

# Layout and Control Co-Design for Floating Offshore Wind Farms with Active and Passive Turbine Repositioning

Konstantinos Chondros



# Layout and Control Co-Design for Floating Offshore Wind Farms with Active and Passive Turbine Repositioning

by

Konstantinos Chondros

Chairman of Thesis Committee: Professor Dr. J.W. (Jan-Willem) van Wingerden  
Thesis Committee Member: Dr. J. (Jenna) Iori  
Daily Supervisor: Matteo Baricchio, PhD Student  
External Supervisor: Dr. Maarten van den Broek, sowento  
Project Duration: November 2025 – July 2026  
Faculty: Faculty of Mechanical Engineering, Delft  
Department: Delft Center for Systems and Control  
Cover: Generated with ChatGPT



# Summary

Floating offshore wind turbines extend wind energy into deep-water sites that bottom-fixed foundations cannot reach. Unlike a bottom-fixed machine, a floating turbine does not hold a fixed position. Under thrust, each platform settles at a mean offset from its nominal location, a drift referred to here as passive repositioning. Wake steering adds a second source of motion. When a turbine yaws to deflect its wake, the cross-wind component of the thrust also pushes the platform sideways, shifting the rotor to a new position, an effect referred to as active repositioning. The layout the wind actually sees is therefore not the one drawn at nominal positions, and the wake interactions across the farm move with it. This motivates accounting for platform displacement inside both the layout and the control optimization. Building on the integrated layout and control co-design frameworks developed for bottom-fixed farms, this thesis extends the approach to floating wind turbines.

The resulting framework accounts for both passive and active repositioning within a single optimization loop. One XGBoost surrogate predicts the surge and sway displacement of each platform, and a second predicts the optimal yaw angle directly, which removes the nested control optimization that would otherwise run at every candidate layout. The framework is assessed on the Kriti 3 site off Crete, with 12 IEA 15 MW turbines on the VolturnUS-S semi-submersible and three mooring designs ranging from stiff to highly compliant.

Accounting for passive repositioning inside the layout optimization yields only a small median gain in annual energy production (AEP) over fixed-position optimization, though the gain increases with mooring compliance, and on the most compliant design the displacement-aware approach reaches a distinctly better layout than fixed-position optimization finds. The stronger effect is on the control side. On that same compliant design, yaw setpoints optimized under a fixed-position assumption turn a predicted gain into a net loss once applied to the moving platform, since the yaw-induced drift carries the rotors into wakes the fixed-position optimizer never evaluated. A displacement-aware optimization not only avoids this loss but unlocks further gain, using the same yaw-induced motion to steer the platforms clear of upstream wakes. The displacement response must therefore be accounted for inside the control optimization, and the requirement grows stronger the more compliant the mooring.

The two surrogates keep the layout-control co-design computationally tractable, where a nested formulation running the actual yaw optimizer and platform simulations inside the loop would be prohibitive. At this site and farm size, fixed-position, displacement-aware, and co-design optimization reach nearly the same AEP, with co-design adding little over displacement-aware layout optimization, in line with co-design results reported for bottom-fixed farms. Because many distinct layouts reach a near-equal AEP, the designer keeps the freedom to choose among them on criteria beyond energy capture, such as cabling or structural loads.

# Contents

<b>Summary</b>	<b>i</b>
<b>Nomenclature</b>	<b>viii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Background and Motivation . . . . .	1
1.2 Problem Statement . . . . .	2
1.3 Research Objective . . . . .	3
1.4 Thesis Outline . . . . .	3
<b>2 Literature review</b>	<b>5</b>
2.1 Wind Farm Layout Optimization Fundamentals . . . . .	5
2.1.1 Wake Effects and Wake Modeling . . . . .	5
2.1.2 Layout Optimization Problem Formulation . . . . .	6
2.2 Floating Wind Farm Specific Considerations . . . . .	7
2.2.1 Platform Types and Dynamics . . . . .	7
2.2.2 Mooring System Design . . . . .	9
2.2.3 Impact on Wake Interactions . . . . .	10
2.2.4 Modeling Approaches for FOWF . . . . .	10
2.2.5 Layout Optimization for FOWF . . . . .	12
2.3 Wind Farm Control Strategies . . . . .	13
2.3.1 Control Techniques Overview . . . . .	13
2.3.2 Wake Steering / Yaw Control . . . . .	13
2.3.3 Control for Floating Turbines . . . . .	15
2.3.4 Combined Control Strategies . . . . .	16
2.4 Layout and Control Co-Design . . . . .	16
2.4.1 Co-Design Problem Definition . . . . .	16
2.4.2 Single-Loop vs Nested Approaches . . . . .	16
2.4.3 Layout Control Co-Design for Floating Wind Farms . . . . .	17
2.5 Research Gaps and Thesis Positioning . . . . .	18
2.5.1 Identified Research Gaps . . . . .	18
2.5.2 Detailed Research Questions . . . . .	18
2.5.3 Thesis Contribution . . . . .	19
<b>3 Methodology</b>	<b>20</b>
3.1 Methodology Overview . . . . .	20
3.2 Floating Wind Turbine System . . . . .	21
3.2.1 Turbine and Platform Specifications . . . . .	21
3.2.2 Mooring Designs . . . . .	22
3.3 Platform Displacement Modelling . . . . .	23
3.3.1 Simulation Framework (SLOW) . . . . .	23

3.3.2	Dataset Generation and Symmetry Exploitation . . . . .	24
3.3.3	Displacement Surrogate (XGBoost) . . . . .	25
3.4	Wind Farm Model . . . . .	25
3.4.1	Wake Model . . . . .	26
3.4.2	Turbine Power Model . . . . .	26
3.5	Displacement-Aware Yaw Optimization . . . . .	26
3.5.1	The <code>m2sr_floating</code> Optimizer . . . . .	26
3.5.2	Role in the Thesis . . . . .	27
3.6	Displacement-Aware Layout Optimization . . . . .	29
3.6.1	Layout Optimization Genetic Algorithm . . . . .	29
3.6.2	Objective Function: Displacement-Aware AEP Wrapper . . . . .	29
3.7	Integrated Yaw Surrogate . . . . .	31
3.7.1	Dataset Generation . . . . .	31
3.7.2	Feature Engineering and Dataset Preprocessing . . . . .	33
3.7.3	XGBoost Surrogate: Training and Validation . . . . .	34
3.8	Integrated Single-Loop Co-Design . . . . .	35
3.8.1	Co-Design Pipeline . . . . .	35
3.9	Layout Comparison: the Layout Field Error . . . . .	36
3.10	Case Study: Kriti 3 . . . . .	37
3.10.1	Site Description . . . . .	37
3.10.2	Wind Resource Characterisation . . . . .	37
3.11	Computational Setup . . . . .	39
<b>4</b>	<b>Results</b>	<b>41</b>
4.1	Platform Displacement Patterns - Watch Circles . . . . .	41
4.2	Displacement-Aware Yaw Optimization on the Reference Layout . . . . .	43
4.2.1	Per-Mooring Yaw Optimization Gains . . . . .	43
4.2.2	Active Repositioning in the Semitaut_v2 Case . . . . .	45
4.3	Displacement Surrogate Validation . . . . .	47
4.4	Displacement-Aware Layout Optimization . . . . .	48
4.4.1	Layout Optimization Results . . . . .	48
4.4.2	Statistical AEP Comparison: Fixed vs. Displacement-Aware . . . . .	49
4.5	Integrated Yaw Surrogate Validation . . . . .	51
4.5.1	Yaw Surrogate Diagnostics . . . . .	51
4.5.2	Two-Turbine Study . . . . .	55
4.5.3	Investigating the Two-Turbine Ground-Truth Asymmetry . . . . .	59
4.5.4	Yaw Surrogate Farm-Scale AEP Validation . . . . .	62
4.6	Integrated Single-Loop Co-Design . . . . .	64
4.6.1	Statistical AEP Comparison: Fixed, Displacement-Aware, and Co-Design . . . . .	64
4.6.2	Spatial Layout Comparison via the Layout Field Error . . . . .	66
4.7	Summary of Results . . . . .	70
<b>5</b>	<b>Discussion</b>	<b>71</b>
5.1	Interpretation of Key Findings . . . . .	71
5.2	Limitations . . . . .	73
<b>6</b>	<b>Conclusions and Recommendations</b>	<b>75</b>
6.1	Answers to the Research Questions . . . . .	75
6.2	Contributions . . . . .	77
6.3	Recommendations for Future Work . . . . .	77

<b>Acknowledgements</b>	<b>79</b>
<b>A Surrogate Model Hyperparameters</b>	<b>80</b>
<b>B Displacement Surrogate Diagnostics</b>	<b>81</b>
<b>C Optimization Run Statistics</b>	<b>83</b>
<b>D Underlying Layouts of the Layout Field Error Maps</b>	<b>84</b>

# List of Figures

1.1	Bathymetry of the European seas . . . . .	1
1.2	IEA 15 MW turbine on the VoltornUS-S platform . . . . .	2
2.1	Wake effect at the Horns Rev wind farm . . . . .	5
2.2	The four main floating platform concepts . . . . .	8
2.3	Six degrees of freedom of a floating wind turbine . . . . .	9
2.4	Iterative wake–displacement coupling procedure . . . . .	12
2.5	Wake steering versus greedy operation . . . . .	14
2.6	Control-related effects for floating wind farms . . . . .	15
2.7	Nested versus integrated single-loop co-design . . . . .	17
3.1	Methodology overview and data flow . . . . .	20
3.2	VoltornUS-S platform and three-line mooring configuration . . . . .	22
3.3	Reference mooring designs adapted for Kriti 3 . . . . .	23
3.4	Example SLOW platform-response time series . . . . .	24
3.5	Symmetry-based reconstruction of the semिताut displacement dataset . . . . .	25
3.6	Displacement-aware yaw optimizer (m2sr_floating) schematic . . . . .	28
3.7	Displacement-aware AEP wrapper schematic . . . . .	30
3.8	Training-database generation for the yaw surrogate . . . . .	32
3.9	Feature engineering for the integrated yaw surrogate . . . . .	33
3.10	Integrated single-loop co-design pipeline . . . . .	35
3.11	Candidate offshore wind areas in Greece by foundation technology . . . . .	38
3.12	Environmental documentation for the Crete candidate areas . . . . .	38
3.13	The Kriti 3 site and Poseidon measurement buoy . . . . .	39
3.14	Kriti 3 wind rose at 150 m hub height . . . . .	40
4.1	Watch circles for the three mooring configurations . . . . .	42
4.2	Optimized-yaw wake maps for the three moorings . . . . .	44
4.3	Three-way wake-map comparison (semिताut_v2) . . . . .	46
4.4	Displacement surrogate diagnostics (semिताut_v2) . . . . .	48
4.5	Best displacement-aware layout and GA convergence (semिताut) . . . . .	49
4.6	Layout-only AEP: fixed-bottom versus displacement-aware . . . . .	50
4.7	Predicted versus actual optimal yaw, per mooring . . . . .	53
4.8	Yaw-surrogate error and power gain by wind speed . . . . .	54
4.9	Two-turbine optimal-yaw comparison (taut) . . . . .	56
4.10	Two-turbine optimal-yaw comparison (semिताut) . . . . .	57
4.11	Two-turbine optimal-yaw comparison (semिताut_v2) . . . . .	58
4.12	Control watch circles: shape versus magnitude . . . . .	60
4.13	Two-turbine ground-truth yaw for the control moorings . . . . .	61
4.14	Farm-scale AEP-gain validation of the yaw surrogate . . . . .	63
4.15	Three-way AEP comparison across moorings . . . . .	65

4.16	Layout field error (taut) . . . . .	67
4.17	Layout field error (semitaut) . . . . .	68
4.18	Layout field error (semitaut_v2) . . . . .	69
B.1	Displacement surrogate diagnostics (taut) . . . . .	81
B.2	Displacement surrogate diagnostics (semitaut) . . . . .	82
D.1	Layout field error with all turbines overlaid (semitaut) . . . . .	84

# List of Tables

2.1	Simulation tools for floating offshore wind turbines . . . . .	11
4.1	Per-mooring yaw-optimization results on the reference layout . . . . .	43
4.2	Farm power for the three semitaut_v2 yaw-control cases . . . . .	45
4.3	Displacement surrogate accuracy per mooring . . . . .	47
4.4	Layout-only AEP: distributions and median gains . . . . .	51
4.5	Integrated yaw surrogate accuracy per mooring . . . . .	52
4.6	Farm-scale AEP-gain validation statistics . . . . .	62
4.7	Three-way AEP comparison: median and spread . . . . .	65
4.8	Median AEP gains for the three-way comparison . . . . .	66
A.1	XGBoost hyperparameters for the two surrogates . . . . .	80
C.1	Full layout-only statistics . . . . .	83
C.2	Full three-way statistics . . . . .	83

# Nomenclature

## Abbreviations

AEP	Annual Energy Production
CAPEX	Capital Expenditure
CV	Cross-Validation
DOF	Degrees of Freedom
DTU	Technical University of Denmark
FOWF	Floating Offshore Wind Farm
FOWT	Floating Offshore Wind Turbine
GA	Genetic Algorithm
HCMR	Hellenic Centre for Marine Research
HPC	High-Performance Computing
IEA	International Energy Agency
LCOE	Levelized Cost of Energy
LES	Large-Eddy Simulation
LFE	Layout Field Error
LO-GA	Layout Optimization Genetic Algorithm
MAE	Mean Absolute Error
ML	Machine Learning
NLR	National Laboratory of the Rockies (formerly National Renewable Energy Laboratory, NREL)
NTUA	National Technical University of Athens
SLOW	Simplified Low-Order Wind turbine (simulator)
TLP	Tension-Leg Platform
WFLO	Wind Farm Layout Optimization
WMO	World Meteorological Organization
XGBoost	eXtreme Gradient Boosting
YITuR	Yaw and Induction-based Turbine Repositioning

## Symbols

$(x_i, y_i)$	Coordinates of turbine $i$ (m)
$\alpha$	Line length factor (-)
$\beta$	Anchor angle ( $^\circ$ )
$\gamma^*$	Optimal yaw angle vector ( $^\circ$ )
$\Delta x, \Delta y$	Platform displacement in the global farm frame (m)
$\gamma$	Yaw angle ( $^\circ$ )
$\gamma_\Delta$	Yaw misalignment angle, SLOW-frame displacement-surrogate input ( $^\circ$ )
$\gamma_{nacelle}$	Nacelle yaw angle ( $^\circ$ )
$\mathbf{F}_{mooring}$	Net mooring restoring force (N)
$\mathbf{F}_{thrust}$	Horizontal aerodynamic thrust force (N)
$\mathbf{x}, \mathbf{y}$	Turbine coordinate vectors (m)
$\mathcal{A}$	Feasible site area
$\mathcal{E}$	Exclusion zones
LF	Layout field (-)
TI	Ambient turbulence intensity (-)
$\phi_p$	Platform orientation ( $^\circ$ )
$\psi$	Rotation offset in the mooring-symmetry reconstruction ( $^\circ$ )
$\sigma$	Smoothness parameter of the layout field (rotor diameters)
$\theta$	Wind direction relative to the platform orientation ( $^\circ$ )
$A$	Weibull scale parameter ( $\text{m s}^{-1}$ )
$C_T$	Thrust coefficient (-)
$D$	Rotor diameter (m)
$d_x$	Streamwise inter-turbine separation (rotor diameters)
$d_y$	Crosswind inter-turbine separation (rotor diameters)
$d_{min}$	Minimum inter-turbine spacing (m)
$f_{ij}$	Joint occurrence probability of wind condition $(wd_i, ws_j)$ (-)
$k$	Weibull shape parameter (-)
$L$	Anchor distance from the platform (m)
$N_{wd}$	Number of wind-direction bins (-)
$N_{ws}$	Number of wind-speed bins (-)
$N_{wt}$	Number of turbines (-)
$P_k$	Power of turbine $k$ (W)

$P_{ij}$	Total farm power in wind condition $(i, j)$ (W)
$r$	Pearson correlation coefficient (-)
$R^2$	Coefficient of determination (-)
$U$	Wind speed ( $\text{m s}^{-1}$ )
$U_{\text{eff}}$	Effective (wake-affected) wind speed at the rotor ( $\text{m s}^{-1}$ )
$U_{\text{free}}$	Free-stream wind speed ( $\text{m s}^{-1}$ )
$W$	Neighbor-slot occupancy flag (-)
$z_0$	Surface roughness length (m)

# 1

## Introduction

### 1.1 Background and Motivation

Offshore wind has become a central component of decarbonization strategies, driven by rising turbine ratings, maturing supply chains, and a wind resource that is stronger and steadier than on land [1]. Deployment so far, however, has been concentrated in a narrow band of shallow water. Bottom-fixed foundations such as monopiles and jackets become technically challenging and economically unattractive beyond depths of roughly 50–60 m [1], which has confined installed capacity to shelf seas such as the North Sea and the Baltic. Figure 1.1 shows why this constraint matters at the European scale. Outside the northern shelf the seabed drops away within a few kilometres of the coast, and the Mediterranean basins reach depths of several hundred to several thousand metres. A large share of Europe’s offshore wind resource therefore sits over water too deep for bottom-fixed foundations [2, 3].




Figure 1.1: Bathymetry of the European seas, from the EMODnet Digital Bathymetry (DTM 2024) [4]. The shallow shelf seas where bottom-fixed offshore wind has been deployed, such as the North Sea and the Baltic, contrast with the much deeper Mediterranean basins, where floating support structures are required.

Floating offshore wind turbines (FOWTs) remove this depth constraint by decoupling the support structure from the seabed. Instead of resting on a foundation, the turbine is mounted on a buoyant platform held on station by mooring lines, an arrangement that remains feasi-

ble at depths of several hundred metres or more [1, 2]. Figure 1.2 shows the reference system used throughout this thesis: the IEA 15 MW reference wind turbine mounted on the UMaine VoltturnUS-S semi-submersible platform. Floating technology thereby opens large new areas for offshore wind development, including the deep basins visible in Figure 1.1 [2,3]. The same design choice that enables deep-water deployment, however, also changes the nature of the machine. A moored platform is compliant by design, and the position of the turbine is no longer a fixed quantity.



Figure 1.2: Rendering of the IEA 15 MW reference wind turbine mounted on the UMaine VoltturnUS-S semi-submersible floating platform [5], the reference system used throughout this thesis. Image: University of Maine Advanced Structures and Composites Center [6].

At the wind farm scale, performance is governed by more than the individual machine. Each rotor extracts momentum from the flow and leaves behind a wake, a region of reduced wind speed and increased turbulence that penalizes the turbines operating downstream. Wake effects can reduce the power of individual downstream turbines by up to 60%, and wake-induced losses at standard inter-turbine spacings are estimated at 20–30% of annual revenue [7]. The wake interactions also increase fatigue loading on the affected machines. These losses motivate two well-established families of methods for bottom-fixed farms. Wind farm layout optimization (WFLO) selects turbine positions that balance energy capture against wake losses within a site boundary [8], and farm-level control strategies such as wake steering intentionally yaw upstream turbines to deflect their wakes away from downstream rotors [9, 10]. Both approaches recover a meaningful fraction of the wake losses, typically a few percent of annual energy production (AEP) [7].

For floating wind, these farm-level gains carry particular economic weight. Capital expenditures and operational uncertainties remain considerably higher than for mature bottom-fixed projects, and the levelized cost of energy (LCOE) of floating wind must fall substantially for the technology to compete at scale [2]. Percentage-level improvements in AEP obtained through better layout and control decisions translate directly into revenue, and therefore into LCOE reductions, without additional hardware. There is consequently a strong incentive to develop layout and control optimization approaches that account for the motion of the floating platform. What that motion implies for the design problem is the subject of the next section.

## 1.2 Problem Statement

Unlike its bottom-fixed counterpart, a floating wind turbine does not hold its nominal position. The platform responds to the loads acting on it: the aerodynamic thrust on the rotor, wave

loads, and current loads. Balanced against the restoring force of the mooring system, these loads settle the platform at a mean offset from its anchored reference position [11]. This drift-induced displacement is referred to as passive repositioning throughout this thesis. The offset is not constant. The thrust-driven component varies with the operating point of the turbine, changing with wind speed, wind direction, and yaw misalignment [12], while the sea state and currents shift it further, and for compliant mooring designs the resulting excursions reach tens of metres [13]. Turbine separations, and with them the wake interactions across the farm, therefore become condition-dependent. A layout drawn at nominal positions is not the layout the wind actually sees.

Farm-level control deepens this coupling. On a bottom-fixed turbine, yaw misalignment serves a single purpose: the thrust vector rotates, and the wake is deflected away from downstream rotors [7]. On a floating turbine the same action has a second consequence. The cross-wind component of the thrust pushes the platform sideways, shifting the rotor itself to a new position [14]. This yaw-induced lateral displacement is referred to as active repositioning. Layout, motion, and control are thus linked in a closed chain: the layout sets the wake field, the wakes set the thrust on each rotor, the thrust sets the displacement of each platform, and the displaced positions form the effective layout in which the wakes must be re-evaluated [15].

Fixed-position assumptions therefore break down for floating wind farms. A layout optimizer that scores candidate layouts at their nominal coordinates evaluates configurations the turbines will not occupy during operation [12], and a controller optimized for a fixed geometry acts on a geometry that its own actions have changed. Treating layout and control as separable design problems therefore risks neglecting these coupled effects and missing potential gains [16]. An optimization framework that accounts for both passive and active repositioning inside the layout loop, evaluating every candidate at the positions the turbines will actually occupy under the control applied to them, can instead improve the achievable performance. Developing and assessing such a framework is the subject of this thesis.

## 1.3 Research Objective

**Main Research Question:** How does integrated layout and control co-design, accounting for passive and yaw-induced turbine repositioning, change the optimal layout and AEP compared to fixed-position approaches?

The research objective of this thesis is to develop an integrated layout and control co-design framework that accounts for the passive and yaw-induced repositioning of the turbines, and to assess it in a real-world case study by quantifying the changes in optimal layout and AEP relative to fixed-position approaches. The framework extends layout optimization so that yaw control and platform displacement are evaluated inside a single optimization loop, rather than treated as separate design stages. Following the literature review, the main research question is refined into four supporting subquestions at the end of Chapter 2.

## 1.4 Thesis Outline

The remainder of this thesis is organized in five chapters. Chapter 2 reviews the literature this work builds on: the fundamentals of wind farm layout optimization, the considerations specific to floating platforms, farm-level control strategies, and the co-design approaches that combine layout and control decisions. The chapter closes by identifying the research gap and refining the

main research question into four supporting subquestions. Chapter 3 develops the methodology. It sets out the modeling scope, defines the floating wind turbine system and its three mooring designs, and builds the platform displacement model and its surrogate. It then extends yaw and layout optimization to account for platform displacement and introduces the integrated yaw surrogate that enables single-loop co-design. The chapter closes with the Kriti 3 case study and its wind resource.

Chapter 4 presents the results, building from isolated components to the complete co-design framework: platform displacement patterns, yaw optimization gains on the reference layout, validation of both surrogates, displacement-aware layout optimization, and the full co-design comparison. Chapter 5 discusses the findings and the limitations of the modeling choices. Chapter 6 answers the research questions, states the contributions, and outlines directions for future work.

# 2

## Literature review

This chapter reviews the body of work that this thesis builds on. The review covers the fundamentals of wind farm layout optimization, the considerations specific to floating wind farms, and the control strategies available at the farm level. These threads converge in the literature on layout and control co-design, the framework this thesis extends to the floating case. The chapter closes by identifying the open research gap and positioning the thesis contribution against it.

### 2.1 Wind Farm Layout Optimization Fundamentals

#### 2.1.1 Wake Effects and Wake Modeling



Figure 2.1: A light fog reveals the wake effect behind turbines at Vattenfall’s Horns Rev wind farm off Denmark. Photo: Vattenfall

When a wind turbine extracts energy from the wind, it exerts a thrust force on the incoming flow. This creates a region of reduced wind speed and increased turbulence behind the rotor, known as the wake. The wake expands and recovers gradually with downstream distance as turbulent mixing entrains surrounding higher-speed air back into the deficit region. In a wind farm, wakes from upstream turbines propagate onto downstream machines, reducing their power output and increasing structural loading. Wake effects can reduce the power of individual downstream turbines by up to 60%, and wake-induced losses at standard inter-turbine spacings are estimated at 20–30% of annual revenue [7], making wake modeling a central element of layout design and control optimization.

Wake models span a wide range of fidelity. Engineering wake models describe the velocity deficit through fast analytical relationships at the low-cost end, medium-fidelity dynamic farm

models add unsteady wake behaviour, and high-fidelity computational fluid dynamics methods such as Reynolds-averaged Navier–Stokes and large-eddy simulation (LES) resolve the turbulent flow field directly [17]. A single high-fidelity evaluation requires hours of computation, whereas layout and control optimization requires thousands of evaluations per run, so these studies rely on engineering wake models and reserve higher-fidelity tools for verification.

Engineering wake models describe the velocity deficit behind a turbine using analytical relationships derived from momentum conservation. Because they are fast to evaluate, they are the standard choice for optimization problems that require large numbers of evaluations. Two popular formulations are the Jensen (Park) model [18], which assumes a uniform top-hat deficit profile with linear wake expansion, and the Gaussian model [19], which assumes a self-similar Gaussian deficit profile that better captures the radial variation of wind speed within the wake. Both models are parameterized by quantities such as the thrust coefficient, rotor diameter, and an empirical wake expansion rate that depends on atmospheric turbulence.

In addition to the single-wake deficit models, two further components are needed to describe the flow field across a full wind farm. First, a downstream turbine may lie in the overlapping wakes of several upstream machines. Superposition models address this by combining the individual deficits into a single effective velocity at each turbine location, with the root-sum-of-squares method [20] being the most common approach. Second, when a turbine operates at a yaw misalignment angle, the thrust vector develops a lateral component that deflects the wake away from its original centerline trajectory. Deflection models predict this lateral displacement as a function of the yaw angle and thrust coefficient [9]. This mechanism is the physical basis for wake steering control, discussed in Section 2.3.2.

### 2.1.2 Layout Optimization Problem Formulation

Wind farm layout is one of the most important decisions in a wind project’s design phase. Once a site is chosen, turbine locations must respect the site boundaries, exclusion zones, and minimum spacing requirements. Once the turbines are installed, their positions are fixed, and relocating them afterward is costly, if not impossible. The wake interactions described in the previous section couple the performance of every turbine to the positions of all others, making the layout a decisive driver of the farm’s lifetime energy yield. Even small improvements in AEP translate to improved project economics, strengthening the investment case of a wind farm. For these reasons, the wind farm layout design is formulated as a constrained optimization problem. In literature, different objective functions have been considered, like the AEP maximization, leveled cost of energy (LCOE) minimization, and capital expenditure (CAPEX) reduction [21, 22]. This thesis focuses on the maximization of the AEP.

The optimization treats turbine positions as continuous variables within the site area. Some studies discretize the domain into a grid of candidate locations, but continuous formulations have been shown to achieve better solutions and better reflect the physical flexibility available during installation [23]. The problem is stated as follows:

$$\begin{aligned}
\max_{\mathbf{x}, \mathbf{y}} \quad & \text{AEP}(\mathbf{x}, \mathbf{y}) \approx 8760 \sum_{i=1}^{N_{\text{wd}}} \sum_{j=1}^{N_{\text{ws}}} P_{ij}(\mathbf{x}, \mathbf{y}) \cdot f_{ij} \\
\text{subject to} \quad & (\mathbf{x}, \mathbf{y}) = \{(x_k, y_k)\}_{k=1}^{N_{\text{wt}}} \\
& (x_k, y_k) \in \mathcal{A}, \quad \forall k \in \{1, \dots, N_{\text{wt}}\} \\
& \|(x_k, y_k) - (x_l, y_l)\| \geq d_{\text{min}}, \quad \forall k \neq l \\
& (x_k, y_k) \notin \mathcal{E}, \quad \forall k \in \{1, \dots, N_{\text{wt}}\}
\end{aligned} \tag{2.1}$$

Here,  $N_{\text{wt}}$  denotes the number of turbines and  $(x_k, y_k)$  represent the Cartesian coordinates of turbine  $k$  in meters. The vectors  $\mathbf{x}$  and  $\mathbf{y}$  collect all turbine coordinates. The feasible site area is denoted by  $\mathcal{A}$ , while  $\mathcal{E}$  represents any exclusion zones. Wind conditions are discretized into  $N_{\text{wd}}$  direction bins and  $N_{\text{ws}}$  speed bins, indexed by  $i$  and  $j$  respectively. For each wind condition  $(i, j)$ ,  $P_{ij}(\mathbf{x}, \mathbf{y})$  gives the total wind farm power in watts, and  $f_{ij}$  represents its frequency of occurrence. The constant 8760 converts the weighted power to annual energy in watt-hours. The constraint  $d_{\text{min}}$  enforces a minimum inter-turbine distance, typically expressed as a multiple of the rotor diameter.

Lastly, a variety of optimization approaches have been tried to solve this problem. Gradient-based methods offer fast convergence when good initial layouts and smooth wake models are available, but they are susceptible to local optima. Population-based algorithms, genetic algorithms, particle swarm optimization, and their variants, explore the design space more broadly at higher computational cost. Hybrid strategies that combine global exploration with local refinement have also been proposed [21].

## 2.2 Floating Wind Farm Specific Considerations

### 2.2.1 Platform Types and Dynamics

Floating offshore wind turbines are deployed in water depths where fixed-bottom foundations become technically or economically infeasible, typically beyond 50–60 m [5, 24]. The turbine is mounted on a buoyant platform that is station-kept by a mooring system anchored to the seabed. Four main platform concepts have been developed, each with a different mechanism for achieving stability: the barge, the spar buoy, the semi-submersible, and the tension-leg platform [24].

A barge floats at the surface and gains its stability from the width of its hull, a design inherited from the oil and gas industry. Because most of the structure sits at the waterline, it takes the full force of the waves, and recent floating wind projects have largely moved away from the concept [24]. A spar buoy takes the opposite approach: a slender, heavily ballasted cylinder reaches far below the waterline, and the low-hanging weight pulls the system back upright whenever wind or waves tilt it. The penalty is a large material requirement and a draft that rules out shallower sites. A semi-submersible distributes its buoyancy over several columns spaced widely apart, with most of the structure kept below the wave zone. Spars and semi-submersibles are commonly anchored with catenary or semitaught mooring lines, which leave the platform some freedom to drift laterally under wind loading and control actions [24]. A tension-leg platform (TLP) is the exception: vertical tendons, pulled taut by the platform’s surplus buoyancy, pin it in place and permit only small displacements [24]. The four concepts are illustrated in Figure 2.2.

For the purpose of this thesis, the catenary-moored semi-submersible is the platform of interest. The combination of loose mooring restraint and column-based stability produces the largest horizontal offsets of the four concepts under aerodynamic and hydrodynamic loading. These

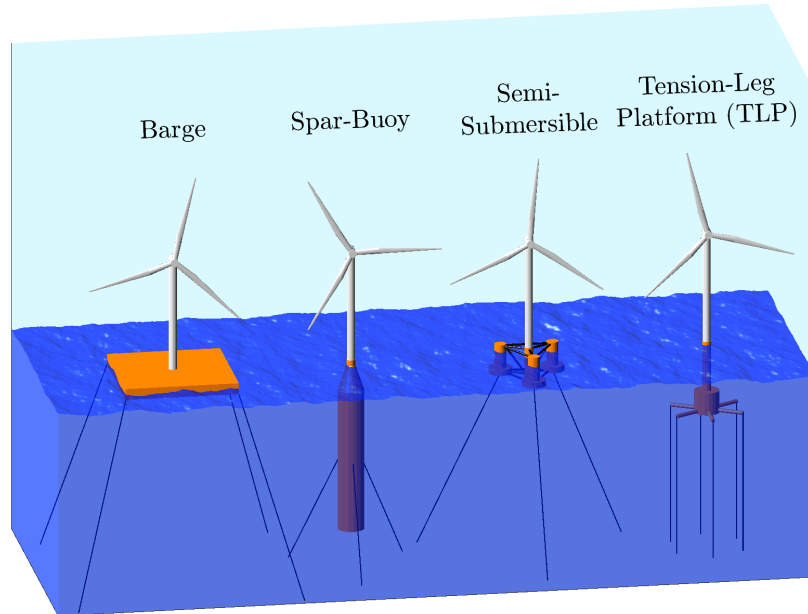


Figure 2.2: The four main floating platform concepts: barge, spar buoy, semi-submersible, and tension-leg platform, shown with a mounted DTU 10 MW reference turbine. Reproduced from Stockhouse et al. [25].

offsets are precisely what makes displacement-aware modeling and optimization relevant. The reference system used throughout this work is the IEA 15 MW offshore reference turbine [26] mounted on the UMaine VoltturnUS-S semi-submersible platform [5], a widely adopted benchmark configuration in floating wind research.

A floating platform moves in six rigid-body degrees of freedom (DOF): three translations (surge, sway, heave) and three rotations (roll, pitch, yaw), as illustrated in Figure 2.3. In operation, all six are excited simultaneously by wind, wave, and current loads, but they are not equally important for wake interactions. Heave, roll, pitch, and platform yaw change the hub height and the rotor orientation, yet they barely move the rotor in the horizontal plane, which is what sets the spacing between turbines.

Surge and sway are the degrees of freedom that directly change the effective spacing between turbines. When a platform displaces in the along-wind direction (surge) or the cross-wind direction (sway), the turbine physically moves to a new location within the farm. This shifts the wake it generates and changes how it interacts with wakes from neighboring turbines. For layout optimization and wake steering, these horizontal offsets are the dominant concern. This thesis therefore focuses on surge and sway displacements as the primary output of the platform motion model.

An important distinction must be made between mean platform offset and dynamic oscillations. Under steady wind loading, a floating turbine settles at a mean displaced position where the aerodynamic thrust is balanced by the mooring restoring force. Dynamic oscillations, driven by turbulence, wave forcing, and platform resonances, fluctuate around this mean position on timescales of seconds to minutes. The mean offset depends primarily on the wind speed, wind direction relative to the mooring system, and yaw misalignment angle. It can range from negligible at low wind speeds to a significant fraction of the rotor diameter near rated conditions.




Figure 2.3: The six rigid-body degrees of freedom of a floating offshore wind turbine: surge, sway, and heave in translation; roll, pitch, and yaw in rotation. Reproduced from the AE4W31 Floating Offshore Wind Energy lecture notes [27].

For the steady-state wake models used in layout optimization, the relevant quantity is the mean equilibrium position, not the instantaneous dynamic state. This quasi-static representation treats the platform as settled at its mean offset for a given set of wind conditions and yaw settings. While this approximation neglects the effect of dynamic motion on wake meandering and time-varying wake overlap, it is consistent with the steady-state assumptions already embedded in the wake models themselves and enables the large number of evaluations required during optimization.

### 2.2.2 Mooring System Design

The semi-submersible platform considered in this thesis is held in place by a catenary mooring system. The horizontal restoring force that the mooring lines exert is nonlinear and depends on the direction of the platform offset. Under steady loading, the platform settles at the quasi-static equilibrium position where the aerodynamic thrust is balanced by this restoring force:

$$\mathbf{F}_{\text{thrust}} + \mathbf{F}_{\text{mooring}} = \mathbf{0}. \quad (2.2)$$

Here,  $\mathbf{F}_{\text{thrust}}$  denotes the horizontal thrust on the rotor and  $\mathbf{F}_{\text{mooring}}$  the net restoring force from the mooring lines. Because the thrust load changes with the wind speed and direction, the platform settles at a different equilibrium position for each condition. The set of positions it can occupy across all conditions is its watch circle.

### 2.2.3 Impact on Wake Interactions

The mean platform offset introduced in the previous subsection shifts the effective turbine position within the farm. Because turbine spacing controls wake overlap, this shift alters how much of each upstream wake reaches downstream rotors, and therefore changes the downstream power capture relative to a fixed-position assumption.

The magnitude and direction of the offset under a given wind condition also depend on the platform orientation  $\phi_p$ , the heading of the mooring system in the horizontal plane. Rotating  $\phi_p$  rotates this watch circle, redirecting the displacement that a given thrust load produces. A layout optimizer for a floating wind farm must therefore evaluate each candidate layout at the positions the turbines actually occupy during operation, not at their nominal locations. This requirement is the starting point for the framework developed in this thesis.

### 2.2.4 Modeling Approaches for FOWF

Simulating a floating offshore wind farm requires coupling phenomena across multiple physical domains. Each turbine interacts with incoming wind (aerodynamics), ocean waves and currents (hydrodynamics), mooring line tensions (statics and dynamics), flexible tower and blade deformation (structural dynamics), and the control system that regulates power and loads. Platform pitch affects rotor loading, which changes thrust, which shifts the platform, which alters mooring forces. This fully coupled aero-servo-hydro-elastic problem defines the modeling challenge for floating offshore wind turbines (FOWTs).

At the farm level, wake interactions add another layer. The wake generated by an upstream turbine propagates downstream, reducing wind speed and increasing turbulence for subsequent turbines. For land-based and fixed-bottom farms, steady-state wake models suffice for many layout and control studies. Floating systems complicate this picture: platform motion changes the rotor position over time, potentially affecting wake trajectories and downstream power capture [24].

### Modeling Tools for FOWTs

Tools for FOWTs span a wide fidelity range [28]. The choice depends on the application. Detailed design validation demands accuracy, while optimization studies that require thousands of evaluations demand speed. Large-eddy simulation (LES) coupled with structural solvers provides the most physically accurate representation of FOWT behavior, resolving turbulent flow structures and their interaction with the rotor and platform. However, computational cost makes LES impractical for iterative design or optimization, and it serves primarily for final verification and for validating lower-fidelity models [24].

Mid-fidelity aero-servo-hydro-elastic tools operate in the time domain and couple engineering models for each physical subsystem. OpenFAST (NLR) [29] has become the standard in both academia and industry [24], connecting submodules for blade element momentum aerodynamics, potential flow or Morison-equation hydrodynamics, finite-element structural dynamics, and control system logic. HAWC2 (DTU) [30] offers similar capabilities [24]. These tools enable nonlinear time-domain simulation and can linearize the dynamics to generate state-space models useful for control design. Computational cost remains significant, since simulating a single turbine through a range of wind conditions takes hours, but it stays tractable for design iteration on individual turbines.

For farm-level wake analysis, FLORIS (NLR) [31] and PyWake (DTU) [32] provide steady-state aerodynamic calculations suitable for layout optimization and control design [24]. These

tools implement analytical wake models (Jensen, Gaussian, and variants) that compute velocity deficits as functions of downstream distance, turbine operating point, and atmospheric conditions. Because they are rooted in steady-state analysis, platform dynamics are not directly captured, though NLR has been incorporating FOWT-specific steady-state effects such as mean platform tilt into FLORIS [24]. Dynamic farm simulations are possible through FAST.Farm [33], which couples multiple OpenFAST instances with a farm-level wake transport model and a supercontroller that dispatches control inputs to each turbine [24], but computational cost scales unfavorably with turbine count and simulation length.

Low-fidelity tools like RAFT (NLR) [34] use frequency-domain methods for rapid platform response estimation [24], and SLOW (sowento [35], originally developed at the University of Stuttgart [36]) provides steady-state solutions for early-stage design exploration. MoorPy (NLR) [37] provides quasi-static mooring analysis and is used extensively in coupled wake–mooring frameworks, as illustrated in Figure 2.4. These tools are particularly useful when coupled with steady-state wake models to approximate floating wind farm behavior without full time-domain simulation.

Table 2.1 summarizes the primary modeling tools used in FOWT research.

Table 2.1: Simulation tools for floating offshore wind turbines.

Tool	Developer	Fidelity	Availability
OpenFAST	NLR	Mid (time-domain)	Open source
HAWC2	DTU	Mid (time-domain)	Free (research)
FLORIS	NLR	Low (steady-state wake)	Open source
PyWake	DTU	Low (steady-state wake)	Open source
FAST.Farm	NLR	Mid (dynamic farm)	Open source
RAFT	NLR	Low (frequency-domain)	Open source
SLOW	sowento	Low (steady-state)	Licensed

### Steady-State Coupling of Wakes and Platform Offsets

Computing power production in a floating wind farm requires determining where each platform settles under the combined action of aerodynamic thrust and mooring restoring forces. This creates a coupled problem: wake deficits depend on turbine positions, but positions depend on thrust forces, which depend on the wind speeds that result from wake interactions.

Figure 2.4 illustrates the standard approach to resolving this coupling [11, 15, 38]. Given wind conditions, initial positions, and yaw angles, a wake model computes the effective wind speed and thrust coefficient at each turbine. Actuator disc theory yields the thrust force magnitude and direction. A quasi-static mooring solver then finds the platform position where aerodynamic thrust balances the mooring restoring force. These updated positions change the wake geometry, so the wake model must be re-evaluated. The loop continues until positions converge.

An early formalization of the iterative coupling approach is the Floating Offshore Wind Farm Simulator (FOWFSim) [11], which couples the steady-state FLORIS wake model with a two-dimensional platform displacement model. A similar scheme coupling FLORIS with MoorPy was adopted in [39], where the coupled model served both for training a reinforcement learning-based yaw control policy and for model-based state estimation during deployment.

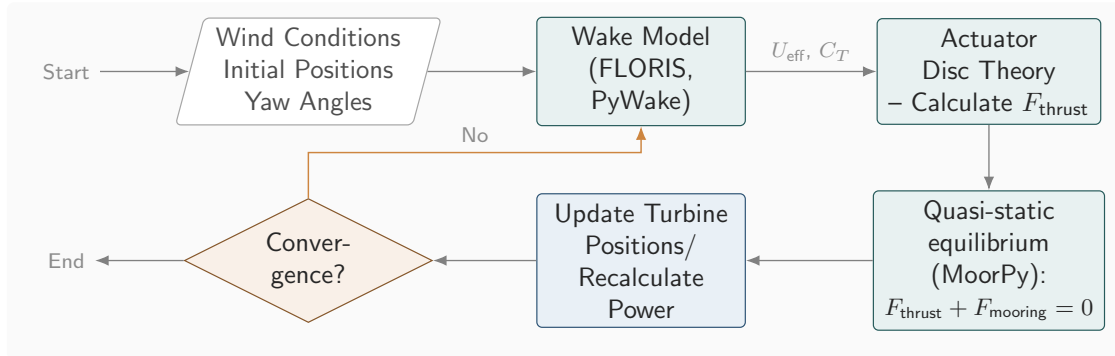


Figure 2.4: Iterative wake-displacement coupling procedure.

Each iteration involves wake deficit calculations across the farm and numerical solution of the mooring equilibrium for every turbine. For layout optimization, this process executes at every objective function evaluation, potentially thousands of times across different layouts, wind directions, and wind speeds. The costly step is computing each platform’s equilibrium displacement, which motivates surrogate models that predict it directly [15, 38].

### Surrogate Models for Wake-Offset Coupling

Surrogate models can replace the equilibrium calculation with a direct prediction of the platform’s displaced state. Riva et al. [15] developed such a surrogate for floating platforms, moving beyond the simplified engineering approximations used in earlier work. They trained neural networks on a large database of HAWC2 simulations spanning the operational envelope of wind speeds, directions, and current conditions, generated with a newly developed static solver to reduce the computational cost. The resulting models predict platform displacements and orientations significant enough to affect wake interactions. Where a physics-based calculation takes seconds, or even minutes, per evaluation, the surrogate returns the same result in microseconds. Feng et al. [38] applied these surrogates to floating layout optimization, optimizing the layout with the turbines treated as fixed and re-evaluating the resulting layouts with the surrogate to keep the search tractable, rather than embedding the displacement calculation inside the optimizer.

### 2.2.5 Layout Optimization for FOWF

The literature on layout optimization for floating wind farms remains limited and consists mostly of adaptations of fixed-turbine methods [40, 41]. Beyond the turbine coordinates  $(x_i, y_i)$ , the floating case introduces additional design variables, in particular the platform orientation  $\phi_p$  of each turbine and the mooring parameters [22, 42], which are often held fixed or treated on a scenario basis. The key challenge is that these design variables affect both the nominal and the operational (displaced) turbine position. Most existing studies assume fixed turbine positions and do not account for platform displacement [40, 41]. Hall et al. neglect it explicitly, on the grounds that taut moorings produce small and uniform offsets across the farm [22], an assumption that does not hold for the compliant mooring designs considered in this thesis. The few studies that do address displacement largely keep it outside the optimization loop, optimizing nominal positions and accounting for the offsets only when re-evaluating the resulting layouts [38, 43], while resolving it inside the loop has so far required a nested optimization [42].

## 2.3 Wind Farm Control Strategies

### 2.3.1 Control Techniques Overview

In most operational wind farms today the dominant control strategy is greedy control. Under greedy control, each turbine maximizes its own power production independently, without consideration of its effects on the wind field or on neighboring turbines. More sophisticated methods have emerged in research, aiming to manage wake interactions so that the wind farm as a whole produces more energy. The main wind farm control methods that have been studied are wake steering, power derating, and dynamic mixing concepts [7, 24, 44].

Wake steering, also known as yaw control, intentionally misaligns the rotor from the incoming wind to deflect the wake away from downstream turbines. Power derating reduces the thrust of upstream turbines by adjusting blade pitch or generator torque, weakening the wake at the cost of reduced upstream power extraction. Dynamic mixing concepts, such as the Helix approach [45], use periodic individual pitch control to enhance wake recovery through increased turbulent mixing, but these methods are still in early experimental stages. The remainder of this section focuses on wake steering, as it is the most mature farm-level control strategy and the one adopted in this thesis.

### 2.3.2 Wake Steering / Yaw Control

Wake steering manipulates the aerodynamic thrust vector by intentionally misaligning the rotor plane from the incoming wind direction. When a turbine operates at a yaw angle  $\gamma \neq 0^\circ$ , the thrust force develops a cross-wind component that deflects the wake laterally, steering it away from downstream turbines positioned in the crosswind direction. The deflection magnitude depends on the yaw angle magnitude, atmospheric stability, and turbulence intensity. Larger yaw misalignments produce stronger deflections but at the cost of reduced power extraction from the yawed turbine itself.

This creates the fundamental trade-off in wake steering: intentionally reducing upstream turbine power to gain disproportionately larger power increases downstream. The net farm benefit depends on wake strength, turbine spacing, and wind rose characteristics [21]. For closely spaced turbines under strong wake conditions, downstream gains can exceed upstream losses [7, 21].

Figure 2.5 illustrates this mechanism. Under baseline greedy control, each turbine maximizes individual power ( $\gamma = 0^\circ$ ), directing full wakes onto downstream turbines. With wake steering applied, the upstream turbine operates at a yaw misalignment, deflecting the wake and exposing the downstream turbine to higher wind speeds, increasing total farm output despite the upstream penalty.

Several practical factors complicate wake steering implementation. Wind direction uncertainty from turbulence and sensor limitations can reverse expected gains if yaw commands are applied incorrectly [7, 46]. Yaw misalignment also increases blade root and tower base loads through asymmetric rotor loading [47]. For floating platforms specifically, the yaw-induced cross-wind thrust component does more than deflect wakes. It displaces the entire platform, coupling aerodynamic control with structural repositioning. As a result, the platform can drift laterally back toward the wake even as the wake deflects away, partially canceling the intended benefit [14, 15].




Figure 2.5: Illustration of wake steering compared to baseline (greedy) operation.

### 2.3.3 Control for Floating Turbines

The mean horizontal offsets introduced earlier in this chapter, which depend on wind speed, wind direction, and yaw misalignment, are referred to in the floating-wind-farm control literature as passive repositioning [11, 43]. The term emphasizes that the displacement is not commanded by an actuator but emerges from the static balance between thrust and the mooring restoring force. The design implication is that the layout and the mooring system are no longer independent choices. By tuning mooring stiffness and anchor geometry together with the nominal turbine positions, the drift that each platform undergoes under typical operating conditions can be steered toward configurations that are more favorable for wake interactions across the farm.

A second repositioning mechanism is already present in the wake-steering picture of the previous subsection but has not yet been named. When a floating turbine yaws to deflect its wake, the same cross-wind thrust component that produces the deflection also shifts the platform laterally on its mooring lines. This yaw-driven lateral shift is called active repositioning, in contrast with the passive case where the offset arises from baseline thrust alone. The two effects are not independent. A fixed-position analysis accounts for only the aerodynamic wake deflection. It cannot capture the platform's resulting displacement, which changes the inter-turbine geometry at the same instant yaw is applied [11]. The net wake-steering gain on a floating turbine therefore differs from what such a model predicts.

Beyond passive offsets and yaw-driven lateral shifts, more direct actuation concepts are being explored. One such proposal is to install winches on each mooring line so that the line lengths, and therefore the platform position, can be commanded directly [17]. Concepts of this kind would in principle decouple repositioning from the aerodynamic trade-offs that yaw control entails, but they are still at an early stage of development. Figure 2.6 summarizes these effects and concepts, which together motivate treating the layout and the control strategy of a floating wind farm as parts of one joint design problem.




Figure 2.6: Control-related effects and concepts for floating wind farms. Illustration generated with Google Gemini.

### 2.3.4 Combined Control Strategies

The strategies described above can also be applied in combination. YITuR (Yaw and Induction-based Turbine Repositioning) is one such combined approach [11, 42]. Wake steering has also been paired with power derating and dynamic wake mixing, with combined strategies yielding gains beyond what any single strategy achieves on its own. Niu and Nagamune [44] demonstrate this for floating farms with reconfigurable layouts, while Baricchio et al. [48] report analogous results for fixed-bottom farms. These combinations are beyond the scope of this thesis, but they represent a natural extension of the co-design framework developed here.

## 2.4 Layout and Control Co-Design

### 2.4.1 Co-Design Problem Definition

Co-design is the design process in which multiple aspects of a system are designed together, accounting for their interdependencies. In contrast with traditional sequential design approaches, where each part of the system is designed independently, co-design exploits the interactions between subsystems so that design choices are made with awareness of their downstream effects.

Conventionally, control strategies are designed for an existing engineering solution, with the controller tasked with stabilizing the plant or tracking a reference. This sits at the final stage of a sequential design process. Control co-design instead is the joint design of the controller and the rest of the system, incorporating control considerations from the preliminary and early phases of the design. This allows engineers to better optimize the system as a whole, accounting for its control capabilities from the start.

At the wind farm level, the layout depends on the control strategy that will be used, and conversely, the control strategy depends on the layout. To date, however, the most common approach is to optimize the layout based on AEP, assuming each turbine operates under a greedy controller, and only then establish a farm-level control strategy. This sequential workflow leads to a suboptimal overall design. Since layout optimization is a system-design task fixed at the start of a project, the control strategy needs to be decided from the beginning. In this way, the layout can be jointly optimized for the specific site resource and for the fact that the wind farm will be actively controlled [16, 21]. This problem is high-dimensional, especially in real-world scenarios with large numbers of turbines, complex site boundaries, and exclusion zones.

### 2.4.2 Single-Loop vs Nested Approaches

Two approaches to layout-control co-design are found in the literature: the nested (also called two-stage) approach and the integrated (single-loop) approach.

In the nested approach, the problem is formulated as a two-level optimization. At every iteration of the outer layout loop, computing the objective function (typically the AEP) requires knowing the optimal control variables for the proposed layout. An inner control optimization is therefore called at each step to deliver these variables, for example the optimal yaw angles in the case of wake steering. Although this scheme is correct and accounts for control during layout optimization, it carries a heavy computational cost. For floating wind farms, this cost is compounded by the additional calculations required to update the turbine positions at each evaluation, making efficiency a critical concern.

The integrated, or single-loop, approach addresses this efficiency problem by replacing the expensive inner control optimization with a cheaper substitute. Common choices are geometric

yaw and machine-learning surrogates, both of which map a proposed layout directly to a set of near-optimal yaw angles. Geometric yaw, introduced by Stanley et al. [16], derives the angles analytically from normalised inter-turbine distances, so that yaw becomes a function of the positions being evaluated with no inner optimization call; it is adopted by Baricchio et al. [21] within the fixed-bottom layout co-design framework that this thesis extends. A machine-learning surrogate [49] instead predicts the near-optimal yaw directly from layout and wind-state features. The nested problem therefore collapses into a single loop: at every AEP evaluation, the surrogate supplies the near-optimal control variables in place of the inner optimization, speeding up the procedure substantially.

This efficiency argument carries additional weight for floating wind farms. As shown in Figure 2.4, computing the AEP for a floating layout already requires an iterative wake-displacement loop, since passive repositioning shifts each platform from its nominal position depending on the wind conditions and the yaw setting. Adding an inner control optimization on top of this loop produces a nested structure with three levels: the outer layout loop, the control optimization, and the wake-displacement equilibrium that must be re-solved at every yaw evaluation. Active repositioning strengthens this coupling further, because the yaw angle chosen by the controller directly modifies the platform position that the layout optimizer is targeting. The cost of the nested formulation grows accordingly, and a single-loop scheme in which both the optimal yaw and the resulting displacement are predicted by surrogates becomes a practical necessity rather than just an efficiency improvement.

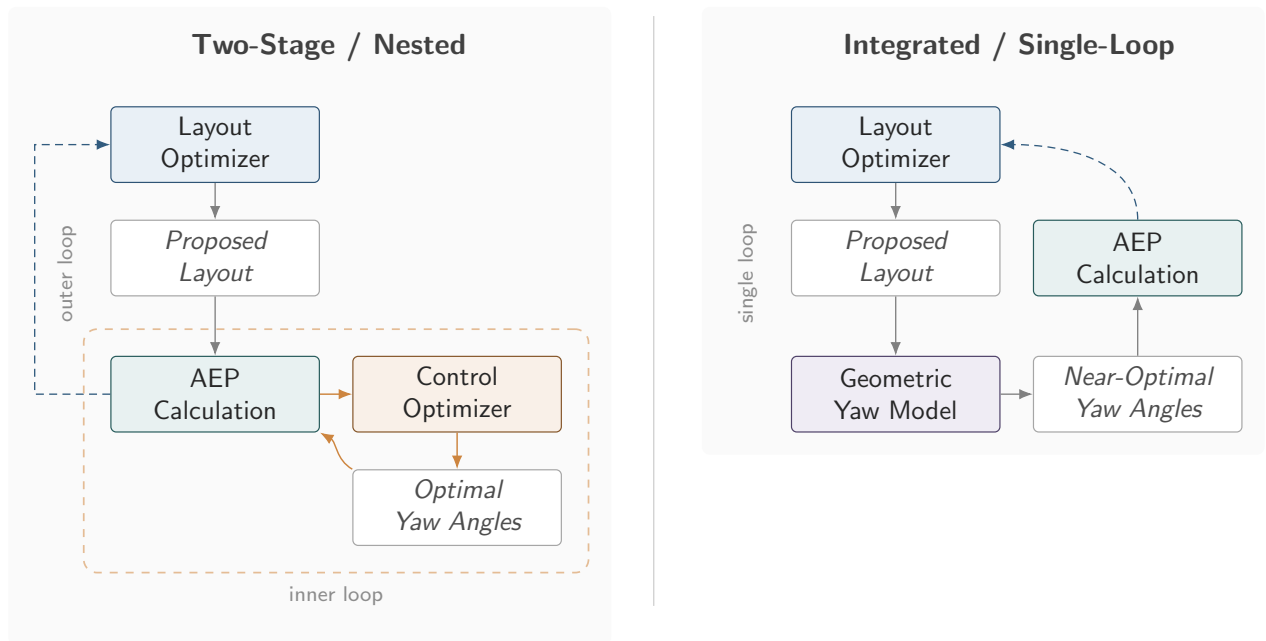


Figure 2.7: Schematic comparison of a two-stage/nested co-design approach and an integrated (single-loop) approach.

### 2.4.3 Layout Control Co-Design for Floating Wind Farms

The literature on layout-control co-design for floating wind farms remains very limited, and the work that does exist relies almost exclusively on nested optimization. A recent example is the two-stage scheme of Tao et al. [50], in which an outer layout optimizer and an inner control optimizer (acting on yaw angles and axial induction factors) are called in sequence at every evaluation. The model also accounts for passive platform displacement under varying wind conditions, reporting AEP improvements relative to a layout-only baseline. Despite this

advance, the underlying structure remains nested, with the inner optimization called inside every layout evaluation, and the resulting computational cost limits the scale of the studies that can be performed.

Floating wind farms introduce a difficulty that does not arise in fixed-bottom co-design: the layout and the control variables are coupled through platform motion. In a fixed-bottom farm, a chosen layout fixes the geometry that the wake model sees, and the control optimizer simply finds the yaw angles that maximize the AEP for that geometry. In a floating farm, this geometry is itself a function of the wind conditions and the yaw angles. The chosen layout sets the nominal positions, but the wakes that develop drive the thrust on each platform, the thrust drives the displacement, and the displaced positions are what the wake model must use at the next iteration. This cyclic dependency has to be resolved at every operating point before the AEP can be evaluated. Yaw control acts at two points in this loop simultaneously, deflecting the wake away from downstream turbines and shifting the platform laterally on its mooring lines through the cross-wind component of the thrust. Passive and active repositioning therefore combine with the wake–displacement coupling, and floating co-design has to address all three effects at once.

The gap, then, is the absence of an integrated single-loop scheme for floating co-design that handles both passive and active repositioning consistently. Existing nested formulations such as that of Tao et al. [50] account for passive displacement but still call an inner control optimization at every layout evaluation. Single-loop co-design is by now well established for fixed-bottom farms. Stanley et al. replace the inner control optimization with an analytical geometric-yaw relation [16], and later with a learned surrogate that predicts the optimal yaw directly [49]. A DTU study uses the same single-loop structure with a learned yaw surrogate, extending the co-design objective from AEP to market revenue [51]. All of these schemes assume the turbine positions stay fixed, so none transfers directly to a floating farm. This thesis addresses that gap by extending to the floating case the single-loop scheme of Baricchio [21], in which the inner control optimization is replaced by an analytical geometric-yaw relation that maps the layout directly to near-optimal yaw angles. This relation was developed for fixed-bottom turbines and assumes the turbine positions remain fixed, which does not hold on floating platforms. In the floating extension proposed here, a first surrogate predicts the platform displacement induced by the wind conditions and the yaw setting, so that the wake–displacement coupling is resolved at every AEP evaluation. A second surrogate then plays the role of the geometric-yaw relation, predicting the optimal yaw angles for each candidate layout directly. In this way, the cyclic dependency identified above is closed inside a single optimization loop, without the cost penalty of a nested formulation.

## 2.5 Research Gaps and Thesis Positioning

### 2.5.1 Identified Research Gaps

Across the literature reviewed above, one gap remains open. No published framework brings together integrated single-loop co-design with passive and active turbine repositioning for floating wind farms. Single-loop co-design exists, but only for fixed-bottom farms, where geometric-yaw or surrogate-based relations replace the inner control loop. Floating co-design exists, but only in nested form, and it accounts for passive displacement alone.

### 2.5.2 Detailed Research Questions

The gap identified above motivates the central research question of this thesis:

**Main Research Question:** How does integrated layout and control co-design, accounting for passive and yaw-induced turbine repositioning, change the optimal layout and AEP compared to fixed-position approaches?

The central question is addressed through four subquestions:

1. **RQ1 (Active repositioning):** How does accounting for active (yaw-induced) repositioning in yaw optimization change the corresponding yaw angles and the achievable wake-steering gains, compared with fixed-position yaw optimization?
2. **RQ2 (Passive repositioning):** How does accounting for passive repositioning alone (displacement-aware) change the optimal layout and the achievable AEP compared to fixed-position layout optimization?
3. **RQ3 (Optimized yaw surrogate):** How can geometric yaw be extended to floating turbines so that layout and control co-design becomes displacement-aware and remains single-loop?
4. **RQ4 (Mooring sensitivity):** How do different mooring designs affect the AEP gains and the optimal layouts?

The first question isolates the effect of active repositioning, the yaw-induced shift of the platform, on the achievable wake-steering gains. The second isolates the effect of passive repositioning on the layout problem, before control enters the picture. The third develops the framework that brings yaw control inside the optimization loop. The fourth examines how strongly the answers depend on the mooring design, the property that sets how much the platforms move.

### 2.5.3 Thesis Contribution

This thesis makes four contributions to the floating wind farm co-design literature:

1. An implementation of displacement-aware layout optimization for floating wind farms, together with the quantification of the effect of passive repositioning through a comparison against fixed-position optimization.
2. An extension of the geometric-yaw concept to floating turbines, realized through surrogate models trained on the results of the optimal yaw control algorithm.
3. The development of an integrated, single-loop co-design framework for floating offshore wind farms that accounts for both passive and active turbine repositioning, together with the quantification of its AEP gains relative to displacement-aware and fixed-position layout optimization.
4. A sensitivity analysis of the optimal layout and AEP gains with respect to the mooring design.

# 3

## Methodology

### 3.1 Methodology Overview

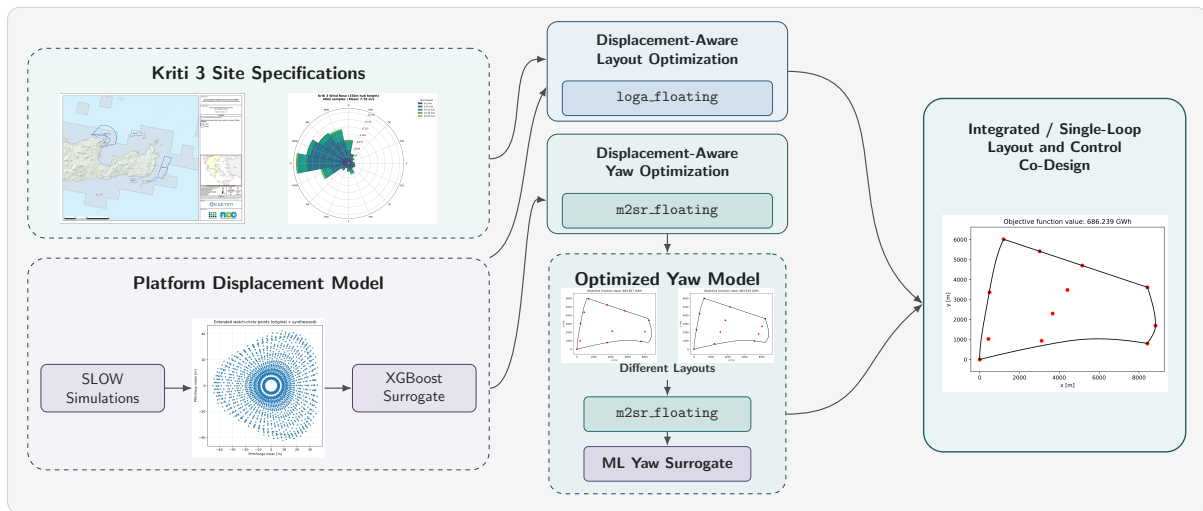


Figure 3.1: Methodology overview and data flow, from the Kriti 3 site and platform displacement model (left), through displacement-aware layout and yaw optimization (centre), to the integrated single-loop co-design optimizer (right).

Figure 3.1 summarises the methodology of this thesis and the data flow that connects its parts. The starting observation is that floating turbines do not hold their nominal positions. Under thrust loading each platform settles at a mean offset that depends on wind speed, wind direction, and yaw misalignment, and this offset feeds back into the wake interactions that determine farm power. Two mechanisms drive it. Passive repositioning is the drift induced by thrust acting against the mooring restoring force. Active repositioning is the additional lateral shift produced when a turbine yaws and the cross-wind component of thrust pushes the platform sideways. Accounting for both inside layout optimization is the central methodological problem of this thesis.

First, a fast predictor of passive repositioning (platform displacement) is built. Steady-state SLOW simulations across a grid of wind speed, wind direction, and yaw misalignment produce the watch-circle datasets shown in Figure 3.1, and a displacement surrogate of three XGBoost models [52], one per mooring, is fitted to predict mean surge and sway fast enough for the repeated evaluations optimization requires (§3.3).

The displacement surrogate enters two optimizers. **m2sr\_floating** computes optimal yaw angles for a fixed layout while accounting for the repositioning each yaw setting induces (§3.5), and **loga\_floating** embeds the surrogate inside the genetic-algorithm objective so that every candidate layout is scored at the positions the turbines will actually occupy, with yaw held at zero to

isolate passive repositioning (§3.6). Bringing active repositioning into the layout loop requires a second, distinct surrogate, since running `m2sr_floating` at every candidate evaluation would be prohibitively expensive. Instead, `m2sr_floating` is run offline to generate optimal-yaw labels, and an integrated yaw surrogate is trained to predict those angles directly from layout geometry and wind state (§3.7). The two surrogates serve distinct purposes and are not interchangeable. The displacement surrogate predicts where a platform sits, while the integrated yaw surrogate predicts how each turbine should be steered.

The final component combines both surrogates in a single optimization loop. In integrated single-loop co-design, each candidate layout is evaluated in three steps. The integrated yaw surrogate predicts its yaw angles, the displacement surrogate predicts the resulting displaced positions, and AEP is computed at those positions, with no nested control optimization (§3.8). This defines a progression of increasing physical fidelity, from fixed-position to displacement-aware to co-design, that also structures the results, with layout-only baselines first, then displacement-aware layouts, and finally the co-design comparison.

The framework itself is not tied to a particular wind farm. It requires only a site, a turbine, a floating platform, a set of mooring designs, and a wind resource, and applies to any floating wind farm defined by these inputs. In this thesis it is applied to the Kriti 3 site, using the IEA 15 MW reference turbine on the UMaine VoltturnUS-S platform and three mooring designs of increasing compliance (§3.2). The site and its wind resource are described in the case study at the end of the chapter (§3.10).

## 3.2 Floating Wind Turbine System

### 3.2.1 Turbine and Platform Specifications

The IEA Wind 15-Megawatt Offshore Reference Wind Turbine (IEA-15-240-RWT) is the hardware baseline for all simulations in this thesis [26]. The turbine has a rated power of 15 MW, a rated wind speed of 10.59 m/s, a rotor diameter of 241.94 m, a hub height of 150 m, and an IEC Class 1B design classification. The IEA-15-240-RWT is the community-standard reference turbine for floating offshore wind research, ensuring direct comparability with related work [53].

The UMaine VoltturnUS-S semi-submersible is the reference floating support structure paired with the IEA 15 MW turbine [5]. The platform is a four-column semi-submersible, comprising three radially offset buoyant columns of 12.5 m diameter with centres 51.75 m from the tower axis, and one central column carrying the turbine tower. The platform was originally designed for approximately 200 m water depth at the Gulf of Maine reference site. Its original mooring system is replaced in this thesis by three designs adapted to the Kriti 3 depth of approximately 550 m, described in §3.2.2.




Figure 3.2: IEA 15 MW reference turbine on VoltornUS-S semi-submersible platform. Left: System overview showing mooring lines. Right: Top view of the three-line mooring configuration with 120° fairlead spacing. Source: [5]

### 3.2.2 Mooring Designs

The UMaine VoltornUS-S was originally equipped with a three-line chain catenary mooring designed for approximately 200 m water depth at the Gulf of Maine. This system cannot be applied at Kriti 3, where the mean site depth is approximately 548.5 m. Three mooring designs adapted from Lozon et al. (2025) [53] are used instead, spanning a range of compliance from stiff to highly compliant. Studying this range enables a systematic examination of how mooring stiffness affects platform displacements, yaw-optimization gains, and the resulting optimal turbine layout. The mooring system enters as these three fixed designs compared in a sensitivity analysis and is not treated as a set of optimization variables.

The **taut** mooring is adapted from the Humboldt Bay reference design of Lozon et al. (2025) [53], originally specified for 800 m water depth, and uses a chain-polyester-chain configuration in tension. It is the stiffest of the three designs and produces the smallest platform displacements for a given thrust force, providing a near-fixed-bottom baseline for the displacement comparisons.

The **semitaut** mooring is adapted from the Gulf of Maine reference design of Lozon et al. (2025) [53], originally specified for 200 m water depth. It uses a polyester-top and chain-bottom configuration, placing it at an intermediate level of compliance between the taut and semitaut\_v2 designs.

The **semitaut\_v2** is a more compliant variant of the semitaut design, developed for this study, to showcase the effect of bigger displacements. It is not a published reference design. Its high compliance makes it the most sensitive of the three designs to yaw-induced repositioning.




Figure 3.3: The two published reference mooring designs adapted in this thesis: (a) the Humboldt design, basis for the taut mooring; and (b) the Gulf of Maine design, basis for the semitaut mooring. Reproduced from Lozon et al. (2025) [53].

### 3.3 Platform Displacement Modelling

The wake–displacement coupling is treated quasi-statically: at each evaluation the displaced equilibrium position is obtained from a single surrogate prediction, rather than by iterating a wake model and a mooring solver to convergence at runtime.

#### 3.3.1 Simulation Framework (SLOW)

All displacement training data are generated using SLOW (Simplified Low-Order Wind turbine), a floating offshore wind turbine simulator developed at the University of Stuttgart and extended by sowento GmbH for aero-hydro-servo-elastic simulation of floating wind systems [35]. The simulated system is the IEA 15 MW reference turbine on the UMaine VoltornUS-S semi-submersible (§3.2.1), and each of the three mooring designs (taut, semitaut, and semitaut\_v2) is treated as a separate simulation campaign producing an independent training dataset. Simulations are conducted under steady, uniform inflow wind with no waves or currents, following the steady-state assumption adopted in the SET4400 prior work [12]. Wave-induced platform motion is outside the scope of this study. Each simulation is parameterised by three inputs (wind speed  $U$ , wind direction  $\theta$ , and yaw misalignment  $\gamma_\Delta$ ), with the input ranges and parametric discretisation described in §3.3.2.

In order to define the nacelle yaw angle, considering the wind direction  $\theta$  is measured clockwise from the axis and yaw misalignment  $\gamma_\Delta$  defined as the counter-clockwise deviation of the nacelle from the wind direction. The nacelle yaw angle,  $\gamma_{\text{nacelle}}$ , is then:  $\gamma_{\text{nacelle}} = -\theta + \gamma_\Delta$ . The primary outputs recorded per simulation are mean surge and mean sway, both in metres: surge is the displacement along the  $\theta = 0$  reference direction and sway is the displacement perpendicular to it. Only these two mean offsets are retained, the degrees of freedom that change a turbine’s horizontal position within the farm. Dynamic oscillation about the mean and the remaining degrees of freedom are neglected for their limited effect on horizontal position, so the full aero-hydro-servo-elastic dynamic response is reduced to its quasi-static mean. Mean displacement is computed from the final 40% of the simulation time series, after the system has settled. For the taut and semitaut moorings, this corresponds to the last 400 s of a 1000 s run. For semitaut\_v2, the simulation is extended to 2500 s and the mean is taken from the final 1000 s, as the high compliance of this design produces substantially longer-settling transients. An example time series from a SLOW simulation is shown in Figure 3.4.




Figure 3.4: Example time series from a SLOW simulation over 1000 seconds. Top: surge. Middle: sway. Bottom: heave. Orange horizontal lines mark the steady-state mean used in the dataset.

### 3.3.2 Dataset Generation and Symmetry Exploitation

Three independent simulation campaigns are run, one per mooring design, over a common three-parameter grid. Wind speed takes nine values from 3 m/s to 11 m/s, wind direction 13 values from  $0^\circ$  to  $120^\circ$  at  $10^\circ$  resolution, and yaw misalignment 13 values from  $-30^\circ$  to  $30^\circ$  at  $5^\circ$  resolution, resulting in 1,521 simulations per mooring. The upper bound of 11 m/s is set just above the rated wind speed of 10.59 m/s, covering the below-rated regime where aerodynamic thrust is highest and platform displacements are most significant. The treatment of above-rated conditions is described further in §3.6.2.

The three-line mooring with fairleads equally spaced at  $120^\circ$  intervals has three-fold rotational symmetry: the displacement response in each  $120^\circ$  sector is a rotated copy of the adjoining sector. Simulations are therefore confined to  $\theta \in [0^\circ, 120^\circ]$ . The remaining two sectors are reconstructed by applying a 2D rotation to the recorded surge–sway displacement vector:

$$\begin{bmatrix} \text{surge}(\theta + \psi) \\ \text{sway}(\theta + \psi) \end{bmatrix} = \begin{bmatrix} \cos \psi & -\sin \psi \\ \sin \psi & \cos \psi \end{bmatrix} \begin{bmatrix} \text{surge}(\theta) \\ \text{sway}(\theta) \end{bmatrix} \quad (3.1)$$

Here,  $\psi$  is the rotation offset:  $\psi = -120^\circ$  for Sector B ( $\theta + 120^\circ$ ) and  $\psi = +120^\circ$  for Sector C ( $\theta + 240^\circ$ ). This reduction of the simulation count by a factor of three follows the procedure adopted in the SET4400 prior work [12] and is exact for these symmetric mooring configurations. Figure 3.5 illustrates this reconstruction for the semitaut dataset.

The three datasets share the same  $(U, \theta, \gamma_\Delta)$  grid, enabling direct displacement comparison across mooring designs. Displacement magnitudes differ substantially at  $U = 11$  m/s, the highest speed in the simulated range. The taut mooring reaches approximately 15.9 m ( $0.07D$ ), the semitaut mooring approximately 45.0 m ( $0.19D$ ), and the semitaut\_v2 mooring approximately 205.7 m ( $0.85D$ ), where  $D = 241.94$  m is the rotor diameter. These ranges span nearly two orders of magnitude and motivate the use of three mooring designs to systematically explore the effect of compliance on layout optimization outcomes.

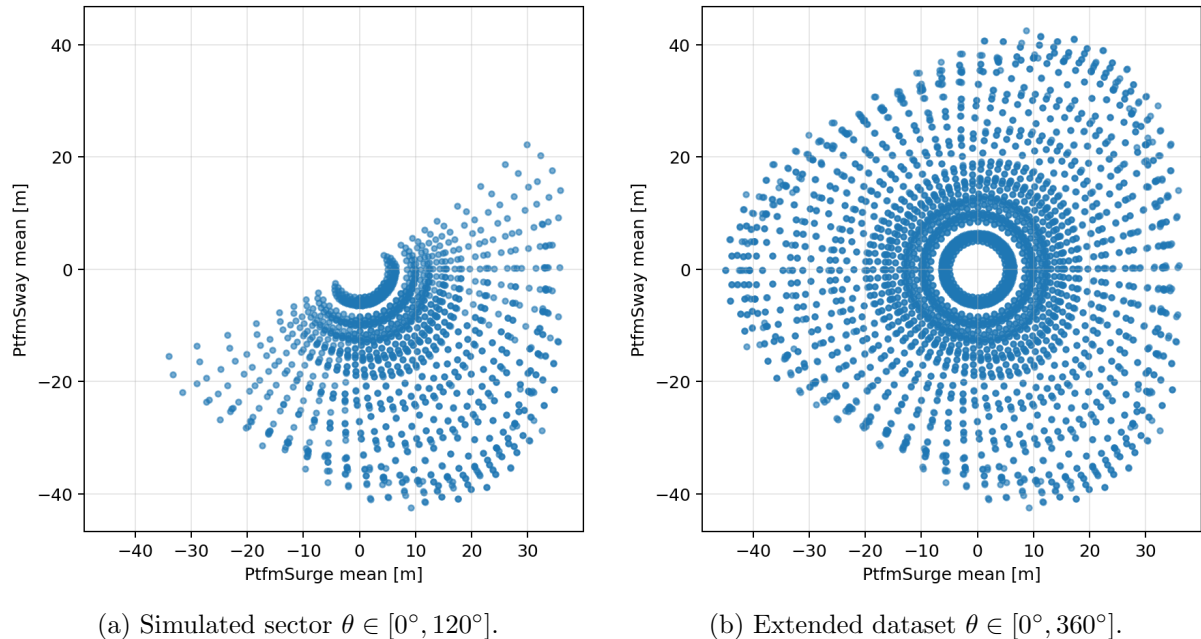


Figure 3.5: Symmetry-based reconstruction of the semitaut displacement dataset in the surge–sway plane. *Left*: the simulated sector  $\theta \in [0^\circ, 120^\circ]$ . *Right*: the full  $[0^\circ, 360^\circ]$  dataset obtained by the  $\pm 120^\circ$  rotations of Equation (3.1).

### 3.3.3 Displacement Surrogate (XGBoost)

Three independent XGBoost (eXtreme Gradient Boosting) displacement-prediction surrogates are trained in this thesis, one per mooring design, following the modelling approach established in the SET4400 prior work [12]. XGBoost constructs an ensemble of regression trees sequentially, with each tree correcting the residuals of its predecessor, yielding a gradient-boosted model well suited to structured tabular data [52]. In SET4400, this model family achieved the lowest prediction errors when compared against polynomial regression and random forest alternatives. The same approach is adopted for the new mooring datasets without repeating the comparison.

Each surrogate takes five input features per observation: wind speed  $U$ , and the trigonometric encodings  $\cos(\theta)$ ,  $\sin(\theta)$ ,  $\cos(\gamma_\Delta)$ , and  $\sin(\gamma_\Delta)$ . Encoding wind direction and yaw misalignment as sine-cosine pairs ensures the model inputs remain continuous. The output is a two-dimensional displacement vector (mean surge and mean sway, both in metres). The same set of hyperparameters is applied to all three surrogates; the values are listed in Appendix A (Table A.1), alongside those of the integrated yaw surrogate.

Each model is trained on the complete extended dataset for its mooring design, after the symmetry-based reconstruction described in §3.3.2. Predictive performance is assessed by five-fold cross-validation on the full dataset, supplemented by an independent 80/20 hold-out split. The resulting metrics for all three models are presented in §4.3.

## 3.4 Wind Farm Model

Farm power is evaluated with a steady-state wind farm model that maps a turbine layout and a wind condition to the power produced at each rotor. The model is assembled in PyWake from the Kriti 3 wind resource (§3.10.2), the IEA 15 MW reference turbine, and a set of engineering wake models. Its output, the farm’s annual energy production, is the objective that every optimizer

in this thesis maximizes.

#### 3.4.1 Wake Model

All farm-level AEP evaluations in this thesis use PyWake (DTU Wind Energy), an open-source Python framework for steady-state engineering wake modelling [32]. Engineering wake models are chosen over higher-fidelity approaches such as computational fluid dynamics or large-eddy simulation because they are fast to evaluate. This is a necessary property given that the genetic algorithm layout optimizer performs thousands of AEP evaluations per optimization run. PyWake is being used in wind farm layout optimization studies [21, 38], enabling direct comparability with related results.

The PyWake configuration used in this thesis comprises the `PropagateDownwind` wind farm model, the Bastankhah Gaussian wake deficit [19] for computing velocity deficits, root-sum-of-squares superposition (`SquaredSum`) for combining individual wakes, the Jiménez analytical wake-deflection model [9] for wake steering effects, and Gaussian overlap rotor averaging (`GaussianOverlapAvgModel`). The Bastankhah Gaussian deficit and the Jiménez deflection are standard choices for engineering-level layout and yaw-control optimization.

#### 3.4.2 Turbine Power Model

The wake model determines the wind speed reaching each rotor, and the turbine’s operating curves convert that wind speed into power and thrust. The power  $P(U)$  and the thrust coefficient  $C_T(U)$  are taken from the official IEA-15-240-RWT curves [26] and supplied to PyWake as a lookup table. The thrust coefficient sets the strength of the wake deficit and deflection of §3.4.1, while the power curve gives the electrical power produced at each rotor’s local wind speed. Power is zero below the cut-in speed of 3 m/s, rises with wind speed to the rated 15 MW, and holds at that value up to the cut-out speed of 25 m/s.

The farm power in a wind condition is the sum of the individual turbine powers. Annual energy production sums this over the wind rose, weighting each condition by its occurrence probability from the site Weibull model (§3.10.2) and scaling the total to the 8760 hours of a year. For a fixed layout this is a single vectorised PyWake evaluation. It is the baseline objective the displacement-aware optimizers in the following sections build on, evaluating the same farm power at the positions each platform settles into.

## 3.5 Displacement-Aware Yaw Optimization

The only control degree of freedom considered in this thesis is wake steering through intentional nacelle yaw misalignment.

#### 3.5.1 The `m2sr_floating` Optimizer

The yaw optimizer adopted in this thesis is `m2sr`, Baricchio et al.’s implementation [21] of the Serial-Refine yaw optimization algorithm [54]. Serial-Refine optimizes yaw angles sequentially, processing turbines from the most upstream to the most downstream. At each step, the yaw angle of one turbine is adjusted while the angles of all others are held fixed. Iterating through the farm in this order exploits the directed structure of wake interactions and produces a converged solution with far fewer function evaluations than a simultaneous search over all angles.

`m2sr_floating` extends `m2sr` to floating wind farms by embedding the displacement surrogate (§3.3.3) inside the Serial-Refine objective function. It was developed in the SET4400 prior re-

search project [12] and is adopted unchanged in this thesis. Two additional inputs are required relative to bottom-fixed `m2sr`: the trained XGBoost displacement surrogate and the platform orientation  $\phi_p$ , which is the meteorological angle of the first mooring line measured clockwise from north, which is held fixed for all turbines in the farm and is not a design variable in any optimization. The platform orientation determines the watch-circle geometry in the global farm frame and maps SLOW-frame displacements to global turbine coordinates. The SLOW-compatible wind angle  $\theta$  at each operating condition is computed as  $\theta = (\text{wd} - \phi_p) \bmod 360^\circ$ , where `wd` is the farm wind direction in meteorological convention (degrees clockwise from north) and  $\phi_p$  is the orientation angle of the first mooring line in the same convention. The modulo operation maps  $\theta$  to  $[0^\circ, 360^\circ)$ . A further consequence of floating operation is that the displacement surrogate requires effective (wake-affected) wind speed  $U_{\text{eff}}$  at each rotor rather than free-stream speed, which mandates an extra `PyWake` call inside each objective function evaluation.

At each wind condition within the Serial-Refine loop, four additional steps are executed before farm power is evaluated (following [12]). First, `PyWake` is called at the nominal (non-displaced) turbine positions to obtain  $U_{\text{eff}}$  at each rotor. Second, the XGBoost displacement surrogate is queried with the triplet  $(U_{\text{eff}}, \theta, \gamma_\Delta)$ , where  $\gamma_\Delta$  is the yaw misalignment angle, to obtain mean surge and sway in the SLOW reference frame. Third, the predicted (surge, sway) pair is rotated into the global farm frame using the platform orientation angle. Fourth, the rotated displacements are added to the baseline turbine coordinates and `PyWake` is re-evaluated at the updated positions to compute farm power. This procedure adds one extra `PyWake` call per wind condition per objective function evaluation relative to bottom-fixed `m2sr`.

#### 3.5.2 Role in the Thesis

`m2sr_floating` plays three roles across the thesis. It first serves as a standalone yaw-optimization baseline, computing displacement-aware optimal yaw angles for fixed Kriti 3 layouts under each of the three mooring configurations, with the resulting AEP gains reported in §4.2. It then acts as the ground-truth label generator for the integrated yaw surrogate, supplying the optimal yaw angles  $\gamma^*$  on which that surrogate is trained (§3.7). Finally, it provides the validation benchmark, since co-design layouts produced with the surrogate are re-evaluated by full `m2sr_floating` optimization and the difference between the two measures the surrogate’s approximation error (§4.5).

Using the same optimizer for all three roles is deliberate. The training labels and the validation benchmark then share one ground truth, so any gap between the surrogate and full optimization is a real approximation error and not just a difference between two optimizers. Running `m2sr_floating` for every candidate layout inside the genetic algorithm would be far too slow, and removing that cost is what the integrated yaw surrogate is for.

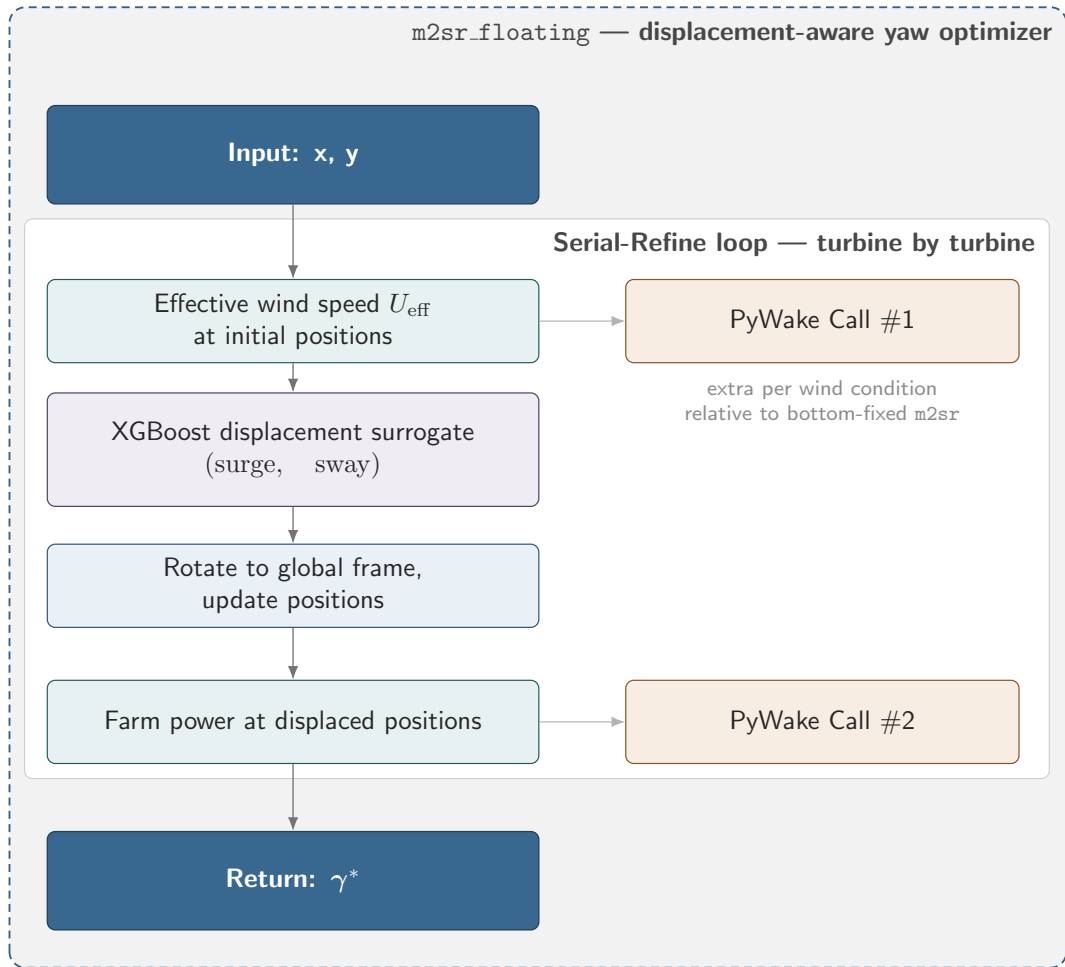


Figure 3.6: Schematic of `m2sr_floating`, the displacement-aware yaw optimizer. The Serial-Refine loop embeds the four-stage displacement-aware evaluation at each candidate yaw angle and returns the optimal yaw vector  $\gamma^*$ .

## 3.6 Displacement-Aware Layout Optimization

### 3.6.1 Layout Optimization Genetic Algorithm

The layout optimizer used throughout this thesis is the Layout Optimization Genetic Algorithm (LO-GA), developed by Baricchio et al. (2024) [21]. The published formulation targets bottom-fixed offshore wind farms. This thesis inherits the GA mechanics without modification and extends the framework to floating wind by replacing the original AEP objective function with the displacement-aware wrapper described in §3.6.2.

The wind farm layout optimization problem is non-convex with many local optima. The optimization variables are the Cartesian coordinates of each turbine, giving  $2N_{\text{wt}}$  continuous variables over the Kriti 3 site boundary. Three algorithmic choices distinguish the LO-GA from a standard continuous-variable GA. Initialization seeds the population from a regular grid layout and then applies targeted mutation to introduce diversity, ensuring the starting population is both feasible and varied. The crossover operator identifies structural similarities between parent layouts and recombines them in a way that preserves high-fitness turbine clusters, rather than combining positional chromosomes arbitrarily. Mutation uses an adaptive step size that decreases over generations, promoting broad exploration early in the run and fine-grained positional refinement as convergence approaches. These properties are inherited unchanged in the floating extension. The computational overhead of displacement-aware evaluation falls entirely on the objective function wrapper (§3.6.2).

### 3.6.2 Objective Function: Displacement-Aware AEP Wrapper

The displacement-aware AEP wrapper is the objective function evaluated for every candidate layout proposed by the genetic algorithm. A standard fixed-position AEP calculation places turbines at their nominal coordinates and calls PyWake once, vectorised over all wind conditions. For floating turbines this is insufficient: each platform settles at a displaced position that varies with wind speed, wind direction, and yaw misalignment, so nominal and actual positions diverge at every operating condition. The wrapper addresses this by substituting condition-specific displaced positions for nominal coordinates before each power evaluation, using the XGBoost displacement surrogate to predict those offsets without invoking SLOW at runtime. The resulting four-step procedure is illustrated in Figure 3.7.

Two geometric constraints are applied to nominal turbine positions before the wrapper evaluates AEP. The minimum spacing constraint assigns a power contribution of zero to any turbine placed within  $0.5D$  of its nearest neighbour [21]. The site boundary constraint requires all turbine positions to lie inside the Kriti 3 polygon, enforced by a polygon-containment test at each candidate evaluation.

After the spacing and boundary checks, the first step is a single vectorised PyWake call at equilibrium (non-displaced) turbine positions. This call evaluates the flow field simultaneously across all wind conditions. Yaw is set to zero for this call, establishing the equilibrium wake state from which the displacement model is subsequently queried.

The second step handles yaw assignment and is conditional on the chosen optimization mode. In displacement-only mode, yaw angles remain zero and this step is skipped entirely. In co-design mode, optimal yaw angles are computed for every wind condition, and a second vectorised PyWake call re-evaluates effective wind speeds at those angles before the displacement model is queried. This second call is physically necessary: the displacement surrogate was trained on SLOW data where the wind-speed input is the effective speed at the operating yaw angle. Using

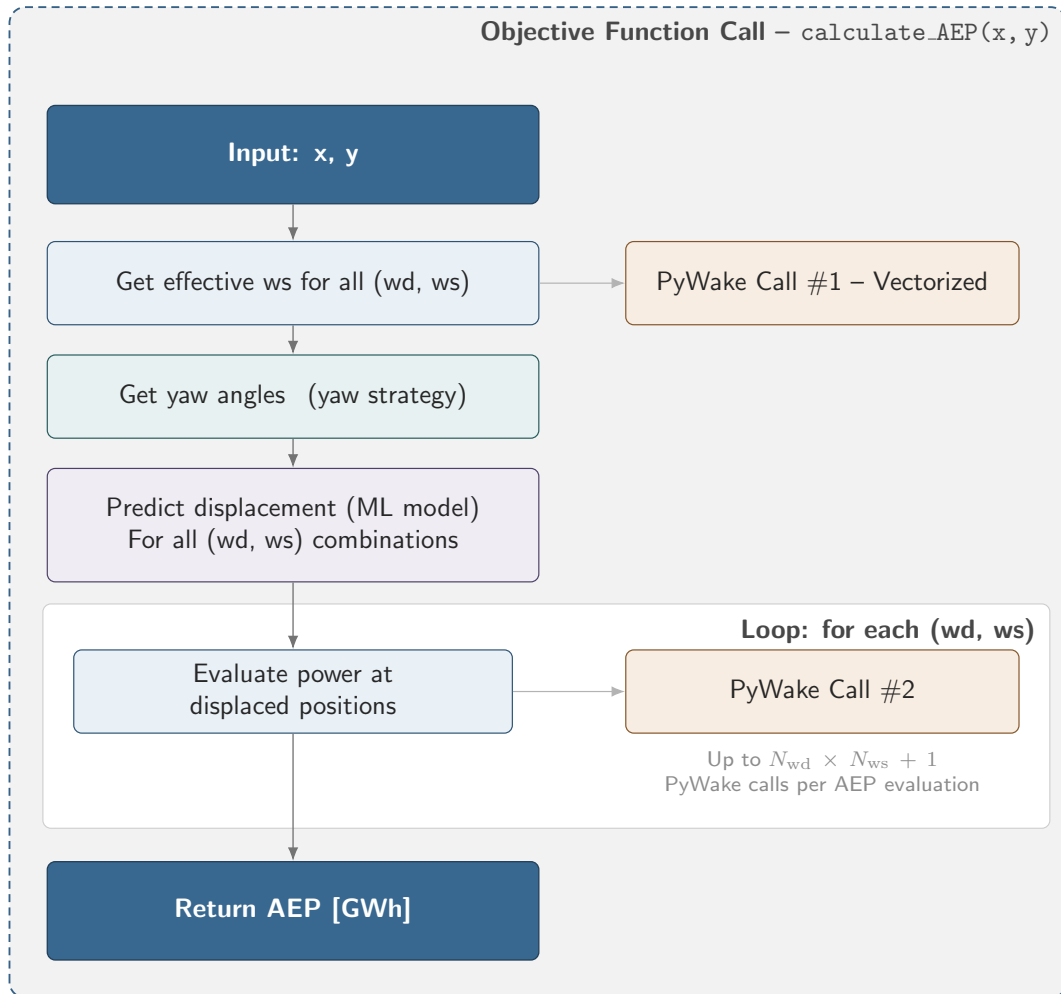


Figure 3.7: Displacement-aware AEP wrapper (`calculate_AEP`), the objective function evaluated for each candidate layout. Farm power is evaluated at surrogate-predicted displaced positions, using up to  $N_{wd} \times N_{ws} + 1$  PyWake calls per AEP evaluation.

the zero-yaw effective speed at a non-zero yaw setting would introduce a systematic error in the predicted platform offsets. The yaw strategy used in co-design mode is described in §3.7, and its integration into the GA objective function is described in §3.8.1.

The third step predicts platform displacements for all turbines and all wind conditions in a single batched XGBoost inference call. For each turbine and each wind condition, the model is queried with five features: effective wind speed (clamped to the rated value of 10.59 m/s for above-rated conditions), and the trigonometric encodings  $\cos(\theta)$ ,  $\sin(\theta)$ ,  $\cos(\gamma_\Delta)$ , and  $\sin(\gamma_\Delta)$ . The model returns surge and sway in the SLOW coordinate frame. These are rotated into global farm-frame offsets  $(\Delta x, \Delta y)$  using the platform orientation angle, following the coordinate transform described in §3.5.1.

With displaced positions precomputed for all wind conditions, a per-condition loop evaluates farm power at the displaced turbine positions. For each wind condition  $(wd_i, ws_j)$ , PyWake is called with the corresponding displaced positions and per-condition yaw angles. Per-turbine annual energy production is given by:

$$\text{AEP}_k = \sum_{i,j} f_{ij} \cdot P_k(\mathbf{x} + \Delta\mathbf{x}^{(i,j)}, \mathbf{y} + \Delta\mathbf{y}^{(i,j)}) \cdot 8760 \times 10^{-9} \quad (3.2)$$

Here,  $f_{ij}$  is the joint occurrence probability of wind condition  $(wd_i, ws_j)$  from the Weibull site model (§3.10.2), and  $P_k(\cdot)$  denotes the power of turbine  $k$  in watts at its displaced position. The vectors  $\mathbf{x}$  and  $\mathbf{y}$  collect the nominal turbine coordinates, and  $\Delta\mathbf{x}^{(i,j)}$  and  $\Delta\mathbf{y}^{(i,j)}$  are the global-frame displacement offsets predicted by the displacement surrogate for that wind condition. The factor  $8760 \times 10^{-9}$  converts the sum to gigawatt-hours, and the summation runs over all conditions. The farm AEP returned to the GA is  $\sum_k \text{AEP}_k$ .

The displacement-aware wrapper incurs approximately  $65\times$  the wall-clock time of an equivalent fixed-position AEP evaluation. This overhead is the direct motivation for cluster-based deployment (§3.11) and for the integrated yaw surrogate (§3.7), which replaces the per-evaluation yaw optimization with a single ML inference call.

## 3.7 Integrated Yaw Surrogate

Geometric yaw is the analytical single-loop substitute developed for fixed-bottom farms (§2.4.2), where it maps a layout directly to near-optimal yaw angles with no inner optimization call. For the floating problem addressed here, this thesis instead follows the machine-learning direction of Stanley et al. [49].

This thesis replaces the analytic relation with a learned surrogate. A set of XGBoost models is trained on `m2sr_floating` outputs to predict near-optimal yaw angles directly. The surrogate preserves the single-loop property, since yaw angles are predicted from wind-state and layout features at evaluation time with no nested optimizer call. A separate surrogate is trained for each mooring design, so that its predictions account for the displaced position each platform settles into. The dataset, feature construction, and training are described in the subsections that follow.

### 3.7.1 Dataset Generation

The integrated yaw surrogate is trained on (features, optimal yaw angle) pairs generated by running `m2sr_floating` across a structured campaign that varies three axes independently, layout

geometry, wind conditions, and mooring design. Each task is a single `m2sr_floating` call for one combination of layout, wind direction, wind speed, and mooring, returning an optimal yaw angle per turbine. A separate dataset is generated for each mooring, and a separate surrogate is trained from each. The campaign is summarised in Figure 3.8.

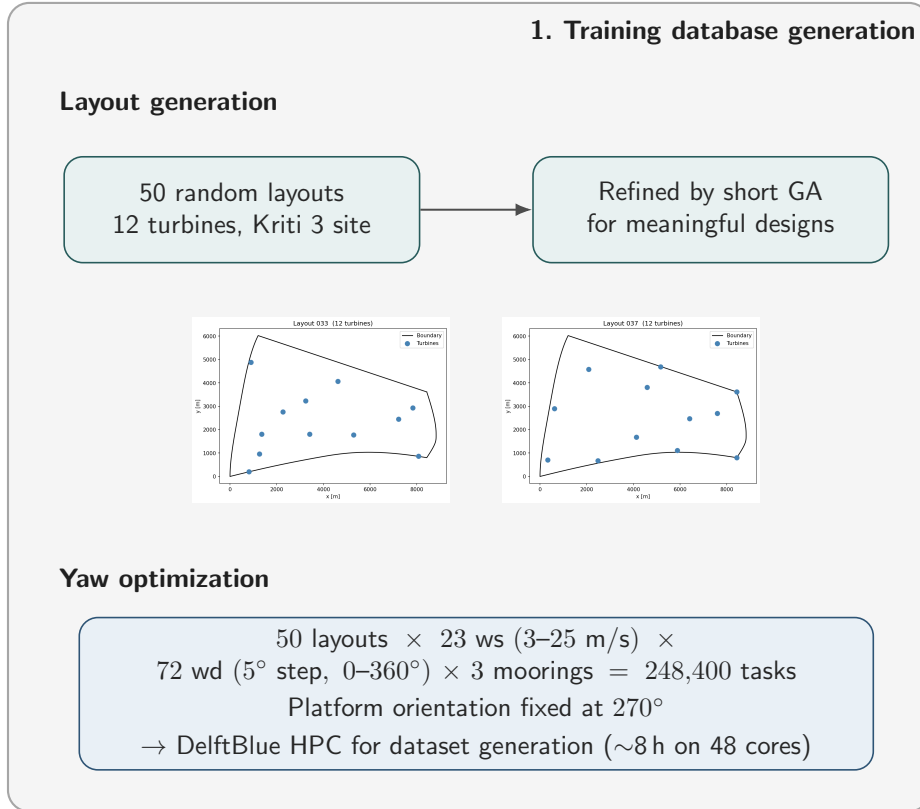


Figure 3.8: Training database generation for the integrated yaw surrogate: 50 randomly generated layouts are passed to `m2sr_floating` over the full wind grid and three mooring designs, giving 248,400 yaw-optimization tasks.

Layout diversity is important for the training dataset. Basically, a surrogate fitted exclusively on densely or uniformly spaced configurations would generalize poorly to the farm geometries the co-design optimizer explores. Fifty 12-turbine layouts are generated inside the Kriti 3 boundary polygon using a two-phase procedure. In the first phase, turbine positions are drawn uniformly at random inside the polygon. Each candidate is accepted only if it lies within the boundary and maintains at least  $2D$  separation from all already-placed turbines. In the second phase, each layout undergoes a short GA refinement of 10 generations with a population of 20, using a fixed-bottom AEP objective. This step nudges turbines toward configurations with meaningful wake interactions while preserving the spatial diversity introduced by random sampling. The platform orientation is fixed at  $270^\circ$  for all 50 layouts and all wind conditions, consistent with the rest of this work.

Wind conditions are discretised on a regular grid: 72 directions from  $0^\circ$  to  $355^\circ$  in  $5^\circ$  increments, and 23 wind speeds from 3 m/s to 25 m/s in 1 m/s steps. The  $5^\circ$  directional resolution matches the evaluation grid used in the displacement-aware AEP wrapper (§3.6.2), ensuring the surrogate is never queried at a wind direction outside its training distribution. The 3–25 m/s speed range spans from cut-in through well above rated (10.59 m/s), covering the below-rated regime where yaw steering is most effective and the above-rated regime where optimal yaw angles are expected to be near zero.

For each (layout, wind direction, wind speed) combination, PyWake is first called at zero yaw to extract the effective wind speed  $U_{\text{eff}}$  per rotor, the wind-speed input the displacement surrogate requires. Zero yaw is the reference state because the optimal angles are not yet known at inference, so matching it at training keeps the wind-speed input consistent across both. `m2srfloating` is then called with the displacement surrogate active, using a yaw search range of  $\pm 30^\circ$ , three refinement steps, and five candidate values per turbine per step, and the optimal yaw angle per turbine is recorded alongside its input features. The procedure runs independently for each mooring, producing one training dataset per mooring.

The campaign was run on DelftBlue [55], with its computational cost reported in §3.11.

### 3.7.2 Feature Engineering and Dataset Preprocessing

The raw training data produced in §3.7.1 pass through two preprocessing stages before training. Input feature construction encodes the local wake geometry and wind state, and output-side filtering zeroes the yaw labels for conditions where steering has no aerodynamic effect. Both stages are summarised in Figure 3.9.

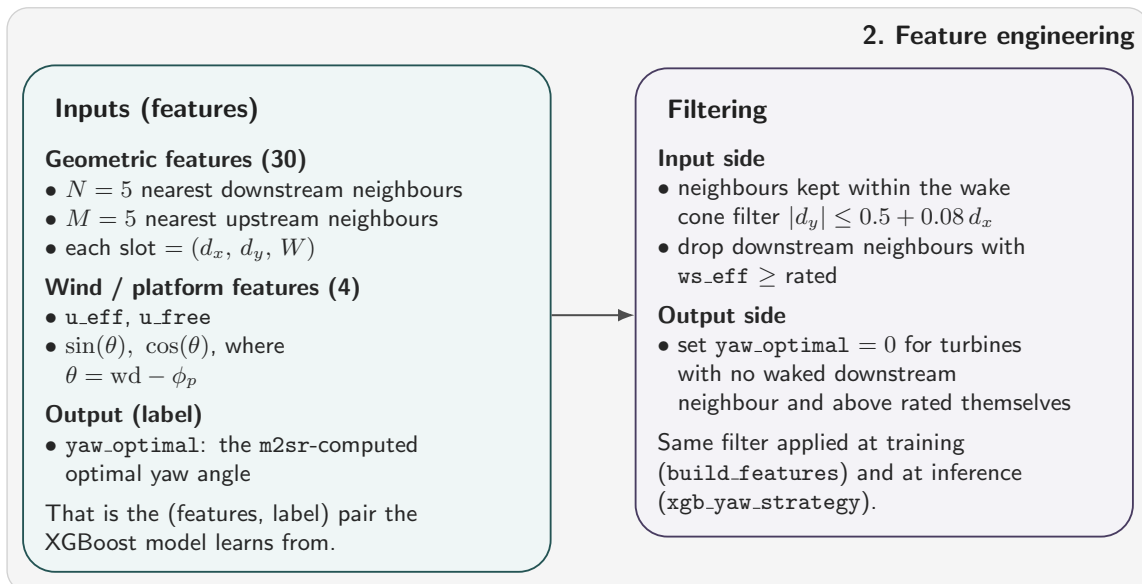


Figure 3.9: Feature engineering for the integrated yaw surrogate: each training row pairs a 34-dimensional input vector with the `m2sr`-computed optimal yaw label, with a two-sided filter applied identically at training (`build_features`) and inference (`xgb_yaw_strategy`).

Each training row encodes the local wind and platform state via four scalar features: the per-rotor effective wind speed at zero yaw ( $U_{\text{eff}}$ ), the free-stream wind speed ( $U_{\text{free}}$ ), and the wind direction encoded as  $\sin(\theta)$  and  $\cos(\theta)$ . Effective wind speed is derived from the zero-yaw PyWake call performed in §3.7.1, not from the optimal-yaw operating point. This matches the inference-time state: when the surrogate is queried inside the co-design loop, the optimal yaw angle is not yet known, so  $U_{\text{eff}}$  must be obtained from the pre-yaw flow field. Free-stream wind speed  $U_{\text{free}}$  provides the global wind magnitude, common to all turbines at a given condition. The wind direction  $\theta$ , defined relative to the platform orientation of  $270^\circ$ , is encoded via its sine and cosine.

The remaining 30 features encode inter-turbine wake geometry through a fixed-size neighbor slot structure. For each turbine  $i$ , the feature vector includes  $N = 5$  downstream slots and  $M = 5$  upstream slots, each carrying three values: the downstream separation  $d_x$  in rotor diameters,

the lateral offset  $d_y$  in rotor diameters, and a binary occupancy flag  $W$ . The total geometric feature count is  $3 \cdot N + 3 \cdot M = 30$ . Downstream slots list the up to five turbines lying within turbine  $i$ 's wake cone, sorted by distance  $d_x$  from nearest to farthest. Each occupied slot records  $d_x$ ,  $d_y$ , and  $W = 1$ ; turbines at or above rated wind speed are excluded from downstream slots, as pitch-regulated turbines are insensitive to yaw steering. Upstream slots list the up to five turbines upstream of turbine  $i$ , sorted by upstream distance.

Neighbors are assigned to downstream or upstream slots only if they satisfy the wake cone filter [21]:

$$|d_y| \leq 0.5 + 0.08 \cdot d_x \quad (3.3)$$

Here  $d_x$  and  $d_y$  are the inter-turbine separations in rotor diameters, measured in the wind-aligned reference frame of turbine  $i$ . The constant offset of 0.5 keeps turbines that are laterally offset but not genuinely downstream out of the slots, and the growth rate of 0.08 per rotor diameter approximates the lateral expansion of the wake with downstream distance. The same filter is applied identically at training and at inference. This consistency is deliberate, since a turbine excluded from a slot during training must be excluded by the same criterion at inference.

Not all yaw labels returned by `m2sr_floating` constitute meaningful training signal. When a turbine is above rated and has no downstream wake neighbor, yaw can do nothing useful. Pitch regulation already holds the turbine at rated power, so misalignment cannot improve its own output, and with no downstream turbine there is no wake to redirect. The farm-power objective is then flat in that turbine's yaw, and Serial-Refine returns an essentially arbitrary angle from its search grid rather than a meaningful optimum. Training on these noise labels would teach the surrogate spurious patterns, so they are set to  $0^\circ$ . A turbine below rated is left unchanged even with no downstream neighbor, since a small benefit from platform repositioning under thrust may still exist, and a turbine with a downstream neighbor is left unchanged even above rated, since the wake deficit it sheds persists at reduced thrust and can still be steered. The result is a heavily zero-dominated label distribution, with roughly 80% of rows carrying a zero label.

#### 3.7.3 XGBoost Surrogate: Training and Validation

Three independent XGBoost regression models are trained for the integrated yaw surrogate, one per mooring configuration (`taut`, `semitaut`, and `semitaut_v2`). XGBoost is chosen for the same reasons as the displacement surrogate (§3.3.3), since it predicts the optimal yaw angle accurately at low inference cost. Three separate models are trained because the optimal yaw angles depend on mooring compliance. A more compliant platform undergoes larger yaw-induced lateral displacement, and the optimizer exploits this effect at magnitudes and angular dependencies that differ from those of a stiffer design.

Each model is trained on the dataset produced in §3.7.1, comprising 993,600 rows per mooring ( $50 \text{ layouts} \times 72 \text{ directions} \times 23 \text{ speeds} \times 12 \text{ turbines}$ ). Each row carries the 34 input features defined in §3.7.2 and the single label `yaw_optimal`. Because roughly 80% of those labels are zero (§3.7.2), the model has to produce near-zero outputs for the large majority while still predicting angles up to  $\pm 30^\circ$  accurately for the active-steering minority.

The hold-out split is layout-stratified. Ten of the 50 layouts (20%) are withheld as the test set, with all wind conditions for those layouts excluded from training, so that held-out metrics reflect generalization to unseen farm geometries rather than interpolation within known configurations. Within the 40-layout training pool, 5-fold GroupKFold cross-validation groups folds by layout

identity, so no layout in a training fold appears in its validation fold. The final deployment model is trained on all 50 layouts, maximizing the training set for inference inside the co-design loop.

Hyperparameters (tree depth, learning rate, and minimum child weight) are tuned per mooring by grid search, using MAE on the layout-stratified validation fold rather than  $R^2$ , which the zero-dominated labels would inflate. The selected values are given in Appendix A (Table A.1).

Surrogate accuracy, including held-out test metrics, a two-turbine controlled validation, and farm-scale AEP comparison against `m2sr_floating`, is reported in Chapter 4.

## 3.8 Integrated Single-Loop Co-Design

### 3.8.1 Co-Design Pipeline

The co-design pipeline implements the single-loop architecture as an extension of the displacement-aware AEP wrapper (§3.6.2). The integrated yaw surrogate replaces the inner `m2sr_floating` call, so near-optimal yaw angles and condition-specific platform displacements are both predicted within a single objective function evaluation. The genetic algorithm therefore assesses each candidate layout without invoking any nested control optimization loop [16,21]. Figure 3.10 illustrates the resulting pipeline.

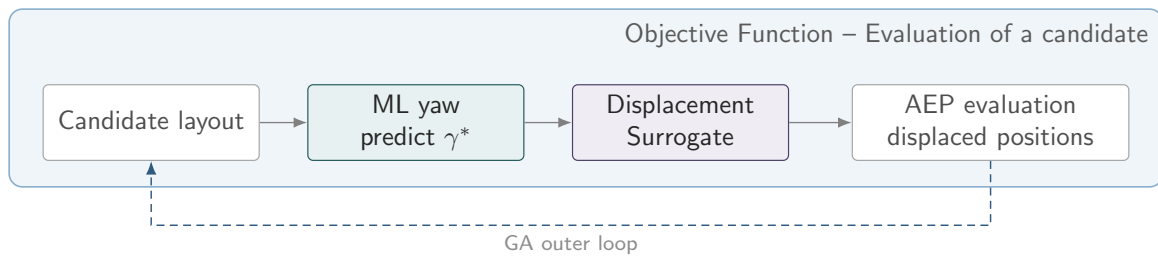


Figure 3.10: Integrated single-loop co-design pipeline: for each candidate layout, the ML yaw surrogate predicts  $\gamma^*$ , the displacement surrogate computes the resulting offsets, and AEP is evaluated at the displaced positions, all within one objective function call.

The co-design AEP evaluation extends the four-step wrapper of §3.6.2 by conditioning the displacement prediction on surrogate-predicted yaw angles. First, a single vectorised PyWake call at nominal positions yields per-turbine effective wind speeds  $U_{\text{eff}}$  across all wind conditions. Second, the XGBoost yaw surrogate (§3.7.3) predicts near-optimal yaw angles  $\gamma_i^*$  per turbine per condition from the 34-feature input vector (§3.7.2), with feature construction at inference, including the wake cone filter, identical to training. Third, the displacement surrogate is queried with the triplet  $(u_{\text{eff}}, \theta, \gamma_i^*)$  to obtain per-turbine surge and sway offsets  $(\Delta x_i, \Delta y_i)$  conditioned on the predicted yaw state. Fourth, farm power is evaluated at the yaw-displaced positions  $(x_i + \Delta x_i, y_i + \Delta y_i)$  under the surrogate-predicted angles. The resulting co-design AEP objective is:

$$\text{AEP}(\mathbf{x}, \mathbf{y}) = \sum_{i,j} f_{ij} \cdot P(\mathbf{x} + \Delta \mathbf{x}(wd_i, ws_j, \gamma^*), \mathbf{y} + \Delta \mathbf{y}(wd_i, ws_j, \gamma^*), \gamma^*) \cdot 8760 \times 10^{-9} \quad (3.4)$$

Here  $\mathbf{x}$  and  $\mathbf{y}$  collect the nominal turbine coordinates,  $\gamma^* = \gamma^*(\mathbf{x}, \mathbf{y}, wd_i, ws_j)$  is the surrogate-predicted yaw vector for condition  $(wd_i, ws_j)$ ,  $\Delta \mathbf{x}$  and  $\Delta \mathbf{y}$  are the displacement vectors returned

by the displacement surrogate conditioned on  $\gamma^*$ , and  $f_{ij}$  is the joint occurrence probability from the Weibull site model (§3.10.2). The predicted yaw vector enters the power evaluation in two distinct ways. It conditions the platform displacements  $\Delta\mathbf{x}$  and  $\Delta\mathbf{y}$ , shifting the turbines to their yaw-induced equilibrium positions, and it acts directly inside  $P$ , where misalignment deflects the wakes downstream and reduces each yawed rotor’s own power capture. The objective remains a function of the layout  $(\mathbf{x}, \mathbf{y})$  alone. The yaw vector  $\gamma^*$  is not an independent decision variable but a dependent quantity predicted from the layout and the wind condition, which is what distinguishes this single-loop formulation from nested co-design.

The ordering of Steps 2 and 3 reflects the physical coupling between yaw and platform displacement. Under nacelle yaw misalignment, the cross-wind component of rotor thrust shifts the platform laterally, so the equilibrium displaced position depends on the yaw angle applied. Displacement must therefore be conditioned on the predicted yaw state, which is why the yaw surrogate is queried before the displacement surrogate in the pipeline. Predicting displacement first and yaw second would be physically inconsistent, since the displaced position used to evaluate farm power would not correspond to the operating yaw condition that produced it.

### 3.9 Layout Comparison: the Layout Field Error

Each optimization method is run ten times and returns ten distinct layouts, as optimal solutions, so comparing two methods means comparing two sets of ten layouts rather than two single designs. A meaningful comparison therefore has to condense each method’s runs into one representative description and show where the methods place turbines differently, rather than contrasting arbitrarily chosen individual runs.

Two layouts can reach almost identical AEP while placing individual turbines at markedly different positions, so a comparison based on energy alone does not reveal whether the optimized layouts are geometrically similar or different. To compare layouts on the basis of turbine positions rather than performance, this thesis uses the layout field error (LFE), a model-agnostic metric introduced by Baricchio et al. [56].

The metric represents a layout as a continuous field rather than a set of discrete points. Each turbine is replaced by a two-dimensional Gaussian distribution centered at its position, and the layout field (LF) is the sum of these contributions over all turbines. The layout field at a coordinate  $(x, y)$  is given by

$$\text{LF}(x, y) = \sum_{i=1}^{N_{\text{wt}}} \frac{1}{2\pi\sigma^2} \exp\left(-\frac{(x - x_{\text{wt},i})^2 + (y - y_{\text{wt},i})^2}{2\sigma^2}\right) \quad (3.5)$$

Here,  $N_{\text{wt}}$  is the number of turbines,  $(x_{\text{wt},i}, y_{\text{wt},i})$  are the coordinates of turbine  $i$ , and  $\sigma$  is a smoothness parameter expressed as a multiple of the rotor diameter  $D$ . A small  $\sigma$  concentrates the field at the turbine positions, while a larger  $\sigma$  spreads the contribution of each turbine over a wider area. As a superposition of normal distributions, the layout field integrates over the domain to the number of turbines  $N_{\text{wt}}$ . The layout field of a set of layouts can then be defined as the average of the individual fields, with the integral remaining equal to  $N_{\text{wt}}$  [56]. In this way, the ten runs of each method are condensed into a single field per method.

The layout field error (LFE) is the difference between two layout fields. For each mooring, the ten layouts per method are first averaged into one field, and the LFE is then computed between method pairs. A positive value marks regions where the first method places more turbines than the second, a negative value marks the opposite, and a value near zero means that both methods

occupy the region to the same extent. The comparison relies on turbine positions alone, so it is unaffected by the performance models and their uncertainties.

## 3.10 Case Study: Kriti 3

### 3.10.1 Site Description

All layout optimization experiments in this thesis use the Kriti 3 floating wind farm site, located off the northern coast of Crete in southern Greece. Kriti 3 is one of the Greek scenarios defined within the SUDOCO project [57], a Horizon Europe collaboration. The site boundary data and wind resource measurements used in this thesis were provided by partners at the National Technical University of Athens (NTUA) as part of the SUDOCO network.

The site boundary is a polygon, with a bounding box of approximately 8.83 km (east-west) by 6.02 km (north-south) and an enclosed area of approximately 33.95 km<sup>2</sup>. The site sits at approximately 550 m water depth, well beyond the roughly 60 m practical limit for fixed-bottom foundations, and therefore requires floating support structures.

Kriti 3 is one of the candidate sites (POAYAP) designated under the Greek National Offshore Wind Farm Development Programme, most of which target floating rather than fixed-bottom technology given the prevailing water depths (Figure 3.11) [58]. Each candidate area is accompanied by detailed environmental assessments, including marine-mammal distribution maps for the Crete sites (Figure 3.12) [59].

### 3.10.2 Wind Resource Characterisation

Wind resource data for the Kriti 3 site were obtained from the Poseidon oceanographic monitoring system [60], operated by the Hellenic Centre for Marine Research (HCMR). Station 61277, located at 35.727°N, 25.135°E (approximately 24 nautical miles north of Crete and 50 nautical miles northwest of Kriti 3), provides wind speed and direction measurements at 3 m above sea surface. The location of the buoy relative to the site is shown in Figure 3.13. The dataset spans two full years, from November 27, 2022 to November 26, 2024. Wave data from the same station are not used in this thesis.

Wind speeds measured at 3 m were scaled to the turbine hub height of 150 m using a logarithmic wind profile. Hub-height wind speed is given by:

$$U(z) = U_{\text{ref}} \cdot \frac{\ln(z/z_0)}{\ln(z_{\text{ref}}/z_0)} \quad (3.6)$$

where  $z = 150$  m is the target height,  $z_{\text{ref}} = 3$  m is the measurement height, and  $z_0$  is the surface roughness length. A value of  $z_0 = 0.001$  m is adopted, standard for open-sea offshore conditions. No atmospheric stability corrections are applied, the profile is evaluated under the neutral-stability assumption, consistent with common practice in offshore wind resource assessment.

The scaled measurements are aggregated into a wind rose with 16 directional sectors of 22.5° each. Per-sector Weibull scale ( $A$ ) and shape ( $k$ ) parameters are estimated by maximum likelihood estimation on the hub-height wind speed samples within each sector. The frequency-weighted mean wind speed at hub height is 7.76 m/s, computed as  $\sum_i p_i A_i \Gamma(1 + 1/k_i)$  where  $p_i$  is the sector frequency of sector  $i$ . The dominant wind direction is 270° (west), with a sector frequency of 18.0%, and the five westerly sectors spanning 247.5°–337.5° together account for approximately

### 3.10. Case Study: Kriti 3




Figure 3.11: Candidate offshore wind development areas (POAYAP) across Greece by foundation technology (cyan: fixed-bottom; blue: floating). The Kriti 3 case study lies off northeastern Crete [58].

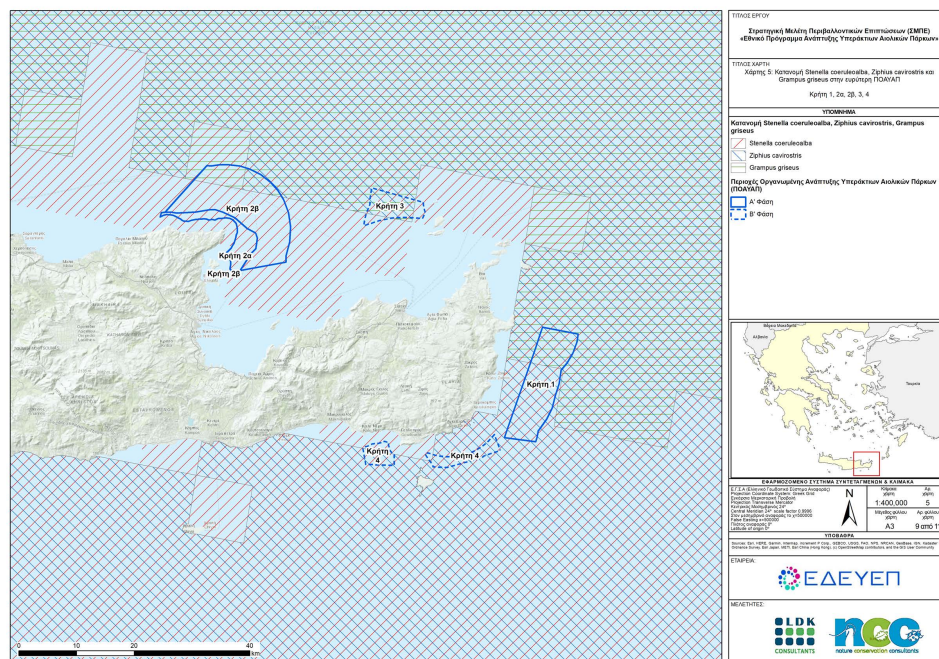


Figure 3.12: Example of the per-area environmental documentation accompanying the programme: marine-mammal distribution across the Crete candidate areas, including Kriti 3 [59].

66.7% of all occurrences. The resulting 16-sector Weibull wind rose, shown in Figure 3.14, serves as the AEP evaluation basis throughout this thesis.



Figure 3.13: The Kreti 3 site (red polygon) north-east of Crete and the Poseidon E1M3A measurement buoy (WMO 61277, yellow triangle), located approximately 97 km (52 nm) to the north-west. Satellite basemap: Esri World Imagery.

## 3.11 Computational Setup

All optimization and dataset-generation campaigns in this thesis are executed on DelftBlue [55], TU Delft’s high-performance computing cluster. Each job is allocated a single node with 48 CPU cores and submitted as a SLURM array, with every run saved to a uniquely tagged output directory. Fitness evaluations are parallelised at the population level: each of the 48 worker processes independently evaluates the AEP objective for its assigned candidate layouts, with no inter-worker communication within a generation. The platform orientation is fixed at  $270^\circ$  throughout, consistent with the AEP wrapper configuration (§3.6.2) and the yaw surrogate training conditions (§3.7.1).

The genetic algorithm shares one configuration across all layout campaigns: a population of 120 individuals, 400 generations, and 12 elites retained per generation. This configuration was locked after preliminary runs with smaller populations and generation counts. The population size follows the guideline of roughly ten individuals per turbine, the elite count was varied alongside it, and the generation budget was set well above the point where the exploratory runs stopped improving. The convergence histories reported in Chapter 4 confirm this margin, with the optimizer effectively converged by generation 200. The same configuration is used for the fixed-position baseline, the displacement-aware, and the co-design runs, and all three campaigns share the same wind condition grid, which matches the training grid of the integrated yaw surrogate (§3.7.1), so the surrogate is never queried at conditions outside its training distribution. Holding the optimizer configuration and the evaluation grid constant means any difference in

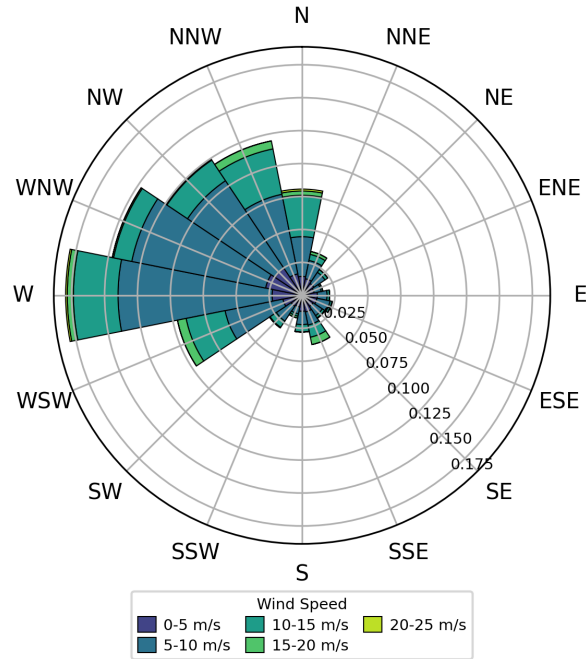


Figure 3.14: Kriti 3 wind rose at 150 m hub height, over 16 directional sectors and five speed bins, generated from Poseidon station 61277 measurements [60]. The resource is strongly westerly.

optimized layouts or AEP between conditions reflects the change in objective function alone, not a difference in search budget or parallelisation.

Three campaigns make up the computational workload. The training database for the integrated yaw surrogate (§3.7.1) completes in approximately 8 hours of wall-clock time ( $\approx 384$  core-hours). A displacement-aware layout optimization run completes in approximately 12.6 hours ( $\approx 605$  core-hours), against roughly 695 s ( $\approx 9$  core-hours) for an equivalent fixed-position run (§3.6.2), and a co-design run in approximately 20 hours ( $\approx 960$  core-hours). Because the layout optimization problem is non-convex and multimodal, with many local optima of similar performance [61, 62], a single optimization reaches only one of these optima. Ten independent runs are therefore launched for each method and mooring configuration, giving a distribution of optimized layouts and AEP values per method. These distributions are reported and compared in Chapter 4. At these per-run costs, the total compute budget makes cluster deployment a practical necessity, since a standard workstation could not sustain the workload.

# 4

## Results

This chapter presents the results, building from isolated components to the complete co-design framework. Section 4.1 presents the watch circles for each mooring design, and Section 4.2 quantifies the gains of displacement-aware yaw optimization on the reference layout. Section 4.3 validates the displacement surrogate, Section 4.4 applies it within displacement-aware layout optimization, Section 4.5 validates the integrated yaw surrogate, and Section 4.6 combines all components in the single-loop co-design comparison. Interpretation of the results is deferred to Chapter 5.

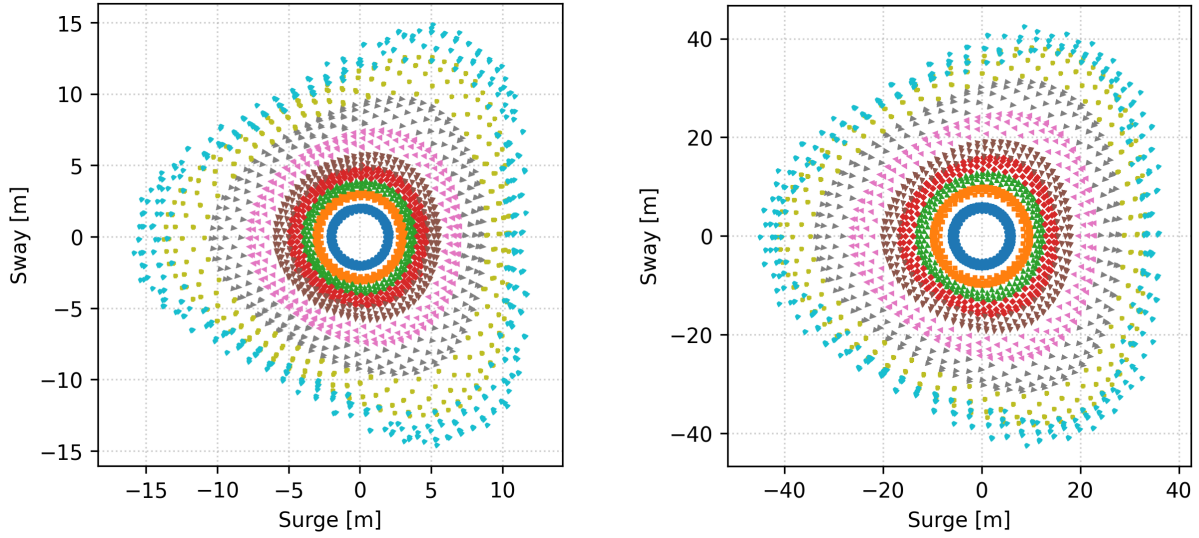
### 4.1 Platform Displacement Patterns - Watch Circles

The three panels in Figure 4.1 show the watch circles for each mooring configuration. Each watch circle shows the surge and sway offsets, collected from the simulations, across all combinations of wind speed, wind direction, and yaw misalignment in the extended dataset, tracing the region of positions a platform occupies around its anchored reference. The watch circles convey how far and in which directions platforms of each type drift for wind speeds from cut-in to rated for all wind directions.

The taut mooring produces a compact, near-circular response. Figure 4.1a shows nine nearly concentric rings, one per simulated wind speed ( $U = 3$  to  $11$  m/s), with axes spanning approximately  $\pm 15$  m in both surge and sway. The rings are cleanly separable and directionally symmetric, reflecting the 3-fold rotational symmetry of the three-line mooring geometry. At  $U = 11$  m/s the outer ring reaches approximately  $15.9$  m ( $0.07D$ , where  $D = 241.94$  m is the rotor diameter).


The semitaught mooring shares the same qualitative structure. Figure 4.1b shows concentric, directionally symmetric rings extending to approximately  $\pm 45$  m in both surge and sway. The maximum displacement at  $U = 11$  m/s is  $45.0$  m ( $0.19D$ ), approximately three times that of the taut design.

The semitaught.v2 mooring departs from this pattern entirely. Figure 4.1c shows two spatially separated directional clusters, one lobe in the upper-right quadrant (surge approximately  $50$  to  $150$  m, sway approximately  $50$  to  $175$  m) and one elongated lobe in the lower-left (surge approximately  $-100$  to  $-250$  m, sway approximately  $-50$  to  $-150$  m). The maximum displacement at  $U = 11$  m/s is  $205.7$  m ( $0.85D$ ). At this offset a platform shifts by nearly one full rotor diameter from its nominal position, giving meaningful displacement for active repositioning. The unique watch-circle structure of the semitaught.v2 design is a direct consequence of its high compliance, which allows the platform to drift into distinct regions of the (surge, sway) plane under different wind directions. This design is theoretical and how it would behave in practice is uncertain, but it serves as a useful test case for the layout optimizer's ability to exploit large displacements.



(a) Taut mooring.

(b) Semitaut mooring.



(c) Semitaut\_v2 mooring.




Figure 4.1: Watch circles for the three mooring configurations: platform surge and sway displacement across all training samples, coloured by wind speed. Displacement magnitude grows with mooring compliance, from the taut design to the highly compliant semitaut\_v2.

## 4.2 Displacement-Aware Yaw Optimization on the Reference Layout

### 4.2.1 Per-Mooring Yaw Optimization Gains

The displacement-aware yaw optimizer, `m2sr_floating`, produces markedly different power gains across the three mooring configurations. The comparison is presented at a single operating condition,  $U = 8$  m/s and  $WD = 270^\circ$ , the dominant wind direction of the site, so the reported values are farm power gains at that condition. Table 4.1 reports gains of 2.83% (44.71 to 45.98 MW), 2.44% (44.65 to 45.75 MW), and 22.16% (44.62 to 54.50 MW) for the taut, semिताut, and semिताut\_v2 moorings, respectively. The taut and semिताut figures are consistent with gains reported for fixed-bottom wake-steering campaigns, indicating that at these compliance levels the gain is principally through wake deflection.

The wake field at  $U = 8$  m/s,  $WD = 270^\circ$  is shown in Figure 4.2 for all three moorings, with open circles marking nominal positions and filled circles marking displaced positions. In the taut and semिताut panels the two sets of markers are nearly coincident, reflecting passive displacements of at most 15.9 m ( $0.07D$ ) and 45.0 m ( $0.19D$ ), respectively. At these offsets the spatial layout is effectively unchanged, and the gains of 2.83% and 2.44% arise from lateral wake deflection.

The semिताut\_v2 panel looks qualitatively different from the other two. Filled circles are visibly separated from open circles, with several mid-farm turbines shifted clear of upstream wake corridors, while optimized farm power reaches 54.50 MW, a gain of 22.16%. At this level of mooring compliance, the yaw optimizer exploits lateral thrust forces to physically reposition platforms out of upstream wakes. These single-condition gains do not translate directly to annual energy. Averaged over the full wind rose, where above-rated and lightly waked conditions carry most of the probability mass, the gain is substantially lower (§4.5.4).

Lastly, the optimal yaw angles of the taut and semिताut cases are similar, while the ones for the semिताut\_v2 case are markedly different, with opposite angles at the middle columns to drive the offset of these out of the upstream wakes. This difference in optimal yaw angles is a direct consequence of the different displacement patterns of the three mooring designs, and it underscores the importance of embedding platform displacement inside the optimization loop.

Table 4.1: Per-mooring yaw optimization results on the reference 12-turbine Kriti 3 layout at  $U = 8$  m/s,  $WD = 270$ . Baseline farm power corresponds to zero yaw misalignment; optimized farm power is obtained with `m2sr_floating`.

Mooring	Baseline [MW]	Optimized [MW]	Gain [%]
Taut	44.71	45.98	2.83
Semिताut	44.65	45.75	2.44
Semिताut_v2	44.62	54.50	22.16




Figure 4.2: Optimized yaw wake maps for the three mooring configurations at  $U = 8$  m/s,  $WD = 270$ . Open circles indicate the nominal turbine positions; filled circles indicate the displaced positions after passive repositioning.

### 4.2.2 Active Repositioning in the Semitaut\_v2 Case

The three-way comparison in Figure 4.3 isolates the semitaut\_v2 mooring case against the optimal yaw angles of the ordinary `m2sr`, evaluated at the same single operating condition of  $U = 8$  m/s and  $WD = 270^\circ$  as Section 4.2.1. The left panel shows fixed-platform yaw optimization, where the optimizer treats the turbines as stationary. The rotors are visibly yawed and the wakes are deflected, with effective wind speeds of 8.0 m/s at the front column and 6.3 m/s at the middle and downstream columns. Farm power reaches 46.08 MW, a gain of 3.0% over the zero-yaw baseline (Table 4.2). This gain comes entirely from wake deflection.

The centre panel applies the same fixed optimal yaw angles to the floating case. The effective wind speeds at the middle and downstream columns fall to 6.2 m/s and 6.1 m/s, slightly below the fixed-platform case, and farm power drops to 44.47 MW, a net loss of 0.3% relative to the no-yaw baseline (Table 4.2). The yaw solution derived for stationary turbines does not transfer to a compliant platform. A yaw strategy computed under a fixed-position assumption can turn a predicted gain into a net loss on a compliant platform, if displacement is not accounted for, which is the strongest argument for accounting for platform displacement inside the control optimization.

The right panel shows displacement-aware floating yaw optimization with `m2sr_floating`, which accounts for the displaced positions. The middle-column rotors recover to an effective wind speed of 7.1 m/s, against 6.2–6.3 m/s in the other two panels, and their positions are shifted clear of the upstream wake corridors. Farm power reaches 54.50 MW, a gain of 22.2% (Table 4.2). At this level of mooring compliance, the lateral thrust component of the yawed turbines physically repositions the platforms out of the wakes shed upstream, and the effective-wind-speed recovery at the middle column is the benefit of that repositioning.

Table 4.2: Farm power for the three yaw-control cases of the semitaut\_v2 wake-map comparison (Figure 4.3) at  $U = 8$  m/s,  $WD = 270$ . Gains are relative to the zero-yaw baseline.

Yaw-control case	Power [MW]	Gain [%]
Fixed-platform optimization	46.08	+3.0
Fixed yaw applied to floating	44.47	-0.3
Displacement-aware floating	54.50	+22.2

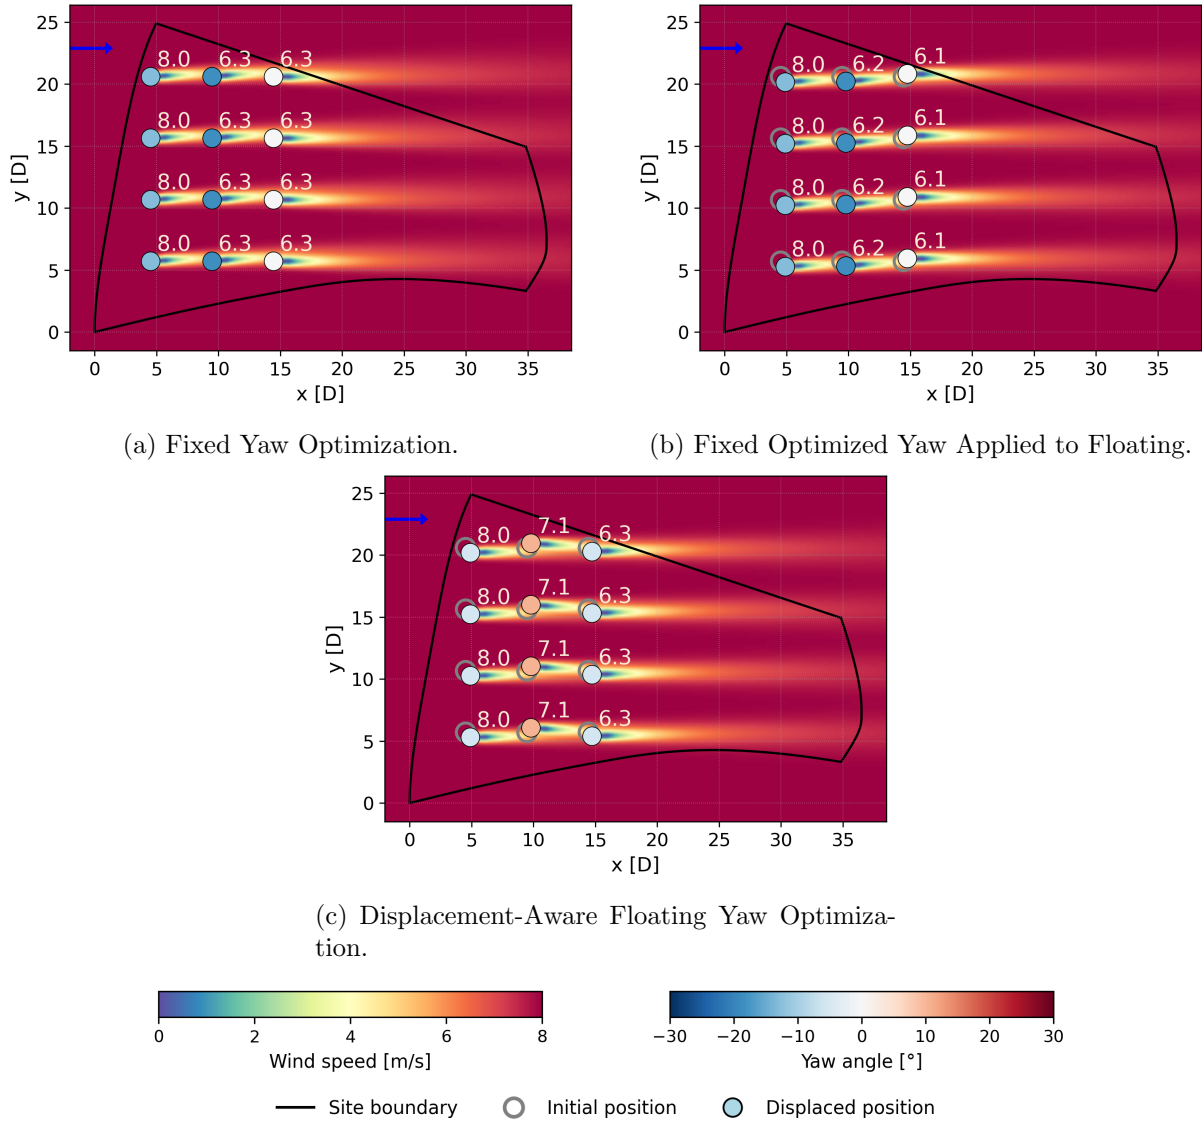


Figure 4.3: Three-way wake map comparison for the semिताut\_v2 mooring at  $U = 8$  m/s,  $WD = 270$ . Numbers indicate the effective wind speed at each rotor. Only the displacement-aware floating optimization (right) recovers the steering gain on this compliant mooring.

### 4.3 Displacement Surrogate Validation

The displacement surrogate must predict the platform offsets inside the layout optimizer, across all three mooring designs and both displacement components, mimicking the behavior of the full simulation. Table 4.3 reports the per-mooring accuracy on the 80/20 hold-out split alongside the five-fold cross-validation results. The coefficient of determination,  $R^2$ , is 0.999 for taut and semिताut in both surge and sway, and 0.996 (surge) and 0.997 (sway) for semिताut.v2, with the fold standard deviation in  $R^2$  below 0.001 throughout. Mean absolute error (MAE) grows with mooring compliance, from roughly 0.10 m for taut to 0.25–0.27 m for semिताut and 2.8–3.2 m for semिताut.v2. This accuracy is in line with the displacement surrogate developed in prior work [12].

The diagnostic plots confirm that this accuracy holds uniformly across the operating space. Figure 4.4 shows the semिताut.v2 case, taken as representative because its compliance produces the largest displacements. The predicted-versus-actual points lie on the dashed 1:1 reference for both surge and sway over the full  $\pm 200$  m range, and the residuals scatter symmetrically about zero with no systematic trend against wind direction  $\theta$  or yaw misalignment  $\gamma_\Delta$ . The most compliant design is thus captured without directional or yaw-dependent bias, which supports using the surrogate in place of SLOW inside the layout optimizer (§3.3.3). The taut and semिताut moorings show the same behaviour and their diagnostics are reported in Appendix B.

Table 4.3: Accuracy of the displacement surrogate per mooring and output.  $R^2$  and MAE are evaluated on the 80/20 held-out split; the CV MAE column is 5-fold cross-validation (mean  $\pm$  std).

Mooring	Output	$R^2$	MAE [m]	CV MAE [m]
Taut	Surge	0.999	0.10	$0.10 \pm 0.01$
	Sway	0.999	0.09	$0.09 \pm 0.00$
Semिताut	Surge	0.999	0.27	$0.27 \pm 0.02$
	Sway	0.999	0.25	$0.24 \pm 0.01$
Semिताut.v2	Surge	0.996	3.19	$3.04 \pm 0.14$
	Sway	0.997	2.81	$2.65 \pm 0.12$

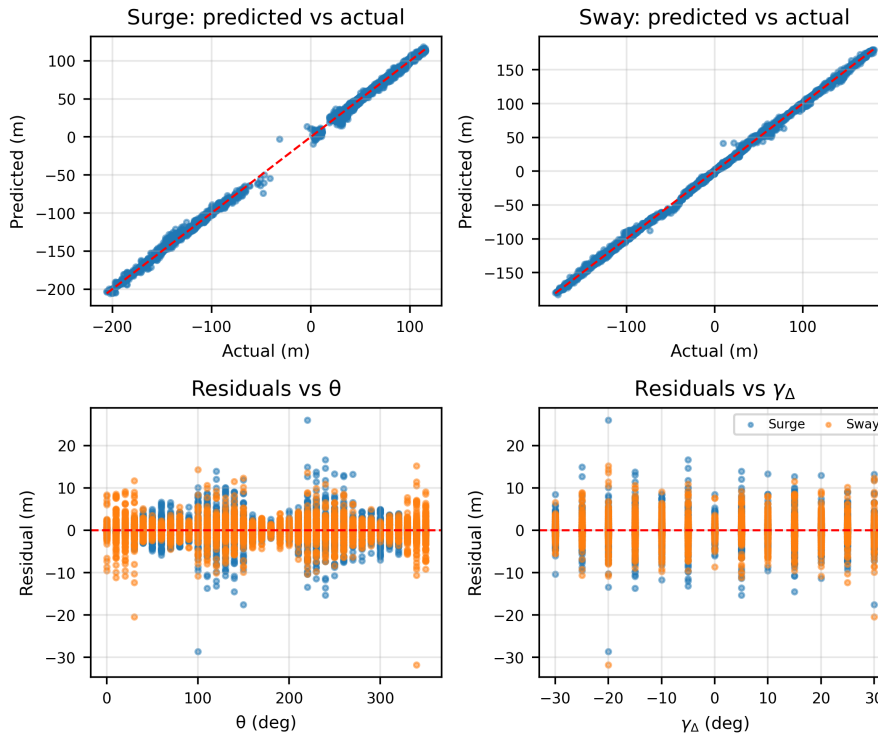


Figure 4.4: Displacement surrogate diagnostics for the semitaut\_v2 mooring, the representative (most compliant) case. Top: surge (left) and sway (right) predicted vs. actual, with the dashed 1:1 reference. Bottom: residuals against wind direction  $\theta$  (left) and yaw misalignment  $\gamma_{\Delta}$  (right).

## 4.4 Displacement-Aware Layout Optimization

### 4.4.1 Layout Optimization Results

The displacement-aware optimizer was run ten times independently for each mooring design, and this subsection inspects one result in detail before the across-run statistics are reported in §4.4.2. The case shown is run 7, the best-performing of the ten semitaut runs, a 12-turbine layout optimized at the fixed platform orientation of  $270^\circ$  using the LO-GA mechanics inherited from Baricchio et al. (2024) [21] with the displacement-aware AEP wrapper (§3.6.2).

Figure 4.5a shows the twelve turbines distributed over the boundary polygon, with several positioned on or close to the site edges, a small group along the northern edge, and the remainder spaced through the interior. This pattern is the expected response to the objective and its constraints. Pushing turbines toward the boundary maximises inter-turbine separation and reduces wake overlap, while the minimum-spacing penalty and the polygon-containment constraint (§3.6.2) keep nominal positions feasible. The layout reaches an AEP of 683.8 GWh, evaluated with the displacement-aware wrapper at zero yaw.

The convergence history confirms that this layout reflects a settled optimum rather than a non converged solution. Figure 4.5b plots the best-AEP trajectory over 400 generations at a population size of 120, rising monotonically toward the final value. Most of the improvement is captured within the first 40 generations or so, after which the gains diminish steadily and the optimizer is effectively converged by generation 200. The early saturation indicates that the 400-generation budget set in §3.11 is more than sufficient for a problem of this size, which supports treating each of the ten runs as fully converged in the statistical comparison that follows.

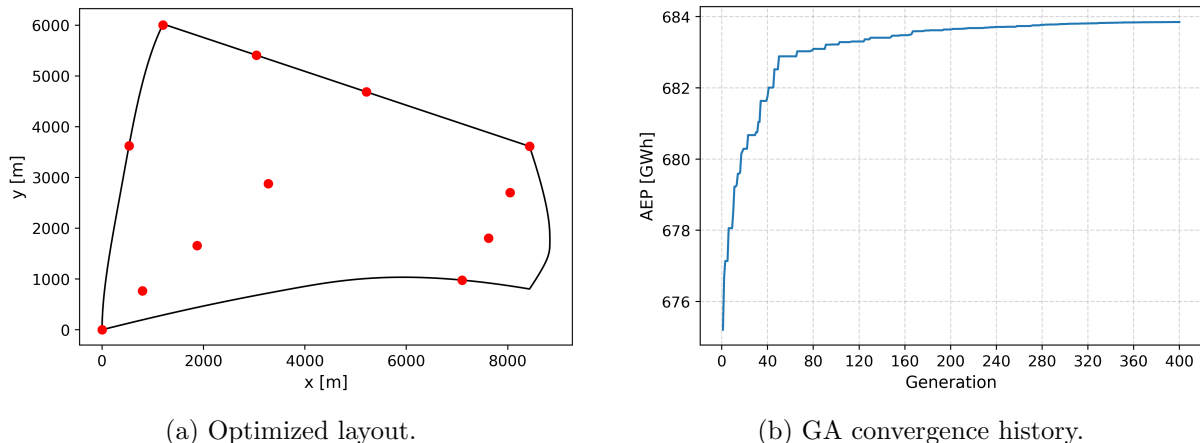


Figure 4.5: Best-performing displacement-aware optimization run for the semitaut mooring (run 7 of 10, AEP = 683.8 GWh). (a) The optimized layout over the Kriti 3 site boundary. (b) GA convergence history over 400 generations at population size 120.

#### 4.4.2 Statistical AEP Comparison: Fixed vs. Displacement-Aware

Following the single run inspected in §4.4.1, the question this subsection addresses is whether accounting for passive repositioning changes the optimized AEP once the stochastic variability across the ten runs is taken into account (see §3.11 for why multiple runs are used). Across all three moorings, displacement-aware optimization shifts the AEP distribution upward, but the median gain stays below 0.1% and is small relative to the run-to-run spread. Figure 4.6 compares the ten-run AEP distributions for fixed-bottom and displacement-aware optimization, one panel per mooring, and Table 4.4 reports the per-condition statistics and median gains. Each of the results shown correspond to the optimized layouts evaluated at zero yaw with the displacement-aware wrapper, so the AEP values reflect the passive repositioning of the turbines under the optimized layouts, and the gain is the improvement in that AEP when the layout optimization accounts for that repositioning.

For the taut mooring the two distributions overlap substantially. Figure 4.6a shows the displacement-aware and fixed-bottom interquartile ranges covering nearly the same band, with a median gain of +0.3 GWh (+0.04%). This is the expected outcome for the stiffest of the three designs. Taut platforms drift little under thrust, so the nominal and displaced positions barely diverge and the space for improvement by the displacement-aware objective is limited.

The semitaut mooring produces a comparable median gain of +0.3 GWh (+0.05%) but a different distribution shape. Figure 4.6b shows the displacement-aware condition with a tighter upper quartile than the fixed-bottom condition, with the higher-AEP runs clustered near the top of the box, driving the median up.

The semitaut\_v2 mooring shows both the largest median gain and the widest spread. Figure 4.6c shows a median gain of +0.4 GWh (+0.06%), the largest of the three, alongside a noticeably wider displacement-aware interquartile range and a long upper whisker. This is the most compliant of the three designs, and its layout outcomes are correspondingly the most sensitive to passive drift. The larger platform offsets give the displacement-aware objective more to act on, which raises the median but also widens the run-to-run scatter. The best individual solution also comes from the displacement-aware condition, at 685.4 GWh against 684.0 GWh for the best fixed-bottom run, a margin of +1.4 GWh (+0.21%, Appendix C). Semitaut\_v2 is the only mooring for which this holds; on taut and semitaut the best individual run comes from the fixed-bottom condition.

Where displacement is significant, the displacement-aware approach can therefore reach solutions the fixed-position optimization does not. In every panel the median gain is of similar size to the run-to-run standard deviation, so the displacement-aware improvement is only marginally above the optimizer scatter.

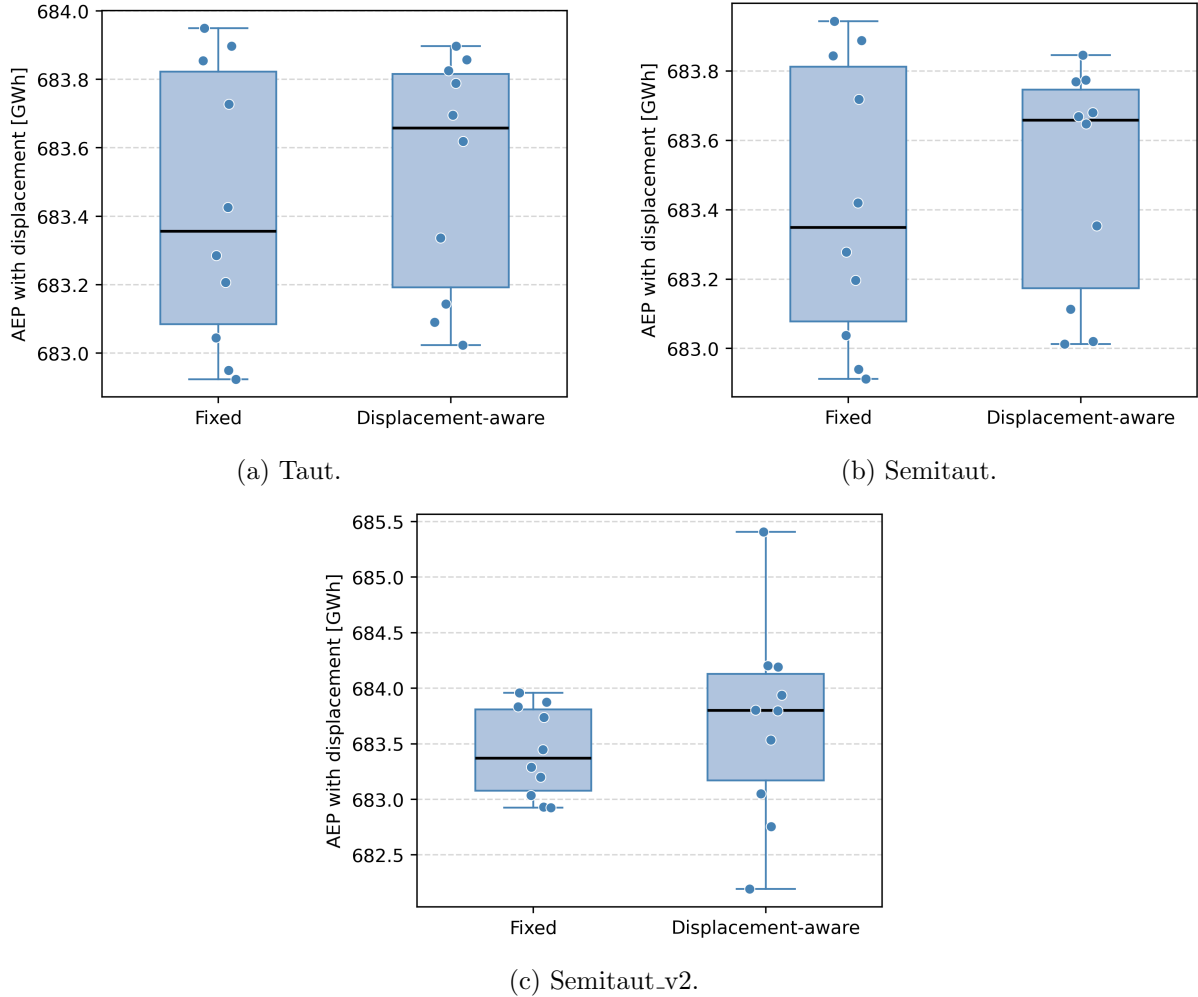


Figure 4.6: AEP distributions across 10 GA runs per mooring: fixed-bottom vs. displacement-aware layout optimization. Each panel retains its own AEP scale. Median gains are listed in Table 4.4.

Table 4.4: Fixed-bottom versus displacement-aware layout optimization per mooring ( $N = 10$  runs per condition, evaluated at zero yaw with the displacement-aware wrapper). Median AEP, run-to-run standard deviation, and the median gain of displacement-aware over fixed-bottom. Full statistics (mean, minimum, maximum) are reported in Appendix C.

Mooring	Method	Median [GWh]	Std [GWh]	Median gain [GWh]	[%]
Taut	Fixed-bottom	683.4	0.38	–	–
	Disp.-aware	683.7	0.33	+0.30	+0.04
Semitaut	Fixed-bottom	683.3	0.38	–	–
	Disp.-aware	683.7	0.31	+0.31	+0.05
Semitaut_v2	Fixed-bottom	683.4	0.38	–	–
	Disp.-aware	683.8	0.84	+0.43	+0.06

## 4.5 Integrated Yaw Surrogate Validation

This section establishes how faithfully the integrated yaw surrogate reproduces the `m2sr_floating` ground truth it was trained against. The validation proceeds in three steps. Training-set accuracy is examined first (§4.5.1), the surrogate’s behaviour is then isolated in a controlled two-turbine study (§4.5.2), and farm-scale annual energy production is assessed last (§4.5.4). Because the surrogate is embedded directly in the genetic algorithm objective during single-loop co-design, the accuracy reported across these steps is the precondition for assessing the co-design results of §4.6.

### 4.5.1 Yaw Surrogate Diagnostics

The first validation step compares predicted against actual optimal yaw on the held-out layouts. Figure 4.7 shows this comparison for each mooring, split into all wind speeds (left) and the below-rated region only (right,  $3 \leq ws \leq 10$  m/s). The all-speeds coefficient of determination is depressed by the many zero-yaw labels above rated, falling to 0.70 for taut, 0.69 for semitaut, and 0.60 for semitaut\_v2 (Table 4.5). The below-rated subset is the more representative measure, since this is the regime where wake steering produces non-trivial angles, and there the agreement is markedly stronger at 0.88, 0.87, and 0.77 respectively. Accuracy tracks mooring compliance. Taut and semitaut are close, and the more compliant semitaut\_v2 is the hardest case, consistent with its larger yaw-induced displacements requiring the optimizer to exploit dependencies a stiffer design does not.


The associated error magnitudes confirm that the surrogate is accurate where it matters. In the below-rated region the mean absolute error stays sub-degree for taut and semitaut at  $0.59^\circ$  and  $0.62^\circ$ , rising to  $0.90^\circ$  for semitaut\_v2 (Table 4.5). The cross-validated values, computed over five layout-grouped folds, sit closely to the held-out figures with small standard deviations, so accuracy does not hinge on which layouts happen to fall in the test set. Because the hold-out split withholds whole layouts rather than individual conditions (§3.7.3), these numbers reflect generalization to unseen farm geometries rather than interpolation within configurations the model has already seen. These values are compared against bottom-fixed yaw surrogates in §5.1.

The accuracy of the predicted yaw angle ultimately matters only through its effect on farm power. Figure 4.8 reports, by wind speed, the held-out yaw MAE as bars and the resulting mean farm

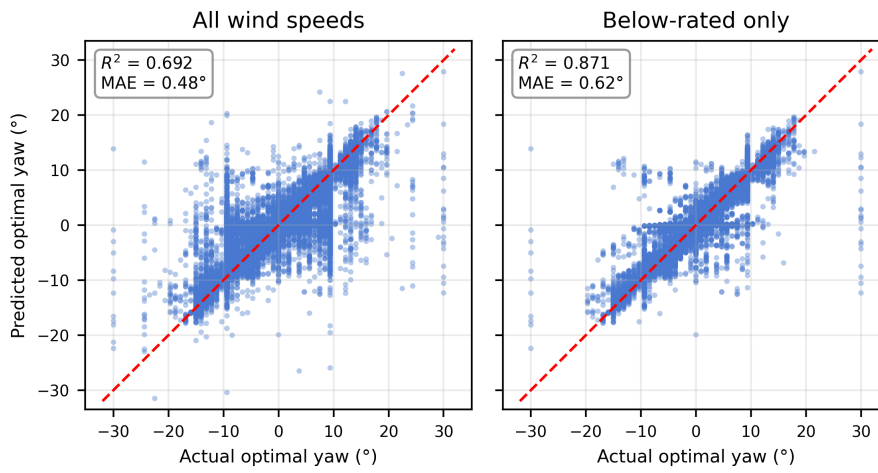
power gain over the yaw = 0 baseline as two lines, one using the `m2sr_floating` optimal yaw and one using the surrogate prediction. The gain is concentrated below the rated wind speed of 10.59 m/s, peaks at low speeds where the relative power gain is largest, and collapses to near zero above 12–13 m/s, where the turbines are power-limited and nacelle misalignment achieves no benefit. Below rated, the surrogate-driven gain follows the `m2sr_floating` curve closely, with a slight underestimate that is smallest for taut and grows with mooring compliance.

Table 4.5: Held-out accuracy of the integrated yaw surrogate per mooring: coefficient of determination  $R^2$  and mean absolute error (MAE) over all wind speeds and over the below-rated region ( $3 \leq ws \leq 10$  m/s), and the cross-validated values (mean  $\pm$  std over 5 layout-grouped folds) as a stability check.

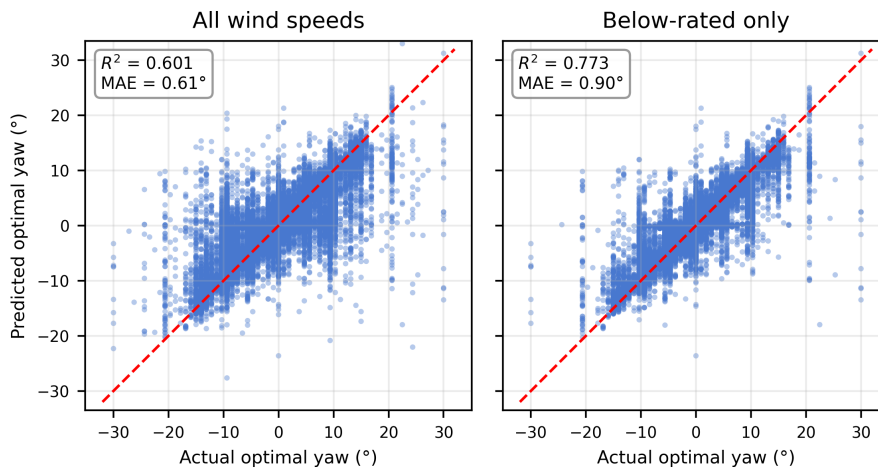
<b>Mooring</b>	<b>All wind speeds</b>		<b>Below-rated</b>		<b>Cross-validation</b>	
	$R^2$	MAE [°]	$R^2$	MAE [°]	$R^2$	MAE [°]
Taut	0.70	0.47	0.88	0.59	$0.70 \pm 0.01$	$0.48 \pm 0.01$
Semitaut	0.69	0.48	0.87	0.62	$0.70 \pm 0.01$	$0.49 \pm 0.01$
Semitaut_v2	0.60	0.61	0.77	0.90	$0.60 \pm 0.01$	$0.63 \pm 0.01$



(a) Taut.

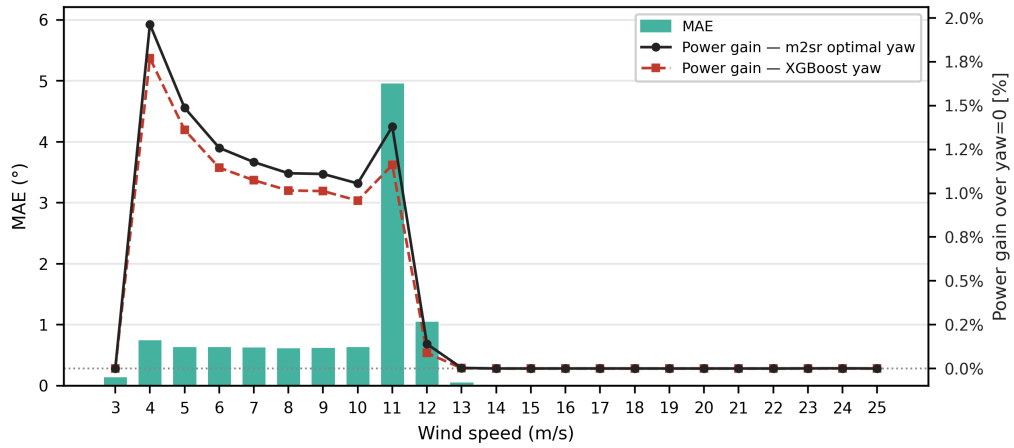


(b) Semitaut.




(c) Semitaut\_v2.


Figure 4.7: Held-out predicted vs. actual optimal yaw for the integrated yaw surrogate, one row per mooring: all wind speeds (left) and the below-rated region only ( $3 \leq ws \leq 10$  m/s, right). Dashed line is 1:1; points at  $\pm 30^\circ$  are optima saturated at the steering bound.



(a) Taut.



(b) Semitaut.



(c) Semitaut\_v2.


Figure 4.8: Surrogate yaw error and the resulting farm power gain by wind speed, per mooring. Bars (left axis) show the held-out MAE of the integrated yaw surrogate; the two lines (right axis) show the mean farm power gain over the yaw = 0 baseline using the m2sr optimal yaw (ground truth) and the XGBoost-predicted yaw.

### 4.5.2 Two-Turbine Study

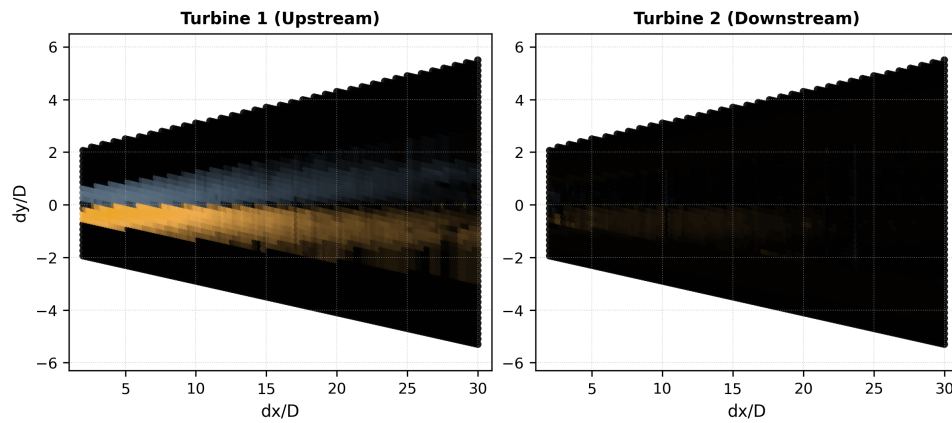
The two-turbine study isolates the integrated yaw surrogate against its `m2sr_floating` ground truth in a simple geometry. An upstream turbine is held fixed and a downstream turbine is swept across streamwise positions up to  $30D$  and crosswind offsets of several rotor diameters, at  $WD = 270^\circ$ , platform orientation  $270^\circ$ , and  $U = 10$  m/s. For each mooring and each swept position, the optimal yaw of both turbines is recorded, and the resulting fields are plotted for the `m2sr_floating` ground truth, for the surrogate prediction, and for their difference `m2sr - XGB`. Where §4.5.1 reported aggregate held-out accuracy, these maps show where in the wake the surrogate disagrees with full optimization. The shared  $\pm 20^\circ$  colour scale on the difference panels makes the magnitudes directly comparable across the three moorings.

For the taut and semitaut moorings the surrogate tracks the optimizer tightly. In Figures 4.9 and 4.10 the predicted fields reproduce the ground-truth steering pattern, and the difference panels stay close to neutral over most of the swept region, within a few degrees on the shared scale. The deviations that do remain are not scattered randomly. They form two narrow bands that follow the edges of the wake cone used in the feature construction (Equation 3.3), the boundary at which the downstream neighbour enters or leaves the feature slots and the optimal yaw transitions between active steering and zero. Errors concentrating at this boundary point to the wake cone filter itself, and a different selection of its width parameters could improve the surrogate further. The agreement over the rest of the field is consistent with the diagnostics of §4.5.1, where the two stiffer moorings carried the highest below-rated accuracy.



The `semitaut_v2` mooring behaves differently. Its ground-truth yaw fields do not follow the pattern shared by the other two designs. The optimizer assigns substantial angles to the downstream turbine as well, and the optimal angle of each turbine depends on the relative position of the pair in a more complex way, because at this compliance the yaw-induced displacement feeds back into the geometry the optimizer is steering. These labels are harder to learn, and the difference field in Figure 4.11 is correspondingly stronger. Its dominant feature is a concentrated patch of under-prediction at the downstream wake centre, near  $dy/D \approx 0$  within the first few rotor diameters, reaching close to  $20^\circ$ , while the upstream panel carries weaker mixed patches near the centreline. The deviation stays confined to the wake-centre core.



(a) `m2sr_floating` (ground truth).



(b) XGBoost surrogate prediction.



(c) Difference (`m2sr` - XGB).



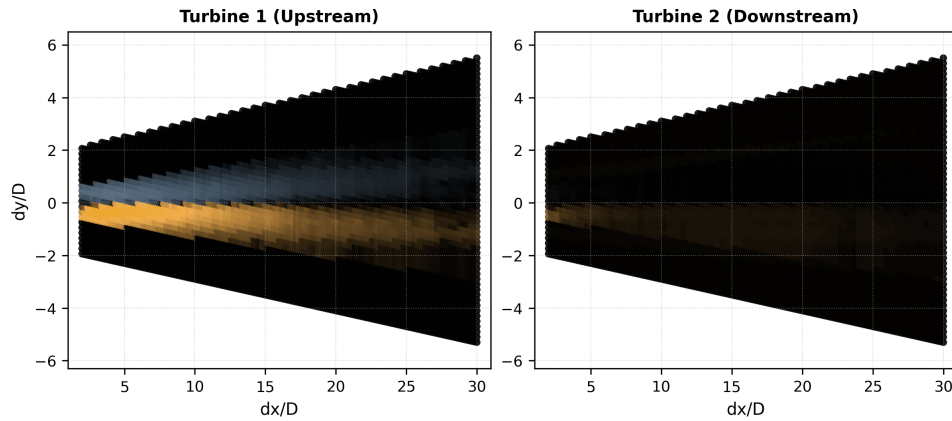




Figure 4.9: Two-turbine optimal-yaw comparison for the `taut` mooring at  $WD = 270$ ,  $U = 10$  m/s, upstream (left) and downstream (right) turbines. Panels (a) and (b) share the upper colour scale (yaw angle,  $\pm 30$ ); panel (c) uses the lower scale (yaw error,  $\pm 20$ ), shared across all three mooring figures.



(a) m2sr\_floating (ground truth).



(b) XGBoost surrogate prediction.



(c) Difference (m2sr - XGB).








Figure 4.10: Two-turbine optimal-yaw comparison for the **semitaut** mooring at  $WD = 270$ ,  $U = 10$  m/s, upstream (left) and downstream (right) turbines. Panels (a) and (b) share the upper colour scale (yaw angle,  $\pm 30$ ); panel (c) uses the lower scale (yaw error,  $\pm 20$ ), shared across all three mooring figures.



(a) m2sr\_floating (ground truth).



(b) XGBoost surrogate prediction.



(c) Difference (m2sr - XGB).




Figure 4.11: Two-turbine optimal-yaw comparison for the **semitaut\_v2** mooring at  $WD = 270$ ,  $U = 10$  m/s, upstream (left) and downstream (right) turbines. Panels (a) and (b) share the upper colour scale (yaw angle,  $\pm 30$ ); panel (c) uses the lower scale (yaw error,  $\pm 20$ ).

### 4.5.3 Investigating the Two-Turbine Ground-Truth Asymmetry

Of the three moorings, only `semitaut_v2` produces the one-sided two-turbine ground-truth yaw field flagged in §4.5.2 (Figure 4.11). The `taut` and `semitaut` cases give the antisymmetric dipole pattern expected when the problem is mirror-symmetric about the wind axis. The `semitaut_v2` watch circle differs from the others in two respects that could each be responsible. The first is its directional asymmetry, the second its non-circular, three-lobed shape. To separate them, two control watch circles are built from the existing data without new simulations, as idealized constructs rather than physical mooring designs. The first is a symmetrized version of the `semitaut_v2` watch circle, with the measured  $0\text{--}120^\circ$  sector made symmetric about its own axis and repeated at  $120^\circ$  intervals, which removes the asymmetry while keeping the lobed shape and the  $0.85D$  amplitude. The second is the `semitaut` watch circle scaled by a factor of 4.57 to the `semitaut_v2` amplitude, which keeps a circular shape while matching the  $0.85D$  magnitude (Figure 4.12). For each construction the displacement surrogate is retrained and the `m2sr_floating` two-turbine ground truth recomputed at  $WD = 270^\circ$  and  $U = 10$  m/s.

The first control isolates the asymmetry. With the asymmetry removed but the lobed shape and  $0.85D$  amplitude retained, the recomputed ground-truth yaw field stays one-sided and is qualitatively identical to the original `semitaut_v2` case (Figure 4.13a against Figure 4.11). The watch-circle asymmetry is therefore not the source of the one-sided field.

The second control matches the magnitude under a circular shape. The `semitaut` watch circle scaled to the same  $0.85D$  amplitude reproduces the antisymmetric dipole field of the native `semitaut` case rather than the one-sided pattern (Figure 4.13b against Figure 4.10). Raising the displacement magnitude under a circular shape therefore does not reproduce the one-sidedness, and the magnitude is ruled out.

Taken together, the two controls locate the cause in the shape of the watch circle. Neither the directional asymmetry nor the displacement magnitude reproduces the one-sided field on its own, whereas the non-circular, three-lobed geometry of the `semitaut_v2` watch circle does. This one-sided ground truth is unlike the symmetric fields of the other designs and of the fixed-bottom case, and the optimal yaw is likely harder to learn. This connects to the lower below-rated accuracy seen for `semitaut_v2` (§4.5.1).

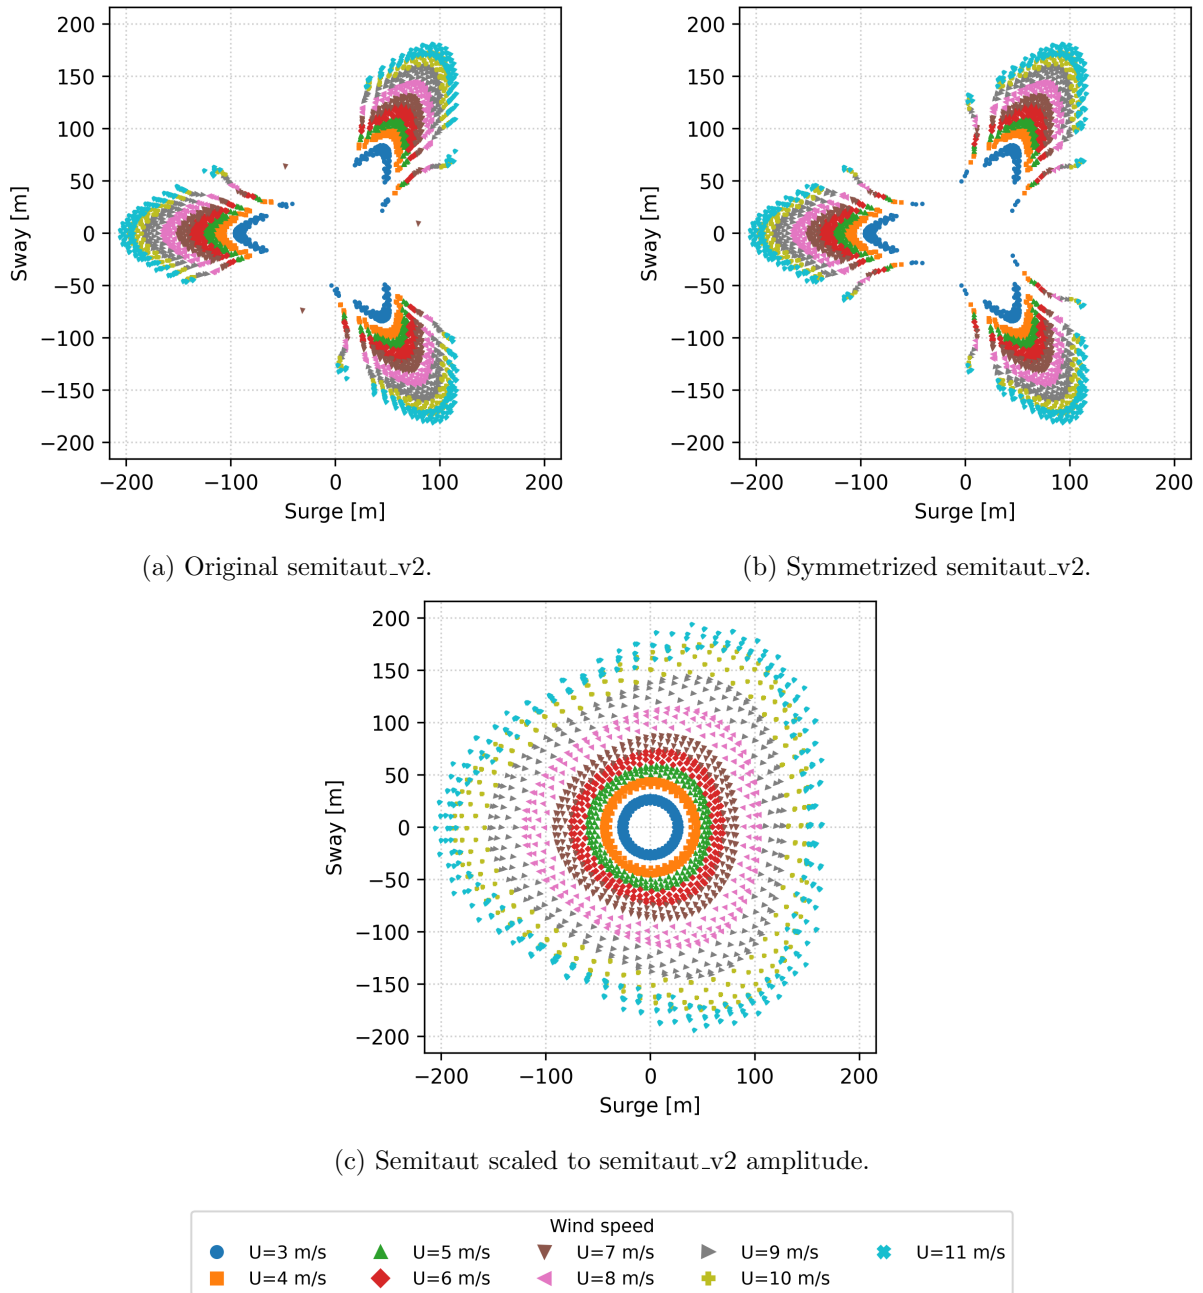


Figure 4.12: Control watch circles used to separate watch-circle shape from displacement magnitude, shown on a common  $\pm 216$  m scale and coloured by wind speed. Panel (a) is the original semitaut\_v2 watch circle, which is directionally asymmetric. Panel (b) is the symmetrized version, built by making the measured  $0-120^\circ$  sector symmetric about its own axis and repeating it at  $120^\circ$  intervals, which keeps the  $0.85D$  amplitude. Panel (c) is the semitaut watch circle scaled by a factor of 4.57 to the semitaut\_v2 amplitude, keeping a symmetric, near-circular shape. Both constructions are idealized and do not correspond to physical mooring designs.

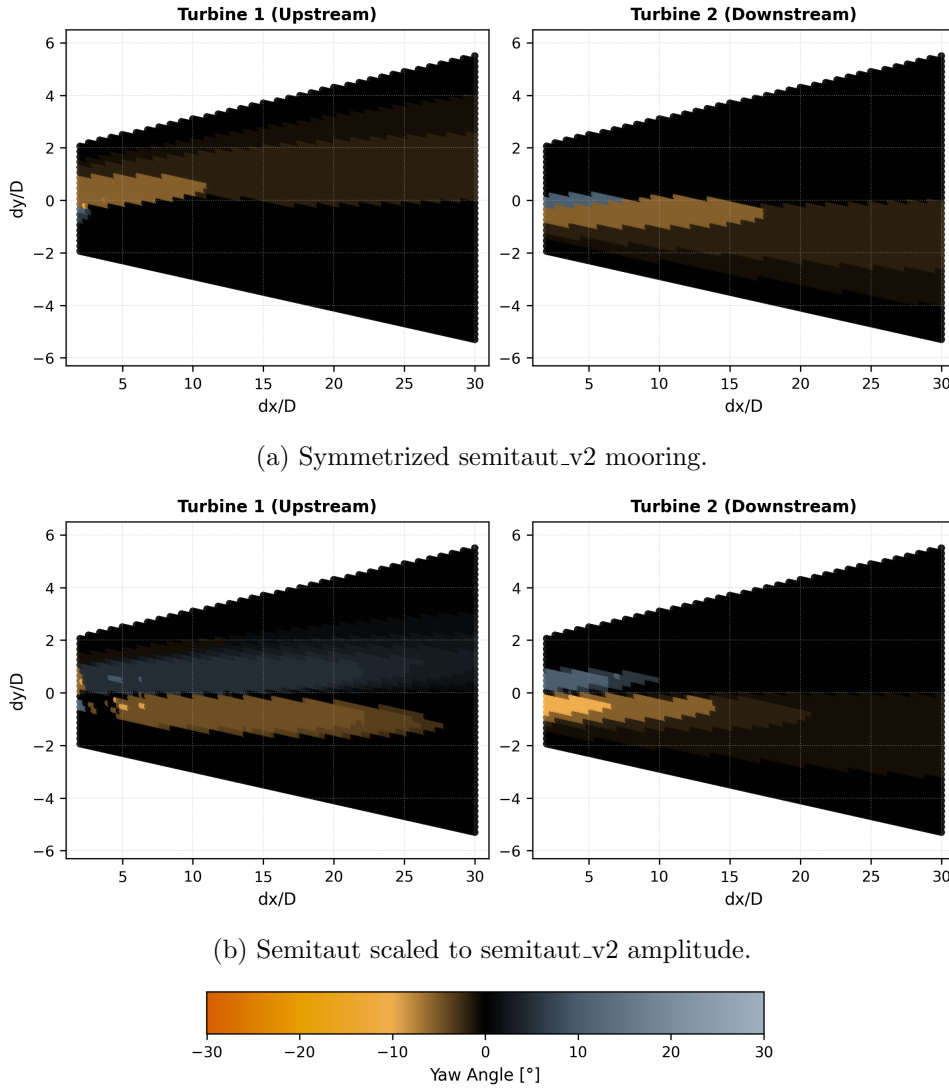


Figure 4.13: Two-turbine `m2sr_floating` ground-truth optimal yaw for the two control moorings at  $WD = 270^\circ$  and  $U = 10$  m/s, upstream (left) and downstream (right) turbines, on the same  $\pm 30^\circ$  scale as Figure 4.11. The symmetrized construction (a) stays one-sided, matching the original `semitaut_v2` ground truth, whereas the scaled-`semitaut` construction (b) keeps the antisymmetric dipole of the native `semitaut` case. The one-sided field persists once the asymmetry is removed and disappears once the shape is made circular at the same amplitude, identifying the watch-circle shape, not its asymmetry or the displacement magnitude, as the origin.

#### 4.5.4 Yaw Surrogate Farm-Scale AEP Validation

The farm-scale assessment is the last of the three validation steps, following the training-set accuracy of Section 4.5.1 and the controlled two-turbine study of Section 4.5.2. It compares the AEP gain delivered by the XGBoost integrated yaw surrogate against `m2sr_floating` across ten layouts, from the test dataset, per mooring, with the gain measured over the yaw= 0 baseline. This is the level at which the surrogate is actually used, since it is embedded in the GA objective during co-design, so its accuracy here is what matters most for the results of Section 4.6. Figure 4.14 and Table 4.6 report the comparison. Across all three moorings the surrogate reproduces the median gain closely, with a consistent slight underestimate relative to `m2sr_floating`.

For the taut and semitaut moorings the agreement is tight. Figures 4.14a and 4.14b show the surrogate and `m2sr_floating` distributions overlapping substantially, with the XGBoost median sitting just below the optimizer in each case. The median gaps are 0.02 percentage points for taut (+0.52% versus +0.54%) and 0.02 percentage points for semitaut (+0.49% versus +0.51%), and the per-run gains correlate at  $r \geq 0.995$  (Table 4.6). For the two stiffer moorings, the surrogate prediction is therefore close to the full optimization in absolute terms.

The semitaut\_v2 mooring is the exception. Figure 4.14c shows the two distributions clearly separated, with the XGBoost median visibly below `m2sr_floating` and little overlap between the boxes. The median gap is 0.12 percentage points (+0.47% versus +0.59%), and the per-run correlation weakens as well (Table 4.6). The larger offset reflects the higher surrogate approximation error at elevated mooring compliance, consistent with the accuracy trend seen in the earlier validation steps.

Table 4.6: Farm-scale AEP gain validation: median gain of the XGBoost yaw surrogate versus `m2sr_floating` across 10 co-design layouts per mooring, with the median gap and Pearson correlation of per-run gains.

Mooring	XGB gain [%]	m2sr gain [%]	Gap [pp]	$r$
Taut	+0.52	+0.54	0.02	0.995
Semitaut	+0.49	+0.51	0.02	0.998
Semitaut_v2	+0.47	+0.59	0.12	0.937

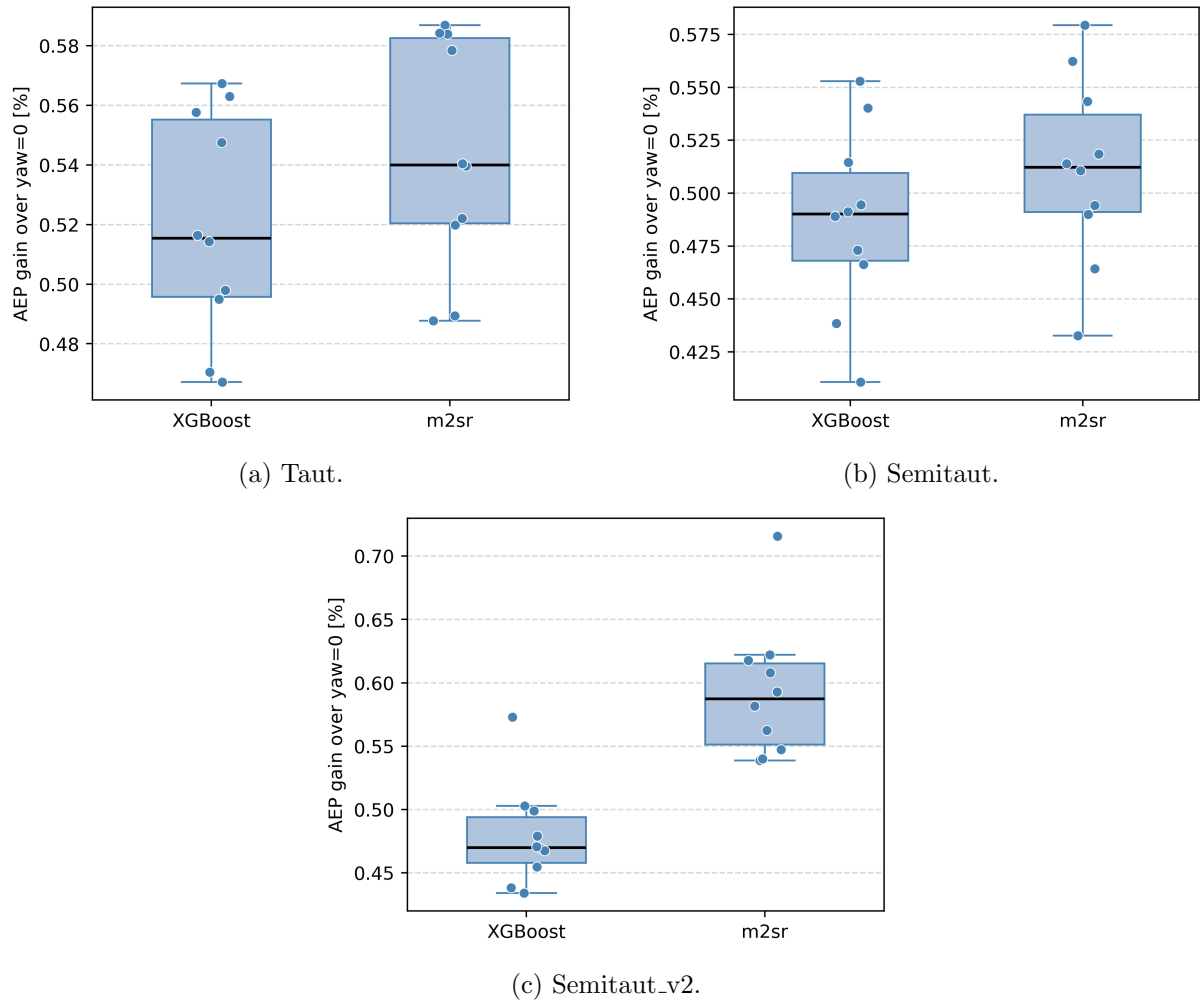


Figure 4.14: Farm-scale AEP gain validation, XGBoost yaw surrogate vs. m2sr\_floating across 10 co-design layouts. The y-axis reports AEP gain over the yaw= 0 baseline; each panel keeps its own scale. Medians, gap, and correlation are in Table 4.6.

## 4.6 Integrated Single-Loop Co-Design

### 4.6.1 Statistical AEP Comparison: Fixed, Displacement-Aware, and Co-Design

Three optimization conditions are compared for each mooring design: the fixed-bottom layout, the displacement-aware layout, and the single-loop co-design layout. Each condition is represented by ten independent GA runs (see §3.11 for why multiple runs are used), and every run is re-evaluated with `m2sr_floating` after optimization, for fair comparison. Because the three conditions optimize against different objective functions, their internal AEP estimates are not directly comparable, so reporting all of them under one displacement-aware yaw-optimization benchmark removes that mismatch. The co-design layouts are produced by the integrated single-loop pipeline, in which the yaw surrogate provides the near-optimal yaw angles for each turbine inside the GA objective (§3.8.1). The GA configuration and wind grid are held identical to the displacement-aware runs (§3.11). Figure 4.15 shows the resulting distributions, with summary statistics in Table 4.7 and median gains in Table 4.8.

For the taut mooring, the three conditions are almost indistinguishable. The median AEP is 686.2 GWh fixed-bottom, 686.4 GWh displacement-aware, and 686.4 GWh co-design (Figure 4.15a). Displacement awareness accounts for essentially the entire gain over the fixed-bottom baseline, at +0.1 GWh (+0.02%), while adding yaw co-design changes the median by less than 0.01 GWh relative to the displacement-aware case (Table 4.8). Both gains are small next to the spread, with standard deviations between 0.28 and 0.36 GWh per condition. The medians do not capture everything, however. The single best taut solution across all thirty runs comes from co-design, at 687.2 GWh (Appendix C), 0.4 GWh above the best displacement-aware run.

For the semitaut mooring, the median slightly increases across the three conditions: 686.1 GWh fixed-bottom, 686.2 GWh displacement-aware, and 686.3 GWh co-design (Figure 4.15b). Co-design gains +0.2 GWh (+0.03%) over fixed-bottom and +0.1 GWh (+0.02%) over displacement-aware, while displacement awareness alone contributes +0.1 GWh (+0.02%) (Table 4.8). This is the only mooring on which adding yaw co-design raises the median above the displacement-aware level. The gains remain small relative to the spread, with standard deviations between 0.26 and 0.33 GWh and overlapping boxes, although here too the best individual run comes from co-design (686.6 GWh, Appendix C).

For the semitaut\_v2 mooring, the ordering changes. The displacement-aware condition reaches the highest median at 687.2 GWh, ahead of both fixed-bottom at 687.0 GWh and co-design at 687.0 GWh (Figure 4.15c). Displacement awareness gains +0.2 GWh (+0.03%) over fixed-bottom, but co-design recovers only +0.03 GWh (+0.00%) over the same baseline and falls -0.2 GWh (-0.03%) below the displacement-aware result (Table 4.8). The displacement-aware distribution is also the widest, with a standard deviation of 0.74 GWh and a high outlier near 688.7 GWh, whereas the co-design distribution sits lower and tighter at a standard deviation of 0.45 GWh. Co-design underperforming displacement-aware despite incorporating yaw is consistent with the 0.12 percentage point surrogate approximation gap reported at farm scale for this mooring (§4.5.4). At this compliance the surrogate predicts the optimal yaw less accurately, so the co-design GA scores its candidates against a slightly distorted objective, and the layouts it converges to fall below the displacement-aware result once re-evaluated with `m2sr_floating`. On every mooring the median differences stay below the run-to-run standard deviation, so the three conditions overlap within the scatter.

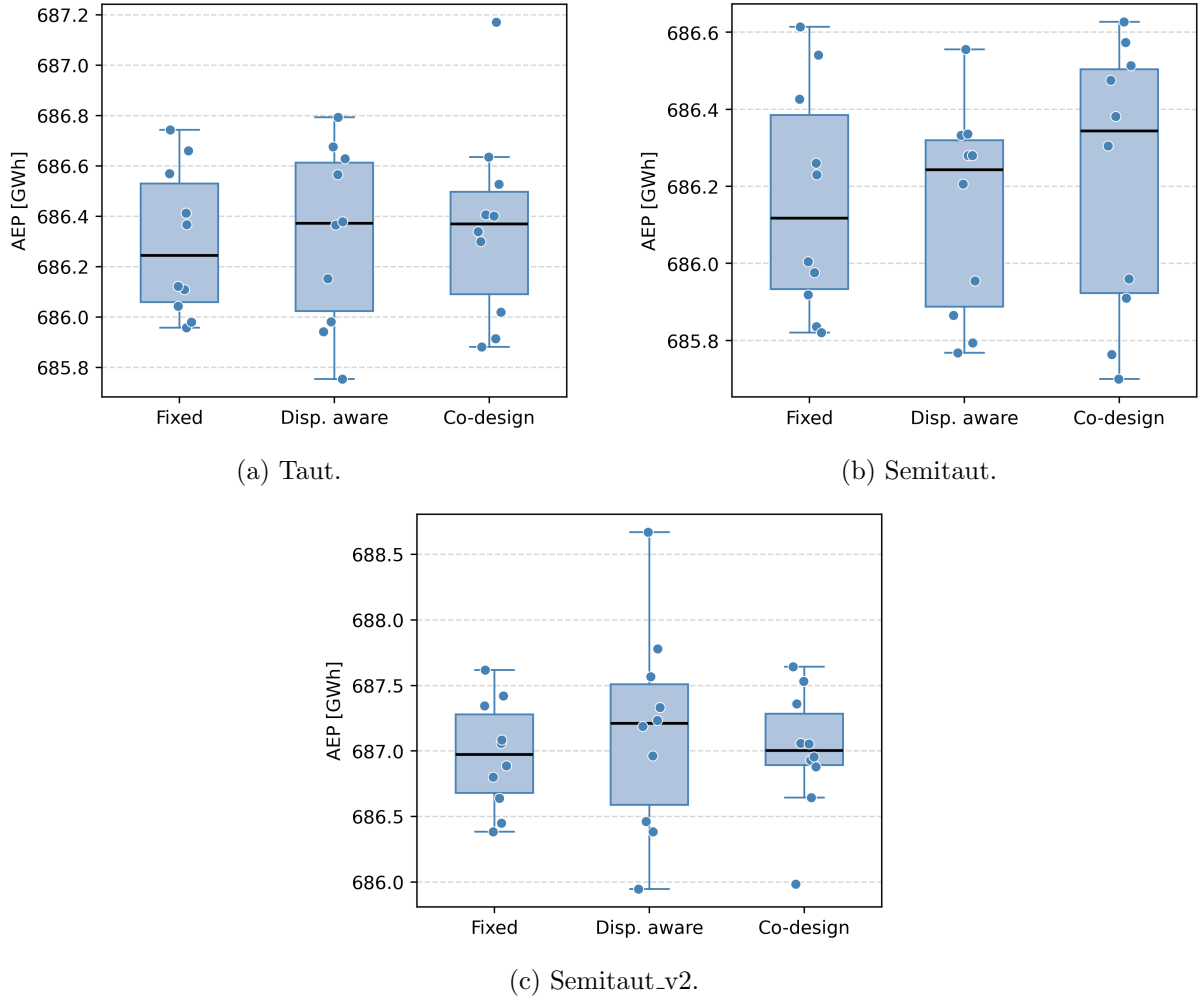


Figure 4.15: Three-way AEP comparison across moorings: fixed-bottom vs. displacement-aware vs. single-loop co-design, across 10 GA runs. All conditions re-evaluated with `m2sr_floating`; each panel keeps its own AEP scale. Median gains are summarized in Table 4.8.

Table 4.7: Median AEP and run-to-run standard deviation for the three-way comparison across all moorings ( $N = 10$  runs per condition, all re-evaluated with `m2sr_floating`). Full statistics (mean, minimum, maximum) are reported in Appendix C.

Mooring	Method	Median [GWh]	Std [GWh]
Taut	Fixed-bottom	686.2	0.28
	Disp.-aware	686.4	0.34
	Co-design	686.4	0.36
Semitaut	Fixed-bottom	686.1	0.28
	Disp.-aware	686.2	0.26
	Co-design	686.3	0.33
Semitaut_v2	Fixed-bottom	687.0	0.39
	Disp.-aware	687.2	0.74
	Co-design	687.0	0.45

Table 4.8: Median AEP gains for the three-way comparison, per mooring ( $N = 10$  runs, re-evaluated with `m2sr_floating`). Displacement-aware and co-design are compared against fixed-bottom, and co-design against displacement-aware.

Mooring	Disp. vs. Fixed		Co-design vs. Fixed		Co-design vs. Disp.	
	[GWh]	[%]	[GWh]	[%]	[GWh]	[%]
Taut	+0.13	+0.02	+0.13	+0.02	0.00	0.00
Semitaut	+0.13	+0.02	+0.23	+0.03	+0.10	+0.02
Semitaut_v2	+0.24	+0.03	+0.03	0.00	-0.21	-0.03

#### 4.6.2 Spatial Layout Comparison via the Layout Field Error

Having compared the three approaches by AEP, this section compares them by where they place turbines, using the layout field error (LFE) defined in §3.9. The metric is applied to the ten optimized layouts of each method, one comparison per mooring.

All fields are computed with  $\sigma = 5D$ . The value is adapted to the 12-turbine layouts of this case study, sparser than the large-scale farm in which the metric was introduced [56], and was selected as the width that brings out the differences between the methods most clearly. At this width the maps do not resolve the placement of individual turbines. Instead, they show in which regions of the site one method places more of its turbines than the other. In the figures that follow, each method keeps a fixed color across every comparison, with orange for the fixed-position layouts, green for the displacement-aware layouts, and purple for the co-design layouts. For a given pair, the error map runs from one method’s color through gray at zero difference to the other method’s color.

Each panel additionally overlays the best layout of each method as filled circles in the method colors. The best layout is the run with the highest AEP under `m2sr_floating` yaw optimization, the ground-truth control against which the integrated yaw surrogate is validated. The colored field, by contrast, aggregates all ten runs of each method; the complete set of turbine positions behind it is shown for the semitaut mooring in Appendix D.

Because each method is represented by ten independent runs of a stochastic genetic algorithm, the optimized layouts already vary from run to run within a single method. The regional patterns described below are therefore read as a qualitative picture of where each method tends to place more of its turbines. Because this variation between methods is comparable to the spread among a single method’s own runs, it is not taken as evidence of a systematic, method-driven difference in layout.

For the taut mooring, the displacement-aware layouts favour the left and centre of the site, while the fixed layouts place more turbines toward the eastern boundary (Figure 4.16a). Against the fixed layouts, co-design places more turbines in a band through the centre-left (Figure 4.16b). Against the displacement-aware layouts, the clearest feature is a displacement-aware concentration along the western edge, with co-design favouring the upper-right of the site, as also a band through the centre (Figure 4.16c). These shifts are modest, and the three AEP medians in §4.6.1 remain nearly identical.

For the semitaut mooring, the displacement-aware layouts favour the centre-left of the site and the fixed layouts the eastern boundary (Figure 4.17a). Against co-design, the fixed layouts again concentrate near the eastern boundary and the upper and lower left corners, with co-design filling

the centre (Figure 4.17b). The displacement-aware versus co-design map repeats the western-edge concentration of the taut case, with a second displacement-aware patch east of the centre, while co-design places more turbines through the upper centre (Figure 4.17c). This western-edge concentration of the displacement-aware layouts is visible on both taut and semitaut.

For the semitaut\_v2 mooring, the displacement-aware layouts place more turbines than the fixed ones in a band through the lower-left and centre of the site (Figure 4.18a). The fixed versus co-design map splits into one broad region per method, with co-design placing more turbines through the centre-left and the fixed layouts toward the eastern part of the site (Figure 4.18b). In the displacement-aware versus co-design comparison the two methods alternate across the site, with displacement-aware concentrations on the western side and east of the centre, and co-design concentrations through the centre and along the eastern edge (Figure 4.18c). This is also the pair for which co-design fell below displacement-aware in §4.6.1.

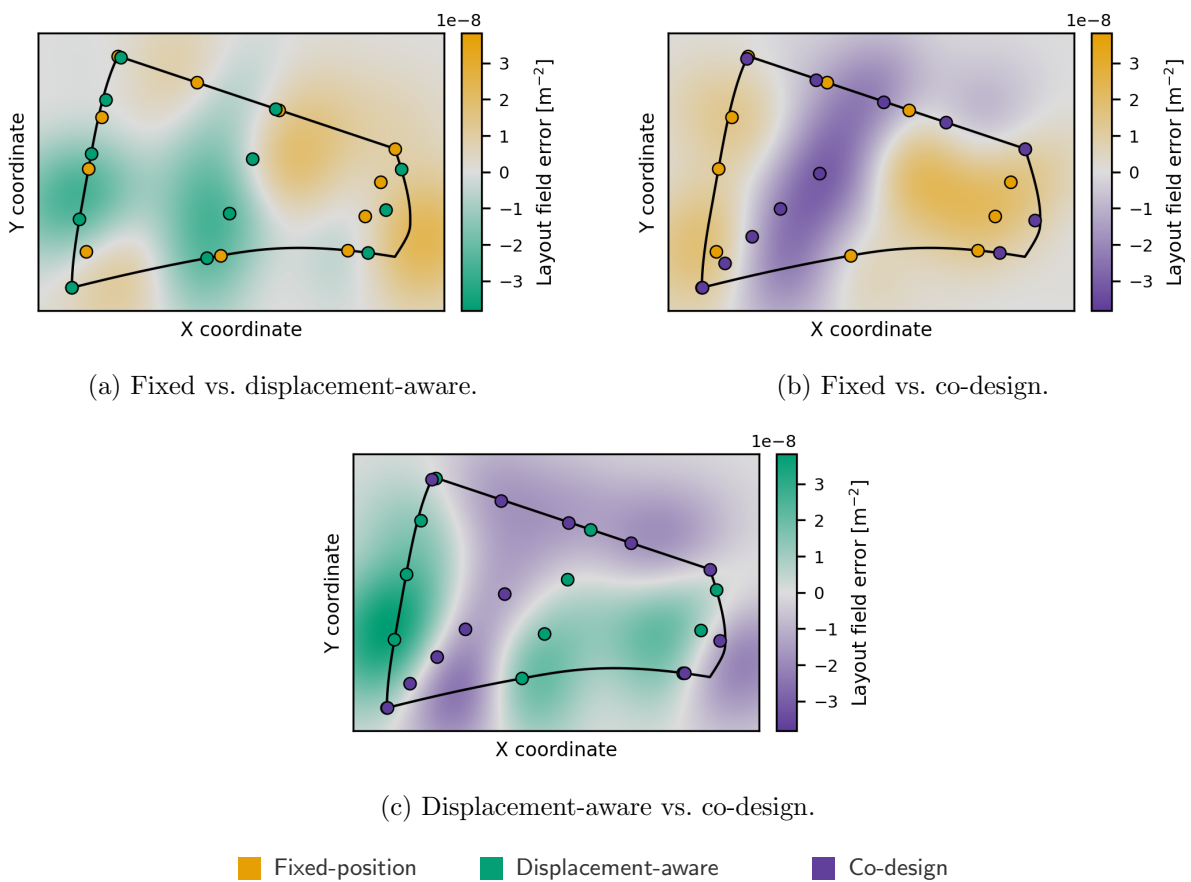


Figure 4.16: Layout field error for the taut mooring at  $\sigma = 5D$ , comparing the three optimization methods pairwise. In each panel the map runs from one method’s color, through gray at zero difference, to the other’s. Filled circles mark the best-performing layout of each method (highest AEP under `m2sr_floating` yaw optimization). The magnitude scale is shared across the three panels.

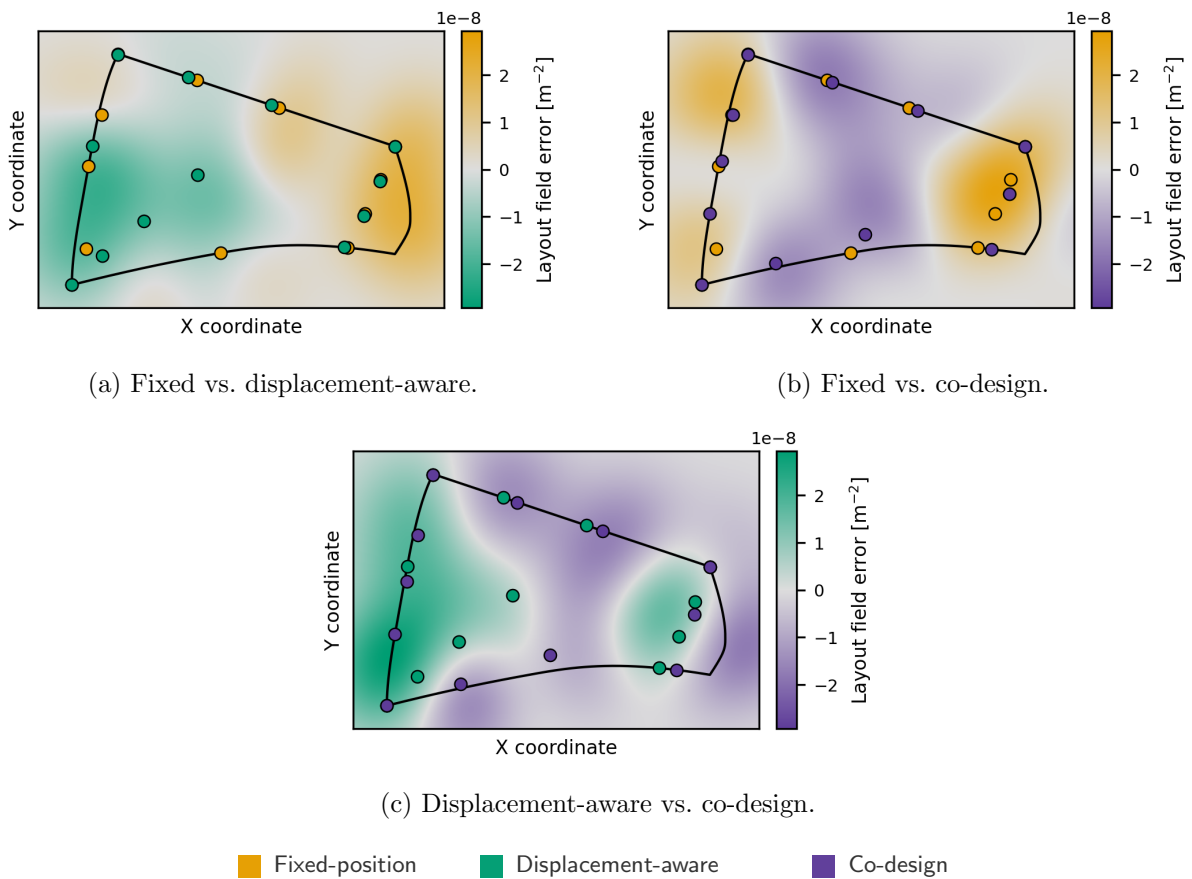


Figure 4.17: Layout field error for the semitaot mooring at  $\sigma = 5D$ , comparing the three optimization methods pairwise, with gray marking zero difference. Filled circles mark the best-performing layout of each method (highest AEP under `m2sr_floating` yaw optimization). The magnitude scale is shared across the three panels.

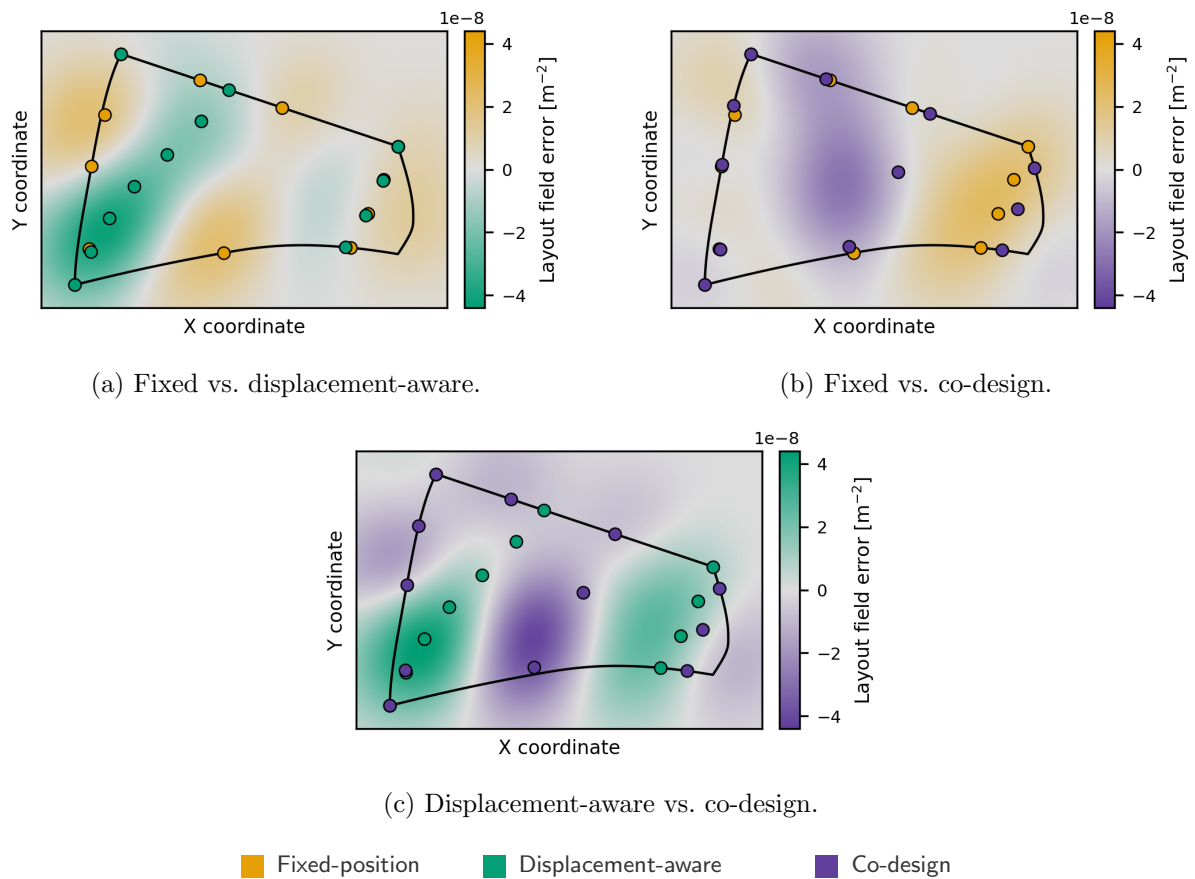


Figure 4.18: Layout field error for the semिताut.v2 mooring at  $\sigma = 5D$ , comparing the three optimization methods pairwise, with gray marking zero difference. Filled circles mark the best-performing layout of each method (highest AEP under `m2sr_floating` yaw optimization). The magnitude scale is shared across the three panels.

## 4.7 Summary of Results

This chapter presented the results of the framework, from isolated components to the full co-design comparison. The watch circles established mean platform offsets ranging from  $0.07D$  for the taut mooring to  $0.85D$  for semिताut\_v2. On the reference layout, displacement-aware yaw optimization reached farm power gains of 2.83% and 2.44% for the taut and semिताut moorings and 22.16% for semिताut\_v2. These gains hold at a single representative wind condition, not at the AEP level, and at the same condition fixed-position yaw angles produced a net loss on the semिताut\_v2 design. Both surrogates reproduced their ground truth closely, the displacement surrogate with  $R^2 \geq 0.996$  and the integrated yaw surrogate with a farm-scale gain gap of at most 0.02 percentage points on the two stiffer moorings and 0.12 on semिताut\_v2. Accounting for passive repositioning in layout optimization shifted the median AEP by less than 0.1% on every mooring. In the three-way comparison, co-design raised the median only on semिताut, while on semिताut\_v2 it fell below the displacement-aware result, although the best individual taut and semिताut layouts came from co-design. The layout field error maps compared where the three methods place their turbines, with differences that fall within the run-to-run spread of a single method and are consistent with their near-equal AEP. The interpretation of these findings and the limitations behind them are taken up in Chapter 5.

# 5

## Discussion

This chapter interprets the results of Chapter 4 against the research questions, places them next to the literature of Chapter 2, and states the limitations of the modeling choices.

### 5.1 Interpretation of Key Findings

#### Yaw Optimization on the Reference Layout

The yaw optimization results on the reference layout (§4.2) separate by mooring design. For the taut and semitaut designs, the displacement-aware optimizer reaches farm power gains of 2.83% and 2.44% at  $U = 8$  m/s and  $WD = 270^\circ$ , consistent with the gains obtained for fixed-bottom turbines with the same algorithm. At mean offsets of  $0.07D$  and  $0.19D$  the spatial layout is effectively unchanged, so the optimizer acts through wake deflection alone. At low compliance, the floating yaw problem is therefore similar to its fixed-bottom counterpart.

The semitaut\_v2 gain of 22.16% reflects a different mechanism. At offsets reaching  $0.85D$ , the optimizer uses the cross-wind component of the thrust force to physically reposition platforms out of upstream wakes. Active repositioning becomes a form of actuation, one the optimizer discovered on its own once the displacement model sat inside the loop, without any modification to the optimizer itself. The same thrust-induced repositioning underlies the yaw and induction-based repositioning of Kheirabadi and Nagamune [13], whose single-condition farm-efficiency gains rise from 16.5% to 53.5% as the moorings are made longer and more compliant. The 22% reached here, through yaw alone, falls within that range, and shows the same dependence on compliance: the large gain appears only for the most compliant design.

These single-condition gains do not translate directly to annual energy. Wake steering helps only below rated wind speed and in waked directions. In the dominant wind directions, the layout optimizer has already spaced the turbines to reduce wake overlap, so less wake loss remains for steering to recover. Averaged over the wind rose, the semitaut\_v2 gain is therefore +0.59% at AEP level (§4.5.4).

The fixed-yaw-applied-to-floating case is another important finding. A yaw solution computed for stationary turbines transfers to the compliant platform as a net loss of 0.3% (Table 4.2), because the yaw-induced drift moves rotors into wakes the fixed solution never saw. Displacement must therefore sit inside the control optimization rather than be applied after it, and the more compliant the mooring, the stronger this requirement becomes.

#### Effect of Displacement on the Layout Optimum

Accounting for passive repositioning in the layout optimization barely changes the achievable AEP (§4.4). The median gains of +0.04%, +0.05%, and +0.06% for the taut, semitaut, and semitaut\_v2 designs are smaller than the spread of AEP across the ten runs of each approach. This

magnitude is consistent with the displacement-aware PyWake studies [15, 38], which reported sub-percent AEP differences from enabling floater motion.

Two mechanisms can explain why the effect is so small. First, the mean offsets are small against inter-turbine spacings of several rotor diameters, so the drift perturbs the wake geometry of a given layout only marginally. Second, under a given wind condition all platforms drift in broadly similar directions. The relative geometry between turbines, and with it the wake overlap, therefore changes less than the absolute positions do.

The ordering of the gains, however, tracks mooring compliance ( $\text{taut} < \text{semitaut} < \text{semitaut\_v2}$ ), consistent with the watch-circle magnitudes of §4.1. The best individual `semitaut_v2` layout also comes from the displacement-aware approach, at 685.4 GWh against 684.0 GWh for the best fixed-bottom run (+1.4 GWh, +0.21%, Appendix C), which is not true for the other moorings. Where displacement is significant, displacement-aware optimization can therefore reach solutions the fixed-position approach does not, even though the median gain stays small.

### Surrogate Behavior inside the Optimization Loop

The displacement surrogate is near-exact, with  $R^2 \geq 0.996$  and a mean absolute error that is a small fraction of each design’s maximum displacement (§4.3), consistent with the high accuracy reported by Riva et al. [15], who introduced this kind of learned displacement surrogate. Feng et al. [38] used it for layout optimization, although they optimized on the fixed-position model and re-evaluated the final layouts with displacement only afterward, to keep the cost manageable. The framework developed here instead keeps the displacement prediction inside the optimization loop, evaluating every candidate layout at its displaced positions.

Of the two surrogates, the integrated yaw surrogate is therefore the one that limits the accuracy of the framework (§4.5). Its accuracy is lower than that of bottom-fixed yaw surrogates, but the two are not measured on the same problem. The  $R^2$  of 0.98 reported by Stanley et al. [49] covers all conditions, including the many easy zero-yaw cases, and the  $R^2$  of 0.990 and mean absolute error of  $0.148^\circ$  from Dirik et al. [51] come from a two-turbine case. In contrast, the full farms evaluated here involve more complex, overlapping wake interactions than a single turbine pair. The remaining difference reflects the floating problem itself, where platform drift and yaw-induced repositioning shift the turbine positions and add complexity to the wake interactions that a bottom-fixed surrogate does not face.

What the GA requires from that surrogate is not absolute accuracy but a consistent ranking of candidate layouts. Across the test layouts, the gains predicted by the surrogate correlate with the gains computed by `m2sr_floating` at  $r = 0.995$  and  $0.998$  for the `taut` and `semitaut` designs and  $0.937$  for `semitaut_v2` (§4.5.4), so the layouts are ranked in nearly the same order. Although the surrogate slightly underestimates the yaw gains, the GA can still identify the better layouts. The correlation is lowest for `semitaut_v2`, so at high compliance the ranking is preserved less well. The yaw surrogate’s accuracy also degrades systematically with mooring compliance, with the below-rated  $R^2$  near  $0.88$  for the `taut` and `semitaut` designs and dropping to  $0.77$  for `semitaut_v2`: the more the platform moves, the harder the optimal-yaw mapping is to learn. The two-turbine maps show that for the most compliant design the optimal yaw patterns differ from those of the fixed and low-compliance cases, concentrating the error at the downstream wake centre. The consequence at co-design level is taken up in the three-way comparison below. The surrogate could be improved by adding training data in structured layouts, such as the two-turbine maps, or by a different model architecture.

### The Three-Way Co-Design Comparison

In the three-way comparison (§4.6), the added value of co-design over displacement-aware optimization is small on every mooring. Semitaut is the only design with the expected ordering, where co-design adds +0.02% over the displacement-aware approach. The margin is measurable but small. On taut, co-design adds nothing over displacement awareness (less than 0.01 GWh).

For semitaut\_v2 the ordering changes, with co-design falling  $-0.03\%$  below the displacement-aware result. A possible explanation is the farm-scale approximation gap of 0.12 percentage points: the objective no longer reflects the true AEP faithfully at high compliance, and an optimizer driven by an objective that departs from the truth converges toward solutions that are suboptimal under the true `m2sr_floating` evaluation.

These findings are consistent with the fixed-bottom counterpart of this comparison by Baricchio et al. [63], which embeds an approximation of the optimal control in the layout GA under an equivalent protocol. Their sequential approach, in which the layout is optimized first and the control is optimized on the result, corresponds to the assessment of both the fixed and displacement-aware approaches in the three-way comparison of this thesis, and they likewise report sequential and co-design AEP gains that are effectively identical.

### Layout Comparison at Near-Equal AEP

The LFE maps (§4.6.2) compare where the three methods place their turbines. As shown there, the differences between methods fall within the run-to-run scatter of a single method's ten optimization runs. This is consistent with a multimodal AEP landscape, in which many different layouts reach a similar optimum [61, 62]. For this reason, no systematic conclusions about the layouts are drawn from these difference maps.

The displacement-aware layouts tend to lie towards the upwind, western side of the site, while the co-design layouts place more turbines through the centre. This is consistent with the information each method uses, since the displacement-aware optimizer accounts only for passive drift and favours upwind positions, while co-design anticipates the yaw control that recovers part of the wake cost and can accept more central placements. Because this stays within the run-to-run scatter, it is read as a description of the maps rather than an established difference between the methods.

Because many different layouts reach a similar AEP, no single layout is clearly best. The designer is therefore free to choose among the near-optimal layouts on criteria beyond energy, such as mooring and cabling cost or structural loads. An objective that includes cost, such as the levelized cost of energy (LCOE), would fold that choice into the optimization itself. With AEP nearly equal across these layouts, it would be driven by cost, favoring the cheaper ones.

## 5.2 Limitations

The findings of this chapter hold within a set of modeling choices, stated below in the order of the modeling pipeline, from the wake model to the experimental design.

**Wake model fidelity:** All AEP evaluations use a steady-state Gaussian deficit with Jiménez deflection. Dynamic wake meandering and the interaction between platform motion and wake dynamics are not represented, and the absolute AEP uncertainty of engineering wake models likely exceeds the AEP differences between the optimization approaches reported in Chapter 4.

**Simplified mooring and control assumptions:** The displacement model covers quasi-static mean offsets only, without dynamic oscillations. Control is limited to yaw with an AEP-only objective, and the optimal yaw is a quasi-static `m2sr_floating` solution, so the load penalty of sustained misalignment, the dynamic behaviour of these setpoints, and their wider effects on the farm are not assessed. The mooring designs are fixed inputs rather than design variables, so compliance is sampled through three discrete cases rather than optimized, and the mooring is not co-designed with the layout and control.

**Unvalidated high-compliance case (`semitaut_v2`):** The `semitaut_v2` mooring is an in-house variant built to reach high compliance, and it is not validated against a real mooring design. Its large offsets, up to  $0.85D$ , produce the largest single-condition yaw gain of the three moorings, near 22%. Its distinct watch-circle shape gives the one-sided control ground truth examined in §4.5.3. The integrated yaw surrogate is also least accurate for this design, since that one-sided ground truth is harder to learn. The `semitaut_v2` case therefore demonstrates the importance of accounting for displacement in highly compliant designs, but its specific numbers should not be read as quantitative predictions.

**Weak wake–displacement coupling:** The displacement of each platform is predicted once, from the effective wind speed at its nominal position, and the wake field is not recomputed at the displaced positions. For large offsets, the wake a turbine actually sits in can therefore differ from the one used to predict its displacement. Iterating the two to a steady state is feasible [15, 50], but costly inside an optimizer: Riva et al. recommend enabling floater motion only in the final phase of layout optimization, Feng et al. [38] optimize with the fixed-position model and re-evaluate with the floating one. The single-pass choice trades full consistency for a displacement model cheap enough for every GA evaluation.

**Constraints enforced on nominal positions only:** Displaced coordinates are not checked against the spacing threshold or the site boundary at runtime. Turbines nominally separated beyond the  $0.5D$  threshold may drift closer during operation, and platform offsets may carry displaced positions outside the site polygon. The risk grows with compliance: it is negligible for `taut`, with offsets of  $0.07D$ , but plausible for `semitaut_v2`, whose offsets reach  $0.85D$ . Further constraints could be imposed inside the optimizer.

**Statistical power of the comparisons:** Each approach is represented by ten independent runs, and most median gains are smaller than the standard deviation of the AEP across the ten runs of a single approach. The remaining differences are therefore difficult to separate from the noise the stochastic GA itself produces.

**Single-site test case:** All results are obtained for one site polygon, one strongly westerly wind rose, one turbine and platform combination, and 12 turbines. The finding that displacement matters more for yaw as the mooring becomes more compliant follows from turbine repositioning, not from this particular site, so it is expected to hold elsewhere. The layout and co-design results are more site-conditional, and how they would change under a more multidirectional or heterogeneous wind resource is untested. The integrated yaw surrogate is trained on this single configuration, and whether it generalizes to other farm sizes or sites was not investigated and remains unknown.

# 6

## Conclusions and Recommendations

This thesis developed and assessed an integrated, single-loop layout and control co-design framework for floating offshore wind farms. The framework extends the fixed-bottom scheme of Baricchio et al. [21] to account for passive (drift-induced) and active (yaw-induced) turbine repositioning. Two XGBoost surrogates make the extension tractable: one predicts the platform displacement of each turbine, and one predicts its optimal yaw angle. The framework was assessed on the Kriti 3 case study, with 12 IEA 15 MW turbines and three mooring designs spanning stiff to highly compliant. No published framework had combined single-loop co-design with both repositioning mechanisms: single-loop co-design existed only for fixed-bottom farms, and floating co-design only in nested form. This chapter answers the research questions posed in Chapter 2, states the contributions, and closes with recommendations for future work.

### 6.1 Answers to the Research Questions

#### RQ1 (Active repositioning)

How does accounting for active (yaw-induced) repositioning in yaw optimization change the corresponding yaw angles and the achievable wake-steering gains, compared with fixed-position yaw optimization?

Accounting for active repositioning changes both the optimal yaw angles and the gains they deliver, and its importance scales with mooring compliance. On the stiff taut and semitaut designs the displacement-aware optimal angles and the resulting farm-power gains stay close to fixed-bottom wake steering (+2.83% and +2.44% at  $U = 8$  m/s and  $270^\circ$ ). On the highly compliant semitaut\_v2 design the optimal angles differ markedly, taking opposite signs at the middle column to drive the platforms clear of the upstream wakes, and the gain at that condition rises to +22.16%. Critically, applying the fixed-position optimal yaw to the same floating platform turns the expected gain into a net loss of 0.3%, because the yaw-induced drift moves the rotors into wakes the fixed-position optimizer never evaluated. Accounting for active repositioning inside the yaw optimization is therefore necessary on compliant platforms.

#### RQ2 (Passive repositioning)

How does accounting for passive repositioning alone (displacement-aware) change the optimal layout and the achievable AEP compared to fixed-position layout optimization?

Accounting for passive repositioning has a small but consistent effect on the achievable AEP. Across the three moorings, the median gains of displacement-aware over fixed-position optimization stay below 0.1% and are smaller than the standard deviation of the AEP across the ten runs

of a single approach, with the gain ordering following mooring compliance (§4.4.2). The medians are not the whole answer, however. On `semitaut_v2`, the design with the biggest displacement, the best individual layout comes from the displacement-aware approach, at +0.21% over the best fixed-position run (Appendix C). Where displacement is significant, displacement-aware optimization reaches solutions that the fixed-position approach does not.

The optimized layouts themselves vary between runs. Each method returns ten different layouts of near-equal AEP, and the displacement-aware layouts differ from the fixed-position ones by an amount comparable to the difference between two runs of the same method. The fixed and displacement-aware maps do show the displacement-aware layouts leaning towards the upwind side of the site, the side of the dominant wind direction (§4.6.2). Because this lean stays within that run-to-run variation, it is reported as an observation and not as an established difference.

### RQ3 (Optimized yaw surrogate)

How can geometric yaw be extended to floating turbines so that layout and control co-design becomes displacement-aware and remains single-loop?

Geometric yaw transfers to floating turbines by replacing the analytic relation with a learned surrogate. Three XGBoost models, one per mooring, are trained on the optimal yaw angles computed by `m2sr_floating` (§3.7). At every AEP evaluation the surrogate predicts the yaw angles directly from the layout geometry and the wind state, with no nested control optimization, so the co-design remains single-loop. Because the training labels come from the displacement-aware optimizer, the predicted angles account for the repositioning that each yaw setting induces.

On `taut` and `semitaut` the surrogate reproduces the farm-scale AEP gain of `m2sr_floating` to within 0.02 percentage points and preserves the ranking of candidate layouts, which is the property the genetic algorithm needs (§4.5.4). On `semitaut_v2`, the highest compliance design, the farm-scale gap widens to 0.12 percentage points. The optimal-yaw mapping becomes more complex there, with each turbine’s optimal angle depending on the positions and optimal angles of its upstream and downstream neighbours, and is therefore harder to learn. Geometric yaw can thus be extended to floating turbines through a learned surrogate that keeps co-design single-loop and displacement-aware.

### RQ4 (Mooring sensitivity)

How do different mooring designs affect the AEP gains and the optimal layouts?

The three mooring designs produce mean platform offsets spanning nearly two orders of magnitude, from  $0.07D$  for `taut` to  $0.85D$  for `semitaut_v2` (§4.1). On the reference layout, the `taut` and `semitaut` designs yield wake-steering gains of 2.8% and 2.4% at the dominant wind condition, consistent with fixed-bottom cases. The `semitaut_v2` gain of 22.2% at the same condition arises from a different mechanism, active repositioning out of upstream wakes, and is a consequence of the compliant `semitaut_v2` design. Averaged over the wind rose, it compresses to roughly half a percent of AEP (§4.2, §4.5.4).

On a compliant design, ignoring displacement in the yaw optimization is no longer merely sub-optimal but counterproductive. Fixed-position optimal yaw applied to the floating `semitaut_v2` platform turns a predicted +3.0% gain into a -0.3% net loss (§4.2.2). Adding yaw co-design on top of displacement awareness changes the median AEP only marginally, and the direction

differs per mooring: a small gain on semitaut, none on taut, and a loss on semitaut\_v2 (§4.6.1). The best individual taut and semitaut layouts nonetheless come from co-design runs. The mooring design therefore shapes every result in this thesis: the yaw-optimization gains, the effect of displacement awareness, and the outcome of co-design all differ across the three designs.

### Main Research Question

How does integrated layout and control co-design, accounting for passive and yaw-induced turbine repositioning, change the optimal layout and AEP compared to fixed-position approaches?

Integrated single-loop co-design accounting for passive and active repositioning is feasible and tractable. A co-design run completes in approximately 20 hours on 48 cores, where a nested formulation calling `m2sr_floating` at every candidate layout would be computationally prohibitive (§3.11).

At this site and farm size, the AEP is largely insensitive to the choice of method. The median AEP differences between fixed-position, displacement-aware, and co-design optimization stay below 0.1%, and the optimized layouts differ by about as much as the run-to-run scatter of a single method, so the framework does not produce a systematically different layout. The clearest practical finding sits on the control side rather than the layout side. Yaw setpoints computed under a fixed-position assumption can transfer to a compliant floating platform as a net power loss, so displacement must be accounted for inside the control optimization, and the more compliant the mooring, the stronger this requirement. Because many different layouts reach a near-equal AEP, the designer is left free to select among them on criteria beyond AEP, such as cabling or loads.

## 6.2 Contributions

This thesis made four contributions to the floating wind farm co-design literature:

1. The development of a displacement-aware layout optimization methodology for floating wind farms, embedding an ML displacement surrogate inside the GA objective, which enabled the quantification of the effect of passive repositioning against fixed-position optimization.
2. An extension of the geometric-yaw concept to floating turbines through surrogate models trained on displacement-aware optimal-yaw labels, keeping co-design single-loop.
3. An integrated single-loop co-design framework accounting for both passive and active repositioning, with its AEP gains quantified against displacement-aware and fixed-position optimization under a common displacement-aware optimal layout re-evaluation.
4. A sensitivity analysis across three mooring designs with different displacements, establishing compliance as an important parameter.

## 6.3 Recommendations for Future Work

The directions below follow from the limitations and findings of this work:

- The integrated yaw surrogate can be improved with training data enriched by structured layouts, such as two-turbine sweeps, and by revisiting the width parameters of the wake cone filter.
- A graph neural network (GNN) could replace tabular models, such as the gradient-boosted trees used here or a feed-forward neural network. Representing the farm as a graph captures wake interactions natively and removes the encoding of the neighboring turbines.
- The per-condition AEP calculations are currently executed sequentially, since each call requires independently computed displaced positions. Computational cost could be reduced if the proposed displacement-aware approach were integrated inside PyWake, with the AEP calculation vectorized instead of looping per condition.
- Constraints could be enforced on displaced rather than nominal positions, since operational spacing violations and boundary exceedance are plausible at `semitaut_v2` compliance, with offsets of  $0.85D$  against a  $0.5D$  spacing threshold (§5.2). The objective function could be constrained to penalize such positions, or the boundary shrunk by the maximum displacement.
- The behaviour of the yaw optimizer on the most compliant design starts to deviate from the two stiffer moorings, and should be better understood.
- Other or combined control strategies (derating, dynamic wake mixing such as the Helix) could be tested in the current co-design framework, as they would also allow the comparison between co-design and displacement-aware optimization to be repeated under a control strategy with a larger AEP contribution, where the margin between the two is expected to widen.
- The platform orientation could be treated as a design variable rather than a fixed input.
- The experimental basis can be broadened to larger farms and to other sites and wind roses, including heterogeneous inflow.
- The wake-displacement coupling can be strengthened: the single surrogate pass takes  $U_{\text{eff}}$  at nominal positions, and iterating to a self-consistent equilibrium between displaced positions and the wake field would remove this inconsistency at extra cost (§5.2).

# Acknowledgements

This thesis marks the end of my MSc in Sustainable Energy Technology at TU Delft, and the people below made it possible.

First, I would like to thank my supervisor, Prof. Jan-Willem van Wingerden, for giving me the opportunity to work on this project and for allowing me the freedom to shape it in my own way. His guidance helped keep the work on track, and his questions during our meetings often made me rethink things I thought I had already understood.

I am especially grateful to Matteo Baricchio, my daily supervisor. Matteo was there throughout the day-to-day work, always willing to help whenever I was stuck, and our collaboration was one of the most enjoyable parts of this thesis. Much of what I now understand about this topic I learned from him. Since this thesis also builds on his work, it feels only right that he played such an important role in shaping it.

I would also like to thank Maarten van den Broek from sowento for providing the SLOW simulator and for his support with the floating modelling aspects of the work. His insights and suggestions played an important role in shaping several key decisions in this thesis.

My thanks also go to the team at NTUA, for providing the site data that grounded this work in a realistic case study, and to Jenna Iori, for serving on my thesis committee and taking the time to read and evaluate this thesis.

Beyond the thesis itself, I am grateful for the two years I spent in the Netherlands. Living in Delft and working within the academic environment at TU Delft gave me the space to grow, both personally and academically.

Finally, I want to thank the people closest to me. To my parents, my brother, and my sister, thank you for supporting me from Greece. To Foteini, thank you for your patience, support, and for always being there for me throughout this process. And to my friends, both here in the Netherlands and back in Greece, thank you for making these years so enjoyable and memorable.

Konstantinos Chondros  
Delft, July 2026

# A

## Surrogate Model Hyperparameters

Both the displacement surrogate (§3.3.3) and the integrated yaw surrogate (§3.7) are XGBoost gradient-boosted regression models. Their hyperparameters are collected in Table A.1. The displacement values are shared across all three mooring models. The yaw values are likewise shared, with the single exception of the tree depth, which differs per mooring (§3.7.3).

Table A.1: XGBoost hyperparameters for the displacement and integrated yaw surrogates. Displacement values are shared across all three mooring models (taut, semitaut, semitaut.v2); yaw values are shared except for the per-mooring tree depth. The displacement surrogate maps five input features to a two-dimensional output (mean surge and sway); the yaw surrogate maps 34 input features to a single output (optimal yaw angle).

Hyperparameter	Displacement	Yaw
n_estimators	1200	1200
max_depth	5	7–9 <sup>1</sup>
learning_rate	0.05	0.1
subsample	0.8	0.8
colsample_bytree	0.8	0.8
reg_lambda	3.0	3.0
reg_alpha	0.1	0.1
min_child_weight	5	5
tree_method	hist	hist

<sup>1</sup> Per-mooring tree depth: taut 9, semitaut 7, semitaut.v2 9.

# B

## Displacement Surrogate Diagnostics

Figure 4.4 in §4.3 shows the displacement surrogate diagnostics for the semitaut.v2 mooring, taken as the representative case because its compliance produces the largest displacements. The diagnostics for the taut and semitaut moorings are reported here for completeness. In both, the predicted surge and sway lie on the dashed 1:1 reference and the residuals scatter symmetrically about zero with no systematic trend against wind direction  $\theta$  or yaw misalignment  $\gamma_\Delta$ , consistent with the near-unity  $R^2$  values in Table 4.3.

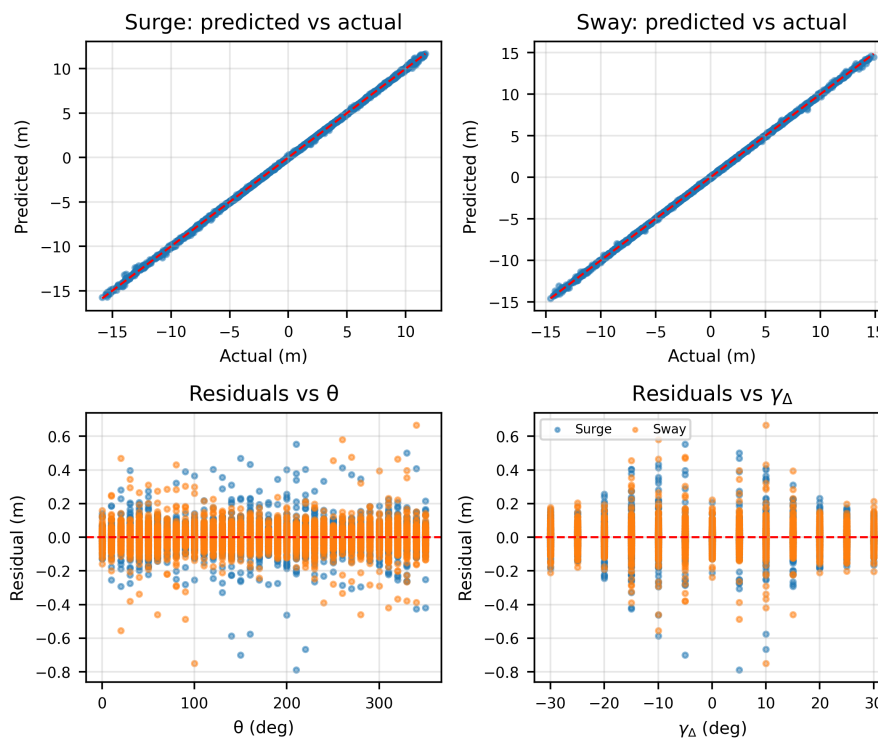


Figure B.1: Displacement surrogate diagnostics for the taut mooring. Top: surge (left) and sway (right) predicted vs. actual, with the dashed 1:1 reference. Bottom: residuals against wind direction  $\theta$  (left) and yaw misalignment  $\gamma_\Delta$  (right).

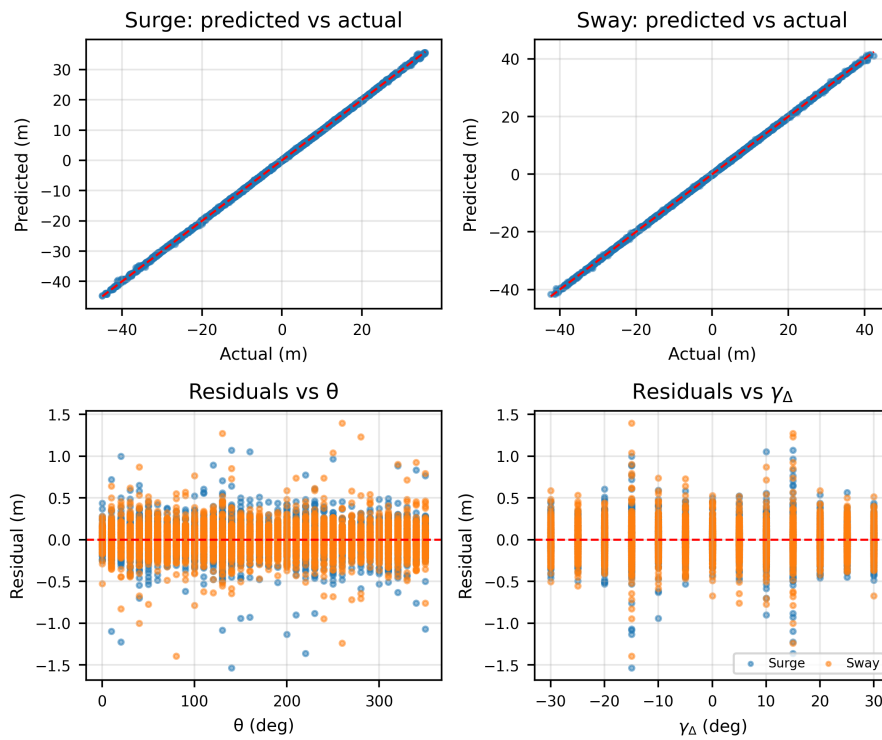


Figure B.2: Displacement surrogate diagnostics for the semitaot mooring. Top: surge (left) and sway (right) predicted vs. actual, with the dashed 1:1 reference. Bottom: residuals against wind direction  $\theta$  (left) and yaw misalignment  $\gamma_{\Delta}$  (right).

# C

## Optimization Run Statistics

This appendix reports the complete run-to-run statistics for the layout-only (§4.4.2) and three-way (§4.6.1) AEP comparisons. The main text (Tables 4.4 and 4.7) lists the median and standard deviation; the mean, minimum, and maximum across the ten runs per condition are collected here for completeness. Values follow the same rounding as the main text, with AEP to one decimal place and the standard deviation to two.

Table C.1: Full statistics for the fixed-bottom versus displacement-aware layout comparison ( $N = 10$  runs per condition, evaluated at zero yaw with the displacement-aware wrapper). Summarised in Table 4.4.

<b>Mooring</b>	<b>Method</b>	<b>Median</b> [GWh]	<b>Mean</b> [GWh]	<b>Std</b> [GWh]	<b>Min</b> [GWh]	<b>Max</b> [GWh]
Taut	Fixed-bottom	683.4	683.4	0.38	682.9	684.0
	Disp.-aware	683.7	683.5	0.33	683.0	683.9
Semitaut	Fixed-bottom	683.3	683.4	0.38	682.9	683.9
	Disp.-aware	683.7	683.5	0.31	683.0	683.8
Semitaut_v2	Fixed-bottom	683.4	683.4	0.38	682.9	684.0
	Disp.-aware	683.8	683.7	0.84	682.2	685.4

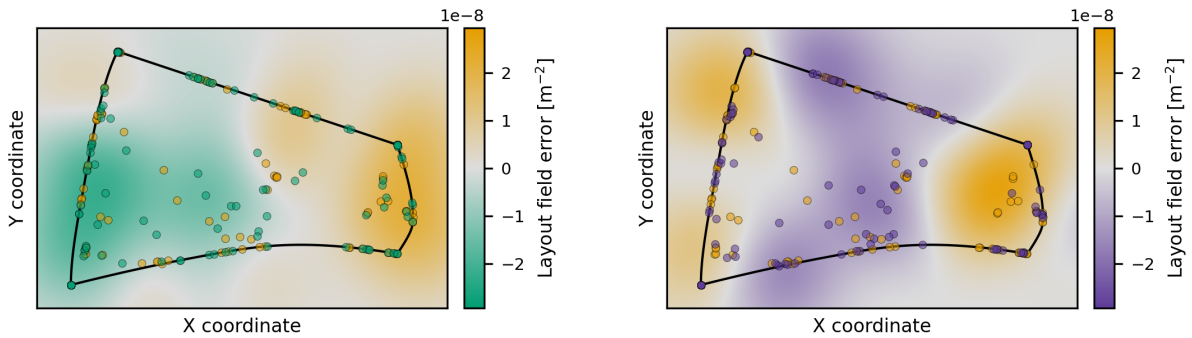
Table C.2: Full statistics for the three-way AEP comparison ( $N = 10$  runs per condition, all re-evaluated with `m2sr_floating`). Summarised in Table 4.7.

<b>Mooring</b>	<b>Method</b>	<b>Median</b> [GWh]	<b>Mean</b> [GWh]	<b>Std</b> [GWh]	<b>Min</b> [GWh]	<b>Max</b> [GWh]
Taut	Fixed-bottom	686.2	686.3	0.28	686.0	686.7
	Disp.-aware	686.4	686.3	0.34	685.8	686.8
	Co-design	686.4	686.4	0.36	685.9	687.2
Semitaut	Fixed-bottom	686.1	686.2	0.28	685.8	686.6
	Disp.-aware	686.2	686.1	0.26	685.8	686.6
	Co-design	686.3	686.2	0.33	685.7	686.6
Semitaut_v2	Fixed-bottom	687.0	687.0	0.39	686.4	687.6
	Disp.-aware	687.2	687.2	0.74	685.9	688.7
	Co-design	687.0	687.0	0.45	686.0	687.6

# D

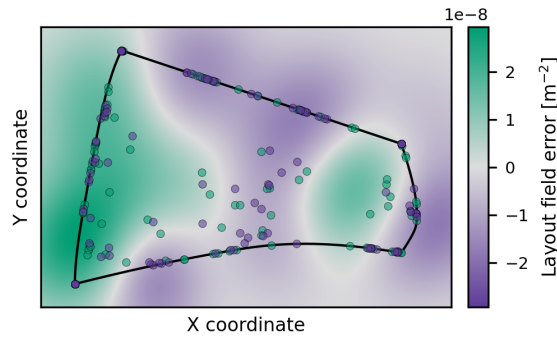
## Underlying Layouts of the Layout Field Error Maps

The layout field error maps of §4.6.2 present each method as a smoothed field rather than as discrete turbines. Figure D.1 overlays, for the semitaot mooring, every turbine of all ten runs of each method on the corresponding difference map, making explicit that the colored field is a smoothed density of these positions.



(a) Fixed vs. displacement-aware.

(b) Fixed vs. co-design.



(c) Displacement-aware vs. co-design.

■ Fixed-position    ■ Displacement-aware    ■ Co-design

Figure D.1: Layout field error for the semitaot mooring with every turbine of all ten runs per method overlaid. The colored field is a smoothed density of these positions: a region takes a method's color where that method places, on average, slightly more turbines than the other.

# References

- [1] International Energy Agency, “Offshore Wind Outlook 2019: World Energy Outlook Special Report,” International Energy Agency, Paris, Tech. Rep., 2019. [Online]. Available: <https://www.iea.org/reports/offshore-wind-outlook-2019>
- [2] A. McCoy, W. Musial, R. Hammond, D. M. Hernando, P. Duffy, P. Beiter, P. Pérez, R. Baranowski, G. Reber, and P. Spitsen, “Offshore Wind Market Report: 2024 Edition,” National Renewable Energy Laboratory, Golden, CO, Tech. Rep. NREL/TP-5000-90525, 2024. [Online]. Available: <https://docs.nrel.gov/docs/fy24osti/90525.pdf>
- [3] WindEurope, “Unleashing Europe’s offshore wind potential: A new resource assessment,” WindEurope, Brussels, Tech. Rep., 2017. [Online]. Available: <https://windeurope.org/data/products/unleashing-europe-s-offshore-wind-potential/>
- [4] EMODnet Bathymetry Consortium, “EMODnet digital bathymetry (DTM 2024),” 2024. [Online]. Available: <https://emodnet.ec.europa.eu/geoviewer/>
- [5] C. Allen, A. Viselli, H. Dagher, A. Goupee, E. Gaertner, N. Abbas, M. Hall, and G. Barter, “Definition of the UMaine VoltturnUS-S reference platform developed for the IEA wind 15-megawatt offshore reference wind turbine,” National Renewable Energy Laboratory, Golden, CO, USA, Tech. Rep. NREL/TP-5000-76773, 2020, iEA Wind TCP Task 37. [Online]. Available: <https://www.nrel.gov/docs/fy20osti/76773.pdf>
- [6] University of Maine Advanced Structures and Composites Center, “UMaine to expand offshore wind curriculum with new course offering for spring 2023,” University of Maine, Advanced Structures and Composites Center, 2022. [Online]. Available: <https://composites.umaine.edu/2022/10/19/umaine-to-expand-offshore-wind-curriculum-with-new-course-offering-for-spring-2023/>
- [7] A. C. Kheirabadi and R. Nagamune, “A quantitative review of wind farm control with the objective of wind farm power maximization,” *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 192, pp. 45–73, Sep. 2019.
- [8] G. Mosetti, C. Poloni, and B. Diviacco, “Optimization of wind turbine positioning in large windfarms by means of a genetic algorithm,” *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 51, no. 1, pp. 105–116, Jan. 1994.
- [9] Á. Jiménez, A. Crespo, and E. Migoya, “Application of a LES technique to characterize the wake deflection of a wind turbine in yaw,” *Wind Energy*, vol. 13, no. 6, pp. 559–572, Sep. 2010.
- [10] P. Fleming, J. Annoni, J. J. Shah, L. Wang, S. Ananthan, Z. Zhang, K. Hutchings, P. Wang, W. Chen, and L. Chen, “Field test of wake steering at an offshore wind farm,” *Wind Energy Science*, vol. 2, no. 1, pp. 229–239, May 2017.

- [11] A. C. Kheirabadi and R. Nagamune, “Real-time relocation of floating offshore wind turbine platforms for wind farm efficiency maximization: An assessment of feasibility and steady-state potential,” *Ocean Engineering*, vol. 208, p. 107445, Jul. 2020.
- [12] K. Chondros, “Displacement-aware yaw optimization for floating wind farms,” Delft University of Technology, Delft, the Netherlands, University Research Project (SET4400), Nov. 2025, Delft Center for Systems and Control.
- [13] A. C. Kheirabadi and R. Nagamune, “Modeling and Power Optimization of Floating Offshore Wind Farms with Yaw and Induction-based Turbine Repositioning,” in *2019 American Control Conference (ACC)*, Jul. 2019, pp. 5458–5463.
- [14] E. Lozon, M. Hall, and M. Y. Mahfouz, “Coupled modeling of wake steering and platform offsets for floating wind arrays,” *Journal of Physics: Conference Series*, vol. 2767, no. 6, p. 062035, Mar. 2024.
- [15] R. Riva, M. M. Pedersen, G. Pirrung, H. Bredmose, and J. Feng, “Incorporation of floater rotation and displacement in a static wind farm simulator,” *Journal of Physics: Conference Series*, vol. 2767, no. 6, p. 062019, Mar. 2024.
- [16] A. P. J. Stanley, C. J. Bay, and P. Fleming, “Enabling control co-design of the next generation of wind power plants,” *Wind Energy Science*, vol. 8, no. 8, pp. 1341–1350, Aug. 2023.
- [17] Y. Niu, A. Dwivedi, J. Sathiaraj, P. P. Lathi, and R. Nagamune, “Floating Offshore Wind Farm Control via Turbine Repositioning: Unlocking the Potential Unique to Floating Offshore Wind,” *IEEE Control Systems*, vol. 44, no. 5, pp. 106–129, Oct. 2024.
- [18] N. Jensen, “A note on wind generator interaction,” Risø National Laboratory, Roskilde, Report 87-550-0971-9, 1983.
- [19] M. Bastankhah and F. Porté-Agel, “A new analytical model for wind-turbine wakes,” *Renewable Energy*, vol. 70, pp. 116–123, Oct. 2014.
- [20] I. Katic, J. Højstrup, and N. O. Jensen, “A Simple Model for Cluster Efficiency,” in *EWEC’86. Proceedings. Vol. 1*, W. Palz and E. Sesto, Eds. Rome, Italy: A. Raguzzi, 1987, pp. 407–410.
- [21] M. Baricchio, P. M. O. Gebraad, and J.-W. van Wingerden, “Evaluating the potential of a wake steering co-design for wind farm layout optimization through a tailored genetic algorithm,” *Wind Energy Science*, vol. 9, no. 11, pp. 2113–2132, Nov. 2024.
- [22] M. Hall, M. Biglu, S. Housner, K. Coughlan, M. Y. Mahfouz, and E. Lozon, “Floating Wind Farm Layout Optimization Considering Moorings and Seabed Variations,” *Journal of Physics: Conference Series*, vol. 2767, no. 6, p. 062038, Jun. 2024.
- [23] L. Wang, A. Tan, and Y. Gu, “A novel control strategy approach to optimally design a wind farm layout,” *Renewable Energy*, vol. 95, pp. 10–21, Sep. 2016.
- [24] D. Stockhouse, M. Phadnis, A. Henry, N. J. Abbas, M. Sinner, M. Pusch, and L. Y. Pao, “A Tutorial on the Control of Floating Offshore Wind Turbines: Stability Challenges and Opportunities for Power Capture,” *IEEE Control Systems*, vol. 44, no. 5, pp. 28–57, Jul. 2024.

- [25] D. Stockhouse, M. Phadnis, A. Henry, N. Abbas, M. Sinner, M. Pusch, and L. Y. Pao, “Sink or Swim: A Tutorial on the Control of Floating Wind Turbines,” in *2023 American Control Conference (ACC)*. San Diego, CA, USA: IEEE, May 2023, pp. 2512–2529.
- [26] E. Gaertner, J. Rinker, L. Sethuraman, F. Zahle, B. Anderson, G. Barter, N. Abbas, F. Meng, P. Bortolotti, W. Skrzypinski, G. Scott, R. Feil, H. Bredmose, K. Dykes, M. Shields, C. Allen, and A. Viselli, “Definition of the IEA wind 15-megawatt offshore reference wind turbine,” National Renewable Energy Laboratory, Golden, CO, USA, Tech. Rep. NREL/TP-5000-75698, 2020, iEA Wind TCP Task 37. [Online]. Available: <https://www.nrel.gov/docs/fy20osti/75698.pdf>
- [27] D. De Tavernier and A. Viré, “Floating offshore wind energy (AE4W31), lecture 1: Introduction,” Lecture slides, Delft University of Technology, 2024, unpublished course material.
- [28] A. Otter, J. Murphy, V. Pakrashi, A. Robertson, and C. Desmond, “A review of modelling techniques for floating offshore wind turbines,” *Wind Energy*, vol. 25, no. 5, pp. 831–857, 2022.
- [29] National Renewable Energy Laboratory, “OpenFAST,” 2024, software, accessed 2024-08-09. [Online]. Available: <https://github.com/OpenFAST/openfast>
- [30] T. J. Larsen and A. M. Hansen, “How 2 HAWC2, the user’s manual,” Risø National Laboratory, Technical Report Risø-R-1597(ver.3-1)(EN), 2007.
- [31] NLR, “FLORIS. version 4.6.6,” *GitHub repository*, 2026. [Online]. Available: <https://github.com/NatLabRockies/floris>
- [32] M. M. Pedersen, A. M. Forsting, P. van der Laan, R. Riva, L. A. A. Romàn, J. C. Risco, M. Friis-Møller, J. Quick, J. P. S. Christiansen, R. V. Rodrigues, B. T. Olsen, and P.-E. Réthoré, “PyWake 2.5.0: An open-source wind farm simulation tool,” 2 2023. [Online]. Available: <https://gitlab.windenergy.dtu.dk/TOPFARM/PyWake>
- [33] J. Jonkman and K. Shaler, “FAST.Farm user’s guide and theory manual,” National Renewable Energy Laboratory, Golden, CO, USA, Technical Report NREL/TP-5000-78485, 2021.
- [34] M. Hall, S. Housner, D. Ogden, D. Zalkind, G. Barter, and P. Bortolotti, “RAFT (Response Amplitudes of Floating Turbines),” National Renewable Energy Laboratory, Golden, CO, USA, Technical Report NREL SWR-22-40, 2022.
- [35] sowento GmbH, “SLOW: Simulation tool for floating offshore wind turbines,” <https://www.sowento.com/sowento-slow/>, 2024, floating offshore wind turbine simulation software.
- [36] F. Lemmer, W. Yu, B. Luhmann, D. Schlipf, and P. W. Cheng, “Multibody modeling for concept-level floating offshore wind turbine design,” *Multibody System Dynamics*, vol. 49, Jun. 2020.
- [37] M. Hall, S. Housner, S. Srinivas, and S. Wilson, “MoorPy: Quasi-Static Mooring Analysis in Python,” 2021, software.
- [38] J. Feng, M. M. Pedersen, R. Riva, H. Bredmose, and P. Santos, “Design optimization of floating offshore wind farms using a steady state movement and flow model,” *Journal of Physics: Conference Series*, vol. 2875, no. 1, p. 012039, Aug. 2024.

- [39] M. Mei, P. Kou, Y. Xu, Z. Zhang, R. Tian, and D. Liang, “Floating Offshore Wind Farm Yaw Control Via Model-based Deep Reinforcement Learning,” in *2025 IEEE Power & Energy Society General Meeting (PESGM)*, Jul. 2025, pp. 1–5.
- [40] N. Tsarknias, M. Friis-Møller, and N.-E. Clausen, “Optimizing the layout of floating wind farms in Crete: A combined LCOE and visual impact minimization,” *Energy*, vol. 337, p. 138645, Nov. 2025.
- [41] A. I. Hietanen, T. H. Snedker, K. Dykes, and I. Bayati, “A novel techno-economical layout optimization tool for floating wind farm design,” *Wind Energy Science*, vol. 9, no. 2, pp. 417–438, Feb. 2024.
- [42] G. Froese, S. Y. Ku, A. C. Kheirabadi, and R. Nagamune, “Optimal layout design of floating offshore wind farms,” *Renewable Energy*, vol. 190, pp. 94–102, May 2022.
- [43] M. Mahfouz, E. Lozon, M. Hall, and P. Cheng, “Dynamic Performance of a Passively Self-Adjusting Floating Wind Farm Layout to Increase the Annual Energy Production,” *Wind Energy Science*, vol. 9, no. 7, pp. 1595–1615, 2024.
- [44] Y. Niu and R. Nagamune, “Integrated Control of Floating Offshore Wind Farms with Reconfigurable Layouts,” *Wind Energy Science Discussions*, pp. 1–23, Nov. 2025.
- [45] J. A. Frederik, B. M. Doekemeijer, S. P. Mulders, and J.-W. van Wingerden, “The Helix approach: Using dynamic individual pitch control to enhance wake mixing in wind farms,” *Wind Energy*, vol. 23, no. 8, pp. 1739–1751, 2020.
- [46] E. L. Hodgson, M. Baricchio, M. B. Nilsen, N. Troldborg, and S. J. Andersen, “Wake steering and individual turbine yaw control in large wind farms with complex inflows,” *Journal of Physics: Conference Series*, vol. 3224, no. 3, p. 032013, May 2026.
- [47] T. Zhang, X. Xu, S. Wang, Y. Xing, P. Dou, R. Ji, and P. Yang, “Investigation of wake steering control effects on the dynamic responses of 15 MW semi-submersible floating wind farms,” *Renewable Energy*, vol. 254, p. 123704, Dec. 2025.
- [48] M. Baricchio, D. van der Hoek, T. Dammann, P. M. O. Gebraad, J. Iori, and J.-W. van Wingerden, “Combining wake steering and active wake mixing on a large-scale wind farm,” *Wind Energy Science Discussions*, pp. 1–36, Dec. 2025.
- [49] A. P. J. Stanley, T. Mulder, B. Doekemeijer, and J. Kreeft, “Machine learning to rapidly predict turbine yaw angles for wake steering,” *Journal of Physics: Conference Series*, vol. 2767, no. 8, p. 082011, Mar. 2024.
- [50] S. Tao, J. Yang, G. Zheng, R. He, and A. E. Feijóo-Lorenzo, “A two-stage optimization model for floating offshore wind farm layout and control,” *Renewable Energy*, vol. 253, p. 123614, Nov. 2025.
- [51] D. G. Dirik, J. Quick, P.-E. Réthoré, J. P. Murcia Leon, and R. Riva, “Revenue-Focused Wind Farm Control Co-Design for Future Electricity Markets Scenarios,” *Journal of Physics: Conference Series*, vol. 3016, no. 1, p. 012024, Feb. 2025.
- [52] T. Chen and C. Guestrin, “XGBoost: A Scalable Tree Boosting System,” *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–794, 2016.

- [53] E. Lozon, M. R. Lekkala, L. Sirkis, and M. Hall, “Reference mooring and dynamic cable designs for representative U.S. floating wind farms,” *Ocean Engineering*, vol. 322, p. 120473, Apr. 2025.
- [54] P. A. Fleming, A. P. J. Stanley, C. J. Bay, J. King, E. Simley, B. M. Doekemeijer, and R. Mudafort, “Serial-Refine Method for Fast Wake-Steering Yaw Optimization,” *Journal of Physics: Conference Series*, vol. 2265, no. 3, p. 032109, Feb. 2022.
- [55] Delft High Performance Computing Centre (DHPC), “DelftBlue Supercomputer (Phase 2),” <https://www.tudelft.nl/dhpc/ark:/44463/DelftBluePhase2>, 2024.
- [56] M. Baricchio, J. Iori, T. Lokken, P. Gebraad, and J.-W. Wingerden, “A novel metric to compare wind farm layouts and identify trends in the layout optimization problem,” *Journal of Physics: Conference Series*, vol. 3224, p. 032020, May 2026.
- [57] “SUDOCO – Sustainable Resilient Data-enabled Offshore Wind Farm and Control CO-design.” [Online]. Available: <https://sudoco.eu/>
- [58] Hellenic Hydrocarbons and Energy Resources Management Company (HEREMA), “Draft National Offshore Wind Farm Development Programme (NDP-OWF),” Athens, Greece, National Programme (draft), Sep. 2023, in Greek. [Online]. Available: [herema.gr](http://herema.gr) (NDP-OWF).
- [59] Hellenic Hydrocarbons and Energy Resources Management Company (HEREMA) , “Strategic Environmental Impact Assessment (SEIA) of the Draft National Offshore Wind Farm Development Programme,” Athens, Greece, Strategic Environmental Impact Assessment, Sep. 2023, annex 15.3, Map 5: distribution of striped dolphin, Cuvier’s beaked whale and Risso’s dolphin across the Crete candidate areas. Prepared by LDK Consultants and NCC. In Greek. [Online]. Available: [herema.gr](http://herema.gr) (Annex 15.3).
- [60] “POSEIDON forecasting system,” <https://poseidon.hcmr.gr/services/ocean-data/situ-data>.
- [61] B. Allen, L. F. Cameron, T. R. O. Wainwright, and D. J. Poole, “An Initial Study of Multimodality in Wind Farm Layout Optimization Problems.”
- [62] D. Poole, “Characterization of Multimodality in Wind Farm Layout Optimization,” *Energy Science & Engineering*, vol. 14, pp. 737–751, Nov. 2025.
- [63] M. Baricchio, D. van der Hoek, T. Dammann, J. Iori, P. Gebraad, and J.-W. Wingerden, “An approximation of the optimal combined Helix and yaw control for wind farm co-design applications,” *Journal of Physics: Conference Series*, vol. 3224, p. 032019, May 2026.