrain. This surrogate-assisted approach enables broader exploration of design spaces whilst preserving the physical insights that RANS provides over analytical models.

Overall, RANS approaches remain a valuable mid-fidelity option; however, their steady-state nature means they cannot resolve unsteady wake meandering, dynamic terrain-wake interactions, or flow separation in steep terrain. These limitations become critical in complex topographies where such phenomena dominate performance.

1.3.3 High-Fidelity Large-Eddy Simulation

LES has become the benchmark for wind-farm flow modelling due to its ability to resolve the dominant turbulent scales whilst parameterising only the subgrid contributions. Compared with low-order models and RANS, LES can capture unsteady wake dynamics, meandering, terrain-induced flow separation, and recirculation at high fidelity, though at a substantially higher computational cost (typically $\mathcal{O}(10^3$–$10^4)$ CPU-h per evaluation; i.e. hours–days on HPC clusters).

LES has been applied extensively to investigate turbine siting over idealised topographies. Shamsoddin and Porté-Agel (2017) studied Gaussian hills and highlighted the role of flow acceleration and recirculation in downstream wake development. Yang et al. (2015) performed LES of a wind turbine operating in the wake induced by a three-dimensional hill and demonstrated good agreement with experimental measurements. They showed that the interaction between the hill-induced flow and the turbine wake leads to enhanced entrainment and a faster wake recovery downstream of the turbine. More recently, Liu and Stevens (2020a) and Zhang et al. (2022) validated immersed-boundary LES against the wind-tunnel experiments of Cao and Tamura (2006) around a Gaussian hill, demonstrating the trade-offs of hilltop, windward, and leeward siting. Wang et al. (2023) extended these analyses to irregular ridge topographies. These studies establish LES as the reference tool for understanding topographic effects on wake behaviour, despite the associated computational burden.

LES studies have demonstrated the importance of site-specific analysis: Dar et al. (2019) showed that wake self-similarity deteriorates more rapidly in complex terrain due to terrain-induced shear variations, and Troldborg et al. (2022) showed that turbine power curves in complex terrain can significantly differ from flat-terrain expectations, both requiring individual site verification using high-fidelity simulations. Real-world LES validations confirm model fidelity across diverse environments, from utility-scale wind farms in hilly terrain (Yang et al., 2018) to forested, terrain-influenced installations where SCADA data validation has been achieved for multiple wind directions (Abedi, 2023). Bernardoni et al. (2023) demonstrated that even modest upstream terrain undulations can significantly increase turbine load fluctuations and power variability, effects that simplified models may underestimate.

LES for WFLO has only recently become computationally tractable, owing to advances in HPC and efficient optimisation algorithms (Breton et al., 2017). Because of the high cost of each simulation, coupling LES with efficient optimisation strategies is essential. Bayesian optimisation (BO) has emerged as a particularly suitable candidate, since it is well-suited for expensive black-box functions where the number of evaluations must be minimised (Diessner et al., 2022; Bempedelis and Magri, 2023). BO constructs a Gaussian process (GP) surrogate from available LES data, providing both predictions and quantified uncertainties across the design space. An acquisition function then selects new evaluation points by balancing exploitation of promising regions with exploration of uncertain areas. This sample efficiency, typically requiring tens to hundreds rather than thousands of evaluations, makes BO feasible for LES-driven WFLO.

Recent work has demonstrated the advantages of combining BO with high-fidelity simulation. Bempedelis et al. (2024) first showed that while analytical wake models are adequate for flat-terrain optimisation, LES becomes essential for wake steering problems involving complex flow physics. Mole and Laizet (2024) extended this concept through a multi-fidelity BO framework, coupling LES with analytical wake models to reduce cost in yaw-control optimisation while capturing key wake–ABL interactions missed at lower fidelity. This motivates the need to extend LES-driven BO into complex terrain, where the physics demand high fidelity and where industry-ready optimisation methods are still lacking.

1.4 Research Gaps and Objectives

To summarise, WFLO has matured on the basis of analytical wake models, which remain effective in flat terrain but lose reliability in complex topographies. Mid-fidelity RANS methods extend applicability to such sites, yet their steady-state closures limit their ability to capture unsteady flow physics. LES resolves the necessary physics, though at higher computational cost unless coupled with efficient optimisation strategies. Bayesian optimisation offers this efficiency, but its use alongside LES has so far been restricted to simplified test cases. Together, these developments define a set of unresolved research gaps that this thesis addresses:

- **Gap 1:** Existing WFLO methods rely on analytical wake or RANS-based models that cannot capture unsteady turbulence or nonlinear terrain–wake interactions.
- **Gap 2:** LES predictions are rarely validated against operational wind-farm data, and the scarcity of such data limits their credibility for design applications.
- **Gap 3:** The influence of modelling choices (such as inflow specification, domain design, and turbine parameterisation) on predictive skill at utility scale remains underexplored due to limited real-world validation data, particularly in operational wind-farm contexts.
- **Gap 4:** Only a handful of LES-driven optimisation studies exist, and these are restricted to flat terrain problems, leaving complex terrain unexplored.
- **Gap 5:** The applicability of BO to utility-scale wind-farm layout problems has not been sufficiently demonstrated.

This thesis addresses these gaps in a staged manner, with objectives that progressively build from canonical validation to utility-scale optimisation:

Computational framework validation. The first objective is to establish and validate a high-fidelity computational framework for wind farms in complex terrain. This involves extending the LES solver `WInc3D` with immersed boundary and wall models, validating against

canonical wind-tunnel experiments of flow over hills, benchmarking turbine performance in idealised terrain, and quantifying the accuracy of LES predictions for terrain-induced flow features and wake dynamics. This ensures that subsequent optimisation studies are based on reliable physical representations of terrain–wake interactions.

Terrain–wake coupling mechanisms. The second objective is to investigate fundamental coupling mechanisms between terrain and wake dynamics that govern wind-farm performance in complex topography. This includes characterising terrain-induced wake–wake interactions, quantifying how turbine siting affects wake recovery and power generation in hilly terrain, and identifying the physical processes responsible for wake deflection, acceleration, and turbulence enhancement. Understanding these mechanisms is essential for developing physics-informed optimisation strategies.

LES-driven optimisation development. The third objective is to demonstrate LES-driven optimisation for idealised terrain configurations by integrating LES with BO in a two-turbine configuration around a 2D hill (Chapter 4). This proof-of-concept study explores turbine position, hub height, and tilt angle as design variables, and assesses the feasibility and convergence of the LES–BO methodology. It represents the first demonstration of LES-based optimisation in complex terrain.

Industry-scale simulation and analysis. A critical fourth objective is to increase the technology readiness level through operational data validation. This PhD has been conducted as a CASE studentship in partnership with SSE Renewables Services (UK) Limited, with the explicit aim of ensuring that the research is informed by real operational challenges. SSE Renewables contributed to defining the research problem, engaged in regular project discussions, and provided SCADA data from one of their managed wind farms. This collaboration enabled the extension of the LES methodology to utility-scale farms, allowing direct comparison of LES power predictions against operational SCADA records, quantification of discrepancies, and identification of uncertainty sources. By grounding the methodology in industrial data and

practice, this objective is essential for establishing the reliability and relevance of high-fidelity simulations for wind-farm design.

Utility-scale methodology scaling. The fifth objective is to scale the methodology to realistic design problems by applying the LES–BO framework to a 16-turbine farm in synthetically generated terrain comprising ridges, valleys, and elevated zones (Chapter 5). The study compares black-box BO with physics-informed strategies that exploit terrain-induced flow features, across both single- and multi-directional optimisations. This demonstrates the practical viability of LES-driven optimisation at scales relevant to industry.

Assessment of high-fidelity value. The final objective is to assess when high-fidelity LES is necessary by identifying flow features and optimisation mechanisms that cannot be captured by conventional models, thereby clarifying the added value of LES for wind-farm design in complex terrain.

The integration of these objectives provides a comprehensive foundation for advancing wind-farm design in complex terrain, combining rigorous computational methods with practical validation and demonstrating clear pathways toward improved energy capture and reduced costs in challenging topographic conditions.

1.5 Research Innovation and Contributions

This thesis addresses the identified research gaps through three principal contributions that advance the state of the art in LES-driven wind-farm optimisation for complex terrain.

First, a validated high-fidelity simulation framework is established by extending the LES solver `WInc3D` with an immersed boundary method (IBM), wall-stress modelling, and advanced turbine parameterisations for yawed and tilted rotors. Systematic validation against canonical experiments and operational SCADA data establishes confidence bounds for practical applications (Gaps 1–3).

Second, the systematic scaling of LES-driven optimisation from proof-of-concept to utility-scale applications is demonstrated. By developing workflows that efficiently couple high-fidelity simulation with BO on HPC platforms, this work establishes that LES-based design can transition from academic validation studies to practical methodology. The framework enables systematic exploitation of terrain features whilst avoiding adverse regions, capabilities that remain inaccessible to conventional design tools (Gap 4).

Third, physics-informed optimisation strategies are introduced to address the curse of dimensionality in high-dimensional wind-farm design. By leveraging engineering knowledge of terrain-induced flow patterns, these strategies accelerate the search process and enable efficient optimisation at utility scale where conventional black-box approaches struggle (Gap 5).

Collectively, these contributions establish LES-driven BO as a viable methodology for wind-farm layout design in complex terrain, representing a fundamental step forward in enabling more reliable and practically feasible optimisation where conventional methods fail.

1.6 Thesis Structure

The thesis is organised in six chapters. Chapter 2 introduces the computational framework, describing the LES solver `WInc3D` and the BO methodology that together form the backbone of the workflow. Chapter 3 establishes confidence in the LES framework through a progressive validation campaign, beginning with controlled wind-tunnel benchmarks and extending to an operational onshore wind farm with SCADA data. Chapter 4 then applies the validated LES–BO framework to a two-turbine configuration over an idealised Gaussian hill, providing a proof-of-concept optimisation study and clarifying terrain–wake coupling mechanisms. Chapter 5 scales the methodology to a 16-turbine farm over synthetically generated three-dimensional terrain, addressing the challenges of high-dimensional optimisation and introducing physics-informed seeding strategies. Finally, Chapter 6 synthesises the findings, discusses scientific and practical impacts, acknowledges limitations, and outlines directions for future work.

1.7 List of Publications

The work conducted during this PhD has been disseminated through a number of journal publications and conference presentations. Material from these manuscripts has been incorporated throughout the thesis, with references to the corresponding papers given at the beginning of each relevant chapter. A summary is provided below.

Journal Publications

- C. **Jané-Ippel**, R. Palacios, S. Laizet (2025). Physics-Informed Bayesian Optimisation of Utility-Scale Wind Farms in Synthetically Generated Complex Terrain Using LES. *Wind Energy*, accepted for publication; available as a preprint at Authorea. *Content included in Chapter 5.*

- K. Steiros, C. **Jané-Ippel**, F. Llewellyn, R. Palacios, R. Bayoan Cal, S. Laizet (2025). Improved wind power extraction downstream of a 3D hill using a tilted actuator disk. *Journal of Renewable and Sustainable Energy* **17**(5), 053303. *Not included in the thesis.*

- C. **Jané-Ippel**, N. Bempedelis, R. Palacios, S. Laizet (2024). Bayesian Optimisation of a Two-Turbine Configuration Around a 2D Hill Using Large Eddy Simulations. *Wind Energy* **27**.11, 1412–1426. *Content included in Chapters 3 and 4.*

Conference Proceedings

- C. **Jané-Ippel**, N. Bempedelis, R. Palacios, S. Laizet (2024). Optimising wind turbine placement in hilly terrains: insights from Bayesian optimisation based on large-eddy simulations. *Proceedings of TSFP-13 (Montréal). Content included in Chapter 4.*

- C. **Jané-Ippel**, N. Bempedelis, R. Palacios, S. Laizet (2023). High-fidelity simulations of wake-to-wake interaction in an atmospheric boundary layer over complex terrain. *Journal of Physics: Conference Series* **2505**.1, 012033. *Content included in Chapter 3.*

Conference Presentations

- APS Division of Fluid Dynamics Meeting, Salt Lake City, USA (2024). "Predicting Wind Power of an Onshore Wind Farm in Complex Terrain with Large-Eddy Simulations."
- ECCOMAS, Lisbon, Portugal (2024). "Enhancing Wind Farm Efficiency in Complex Terrain: A Bayesian Optimisation Based on LES."
- 13th International Symposium on Turbulence and Shear Flow Phenomena (TSFP13), Montréal, Canada (2023). "Optimising Wind Turbine Placement in Hilly Terrains: Insights from Bayesian Optimisation."
- Wake Conference, Visby, Sweden (2023). "High-Fidelity Simulations of Wake-to-Wake Interaction in an ABL Over Complex Terrain."
- EAWE PhD Seminar, Hannover, Germany (2023). "Bayesian Optimisation of a Two-Turbine Layout Around a 2D Hill Using Large Eddy Simulations."
- UK Fluids Conference, University of Glasgow, UK (2023). "Bayesian Optimisation of a Two-Turbine Layout Around a 2D Hill Using Large Eddy Simulations."
- UK Fluids Conference, University of Sheffield, UK (2022). "Optimising Wind Farm Performance in Complex Terrain with LES and Bayesian Optimisation."

Chapter 2

Computational Methods

This chapter introduces the computational framework that underpins the work: a LES solver for modelling atmospheric boundary layer flows and turbine–terrain interactions, coupled with a BO methodology for navigating high-dimensional design spaces.

The LES solver `WInc3D` provides the physical fidelity required to capture terrain-induced accelerations, wake deflection, and turbulent interactions that are not represented in lower-order models. The BO framework enables efficient search through large design spaces where exhaustive evaluation would be computationally prohibitive. Together, these methods address the research gaps identified in Chapter 1 by combining turbulence-resolving simulations with sample-efficient optimisation.

The chapter is structured as follows. Section 2.1 describes the LES framework, including atmospheric boundary layer generation, immersed boundary treatment for complex terrain, and actuator-disc turbine representation. Section 2.2 outlines the BO methodology, covering design of experiments, surrogate modelling, and acquisition strategies. Together, these frameworks provide the methodological foundation for the validation and optimisation studies in subsequent chapters.

2.1 Large-Eddy Simulation Framework

LES resolves the energy-carrying motions of turbulence while modelling only the smallest scales. This makes it particularly well suited to wind-farm aerodynamics, where wakes, terrain-induced speed-up, and shear-layer dynamics span a wide range of spatial and temporal scales. Early turbulence-resolving strategies for wind energy relied on spectral discretisations based on Fourier or Chebyshev representations. These methods deliver excellent accuracy and have been widely used to study canonical turbulence problems and idealised wind farms (Calaf et al., 2010; Wu and Porté-Agel, 2013). In wind-energy applications they are often applied in the horizontal directions and combined with energy-conserving finite differences in the vertical (Porté-Agel et al., 2011; Stevens et al., 2014; Munters et al., 2016). However, the reliance on periodic boundary conditions restricts their direct use to infinitely repeating layouts or requires specialised boundary treatments.

To overcome these constraints, high-order alternatives have been developed. Spectral element methods provide one route and have gained popularity in wind-energy studies (Kleusberg et al., 2017; Chatterjee and Peet, 2018). High-order compact finite-difference schemes offer another practical approach to achieve "spectral-like" accuracy (Lele, 1992; Lamballais et al., 2021). By coupling derivatives at near-neighbour points within a compact stencil, these schemes retain excellent dispersive and dissipative properties while allowing a wider range of boundary conditions and robust near-wall treatment. This flexibility is particularly advantageous for ABL flows over complex terrain, where boundary conditions depart significantly from the periodic setting assumed in spectral methods.

The `Xcompact3D` framework (Bartholomew et al., 2020) exploits such compact schemes on structured Cartesian grids and efficient domain decomposition to run on modern supercomputers. These characteristics make `Xcompact3D` a natural backbone for the simulations and parameter sweeps pursued in this thesis.

A number of LES frameworks have been developed for wind-energy applications, including `EllipSys3D` (Sørensen, 1995; Nilsson et al., 2014), `EPFL` (Wu and Porté-Agel, 2015), `JHU`

(Stevens et al., 2014), SOWFA (Churchfield et al., 2012), WiTTS (Xie and Archer, 2014), SnS (Storey et al., 2012), YALES2 (Benard et al., 2018), PALM (Witha et al., 2014), and SP-Wind (Meyers and Meneveau, 2011). These solvers differ in their discretisation strategies, inflow-generation techniques, and turbine parameterisations, and have been validated against canonical benchmarks such as the Blind Test campaign (Krogstad and Eriksen, 2013; Pierella et al., 2014; Krogstad et al., 2015) and SCADA data from offshore wind farms such as HornsRev (Barthelmie et al., 2009) and Lillgrund (Dahlberg, 2009). Reviews of LES methodologies for wind energy highlight both the predictive power and the computational demands of such approaches (Breton et al., 2017). Within this landscape, the present work employs WInc3D, the wind-energy module of Xcompact3D, which has been further developed in this thesis to address the challenges of simulating wind farms in complex terrain.

2.1.1 Introduction to WInc3D

WInc3D builds on the high-order compact finite-difference schemes and parallelisation strategies of Xcompact3D. It employs up to sixth-order spatial discretisation on structured Cartesian grids, with incompressibility enforced to machine precision via an FFT-based Poisson solver. The computational domain of dimensions $L_x \times L_y \times L_z$ is discretised into $N_x \times N_y \times N_z$ nodes, yielding uniform grid spacings $\Delta x = L_x/N_x$, $\Delta y = L_y/N_y$, and $\Delta z = L_z/N_z$, where x, y, and z denote the streamwise, wall-normal (vertical), and spanwise directions, respectively. The grid flexibility is demonstrated by three distinct boundary conditions that can be applied along all three directions, providing multiple combinations to accommodate different flow configurations. A two-dimensional pencil decomposition enables scalability to $\mathcal{O}(10^4$–$10^5)$ processor cores, making large-scale wind-farm simulations feasible on modern HPC systems. The solver supports multiple time-integration schemes and incorporates modules for ABL inflow, turbine parametrisation, and immersed-boundary treatment.

Since its development, WInc3D has facilitated a diverse range of research efforts, demonstrating its versatility across the wide spectrum of length scales present in wind-energy flows, from large wind farms spanning several kilometres to wind tunnel scale turbines. The simulator

has supported wind-tunnel and utility-scale wake studies for engineering-model development (Bempedelis and Steiros, 2022), investigations of tip-vortex stability and breakdown in atmospheric boundary layers (Hodgkin et al., 2022; Hodgkin et al., 2023), analytical blockage corrections for high-solidity turbines (Steiros et al., 2022), and entrainment-based modelling of finite-length wind-farm flow development (Bempedelis et al., 2023). It has also served as a tool for the optimisation of wind-farm power output in flat terrain, including micro-siting and wake-steering strategies (Bempedelis et al., 2024) and multi-fidelity BO of yaw configurations (Mole and Laizet, 2024). Beyond wind-energy applications, the actuator-line module has been used to simulate UAV aerodynamics (Giannenas and Laizet, 2021), demonstrating the solver's versatility across rotorcraft flow problems.

The present thesis extends `WInc3D` to enable simulations of onshore wind farms over complex terrain. We couple the immersed-boundary method with wall-stress modelling to allow neutral ABL flows over steep hills to be represented at high Reynolds number. The actuator-disc representation was also enhanced with super-Gaussian kernels, ensuring that yawed and tilted turbines could be modelled accurately and without numerical artefacts despite the Cartesian grid. These developments, combined with improvements to the ABL modelling tools for terrain-inclusive simulations, have been integrated into `Xcompact3D` and made openly available, ensuring reproducibility and providing a foundation for future work. Together, these developments extend the simulator's capability from idealised flat terrain to realistic complex topographies.

The subsections that follow present the governing equations, numerical discretisation, and parallelisation approach of `WInc3D`, followed by the ABL inflow model, the immersed-boundary method for complex terrain, and the turbine parametrisation strategies employed in this work. Validation of these methods is reported in Chapter 3.

2.1.2 Governing Equations

The governing equations used in the wind farm simulations are the unsteady, incompressible, filtered Navier–Stokes (N–S) equations, based on an explicit LES formulation:

$$\frac{\partial \widetilde{u}_i}{\partial t} + \frac{1}{2}\left(\widetilde{u}_j \frac{\partial \widetilde{u}_i}{\partial x_j} + \frac{\partial \widetilde{u}_i \widetilde{u}_j}{\partial x_j}\right) = -\frac{1}{\rho}\frac{\partial \widetilde{p}}{\partial x_i} - \frac{\partial \tau_{ij}}{\partial x_j} + \nu \frac{\partial^2 \widetilde{u}_i}{\partial x_j \partial x_j} + \frac{F_i}{\rho}, \tag{2.1a}$$

$$\frac{\partial \widetilde{u}_i}{\partial x_i} = 0. \tag{2.1b}$$

Here, $\widetilde{u}_i$ $(i = 1, 2, 3)$ are the spatially filtered velocity components and $\widetilde{p}$ is the modified pressure, $\rho = 1.2\,\mathrm{kg/m^3}$ denotes the constant fluid density, τ_{ij} is the subgrid-scale stress tensor, and $\nu = 1.5 \times 10^{-5}\,\mathrm{m^2/s}$ is the kinematic viscosity. The forcing term F_i represents body forces used both for terrain representation via the IBM and for turbine modelling. The convective terms are written in skew-symmetric form, which reduces aliasing errors whilst preserving kinetic energy (Kravchenko and Moin, 1997). Here, x_i are the Cartesian coordinates and repeated indices imply summation (Einstein convention).

The subgrid-scale stresses are computed using the standard Smagorinsky model (Smagorinsky, 1963), which employs the Boussinesq hypothesis to relate them to the resolved strain rate. The closure is defined as:

$$\tau_{ij} = -2\nu_{\mathrm{sgs}}\widetilde{S}_{ij}, \tag{2.2}$$

where $\nu_{\mathrm{sgs}} = (C_S\Delta)^2|\widetilde{S}|$ is the subgrid-scale viscosity, C_S is the Smagorinsky constant, $\Delta = (\Delta x \Delta y \Delta z)^{1/3}$ is the geometric mean grid spacing, and $\widetilde{S}_{ij}$ is the resolved strain-rate tensor:

$$\widetilde{S}_{ij} = \frac{1}{2}\left(\frac{\partial \widetilde{u}_i}{\partial x_j} + \frac{\partial \widetilde{u}_j}{\partial x_i}\right). \tag{2.3}$$

The magnitude of the strain-rate tensor is computed as $|\widetilde{S}| = \sqrt{2\widetilde{S}_{ij}\widetilde{S}_{ij}}$. In the incompressible LES formulation adopted here, the isotropic component of the subgrid-scale stress, $\frac{1}{3}\,\tau_{kk}\,\delta_{ij}$ (where δ_{ij} is the Kronecker delta), is absorbed into the modified pressure $\widetilde{p}$ and is therefore not modelled explicitly.

Alternative SGS closures such as the dynamic Smagorinsky model (Germano et al., 1991) and the scale-dependent dynamic model (Porté-Agel et al., 2000) adapt the Smagorinsky coefficient locally to reduce dissipation and improve near-wall behaviour (Mason and Brown, 1999; Bou-Zeid et al., 2005), but at additional computational cost (Mehta et al., 2014). By contrast, the standard Smagorinsky model is significantly simpler, requiring only a fixed coefficient (modified near the wall as described in Section 2.1.4), and has a long record of robust performance in wind-energy LES. Given that this work introduces, for the first time in `WInc3D`, a combination of immersed boundary and wall-stress models, the standard Smagorinsky model was adopted as the most robust and computationally efficient choice compatible with the wall-stress treatment.

2.1.3 Numerical Solver

The numerical solution of the filtered incompressible N-S equations presented in the previous section requires robust discretisation schemes capable of accurately resolving the wide range of spatial and temporal scales present in wind-farm flows. The coupled nature of the momentum and continuity equations, where the pressure field acts to enforce the incompressibility constraint, necessitates a carefully designed solution strategy that maintains both numerical stability and computational efficiency.

To address these challenges, `WInc3D` employs a fractional step method (Kim and Moin, 1985) in which the momentum equations are advanced explicitly and the incompressibility condition is enforced through an implicit pressure projection. High-order compact finite-difference schemes are used for spatial discretisation, while the pressure Poisson equation is solved efficiently in spectral space. These methods, described in detail in the following sections, deliver the "spectral-like" accuracy required for turbulence-resolving simulations whilst retaining the flexibility to implement a variety of boundary conditions essential for realistic wind farm configurations. For notational convenience, the tilde is dropped in the numerical discretisation.

Temporal Discretisation and Fractional Step Method

Detailed descriptions of the general numerical framework can be found in (Laizet and Lamballais, 2009; Bartholomew et al., 2020). The temporal discretisation employs a three-step fractional step method (Kim and Moin, 1985), where the time advancement from t^n to t^{n+1} can be expressed as:

$$\frac{u_i^* - u_i^n}{\Delta t} = a_n\, H_i^n + b_n\, H_i^{n-1} + d_n\, H_i^{n-2} - c_n\, G_i^{n-\frac{1}{2}}, \tag{2.4a}$$

$$\frac{u_i^{**} - u_i^*}{\Delta t} = c_n\, G_i^{n-\frac{1}{2}}, \tag{2.4b}$$

$$\frac{u_i^{n+1} - u_i^{**}}{\Delta t} = -c_n\, G_i^{n+\frac{1}{2}}. \tag{2.4c}$$

where Δt is the simulation time step, H_i^n is a discrete operator for the sum of the convective (skew-symmetric form as in Equation (2.1a)), viscous, and body force terms of the N-S equations at time step n:

$$H_i^n = -\frac{1}{2}\left(u_j^n \frac{\partial u_i^n}{\partial x_j} + \frac{\partial u_i^n u_j^n}{\partial x_j}\right) + \nu \frac{\partial^2 u_i^n}{\partial x_j \partial x_j} + b_i^n, \tag{2.5}$$

where b_i^n denotes the discretised body force at time step n, and

$$G_i^m = \frac{1}{\rho}\frac{\partial p^m}{\partial x_i} \tag{2.6}$$

is the discrete gradient operator applied to the pressure field at sub-step m. In the fractional step method, the pressure is advanced at half time steps (e.g. $n \pm \frac{1}{2}$).

The coefficients a_n, b_n, and d_n depend on the chosen time-integration scheme. In multi-stage Runge–Kutta (RK) methods these coefficients vary from one sub-step to the next, whereas for explicit Adams–Bashforth (AB) schemes they are fixed once the scheme order is selected. In the present work, a third-order AB scheme is used for the convective and diffusive terms, giving:

$$a_n = \frac{23}{12}, \qquad b_n = -\frac{16}{12}, \qquad d_n = \frac{5}{12},$$

whilst the pressure gradient is advanced using $c_n = 1$ following the standard fractional step

procedure.

Pressure Poisson Equation and Solver

After the predictor step (Equation (2.4a)), a pressure Poisson equation is solved to enforce the incompressibility constraint, enabling the velocity correction in Equation (2.4c). In the absence of immersed boundaries, the equation reads:

$$\frac{\partial^2 p^{n+\frac{1}{2}}}{\partial x_j \partial x_j} = \frac{\rho}{c_n \, \Delta t} \frac{\partial u_i^*}{\partial x_i}, \tag{2.7}$$

where u_i^* is the intermediate velocity field from Equation (2.4a).

The Poisson equation is solved in spectral space using Fast Fourier Transforms (FFT) in periodic directions, exploiting the uniform Cartesian mesh structure. This diagonalises the Laplacian operator in those directions, reducing the three-dimensional problem to a set of independent one-dimensional Helmholtz equations along the remaining coordinate. When one or more directions are non-periodic, the Fourier basis is replaced by alternative discrete transforms to enforce the correct physical boundary conditions: discrete cosine transforms (DCT) for homogeneous Neumann boundaries (zero normal gradient), and discrete sine transforms (DST) for homogeneous Dirichlet boundaries (zero value). Mixed configurations combine Fourier expansions in periodic directions with cosine or sine transforms in non-periodic ones.

The uniform mesh spacing ensures that each transformed one-dimensional Helmholtz equation retains constant coefficients. This property is exploited to minimise computational cost, as described in Section 2.1.3.

To reduce spurious oscillations in the pressure solution, the pressure field is stored on a half-staggered grid relative to the velocity nodes (Figure 2.1). Before solving, velocity values are interpolated to these staggered pressure nodes using sixth-order compact interpolation schemes following Lele (1992).

The parallel implementation of the Poisson solver, including domain decomposition, data trans-

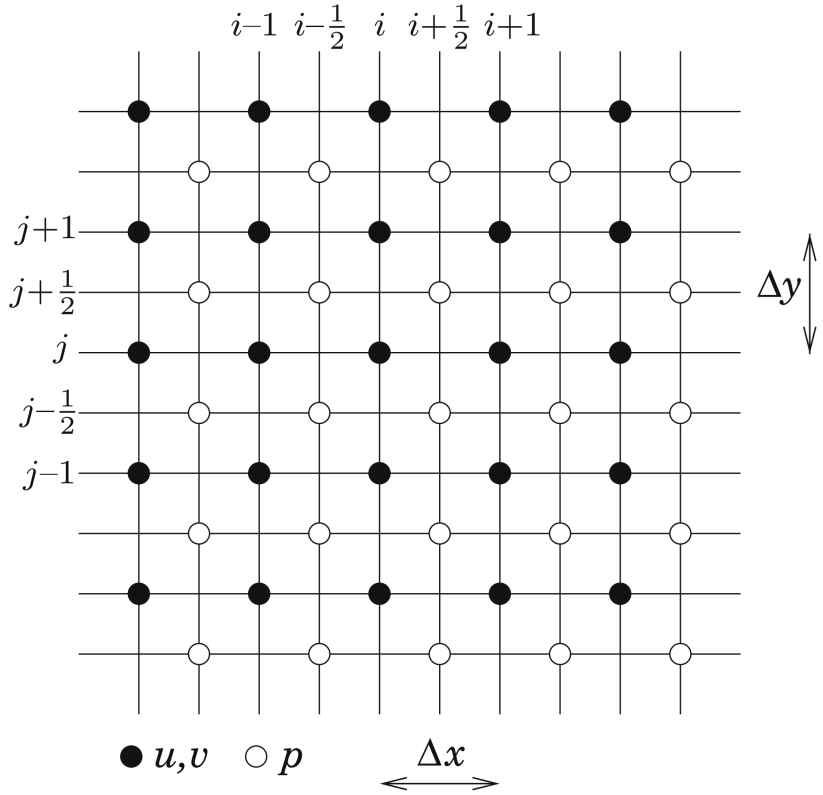


Figure 2.1: Half-staggered pressure field, showing the location of velocity and pressure nodes. Figure reproduced from Laizet and Lamballais (2009), with permission from Elsevier.

positions, and the efficient solution of the reduced systems, is described in 2.1.3.

Spatial Discretisation

The spatial discretisation employs high-order compact finite-difference schemes (Lele, 1992). For the present work, the first derivative is computed using the sixth-order implicit compact scheme. Let $f_i = f(x_i)$ denote the discrete value of a function f at node x_i along the uniform Cartesian mesh, and let f'_i denote the approximate first derivative at x_i. The scheme is expressed as:

$$\alpha f'_{i-1} + f'_i + \alpha f'_{i+1} = a\frac{f_{i+1} - f_{i-1}}{2\Delta x} + b\frac{f_{i+2} - f_{i-2}}{4\Delta x}, \tag{2.8}$$

where the coefficients α, a, b are given in Table 2.1. The scheme couples unknown derivative values at neighbouring points through a tridiagonal system, yielding a compact stencil and enabling accurate evaluation near boundaries.

Table 2.1: Sixth-order discrete operators. The column α/α_f gives the implicitness coefficient α for derivative and interpolation stencils, and the damping coefficient α_f for the filter.

Operator	α / α_f	a	b	c	d
First derivative	1/3	14/9	1/9	–	–
Second derivative	2/11	12/11	3/11	–	–
Interpolation	9/62	63/62	17/62	–	–
Filtering	0.49	$(11 + 10\alpha_\mathrm{f})/16$	$(15 + 34\alpha_\mathrm{f})/32$	$(-3 + 6\alpha_\mathrm{f})/16$	$(1 - 2\alpha_\mathrm{f})/32$

Beyond the formal truncation error, the key advantage of high-order compact schemes lies in their ability to maintain high accuracy across a much broader range of wavenumbers than standard finite-difference methods. While low-order schemes rapidly lose accuracy as wavelengths approach the grid cut-off, high-order compact schemes preserve both large-scale motions and a greater portion of the small-scale turbulent content. The implicit formulation further improves this spectral resolution compared to explicit schemes of the same order, allowing more of the turbulent kinetic energy spectrum to be resolved for a given grid spacing. This reduces numerical dissipation in the inertial subrange and enhances the fidelity of turbulence-resolving simulations such as LES.[1]

[1]Here, "implicit" refers to the spatial discretisation (coupling of derivative values through tridiagonal systems), not to the temporal integration scheme.

Second-order derivatives are computed using an analogous sixth-order implicit compact formulation (Table 2.1):

$$\alpha f''_{i-1} + f''_i + \alpha f''_{i+1} = a\frac{f_{i+1} - 2f_i + f_{i-1}}{\Delta x^2} + b\frac{f_{i+2} - 2f_i + f_{i-2}}{4\Delta x^2}, \tag{2.9}$$

The same spectral-resolution benefits discussed for the first derivative apply directly here, extending the range of accurately represented wavenumbers and improving the resolution of small-scale gradients in diffusive and Laplacian terms. This improves the representation of subgrid-scale energy transfer in LES without incurring the excessive dissipation typical of lower-order schemes.

For the partially staggered mesh approach employed in `WInc3D`, mid-point interpolation schemes are required to evaluate variables at staggered node locations. The interpolation operator is expressed as:

$$\alpha f_{i-1/2} + f_{i+1/2} + \alpha f_{i+3/2} = a\frac{f_{i+1} + f_i}{2} + b\frac{f_{i+2} + f_{i-1}}{2}, \tag{2.10}$$

with coefficients provided in Table 2.1. This interpolation procedure is computationally efficient whilst maintaining the high-order accuracy of the overall discretisation.

To control aliasing errors inherent in high-order schemes, a spatial filter is applied in each coordinate direction to the velocity field at every time step; the discrete low-pass filter adopts the tridiagonal formulation of Motheau and Abraham (2016) and is designed to retain sixth-order accuracy:

$$\alpha_\mathrm{f}\widehat{f}_{i-1} + \widehat{f}_i + \alpha_\mathrm{f}\widehat{f}_{i+1} = af_i + \frac{b}{2}(f_{i+1} + f_{i-1}) + \frac{c}{2}(f_{i+2} + f_{i-2}) + \frac{d}{2}(f_{i+3} + f_{i-3}), \tag{2.11}$$

where the coefficients a, b, c, d are given in Table 2.1, and the damping coefficient is set to $\alpha_\mathrm{f} = 0.49$. This choice introduces artificial dissipation only near the grid cut-off whilst leaving the resolved turbulence essentially unaffected.

For complex terrain applications using the IBM, spatial derivatives must be evaluated throughout the entire computational domain, including both fluid and solid regions. The compact

schemes maintain their homogeneity across the domain, which is crucial for the robustness of the immersed boundary implementation and ensures smooth solutions at fluid-solid interfaces.

Boundary Conditions

An important feature of "spectral-like" compact finite-difference schemes is their ability to handle multiple boundary types: periodic, free-slip (both symmetric and anti-symmetric), and no-slip along each coordinate direction. For periodic and free-slip conditions, the standard interior stencils remain valid throughout the domain with only minor modifications near the edges. These adjustments are introduced via ghost nodes $(f_0, f_{-1}, f_{-2}, f_{N+1}, f_{N+2}, f_{N+3})$, whose values are inferred from the corresponding interior nodes $(f_1, f_2, f_3, f_{N-1}, f_{N-2}, f_{N-3})$, where N denotes the total number of nodes along the relevant coordinate direction.

For periodic boundaries, ghost values for f, f' and f'' are imposed as:

$$f_0 \to f_N, \qquad f_{-1} \to f_{N-1}, \qquad f'_0 \to f'_N, \qquad f''_0 \to f''_N. \tag{2.12}$$

In the case of free-slip boundaries, symmetric conditions use:

$$f_0 \to f_2, \qquad f_{-1} \to f_3, \qquad f'_0 \to f'_2, \qquad f''_0 \to f''_2, \tag{2.13}$$

while anti-symmetric conditions adopt:

$$f_0 \to -f_2, \qquad f_{-1} \to -f_3, \qquad f'_0 \to f'_2, \qquad f''_0 \to f''_2. \tag{2.14}$$

When prescribing no-slip or inflow/outflow boundaries, one-sided compact formulations are applied at the first two nodes near the wall, followed by a standard centred scheme in the interior. These relations are third-order at the boundary and fourth-order at the adjacent

point (Lele, 1992). At the boundary node:

$$f_1' + 2f_2' = \frac{1}{2\Delta x}(-5f_1 + 4f_2 + f_3), \tag{2.15}$$

$$f_1'' + 11f_2'' = \frac{1}{\Delta x^2}(13f_1 - 27f_2 + 15f_3 - f_4), \tag{2.16}$$

and at the neighbouring node:

$$\frac{1}{4}f_1' + f_2' + \frac{1}{4}f_3' = \frac{3}{4\Delta x}(-f_1 + f_3), \tag{2.17}$$

$$\frac{1}{10}f_1'' + f_2'' + \frac{1}{10}f_3'' = \frac{6}{5\Delta x^2}(f_1 - 2f_2 + f_3). \tag{2.18}$$

These formulations enforce the desired boundary behaviour while retaining the solver's high-order accuracy across the domain (Lele, 1992). Boundary conditions for the pressure field follow the same treatment as described in Section 2.1.3.

Immersed Boundary Method (alternating-direction reconstruction)

For simulations over complex terrain, `WInc3D` employs an IBM compatible with the high-order compact finite-difference operators described previously. The method follows the alternating-direction strategy of Gautier et al. (2014), in which the no-slip condition is enforced at immersed surfaces without applying explicit forcing terms throughout the solid region. Instead, the solution is reconstructed locally during each directional sweep of the compact derivative operators, ensuring $u_i = 0$ at the immersed interface while maintaining the smoothness required for the implicit stencils.

A binary mask field ε is defined to distinguish between fluid and solid cells, with $\varepsilon = 0$ in the fluid and $\varepsilon = 1$ in the solid. The immersed interface Γ corresponds to locations where $\varepsilon - \varepsilon^- = 1$, where ε^- denotes a one-cell retraction of the solid mask into the fluid. This retracted mask is used to identify the immediate vicinity of the interface for enforcing boundary conditions and for modifying the pressure equation as described below.

During the predictor stage of the fractional step method (see Section 2.1.3), the intermediate velocity u_i^* is set to zero only at interface points Γ, leaving interior solid points unchanged. This avoids imposing discontinuities deep inside the solid region, which would degrade the accuracy of high-order compact schemes. To evaluate derivatives near the interface, the missing stencil values inside the solid are reconstructed using a one-dimensional polynomial extension along the current sweep direction. A Lagrangian polynomial is fitted to a set of fluid reference points on each side of the interface, excluding the first fluid point adjacent to the interface to improve conditioning. The polynomial is constrained to satisfy the no-slip condition exactly at the interface location. The reconstructed profile is then used to populate the stencil points inside the solid region for that directional sweep. This process is repeated sequentially in each coordinate direction (x, y, and z) during the corresponding scheme sweep. It is important to note that the no-slip condition is imposed only at the immersed interface Γ through this reconstruction procedure; no additional enforcement of no-slip is applied in subsequent stages of the algorithm.

Figure 2.2 illustrates the one-dimensional reconstruction for a fully immersed region. Red squares denote reconstructed values inside the solid, blue filled circles are the fluid reference points used for the polynomial fit, and black circles indicate the interface location.

The polynomial order is chosen to balance accuracy and stability. In the present work, three fluid reference points on each side are used when available, yielding a moderate-order fit that remains robust for curved or complex surfaces. In cases where the immersed body is very thin (thickness comparable to a few mesh spacings), the available reference points may be insufficient for the nominal polynomial order; in such cases, the order is locally reduced to maintain stability.

The pressure projection stage is adapted to account for the presence of solid regions by modifying the Poisson equation described in Section 2.1.3. The divergence correction is applied only in fluid cells, yielding:

$$\frac{\partial^2 p^{n+\frac{1}{2}}}{\partial x_j \partial x_j} = \frac{\rho}{c_n \, \Delta t} \frac{\partial\big[(1-\varepsilon^-)u_i^*\big]}{\partial x_i}. \tag{2.19}$$

When $\varepsilon^- = 0$, this equation reduces to the classical Poisson equation, recovering the standard

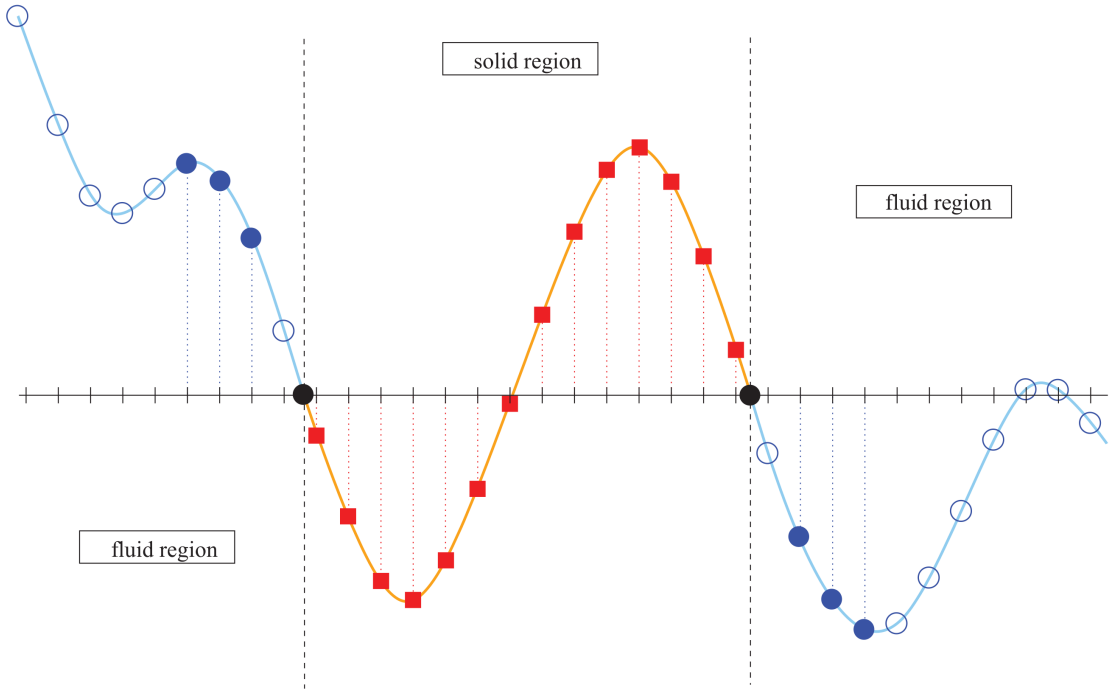


Figure 2.2: Schematic of the polynomial reconstruction in the alternating-direction IBM. Along each sweep direction, a one-dimensional Lagrangian polynomial is fitted to fluid reference points (filled blue circles) and constrained to satisfy $u_i = 0$ at the interface (black circles), providing smooth extended values (red squares) for the compact derivative stencil inside the solid region. Figure reproduced from Gautier et al. (2014), with permission from Informa UK Limited (Taylor & Francis Group).

incompressibility constraint on u_i^{n+1}. When $\varepsilon^- = 1$, it reduces to a Laplace equation that has no physical meaning inside the solid but remains consistent with the pressure field in the neighbouring fluid region. This formulation ensures good mass conservation at immersed boundaries while maintaining the regularity of the solution across the computational domain. Following Gautier et al. (2014), the retracted mask ε^- is used in the Poisson equation as it was found to provide improved mass conservation in the vicinity of immersed boundaries. No explicit boundary condition or additional treatment is applied in the pressure correction to enforce impermeability; instead, impermeability is retained implicitly through the kinematic enforcement at the interface, together with the exclusion of divergence contributions inside the solid.

Parallelisation Strategy

The compact finite-difference schemes in `WInc3D` define derivatives implicitly, and their evaluation involves solving systems of equations that can be written as:

$$\mathbf{A_1}\,\mathbf{f}' = \mathbf{B_1}\,\mathbf{f}, \tag{2.20}$$

$$\mathbf{A_2}\,\mathbf{f}'' = \mathbf{B_2}\,\mathbf{f}, \tag{2.21}$$

where $\mathbf{f}$, $\mathbf{f}'$, and $\mathbf{f}''$ are the vectors of nodal values and their first and second derivatives, respectively. The matrices $\mathbf{A_1}$ and $\mathbf{A_2}$ are tridiagonal and contain coefficients from the implicit left-hand side of the scheme (α) together with contributions from the boundary conditions. $\mathbf{B_1}$ and $\mathbf{B_2}$ are banded matrices that include the explicit right-hand side coefficients (a, b, c, d) scaled by Δx or Δx^2, also incorporating the relevant boundary-condition information. Since these matrices remain constant throughout the simulation, their LU decompositions are computed once during initialisation and reused at every time step, significantly reducing computational cost.

Large-scale simulations are parallelised using a two-dimensional pencil domain decomposition based on the `2DECOMP&FFT` library (Laizet and Li, 2011; Bartholomew et al., 2020). The com-

putational domain is divided into subdomains ("pencils"), each assigned to a separate message passing interface (MPI) process, as illustrated in Figure 2.3. When computing derivatives along a given direction (for example x), the other two coordinates (y and z in this case) are decomposed into a grid of $P_{\text{row}} \times P_{\text{col}}$ pencils, ensuring an even distribution of the workload across processes. Spatial operations, such as implicit derivative evaluations, interpolations, and filtering, are performed sequentially along one coordinate direction at a time. After completing the sweeps in one direction, the data are transposed globally to reorient the pencils for the next direction, following the sequence $X \to Y \to Z \to Y \to X$. These transpositions are implemented using `MPI_ALLTOALL` collective communications, which are handled internally by `2DECOMP&FFT` in a highly scalable manner.

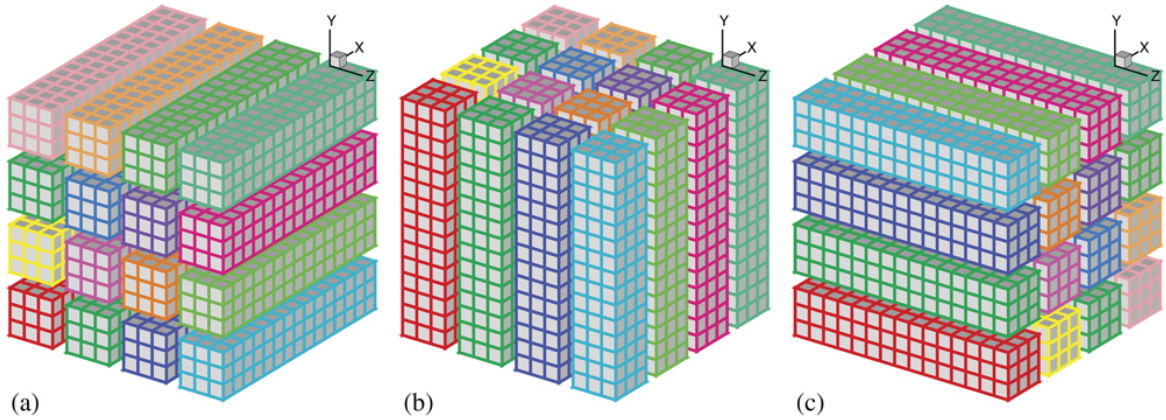


Figure 2.3: Two-dimensional pencil decomposition for $P_{\text{row}} = 4$ and $P_{\text{col}} = 4$, illustrating the three pencil orientations: (a) X-pencil, (b) Y-pencil, and (c) Z-pencil. Each colour represents a distinct MPI process. Figure reproduced from Laizet and Li (2011), with permission from John Wiley & Sons.

The same parallel decomposition is applied to the pressure Poisson solver. In this case, FFTs or discrete sine/cosine transforms are computed along one coordinate direction at a time. After transformation, the problem reduces to a set of independent one-dimensional Helmholtz equations. The tridiagonal systems resulting from these equations are solved locally within each pencil. Once each direction has been processed, the inverse transforms reconstruct the pressure field in physical space, followed by a global transpose to reorient the pencils for the next set of transforms or physical-space operations. This approach allows the solver to exploit the mathematical diagonalisation afforded by spectral methods while maintaining full scalability across distributed-memory architectures.

The scalability of this approach has been demonstrated extensively, with near-ideal strong and weak scaling achieved on up to $\mathcal{O}(10^5)$ cores (Laizet and Li, 2011; Deskos et al., 2020). This parallelisation framework enables `WInc3D` to handle domains containing billions of grid points, making it suitable for high-resolution LES of wind farms and atmospheric flows over complex terrain.

2.1.4 Atmospheric Boundary Layer Modelling

The ABL is modelled by solving the filtered incompressible N-S equations (2.1). In this work we focus on neutral ABL flows that are fully developed and statistically homogeneous in the streamwise and spanwise directions. Periodic boundary conditions are therefore imposed in x and z on the lateral boundaries of the horizontal plane. The top boundary is treated with a free-slip condition to represent the top of the boundary layer. The bottom boundary requires special treatment due to the high Reynolds numbers involved: resolving the near-wall region explicitly would require prohibitively fine grids to capture the viscous sublayer. Instead, a no-slip wall condition is combined with a wall-stress model following the methodology of Jaegle et al. (2010). The no-slip condition avoids interference with the IBM, which also enforces no-slip at terrain boundaries, whilst the wall model provides the correct surface stress without resolving the viscous layer.

A constant body-force term is applied in the streamwise momentum equation to maintain the desired mean pressure gradient and drive the flow. This forcing is set to:

$$F_x = \frac{u_*^2}{\delta},$$

where u_* is the friction velocity and δ is the ABL height. This choice follows from the momentum balance in a steady, horizontally homogeneous boundary layer, in which the total streamwise shear stress (comprising resolved and subgrid-scale contributions) varies linearly with wall-normal coordinate y:

$$\tau_{xy}(y) = \rho u_*^2 \left(1 - \frac{y}{\delta}\right). \tag{2.22}$$

This choice of F_x ensures that the vertical integral of the body force exactly balances the wall shear stress, allowing the simulated ABL to develop toward the logarithmic law-of-the-wall profile.

Log-law mismatch

A well-known challenge in LES of high-Reynolds-number ABL flows is the so-called log-law mismatch. This refers to a spurious overshoot of the velocity profile relative to the theoretical logarithmic law of the wall, typically observed in the near-surface region. The main cause is the excessive dissipation introduced by the subgrid-scale model, especially when a relatively large Smagorinsky coefficient is used to ensure numerical stability (Calaf et al., 2010; Meyers and Meneveau, 2011; Goit and Meyers, 2015).

To mitigate this issue, the present work adopts the wall-damping formulation of Mason and Thomson (1992), which reduces the Smagorinsky constant in the vicinity of the wall, thereby lowering SGS dissipation where it would otherwise disrupt the logarithmic layer. The damping is applied through a spatially varying Smagorinsky coefficient:

$$C_S = \left(C_0^p + \left\{ \kappa \left(\frac{y}{\Delta} + \frac{y_0}{\Delta} \right) \right\}^{-p} \right)^{-1/p}, \tag{2.23}$$

where Δ is the geometrically averaged grid size, y_0 is the roughness length, $\kappa = 0.4$ is the von Kármán constant, $C_0 = 0.14$ is the asymptotic Smagorinsky constant far from the wall, and $p = 3$ controls the rate of variation from the surface to the outer layer. This approach has been extensively used in ABL simulations (Calaf et al., 2010; Meyers and Meneveau, 2011; Goit and Meyers, 2015) and was validated in this work against reference neutral ABL profiles (see Section 3.1.1). The results confirmed that, although the damping improves the near-wall behaviour, the use of a no-slip on the bottom boundary in combination with the rough-wall stress model plays a more significant role in recovering the correct velocity profile.

Wall Modelling in the ABL

In standard flat-terrain LES configurations, the lower boundary is often treated with a free-slip condition combined with a wall-stress model based on horizontally averaged velocities or spatial filtering (Bou-Zeid et al., 2005). However, such approaches are not directly compatible with IBM for complex terrain, where the local wall orientation and roughness distribution must be accounted for without assuming horizontal homogeneity.

To ensure methodological consistency between the precursor ABL simulations and the subsequent complex-terrain simulations, a no-slip condition is used at the bottom wall together with a locally evaluated wall-stress model. This combination follows the methodology of Jaegle et al. (2010), in which the wall stresses are imposed at the first off-wall grid node while retaining the no-slip condition at the boundary.

The wall stress is computed using a rough-wall logarithmic formulation:

$$\tau_w = -\left[\frac{\kappa U_1}{\ln\left(y_1/y_0\right)}\right]^2, \tag{2.24}$$

where y_1 is the height of the first off-wall grid point, $U_1 = \sqrt{u_x^2 + u_z^2}$ is the local instantaneous velocity magnitude at y_1. No spatial averaging or filtering is applied in the velocity sampling, ensuring that the wall model remains applicable in non-homogeneous flow regions such as those encountered in complex terrain.

This local formulation allows unified treatment of wall stresses in both flat-terrain precursor simulations and complex-terrain successor simulations.

Shifted Periodic Boundary Condition

To prevent persistent alignment of large turbulent structures that can arise in short precursor domains, the shifted periodic boundary condition (sPBC) of Munters et al. (2016) is employed. In this approach, the flow exiting the domain is reintroduced at the inlet after a prescribed spanwise displacement, here set to $L_z/8$, where L_z is the spanwise domain length. This choice

was found to be effective, as velocity statistics showed no indication of structural locking.

For the wind farm optimisation studies in this thesis, avoiding such structural locking is essential: without it, regions of persistently higher mean velocity could form in the inflow, leading the optimiser to position turbines preferentially in these artificially favourable locations. The use of the sPBC ensured a more homogeneous statistical distribution of the incoming flow, removing this bias.

In practice, the sPBC also accelerated the convergence of ABL statistics, likely due to the more effective mixing of large-scale turbulent motions introduced by the spanwise shift. This reduction in spin-up time was advantageous, as in many cases the successor (complex terrain or wind farm) simulations reached statistical convergence faster than the precursor, reducing the overall computational cost of the optimisation workflow. The absence of statistical inhomogeneities was confirmed by comparing time-averaged velocity profiles across the spanwise direction.

2.1.5 Complex Terrain Simulations

The ABL simulations described in Section 2.1.4 provide the inflow conditions for the complex-terrain cases considered in this work. Using a precursor simulation ensures that the inflow reproduces the correct energy distribution and large-scale coherent structures of the atmospheric boundary layer, so that flow features in the successor simulations can be confidently attributed to the interaction between incoming turbulence and the topography. Synthetic turbulence methods can match certain statistical properties but often lack phase coherence and realistic large-scale organisation, leading to unphysical wake development in terrain-induced flows (Munters et al., 2016). Precursor inflows therefore provide a physically consistent and more reliable basis for these simulations.

Inflow–Outflow Conditions

Turbulence-rich inflow is generated via a precursor–successor approach. Although this method requires an additional high-fidelity simulation, in the present study it is advantageous because

the same inflow is reused across hundreds of terrain cases, ensuring consistent upstream conditions.

After statistical convergence is achieved in the precursor simulation, two-dimensional velocity planes at the outflow are stored at each time step. These planes are subsequently imposed at the inlet of the complex-terrain domain in a time-synchronised manner. The downstream boundary in the successor domain employs a one-dimensional convective outflow condition, with the convective velocity computed as the time- and grid-averaged mean of the maximum and minimum streamwise velocities at the last interior grid point.

Wall Model over Terrain

The wall-stress formulation described in Section 2.1.4 is retained for the complex-terrain simulations with only minor modifications. Our implementation follows the general philosophy of Liu and Stevens (2020b), whose approach inspired two aspects of the present work. As in their method, the stresses are imposed at the first off-wall grid point following the terrain surface, corresponding to all nodes inside a band of thickness Δy above the terrain. Although this location is not exactly at the same wall–normal distance as in the flat-terrain case, it is treated as representative of the local wall shear stress.

Liu and Stevens (2020b) simplified earlier immersed boundary wall-stress formulations (e.g. Chester et al. (2007)) by avoiding the construction of an interior velocity field within the solid region. Instead, all stress components are set to zero inside the immersed boundary region, while the dominant components are imposed at the band nodes. They demonstrated that this treatment preserves the vertical stress balance in the turbulent boundary layer (TBL) and provides good agreement with wind–tunnel data for smooth topography. More importantly, this simplification also reduces memory and communication overhead, enabling simulations at higher grid resolutions on modern supercomputers.

In their method, velocities are sampled along the local wall-normal direction, and the resulting stresses are rotated into the global coordinate frame before being applied. In `WInc3D`, this

strategy was implemented and compared against a further simplification in which velocities are sampled vertically, and stresses are applied directly in the global coordinate frame using the same formulation as for flat terrain. This treatment effectively approximates the terrain as a series of horizontal steps at grid resolution (staircase approximation), avoiding wall-normal reconstructions. No appreciable differences were observed between the two approaches, supporting the conclusion that the additional complexity of wall-normal sampling and stress rotation offers no advantage in this context.

Adopting the purely vertical sampling approach integrates naturally with the domain-decomposed structure of `WInc3D`, reducing inter-processor communication and improving computational performance. Validation against wind-tunnel experiments (see Chapter 3) confirms that this simplified strategy maintains accuracy for smooth, moderately sloped terrains whilst better aligning with the efficiency goals of large-scale high-resolution simulations.

2.1.6 Turbine Modelling

The ADM employed in this study treats the wind turbine rotor as a permeable disc that extracts momentum from the flow. This approach has its theoretical foundation in classical momentum theory, first developed by Rankine (1865) and Froude (1889) and subsequently refined by Betz (1920). The one-dimensional theory assumes that the cross-sectional area of the stream-tube contracts upstream of the rotor and expands downstream due to mass conservation, whilst the pressure drop across the disc accounts for the momentum extraction.

Under the assumptions of steady, incompressible, inviscid flow and uniform loading across the disc, the theory establishes a relationship between the upstream undisturbed velocity U_0 and the velocity at the rotor plane u_d. The axial induction factor, α_I, is defined as:

$$\alpha_I = 1 - \frac{u_d}{U_0}, \tag{2.25}$$

where u_d represents the axial velocity at the rotor location. From momentum conservation

principles, the dimensionless power coefficient C_P can be expressed as:

$$C_P = 4\alpha_I(1 - \alpha_I)^2. \tag{2.26}$$

Differentiation of this expression with respect to α_I yields the theoretical maximum power coefficient, known as the Betz limit: $C_{P,\max} = 16/27 \approx 0.593$, occurring at $\alpha_I = 1/3$. This represents the theoretical upper bound for kinetic energy extraction from a moving fluid stream, providing a fundamental benchmark for turbine performance assessment. A detailed derivation of these expressions can be found in standard references such as Burton et al. (2011).

Within this theoretical framework, the thrust coefficient relates to the axial induction factor through $C_T = 4\alpha_I(1 - \alpha_I)$, whilst the rotor velocity is equivalently given by Equation (2.25).

The implementation of these theoretical concepts within computational fluid dynamics frameworks was introduced later, first by Sørensen and Myken (1992), who developed the ADM for numerical simulations. This pioneering work established the foundation for incorporating wind turbine effects in CFD codes, enabling the study of wind turbine wakes and their interactions within numerical simulations.

Actuator Disc Implementation

In the present work, wind turbines are represented using a non-rotating actuator disc model. The ADM constitutes a reduced-order parametrisation that models the rotor as a permeable disc exerting a body force on the flow within its swept area. This approach circumvents the computational expense associated with resolving individual blade geometries and their rotation whilst preserving the correct mean loading characteristics. The method is computationally less expensive than the actuator line or the full rotor geometry simulation and has been widely used in the literature for utility-scale wind farms immersed in an atmospheric boundary layer (Calaf et al., 2010; Wu and Porté-Agel, 2011; Wu and Porté-Agel, 2013; Yang et al., 2014; Stevens et al., 2018).

The implementation follows the classical momentum theory formulation, wherein the thrust force exerted by the rotor on the flow is expressed as:

$$F_{\mathrm{wt}} = -\frac{1}{2}\rho A_D C_T U_0^2, \tag{2.27}$$

where A_D represents the rotor-swept area, and C_T is the thrust coefficient. Since the upstream undisturbed velocity U_0 is not directly accessible within the computational domain, the formulation is recast in terms of the velocity at the disc plane u_d using the local thrust coefficient C_T' as defined by Calaf et al. (2010):

$$F_{\mathrm{wt}} = -\frac{1}{2}\rho A_D C_T' u_d^2, \tag{2.28}$$

with

$$C_T' = \frac{C_T}{(1-\alpha_I)^2}. \tag{2.29}$$

This definition allows prescribing C_T and α_I based on turbine operating conditions while applying forces using the locally available u_d.

In the present implementation, u_d is computed as the spatially averaged velocity across the rotor-swept area and temporally filtered using a one-sided exponential filter with a relaxation time $T_{\mathrm{relax}} = 0.27\,\delta/u_*$ following Calaf et al. (2010). The instantaneous turbine power output P_{wt} is subsequently evaluated as

$$P_{\mathrm{wt}} = -F_{\mathrm{wt}} u_d. \tag{2.30}$$

In addition to the constant-C_T formulation, the model was extended to incorporate tabulated vendor curves for thrust and power as functions of the free-stream velocity U_0. Since U_0 is not directly accessible in the LES, it is reconstructed from the disc-averaged filtered velocity u_d using the induction relation in Equation (2.25), with α_I obtained from the tabulated $C_T(U_0)$. At each time step, an iterative update of U_0 is performed until a consistent solution for U_0, C_T, and α_I is achieved. The thrust force is then evaluated using Equation (2.27), with $C_T(U_0)$ interpolated from the table, thereby ensuring consistency with vendor data. When available, turbine

power is obtained directly from the tabulated $P(U_0)$ curve by linear interpolation; otherwise, power is computed using Equation (2.30). This extension ensures that the ADM response reflects operational turbine behaviour, including the decrease in C_T near rated conditions, while retaining the efficiency of the disc-based parametrisation.

Force Distribution via Super-Gaussian Kernels

To apply the thrust force within the Cartesian computational mesh, the load is distributed over the rotor-swept volume using super-Gaussian distribution kernels as proposed by King et al. (2017). This approach facilitates the representation of yawed and/or tilted rotors without requiring grid alignment to the rotor plane. The distribution kernel is defined as:

$$\Gamma(\delta_n, \delta_r) = \exp\left[-\left(\frac{\delta_n}{t_D/2}\right)^{\alpha_n} - \left(\frac{\delta_r}{D/2}\right)^{\alpha_r}\right], \tag{2.31}$$

where δ_n and δ_r represent the distances from the turbine centre in the normal and radial direction of the disc plane, respectively, α_n and α_r are the exponents that control the sharpness of the Gaussian profiles across the normal and radial directions, t_D is the disc thickness and D is the rotor diameter. The distributed thrust force is normalised with the volumetric integral of $\Gamma(\delta_n, \delta_r)$. The exponents used are $\alpha_n = 2$ and $\alpha_r = 8$, validated against the experimental data of Chamorro and Porté-Agel (2010). The exponent α_n ensures a smooth distribution of forces along the disc thickness, while α_r ensures a sharp distribution across the disc edge (see also King et al. (2017)).

Advantages and Limitations

The ADM offers several distinct advantages for wind farm simulations despite its inherent simplifications. A key advantage is that the ADM provides reliable power prediction in the far field, typically beyond 2–3 rotor diameters downstream, where wake recovery and turbine interactions become critical for wind farm performance assessment. The model effectively captures wake meandering phenomena, which are essential for accurately representing the complex wake-

to-wake interactions that govern the collective behaviour of turbines within large wind farms operating in atmospheric boundary layer flows. Furthermore, the computational efficiency of the ADM enables the comprehensive optimisation investigations pursued in this work.

However, it is acknowledged that the ADM has limitations, as it assumes uniform flow through the disc, which does not represent real-world conditions. The model does not account for three-dimensional effects caused by the rotation of the blades or the radial distribution of forces, which can impact the near-wake properties, typically within 1–2 rotor diameters downstream. Additionally, the current implementation does not model the effects of the nacelle and tower, which can influence wake asymmetry and turbulence characteristics, particularly in the near-wake region. Whilst more advanced methods such as the actuator line or actuator surface models may provide a more accurate representation of the wind turbine wake, they are computationally expensive and not feasible within this study's current spatial and temporal resolution requirements for optimisation applications (Li et al., 2022). For the purposes of this investigation, the ADM was deemed sufficient to assess the main trade-offs in wind farm optimisation whilst maintaining computational tractability.

2.2 Bayesian Optimisation Framework

BO was employed in this thesis to optimise wind-farm layouts. This section details the implementation choices adopted: the initial Design of Experiments (DoE), the GP surrogate model and its hyperparameters, the acquisition functions and batch strategy, and the constraint-handling approach.

The BO workflow follows a systematic cycle. For intuition, this process can be illustrated by a toy one-dimensional optimisation problem (Figure 2.4), where the design variable x represents, for example, the position of a turbine and the objective function y its associated power output. Initially, only a few observations of the function are available, corresponding to trial turbine placements selected through a DoE to provide broad coverage of the design space (a). These points are then evaluated, and a GP surrogate model is fitted to the data, yielding a

predictive mean and uncertainty band (b). From this surrogate, an acquisition function is constructed to indicate where the next evaluation should occur, balancing promising predicted performance with regions of high uncertainty (c). The function is evaluated at this new location, the surrogate is updated, and the process repeats (d). In practice, the same iterative loop is applied to high-dimensional problems such as wind-farm layouts, where the design variables span turbine positions, hub heights, and control parameters, and the objective is the total farm power predicted by high-fidelity LES. To leverage parallel computing resources, multiple candidate layouts can be selected simultaneously through batch acquisition strategies, and the loop continues until convergence criteria are met or the computational budget is exhausted.

In the context of this thesis, BO is applied to both wind farm layouts and passive control parameters. The objective function $f(\mathbf{x})$ represents the total farm power production, where $\mathbf{x}$ denotes the design variables: turbine positions, hub heights, tilt angles.

The framework is implemented in Python using `GPyOpt` (The GPyOpt authors, 2016), a library built on top of `GPy` that provides GP modelling, hyperparameter optimisation, acquisition-function maximisation, and constraint handling within a unified workflow.

The following subsections detail the key components of this implementation: the initial sampling strategy, the GP surrogate model, the acquisition functions and batch selection methods, and the constraint-handling approach.

2.2.1 Design of Experiments

BO begins with the selection of the initial points to be evaluated. These points must provide broad coverage of the design space whilst remaining computationally affordable, since each evaluation corresponds to a high-fidelity LES. Among the various DoE strategies, Latin Hypercube Sampling (LHS) (McKay et al., 1979) is widely used for high-dimensional problems due to its favourable space-filling property and low sampling cost.

Mathematically, for a design space $\mathcal{X} = [0, 1]^d$, where d is the number of design variables (which can be linearly scaled to any bounded region), the LHS procedure generates n sample vectors

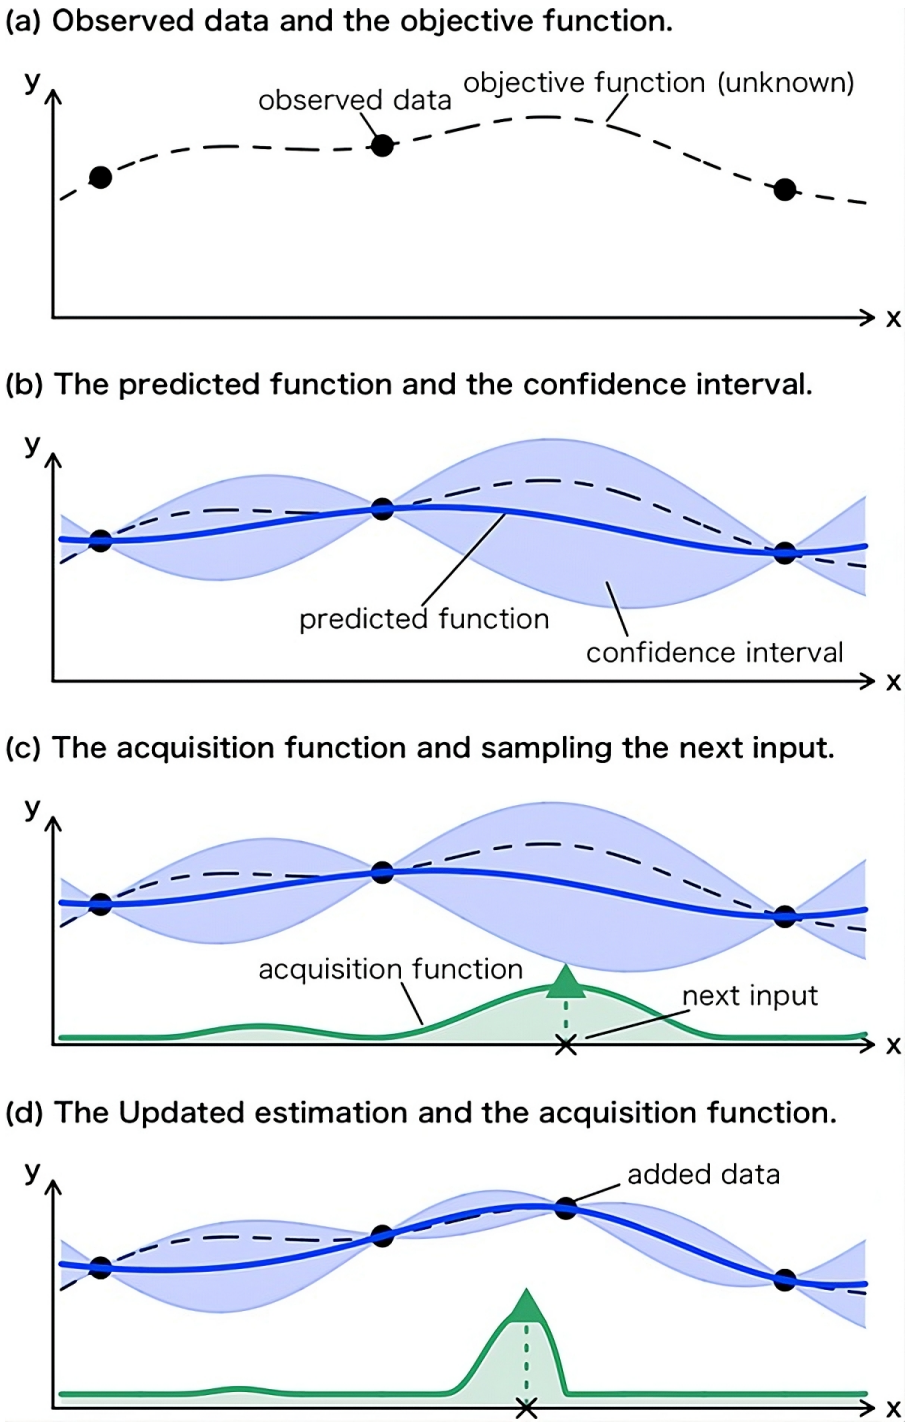


Figure 2.4: An example of BO on a toy one-dimensional problem. Figure reproduced from Tanaka and Iwata (2018), with permission from Springer Nature.

$\mathbf{x}_i \in \mathbb{R}^d$ $(i = 1, \ldots, n)$ as:

$$x_{ij} = \frac{\pi_j(i) - U_{ij}}{n}, \qquad i = 1, \ldots, n,\ j = 1, \ldots, d, \tag{2.32}$$

where x_{ij} is the j-th component of the i-th sample vector $\mathbf{x}_i = (x_{i1}, \ldots, x_{id})$, π_j is a random permutation of $\{1, 2, \ldots, n\}$ for dimension j, and $U_{ij} \sim \mathcal{U}(0, 1)$ are independent uniform random variables. This construction guarantees that, when projected onto any single dimension, the samples are uniformly distributed with exactly one point in each interval $[(i-1)/n, i/n]$.

The space-filling property of LHS is particularly important for GP surrogate training, since it ensures that the initial model captures the global behaviour of the objective function rather than concentrating samples in local regions. In this thesis, the number of initial evaluations, n, was selected as a compromise between the need for sufficient exploration of the design space and the high computational cost of LES. A rule of thumb is to initialise the surrogate with between one and ten times the number of design variables. Diessner et al. (2022) explicitly compared cases with one, five, and ten initial points per input dimension, showing that reducing the number of initial samples provides the BO algorithm with less information during the early iterations, but that the sequential updates quickly compensate for this lack of information. Conversely, increasing the number of initial samples beyond a moderate level does not necessarily yield improved performance.

In practice, $n = 20$ initial samples were used for the four-variable case in Chapter 4. For the thirty-two-variable case in Chapter 5, an initial design with $n = 32$ samples was tested, relying on the sequential BO process to build accuracy over time. However, this setup showed no convergence within the available computational budget. Consequently, the DoE was extended to incorporate domain knowledge (physics informed) of favourable turbine positions based on terrain-induced flow features. This approach retains the space-filling property of LHS whilst biasing samples towards promising regions, thereby accelerating convergence in high-dimensional spaces. The detailed methodology is deferred to Chapter 5, where the limitations of using LHS alone are discussed and the Physics-Informed BO is presented.

2.2.2 Gaussian Process Surrogate and Hyperparameter Strategy

The GP serves as the probabilistic surrogate model at the heart of the BO framework, providing both predictions and uncertainty quantification for unevaluated design points. A GP is a stochastic process that defines a distribution over functions, characterised entirely by its mean function $\mu(\mathbf{x})$ and covariance function (kernel) $k(\mathbf{x}, \mathbf{x}')$ (Rasmussen and Williams, 2006).

Given a set of training data $\mathcal{D}_n = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ of n samples, where $\mathbf{x}_i \in \mathbb{R}^d$ are design points and $y_i = f(\mathbf{x}_i) \in \mathbb{R}$ are the corresponding objective values from LES evaluations, the GP prior is specified as:

$$f(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')).$$

We adopt a zero-mean prior, $\mu(\mathbf{x}) = 0$, which is common when no strong prior information about the mean is available.

For prediction at a new test point $\mathbf{x}_*$, the posterior distribution remains Gaussian with mean and variance:

$$\mu(\mathbf{x}_*) = \mathbf{k}_*^\top (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y}, \tag{2.33a}$$

$$\sigma^2(\mathbf{x}_*) = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\top (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{k}_*, \tag{2.33b}$$

where $\mathbf{k}_* = [k(\mathbf{x}_*, \mathbf{x}_1), \ldots, k(\mathbf{x}_*, \mathbf{x}_n)]^\top$ is the covariance vector between the test and training points, $\mathbf{K} \in \mathbb{R}^{n \times n}$ has entries $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$, $\mathbf{y} = [y_1, \ldots, y_n]^\top$ collects the observed values, $\mathbf{I}$ is the $n \times n$ identity matrix, and σ_n^2 denotes the observational noise variance.

Given the deterministic nature of LES evaluations and the statistical convergence of time-averaged quantities such as turbine power, we employ a noise-free GP formulation ($\sigma_n^2 = 0$). This is justified by the controlled numerical environment and the negligible variation between repeated simulations with identical parameters. Modelling approximations (e.g., actuator-disc representation, omission of tower/nacelle effects) introduce systematic rather than random discrepancies and are therefore not treated as observation noise.

Matérn Covariance Function

The choice of covariance function strongly influences the GP's ability to represent the objective function. We employ the Matérn-5/2 kernel, which has proven effective for modelling physical systems and in wind-farm optimisation (Diessner et al., 2022; Bempedelis et al., 2024). The Matérn family offers a parameterised range of smoothness levels; the 5/2 variant yields sample paths that are twice mean-square differentiable. This kernel is commonly adopted in fluid-dynamics optimisation problems as it provides sufficient flexibility to model moderately smooth but potentially complex response surfaces. It is defined as:

$$k_{\text{Matérn-5/2}}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \left(1 + \sqrt{5}\, r + \frac{5r^2}{3}\right) \exp\left(-\sqrt{5}\, r\right), \tag{2.34}$$

where σ_f^2 is the signal variance controlling overall amplitude, and r is a distance measure discussed below.

Hyperparameter Optimisation

The kernel hyperparameters are collected in:

$$\boldsymbol{\theta} = \begin{cases} \{\sigma_f^2, \ell_1, \ldots, \ell_d\}, & \text{(ARD: Automatic Relevance Determination)}, \\ \{\sigma_f^2, \ell\}, & \text{(isotropic)}. \end{cases}$$

They are estimated by maximising the log marginal likelihood:

$$\log p(\mathbf{y} \mid \mathbf{X}, \boldsymbol{\theta}) = -\frac{1}{2}\mathbf{y}^\top \mathbf{K}_{\boldsymbol{\theta}}^{-1}\mathbf{y} - \frac{1}{2}\log|\mathbf{K}_{\boldsymbol{\theta}}| - \frac{n}{2}\log(2\pi), \tag{2.35}$$

where the subscript on $\mathbf{K}_{\boldsymbol{\theta}}$ emphasises its dependence on the hyperparameters. The first term measures data fit, the second penalises model complexity, and the third is a normalisation constant.

We optimise Equation (2.35) using L-BFGS-B (Liu and Nocedal, 1989), which efficiently handles box constraints to enforce positivity of variances and lengthscales. The gradients are computed

analytically:

$$\frac{\partial}{\partial \theta_j} \log p(\mathbf{y} \mid \mathbf{X}, \boldsymbol{\theta}) = \frac{1}{2}\mathbf{y}^\top \mathbf{K}^{-1} \frac{\partial \mathbf{K}}{\partial \theta_j} \mathbf{K}^{-1}\mathbf{y} - \frac{1}{2} \operatorname{tr}\left(\mathbf{K}^{-1} \frac{\partial \mathbf{K}}{\partial \theta_j}\right). \tag{2.36}$$

Hyperparameters are re-optimised at every BO iteration as new data become available. Early iterations reflect global behaviour based on sparse coverage; later iterations, enriched with data from promising regions, can capture more local structure. The computational cost of this re-optimisation is negligible relative to LES simulations, making frequent updates both feasible and beneficial. To mitigate local optima in the marginal-likelihood surface, we perform multiple random restarts, drawing initial values from ranges informed by design-space scaling and expected variability.

Lengthscale Parameterisation Strategies

The distance r in the kernel incorporates lengthscale parameters that govern how rapidly correlation decays with separation in the design space. We employ two strategies, depending on dimensionality and data availability.

For low-dimensional problems with sufficient data (e.g., the four-variable optimisation in Chapter 4), we use Automatic Relevance Determination (ARD), which assigns an independent lengthscale ℓ_j to each dimension j:

$$r_{\text{ARD}}(\mathbf{x}, \mathbf{x}') = \sqrt{\sum_{j=1}^{d} \left(\frac{x_j - x'_j}{\ell_j}\right)^2}, \tag{2.37}$$

where x_j and x'_j denote the j-th components of design vectors $\mathbf{x}$ and $\mathbf{x}'$, respectively. This anisotropic formulation lets the GP learn the relative importance of each variable: large ℓ_j implies slow variation (low sensitivity) along dimension j, whereas small ℓ_j indicates rapid variation and high sensitivity. For problems with four design variables, ARD can distinguish sensitivities between different parameter types. With 40 initial samples and additional BO samples, the data density supports robust estimation of these $d + 1$ hyperparameters.

For high-dimensional problems (e.g., the thirty-two-variable optimisation in Chapter 5), we

adopt an isotropic kernel with a single shared lengthscale:

$$r_{\text{iso}}(\mathbf{x}, \mathbf{x}') = \frac{\|\mathbf{x} - \mathbf{x}'\|}{\ell}, \tag{2.38}$$

where ℓ is the common lengthscale and $\|\cdot\|_2$ denotes the Euclidean norm. This reduces the number of hyperparameters from $d + 1$ (ARD) to just two, making optimisation tractable with limited data and avoiding overfitting in high-dimensional spaces where sample density is necessarily sparse.

2.2.3 Acquisition Functions and Batch Strategy

The acquisition function is the decision-making component of BO, quantifying the utility of evaluating the objective function at any given point in the design space. In our context of expensive LES evaluations, careful choice of acquisition function is critical, as it guides sampling by combining the GP's mean and uncertainty to balance exploration and exploitation.

Lower Confidence Bound Acquisition Function

For this work we employ the Lower Confidence Bound (LCB) acquisition function, well-suited to deterministic maximisation problems such as wind-farm power optimisation. It is defined as:

$$\text{LCB}(\mathbf{x}) = \mu(\mathbf{x}) - w(\sigma(\mathbf{x}) + \epsilon), \qquad \epsilon = 0.01, \tag{2.39}$$

where $w \geq 0$ is an exploration weight, and $\epsilon = 0.01$ is a small jitter term that prevents the acquisition surface from becoming flat where $\sigma(\mathbf{x}) \approx 0$.

The parameter w controls this trade-off: large values emphasise exploration (favouring uncertain regions), whilst small values emphasise exploitation (favouring high-mean regions). We use a

constant value in Chapter 4 ($w = 2.0$) and an adaptive two-stage schedule in Chapter 5:

$$w = \begin{cases} 2.0, & \text{exploration phase,} \\ 1.0, & \text{exploitation phase,} \end{cases}$$

encouraging broad search early and refinement later.

Because we seek to maximise the objective, points with large LCB values (i.e. large mean value $\mu(\mathbf{x})$ or large uncertainty $\sigma(\mathbf{x})$) are preferred. The next evaluation is selected by maximising the LCB over the design space:

$$\mathbf{x}_{n+1} = \arg\max_{\mathbf{x} \in \mathcal{X}} \text{LCB}(\mathbf{x} \mid \mathcal{D}_n). \tag{2.40}$$

The LCB is maximised using constrained optimisation to respect design variable bounds (e.g., turbine position limits, minimum spacing requirements).

Batch Acquisition Strategy

To reduce wall-clock time, multiple candidate layouts are selected at each iteration and evaluated in parallel. In sequential BO, the next candidate is denoted $\mathbf{x}_{n+1}$. In the batch setting, this generalises to a set of b points at iteration t:

$$\mathbf{B}_t = \{\mathbf{x}_1^{(t)}, \ldots, \mathbf{x}_b^{(t)}\}, \tag{2.41}$$

where b is the batch size.

We adopt the local-penalisation approach of González et al. (2016) to construct diverse batches whilst maintaining the exploration-exploitation balance of the LCB acquisition function:

1. First point:

$$\mathbf{x}_1^{(t)} = \arg\max_{\mathbf{x} \in \mathcal{X}} \text{LCB}(\mathbf{x} \mid \mathcal{D}_t).$$

2. Subsequent points $j = 2, \ldots, b$:

$$\mathbf{x}_j^{(t)} = \arg\max_{\mathbf{x} \in \mathcal{X}} \mathrm{LCB}(\mathbf{x} \mid \mathcal{D}_t) \prod_{i=1}^{j-1} \varphi(\mathbf{x}, \mathbf{x}_i^{(t)}),$$

where

$$\varphi(\mathbf{x}, \mathbf{x}_i) = 1 - \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}_i\|^2}{2\rho_{\mathrm{pen}}^2}\right) \tag{2.42}$$

penalises regions near previously chosen points. The radius ρ_{pen} is set adaptively from the GP lengthscales so that the exclusion zone adapts to the function's typical variation.

This local-penalisation strategy discourages clustering within batches, ensuring that evaluations probe distinct regions and maximise information gain. The combination of LCB acquisition, local-penalisation batch selection, and constrained optimisation provides a robust framework for wind-farm design that consistently converges to high-quality configurations within tight evaluation budgets, as demonstrated in the following chapters.

Chapter 3

Validation of the LES Framework

Before applying the computational framework to optimisation problems, it is essential to establish confidence in its predictive capabilities. This chapter therefore presents the validation of the LES framework introduced in Chapter 2, ensuring that the methodology reliably captures the key flow physics of terrain–wake interactions and turbine performance. The work presented here has been published in Jané-Ippel et al. (2023) and Jané-Ippel et al. (2024). Validation is structured progressively, beginning with controlled wind-tunnel benchmarks and culminating in comparisons against operational wind-farm data.

The first part of the chapter focuses on wind-tunnel experiments and other numerical studies, which provide high-quality reference data under controlled conditions. Three complementary steps are considered: (i) assessment of the ABL precursor inflow through the two-dimensional Gaussian hill benchmark, (ii) evaluation of terrain-induced flow features over both 2D and 3D Gaussian hills, and (iii) validation of turbine performance predictions in complex terrain and under yawed conditions. These cases test all core elements of the solver: precursor generation, the IBM with wall modelling, and the actuator-disc turbine model.

The second part of the chapter extends validation to an operational onshore wind farm comprising 12 turbines sited on real topography. LES under neutral atmospheric conditions are compared directly against SCADA power data for a single wind direction and for a 30° sector, assessing the predictive skill of the framework at utility scale. In addition to the power

comparison, flow-field diagnostics of velocity and turbulence intensity are examined to evaluate terrain-induced effects. Additional tests on domain size, inflow profiles, turbine thrust and power curves, and variable surface roughness are carried out as targeted robustness checks, rather than as an exhaustive sensitivity study, to identify the modelling choices most influential for farm-scale predictions. This part of the study was made possible through collaboration with SSE Renewables, who provided the operational SCADA dataset and contributed to the definition of the validation objectives.

The chapter is structured as follows. Section 3.1 presents the wind-tunnel benchmarks and idealised cases, and Section 3.2 covers the operational wind-farm case. Together, these validation stages demonstrate that the framework can accurately reproduce canonical flows, turbine siting effects, and utility-scale farm performance, providing a robust foundation for the optimisation studies in Chapters 4 and 5.

3.1 Idealised Benchmark

The validation benchmarks presented here are directly connected to the optimisation studies of this thesis. The two-dimensional Gaussian hill of Cao and Tamura (2006) is the terrain for the two-turbine optimisation in Chapter 4, while the three-dimensional Gaussian hill of Ishihara et al. (1999) underpins the generation of the synthetic complex terrain employed in Chapter 5. Validating the LES framework against these wind-tunnel cases ensures that the same numerical set-up used in subsequent optimisations has been rigorously tested against controlled experiments.

3.1.1 Precursor and 2D Hill Validation

The first validation case considers the canonical two-dimensional Gaussian hill investigated experimentally by Cao and Tamura (2006). This terrain is defined by a cosine-squared function:

$$z_s(x) = H \cos^2\left(\frac{\pi x}{2L}\right), \tag{3.1}$$

with maximum height $H = 0.04$ m and half-width $L = 0.1$ m, constant in the spanwise direction, hence referred to as a two-dimensional hill. In the wind tunnel experiments, the inflow turbulence was generated physically using flat boards covered with roughness elements placed on the tunnel floor upstream of the test section. This setup produced a fully developed TBL with a height of $\delta \approx 0.25$ m. The flow was characterised by a friction velocity $u_* = 0.1926$ m/s and a roughness length $y_0 = 0.004$ mm, derived from the experimental log-law fit. To replicate these conditions in the LES, a precursor simulation is employed to generate a physically consistent turbulent inflow. The target flow parameters are set directly to these experimentally reported values, ensuring the numerical boundary layer matches the physical one.

The computational domain has size $L_x \times L_y \times L_z = (5 \times 1 \times 1)\delta$, with periodic boundary conditions in the spanwise direction. A precursor TBL simulation is used to provide physically consistent inflow to the hill configuration. Statistics are collected for a total non-dimensional time of $\hat{T} = 20$ after discarding an initial transient of $\hat{T} = 2$, where $\hat{T} = u_* t/\delta$.

Two grids are considered. The coarser grid, hereafter G1, comprises $N_x \times N_y \times N_z = 193 \times 97 \times 64$ nodes, corresponding to 16 nodes across the hill height. The finer grid, G2, doubles the resolution in each direction ($385 \times 193 \times 128$ nodes), providing 32 nodes across the hill height. The respective non-dimensional time steps are $\Delta\hat{T} \approx 7.7 \times 10^{-5}$ for G1 and 3.85×10^{-5} for G2.

Figure 3.1 shows the computational domain, the hill geometry, and boundary conditions used in the 2D hill validation set-up. It also includes an actuator disc used in Section 3.1.3.

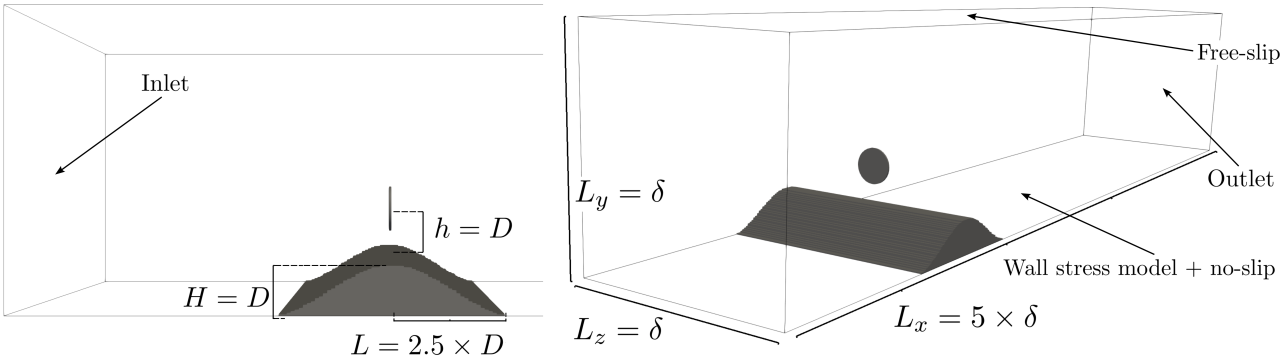


Figure 3.1: Computational domain representation of the validation set-up, featuring the constant section hill and an actuator disc at the hilltop, along with the boundary conditions.

Precursor validation

We perform precursor simulations of the TBL from the experiments of Cao and Tamura (2006) for both G1 and G2 resolutions, with domain size $L_x \times L_y \times L_z = (3 \times 1 \times 1)\delta$.

Figure 3.2a shows the horizontally- and time-averaged streamwise velocity profiles compared against the experimental data and the log-law profile:

$$\frac{\langle u_x \rangle}{u_*} = \frac{1}{\kappa} \ln\left(\frac{y}{y_0}\right), \tag{3.2}$$

with $\kappa = 0.4$. Due to the no-slip wall condition and the use of high-order schemes, the log-law profile is recovered after the first 3–5 nodes off the wall.

Figure 3.2b shows the resolved stresses $-\langle u'_x u'_y \rangle$ from both grids. G1 and G2 are in good agreement with the experimental data and theoretical stress profile down to the wall. As expected, G2 resolves near-wall turbulence slightly better, though the differences are marginal.

Hill-flow validation

Figure 3.3a shows the time-averaged streamwise velocity $\langle u_x \rangle$ colour map obtained with G2. Additionally, it displays the time-averaged streamwise velocity profiles from the experimental data and the two simulations with `WInc3D`, all normalised by the free stream velocity U_∞ at

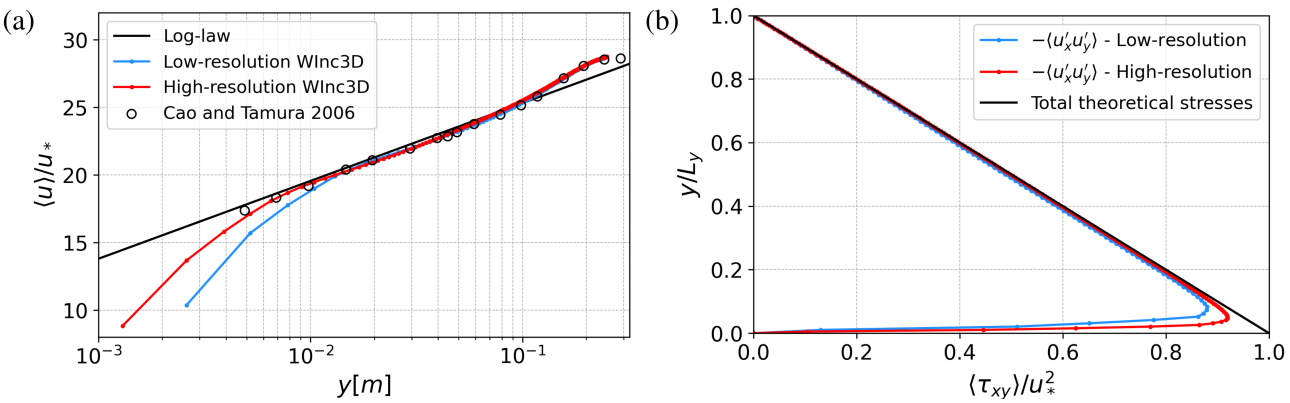


Figure 3.2: Horizontally- and time-averaged streamwise velocity (a) and resolved stress profiles (b) of the TBL in the precursor simulation compared against analytical profiles and experimental data of Cao and Tamura (2006).

$L_z = \delta/2$. The contour of zero velocity in the streamwise direction, highlighting the edge of reverse flow, is demarcated by the white line. Both grid resolutions align closely with the experimental data across all streamwise locations, accurately capturing the speed-up in front of the hill, separation post-crest, and reattachment at around $x/H = 5.54$. In the experimental campaign, the reattachment is reported to be around $x/H = 5.4$.

Figure 3.3b shows the time-averaged streamwise turbulence intensity σ_{u_x} colour map, normalised by U_∞, from G2. Profiles from both resolutions (G1 and G2), set against the experiments and normalised by $U_\infty/5$, are also shown. Unlike the streamwise velocity profiles, differences may be observed between the two grid resolutions. G2 generally concurs with the experimental data, whereas G1 demonstrates some discrepancies in the near wake of the hill. Nonetheless, predictions from both simulations realign downstream of $x/H \approx 5.0$.

To quantify errors, we compute

$$e_{u_x} = \frac{|u_{x,\mathrm{exp}} - u_{x,\mathrm{sim}}|}{U_\infty}, \tag{3.3a}$$

$$e_{\sigma_x} = \frac{|\sigma_{x,\mathrm{exp}} - \sigma_{x,\mathrm{sim}}|}{U_\infty}. \tag{3.3b}$$

Maximum discrepancies are of order 5–16% for velocity and 2–9% for turbulence intensity. The largest errors occur near separation and reattachment ($x/H = 0$ to $x/H = 2.5$), while regions upstream of the crest and farther downstream are typically within 5–8%. G2 shows

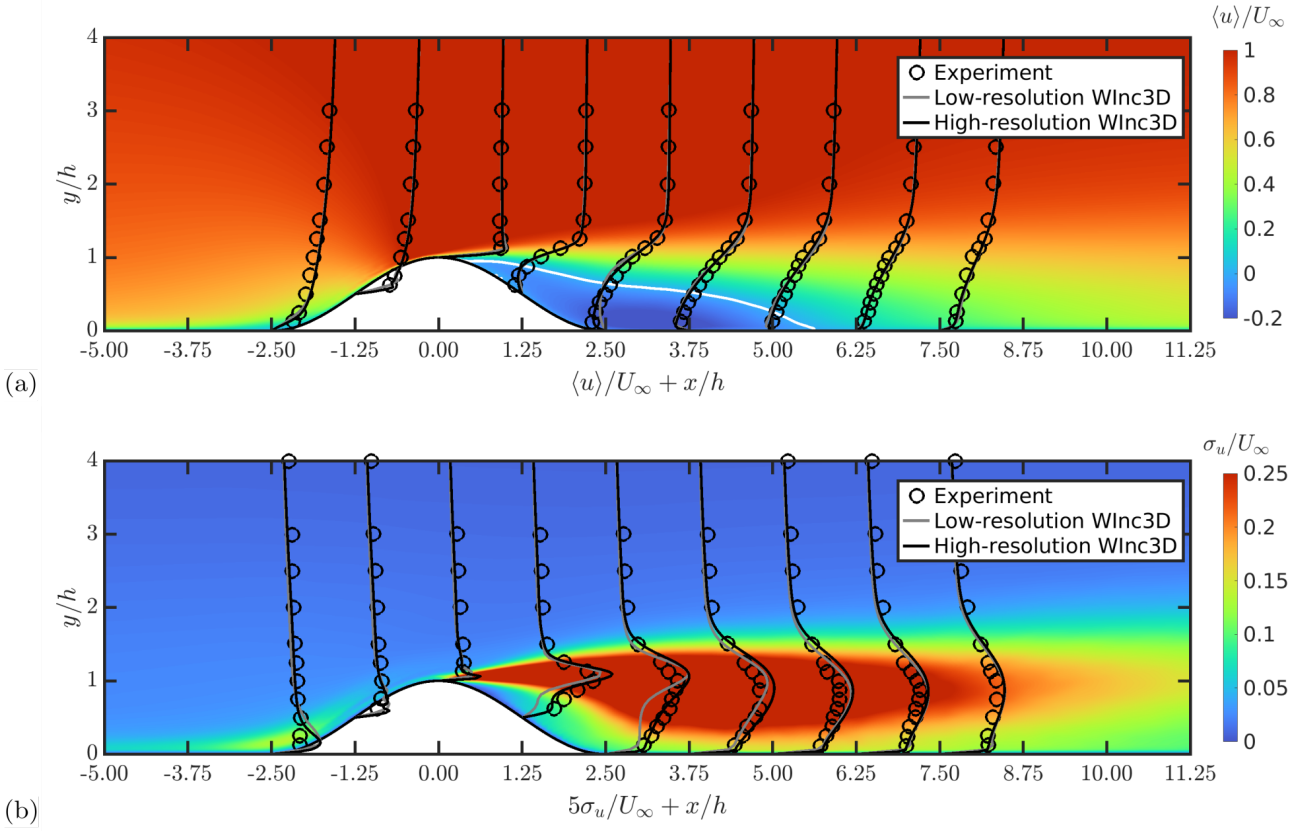


Figure 3.3: Comparison between `WInc3D` simulations of the 2D hill and the experimental data of Cao and Tamura (2006) for (a) normalised streamwise velocity and (b) turbulence intensity. The white line in (a) shows the edge of reverse flow. Vertical profiles are plotted with a horizontal offset corresponding to their streamwise location for visualisation.

systematically better agreement, particularly in capturing reattachment length and turbulence levels in the lee.

These results confirm that the LES framework reproduces the key features of the 2D hill benchmark within experimental uncertainty. The resolution study indicates that 16–32 nodes across the hill height are sufficient for accurate prediction, providing a practical baseline for the more complex 3D terrain considered in the following section. For Chapter 4, the finer G2 grid is adopted despite its higher computational cost, since turbine positions are optimised in close proximity to each other and to the recirculation region where higher resolution is required.

3.1.2 Extension to 3D Hill

The second benchmark considers the three-dimensional Gaussian hill investigated experimentally by Ishihara et al. (1999). The hill geometry is defined by the cosine-squared function:

$$z_s(x, y) = H \cos^2 \left(\frac{\pi (x^2 + y^2)^{1/2}}{2L} \right), \tag{3.4}$$

where $H = 40$ mm is the hill height and $L = 100$ mm is the half-width. This benchmark has been widely used in validation studies of immersed boundary and wall-modelling approaches (Diebold et al., 2013; Fang and Porté-Agel, 2016; Liu and Stevens, 2020b).

The inflow corresponds to a neutral TBL generated experimentally using a combination of 60 mm cubic arrays and smaller roughness elements (20 mm and 10 mm) followed by a smooth plywood floor. This configuration produced a boundary layer depth of $\delta = 0.36$ m. The inflow parameters derived from the experimental log-law fit are $u_* = 0.212$ m/s and $y_0 = 1.0 \times 10^{-5}$ m. The hill surface roughness is prescribed as $y_{0,\text{hill}} = 2y_0$, following established practice. The computational domain size is $L_x \times L_y \times L_z = (4 \times 1 \times 2)L_z$, with $L_z = 0.32$ m. Each case is run for a total non-dimensional time of $\hat{T} = 20$, discarding an initial transient of $\hat{T} = 2$.

As in the 2D hill validation, precursor simulations driven by the experimental TBL parameters were performed for each resolution to provide physically consistent inflow. To avoid repetition,

results are not shown here, since they lead to the same conclusions as in Section 3.1.1.

Two discretisations are assessed. The coarser grid, H1, employs $N_x \times N_y \times N_z = 385 \times 193 \times 192$ nodes, corresponding to 24 points across the hill height. The finer grid, H2, doubles the resolution in each direction ($769 \times 385 \times 384$ nodes), providing 48 points across the hill height. The respective non-dimensional time steps are $\Delta\hat{T} \approx 3.30 \times 10^{-5}$ (H1) and 1.65×10^{-5} (H2).

Figure 3.4 compares time-averaged streamwise velocity profiles at several streamwise locations along the mid-plane, as well as a representative spanwise cut. Both grids reproduce the main features of the hill flow: acceleration at the crest, lateral flow divergence, separation in the lee, and gradual wake recovery downstream. Differences between H1 and H2 are confined to the near-wake region ($x/H = 1.25$–3.75), where H2 better resolves the steep velocity gradients near the surface. Farther downstream and in the outer flow, both grids converge to nearly identical predictions.

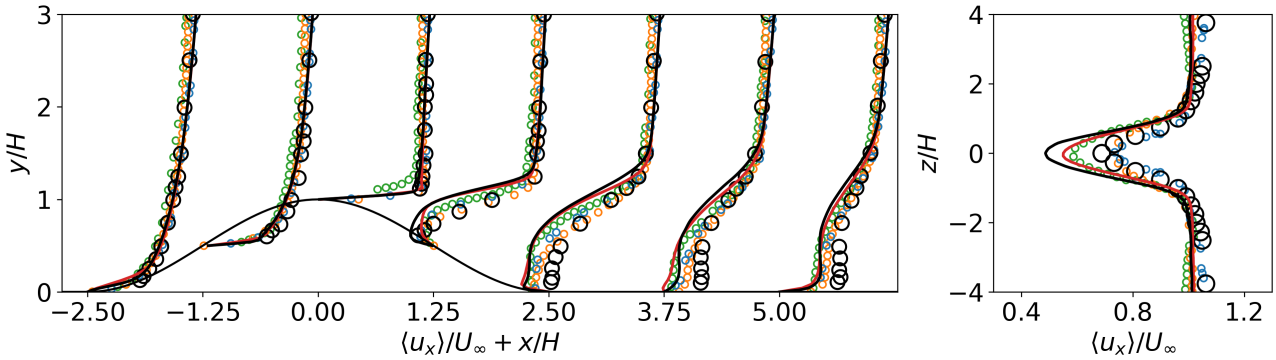


Figure 3.4: Time-averaged streamwise velocity profiles from simulations over the 3D hill compared with the experimental data of Ishihara et al. (1999). Vertical profiles are shown at the mid-plane for various streamwise locations (left), plotted with a horizontal offset corresponding to their streamwise position, and a spanwise profile at $x/H = 3.75$, $y/H = 1$ (right). `WInc3D` simulations (H1: red, H2: black) are compared with Diebold et al. (2013) (green circles), Fang and Porté-Agel (2016) (blue circles), and Liu and Stevens (2020b) (orange circles).

Figure 3.5 shows the corresponding streamwise turbulence intensity distributions. Both grids reproduce the turbulence amplification near the crest and the wake broadening downstream. H2 predicts slightly higher turbulence intensities at crest height in the near wake, closer to the experimental values, while H1 tends to under-predict in this region. Nevertheless, both grids capture the overall magnitude and shape of the turbulence intensity profiles, with differences diminishing in the far wake.

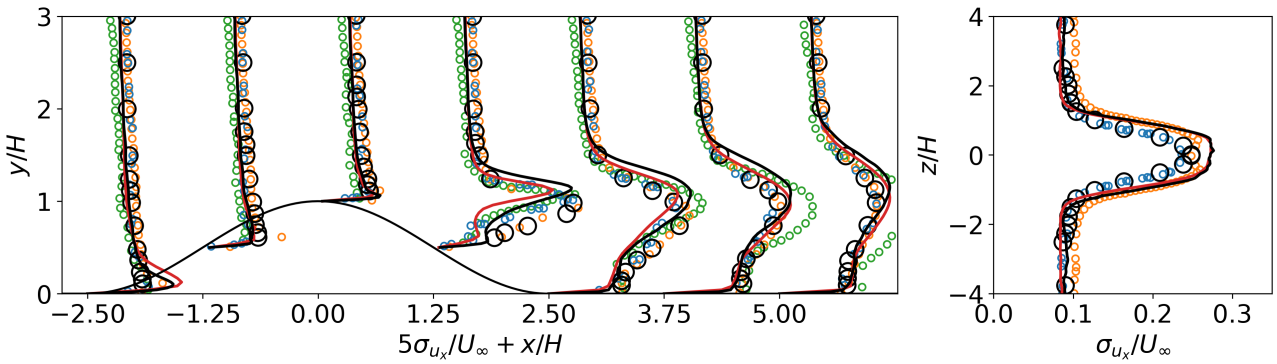


Figure 3.5: Streamwise turbulence intensity profiles from simulations over the 3D hill compared with the experimental data of Ishihara et al. (1999). Vertical profiles are shown at the mid-plane for various streamwise locations (left), plotted with a horizontal offset corresponding to their streamwise position, and a spanwise profile at $x/H = 3.75$, $y/H = 1$ (right). `WInc3D` simulations (H1: red, H2: black) are compared with Diebold et al. (2013) (green circles), Fang and Porté-Agel (2016) (blue circles), and Liu and Stevens (2020b) (orange circles).

Quantitatively, maximum velocity errors are around 12–15% in the separation region ($x/H = 1.25$–2.5), reducing to below 7% in the outer flow and far wake. Turbulence intensity errors remain within 10% across all locations, with H2 providing closer agreement in the near wake.

These results confirm that the LES framework reproduces the key flow features of the 3D hill benchmark within experimental uncertainty. The resolution study further establishes that 24–48 nodes across the hill height are sufficient for accurate predictions. In the optimisation studies of Chapter 5, we therefore adopt the H1 resolution (24 nodes across the hill height) as a practical compromise between accuracy and cost, with further details of the terrain scaling and resolution strategy provided in the corresponding chapter.

3.1.3 Turbine Performance around the 2D Hill

The power prediction capabilities of the ADM are examined using the 2D hill set-up, placing a single turbine with $C_T' = 4/3$ at different streamwise locations with hub height $h = D$. To characterise resolution, we define $n_D = D/\Delta y$ as the number of grid nodes across the rotor diameter in the vertical direction, which provides a convenient measure of how finely the disc is resolved. Figure 3.6 shows the normalised power output at several positions, obtained with the G1 ($n_D = 16$) and G2 ($n_D = 32$) grids, and compared against reference numerical studies (Liu

and Stevens, 2020a; Zhang et al., 2022). The two grid resolutions provide similar results, with only minor discrepancies of around 3.3% and 8.4% at $x/D = 1.25$ and $x/D = 5$, respectively. Overall, the agreement with the literature is good, confirming that the ADM reliably captures siting effects in complex terrain. The expected behaviour is reproduced: reduced performance upstream and in the lee, and significant enhancement at the crest due to flow acceleration.

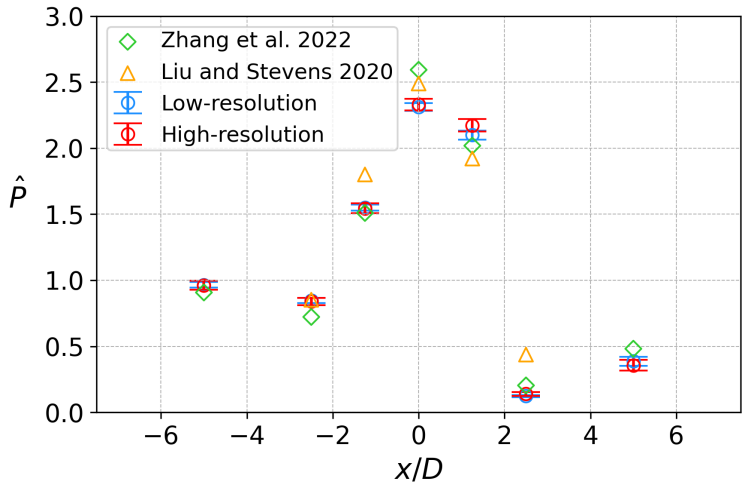


Figure 3.6: Normalised power output of a wind turbine placed at various streamwise positions relative to the 2D hill centred at $x/D = 0$, as computed by `WInc3D`, Liu and Stevens (2020a), and Zhang et al. (2022).

Wake recovery

To further assess the framework, wake recovery is analysed using the velocity deficit field and wake-centre trajectory (Shamsoddin and Porté-Agel, 2018; Zhang et al., 2022). The velocity deficit is defined as:

$$\langle u_{\mathrm{def}} \rangle = \langle u_{\mathrm{nt}} \rangle - \langle u_{\mathrm{t}} \rangle, \tag{3.5}$$

where $\langle u_{\mathrm{nt}} \rangle$ is the mean velocity field without a turbine and $\langle u_{\mathrm{t}} \rangle$ the mean field with a turbine included. The wake-centre trajectory is identified from the location of maximum deficit at hub height:

$$y_{\mathrm{wc}}(x) = \arg\max_{y} \langle u_{\mathrm{def}}(x, y, z_{\mathrm{t}}) \rangle. \tag{3.6}$$

Representative cases are shown in Figure 3.7, comparing the flat terrain reference with turbines

placed upstream of the hill ($x/D = -5$), at the hill crest ($x/D = 0$), and just behind the crest ($x/D = 1.25$). The hill strongly alters the trajectory and recovery of the turbine wake. For the case when the turbine is located upstream of the hill, the adverse blockage reduces turbine performance but promotes faster wake recovery, consistent with the acceleration induced on the windward slope. For the case when the turbine is located at the crest, the turbine produces maximum power due to the speed-up effect, but the wake recovery is slowed by the adverse pressure gradient in the lee of the hill. For the case when the turbine is located downstream at $x/D = 1.25$, the wake is initially hindered by the hill-induced recirculation, but enhanced turbulent mixing between the hill and turbine wakes accelerates recovery further downstream. In all cases, the turbine and hill wakes interact, with each modifying the development of the other.

A quantitative comparison of wake recovery along the trajectories of turbines at different hill positions, benchmarked against the flat terrain case, is shown in Figure 3.8. The velocity deficit is normalised by the disc-averaged velocity of the respective turbine, and the streamwise axis is expressed as $(x - x_t)/D$, where x_t denotes the turbine position, to allow comparison relative to each turbine location. The upstream turbine ($x/D = -5$) recovers faster than the flat terrain case, reaching a deficit of 0.4 within $4D$, compared to $6D$ for the reference. This is consistent with the favourable pressure gradient observed in front of the hill. By contrast, the turbine on the crest exhibits a stronger, longer-lasting wake due to the adverse pressure gradient in the lee, reaching a deficit of 0.4 only after $7D$. The near-lee case ($x/D = 1.25$), however, departs from this trend: although the adverse pressure gradient initially delays recovery in the near wake, enhanced turbulent mixing between the hill-induced flow and the turbine wake subsequently accelerates momentum diffusion, leading to faster recovery than the crest case (Yang et al., 2015). This behaviour is consistent with the downward deflection of the wake-centre trajectory and the elevated turbulence levels induced by the hill, both of which promote mixing and accelerate recovery.

These results demonstrate that the LES framework not only reproduces power variations due to hill siting, but also the key wake-recovery mechanisms reported in the literature (Liu and Stevens, 2020a; Zhang et al., 2022). This provides further confidence in the actuator-disc

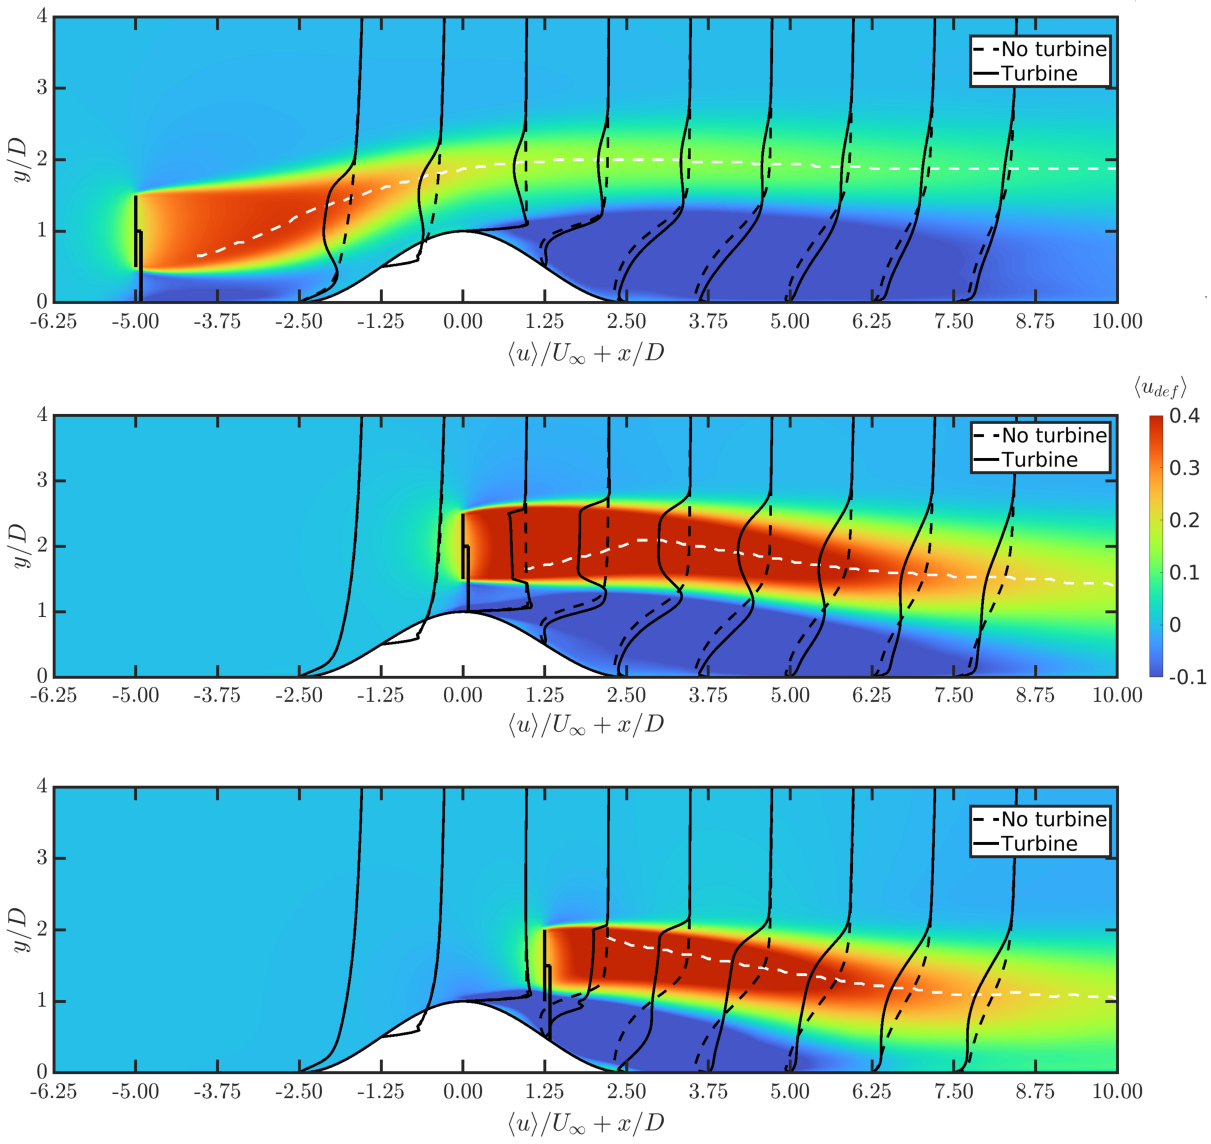


Figure 3.7: Velocity deficit fields and wake-centre trajectories (white dashed lines) for wind turbines placed at different streamwise positions relative to the 2D hill: upstream ($x/D = -5$), crest ($x/D = 0$), and near-lee ($x/D = 1.25$). Vertical profiles are plotted with a horizontal offset corresponding to their streamwise location for visualisation.

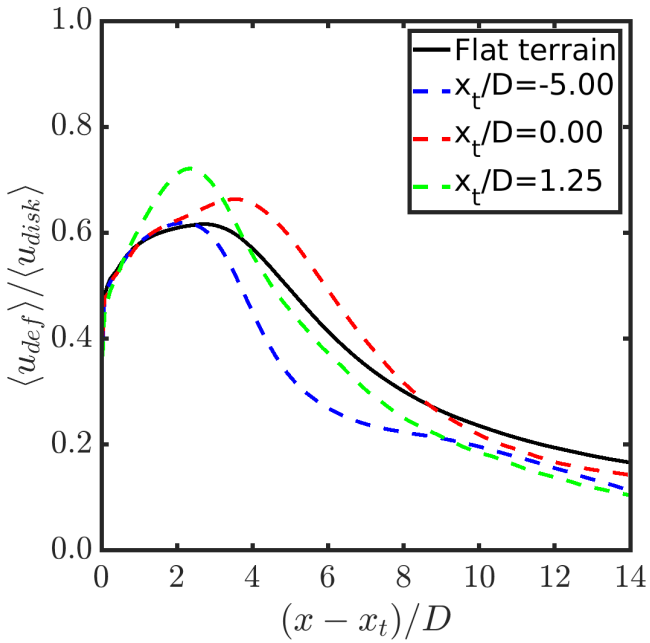


Figure 3.8: Normalised velocity deficit along the wake-centre trajectory for wind turbines placed at different streamwise positions relative to the 2D hill and in flat terrain.

implementation for subsequent optimisation studies.

3.1.4 Yawed Turbine Validation

Since tilted turbines are used in Chapter 4, an additional validation is performed for yawed inflow. The ADM is tested against the LES of Heck et al. (2023), where the power of a single turbine in uniform flow at different yaw angles was computed by rotating the inflow angle and fixing the disc to the Cartesian grid. The present method is instead tested by keeping the inflow angle fixed and yawing the turbine. Various grid resolutions have been assessed to determine the minimum requirements for representing a grid-misaligned disc. These are characterised by the parameter n_D and by the cell aspect ratio in the streamwise direction, $AR = \Delta x/\Delta y$. Figure 3.9 shows the computed power for a yawed wind turbine with $C_T' = 4/3$ normalised by the power production of a turbine at zero yaw ($\gamma = 0°$), denoted by $\dot{P}$, for various resolutions and the reference of Heck et al. (2023). `WInc3D` accurately predicts the power as the yaw angle increases for all tested resolutions except when using very few points across the diameter ($n_D = 6.4$) combined with a high aspect ratio ($AR = 2.5$). The predictions of two low-order

models are also plotted for reference. The resolution used in the optimisations of Chapter 4 is $n_D = 24$ with $AR = 2.5$, which is sufficient for predicting the power of yawed or tilted turbines.

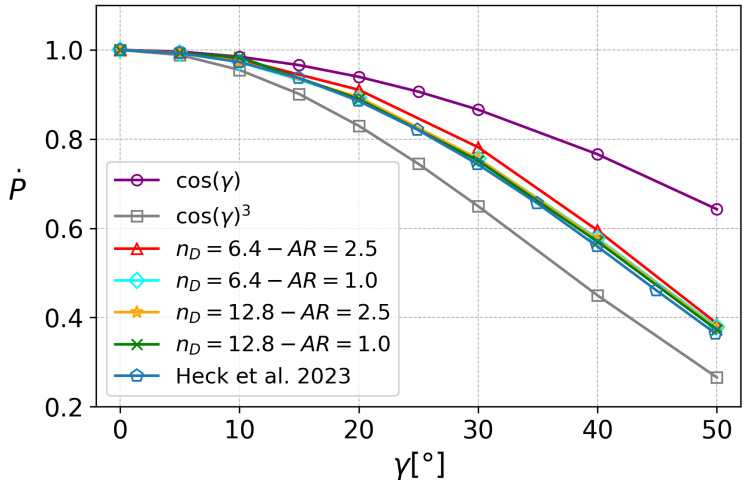


Figure 3.9: Power output of a yawed wind turbine normalised by the power output of a turbine at zero yaw angle. Comparison between different resolutions with `WInc3D` and the LES of Heck et al. (2023). The predictions of the low-order $\cos(\gamma)$ and $\cos(\gamma)^3$ models are also shown for reference.

3.2 Operational Wind Farm Validation

This section extends the validated LES framework to an operational onshore wind farm in moderately complex terrain. In contrast to the idealised benchmarks, application at utility scale requires additional modelling choices, including the computational domain, realistic inflow specification, heterogeneous surface roughness, and turbine characterisation. These elements are incorporated progressively to increase physical realism and their influence on farm-level predictions is assessed in a concise manner to demonstrate robustness, rather than through an exhaustive sensitivity analysis. The reference dataset was provided by SSE Renewables, enabling direct comparison between LES predictions and operational turbine performance. The principal objective is to evaluate the predictive skill of the framework against these SCADA records while also identifying the modelling choices most influential for farm-scale predictions.

The SCADA reference was provided in the form of relative turbine power distributions. Absolute production values were not available; instead, the comparison is restricted to normalised

power outputs, which emphasises the spatial variability of turbine performance across the farm. This approach isolates the flow–physics effects of terrain and wakes from calibration differences in absolute turbine performance.

Simulations are carried out for the predominant inflow direction at 210°, consistent with the main SCADA dataset. Two additional cases at 200° and 220° are also performed, enabling a 30° directional average for comparison with sector-aggregated SCADA statistics. For each direction, the computational domain is rotated such that the inflow aligned with the streamwise axis. The site comprises 12 turbines of rotor diameter $D = 82$ m and hub height $H_h = 62.5$ m. To preserve confidentiality, the farm is anonymised. Figure 3.10 shows the terrain elevation with the three computational domains used for the simulations.

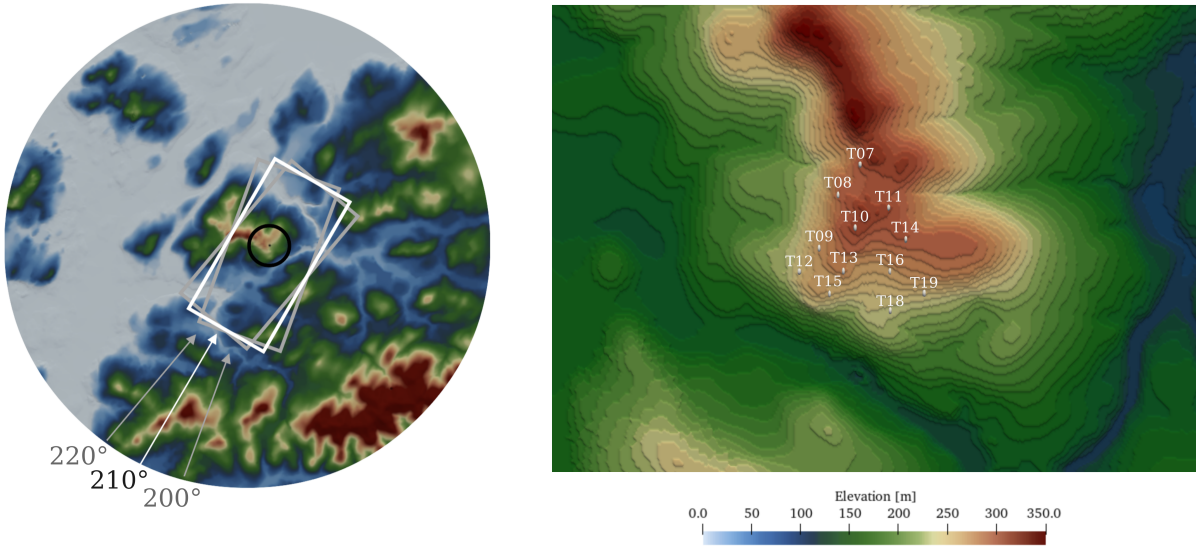


Figure 3.10: Top view of the terrain elevation map of the operational wind farm with indicative computational domains aligned with inflow directions at 200°, 210°, and 220° (left), and close view of the turbine layout with IDs at 210° (right).

3.2.1 Computational Domain and Terrain Representation

The computational domain is designed to capture terrain-induced modifications of the atmospheric boundary layer while minimising artificial blockage. The final configuration have dimensions $L_x \times L_y \times L_z = 12 \times 3 \times 6$ km, as illustrated in Figure 3.11. Boundary conditions include a no-slip wall-stress model at the bottom surface, free-slip at the top, convective outflow

downstream, and free-slip sidewalls to accommodate the lack of spanwise symmetry following the work of Yang et al. (2018).

Initial tests with a smaller domain of $8\times1\times4$ km revealed two key limitations. First, the vertical ceiling of 1 km was insufficient: with a boundary-layer depth of approximately 1 km and hills rising up to 300 m, the ratio of terrain frontal area to the domain cross-section was too large, causing more than 15% acceleration at the top boundary. Enlarging the vertical extent to 3 km reduced this artificial speed-up, ensuring that entrainment at the outer layer was represented without confinement. Second, the enlargement of the upstream terrain affected the predicted farm power. With an 8 km fetch the farm power was 6.9% higher than for the 12 km case, indicating that the inflow did not have sufficient distance to adapt to the terrain. Extending the domain to 12 km stabilised the prediction, and a further increase to 16 km altered the mean farm power by less than 1%, suggesting that the chosen 12 km extent provides a reasonable compromise between accuracy and computational cost.

The spanwise extent was found to be less critical in terms of farm-mean power: the difference between 4 km and 6 km width was under 1%. Nevertheless, the 6 km case was adopted to provide extra lateral padding and to avoid artificial sidewall effects.

At the inflow, the terrain elevation was smoothly ramped from zero to the true topography over a 500 m transition zone using a cosine function similarly to Yang et al. (2018), preventing discontinuities at the inlet boundary. Terrain data were otherwise used directly.

3.2.2 Grid Resolution and Computational Cost

Two grids were considered to assess sensitivity to spatial resolution. The coarser grid comprised $N_x \times N_y \times N_z = 577 \times 289 \times 289$ nodes, corresponding to uniform spacings $(\Delta x, \Delta y, \Delta z) =$ (20.8, 10.4, 20.8) m. With rotor diameter $D = 82$ m, this yields $n_D \approx 8$ grid points across the rotor. The finer grid doubled the resolution in each direction, with $1153 \times 577 \times 577$ nodes and spacings (10.4, 5.2, 10.4) m, corresponding to $n_D \approx 16$. Time integration used $\Delta t = 0.2$ s for the coarser grid and $\Delta t = 0.1$ s for the finer grid. In both cases, statistics were collected over

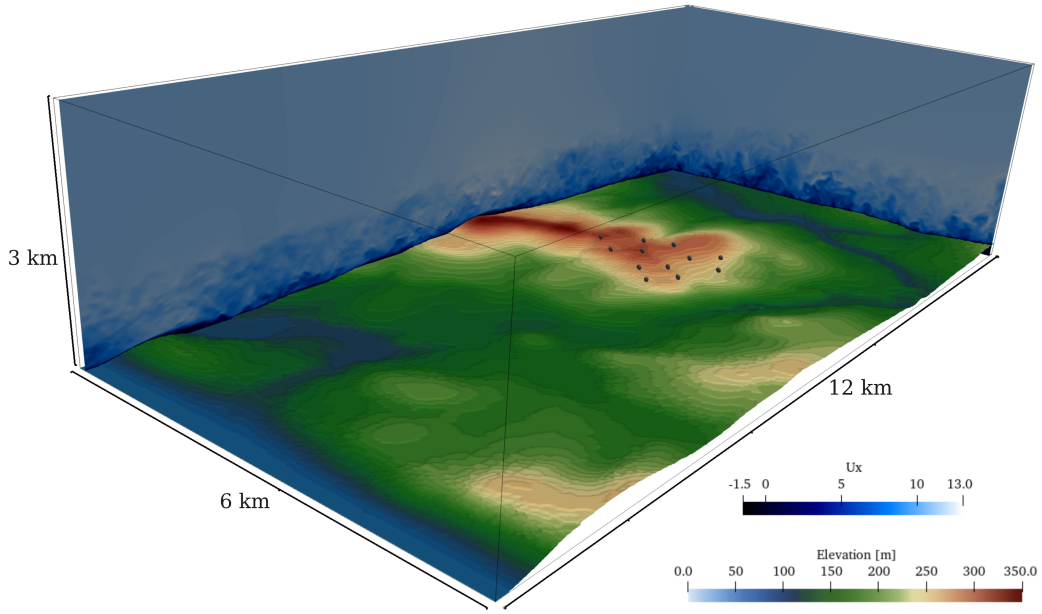


Figure 3.11: Computational domain representation of the operational wind farm simulations, showing terrain elevation, turbine positions, and instantaneous velocity field at lateral and outflow planes.

3.5 h of physical time after discarding an initial 15 min transient.

Given the relatively smooth topography at the site, grid design was governed primarily by rotor resolution rather than by terrain steepness; we therefore quantify resolution using n_D as in the idealised benchmarks. Consistent with the resolution guidelines derived from the idealised benchmarks ($n_D = 16$–32 for complex-terrain flows at aspect ratio $AR = 2$, and $n_D = 7$–13 for yawed discs depending on aspect ratio) the finer mesh offers additional resolution margin, while the coarser mesh sits at the lower end. Both grids yielded indistinguishable turbine-normalised power distributions in the SCADA comparison, so the coarser grid was adopted for production runs and used to present the results, with the finer grid retained for verification checks.

All simulations were executed on the ARCHER2 UK National Supercomputing Service (dual AMD EPYC™ 7742 nodes, 2.25 GHz). The production configuration on 2,048 CPU cores required approximately 1.2 h per simulation, while the high-resolution set-up on 8,192 cores required about 9.6 h.

3.2.3 Inflow Profile

The site is characterised by a neutral ABL of approximately $\delta = 1000$ m. To avoid confinement effects, the computational domain was extended to 3000 m in the vertical, which required an additional modelling treatment in the precursor simulations. A damping layer was applied above the capping inversion to suppress spurious acceleration in the outer flow. The forcing term followed the approach of Bempedelis et al. (2023):

$$F_i^{\text{damp}} = -\gamma\, \psi(y)\, \frac{u_*}{\delta}\, (\tilde{u}_i - G)\,, \tag{3.7}$$

where G is a target velocity profile, $\gamma = 4$ a constant that controls the relaxation strength, and $\psi(y)$ a vertical interpolation function applied near the capping inversion. The damping was introduced gradually within the last 100 m of the boundary layer using a smooth blending function:

$$\psi(y) = \tfrac{1}{2}\left[1 - \cos\left(\pi \tfrac{y - y_d}{d_h}\right)\right], \tag{3.8}$$

so that the forcing transitioned continuously from zero within the boundary layer to its full strength above the inversion. Here, $y_d = 900$ m is the start height of the blending region and $d_h = 100$ m its thickness. This avoided sharp discontinuities between the resolved LES field and the imposed outer-layer profile. In practice, turbulence generated by the surface roughness dominated the flow below the capping inversion, so the damping contributed little dynamically, but ensured numerical stability in tall domains.

Two alternative inflow profiles are tested. The first is a canonical logarithmic profile corresponding to a neutral ABL, with $u_* = 0.42$ and $y_0 = 0.032$ m, referred to as Ideal. The second is a terrain-modified profile extracted 3 km upstream of the farm from a precursor simulation, approximately fitted by a log law with $u_* = 0.55$ and $y_0 = 0.75$ m, referred to as Local. Although the latter is not derived from mast measurements, it can be considered as a hypothetical local inflow representative of upstream conditions. Figure 3.12 compares the horizontally averaged velocity and stress distributions from the precursor with log-law fits. Both inflows reproduced the expected log-law behaviour up to the capping inversion, confirming that the fits correctly

represent the provided data and that the boundary-layer statistics were well captured.

The choice of profile had a substantial impact on predicted farm-level production due to their different velocities at the hub height. The terrain-modified inflow (local profile) yielded an AEP 46% higher than the canonical logarithmic inflow. Since only normalised SCADA data were available, absolute production could not be validated directly. Instead, the analysis focused on whether each inflow reproduced the relative distribution of turbine power across the farm. On this basis, neither inflow showed consistently superior agreement with SCADA, and both fell within the range of natural variability. For subsequent simulations, the local inflow was adopted, as it better reflects upstream adjustment to terrain despite its less canonical shape.

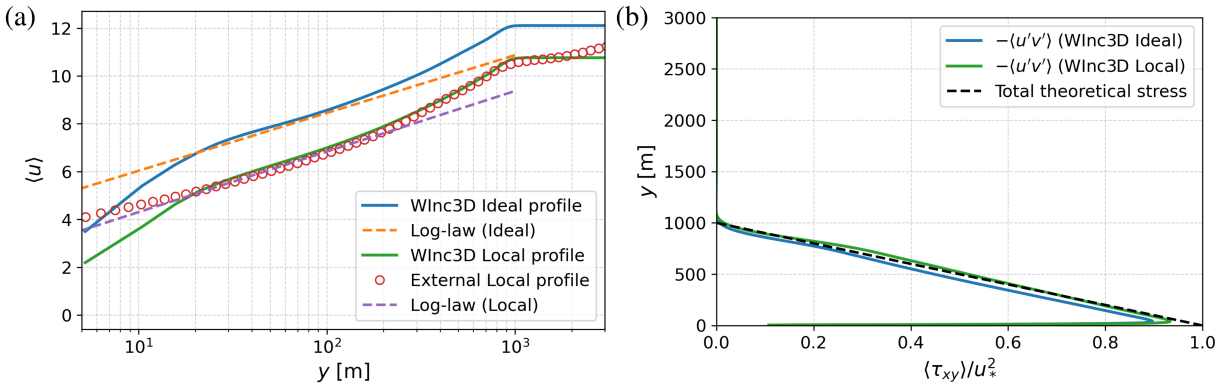


Figure 3.12: Horizontally- and time-averaged streamwise velocity of the ABL in the operational wind farm precursor simulation compared with log-law fits (a), and resolved stress profiles compared against theoretical stress (b).

Overall, the results confirm that while ABL statistics were matched accurately in the precursor, the inflow profile remains a dominant source of uncertainty for farm-scale predictions.

3.2.4 Roughness Map

Surface heterogeneity was represented by prescribing spatially varying roughness lengths in the wall-stress model. The roughness map used for the simulations is shown in Figure 3.13. Two values were applied: $y_0 = 0.032$ m for open terrain and $y_0 = 0.5$ m in regions corresponding to forestry. The distribution was obtained from available land-use data and mapped onto the Cartesian grid.

The difference between homogeneous and heterogeneous roughness was found to be negligible in terms of farm-level power output with our current modelling, with ΔAEP below 0.2%. This outcome is inconsistent since there are forested areas upstream of some turbines. Although they occupy a relatively small fraction of the domain this suggests that the wall-stress formulation employed here cannot substitute models where canopy-scale turbulence or momentum absorption processes are modelled or resolved. In reality, forestry modifies both the mean velocity profile and turbulence intensity through drag and wake recovery mechanisms that are not captured by prescribing an elevated y_0 alone and therefore modifying the wall modelling locally. A more detailed representation, for instance via explicit canopy models such as in Arthur et al. (2019), would therefore be required to assess the impact of forestry on turbine performance.

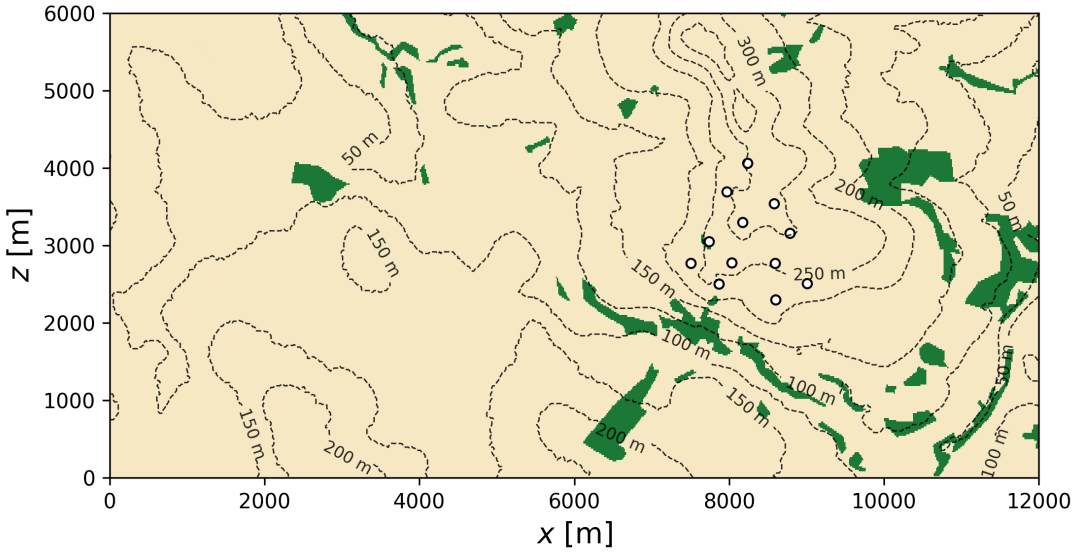


Figure 3.13: Surface roughness map prescribed for the operational wind farm: $y_0 = 0.032$ m for open terrain (beige) and $y_0 = 0.5$ m for forested regions (green).

3.2.5 ADM Configuration

In contrast to the idealised validation cases, where turbines were represented with simplified ADM, real turbines are operated with more complexity. First, their aerodynamic response is governed by performance curves for thrust and power rather than a constant C_T. Second, each turbine yaws to align with the local inflow direction. To reflect these features, the present simulations incorporate tabulated thrust and power curves derived from vendor data, and yaw

alignment was tested explicitly.

The ADM was coupled to tabulated thrust and power curves derived from vendor data, implemented as described in section 2.1.6. This ensured that turbine forces and power were consistent with the operational curves, rather than assuming a constant C_T. Figure 3.14 shows the thrust coefficient $C_T(u_0)$ and power coefficient $C_P(u_0)$ as functions of the undisturbed inflow velocity u_0. The impact on farm-level AEP was non-negligible: the use of tabulated curves reduced the predicted production by 3.3% compared to a constant-C_T ADM, consistent with the simulations operating partly in the quadratic power-growth regime and partly near rated conditions, where C_T decreases with wind speed. This highlights the importance of employing realistic turbine parameterisation for predictive studies.

Yaw angles were also tested, with one configuration in which all turbines were aligned to the global inflow and another where each unit was yawed to face the local wind direction. The resulting difference in AEP was minor (–0.3%), indicating that yaw correction was not a dominant factor under the present terrain and flow conditions. Nevertheless, local misalignment due to terrain-induced flow turning can be relevant in other sites.

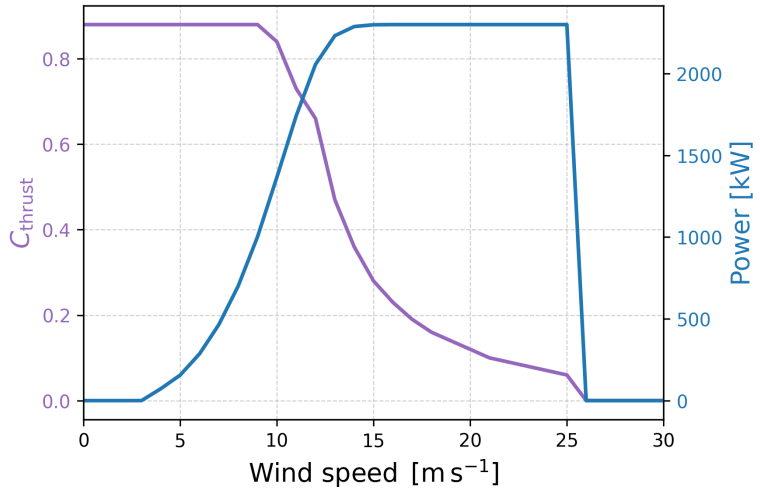


Figure 3.14: Turbine thrust coefficient and power curve, based on vendor data and used in the ADM for the operational wind farm.

3.2.6 Power Comparison with SCADA

The comparison with SCADA data is restricted to normalised turbine power, since only relative production values were available. Absolute AEP validation was therefore not possible. Instead, the focus is on whether the LES reproduces the spatial distribution of turbine performance across the farm, i.e. which turbines generate above or below the farm average. This approach emphasises the influence of flow physics.

Normalised power of the ith turbine (in %) is defined as:

$$P_{\text{norm},i} = \frac{P_{\text{turbine},i}}{P_{\text{farm}}/N_{\text{farm}}} \times 100, \tag{3.9}$$

where P_{farm} is the total farm power and $N_{farm} = 12$ is the number of turbines. In the SCADA reference, error bars correspond to the standard deviation of the power distributions within each directional bin (210° or 200°–220°). These reflect the natural variability of conditions (wind speed, turbulence, ABL stability) sampled across the dataset, rather than measurement uncertainty or intra-record fluctuations.

Figure 3.15 compares LES and SCADA turbine power for the 210° inflow direction. The LES reproduces the overall distribution of relative power across the farm, with most predictions lying within the SCADA variability bands. Some turbines, however, fall outside this range, with both overpredictions and underpredictions present (notably turbines 9, 12, 16 and 18). These discrepancies are interpreted in Section 3.2.7 using the simulated velocity and turbulence fields.

When averaging across the broader 30° sector (200°–220°), shown in Figure 3.16, the agreement improves further. All turbines fall within the SCADA variability, and the mean distribution is reproduced with better accuracy. Sector averaging smooths out case-specific sensitivities, reducing the impact of small angular shifts in inflow, intermittent canopy effects, and ABL stability variations. Nevertheless, some deviations remain, particularly for turbines 14 and 18, indicating the persistence of local terrain and roughness influences even under directional averaging.

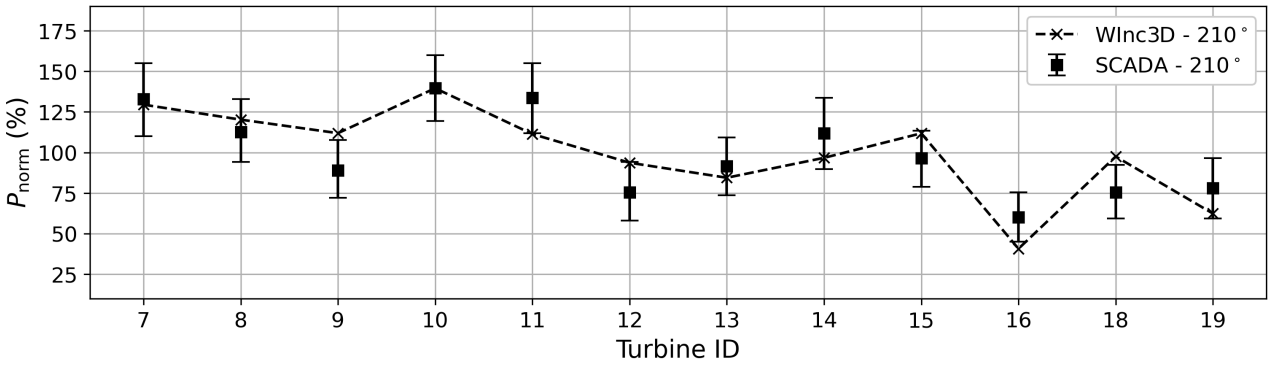


Figure 3.15: Normalised turbine power at 210° inflow for the operational wind farm. LES predictions compared with SCADA mean values and standard deviations.

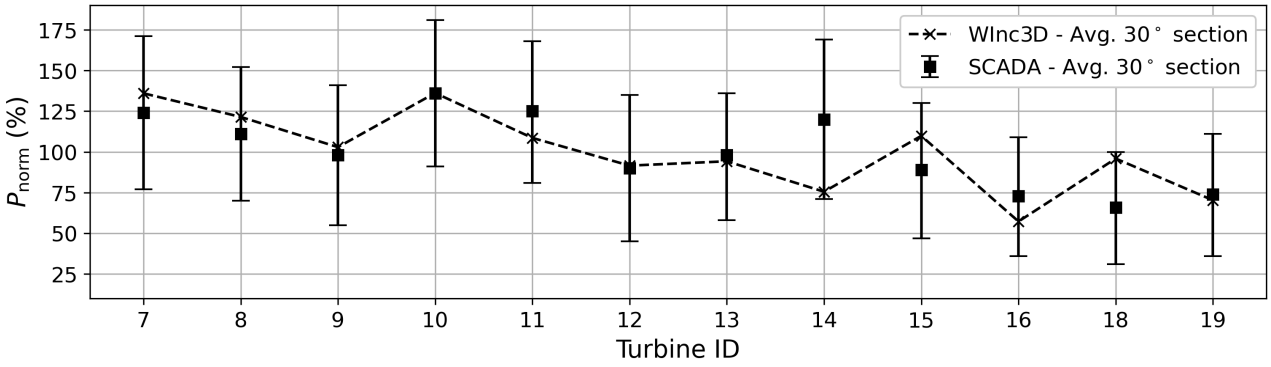


Figure 3.16: Normalised turbine power averaged over the 200°–220° sector for the operational wind farm. LES predictions compared with SCADA mean values and standard deviations.

Overall, the framework reproduces farm-level power distribution within the observed variability and captures the spatial trends across the farm, with specific mismatches discussed in the next section. These results confirm that neutral-ABL LES with the present modelling choices is capable of predicting utility-scale performance within operational uncertainty. Further improvements could be expected by incorporating additional atmospheric stability states and more detailed surface representations (e.g. explicit canopy models), which would help to bring the predictions closer to the long-term mean values observed in SCADA data.

3.2.7 Flow-Field Interpretation

To interpret the 210° power distribution, the simulated flow fields are examined at hub height together with streamwise turbulence intensity. These provide physical insight into terrain-induced acceleration, wake interactions, and roughness effects that drive the observed discrepancies.

Figure 3.17 shows the hub-height velocity contours. The behaviour mirrors the idealised hill cases presented earlier: acceleration occurs on windward slopes and hilltops, while reduced velocities appear in the lee-side wakes. This is most evident at the upper ridge of the domain, where the flow transitions from peak speed on the crest to strong deceleration downstream. No regions of reverse flow are observed at hub height. Superimposed on these terrain-driven patterns are the wakes of individual turbines, clearly visible as elongated streaks of momentum deficit. Higher inflow upstream of front-row turbines (7–11) is consistent with their above-average power production, while turbines strongly impacted by upstream wakes (e.g. 13, 16 and 19) show reduced output. Notably, turbines 12, 15 and 18 appear unobstructed in the velocity field, yet their LES power is overpredicted relative to SCADA. The roughness map (Figure 3.13) shows that these turbines are downstream of forested areas, where canopy drag induces momentum loss not captured at hub height. This strongly suggests that surface heterogeneity contributes to the mismatch. In addition, some wakes (e.g. from turbines 9 and 12) show clear terrain-driven deflection, highlighting the added value of LES in resolving flow–terrain–wake interactions.

The corresponding turbulence field is shown in Figure 3.18. Turbulence intensity is expressed as

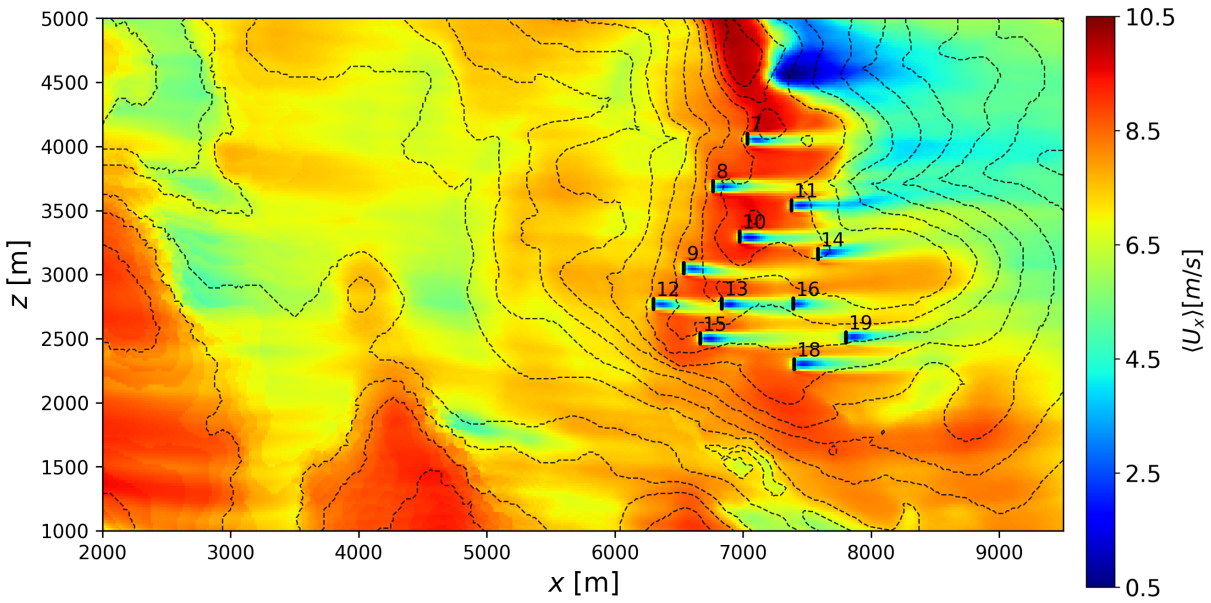


Figure 3.17: Time-averaged hub-height streamwise velocity field at 210° inflow over the operational wind farm. Terrain contours (dashed lines) are shown at 25-m intervals, and turbine IDs are indicated.

σ_u/U_∞, where σ_u is the standard deviation of the streamwise velocity fluctuations and U_∞ is the reference velocity at the top of the atmospheric boundary layer. Terrain-induced turbulence is mainly confined to the upstream hills and the main ridge, with limited direct impact on the turbines, although slightly elevated values appear ahead of the front row. Within the farm, elevated TI values coincide with the turbine discs, particularly at their edges, and are amplified for units immersed in downstream wakes. Turbines exposed to higher turbulence intensity are subject to increased structural loading and accelerated fatigue. While canopy patches are expected to enhance turbulence in addition to reducing mean velocity, this effect is not represented in the present modelling.

The flow-field analysis demonstrates that the LES captures the essential terrain–wake interactions: acceleration and sheltering induced by hills, deflection of wakes by terrain features, and turbulence build-up within the farm. These mechanisms provide a consistent explanation for the spatial distribution of power discussed in Section 3.2.6. Importantly, no large-scale recirculation zones are observed at hub height, suggesting that in this case the main performance drivers are ridge-induced acceleration, wake overlap, and canopy effects rather than terrain-induced reverse flow.

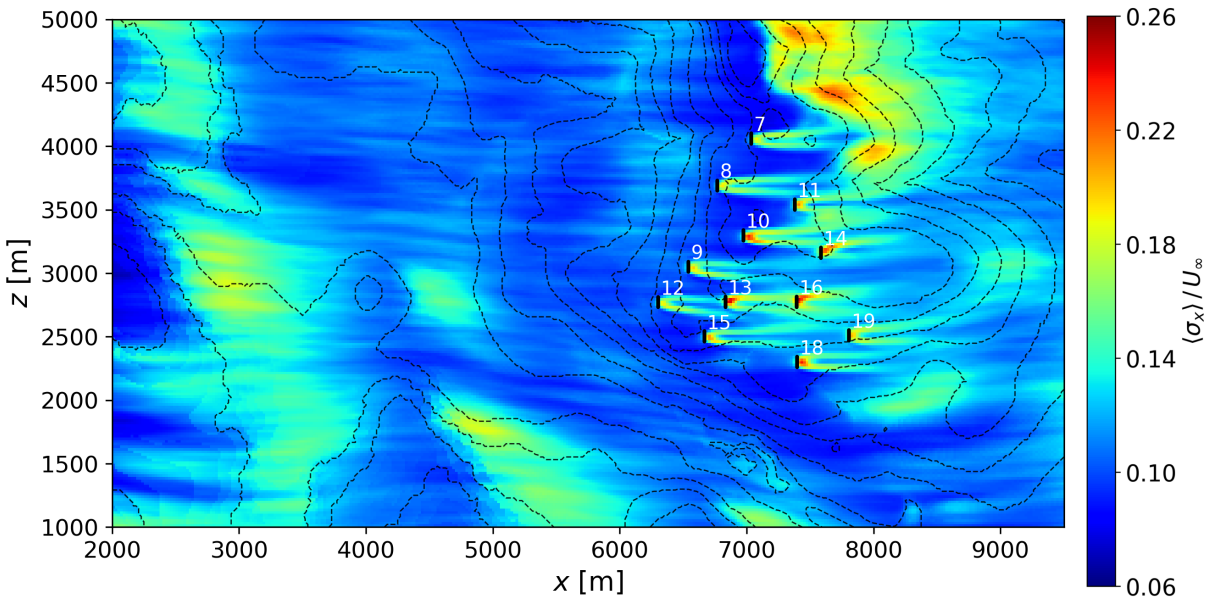


Figure 3.18: Time-averaged hub-height streamwise turbulence intensity field at 210° inflow over the operational wind farm. Terrain contours (dashed lines) are shown at 25-m intervals, and turbine IDs are indicated.

3.3 Conclusions

The purpose of this chapter was to establish the validity of the LES framework introduced in Chapter 2, ensuring that it can be used with confidence in the optimisation studies that follow. Validation was approached progressively, starting from controlled wind-tunnel experiments and culminating in an operational wind-farm case with SCADA data. Together, these benchmarks provide a comprehensive assessment of the modelling components underpinning the optimisation framework.

The controlled validation cases confirmed that the framework accurately captures terrain-induced flow features (acceleration, separation, reattachment) and turbine–terrain interactions across canonical 2D and 3D hills. Grid-refinement studies established resolution requirements, whilst tests on yawed configurations validated the super-Gaussian force distribution for misaligned rotors.

The operational validation against a 12-turbine onshore farm demonstrated that LES reproduces observed spatial performance trends under neutral conditions, despite uncertainties in absolute predictions. Brief sensitivity studies identified domain design as the primary factors

influencing predictive skill, whilst inflow specification, turbine parameterisation, surface roughness variations and yaw alignment had secondary effects in the normalised power production.

Taken together, these results show that the LES framework consistently reproduces boundary-layer development, terrain-induced flow features, and turbine performance across both controlled benchmarks and utility-scale conditions. The elements that most strongly influence predictive skill (domain design, inflow specification, and turbine parameterisation) were clearly identified, while other refinements exert only secondary effects. Within the adopted assumptions of neutral ABL conditions and actuator-disc representation, the framework is sufficiently robust for comparative and optimisation studies. In this sense, the validation provides not only technical reassurance but also a methodological foundation for the rest of the thesis.

The next chapter applies the validated LES framework, combined with BO, to a controlled two-turbine case over a Gaussian hill. This proof-of-concept study isolates terrain–wake interactions and tests whether the optimisation methodology can systematically exploit them to enhance power capture.

Chapter 4

Two-Turbine Optimisation over the 2-D Hill

Following the validation of the LES framework in Chapter 3, we now employ the methodology as a predictive tool to perform turbine layout optimisation. As a first application, this chapter examines the simplified yet representative case of two turbines sited on the two-dimensional Gaussian hill. This configuration allows us to isolate the interplay between terrain-induced flow and turbine–wake interactions in a controlled environment. The work presented here has been published in Jané-Ippel et al. (2024).

The optimisation is formulated to maximise the combined power extraction of the two turbines, subject to constraints on turbine spacing and domain boundaries. BO, introduced in Chapter 2, is employed to efficiently explore and exploit the design space, which includes the relative streamwise positions of the turbines with respect to the hill crest and their hub heights. The use of a reduced, two-turbine setup ensures computational tractability while retaining the essential physics of siting decisions in complex terrain.

The chapter is structured as follows. Section 4.1 introduces the problem configuration and optimisation setup. Section 4.2 presents the baseline and optimised layouts, along with flow-field diagnostics and discusses the robustness of the findings and their implications for turbine siting in hilly terrain. Finally, Section 4.3 summarises the outcomes and provides a link to

the following chapter, which extends the optimisation framework to a full wind-farm case over three-dimensional terrain.

4.1 Problem Setup

The setup follows the work introduced in the previous chapter (Section 3.1.1), where the hill geometry and inflow conditions were defined. As concluded, the finer G2 grid is adopted despite its higher computational cost, since turbine positions are optimised in close proximity to each other and to the recirculation region where higher resolution is required. Here, the focus is placed on specifying the design variables, objective function, and the configuration of the optimisation runs.

4.1.1 Design Variables and Objective Function

The two optimisations performed, henceforth referred to as BO1 and BO2, have four design variables each. A schematic of the hill and the design variables used in BO1 and BO2 is illustrated in Figure 4.1. The design variables for BO1 include the wind turbines' streamwise locations, denoted as x_1 and x_2, and their hub heights, represented by h_1 and h_2. These variables are normalised by the turbines' diameter. Specifically, the streamwise location x_1 spans the top of the hill ($x_1 \in [-1.25, 1.25]$), while x_2 spans the leeward side of the hill ($x_2 \in [0.75, 5.00]$). These choices allow for a thorough exploration of the hill's influence on turbine placement and its corresponding performance. Additionally, they aim to exploit the leeward region of the hill, which is not usually explored. Regarding the hub heights, ranges of $h_1, h_2 \in [0.70, 1.50]$ are selected to ensure a diverse set of turbine configurations. To ensure the turbines operate safely, a constraint is imposed in both optimisations to always maintain a minimum distance of around one turbine diameter between their centres ($x_2 - x_1 \geq 1.05$).

In BO2, the design space is modified upon the learnings from BO1 and to explore the possible benefits of wake steering. To this end, the upstream turbine's hub height is fixed to $h_1 = 1.0$,

removing it from the design variables. We incorporate new design variables which can substitute the hub height optimisation with a wake steering strategy. Given the use of periodic boundary conditions in the spanwise direction, the tilt angle is considered as a means of wake steering instead of the more commonly used yaw. This approach allows us to directly influence the mid-span plane as in BO1. The tilt angle, denoted as α_1, is optimised, allowing it to vary between -20° and +20° ($\alpha_1 \in [-20°, +20.0°]$). This range choice is similar to that used by previous studies (Annoni et al., 2017; Cutler et al., 2021; Steiros et al., 2025), which have explored tilt angles at the wind farm level on flat terrain. A negative tilt angle implies potentially aligning the turbine more closely with the inclined mean flow typically observed at the flow in the hilltop, thereby enhancing the power production of the upstream turbine. High values of negative tilt, such as -20°, are more suitable for downwind turbines, as they avoid the risk of tower strikes often associated with such configurations in conventional turbines (Janajreh et al., 2010). The turbine wake is expected to steer downwards for the negative tilt. Conversely, a positive tilt angle steers the wake upwards, which may be strategically used to mitigate adverse effects on downstream turbines while optimising the airflow capture by the upstream turbine. This dual approach allows the optimisation to evaluate whether steering the wake away from downstream turbines (positive tilt) or improving alignment with the local inflow at the hilltop (negative tilt) results in better overall system performance.

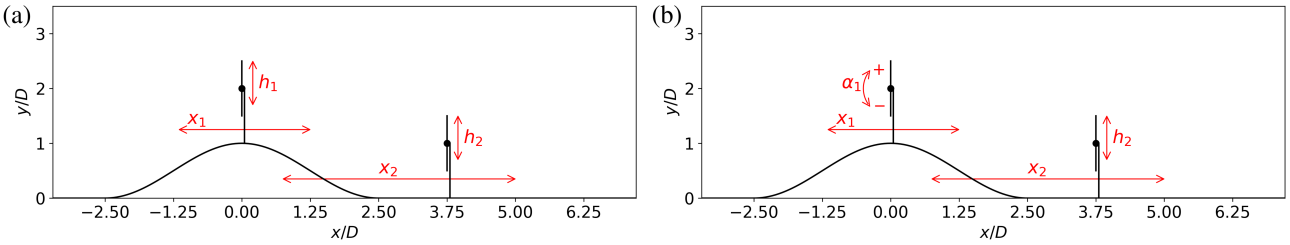


Figure 4.1: Schematic of the 2D hill with two wind turbines placed around it showing the design variables for optimisations BO1 (a) and BO2 (b). The example is for $x_1 = 0.0$, $h_1 = 1.0$, $\alpha_1 = 0°$, $x_2 = 3.75$, and $h_2 = 1.0$.

The objective function for both BO1 and BO2 optimisations is the aggregate power output of the turbine ensemble, defined as the time-averaged power produced by the pair of turbines. The power estimates come from the LES following the methodology presented in Section 2.1.6. The power output of the turbines, represented as $\hat{P}$, is normalised by the average power produced

by a single turbine under similar inflow conditions on a flat terrain (i.e., without the hill, in isolation). To ascertain a short but reliable time window for statistical averaging of the power, an additional investigation is performed utilising two reference simulations: one consisting of a wind turbine on top of the hill ($x/D = 0$) and the other placing the wind turbine in the hill-wake ($x/D = 5.0$). The chosen cut-off for the simulation time is dictated by achieving a variability of less than 1% in the time-averaged power for a minimum of 6 non-dimensional time units and convergence in the CI of the averaged power signal. This threshold aims to ensure statistical convergence in the context of the turbulence-induced chaotic flow behaviour.

Figure 4.2 presents the power signals for the two indicative scenarios, including the cumulative and total averages and 99% CI. The confidence intervals are computed following the methodology detailed in Bergmann et al. (2021). This involves utilising a correlation function to establish confidence intervals that account for the inherent temporal correlation of the Navier-Stokes equations from the LES. The plots confirm that from $\hat{T} = 10$ onward, the cumulative power adheres to the established variability criterion, and that from $\hat{T} = 16$ onward the confidence intervals converge, showing only minimal reductions.

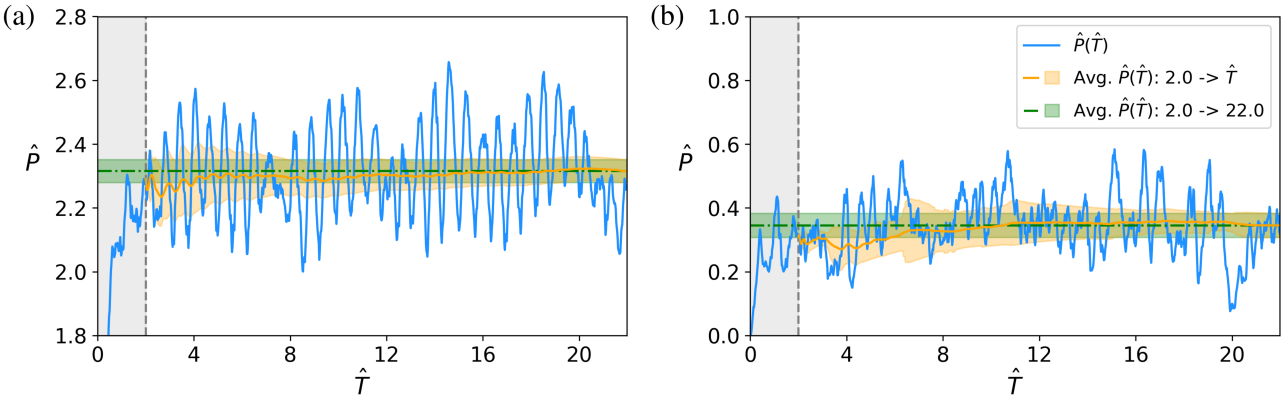


Figure 4.2: Power signals with total and cumulative averages and 99% CI for (a) wind turbine placed at the hilltop ($x/D = 0$) and (b) wind turbine immersed in the hill's wake ($x/D = 5.0$). The grey region indicates the initial transient discarded.

4.1.2 Optimisation Configuration

The BO framework employed here follows the methodology introduced in Chapter 2, and only the case-specific settings are summarised. Both BO1 and BO2 begin with a DoE of $n = 20$ samples, generated using Latin Hypercube Sampling (five samples per design variable). Subsequent evaluations are guided by a GP surrogate model with a Matérn 5/2 kernel with ARD (a lengthscale per design variable) the LCB acquisition function. The acquisition parameters are set to $w = 2$ and $\epsilon = 0.01$, and a batch size of four design points is selected at each iteration using local penalisation.

We considered a budget of 100 evaluations for each optimisation. Convergence is assessed by determining the cumulative number of simulations at which the improvement of the maximum value of power found falls below a certain threshold, in this case set to 1%, similar to the criterion selected for the end of the time-averaging period of the power statistics.

Figure 4.3 illustrates the progression of BO1 and BO2, showing the power deficit ΔP with respect to the global optimum found for all the simulations of the optimisation (in blue) and the optimal values found up to that point as the optimisation progressed (in black). A horizontal red line has been added marking the -1% deficit. For BO1, the optimal power converged after 56 simulations. Similarly, BO2 reached this convergence threshold after 40 simulations. Beyond this point, the lower performance points along the blue line, explore and refine the model's understanding of less optimal regions. This exploration enhances the robustness of the GP model, particularly by confirming the boundaries of lower performance and better defining the gradient towards the optimum. Moreover, the clustering of points near the optimum shown in the blue line, increases confidence in the region around the optimum.

Each evaluation based on one LES leveraged 2048 cores from the ARCHER2 UK National Supercomputing Service, each with dual AMD EPYC^{TM} 7742 64-core processors at 2.25GHz. The wall clock time per LES is approximately 3.6 hours. Across 100 evaluations, the cumulative computational cost is around 737,280 CPU-h. Each optimisation requires approximately five days of wall clock time to complete, accounting for parallel initial sampling and the use

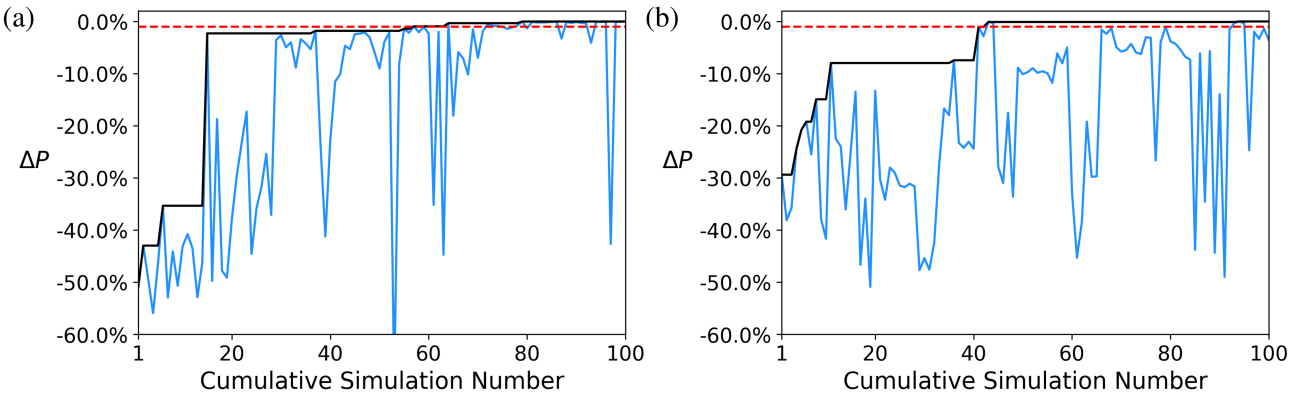


Figure 4.3: Deficit from the optimum power as a function of the cumulative simulation number for (a) BO1 and (b) BO2. The red dashed line indicates the -1% power deficit.

of batches of four evaluations at each surrogate model update. Note that the computational demands and time associated with the BO algorithm are marginal compared to the LES evaluations.

4.2 Results

This section presents the findings from BO1 and BO2. The results compare the best designs from both optimisations with a reference scenario, henceforth referred to as the Reference case, where a wind turbine is positioned on top of the hill, with a second turbine placed at a distance of 5D downstream of the first turbine. An overview of the optimum design variables and the corresponding values of the objective function for both BO1 and BO2, alongside those of the Reference case, is provided in Table 4.1.

Table 4.1: Summary of optimum design variables and objective functions obtained from BO1 and BO2, alongside the Reference case.

Case	x_1	h_1	$\alpha_1[^\circ]$	x_2	h_2	Total Power
Reference case	0.00	1.00	0.0	5.00	1.00	2.70
BO1	**0.12**	**1.50**	0.0	**1.17**	**1.02**	4.66
BO2	**-0.12**	1.00	**8.72**	**1.39**	**0.98**	3.61

Figure 4.4 illustrates the turbine placement in a flow visualization of the two optimal scenarios, BO1 and BO2. The visualisation captures the instantaneous normalised velocity magnitude

($\|\mathbf{u}\|/U_\infty = \sqrt{u_x^2 + u_y^2 + u_z^2}/U_\infty$) across the planes and a vorticity contour ($\omega = 500\,\mathrm{s}^{-1}$) in white, along with the hill and the turbine pair. This snapshot shows the interaction between the TBL and the hill, as well as the strategic positioning of the turbines and their interrelated dynamics.

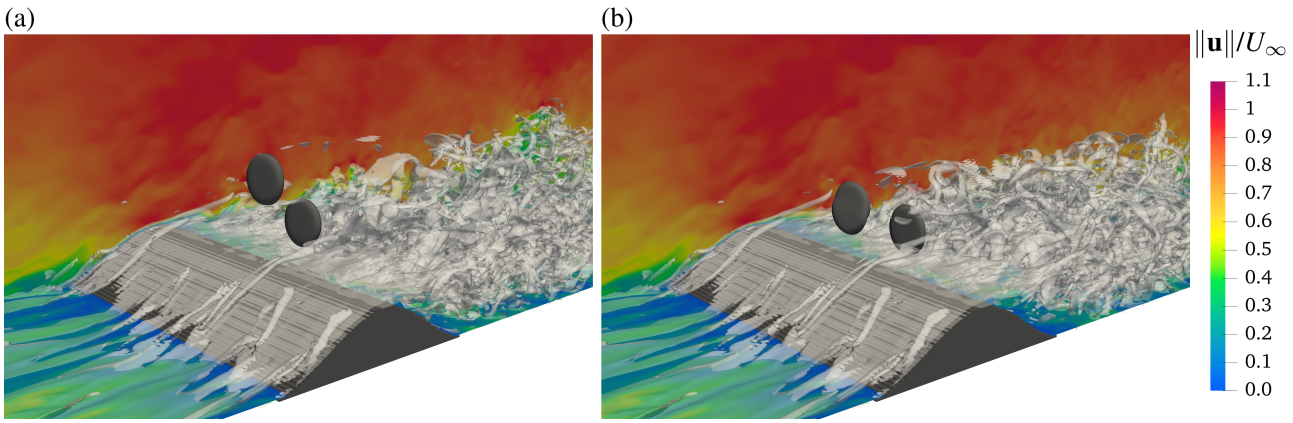


Figure 4.4: Instantaneous flow visualisation of the optimum cases BO1 (a) and BO2 (b), showing normalised velocity magnitude and vorticity contours (white).

In BO1, there is a significant increase in power, approximately 73% more than in the Reference case. The upstream wind turbine is located near the hill's centre, slightly shifted to its leeward side, and at an elevated hub height. The second turbine, placed at a distance of 1.05D satisfying the predefined constraint, benefits from the undisturbed wind left by the upstream turbine, as depicted in Figure 4.4a. BO2 was conceived in an attempt to reproduce the same behaviour through changes in the tilt angle of the upstream wind turbine instead of its hub height, thereby exploring the merits of wake steering. The best case of BO2 shows a 34% power improvement compared to the Reference case. The upstream turbine is positioned near the hill's centre, slightly towards the front, at the fixed hub height of 1D, as seen in Figure 4.4b, with a tilt angle of 8.72°. This indicates that the upstream turbine is sacrificing its power generation for an overall improvement. The second turbine in BO2 is placed further from the upstream turbine compared to BO1, at a distance of 1.51D, potentially leaving more space for the wake steering to deflect vertically the wake and have a greater effect on the downstream turbine. The difference in total power between BO1 and BO2 indicates that the limitation in tower height for the upstream wind turbine potentially limits the achievable maximum power output of BO2.

To comprehend the improved performance of the optimised configurations, we analyse the power series predicted by the ADM of the turbine pairs, as well as the mean velocity and turbulence intensity fields.

4.2.1 Power Production

Figure 4.5 illustrates the time history of power production for each turbine, along with their cumulative averages and 99% CI for the Reference ($\hat{P}_{Ref}$), BO1 optimum ($\hat{P}_{BO1}$), and BO2 optimum ($\hat{P}_{BO2}$) cases.

In the Reference configuration, the resultant total normalised power stands at $\hat{P}_{Ref} = 2.7$. The upstream turbine generates a significant portion, $\hat{P}_{Ref_{Up}} = 2.33$, primarily due to the speed-up effect caused by the hill's presence. Conversely, the downstream turbine's contribution is only $\hat{P}_{Ref_{Down}} = 0.37$. Its reduced output is attributed to its immersion in the wake of both the hill and the upstream turbine, which adversely affects its power generation. This power distribution is consistent with observations made when positioning the turbines individually around the hill, as illustrated in Figure 3.6. Furthermore, there are notable differences in the power series of the two turbines, primarily influenced by their distinct operating environments. The upstream turbine, directly impacted by the TBL, exhibits significant variability in its power output. The TBL's dynamic and spatially inhomogeneous nature, though statistically homogeneous in the spanwise direction, is reflected in the power evolution of the upstream turbine, characterised by larger amplitude fluctuations and larger time-scale dominance. In contrast, the downstream turbine, positioned within the wake of the hill and the upstream turbine, demonstrates a signal characterised by higher frequency and lower amplitude fluctuations.

The BO1 optimal scenario achieves a total power of $\hat{P}_{BO1} = 4.66$, with the upstream and downstream turbines generating $\hat{P}_{BO1_{Up}} = 2.31$ and $\hat{P}_{BO1_{Down}} = 2.35$, respectively. Elevating the hub height of the upstream turbine and optimising the downstream turbine's placement results in a marked power surge in the latter, as evidenced by $\hat{P}_{BO1_{Down}} >> \hat{P}_{Ref_{Down}}$. Remarkably, the power yield from the downstream turbine in BO1 is slightly above the upstream turbine power and presents similar signal characteristics to the upstream turbine. Furthermore, it is

noteworthy that the power of the upstream turbine does not increase as it might be expected with the elevation of its hub height. This suggests that by elevating the upstream turbine, it potentially misses out on the benefits of the speed-up effect caused by the hill's presence and as seen in Figure 3.3, the vertical gradient in the streamwise velocity is limited.

The BO2 configuration garners a total power of $\hat{P}_{BO2} = 3.61$, divided between the upstream and downstream turbines as $\hat{P}_{BO2_{Up}} = 2.14$ and $\hat{P}_{BO2_{Down}} = 1.47$, respectively. The reduction in the upstream turbine of BO2 can be explained by the use of the tilt angle of the disc. The downstream turbine continues to exhibit a substantial boost in its power generation, although lower compared to BO1 ($\hat{P}_{BO1_{Down}} > \hat{P}_{BO2_{Down}}$) due to the limitations in the upstream turbine's hub height. The time history of $\hat{P}_{BO2_{Down}}$ displays similar behaviour to that of $\hat{P}_{Ref_{Down}}$ but with higher magnitude and amplitude. This could be the imprint of the upstream TBL and its intermittent interaction with the upstream ramp of the hill, since both signals are affected by it, as discussed in the following section.

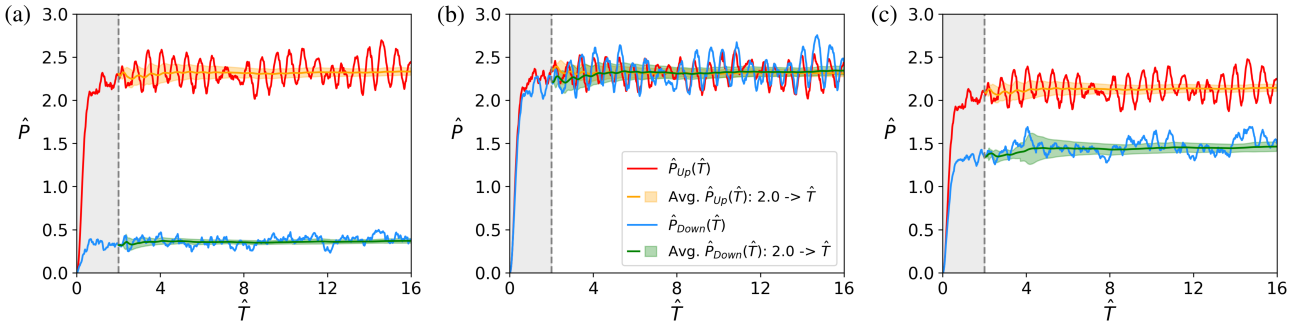


Figure 4.5: Power series with cumulative averages and 99% CI of (a) the Reference case and the optimum cases of (b) BO1 and (c) BO2.

4.2.2 Velocity and Turbulence Intensity Fields

The difference in power generated by the downstream wind turbine across the three scenarios can be better understood by examining the time-averaged velocity fields. Figures 4.6a, 4.6c, and 4.6e depict the normalised time-averaged streamwise velocity field in the mid-plane for each scenario, with streamlines represented by black lines. The contour of zero streamwise velocity, highlighting the edge of reverse flow, is demarcated by the white line. For clarity, Figure 4.6

includes the tower of the turbines; however, these are not modelled in the simulations.

In the Reference scenario, the downstream turbine is adversely influenced by the low velocity imparted by the wake of both the hill and the upstream turbine, explaining its diminished power output. Conversely, the BO1 optimum configuration positions the downstream turbine to circumvent both the hill and turbine wake, attempting to exploit the local flow acceleration between the hill and the upstream wind turbine. In the BO2 optimum case, due to the limitations in hub height, the downstream turbine cannot avoid the upstream wakes, and therefore, it overlaps with the upstream turbine wake. The wake steering performed by the upstream turbine allows an increase in the kinetic energy available downstream by deflecting the wake vertically at the cost of producing less power upstream. As shown in the next section, this trade-off is not trivial. In both optimum cases, power increases are due to the second turbine placed very close to the first one. This is owed to the blockage of the upstream turbine (see Section 4.2.3). Although this turbine acts as a sink of momentum overall, it diverts some of the flow momentum to the region below it, which can be harvested by the second turbine, but only if placed close enough. Also, it is clear that in all cases the presence of the turbines influences the recirculation region of the hill. Both effects cannot be fully captured by the superposition technique, showcasing the merits of using a LES-based optimisation.

Further insights can be obtained from the second-order statistics. Figures 4.6b, 4.6d, and 4.6f display the normalised time-averaged streamwise turbulence intensity in the mid-plane for each scenario. In the Reference case, the downstream turbine experiences a highly turbulent flow, triggered by the flow separation at the hill, combined with the turbulence generated by the upstream turbine. In contrast, the BO1 optimum operates under relatively clean conditions. The BO2 optimum case is affected marginally by the wakes of the hill and the upstream turbine, albeit less compared to the Reference case. High levels of turbulence intensity appear when the upstream and downstream turbine wakes merge in BO2. In both the Reference case and BO2, the downstream turbine would be negatively impacted in terms of dynamic loadings due to turbulence. Overall, the optimally configured turbines reduce the turbulent kinetic energy of the hill wake when compared to the Reference case.

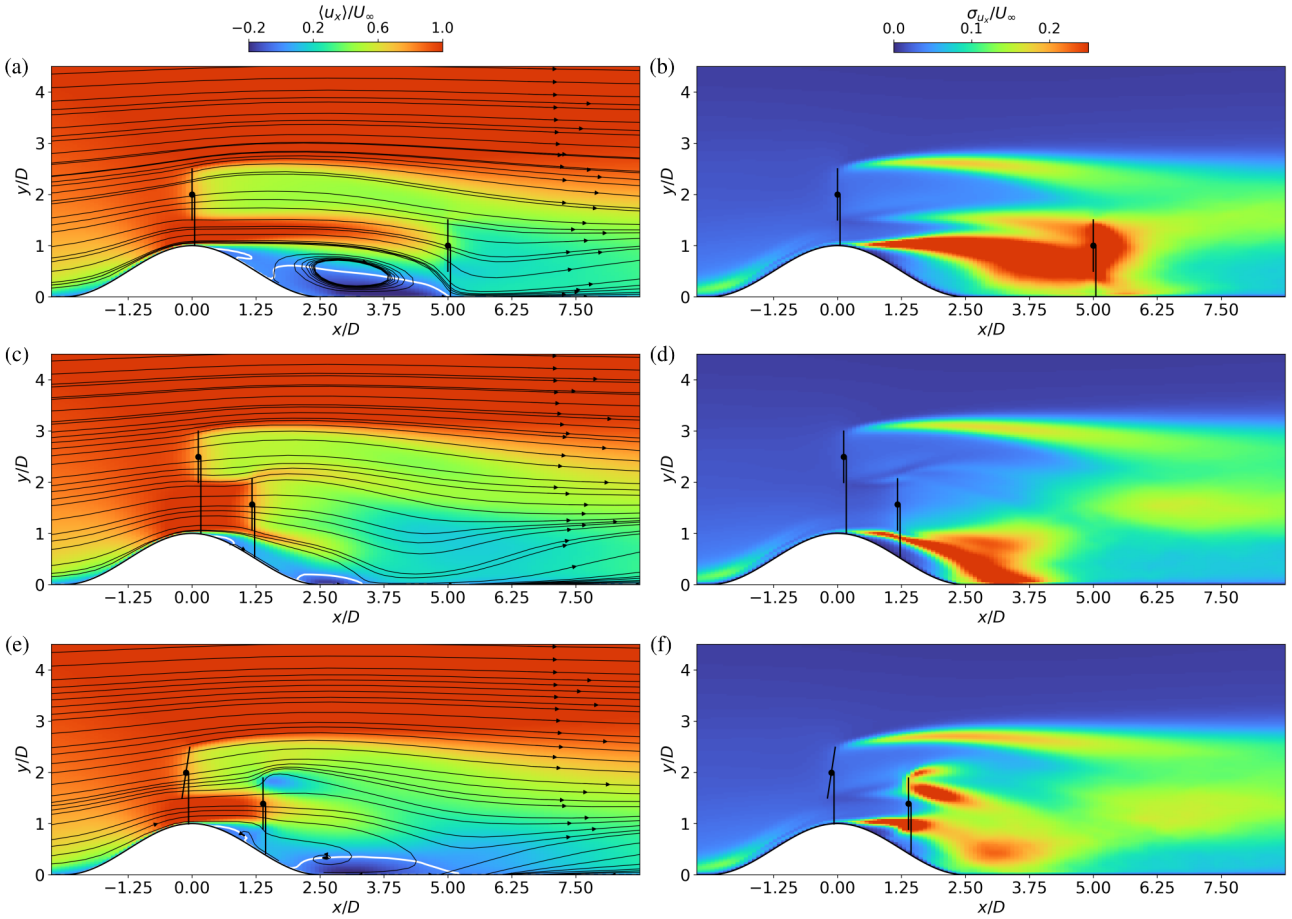


Figure 4.6: Normalised time-averaged streamwise velocity fields (left column) and streamwise turbulence intensity fields (right column) at the symmetry plane for the Reference case (top row) and both optimised scenarios, BO1 (middle row) and BO2 (bottom row). The black lines are streamlines of the velocity field, while the white lines indicate the contours of zero streamwise velocity.

4.2.3 Influence of the upstream wind turbine

The optimisation results indicate that the wind turbines should be placed close to each other, prompting a closer look at the influence of the upstream turbine in terms of blockage and wake steering.

In BO1 we compare the flow fields from simulations with and without the upstream turbine to analyse this influence. The presence of the upstream turbine leads to a 5% increase in the power of the downstream turbine, a result of the flow being redirected by the upstream turbine's blockage. Figure 4.7a illustrates the velocity deficit evaluated in a streamwise plane 0.25D upstream of the downstream turbine ($x/D = 0.92$). The velocity deficit in BO1 is defined as $\langle u_{D-BO1} \rangle = \langle u_{\mathrm{nt}} \rangle - \langle u_{\mathrm{t}} \rangle$, where $\langle u_{\mathrm{nt}} \rangle$ is the time-averaged streamwise velocity field in the absence of the upstream turbine and $\langle u_{\mathrm{t}} \rangle$ is the mean streamwise velocity with the presence of the upstream turbine, both normalised by U_∞. The horizontal line at $y/D = 0.7$ marks the hill's edge at this longitudinal location, and the dashed lines show the contours of velocity deficit at -0.02 and -0.04. These contours indicate regions where the flow speed is increased by 2% and 4% of U_∞ respectively. The turbines' centres and the extent of their rotor discs are projected onto the plane for scale. The influence of the upstream turbine is visible as a positive velocity deficit, slightly covering the downstream turbine's projected area. The contours reveal areas of faster flow close to the upstream turbine and along the hill edge. This flow acceleration spreads across a large area impacting the downstream turbine, which explains the observed performance boost. The acceleration is particularly strong near the hill, indicating that the upstream turbine's presence has altered the hill's recirculation zone across the span of the hill.

In BO2, the upstream turbine has a positive tilt angle, which, despite reducing its power output by 6% if compared to the no-tilt case, leads to a substantial increase of 23.5% for the downstream turbine's power. Figure 4.7b illustrates the alterations in the flow pattern due to the upstream turbine's tilt, applying the velocity deficit concept as defined previously. The velocity deficit in BO2 $\langle u_{D-BO2} \rangle$ is computed by contrasting the flow with and without the upstream turbine's tilt at a plane 0.25D upstream of the downstream turbine ($x/D = 1.139$). The horizontal line at $y/D = 0.57$ designates the hill's edge, and the dashed lines delineate the contours of

velocity deficit at -0.12, -0.02, 0.02, and 0.12. Tilting diverts the wake upwards, creating the kidney-shaped accelerated flow in the lower region of the wake's trajectory, which benefits the downstream turbine over a considerable area. Concurrently, as also seen in Figure 4.7a, the tilt induces modifications in the hill's recirculation, notably a horizontal deceleration over the hill edge, which has a marginal adverse effect on the downstream turbine's performance.

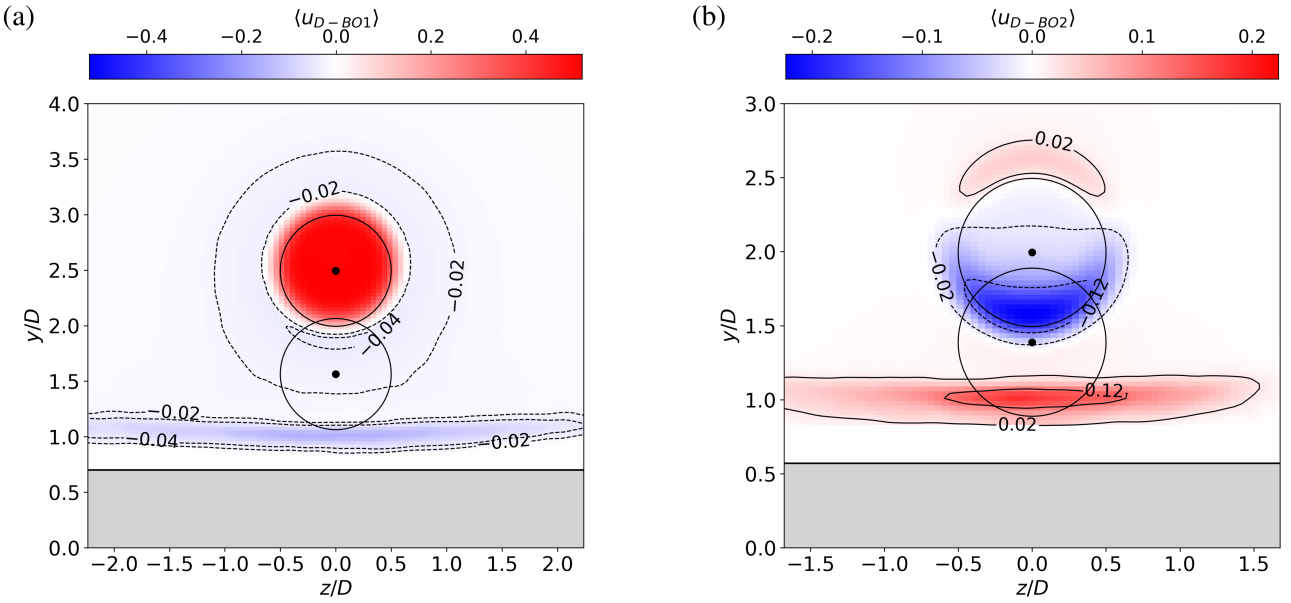


Figure 4.7: Velocity deficit caused by the upstream turbine in (a) BO1 and (b) BO2 evaluated 0.25D upstream of the downstream turbine. Contours are shown with dashed lines and the grey area indicates the hill's section.

To further understand the impact of the tilt angle in BO2 on the total power, we turned to the insights provided by the GP surrogate model constructed at the end of BO2. The 4-dimensional surrogate model can be sliced and analysed in a simplified manner by fixing three of its design variables at the BO2 optimum, x_1, x_2 and h_2, and observing variations solely due to the tilt angle, α_1. The result is a 1-dimensional function, which presents both the mean value predicted by the surrogate model and the accompanying standard deviation, which captures the associated uncertainty. For the purpose of this section, an additional training point has been considered in the surrogate model, corresponding to the BO2 optimum with zero tilt. As depicted in Figure 4.8, introducing a tilt angle of approximately $\alpha_1 \approx 9°$ yields a 4.3% increase in total power in comparison to configurations without tilt. This result is especially noteworthy given that the two turbines are positioned in close proximity to each other, underscoring the

tilt angle's significant potential to enhance total power generation for this case. As indicated by the surrogate model, using negative tilt angles reduces the performance of the turbine pair. This indicates the importance of allowing the downstream turbine to improve its conditions rather than harvesting more power upstream with an alignment to the inclined flow in front of the hill.

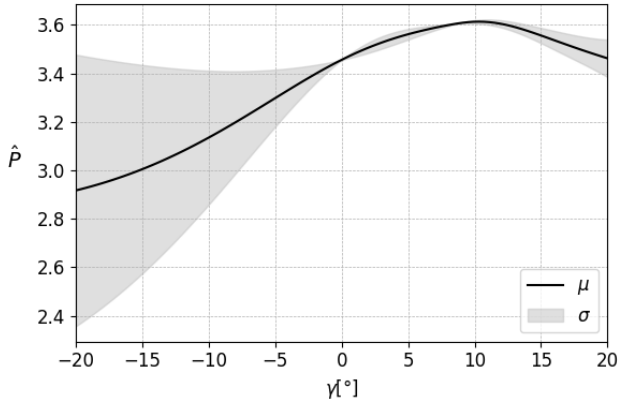


Figure 4.8: Mean and standard deviation power prediction of the GP surrogate model against the tilt angle by fixing x_1, x_2 and h_2 to the BO2 optimum values. The GP surrogate model is built with the 100 evaluations of BO2, plus the BO2 optimum without tilt.

4.2.4 Hill Recirculation Changes

The wind tunnel experiments by Cao and Tamura (2006) denote the hill's natural separation point occurring shortly after the crest in the absence of turbines. Figure 4.6 illustrates how turbine placement substantially influences the flow patterns around the hill in the mid-span plane. The interaction with turbine wakes not only accelerates the flow but also alters the hill's recirculation characteristics when turbines are placed near the separation point. Figure 4.9 provides a 3D representation of the recirculation edge, delineated by the zero-velocity contour and juxtaposed with the actuator discs for each case.

In the Reference case (Figures 4.6a and 4.9a), the upstream turbine's positioning on the hill's summit causes the flow to reattach in a small area of the leeward side of the hill. The introduction of a downstream turbine, as seen in BO1 and BO2, enlarges the reattachment zone owing to the blockage effects exerted by both turbines. However, as evident in Figure 4.7,

changes in the recirculation pattern arise not solely due to the downstream turbine but are also a consequence of the upstream turbine's blockage. It is noteworthy that despite the similarities in turbine placement between BO1 and BO2, the extent of the recirculation edge near the mid-span region varies significantly.

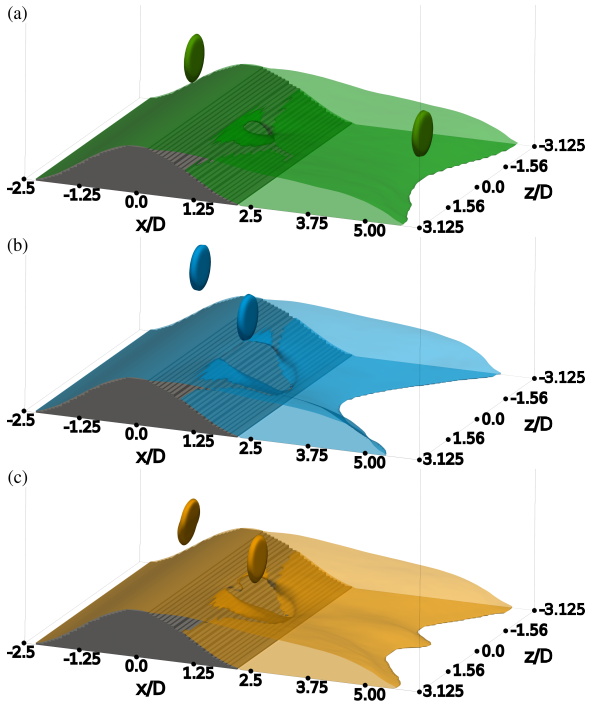


Figure 4.9: Visualisation of the hill recirculation edge depicted by the contour of zero-velocity for the Reference (a), BO1 (b), and BO2 (c) cases.

This complex interaction is critical to understand as it is often not captured by lower-fidelity models such as the superposition method, leading to inaccuracies in power prediction for turbines influenced by flow separation (Liu and Stevens, 2020a). LES becomes indispensable in these scenarios, capturing non-linear effects and providing the necessary fidelity for the optimisation of turbine positioning within such complex terrains with flow separation.

4.3 Conclusions

The purpose of this chapter was to demonstrate the application of the validated LES framework, combined with BO, to a simplified but representative siting problem: two wind turbines positioned around a two-dimensional Gaussian hill. This configuration provided a controlled test

case to assess whether high-fidelity simulations, coupled with BO, can systematically exploit terrain-induced acceleration while mitigating wake losses. The hill and inflow set-up followed the validation work of Chapter 3, with the finer G2 grid adopted to resolve the recirculation region and turbine proximity. Statistical convergence tests confirmed that power signals stabilised within 1% variability, while BO convergence was reached after 56 (BO1) and 40 (BO2) evaluations, establishing the robustness of the optimisation results.

The optimisation studies revealed two distinct strategies for enhancing power extraction. In BO1, with both hub heights and streamwise locations varied, the optimum placed the upstream turbine near the hill centre with an elevated hub height and the downstream turbine only one diameter downstream. This configuration achieved a 73% increase in total power relative to the Reference layout, with the downstream turbine benefiting from both crest speed-up and the accelerated inflow between the hill and the upstream rotor. In BO2, hub height of the upstream turbine was fixed and replaced by tilt angle as a design variable to explore wake steering. The optimum introduced a positive tilt of about 9°, which deflected the wake upwards, reduced upstream output, but improved downstream inflow sufficiently to yield a 34% gain in total power. Both cases highlighted the benefit of close turbine placement, where upstream blockage diverted momentum into the downstream rotor plane and altered the hill's recirculation, generating flow structures not captured by lower-order superposition models.

Taken together, these results underline the value of LES-based optimisation in complex terrain. By resolving non-linear interactions between hill-induced acceleration, turbine wakes, and blockage effects, the framework identified siting strategies inaccessible to conventional engineering tools. At the same time, the study clarified the relative merits of different approaches: hub-height optimisation provided the largest overall gains but may be constrained in practice by structural or cost considerations, while tilt-based steering offered a viable alternative when hub height cannot be varied. Both cases demonstrated that terrain features can be exploited to recover power losses typically associated with wake interactions.

Overall, this chapter confirmed that the LES–BO framework can deliver meaningful siting strategies even in reduced two-turbine scenarios. It provided both a proof-of-concept for the

methodology and physical insight into the mechanisms driving optimal configurations. The next chapter extends the framework to a utility-scale farm in fully three-dimensional terrain, addressing a larger design space and multi-directional inflows representative of more realistic wind-farm scenarios.

Chapter 5

Utility-Scale Wind-Farm Optimisation over Synthetic Terrain

Building on the validation of the LES framework in Chapter 3 and the two-turbine optimisation over a 2D hill in Chapter 4, this chapter scales the methodology to a utility-scale wind farm operating over a synthetically generated three-dimensional terrain. The work presented here is based on Jané-Ippel et al. (2025) accepted for publication and currently available as a preprint. The aim is to test whether high-fidelity LES, coupled with BO, can identify robust siting strategies when the design space expands and the flow physics are shaped by multiple interacting hills, valleys, and ridges. The use of synthetic, analytically defined terrain ensures reproducibility while retaining the essential topographic features that drive speed-up, separation, and terrain-modified wake development.

The case study involves a 16-turbine wind farm constrained within a complex central terrain region, with layouts parametrised by horizontal turbine coordinates and subject to spacing constraints. Several optimisation scenarios are considered to probe spacing effects and directional sensitivity, including single- and multi-directional inflows representative of a typical wind rose. The key novelty of this chapter lies in the use of physics-informed initialisation to guide BO, showing that while BO struggles to discover new high-performing regions independently, it serves as an effective refinement tool once seeded with physically motivated designs.

The chapter is structured as follows. Section 5.1 introduces the terrain, wind-farm configuration, and optimisation framework. Section 5.2.2 presents the optimisation results for single- and multi-directional scenarios, while Section 5.2.3 analyses the flow physics underlying the optimal layouts. Together, these results demonstrate how high-fidelity optimisation can exploit terrain-induced flow features at utility scale, highlight the challenges posed by high-dimensional design spaces, and provide the basis for the final conclusions of this thesis.

5.1 Problem Setup

5.1.1 Site and Terrain Model

To investigate WFLO in realistic complex terrain, we consider a computational domain of dimensions $L_x \times L_y \times L_z = 5.4 \times 1.2 \times 3.6$ km. The coordinate system is oriented with x pointing east, z pointing south, and y vertically upward. The farm is located within a three-dimensional terrain in the central region of the domain.

The complex terrain is constructed by superimposing twelve analytically defined hills using cosine-squared elevation profiles:

$$h(x,z) = \sum_{i=1}^{12} h_{\max,i} \cos^2\!\left(\frac{\pi(x-x_i)}{W_{x,i}}\right) \cos^2\!\left(\frac{\pi(z-z_i)}{W_{z,i}}\right), \qquad |x-x_i| \le \tfrac{W_{x,i}}{2},\ |z-z_i| \le \tfrac{W_{z,i}}{2}, \tag{5.1}$$

where $h_{\max,i}$, $W_{x,i}$, and $W_{z,i}$ are the height and horizontal extents of the i-th hill, and (x_i, z_i) is its centre. The constants have been chosen to ensure smooth hill boundaries for numerically stable LES and to create realistic topographic variability. Several hills have sufficient steepness to induce flow separation, generating complex terrain–wake interactions. The parameters defining all twelve hills are summarised in Table 5.1. Figure 5.2 illustrates the resulting composite terrain with major terrain features A–F (labelled in Figure 5.2). Our parameter choice facilitates reproducibility of the terrain definition without reducing the physical complexity of the problem.

The terrain spans a 3 km × 2 km region and exhibits diverse topographic features commonly

Table 5.1: Characteristics of the individual hills forming the composite terrain.

Hill (i)	x_i (m)	z_i (m)	$h_{\max,i}$ (m)	$W_{x,i}$ (m)	$W_{z,i}$ (m)
1	1303	1551	117	235.1	271.8
2	1467	2002	111	426.3	202.5
3	1722	1867	92	280.6	228.9
4	1893	1468	109	310.0	620.0
5	2025	1994	118	533.3	861.5
6	2205	2209	99	908.3	403.7
7	2426	1601	113	487.5	325.0
8	2649	2126	81	286.4	623.5
9	2909	1751	98	807.1	1130.0
10	3115	2313	103	807.1	491.3
11	3215	1336	114	375.8	302.7
12	3509	1201	96	500.0	818.1

found in real mountainous regions. Hill A represents an isolated upstream hill. Hill B forms an elevated promontory connected to the main ridge. Feature C constitutes the primary ridge spanning the full farm width with varying elevation along its crest. Valley D forms a sheltered depression in the central farm area, and hill E represents the highest elevation feature in the domain. Hill F comprises an isolated downstream hill.

5.1.2 Atmospheric Boundary Layer Inflow

The wind farm operates within a neutral ABL representative of typical onshore conditions, characterised by $u_* = 0.4$ m/s, $\delta = 1200$ m, and $y_0 = 0.1$ m (Kelly et al., 2014; Wiernga, 1993). Under these conditions, the hub-height velocity (at 100 m above flat terrain) is approximately 6.8 m/s with streamwise turbulence intensity of 12%.

The inflow profiles are generated through a precursor simulation with flat terrain and the same domain dimensions, following the approach detailed in Section 2.1.4. Figure 5.1 presents the horizontally- and temporally-averaged streamwise velocity and streamwise turbulence intensity profiles. The mean velocity profile shows excellent agreement with the theoretical logarithmic law, consistent with the precursor validations in Chapter 3.

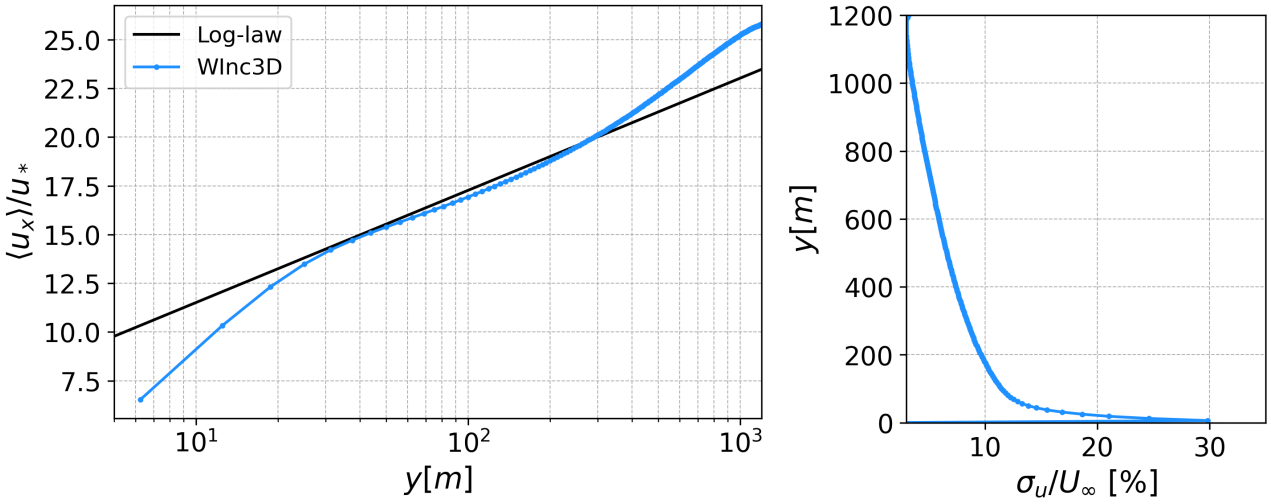


Figure 5.1: Horizontally- and time-averaged streamwise velocity (left) and streamwise turbulence intensity profile (right) of the ABL in the precursor simulation of the utility-scale wind farm, compared against analytical profiles.

5.1.3 Design Space, Constraints and Optimisation Cases

The wind farm consists of 16 identical turbines, each with rotor diameter $D = 100$ m and hub height equal to D, and operate with a local thrust coefficient of $C'_T = 4/3$. The hub height is defined locally relative to the terrain elevation, such that turbine hubs are positioned at a vertical distance D above the local ground. Turbines are constrained to be placed within a 1.2 km $\times$ 2.4 km rectangular region ($12D \times 24D$) positioned over the central terrain, as shown in Figure 5.2. This placement region encompasses the most topographically complex area of the farm.

The turbine layout is parameterised by 32 design variables representing the horizontal coordinates (x_i, z_i) of each turbine $i = 1, \ldots, 16$, as illustrated by the markers in Figure 5.2. Following the approach of Bempedelis et al. (2024), a spanwise ordering constraint is imposed such that $z_i > z_{i-1}$ for $i = 2, \ldots, 16$ to reduce dimensionality and eliminate redundant permutations (see Section 5.1.5 for detailed explanation).

To assess the influence of inflow direction and spacing constraints, four layout optimisation cases are examined, where 1D and 2D refer to minimum turbine spacing of 1 and 2 rotor diameters, respectively:

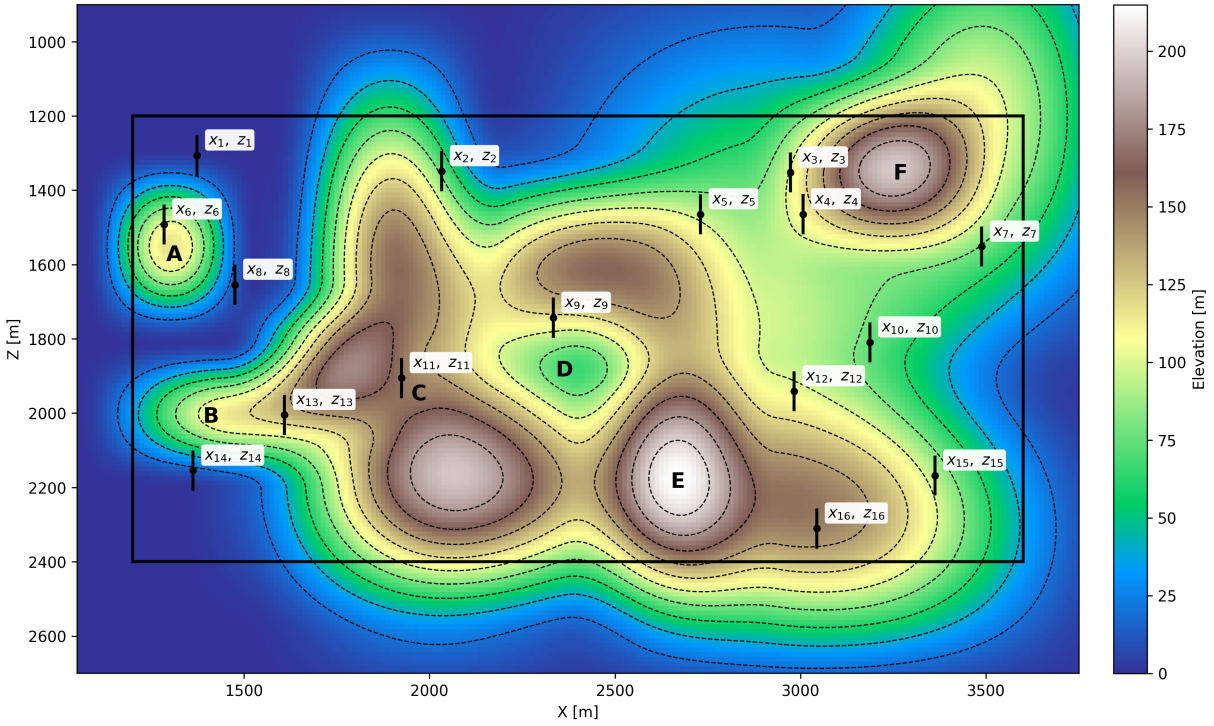


Figure 5.2: Synthetic complex terrain elevation map with example turbine layout showing the 32 design variables. Labels (A–F) correspond to major terrain features. The dashed lines represent 20-m contours of the terrain elevation. The black rectangle is the available land to place turbine centres.

1. **WFLO-1D-270°:** Single wind direction at 270° with 1D spacing constraint.

2. **WFLO-2D-270°:** Single wind direction at 270° with 2D spacing constraint.

3. **WFLO-2D-210°:** Single wind direction at 210° with 2D spacing constraint.

4. **WFLO-2D-MWD:** Multi-directional inflow (300°, 270°, 240°, 210°) with weighted objective and 2D spacing constraint.

The 1D spacing constraint enables higher turbine density, while the 2D constraint represents more realistic operational spacing requirements. Note that this constraint refers solely to the minimum distance between turbines and not to the dimensionality of the model, as turbine wakes are resolved using the ADM within the LES. The multi-directional case reflects practical wind rose considerations where wind farms must perform well across multiple predominant wind directions, with the selected range representing typical westerly to southwesterly winds.

The objective function maximises total wind farm power output over the specified wind direc-

tions:

$$f(\mathbf{c}) = \sum_{k=1}^{K} w_k \sum_{n=1}^{N} P_{n,k}(\mathbf{c}), \tag{5.2}$$

subject to:

$$\mathbf{c} \in \mathcal{X} \subset \mathbb{R}^{2N}, \tag{5.3}$$

$$\|\mathbf{c}_i - \mathbf{c}_j\| > \alpha D, \quad \text{for all } i \neq j, \tag{5.4}$$

$$z_i > z_{i-1}, \quad \text{for } i = 2, \ldots, N, \tag{5.5}$$

where $\mathbf{c} = (x_1, \ldots, x_N, z_1, \ldots, z_N)^T$ represents turbine positions, with individual turbine centre position vectors defined as $\mathbf{c}_i = (x_i, z_i) \in \mathbb{R}^2$, $P_{n,k}(\mathbf{c})$ is the LES-estimated power of turbine n under wind direction k, and w_k are directional weights satisfying $\sum_{k=1}^{K} w_k = 1$, and α is the spacing constraint parameter (equal to 1 or 2). The number of wind directions for the multi-directional case is determined by $K = 4$. For the multi-directional optimisation, the directional weights are $w_{300^\circ} = 0.2$, $w_{270^\circ} = 0.3$, $w_{240^\circ} = 0.26$, and $w_{210^\circ} = 0.24$, reflecting the relative frequency of these wind directions. These constraints ensure that turbines are positioned within the feasible region $\mathcal{X}$ and maintain minimum separation between turbine centres, with the ordering constraint eliminating redundant permutations of turbine layouts.

Power outputs are normalised for comparison across configurations and wind conditions. We define the normalised individual turbine power, $\hat{P}_T$, as the ratio of a turbine's power to that of a single turbine operating on flat terrain under identical inflow where no wake-to-wake interactions or blockage effects are present. The normalised total farm power, $\hat{P}_F$, is the ratio of the farm's actual power output to 16 times this single-turbine reference power. This normalisation approach allows assessment of terrain-induced performance enhancements, where values greater than unity indicate power gains relative to flat terrain expectations.

5.1.4 Simulation Configuration

As introduced in the validation of the 3D hill in Section 3.1.2, the reference grid H1 is chosen to be used as reference here. The synthetic terrain comprises multiple hills with elevations ranging from 117 m (hill A) to 215 m (hill E). To establish consistent grid resolution across the domain, a representative height of 150 m is selected, corresponding approximately to the mean elevation of the terrain features. Scaling the validated 24-node discretisation to this characteristic height yields a vertical resolution of $\Delta y = 150/24 = 6.25$ m. Maintaining an $AR = 2$ results in horizontal resolutions of $\Delta x = \Delta z = 12.5$ m, ensuring adequate resolution of terrain-induced flow features whilst maintaining computational feasibility for the extensive layout optimisation framework. This resolution ensures 16 nodes across the diameter in the vertical direction and 8 in the spanwise direction, providing sufficient resolution to the ADM.

The time step used is $\Delta t = 0.2$ s, ensuring that the maximum Courant-Friedrichs-Lewy (CFL) number remains below 0.2. The simulation data are averaged over a physical time of 3.5 hours after discarding 15 minutes of initial transient, which is found sufficient for the flow to develop and obtain the wind farm power average. For cases with different wind directions, the terrain is rotated about its centre to align the inflow with the desired direction while preserving the domain geometry. Each direction is simulated independently, with the turbines facing the inlet. Boundary conditions are applied as in the operational wind farm validation case in Section 3.2.

All simulations were executed on the ARCHER2 UK National Supercomputing Service, utilising 2048 cores per LES run. Each LES used dual AMD EPYC™ 7742 64-core processors at 2.25 GHz with a wall-clock time of approximately 1.3 hours per simulation. The three single-direction optimisations (512 evaluations each) required ≈ 1.36 million CPU-h, and the multi-directional optimisation (4×512 evaluations) used ≈ 5.45 million CPU-h. The computational demands of the BO algorithm itself are negligible relative to the LES cost, as the surrogate model fitting and acquisition function maximisation require minimal wall-clock time in comparison to simulation runtime.

5.1.5 Optimisation Configuration

The optimisation aims to find the turbine layout $\mathbf{c}^* \in \mathcal{X} \subset \mathbb{R}^{2N}$ that maximises the wind farm power objective $f(\mathbf{c})$, as introduced in Section 5.1.3, using the methodology presented in Chapter 2. Here, we use an isotropic Matérn 5/2 kernel, unlike in Chapter 4, where a separate lengthscale per design variable was employed. This choice is necessary due to the increased number of design variables and the limited evaluation budget. To guide the optimisation, we employ the LCB acquisition function, beginning with $w = 2$ to encourage exploration and later reducing it to $w = 1$ to focus on exploitation (see Section 5.2.1). The jitter term is set to $\epsilon = 0.01$. Each optimisation involves a total of 512 layout evaluations, determined by the available computing resources and the optimisation evolution discussed in Section 5.2.1. To maximise parallel efficiency, BO is performed in batches of eight designs per iteration.

The optimisation is subject to the turbine spacing and domain constraints previously described in Section 5.1.3. In addition, an ordering constraint is imposed along the spanwise direction such that $z_i > z_{i-1}$ for $i = 2, \ldots, N$ (with $N = 16$ turbines in this case), following the method proposed by Bempedelis et al. (2024). This constraint addresses a fundamental challenge in WFLO. Since all turbines are identical, the objective function is invariant under permutations of turbine labelling, meaning that $N!$ different combinations of design variables produce identical outputs. Rather than augmenting the training dataset with all permutations (which scales as $O(m^3)$ with dataset size m), the ordering constraint reduces the design space by a factor of $1/N!$ and maintains solution equivalence. This approach makes the problem more tractable for GP regression and acquisition function optimisation without altering the nature of the solution space, thereby improving optimisation efficiency without loss of generality.

The BO process is initialised with a DoE composed of Latin Hypercube Sampling (LHS) and a set of informed layouts derived from terrain base flow features (detailed in Section 5.2.1). After each batch of LES evaluations, the GP surrogate is updated with the new data and the process is repeated until the maximum number of evaluations (the computational budget) is reached.

5.2 Results

5.2.1 Physics-Informed BO: Motivation and Performance

This section evaluates the limitations of uninformed BO for WFLO in complex terrain and presents a physics-guided initial sampling strategy to improve optimisation by embedding engineering knowledge of flow patterns. We quantify the performance of this approach across four WFLO cases, analysing the resulting optimisation evolution and computational demands.

In high-dimensional WFLO problems, standard BO with LHS initial sampling may require prohibitively many evaluations to discover high-performing designs. As Bempedelis et al. (2024) showed, low-fidelity models can outperform high-fidelity BO frameworks when allowed to explore more layouts due to their lower computational cost.

We evaluated standard BO performance using a baseline optimisation with only LHS for the WFLO-2D-270° case. As shown in Figure 5.3, power gains after 512 LES evaluations were modest, with a 10% improvement relative to the initial samples ($\hat{P}_F = 1.1$ after 480 additional evaluations beyond the initial 32 LHS samples). The resulting layout showed suboptimal turbine placement, confirming that standard BO failed to identify high-performing configurations within a practical evaluation budget. We systematically tested the BO framework under various configurations, including different acquisition functions (such as Expected Improvement), exploration–exploitation weights, single versus dimension-specific length-scale settings in the kernel, and variations in anchor point strategies, consistently observing similar optimisation behaviour. These results highlight how the curse of dimensionality and the absence of physics-informed guidance during early sampling hinder the discovery of optimal regions in the design space.

Given the well-defined structure of the velocity field in complex terrain, we hypothesised that incorporating engineering knowledge about flow patterns could offer more informative guidance to the optimisation process. To test this hypothesis, we used the hub-height velocity field from the LES base flow as a reference for generating informed layout candidates. The base flow

captures the terrain-induced flow acceleration, flow separation, and wake behaviour specific to the complex topography, offering valuable physical insight to guide the sampling process.

We then developed four sampling strategies using terrain flow knowledge (detailed in Section 5.2.1). Our results show that the best physics-informed layout achieved a 50% power improvement ($\hat{P}_F = 1.5$), vastly outperforming the 480-evaluation BO process. This performance demonstrates that uninformed BO is not effective in complex terrain for the given computational budget. Based on these findings, we develop an informed BO approach incorporating physics-guided layouts into the initial sampling. Instead of using 32 LHS samples alone, we combine 32 LHS layouts with 32 physics-informed layouts to form the initial DoE, totalling 64 layouts per optimisation.

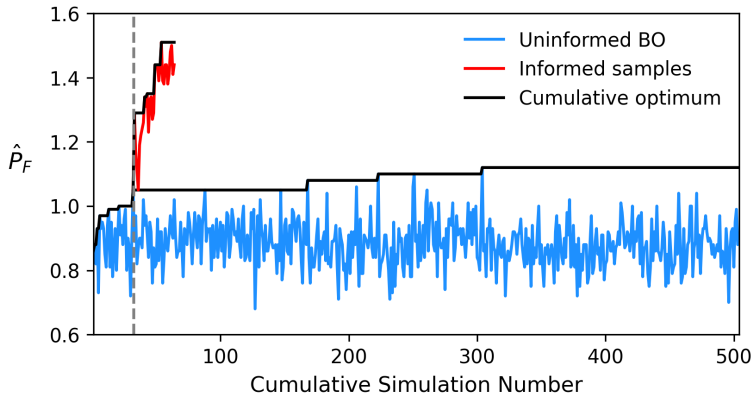


Figure 5.3: Normalised power output versus cumulative simulation number for WFLO-2D-270°, comparing uninformed BO with physics-informed cases. The first 32 LHS cases are marked by the vertical grey dashed line.

Physics-Informed Initial Sampling Strategies

This section describes the four physics-guided sampling strategies developed to generate informed initial layouts, with each strategy capturing a different aspect of flow-turbine interaction. Each strategy produces eight distinct layouts, yielding 32 informed samples in total.

The strategies are applied to all optimisation cases, considering relevant spacing constraints and inflow directions. For multi-directional cases, two layouts per wind direction are generated under each strategy, maintaining the total of eight layouts per strategy. Figure 5.4 shows an

example layout from each strategy overlaid on the LES time-averaged base-flow velocity field at hub height.

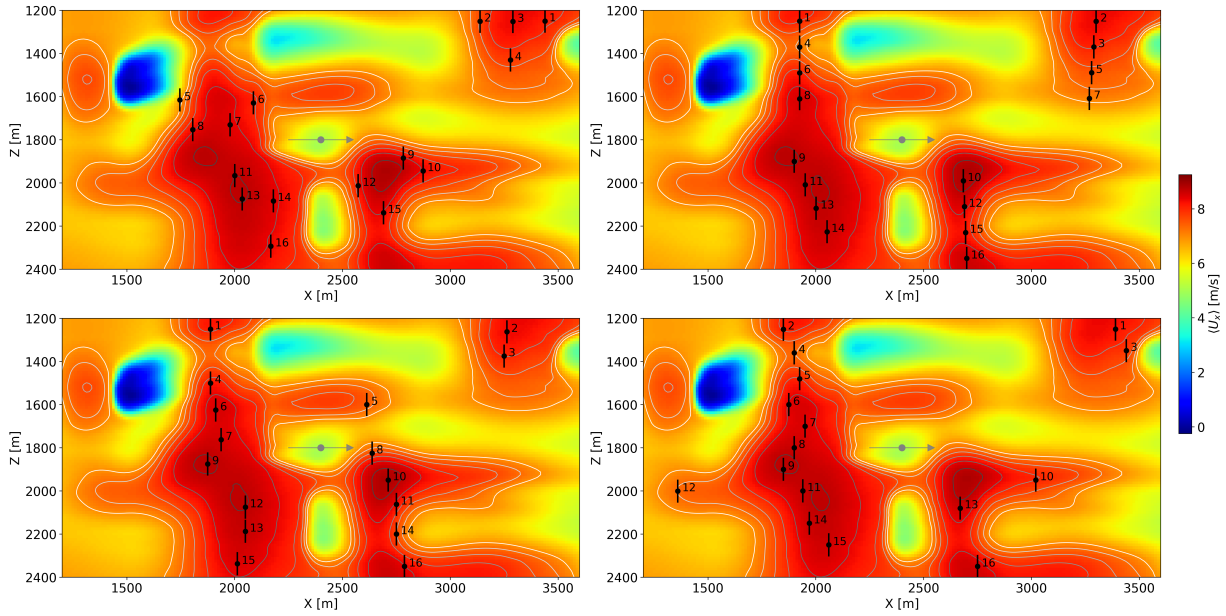


Figure 5.4: Farm layouts from each of the four sampling strategies, overlaid on the time-averaged hub-height velocity field from the base flow (no turbines). Top left: Cluster-Based Sampling; top right: Ridge-Based Layouts; bottom left: Velocity-Ordered Zone Sampling; bottom right: Physics-Informed Custom Layouts. Contours indicate increasing velocity regions.

Strategy 1: Cluster-Based Sampling in High-Velocity Regions This strategy targets hub-height regions with the highest velocities, which are unlikely to be sufficiently covered by LHS. K-means clustering extracts multiple high-velocity regions from the base-flow field, and turbines are randomly distributed within these clusters with spacing constraints enforced. Variations in the number of clusters, turbines per cluster, and random placement within clusters provide broad exploration within high-speed zones.

Strategy 2: Ridge-Based Layouts This strategy places turbines in linear chains along terrain-aligned ridges, where flow accelerates and wake interference is minimised. By varying the number of turbines per ridge and spanwise spacing, we capture multiple levels of ridge exploitation. Unlike the cluster-based strategy, sufficient streamwise spacing between turbines is maintained to mitigate immediate wake overlap.

Strategy 3: Velocity-Ordered Zone Sampling considering Wake Recovery This strategy combines high-velocity placement with strict spacing constraints to ensure wake recovery, defined as 7D in the streamwise direction. Potential turbine positions are ranked by hub-height velocity from the base flow, and turbines are iteratively placed at the highest-velocity locations. After each placement, a blocking mask excludes all remaining candidate points within a circle (spacing constraint) and a 7D streamwise rectangle, ensuring no immediate wake interference.

To introduce layout diversity, the domain is partitioned into horizontal spanwise zones of equal z-extent. Each layout configuration uses a different number of zones (e.g., 1, 2, 4, or more), with turbines distributed across them. Within each zone, points are re-ranked by velocity, and the same 7D exclusion logic is applied. This zonal division generates multiple layouts that balance turbine concentration and spatial distribution.

Strategy 4: Physics-Informed Custom Layouts This strategy explores secondary high-velocity regions not captured by the previous methods, guided by flow understanding and expert judgement. Turbines are placed in areas where terrain features are expected to induce beneficial effects such as flow acceleration or wake deflection. Terrain-driven pressure gradients are leveraged to enhance wake recovery, and flow blockage is exploited to increase velocities. These designs aim to capture complex flow-turbine interactions that are difficult to encode in automated strategies, providing valuable diversity and physically grounded high-performance candidates.

Optimisation Evolution

Each optimisation used a budget of 512 function evaluations. To assess BO performance, we analyse the evolution of the normalised power deficit ΔP across cumulative simulations for each case. Figure 5.5 shows the progression of the four optimisations, with blue markers indicating the power deficit of each evaluated layout and the black line tracking the best layout found up to that point.

Each optimisation is divided into three distinct phases:

- **Initial Sampling (64 evaluations):** The first 64 layouts comprise the initial DoE from Section 5.2.1. The informed samples consistently outperform the LHS across all cases.

- **Exploration Phase (128 evaluations):** The BO proceeds with a LCB acquisition function using a high exploration weight ($w = 2$). This promotes sampling in uncertain regions of the design space. During this phase, new evaluations yield power levels comparable to the LHS samples, indicating that the optimiser is probing less promising or uncharted regions far from the high-performing layouts seeded by the informed strategies.

- **Exploitation Phase (320 evaluations):** After 128 BO evaluations, the LCB weight is reduced to $w = 1$ to favour exploitation. This results in an improvement in power output across all cases, as the optimiser concentrates its search around the promising regions discovered in the informed layouts.

The adaptive weighting strategy proved effective across all optimisations. After switching from the exploration weight to the exploitation weight, we observe a consistent performance improvement, with the optimiser identifying layouts with power outputs comparable to the informed samples. In each case, the exploration phase samples many layouts with power levels similar to the LHS, reducing uncertainty in unexplored regions of the surrogate model. Following the transition to exploitation, the optimiser focuses on areas around the high-performing informed layouts, discovering new configurations with similar or slightly higher power outputs. Over time, the frequency of testing high-performing layouts decreases, and the optimiser reverts to evaluating designs with performance levels between those found during the exploration and initial sampling phases. This behaviour suggests that the change in weighting enables targeted exploitation around informed regions initially, followed by a localised exploration phase that reduces surrogate uncertainty near promising areas and tests high-quality layouts. This adaptive behaviour underscores the practical value of using physics-informed initialisation to guide BO in high-dimensional WFLO problems.

These results confirm that BO struggles to discover high-performing regions within a reasonable number of evaluations but serves effectively as a refinement tool once the search space is seeded with physics-informed designs. BO can exploit and marginally improve upon these initial designs, but further evaluations beyond the 512 budget do not yield better layouts. In the WFLO-2D-270° case, an additional 80 evaluations were performed without observing any improvement, with the optimiser continuing the exploration but failing to identify superior configurations. The lack of improvement indicates a stagnation of the optimisation process, with no clear evidence of convergence towards better solutions. This highlights the need to couple domain knowledge with statistical learning for WFLO in complex terrain, where high dimensionality and non-linear interactions challenge exploration-based methods.

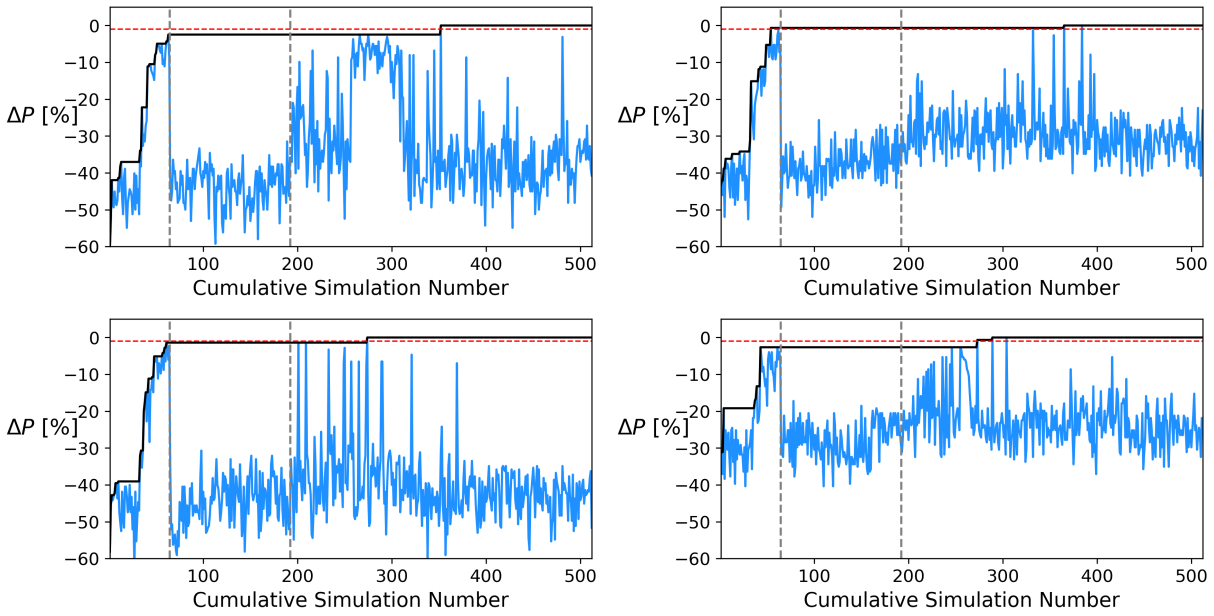


Figure 5.5: Deficit from the optimum power as a function of the cumulative simulation number for (a) WFLO-1D-270°, (b) WFLO-2D-270°, (c) WFLO-2D-210°, and (d) WFLO-2D-MWD. Blue points show the deficit of each evaluated layout relative to the best case; the black line tracks the best layout found so far. Vertical dashed lines mark the end of the initial sampling phase (64 layouts) and the transition from exploration to exploitation (192 layouts). The red dashed line indicates a 1% power deficit.

The WFLO-2D-270° and WFLO-2D-210° cases, despite using the same topography, effectively represent different terrains due to wind direction changes that alter flow patterns. The consistent BO behaviour observed across these rotated cases reinforces the generality of the challenges in complex terrain and supports the robustness of our conclusions.

5.2.2 Optimal Layouts over the Synthetic Multi-Hill Topography

This section presents the optimal wind farm layouts obtained from the four WFLO cases, analysing how terrain complexity, turbine spacing constraints, and wind directionality affect layout topologies and power output.

Figure 5.6 shows the optimal layouts overlaid on the terrain elevation map with wind directions indicated. The WFLO-1D-270° layout has turbines densely aligned along the central ridge, maximising use of high-elevation zones across the farm span. In contrast, the WFLO-2D-270° layout adopts a staggered configuration to satisfy two-diameter spacing while still clustering turbines along the ridge without immediate wake interactions. Both layouts place turbines downstream on hill F and valley D, but only in WFLO-2D-270° the use of hill E is required due to spacing constraints. The WFLO-2D-210° case reorients the staggered layout perpendicular to the inflow, with a broader lateral spread using hills A, B, parts of C, and D, avoiding the downstream hill F. The WFLO-2D-MWD layout presents a more balanced distribution across the domain, distributing turbines across hills A, B, and a single row along ridge C. Due to the trade-offs inherent in accommodating multiple wind directions, we see more turbines in the hill E area, as used in the 210° case.

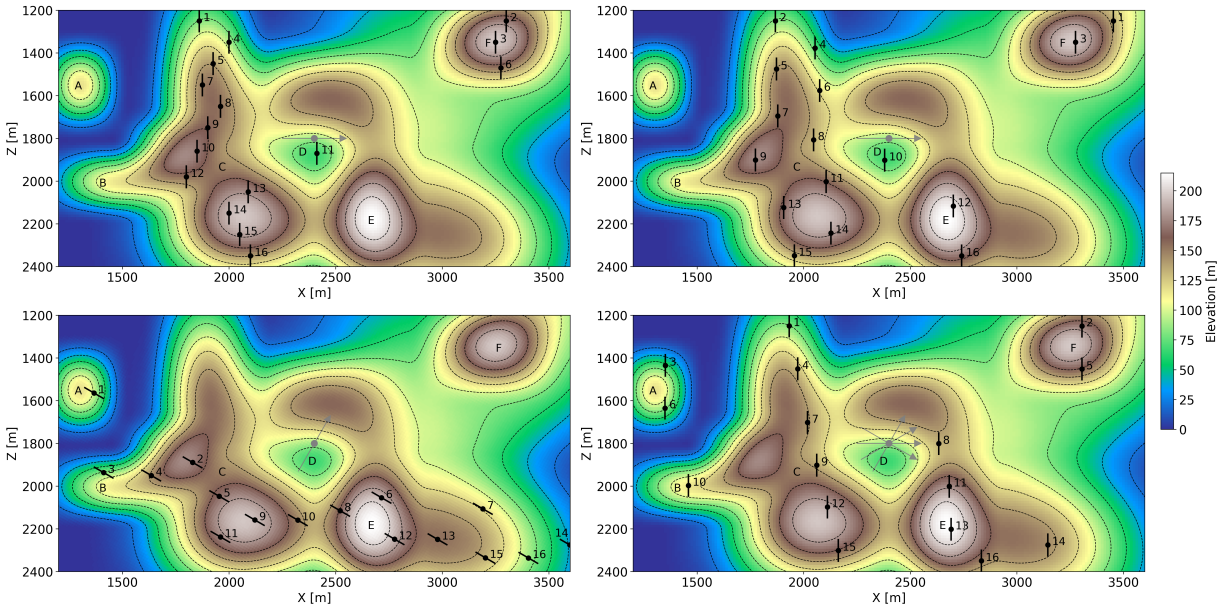


Figure 5.6: Optimal layouts for the four WFLO cases overlaid on terrain height. Each layout was obtained after 512 LES evaluations using physics-guided BO.

Table 5.2 summarises the normalised power outputs for the optimal layouts. All values exceed 1.0, indicating effective exploitation of terrain-induced flow acceleration and/or wake interactions to increase power relative to a single turbine on flat terrain under the same inflow conditions.

Table 5.2: Normalised power outputs for the optimal layouts across WFLO cases. For WFLO-2D-MWD, individual inflow angles are also reported.

	1D-270°	**2D-270°**	**2D-210°**	**2D-MWD**	300°	270°	240°	210°
$\hat{P}_F$	**1.62**	**1.52**	**2.15**	**1.51**	1.35	1.36	1.63	1.72

The WFLO-1D-270° layout achieves higher power than the WFLO-2D-270° layout due to denser turbine placement, benefiting from local flow acceleration and possible inter-turbine blockage effects (see Section 5.2.2). However, such compact layouts may face practical limitations in real deployments due to structural loading, access, and land use constraints.

Among all cases, WFLO-2D-210° delivers the highest power output, benefiting from the ability to span largely perpendicular to the inflow to minimise turbine wake interactions. The WFLO-2D-MWD case achieves a weighted-average performance similar to WFLO-2D-270°, though individual directional evaluations are lower than single-direction optimised layouts, reflecting the inherent trade-offs in achieving robust performance across varying wind conditions.

Single-Direction Optimisations: Performance Analysis

This section compares the performance of optimal layouts from single wind direction optimisations. We analyse individual turbine power trends and the underlying flow mechanisms using hub-height velocity and turbulence intensity fields.

Figure 5.7 shows the normalised power output for each turbine across the three optimal layouts, ranked from highest to lowest. Figure 5.8 shows the corresponding hub-height flow fields, with time averaged streamwise velocity (U_x) and turbulence intensity (σ_u/U_∞) overlaid with turbine locations and terrain contours. Here, U_∞ denotes the reference velocity at the top of the atmospheric boundary layer. These visualisations illustrate how turbine placement and terrain features drive the observed performance differences.

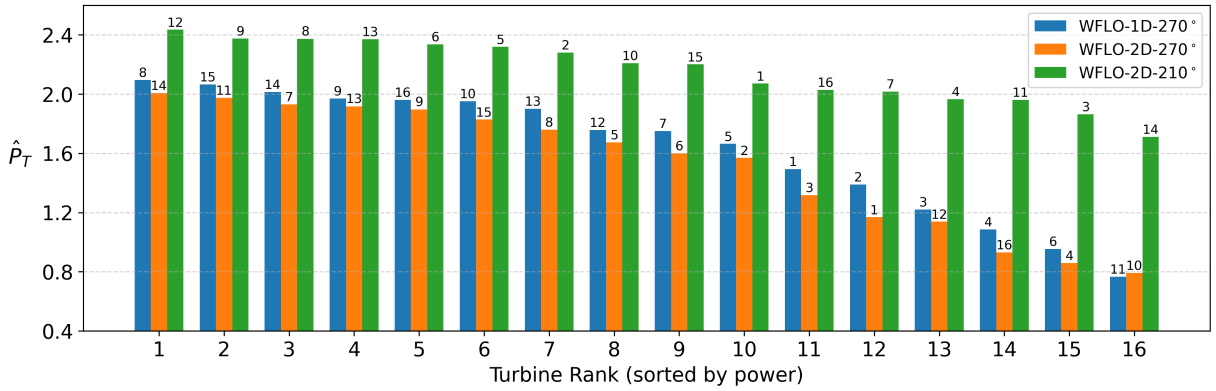


Figure 5.7: Turbine-by-turbine normalised power output ($\hat{P}_T$) for the three optimal layouts: WFLO-1D-270°, WFLO-2D-270°, and WFLO-2D-210°. Turbines are sorted by power from highest to lowest for each layout, with original turbine IDs annotated above the bars.

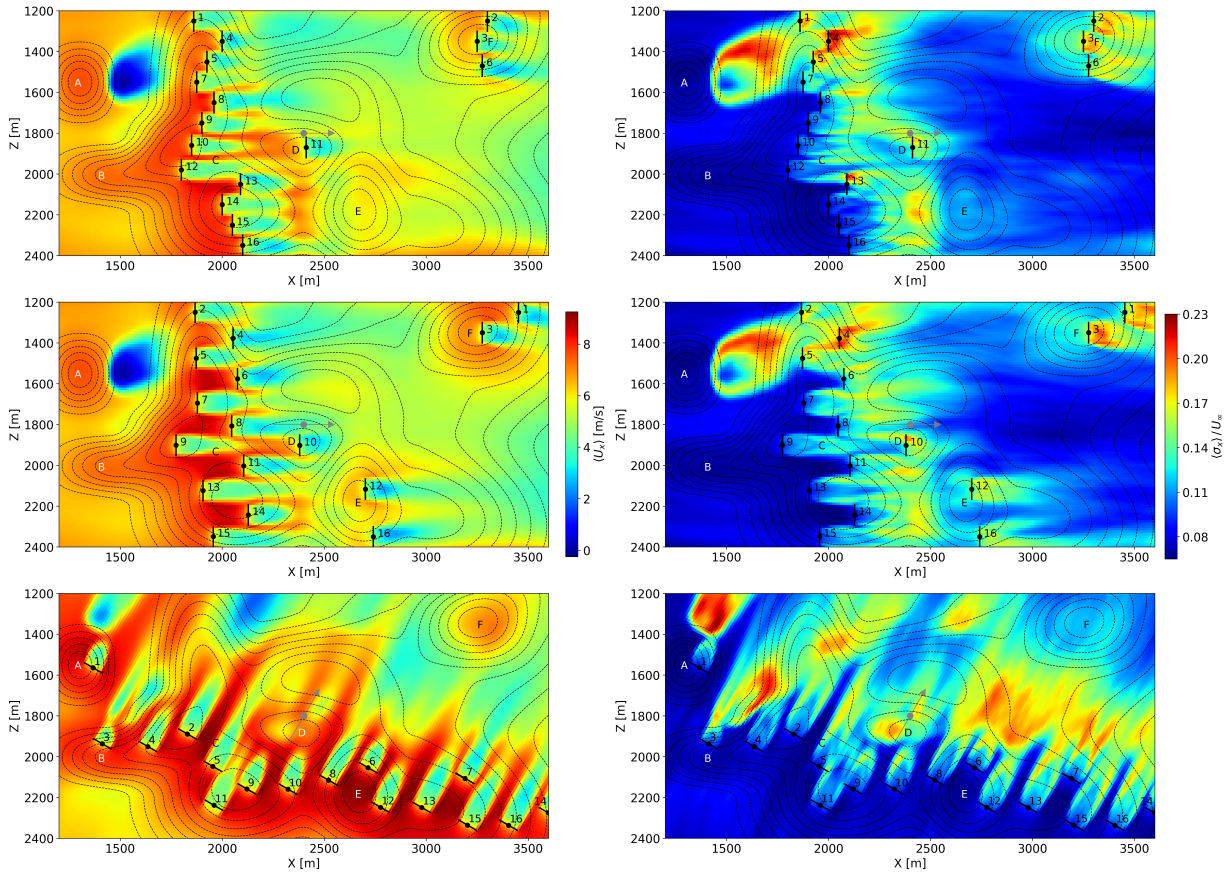


Figure 5.8: Time-averaged hub-height flow fields for the optimal single-direction layouts. Left column: streamwise velocity. Right column: turbulence intensity. Top to bottom: WFLO-1D-270°, WFLO-2D-270°, WFLO-2D-210°. Terrain contours (dashed lines) are shown at 20-m intervals. Flow fields are sampled at turbine hub height, defined locally as a constant vertical distance of D above the terrain.

We first compare the 270° single-direction cases to assess the impact of 1D versus 2D spacing constraints, before discussing the 210° case.

Examining Figure 5.7, WFLO-1D-270° consistently outperforms WFLO-2D-270° due to denser turbine placement along the ridge, benefiting from local flow acceleration and inter-turbine blockage effects. This performance differential suggests that the 2D spacing constraint prevents fully optimal placement but still enables effective ridge exploitation with a staggered configuration. The hub-height velocity fields (Figure 5.8) confirm that turbines are primarily placed in high-velocity regions and avoid immediate wake interactions, ensuring access to stronger inflow and reducing losses. In both 1D and 2D layouts, turbine placement also redirects flow towards downstream units, supporting their performance by channelling accelerated flow near rotor edges (Jané-Ippel et al., 2023; Bempedelis et al., 2024). Turbulence intensity levels are higher downstream of turbines and terrain features, marking regions of wake exposure where velocity deficits reduce power output. Elevated turbulence is particularly evident downstream of hill A, across hill F, and in the valley between ridge C and hill E, indicating separation and wake-affected regions that limit turbine performance. We also observe turbine wake-to-wake interactions at turbines 4 and 13 in WFLO-1D-270°, and turbines 4 and 10 in WFLO-2D-270°, where combined wakes from upstream turbines and terrain features lead to significant inflow velocity deficits.

Analysis of underperforming turbines (with $\hat{P}_T < 1.2$) shows these are typically located in lower terrain areas or exposed to upstream wakes. In both 1D and 2D layouts, turbines in valley D (e.g., turbines 11 for 1D and 10 for 2D) underperform due to reduced inflow velocities associated with wake interactions and terrain-induced flow, which are accompanied by elevated turbulence levels. Despite ridge C being generally high-performing, a single turbine positioned along its lower section (turbine 4 for 1D and 2D) underperforms in both layouts due to partial wake exposure. In the 2D layout, additional underperforming turbines are found on hill E (e.g., turbines 12 and 16), where upstream wake interactions reduce inflow quality. Turbines near hill F further illustrate this trend, with high turbulence intensity acting as an indicator of wake-affected regions.

In contrast, WFLO-2D-210° achieves consistently higher power outputs across all turbines (at least $\hat{P}_T < 1.6$) due to its perpendicular alignment to inflow, avoiding immediate turbine wake interactions and exploiting terrain-induced acceleration over hills A, B, C, and D. The velocity field show stronger flow speeds over terrain in this case compared to previous layouts, demonstrating visibly more intense high-velocity regions and confirming the absence of direct wake overlap between turbines. The associated turbulence intensity is lower and more uniform upstream of the turbines, supporting stable and high power generation under these conditions.

Multi-Directional Optimisation: Balancing Performance and Robustness

This section examines the WFLO-2D-MWD case, where the optimisation objective aggregates power across four inflow directions (300°, 270°, 240°, 210°) with directional weights, reflecting realistic multi-directional operation.

Figure 5.9 shows the turbine power outputs with directional variability, with Figures 5.10 and 5.11 presenting time-averaged hub-height velocity and turbulence intensity fields under each wind direction.

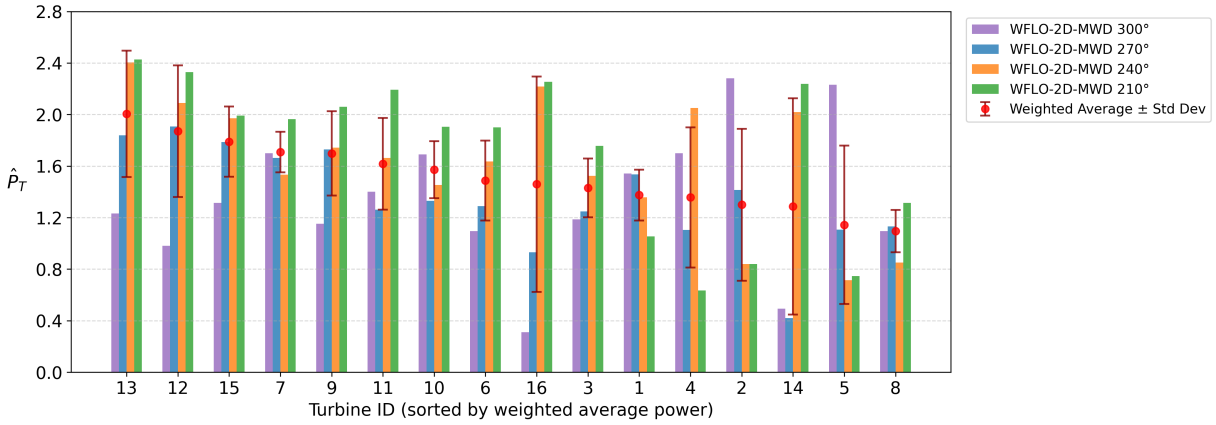


Figure 5.9: Power output for each turbine across four wind directions of the optimal layout for WFLO-2D-MWD. Turbines are sorted by weighted average power, with error bars indicating directional sensitivity.

Examining Figure 5.9, turbines 13, 12, and 15 (located on ridge C and hill E) achieve the highest power outputs under 270°-210° winds. As seen in the velocity fields (Figure 5.10), these directions provide strong inflow with minimal wake overlap, consistent with the single-direction

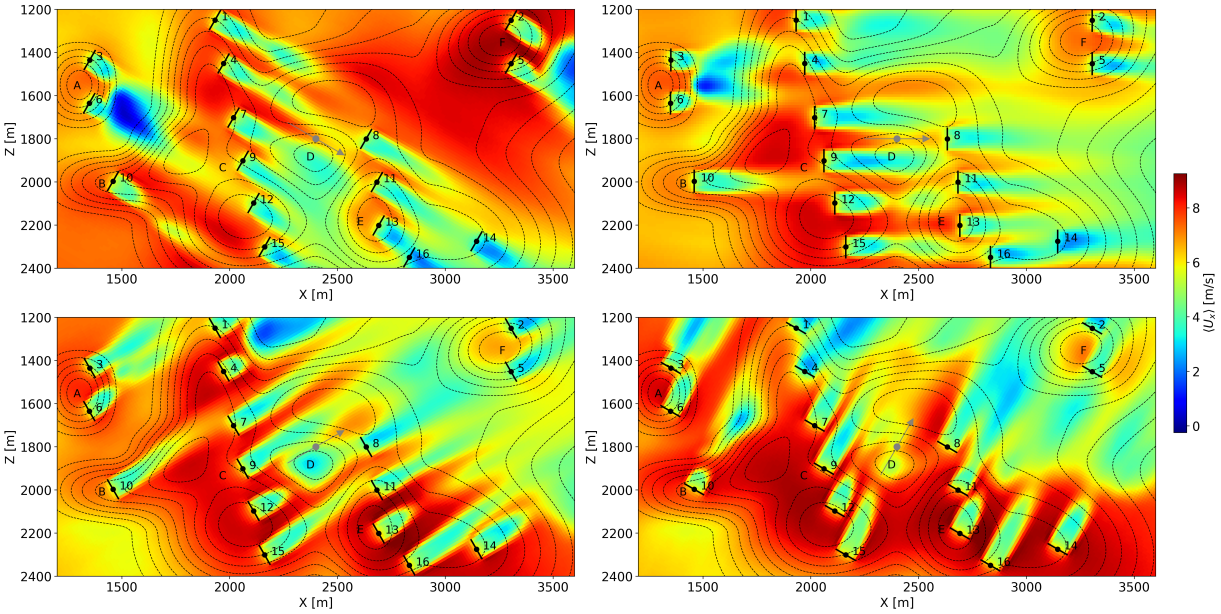


Figure 5.10: Time-averaged hub-height streamwise velocity fields for the optimal multi-directional layout under each wind direction. Top row: 300° (left) and 270° (right). Bottom row: 240° (left) and 210° (right). Terrain contours (dashed lines) are shown at 20-m intervals.

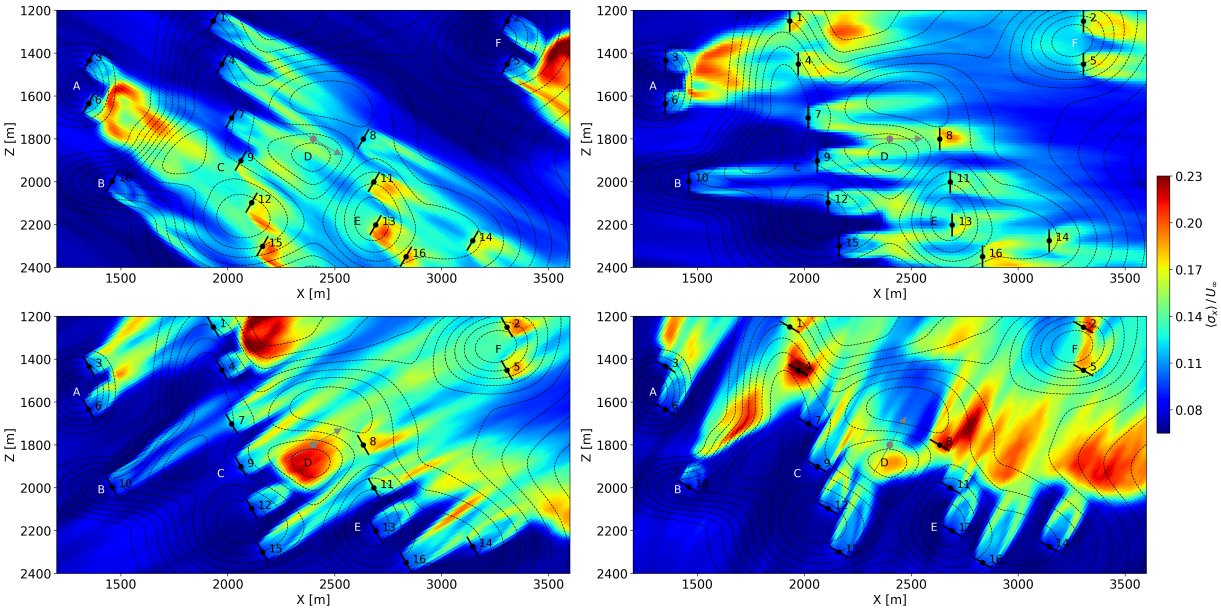


Figure 5.11: Time-averaged hub-height streamwise turbulence intensity fields for the optimal multi-directional layout under each wind direction. Terrain contours (dashed lines) are shown at 20-m intervals.

optimisation findings. Under 300° winds, however, these turbines show reduced performance due to their downstream position relative to the inflow, as visible in the same figure.

In contrast, turbine 8, located centrally near valley D, is the worst-performing unit across all wind directions. This persistent underperformance is due to its susceptibility to upstream wake exposure from multiple directions, as evident in the velocity fields where lower inflow speeds and upstream wakes can be seen around its location.

The directional sensitivity, represented by the error bars in Figure 5.9, reveals distinct patterns of performance variability across the wind farm. Turbines with high sensitivity to wind direction, such as turbine 16 ($\sigma = 0.84$, ranging from 0.31 to 2.25) and turbines 14 and 5 due to their corner positioning, illustrate the challenge of accommodating turbines for varying wind directions. In contrast, robust performers like turbines 7 ($\sigma = 0.16$) and 1 ($\sigma = 0.20$) maintain consistent output across all directions. This indicates that certain terrain locations inherently provide lower variability and more reliable performance, insights that are difficult to predict without high-fidelity flow modelling given the complex interplay between terrain features, wakes, and directional wind patterns.

The MWD layout uses a dispersed turbine distribution to minimise wake alignment across the wind rose while maintaining access to accelerated flow regions. Velocity fields (Figure 5.10) confirm that turbines are positioned to exploit favourable terrain-induced acceleration under at least some wind directions without incurring severe wake losses in others. Turbulence intensity fields (Figure 5.11) show that high levels consistently develop downstream of the first rows of turbines under all wind directions, indicating that downstream turbines are exposed to combined wakes from upstream turbines and terrain features. This supports the observed reduction in power for downstream turbines, which varies with wind direction as different turbines become downstream under different inflow angles.

The multi-directional optimisation provides valuable insights for practical wind farm design in complex terrain. The results demonstrate that optimal layouts should prioritise turbine placement on elevated terrain features (ridges and hilltops) while ensuring adequate separation to accommodate the full wind rose without wake interactions. Secondary locations, such as

valleys, isolated hills, and downstream areas, should be utilised strategically, recognising their inherent performance limitations and higher susceptibility to turbine-induced turbulence and hill wakes. Although these general guidelines are useful, the optimal positioning of turbines remains highly site-specific and challenging to determine a priori due to the complex terrain-flow-wake interactions. Automated optimisation frameworks, therefore, play a crucial role in identifying the best-performing areas within the terrain. The complex flow physics underlying these optimisation results are examined in detail in the following section.

5.2.3 Flow Physics over the Synthetic Multi-Hill Topography

The optimisation outcomes discussed above highlight the critical role of flow-terrain interactions in shaping wind farm performance. These interactions are fundamentally non-linear, exhibiting complex directional sensitivities that are difficult to capture with low-fidelity methods. To demonstrate the value of high-fidelity LES in resolving these dynamics, this section examines the three-dimensional turbulent flow structures, the bidirectional coupling between turbines and terrain, and the complex wake development patterns that arise in this topography.

Three-Dimensional Turbulent Flow Structures

All previous results shown throughout this study are based on statistically converged, time-averaged LES data. In contrast, Figure 5.12 presents an instantaneous three-dimensional visualisation of the flow field over the synthetic complex terrain, illustrating the multi-scale turbulent structures that govern the wind farm operation. Velocity magnitude contours combined with vorticity isosurfaces ($\omega = 0.16\ s^{-1}$) reveal how ABL turbulence, terrain-induced flow features, and turbine wake dynamics interact in ways that require high-fidelity computational resolution.

The boundary layer turbulence, visible in characteristic formations on the velocity planes, becomes progressively modified as it encounters the topography, with smooth terrain features maintaining organised flow patterns whereas sharper features create chaotic three-dimensional structures. Flow separation behind features such as hill A and turbulent mixing in valleys

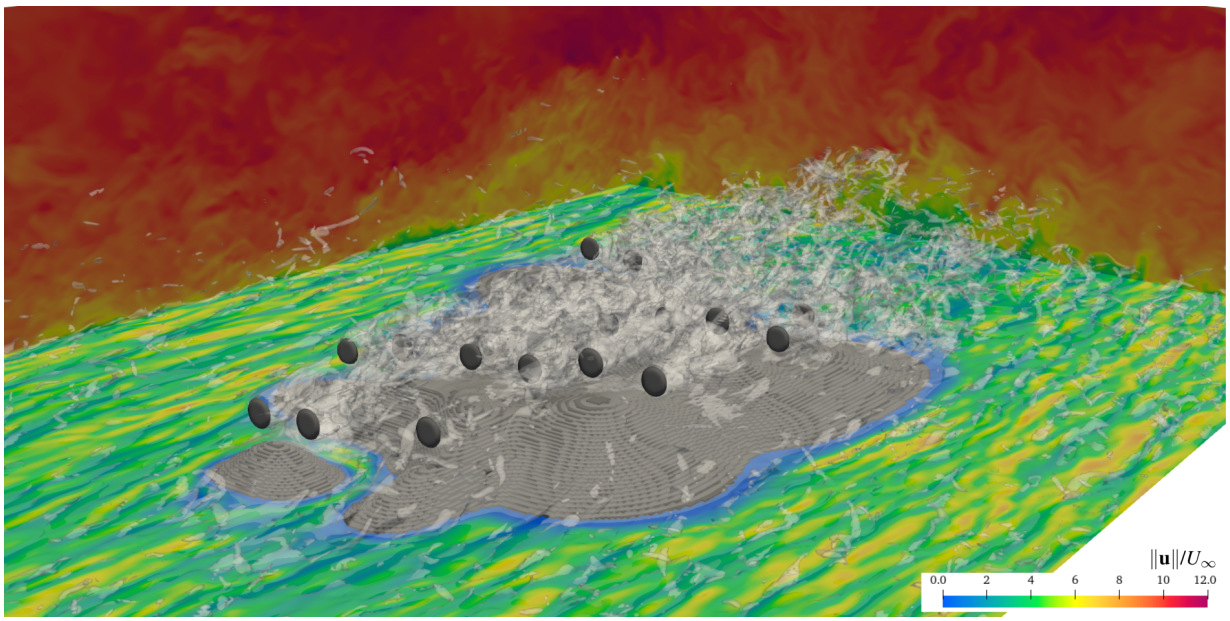


Figure 5.12: Instantaneous flow visualisation showing velocity magnitude contours at the boundaries and vorticity contours (white) for the optimal multi-directional layout under 270° inflow.

downstream of the main ridge produce zones with multiple scales of motion, from large-scale terrain wakes to finer-scale mixing layers.

Turbine-flow interactions add further complexity, as wakes meander and mix with terrain-modified inflow. Several turbines operate within highly turbulent regions formed by upstream terrain, where local conditions deviate significantly from uniform inflow assumptions. The multi-scale turbulence field encompasses large-scale terrain wakes, intermediate-scale turbine wakes, and finer-scale boundary layer turbulence, creating a coupled system where terrain effects, ABL turbulence, and turbine wakes interact inseparably.

The instantaneous nature of this visualisation underscores why steady-state modelling approaches struggle to capture the essential physics governing wind farm performance in complex terrain. The time-dependent turbulent structures contribute to accurate time-averaged flow field statistics, and their unsteady and complex interactions remain difficult to capture with steady-state approximations. These insights reinforce the importance of high-fidelity, time-resolved approaches for WFLO in challenging topographies.

Wake Physics and Terrain Interactions

The terrain-dependent nature of wake development is illustrated in Figure 5.13, showing a cross-sectional view at z = 1865 m through the optimal WFLO-1D-270° layout. It captures two strategically positioned turbines: turbine 10 on a hilltop and turbine 11 in the downstream valley. The velocity field and corresponding velocity deficit patterns (the difference between time-averaged flow with and without turbines, normalised by the hub-height reference velocity) reveal sophisticated wake dynamics from turbine-terrain coupling, demonstrating phenomena that superposition models do not capture.

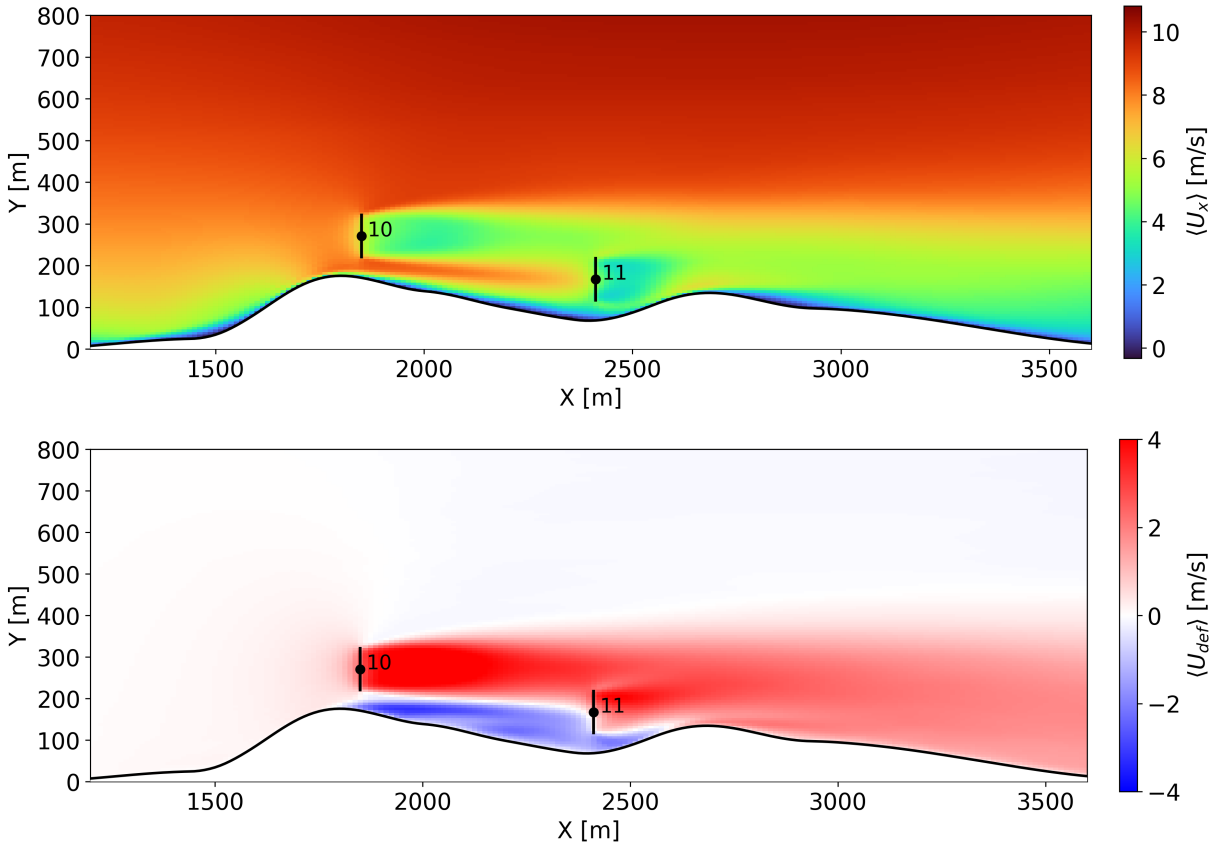


Figure 5.13: Cross-sectional view at z = 1865 m showing time-averaged velocity field (top) and velocity deficit (bottom) for turbines 10 and 11 in the optimal WFLO-1D-270° layout.

The hilltop turbine (turbine 10) demonstrates the well-documented phenomenon of wake elevation over complex terrain (Liu and Stevens, 2020a; Zhang et al., 2022; Jané-Ippel et al., 2023). Its wake remains elevated rather than following the downslope terrain, passing partly over the downstream valley turbine (turbine 11) with reduced direct interaction. This occurs due to the adverse pressure gradient after the hill crest, which slows wake recovery and causes

separation from the surface. In contrast, the valley turbine experiences enhanced mixing and faster wake recovery due to the favourable pressure gradient from the approaching downstream hill, highlighting how terrain-induced pressure gradients fundamentally alter wake physics and recovery rates. The velocity deficit field also shows that turbine 10 generates a flow acceleration zone along its wake boundary, directed toward the downstream turbine, reflecting how careful placement can favourably modify terrain flow dynamics.

These localised effects extend across the wind farm, as shown in Figure 5.14, which presents hub-height velocity deficit fields for the optimal multi-directional layout under four wind directions.

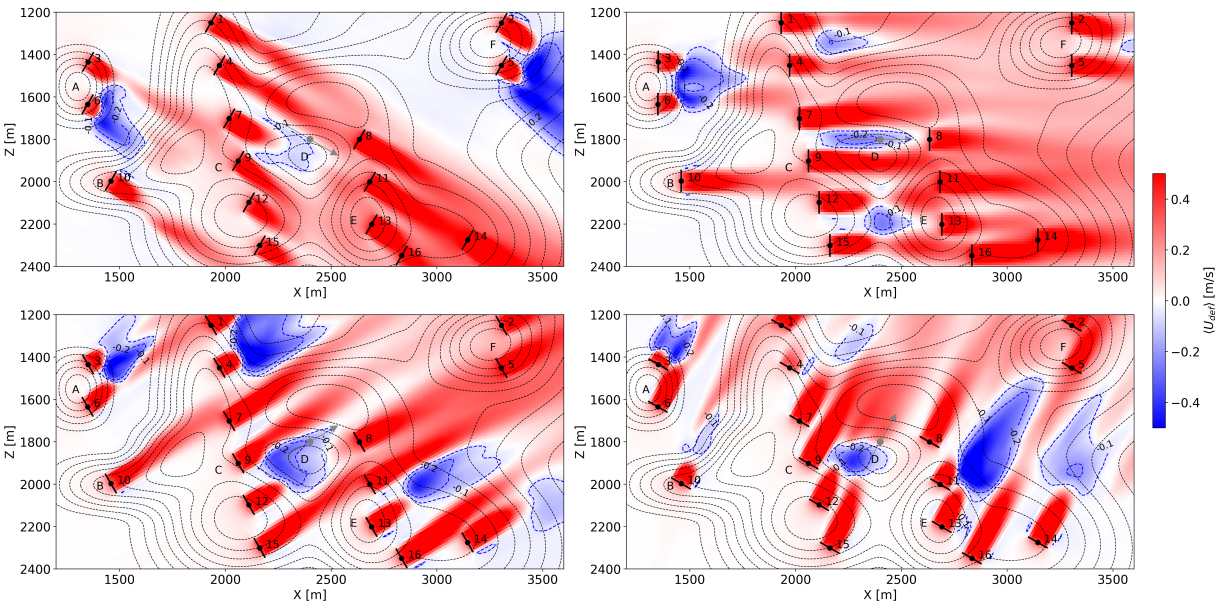


Figure 5.14: Hub-height velocity deficit fields for the optimal multi-directional layout across four wind directions. Red regions indicate velocity deficits (turbine wakes), with blue regions showing flow acceleration induced by turbine presence. Top row: 300° (left) and 270° (right). Bottom row: 240° (left) and 210° (right). Terrain contours (dashed lines) are shown at 20-m intervals.

The multi-directional analysis reveals additional complexity in wake-terrain interactions, including directional variations in wake deflection and recovery patterns. Wake deflection patterns are evident for turbine 10, where the wake trajectory shifts substantially between wind directions due to terrain-induced pressure gradients and cross-flow components. The extensive blue regions within hill wake zones across multiple directions confirm that turbine placement systematically alters terrain flow characteristics, creating downstream conditions that differ fundamentally from baseline terrain flow. This modification of terrain flow patterns by up-

stream turbines invalidates superposition approaches that treat turbine and terrain effects as independent phenomena.

5.3 Conclusions

The aim of this chapter was to extend the LES–BO framework from reduced two-turbine studies to a utility-scale farm operating over synthetically generated complex terrain. By tackling a thirty-two-dimensional optimisation problem under realistic spatial constraints and directional inflows, the study tested whether high-fidelity simulations coupled with physics-informed BO can deliver robust siting strategies in settings representative of practical wind-farm design.

The optimisation results demonstrated that terrain features such as ridges and hilltops can be systematically exploited to enhance power output, with optimal layouts consistently favouring elevated sites and staggered configurations that limit direct wake overlap. Single-direction cases showed how spacing constraints and inflow direction shape the achievable gains, while the multi-directional scenario highlighted the need to balance ridge utilisation with resilience to less favourable conditions. Across all cases, the physics-informed initialisation proved critical: black-box BO alone struggled to identify high-performing regions within the available budget, but when seeded with flow-based candidate layouts, it effectively refined and consolidated high-performing designs.

These findings confirm both the potential and the limitations of LES-driven optimisation in high-dimensional design spaces. On the one hand, the results illustrate how flow physics can be leveraged to produce interpretable, terrain-sensitive layouts that go beyond the reach of lower-order engineering models. On the other hand, they underscore the challenges posed by computational cost and search efficiency, showing that BO must be guided by physical insight if it is to remain tractable at utility scale.

Overall, this chapter demonstrates that the LES–BO framework can scale from proof-of-concept to utility-scale wind-farm problems when domain knowledge is embedded into the optimisation process. The final chapter synthesises the validation, two-turbine, and utility-scale studies

to reflect on the broader scientific and practical contributions of LES-driven optimisation in complex terrain.

Chapter 6

Conclusions and Future Work

This chapter synthesises the findings of the thesis, situates them within the broader context of wind-farm layout optimisation (WFLO), and outlines promising directions for future research. The work has combined high-fidelity large-eddy simulations (LES) with Bayesian Optimisation (BO) to tackle the challenge of wind-farm design in complex terrain, where traditional modelling approaches struggle. The discussion begins by restating the overarching aim and highlighting the methodological advances achieved, before drawing together the results from the validation and optimisation studies, assessing their scientific and practical implications, and concluding with limitations and perspectives for future work.

6.1 Conclusions

The overarching aim of this thesis was to develop and apply an LES–BO framework for WFLO in complex terrain. The motivation was twofold. On one hand, the ambition was to extend optimisation studies beyond flat terrain and simplified wake models, testing whether utility-scale optimisation could be carried out with turbulence-resolving fidelity. On the other, the scientific objective was to clarify how terrain-induced flow features (such as ridge acceleration, recirculation, and wake deflection) can be systematically exploited to improve power capture. A necessary first step was the extension of the `WInc3D` solver, originally designed for flat terrain, to

perform onshore simulations in realistic topographies. This provided the foundation on which the optimisation framework could be built.

6.1.1 Methodological Contributions

Methodologically, the thesis contributed to both the simulation and optimisation components of the workflow. On the simulation side, `WInc3D` was enhanced to support complex terrain through an immersed boundary method coupled with wall-stress modelling, enabling the treatment of steep hills while maintaining numerical stability and physical fidelity. The actuator-disc representation was extended with a super-Gaussian force kernel, allowing yawed and tilted rotors to be represented accurately even when misaligned with the Cartesian grid. These developments extended `WInc3D` from flat-terrain applications to complex topographies, enabling high-fidelity onshore wind-farm simulations.

On the optimisation side, the thesis integrated BO into an HPC-driven workflow in which large-scale LES campaigns could be performed tractably on national supercomputing resources. Standard elements of BO were combined with parallel LES execution to achieve sample-efficient exploration of large design spaces. Building on this, physics-informed seeding strategies were introduced as a key novelty of this work. By constructing initial candidate layouts directly from flow fields, the surrogate models were anchored in physically relevant regions of the design space, substantially improving the efficiency and robustness of optimisation at utility scale.

Together, these methodological contributions establish a methodology that is both physically grounded and computationally feasible, bridging the gap between high-fidelity fluid dynamics and optimisation algorithms. They also lay the groundwork for further developments, such as multi-fidelity approaches and hybrid workflows, which are discussed later in this chapter.

6.1.2 Synthesis of Results Across Chapters

The results presented across the thesis follow a deliberate progression from controlled benchmarks to utility-scale optimisation, each stage adding confidence and complexity. The valida-

tion studies in Chapter 3 demonstrated that the LES solver captures the essential physics of boundary-layer development, terrain-induced flow, and turbine performance. Canonical wind-tunnel benchmarks confirmed the accuracy of inflow generation, terrain-induced separation and reattachment, and wake recovery. Comparison against SCADA data from an operational wind farm extended validation to a more uncertain setting, showing that observed spatial trends were reproduced. These validation studies established the level of reliability required for optimisation to be attempted with confidence.

Building on this foundation, Chapter 4 presented a proof-of-concept optimisation study in which two turbines were placed around a Gaussian hill. This controlled configuration allowed terrain–wake interactions to be isolated and exploited systematically. The results revealed two distinct optimisation pathways: raising hub height to harness crest acceleration, and tilting the upstream rotor to deflect wakes away from the downstream turbine. Both strategies achieved substantial power gains relative to reference layouts, demonstrating that terrain features can be leveraged to mitigate wake losses. These findings provided clear evidence that the LES–BO coupling can identify siting strategies inaccessible to conventional low-order models.

The methodology was then scaled to a utility-scale wind farm in Chapter 5, where 16 turbines were distributed over synthetically generated three-dimensional terrain. Here the design space expanded to thirty-two variables, making the optimisation task substantially more demanding. The study showed that traditional black-box BO was insufficient under such conditions: the surrogate model struggled to identify high-performing regions within the available evaluation budget. In contrast, physics-informed initialisation proved decisive, allowing BO to act effectively as a refinement tool. Optimal layouts consistently exploited ridges and elevated terrain whilst avoiding recirculation zones and direct wake overlap, with directional dependence emerging clearly in multi-directional inflow scenarios. These results demonstrated that the methodology can scale to realistic design problems when guided by physical insight into the underlying flow.

Beyond layout performance metrics, the analysis of velocity and turbulence intensity fields provided insight into the underlying terrain–wake coupling mechanisms. The thesis clarified how

ridge acceleration, flow recirculation, wake deflection, and valley sheltering interact to shape turbine inflow conditions, thereby extending physical understanding of wind-farm performance in complex topography.

Taken together, the progression across chapters demonstrates a coherent narrative. Validation confirmed the solver's reliability; the two-turbine study established proof-of-concept and revealed fundamental mechanisms; and the 16-turbine study showed that physics-informed optimisation enables tractable design outcomes at scale. Each stage therefore reinforced and extended the previous one, building towards a comprehensive demonstration of LES-driven optimisation for onshore wind farms.

6.1.3 Scientific and Practical Impact

Scientifically, this thesis extends WFLO beyond the limits of analytical wake models and steady RANS simulations. By combining high-order LES with BO, it demonstrates for the first time that turbulence-resolving simulations can be used not only for validation but also as predictive tools for optimisation. These flow mechanisms are either absent or strongly underrepresented in lower-fidelity models, underscoring the necessity of high-fidelity methods in complex topography. The work also provided new insight into the behaviour of optimisation algorithms when applied to high-dimensional, physics-dominated problems: while BO was effective in reduced settings, its performance diminished in larger design spaces without appropriate initialisation strategies. This finding highlights both the promise and the limits of black-box optimisation in computational fluid dynamics, with lessons that extend beyond wind energy.

The practical contributions are equally significant. The thesis presents a scalable workflow for wind-farm design in complex terrain, where conventional engineering tools often fail to predict site-specific flow patterns reliably. The demonstration that physics-informed strategies markedly improve optimisation efficiency provides concrete guidance for early-stage project development, where reducing uncertainty is critical to financial and operational decision-making. The results further show that high-fidelity optimisation is feasible with modern computational

resources: large ensembles of LES were successfully executed on national HPC facilities, indicating that such methods are no longer restricted to idealised academic studies. As cloud-based computing resources continue to expand, the integration of LES-driven workflows into industrial practice becomes increasingly realistic. In this way, the thesis not only advances scientific understanding but also outlines practical strategies for how industry could progressively incorporate high-fidelity optimisation into routine wind-farm design.

6.1.4 Limitations

While the thesis demonstrates the feasibility and value of LES-driven optimisation, several limitations must be acknowledged. These arise from modelling assumptions, computational constraints, and the deliberate scope of the study.

The simulations were restricted to neutral atmospheric boundary-layer conditions. Stability effects such as thermal stratification or diurnal variability were not included, even though they can strongly influence flow development and wake recovery. The present approach prescribes inflow conditions at the microscale, without explicit coupling to mesoscale models such as WRFvegetatio. This omission means that larger-scale weather dynamics are not yet reflected in the optimisation process. Similarly, the turbine representation relied on the ADM, which neglects tower and nacelle effects and does not resolve blade-scale dynamics or structural loads. These simplifications are appropriate for layout optimisation, where the focus is on mean power capture, but they limit the ability to assess fatigue, extreme loading, or turbine-specific design implications. The absence of an actuator-line representation means that interactions at very short spacings could not be fully resolved.

The synthetic terrain employed in the large-scale optimisation was chosen to ensure reproducibility and to capture generic features such as ridges, valleys, and slopes. While this design allowed systematic exploration of terrain–wake coupling mechanisms, it was not tied to a specific commercial site. The real-farm validation in Chapter 3 mitigates this limitation by showing that the simulations can reproduce operational trends, but further work is required before the methodology can be deployed directly in industrial projects.

Finally, the optimisation campaigns were constrained by computational cost, with evaluation budgets already at the upper bound of what is currently tractable on national HPC systems. While sufficient to demonstrate convergence and extract clear trends, such budgets restrict the exploration of very large design spaces or simultaneous optimisation of multiple design variables such as yaw, tilt, or hub height. These constraints underline the need for multi-fidelity or smarter surrogate-assisted approaches in future extensions.

These limitations define the scope within which the conclusions of this thesis are valid. They also highlight natural directions for future work, where the inclusion of additional physics, broader objective functions, and more efficient optimisation strategies will allow the methodology to approach the full complexity of real wind-farm design.

6.1.5 Closing Statement

The objectives set out in Chapter 1 have been fulfilled. The coupling of LES with BO has delivered both proof-of-concept and utility-scale demonstrations, establishing this methodology as viable for WFLO in challenging terrain and setting the stage for the future developments outlined below.

6.2 Future Work

The limitations identified above suggest several promising avenues for further research and development.

A first priority is the inclusion of additional atmospheric physics. Stability effects, diurnal cycles, and non-neutral stratification play a major role in shaping flow development and wake recovery in practice. Extending the LES–BO to account for these processes would allow optimisation outcomes to be tested under more realistic operating conditions. Similarly, incorporating canopy or vegetation models would capture the influence of forestry, which emerged as a source of uncertainty in the validation study.

A second direction is the joint optimisation of turbine layout with control parameters such as yaw and tilt. The two-turbine study already demonstrated that wake steering can recover power losses associated with wake interactions, but this was explored in isolation. A natural extension is to treat siting and control variables simultaneously, which could reveal synergies between terrain exploitation and active wake redirection. Alongside this, load-aware objectives should be considered, balancing energy capture against fatigue or extreme loading, so that optimal layouts remain viable from both an aerodynamic and structural perspective.

From a methodological standpoint, efficiency remains a central challenge. Although substantial budgets of LES runs demonstrated feasibility, they represent the upper bound of current HPC capabilities and are unlikely to be practical in routine industrial workflows. Future studies should therefore explore the viability of multi-fidelity strategies in onshore farms, in which reduced-order models (such as RANS, wake models, or surrogates trained on LES data) are used to propose candidate layouts, while LES is reserved for refining and validating the most promising designs. The physics-informed seeding strategies introduced in this work already take a step in this direction, and could be expanded to guide the sampling process through the optimisation. A further extension is the coupling of mesoscale and microscale models. Mesoscale weather prediction tools such as the Weather Research and Forecasting (WRF) model capture synoptic-scale and regional wind patterns that strongly affect inflow conditions at wind-farm sites. Linking WRF outputs to the LES–BO framework would enable optimisation studies to reflect not only local terrain-induced effects, but also the variability imposed by larger-scale atmospheric dynamics. This multi-scale approach would bridge resource assessment and farm-scale optimisation, providing a more complete basis for design decisions in industry practice.

Finally, this methodology should be applied to real sites with complex topography. The synthetic terrain used here ensured reproducibility and captured generic flow mechanisms, but practical deployment requires adaptation to site-specific features, wind roses, and land-use constraints. The validation against an operational wind farm has already shown that the approach can reproduce measured spatial performance trends; the next step is to embed the workflow into full-scale design studies in collaboration with industry partners. Growing access to cloud-based computing and national HPC systems makes this prospect increasingly realistic.

In summary, future work should broaden physical realism, expand the range of design variables and objectives, and improve computational efficiency through advanced informed optimisations and multi-fidelity approaches. Together, these developments will extend the LES–BO methodology closer to an industry-ready framework for WFLO in complex terrain.

Bibliography

Abedi, H. (2023). Assessment of flow characteristics over complex terrain covered by the heterogeneous forest at slightly varying mean flow directions: (a case study of a swedish wind farm). *Renewable Energy* **202**, 537–553. DOI: 10.1016/j.renene.2022.11.030.

Allen, J., R. King and G. Barter (2020). Wind farm simulation and layout optimization in complex terrain. *Journal of Physics: Conference Series* **1452**.1, 012066. DOI: 10.1088/1742-6596/1452/1/012066.

Annoni, J., A. Scholbrock, M. Churchfield and P. Fleming (2017). Evaluating tilt for wind plants. *2017 American Control Conference (ACC)*. IEEE, 717–722. DOI: 10.23919/ACC.2017.7963037.

Antonini, E. G. A., D. A. Romero and C. H. Amon (2020). Optimal design of wind farms in complex terrains using computational fluid dynamics and adjoint methods. *Applied Energy* **261**, 114426. DOI: 10.1016/j.apenergy.2019.114426.

Arthur, R. S., J. D. Mirocha, K. A. Lundquist and R. L. Street (2019). Using a canopy model framework to improve large-eddy simulations of the neutral atmospheric boundary layer in the weather research and forecasting model. *Monthly Weather Review* **147**.1, 31–52. DOI: 10.1175/MWR-D-18-0204.1.

Astolfi, D., F. Castellani and L. Terzi (2018). A study of wind turbine wakes in complex terrain through RANS simulation and SCADA data. *Journal of Solar Energy Engineering* **140**.3, 031001. DOI: 10.1115/1.4039093.

Barthelmie, R. J., K. Hansen, S. T. Frandsen, O. Rathmann, J. G. Schepers, W. Schlez, J. Phillips, K. Rados, A. Zervos, E. S. Politis and P. K. Chaviaropoulos (2009). Modelling and

measuring flow and wind turbine wakes in large wind farms offshore. *Wind Energy* **12**.5, 431–444. DOI: 10.1002/we.348.

Barthelmie, R. J. and S. C. Pryor (2019). Automated wind turbine wake characterization in complex terrain. *Atmospheric Measurement Techniques* **12** (6), 3463–3484. DOI: 10.5194/amt-12-3463-2019.

Bartholomew, P., G. Deskos, R. A. S. Frantz, F. N. Schuch, E. Lamballais and S. Laizet (2020). Xcompact3D: an open-source framework for solving turbulence problems on a cartesian mesh. *SoftwareX* **12**, 100550. DOI: 10.1016/j.softx.2020.100550.

Bastankhah, M. and F. Porté-Agel (2014). A new analytical model for wind-turbine wakes. *Renewable Energy* **70**, 116–123. DOI: 10.1016/j.renene.2014.01.002.

Bempedelis, N. and K. Steiros (2022). Analytical all-induction state model for wind turbine wakes. *Physical Review Fluids* **7**.3 (3), 034605. DOI: 10.1103/physrevfluids.7.034605.

Bempedelis, N., F. Gori, A. Wynn, S. Laizet and L. Magri (2024). Data-driven optimisation of wind farm layout and wake steering with large-eddy simulations. *Wind Energy Science* **9**.4, 869–882. DOI: 10.5194/wes-9-869-2024.

Bempedelis, N., S. Laizet and G. Deskos (2023). Turbulent entrainment in finite-length wind farms. *Journal of Fluid Mechanics* **955**, A12. DOI: 10.1017/jfm.2022.1064.

Bempedelis, N. and L. Magri (2023). Bayesian optimization of the layout of wind farms with a high-fidelity surrogate model. *arXiv preprint*. DOI: 10.48550/arxiv.2302.01071.

Benard, P., A. Viré, V. Moureau, G. Lartigue, L. Beaudet, P. Deglaire and L. Bricteux (2018). Large-eddy simulation of wind turbines wakes including geometrical effects. *Computers & Fluids* **173**, 133–139. DOI: 10.1016/j.compfluid.2018.03.015.

Bergmann, M., C. Morsbach, G. Ashcroft and E. Kügeler (2021). Statistical error estimation methods for engineering-relevant quantities from scale-resolving simulations. *Journal of Turbomachinery* **144**.3, 031005. DOI: 10.1115/1.4052402.

Bernardoni, F., M. Guzman and S. Leonardi (2023). Implications of complex terrain topography on the performance of a real wind farm. *Journal of Physics: Conference Series* **2505**.1, 012052. DOI: 10.1088/1742-6596/2505/1/012052.

Betz, A. (1920). Das Maximum der theoretisch möglichen Ausnutzung des Windes durch Windmotoren. *Zeitschrift für das gesamte Turbinenwesen* **26**, 307–309.

Bou-Zeid, E., C. Meneveau and M. Parlange (2005). A scale-dependent lagrangian dynamic model for large eddy simulation of complex turbulent flows. *Physics of Fluids* **17**.2, 025105. DOI: 10.1063/1.1839152.

Breton, S.-P., J. Sumner, J. N. Sørensen, K. S. Hansen, S. Sarmast and S. Ivanell (2017). A survey of modelling methods for high-fidelity wind farm simulations using large eddy simulation. *Philosophical Transactions of the Royal Society A* **375**.2091, 20160097. DOI: 10.1098/rsta.2016.0097.

Bretos, D., G. Campaña-Alonso, B. Méndez-López and E. Cantero-Nouqueret (2024). CFD wind farm evaluation in complex terrain under free and wake induced flow conditions. *Journal of Physics: Conference Series* **2767**.9, 092099. DOI: 10.1088/1742-6596/2767/9/092099.

Brogna, R., J. Feng, J. N. Sørensen, W. Z. Shen and F. Porté-Agel (2020). A new wake model and comparison of eight algorithms for layout optimization of wind farms in complex terrain. *Applied Energy* **259**, 114189. DOI: 10.1016/j.apenergy.2019.114189.

Burton, T., N. Jenkins, D. Sharpe and E. Bossanyi (2011). *Wind energy handbook*. 2nd ed. Chichester, UK: Wiley. DOI: 10.1002/9781119992714.

Calaf, M., C. Meneveau and J. Meyers (2010). Large eddy simulation study of fully developed wind-turbine array boundary layers. *Physics of Fluids* **22**, 015110. DOI: 10.1063/1.3291077.

Cao, S. and T. Tamura (2006). Experimental study on roughness effects on turbulent boundary layer flow over a two-dimensional steep hill. *Journal of Wind Engineering and Industrial Aerodynamics* **94** (1), 1–19. DOI: 10.1016/j.jweia.2005.10.001.

Chamorro, L. P. and F. Porté-Agel (2010). Effects of thermal stability and incoming boundary-layer flow characteristics on wind-turbine wakes: a wind-tunnel study. *Boundary-Layer Meteorology* **136**.3, 515–533. DOI: 10.1007/s10546-010-9512-1.

Chatterjee, T. and Y. T. Peet (2018). Contribution of large scale coherence to wind turbine power: a large eddy simulation study in periodic wind farms. *Physical Review Fluids* **3**.3 (3), 034601. DOI: 10.1103/physrevfluids.3.034601.

Cheng, X., B. Yan, X. Zhou, Q. Yang, G. Huang, Y. Su, W. Yang and Y. Jiang (2024). Wind resource assessment at mountainous wind farm: fusion of RANS and vertical multi-point on-site measured wind field data. *Applied Energy* **363**, 123116. DOI: 10.1016/j.apenergy.2024.123116.

Chester, S., C. Meneveau and M. B. Parlange (2007). Modeling turbulent flow over fractal trees with renormalized numerical simulation. *Journal of Computational Physics* **225**.1, 427–448. DOI: 10.1016/j.jcp.2006.12.009.

Churchfield, M. J., S. Lee, J. Michalakes and P. J. Moriarty (2012). A numerical study of the effects of atmospheric and wake turbulence on wind turbine dynamics. *Journal of Turbulence* **13**, N14. DOI: 10.1080/14685248.2012.668191.

Clifton, A., S. Barber, A. Stökl, H. Frank and T. Karlsson (2022). Research challenges and needs for the deployment of wind energy in hilly and mountainous regions. *Wind Energy Science* **7**.6, 2231–2254. DOI: 10.5194/wes-7-2231-2022.

Cutler, J. J., A. P. Stanley, A. Ning and J. J. Thomas (2021). Optimization of turbine tilt in a wind farm. *AIAA Scitech 2021 Forum*. American Institute of Aeronautics and Astronautics, 1180. DOI: 10.2514/6.2021-1180.

Dahlberg, J.-A. (2009). Assessment of the Lillgrund windfarm, power performance and wake effects. Lillgrund pilot project.

Dar, A. S., J. Berg, N. Troldborg and E. G. Patton (2019). On the self-similarity of wind turbine wakes in a complex terrain using large eddy simulation. *Wind Energy Science* **4** (4), 633–644. DOI: 10.5194/wes-4-633-2019.

Dar, A. S. and F. Porté-Agel (2022). Wind turbine wakes on escarpments: a wind-tunnel study. *Renewable Energy* **181**, 1258–1275. DOI: 10.1016/j.renene.2021.09.102.

Department for Business, Energy & Industrial Strategy (2019). *UK becomes first major economy to pass net zero emissions law*. URL: https://www.gov.uk/government/news/uk-becomes-first-major-economy-to-pass-net-zero-emissions-law.

Department for Energy Security and Net Zero (2024). *Delivering clean power 2030: action plan*. Tech. rep. Crown Copyright, Open Government Licence v3.0. UK Government.

Deskos, G., S. Laizet and R. Palacios (2020). WInc3D: a novel framework for turbulence-resolving simulations of wind farm wake interactions. *Wind Energy* **23** (3), 779–94. DOI: 10.1002/we.2458.

Diebold, M., C. Higgins, J. Fang, A. Bechmann and M. B. Parlange (2013). Flow over hills: a large-eddy simulation of the bolund case. *Boundary-Layer Meteorology* **148**.1, 177–194. DOI: 10.1007/s10546-013-9807-0.

Diessner, M., J. O'Connor, A. Wynn, S. Laizet, Y. Guan, K. Wilson and R. D. Whalley (2022). Investigating Bayesian optimization for expensive-to-evaluate black box functions: application in fluid dynamics. *Frontiers in Applied Mathematics and Statistics* **8**. DOI: 10.3389/fams.2022.1076296.

Du, B., M. Ge, X. Li and Y. Liu (2024). A momentum-conserving wake superposition method for wind-farm flows under pressure gradient. *Journal of Fluid Mechanics* **999**, A27. DOI: 10.1017/jfm.2024.761.

Elgendi, M., M. AlMallahi, A. Abdelkhalig and M. Y. E. Selim (2023). A review of wind turbines in complex terrain. *International Journal of Thermofluids* **17**, 100289. DOI: 10.1016/j.ijft.2023.100289.

Fang, J. and F. Porté-Agel (2016). Intercomparison of terrain-following coordinate transformation and immersed boundary methods in large-eddy simulation of wind fields over complex terrain. *Journal of Physics: Conference Series* **753**.8, 082008. DOI: 10.1088/1742-6596/753/8/082008.

Farrell, A., J. King, C. Draxl, R. Mudafort, N. Hamilton, C. J. Bay, P. Fleming and E. Simley (2021). Design and analysis of a wake model for spatially heterogeneous flow. *Wind Energy Science* **6**.3, 737–758. DOI: 10.5194/wes-6-737-2021.

Fernando, H. J. S., J. Mann, J. M. L. M. Palma, J. K. Lundquist, R. J. Barthelmie, M. Belo-Pereira, W. O. J. Brown, F. K. Chow, T. Gerz, C. M. Hocut, P. M. Klein, L. S. Leo, J. C. Matos, S. P. Oncley, S. C. Pryor, L. Bariteau, T. M. Bell, N. Bodini, M. B. Carney, M. S. Courtney, E. D. Creegan, R. Dimitrova, S. Gomes, M. Hagen, J. O. Hyde, S. Kigle, R. Krishnamurthy, J. C. Lopes, L. Mazzaro, J. M. T. Neher, R. Menke, P. Murphy, L. Oswald, S. Otarola-Bustos, A. K. Pattantyus, C. V. Rodrigues, A. Schady, N. Sirin, S. Spuler, E.

Svensson, J. Tomaszewski, D. D. Turner, L. van Veen, N. Vasiljević, D. Vassallo, S. Voss, N. Wildmann and Y. Wang (2019). The Perdigão: peering into microscale details of mountain winds. *Bulletin of the American Meteorological Society* **100**.5, 799–819. DOI: 10.1175/BAMS-D-17-0227.1.

Frandsen, S., R. Barthelmie, S. Pryor, O. Rathmann, S. Larsen, J. Højstrup and M. Thøgersen (2006). Analytical modelling of wind speed deficit in large offshore wind farms. *Wind Energy* **9**.1–2, 39–53. DOI: 10.1002/we.189.

Froude, R. (1889). On the Part Played in Propulsion by Differences of Fluid Pressure. *Transactions of the Institution of Naval Architects* **30**, 390–405.

Gautier, R., S. Laizet and E. Lamballais (2014). A DNS study of jet control with microjets using an immersed boundary method. *International Journal of Computational Fluid Dynamics* **28**.6-10, 393–410. DOI: 10.1080/10618562.2014.950046.

Germano, M., U. Piomelli, P. Moin and W. H. Cabot (1991). A dynamic subgrid-scale eddy viscosity model. *Physics of Fluids A: Fluid Dynamics* **3**.7, 1760–1765. DOI: 10.1063/1.857955.

Giannenas, A. E. and S. Laizet (2021). A simple and scalable immersed boundary method for high-fidelity simulations of fixed and moving objects on a cartesian mesh. *Applied Mathematical Modelling* **99**, 606–627. DOI: 10.1016/j.apm.2021.06.026.

Global Wind Energy Council (2024). *Global wind report 2024*. Tech. rep. Brussels, Belgium: Global Wind Energy Council.

Global Wind Energy Council (2025). *Global wind report 2025*. Tech. rep. Brussels, Belgium: Global Wind Energy Council.

Goit, J. P. and J. Meyers (2015). Optimal control of energy extraction in wind-farm boundary layers. *Journal of Fluid Mechanics* **768**, 5–50. DOI: 10.1017/jfm.2015.70.

González, J., Z. Dai, P. Hennig and N. Lawrence (2016). Batch Bayesian optimization via local penalization. *Artificial Intelligence and Statistics*. PMLR, 648–657. DOI: 10.48550/arXiv.1505.08052.

Hansen, K. S., G. C. Larsen, R. Menke, N. Vasiljevic, N. Angelou, J. Feng, W. J. Zhu, A. Vignaroli, W. Liu, C. Xu and W. Z. Shen (2016). Wind turbine wake measurement in complex

terrain. *Journal of Physics: Conference Series* **753** (3). DOI: 10.1088/1742-6596/753/3/032013.

Heck, K. S., H. M. Johlas and M. F. Howland (2023). Modelling the induction, thrust and power of a yaw-misaligned actuator disk. *Journal of Fluid Mechanics* **959**, A9. DOI: 10.1017/jfm.2023.129.

Helmis, C. G., K. H. Papadopoulos, D. N. Asimakopoulos, P. G. Papageorgas and A. T. Soilemes (1995). An experimental study of the near-wake structure of a wind turbine operating over complex terrain. *Solar Energy* **54** (6), 413–428. DOI: 10.1016/0038-092x(95)00009-g.

Hodgkin, A., G. Deskos and S. Laizet (2023). On the interaction of a wind turbine wake with a conventionally neutral atmospheric boundary layer. *International Journal of Heat and Fluid Flow* **102**, 109165. DOI: 10.1016/j.ijheatfluidflow.2023.109165.

Hodgkin, A., S. Laizet and G. Deskos (2022). Numerical investigation of the influence of shear and thermal stratification on the wind turbine tip-vortex stability. *Wind Energy* **25**.7, 1270–1289. DOI: 10.1002/we.2728.

Howard, K. B., L. P. Chamorro and M. Guala (2016). A comparative analysis on the response of a wind-turbine model to atmospheric and terrain effects. *Boundary-Layer Meteorology* **158** (2), 229–255. DOI: 10.1007/s10546-015-0094-9.

Howard, K. B., J. S. Hu, L. P. Chamorro and M. Guala (2015). Characterizing the response of a wind turbine model under complex inflow conditions. *Wind Energy* **18**, 729–743. DOI: 10.1002/we.1724.

Hyvärinen, A. and A. Segalini (2017a). Qualitative analysis of wind-turbine wakes over hilly terrain. *Journal of Physics: Conference Series* **854** (1). DOI: 10.1088/1742-6596/854/1/012023.

Hyvärinen, A., G. Lacagnina and A. Segalini (2018). A wind-tunnel study of the wake development behind wind turbines over sinusoidal hills. *Wind Energy* **21**.8, 605–617. DOI: 10.1002/we.2181.

Hyvärinen, A. and A. Segalini (2017b). Effects from complex terrain on wind-turbine performance. *Journal of Energy Resources Technology, Transactions of the ASME* **139** (5). DOI: 10.1115/1.4036048.

Ishihara, T., K. Hibi and S. Oikawa (1999). A wind tunnel study of turbulent flow over a three-dimensional steep hill. *Journal of Wind Engineering and Industrial Aerodynamics* **83**.1–3, 95–107. DOI: 10.1016/s0167-6105(99)00064-1.

Jaegle, F., O. Cabrit, S. Mendez and T. Poinsot (2010). Implementation methods of wall functions in cell-vertex numerical solvers. *Flow, Turbulence and Combustion* **85**.2, 245–272. DOI: 10.1007/s10494-010-9276-1.

Janajreh, I., R. Qudaih, I. Talab and C. Ghenai (2010). Aerodynamic flow simulation of wind turbine: downwind versus upwind configuration. *Energy Conversion and Management* **51**.8, 1656–1663. DOI: 10.1016/j.enconman.2009.12.013.

Jané-Ippel, C., N. Bempedelis, R. Palacios and S. Laizet (2023). High-fidelity simulations of wake-to-wake interaction in an atmospheric boundary layer over a complex terrain. *Journal of Physics: Conference Series* **2505**.1, 012033. DOI: 10.1088/1742-6596/2505/1/012033.

Jané-Ippel, C., N. Bempedelis, R. Palacios and S. Laizet (2024). Bayesian optimisation of a two-turbine configuration around a 2D hill using large eddy simulations. *Wind Energy* **27**.11, 1412–1426. DOI: 10.1002/we.2946.

Jané-Ippel, C., R. Palacios and S. Laizet (2025). Physics-informed Bayesian optimisation of utility-scale wind farms in synthetically generated complex terrain using LES. *Wind Energy*. Accepted for publication. DOI: 10.22541/au.175273945.57235516/v1.

Jensen, N. O. (1983). *A note on wind generator interaction*. Risø-M 2411. Risø National Laboratory.

Kelly, M., G. Larsen, N. K. Dimitrov and A. Natarajan (2014). Probabilistic meteorological characterization for turbine loads. *Journal of Physics: Conference Series* **524**.1, 012076. DOI: 10.1088/1742-6596/524/1/012076.

Kim, J. and P. Moin (1985). Application of a fractional-step method to incompressible Navier-Stokes equations. *Journal of Computational Physics* **59**.2, 308–323. DOI: 10.1016/0021-9991(85)90148-2.

King, R. N., K. Dykes, P. Graf and P. E. Hamlington (2017). Optimization of wind plant layouts using an adjoint approach. *Wind Energy Science* **2**.1, 115–131. DOI: 10.5194/wes-2-115-2017.

Kleusberg, E., R. F. Mikkelsen, P. Schlatter, S. Ivanell and D. S. Henningson (2017). High-order numerical simulations of wind turbine wakes. *Journal of Physics: Conference Series* **854**.1, 012025. DOI: 10.1088/1742-6596/854/1/012025.

Kravchenko, A. G. and P. Moin (1997). On the effect of numerical errors in large eddy simulations of turbulent flows. *Journal of Computational Physics* **131**.2. DOI: 10.1006/jcph.1996.5597.

Krogstad, P.-Å. and P. E. Eriksen (2013). "Blind test" calculations of the performance and wake development for a model wind turbine. *Renewable Energy* **50**, 325–333. DOI: 10.1016/j.renene.2012.06.044.

Krogstad, P.-Å., L. Saetran and M. S. Adaramola (2015). "Blind test 3" calculations of the performance and wake development behind two in-line and offset model wind turbines. *Journal of Fluids and Structures* **52**, 65–80. DOI: 10.1016/j.jfluidstructs.2014.10.002.

Laizet, S. and E. Lamballais (2009). High-order compact schemes for incompressible flows: a simple and efficient method with quasi-spectral accuracy. *Journal of Computational Physics* **228** (16), 5989–6015. DOI: 10.1016/j.jcp.2009.05.010.

Laizet, S. and N. Li (2011). Incompact3D: A powerful tool to tackle turbulence problems with up to $\mathcal{O}(10^5)$ computational cores. *International Journal for Numerical Methods in Fluids* **67** (11), 1735–57. DOI: 10.1002/fld.2480.

Lamballais, E., R. Vicente Cruz and R. Perrin (2021). Viscous and hyperviscous filtering for direct and large-eddy simulation. *Journal of Computational Physics* **431**, 110115. DOI: 10.1016/j.jcp.2021.110115.

Lele, S. K. (1992). Compact finite difference schemes with spectral-like resolution. *Journal of Computational Physics* **103**, 16–42. DOI: 10.1016/0021-9991(92)90324-r.

Li, Z., X. Liu and X. Yang (2022). Review of turbine parameterization models for large-eddy simulation of wind turbine wakes. *Energies* **15**.18, 6533. DOI: 10.3390/en15186533.

Liu, D. C. and J. Nocedal (1989). On the limited memory BFGS method for large scale optimization. *Mathematical Programming* **45**.3, 503–528. DOI: 10.1007/BF01589116.

Liu, H., J. Fu, Z. Liang, Z. Liang, Y. Zhang and Z. Xiao (2022). A simple method of fast evaluating full-field wake velocities for arbitrary wind turbine arrays on complex terrains. *Renewable Energy* **201**, 961–976. DOI: 10.1016/j.renene.2022.10.124.

Liu, L. and R. J. A. M. Stevens (2020a). Effects of two-dimensional steep hills on the performance of wind turbines and wind farms. *Boundary-Layer Meteorology* **176**, 251–69. DOI: 10.1007/s10546-020-00522-z.

Liu, L. and R. J. A. M. Stevens (2020b). Wall modeled immersed boundary method for high Reynolds number flow over complex terrain. *Computers & Fluids* **208**, 104604. DOI: 10.1016/j.compfluid.2020.104604.

Mann, J., N. Angelou, J. Arnqvist, D. Callies, E. Cantero, R. C. Arroyo, M. Courtney, J. Cuxart, E. Dellwik, J. Gottschall, S. Ivanell, P. Kühn, G. Lea, J. C. Matos, J. M. L. M. Palma, L. Pauscher, A. Peña, J. S. Rodrigo, S. Söderberg, N. Vasiljevic and C. V. Rodrigues (2017). Complex terrain experiments in the new european wind atlas. *Philosophical Transactions of the Royal Society A* **375**.2091, 20160101. DOI: 10.1098/rsta.2016.0101.

Mason, P. J. and A. R. Brown (1999). On subgrid models and filter operations in large eddy simulations. *Journal of the Atmospheric Sciences* **56**.13, 2101–2114. DOI: 10.1175/1520-0469(1999)056<2101:OSMAFO>2.0.CO;2.

Mason, P. J. and D. J. Thomson (1992). Stochastic backscatter in large-eddy simulations of boundary layers. *Journal of Fluid Mechanics* **242**, 51–78. DOI: 10.1017/S0022112092002271.

McKay, M. D., R. J. Beckman and W. J. Conover (1979). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* **21**.2, 239–245. DOI: 10.2307/1271432.

Mehta, D., A. H. van Zuijlen, B. Koren, J. G. Holierhoek and H. Bijl (2014). Large eddy simulation of wind farm aerodynamics: a review. *Journal of Wind Engineering and Industrial Aerodynamics* **133**, 1–17. DOI: 10.1016/j.jweia.2014.07.002.

Menke, R., N. Vasiljević, K. S. Hansen, A. N. Hahmann and J. Mann (2018). Does the wind turbine wake follow the topography? A multi-lidar study in complex terrain. *Wind Energy Science* **3** (2), 681–691. DOI: 10.5194/wes-3-681-2018.

Meyers, J. and C. Meneveau (2011). Optimal turbine spacing in fully developed wind farm boundary layers. *Wind Energy* **15**.2, 305–317. DOI: 10.1002/we.469.

Mole, A. and S. Laizet (2024). Multi-fidelity Bayesian optimisation of wind farm wake steering using wake models and large eddy simulations. *Flow, Turbulence and Combustion*. DOI: 10.1007/s10494-024-00629-0.

Motheau, E. and J. Abraham (2016). A high-order numerical algorithm for DNS of low-mach-number reactive flows with detailed chemistry and quasi-spectral accuracy. *Journal of Computational Physics* **313**, 430–454. DOI: 10.1016/j.jcp.2016.02.059.

Mudafort, R. M., Paulf81, C. Bay, R. Hammond, Ejsimley, Nhamilto, P. Bachant, K. Fleming, B. Doekemeijer, Jrannoni, E. Quon, PJ Stanley, Misi9170, Pjireland, DavidBensason, J. Schreiber, J. Quick, K. S. Seim, P. B. Cia, Sondre Sortland, T. Martinez, Afarrell2, Pduff-Code and Zerweck (2023). Floris, version 3.4, github [code]. DOI: 10.5281/ZENODO.7942258.

Munters, W., C. Meneveau and J. Meyers (2016). Shifted periodic boundary conditions for simulations of wall-bounded turbulent flows. *Physics of Fluids* **28**.2, 025112. DOI: 10.1063/1.4941912.

Nilsson, K., S. Ivanell, K. S. Hansen, R. Mikkelsen, J. N. Sørensen, S.-P. Breton and D. Henningson (2014). Large-eddy simulations of the Lillgrund wind farm. *Wind Energy* **18**.3, 449–467. DOI: 10.1002/we.1707.

Pierella, F., P.-Å. Krogstad and L. Saetran (2014). "Blind test 2" calculations for two in-line model wind turbines where the downstream turbine operates at various rotational speeds. *Renewable Energy* **70**, 62–77. DOI: 10.1016/j.renene.2014.03.034.

Politis, E. S., J. Prospathopoulos, D. Cabezon, K. S. Hansen, P. K. Chaviaropoulos and R. J. Barthelmie (2011). Modeling wake effects in large wind farms in complex terrain: the problem, the methods and the issues. *Wind Energy* **15**.1, 161–182. DOI: 10.1002/we.481.

Porté-Agel, F., M. Bastankhah and S. Shamsoddin (2019). Wind-turbine and wind-farm flows: a review. *Boundary-Layer Meteorology* **174**.1, 1–59. DOI: 10.1007/s10546-019-00473-0.

Porté-Agel, F., C. Meneveau and M. B. Parlange (2000). A scale-dependent dynamic model for large-eddy simulation: application to a neutral atmospheric boundary layer. *Journal of Fluid Mechanics* **415**, 261–284. DOI: 10.1017/S0022112000008776.

Porté-Agel, F., Y.-T. Wu, H. Lu and R. J. Conzemius (2011). Large-eddy simulation of atmospheric boundary layer flow through wind turbines and wind farms. *Journal of Wind Engineering and Industrial Aerodynamics* **99**.4, 154–168. DOI: 10.1016/j.jweia.2011.01.011.

Rankine, W. J. M. (1865). On the mechanical principles of the action of propellers. *Transactions of the Institution of Naval Architects* **6**, 13–39.

Rasmussen, C. E. and C. K. I. Williams (2006). *Gaussian processes for machine learning*. Cambridge, MA, USA: MIT Press. DOI: 10.7551/mitpress/3206.001.0001.

Sanderse, B., S. P. van der Pijl and B. Koren (2011). Review of computational fluid dynamics for wind turbine wake aerodynamics. *Wind Energy* **14**.7, 799–819. DOI: 10.1002/we.458.

Segalini, A. (2017). Linearized simulation of flow over wind farms and complex terrains. *Philosophical Transactions of the Royal Society A* **375**.2091, 20160099. DOI: 10.1098/rsta.2016.0099.

Sessarego, M., W. Z. Shen, M. P. Van der Laan, K. S. Hansen and W. J. Zhu (2018). CFD simulations of flows in a wind farm in complex terrain and comparisons to measurements. *Applied Sciences* **8**.5. DOI: 10.3390/app8050788.

Shamsoddin, S. and F. Porté-Agel (2017). Large-eddy simulation of atmospheric boundary-layer flow through a wind farm sited on topography. *Boundary-Layer Meteorology* **163**, 1–17. DOI: 10.1007/s10546-016-0216-z.

Shamsoddin, S. and F. Porté-Agel (2018). A model for the effect of pressure gradient on turbulent axisymmetric wakes. *Journal of Fluid Mechanics* **837**, R3. DOI: 10.1017/jfm.2017.864.

Smagorinsky, J. (1963). General circulation experiments with the primitive equations. *Monthly Weather Review* **91**.3, 99–164. DOI: 10.1175/1520-0493(1963)091<0099:GCEWTP>2.3.CO;2.

Sørensen, J. N. and A. Myken (1992). Unsteady actuator disc model for horizontal axis wind turbines. *Journal of Wind Engineering and Industrial Aerodynamics* **39**.1, 139–149. DOI: doi.org/10.1016/0167-6105(92)90540-Q.

Sørensen, N. N. (1995). *General purpose flow solver applied to flow over hills*. Published 2003. PhD thesis. Technical University of Denmark.

Steiros, K., N. Bempedelis and M. M. Cicolin (2022). An analytical blockage correction model for high-solidity turbines. *Journal of Fluid Mechanics* **948**, A57. DOI: 10.1017/jfm.2022.735.

Steiros, K., C. Jané-Ippel, F. Llewellyn, R. Palacios, R. B. Cal and S. Laizet (2025). Improved wind power extraction downstream of a 3D hill using a tilted actuator disk. *Journal of Renewable and Sustainable Energy* **17**.5, 053303. DOI: 10.1063/5.0280650.

Stevens, R. J. A. M., J. Graham and C. Meneveau (2014). A concurrent precursor inflow method for large eddy simulations and applications to finite length wind farms. *Renewable Energy* **68**, 46–50. DOI: 10.1016/j.renene.2014.01.024.

Stevens, R. J. A. M., L. A. Martı´nez-Tossas and C. Meneveau (2018). Comparison of wind farm large eddy simulations using actuator disk and actuator line models with wind tunnel experiments. *Renewable Energy* **116**, 470–478. DOI: 10.1016/j.renene.2017.08.072.

Stevens, R. J. A. M. and C. Meneveau (2017). Flow structure and turbulence in wind farms. *Annual Review of Fluid Mechanics, Vol 43* **49**.1, 311–39. DOI: 10.1146/annurev-fluid-010816-060206.

Storey, R. C., S. E. Norris, K. A. Stol and J. E. Cater (2012). Large eddy simulation of dynamically controlled wind turbines in an offshore environment. *Wind Energy* **16**.6, 845–864. DOI: 10.1002/we.1525.

Subramanian, B., N. Chokani and R. S. Abhari (2016). Aerodynamics of wind turbine wakes in flat and complex terrains. *Renewable Energy* **85**, 454–463. DOI: 10.1016/j.renene.2015.06.060.

Tanaka, R. and H. Iwata (2018). Bayesian optimization for genomic selection: a method for discovering the best genotype among a large number of candidates. *Theoretical and Applied Genetics* **131**. DOI: 10.1007/s00122-017-2988-z.

Taylor, G. J. and D. Smith (1991). Wake measurements over complex terrain. *Proceedings of the 13th BWEA Conference: Wind Energy Conversion*, 10–12.

The GPyOpt authors (2016). GPyOpt: A Bayesian optimization framework in Python. `http://github.com/SheffieldML/GPyOpt`.

Tian, W., A. Ozbay, W. Yuan, P. Sarakar and H. Hu (2013). An experimental study on the performances of wind turbines over complex terrain. *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition 2013*. DOI: `10.2514/6.2013-612`.

Troldborg, N., S. J. Andersen, E. L. Hodgson and A. Meyer Forsting (2022). Brief communication: how does complex terrain change the power curve of a wind turbine? *Wind Energy Science* **7**.4, 1527–1532. DOI: `10.5194/wes-7-1527-2022`.

Veers, P., K. Dykes, E. Lantz, S. Barth, C. L. Bottasso, O. Carlson, A. Clifton, J. Green, P. Green, H. Holttinen, et al. (2019). Grand challenges in the science of wind energy. *Science* **366**.6464. DOI: `10.1126/science.aau2027`.

Wang, D., D. Feng, H. Peng, F. Mao, M. H. Doranehgard, V. Gupta, L. K. B. Li and M. Wan (2023). Implications of steep hilly terrain for modeling wind-turbine wakes. *Journal of Cleaner Production* **398**, 136614. DOI: `10.1016/j.jclepro.2023.136614`.

Wang, Z., Y. Tu, K. Zhang, Z. Han, Y. Cao and D. Zhou (2024). An optimization framework for wind farm layout design using CFD-based kriging model. *Ocean Engineering* **293**, 116644. DOI: `10.1016/j.oceaneng.2023.116644`.

Wiernga, J. (1993). Representative roughness parameters for homogeneous terrain. *Boundary-Layer Meteorology* **63**.4, 323–363. DOI: `10.1007/bf00705357`.

Wingerden, J. W. V., P. A. Fleming, T. Göcmen, I. Eguinoa, B. M. Doekemeijer, K. Dykes, M. Lawson, E. Simley, J. King, D. Astrain, M. Iribas, C. L. Bottasso, J. Meyers, S. Raach, K. Kölle and G. Giebel (2020). Expert elicitation on wind farm control. *Yamada Conference Lx On Research in High Magnetic Fields* **1618** (2). DOI: `10.1088/1742-6596/1618/2/022025`.

Witha, B., G. Steinfeld and D. Heinemann (2014). High-resolution offshore wake simulations with the LES model PALM. *Wind Energy – Impact of Turbulence*. Berlin, Heidelberg: Springer, 175–181. DOI: `10.1007/978-3-642-54696-9_26`.

Wu, Y.-T. and F. Porté-Agel (2011). Large-eddy simulation of wind-turbine wakes: evaluation of turbine parametrisations. *Boundary-Layer Meteorology* **138**.3, 345–366. DOI: 10.1007/s10546-010-9569-x.

Wu, Y.-T. and F. Porté-Agel (2013). Simulation of turbulent flow inside and above wind farms: model validation and layout effects. *Boundary-Layer Meteorology* **146**.2, 181–205. DOI: 10.1007/s10546-012-9757-y.

Wu, Y.-T. and F. Porté-Agel (2015). Modeling turbine wakes and power losses within a wind farm using LES: an application to the Horns Rev offshore wind farm. *Renewable Energy* **75**, 945–955. DOI: 10.1016/j.renene.2014.06.019.

Xie, S. and C. Archer (2014). Self-similarity and turbulence characteristics of wind turbine wakes via large-eddy simulation. *Wind Energy* **18**.10, 1815–1838. DOI: 10.1002/we.1792.

Yang, X., K. B. Howard, M. Guala and F. Sotiropoulos (2015). Effects of a three-dimensional hill on the wake characteristics of a model wind turbine. *Physics of Fluids* **27**, 025103. DOI: 10.1063/1.4907685.

Yang, X., M. Pakula and F. Sotiropoulos (2018). Large-eddy simulation of a utility-scale wind farm in complex terrain. *Applied Energy* **229**, 767–77. DOI: 10.1016/j.apenergy.2018.08.049.

Yang, X., F. Sotiropoulos, R. J. Conzemius, J. N. Wachtler and M. B. Strong (2014). Large-eddy simulation of turbulent flow past wind turbines/farms: the Virtual Wind Simulator (VWiS). *Wind Energy* **18**, 2025–2045. DOI: 10.1002/we.1802.

Zhang, Z., P. Huang, G. Bitsuamlak and S. Cao (2022). Large-eddy simulation of wind-turbine wakes over two-dimensional hills. *Physics of Fluids* **34**, 065123. DOI: 10.1063/5.0095348.

Zong, H. and F. Porté-Agel (2020). A momentum-conserving wake superposition method for wind farm power prediction. *Journal of Fluid Mechanics* **889**, A8. DOI: 10.1017/jfm.2020.77.

Appendix A

Copyright Permissions

Table A.1: Permissions summary table for reproduced figures

Fig. no.	Source of figure	Copyright holder	Permission	Licence
1.1	Global Wind Energy Council (2024). *Global wind report 2024*. Tech. rep. Brussels, Belgium: Global Wind Energy Council, p. 53.	GWEC	yes	GWEC permission for non-commercial thesis use.
1.2	iStock photo.	iStock/Getty Images	yes	Standard licence purchased from iStock (asset ID: 2166705101).
1.3	Nordex SE press release image.	Nordex SE	yes	Publication and reprint free of charge; credit Nordex SE.

Continued on next page

Fig. no.	Source of figure	Copyright holder	Permission	Licence
2.1	Laizet, S. and Lamballais, E. (2009). High-order compact schemes for incompressible flows: a simple and efficient method with quasi-spectral accuracy. *Journal of Computational Physics* **228** (16), 5989–6015. DOI: 10.1016/j.jcp.2009.05.010	Elsevier	yes	RightsLink licence 6114131494018
2.2	Gautier, R. et al. (2014). A DNS study of jet control with microjets using an immersed boundary method. *International Journal of Computational Fluid Dynamics* **28**.6-10, 393–410. DOI: 10.1080/10618562.2014.950046	Informa UK Limited (Taylor & Francis Group)	yes	Taylor & Francis free permission for thesis/dissertation use. Ref: gcfd/05604926
2.3	Laizet, S. and Li, N. (2011). Incompact3D: A powerful tool to tackle turbulence problems with up to $\mathcal{O}(10^5)$ computational cores. *International Journal for Numerical Methods in Fluids* **67** (11), 1735–57. DOI: 10.1002/fld.2480	Wiley	yes	RightsLink licence 6114151014348
2.4	Tanaka, R. and Iwata, H. (2018). Bayesian optimization for genomic selection: a method for discovering the best genotype among a large number of candidates. *Theoretical and Applied Genetics* **131**. DOI: 10.1007/s00122-017-2988-z	Springer Nature	yes	RightsLink licence 6115481236782