



Research paper

Uncertainty-aware Bayesian optimisation framework for the design of floating offshore wind turbines

Ander Aristondo ^{a,b,*}, Vincenzo Nava ^c, Markel Penalba ^{b,d}, Miguel Esteras ^a

^a TECNALIA, Basque Research and Technology Alliance (BRTA), Astondo Bidea, Edificio 700, Derio, 48160, Spain

^b Fluid Mechanics Department, Mondragon University, Loramendi 4, Arrasate, 20500, Spain

^c Politecnico di Torino, Department DIATI, Via Duca degli Abruzzi 24, Turin, 10129, Italy

^d Ikerbasque, Basque Foundation for Science, Euskadi Plaza 5, Bilbao, 48011, Spain

ARTICLE INFO

Keywords:

Floating offshore wind
Design optimisation
Mooring
Bayesian
Uncertainty

ABSTRACT

Floating wind platforms are complex systems requiring integrated optimization of interdependent subsystems under highly dynamic conditions. Conventional design approaches often treat the process as a black-box optimization problem, minimizing a cost metric based on simplified assumptions. However, these assumptions evolve during project development, and inaccuracies in early-stage models can significantly affect outcomes. This study introduces an uncertainty-aware design paradigm in which uncertainty is not a byproduct but a central element guiding the optimization process and informing decision-making. The proposed framework embeds Bayesian Optimization within the design workflow, leveraging Gaussian Process regression to model both the objective function and its associated uncertainty. This probabilistic representation enables explicit integration of uncertainty into the optimization strategy, while the Lower Confidence Bound acquisition function balances exploration and exploitation according to user-defined preferences. The approach reduces the number of objective function evaluations and provides valuable insights into cost distributions. Its effectiveness is demonstrated through three case studies of increasing complexity: a multi-modal objective function, the mass optimization of a catenary mooring system, and the same system with penalization terms. Compared to a Genetic Algorithm applied to an equivalent problem, the proposed method achieves convergence with approximately 80% fewer evaluations.

1. Introduction

The global response to mitigate the effects of climate change is pushing the need of developing cost-effective forms of renewable energy. According to IEA (2022), wind energy is expected to be one of the major low-emission energy sources over the period to 2050. Today, onshore installations represent the majority of the installed capacity. However, offshore wind energy is expected to increase global installed capacity from 67.4 GW at the end of 2023 (WFO, 2024) to 2465 GW by 2050 (IRENA and GWEC, 2023). If compared to their onshore counterparts, the development of offshore wind energy entails substantially greater financial commitments. The most substantial cost factors encompass the manufacturing of the platform and the wind turbines, along with the subsequent installation (Nava et al., 2019). The installation phase frequently demands the utilisation of large laying vessels, the placement of anchors, moorings, and the installation of electrical cables and substations. Additionally, although ongoing monitoring and maintenance activities

need to be more frequent, these activities are more complex due to the harsh maritime environment. Consequently, operation and maintenance (O&M) of offshore wind power can incur in expenses about 25–30% of the total cost for bottom fixed turbines and can increase up to a 40% for floating turbines (Centeno-Telleria et al., 2024).

In addition, most of the installed offshore wind capacity is based on bottom-fixed structures (IRENA, 2024), featuring water depths of up to 50–60 m. However, almost 80% of the global wind resource potential is located in water depths beyond 60 m, where floating support platforms are the most cost-effective solution (Pollini et al., 2023). However, the economic feasibility of floating offshore wind turbines (FOWTs) poses a great challenge. Owing to the technology's immaturity, the Levelised Cost of Energy (LCoE) of FOWTs is significantly higher than for bottom-fixed structures (Díaz and Guedes Soares, 2023). Therefore, continuous improvement in the design, manufacturing, installation, and O&M of FOWTs is key to significantly reducing costs.

* Corresponding author.

E-mail address: ander.aristondo@tecnalia.com (A. Aristondo).

<https://doi.org/10.1016/j.oceaneng.2025.123080>

Received 4 April 2025; Received in revised form 19 September 2025; Accepted 3 October 2025

Available online 16 October 2025

0029-8018/© 2025 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1.1. Literature review of FOWT design optimisation approaches

The design procedure of a FOWT is a multidisciplinary problem involving simultaneous optimisation of all its subsystems (Leimeister et al., 2022). The complete system is reduced to a set of relevant parameters that govern the design. Then, an optimiser is used to explore the design space and select the combination of parameters that minimise a pre-defined design objective (Sykes et al., 2023). The FOWT design standards (DNV, 2021) propose a set of Design Load Cases (DLCs) that need to be addressed to evaluate the integrity, stability and overall performance of each design candidate. This requires numerical simulations of the system response under a set of environmental conditions, including operational, extreme and accidental scenarios and operational states. The number of simulations is further increased to account for the randomness of the excitation. The computational burden and the uncertainty associated with these numerical simulations are significant, which make the optimisation process particularly challenging. Therefore, efficient strategies for exploring the design space while considering uncertainty are key to enabling practical and robust optimisation workflows.

Numerous optimisation frameworks have been proposed in the literature, each aiming to improve the performance, reliability and cost-effectiveness of FOWT systems (Patryniak et al., 2022). These frameworks vary in scope, with some focusing on individual subsystems while others adopt a more integrated approach, simultaneously optimising multiple coupled subsystems. A key distinction among these optimisation strategies lies in the type of algorithm employed. Broadly, the literature can be categorised into two main groups: gradient-based and gradient-free methods.

Gradient-based optimisation methods depend on the availability of analytical or numerically approximated gradients of the objective function with respect to the design variables (Rohrer et al., 2023). These approaches are typically computationally efficient and exhibit rapid convergence. However, they often require simplified physical models to enable gradient computation, which can compromise the fidelity of the results. In the context of FOWT design optimisation, gradient-based methods are commonly implemented using linearised models in the frequency domain (FD). The FD formulation facilitates the derivation of analytical gradients, which are subsequently utilised by the optimiser.

Representative applications of FD-based models in FOWT design optimisation include the work of Pollini et al. (2023) and Dou et al. (2020), who employed the QuLAF model (Pegalajar-Jurado et al., 2018). The QuLAF model simplifies the FOWT dynamics to four degrees of freedom (DOFs), representing the mooring system as a linearised stiffness matrix and inferring controller dynamics from a step-wind response obtained via a time-domain simulation. Another example is provided by Hegseth et al. (2020), who developed a reduced-order FD model that integrates the floater, mooring system, tower and blade-pitch controller into a unified optimisation framework involving over 20 design parameters. While their representation of the mooring system and control dynamics is more sophisticated than that of QuLAF, it remains a linearised FD model to preserve the tractability of analytical gradient computation. A gradient-based optimisation strategy was also adopted by Fylling and Berthelsen (2011) to concurrently optimise the design of a spar-type floater and its mooring system. By employing numerical gradients, the authors were able to incorporate time-domain simulations into the objective function, thereby capturing more detailed dynamic responses. Nevertheless, the use of numerical differentiation requires multiple evaluations of the objective function per design point, resulting in a computational cost that increases rapidly with the number of design variables.

Conversely, gradient-free methods do not require gradient information and are therefore more flexible in handling complex, non-linear and multi-modal design spaces. Among these, evolutionary algorithms have gained significant traction within the FOWT design community and are now widely adopted for optimisation tasks (Sykes et al., 2023). Nevertheless, a key limitation of such methods is their high demand of

objective function evaluations to adequately explore the design space and achieve convergence. Given the substantial computational cost associated with these evaluations, gradient-free methods often incorporate supplementary strategies to enhance computational efficiency. These include the application of meta-heuristics to improve optimiser performance, the use of reduced-order models to speed-up dynamic simulations, and the deployment of surrogate models to approximate the objective function.

Meta-heuristic enhanced optimisation was used in López-Queija et al. (2024), who employed a Genetic Algorithm (GA) to simultaneously optimise the floater geometry and turbine controller parameters. It combined the use of the OPTECOT (Echevarrieta et al., 2024) meta-heuristic with a reduced-order time-domain model to accelerate objective function evaluations. However, this comes at the cost of reducing model fidelity. The works by Du et al. (2024) and Faraggiana et al. (2022) employed surrogate-assisted GA optimisation to optimise a semi-sub and spar-buoy geometries, respectively. The main advantage of surrogate-assisted GA is that the GA builds each subsequent population informed by the surrogate's prediction, offering superior performance when compared to conventional GAs. However, all individuals of each population must be evaluated using the actual (and computationally expensive) objective function, therefore requiring an elevated number of function evaluations. The studies conducted by Lemmer et al. (2020) and Ojo et al. (2025) employed a Pattern Search (PS) algorithm to optimise the hull geometry of FOWTs. In addition, Lemmer et al. (2020) incorporated the turbine controller parameters as part of the optimisation variables. PS is a derivative-free optimisation method that generally requires fewer function evaluations compared to GA. However, due to the high computational cost associated with evaluating the objective function, both studies utilised reduced-order models to represent the FOWT dynamics and thereby accelerate the optimisation process.

1.2. Motivation and proposed approach

Conventional optimisation frameworks applied to FOWT design are formulated using a deterministic approach. They generally struggle to account for the uncertainties present in the FOWT design process. These uncertainties arise from a variety of sources, including the stochastic nature of environmental loading, simplifications inherent in numerical modelling, variability in material properties, and fluctuations in supply chain logistics and material costs. Moreover, traditional optimisation methods typically yield a single optimal solution or a Pareto front in the case of multi-objective formulations, with no additional information on how the cost varies with respect to design parameters. However, due to the evolving nature of design assumptions—particularly those made during early-stage development—these solutions may become sub-optimal as the design progresses and more detailed information becomes available.

To address the limitations outlined above, this study introduces a novel FOWT design optimisation framework based on Bayesian Optimisation (BO). BO is a gradient-free, surrogate-based, global optimisation procedure, especially suitable for cases in which the evaluation of the objective function is computationally expensive (Brochu et al., 2010). The optimisation process employs a probabilistic surrogate model to represent the objective function, enabling the application of Bayesian inference at each iteration to determine the next sampling point (Di Fiore et al., 2024). The probabilistic formulation of both the surrogate model and the acquisition strategy facilitates the explicit incorporation of uncertainty into the optimisation process. Furthermore, this approach supports the development of efficient sampling strategies that balance exploitation—refining predictions near potential optima—and exploration—improving model accuracy in regions of high uncertainty.

Numerous studies have demonstrated the application of Bayesian optimisation (BO) to a wide range of engineering design problems (Di Fiore et al., 2024). Within the offshore engineering domain, BO has been

utilised to address various optimisation challenges, including the design of mooring systems for floating platforms employed in the oil and gas sector (Yin et al., 2025; Lim et al., 2023). Additionally, BO has been applied to optimise blade-pitch control strategies for offshore wind turbines (Fu et al., 2025; Yuan et al., 2024; Mulders et al., 2020), to refine hydrofoil geometries (Rafiee et al., 2022), and to assess the fatigue performance of mooring lines under uncertain loading conditions (Cousin et al., 2024; Rafiee and Repalle, 2024).

The optimisation framework proposed in this study distinguishes itself from existing approaches by not only delivering an optimised design solution but also providing a probabilistic characterisation of the cost function distribution. The majority of BO applications in offshore engineering rely on the Expected Improvement (EI) acquisition function (Jones et al., 1998), which aims to minimise the number of iterations required to locate the global optimum. EI achieves this by implicitly balancing exploration of the design space and exploitation of currently promising regions. In contrast, the present framework employs the Lower Confidence Bound (LCB) acquisition function (Cox and John, 1992), which enables explicit control over the exploration/exploitation trade-off. This flexibility allows the designer to tailor the sampling strategy to specific design requirements and to dynamically adjust the exploration/exploitation trade-off during the optimisation process.

This study presents a novel FOWT design optimisation framework that leverages BO to simultaneously identify optimal designs and provide a probabilistic characterisation of the objective function. We define uncertainty-aware design optimisation as a design paradigm in which uncertainty is not only quantified but actively integrated into the optimisation process—both in the selection of new design points (via acquisition functions that balance exploration and exploitation) and in the interpretation of the final design outcome. This goes beyond conventional design optimisation by explicitly acknowledging and managing uncertainty throughout the process. This study contains a proof-of-concept of the framework, in which the uncertainty terms included within the process are the inherent (aleatory) uncertainty of the objective function evaluations and the predictive uncertainty of the surrogate model in under-sampled regions of the design space. Further, even if in the current implementation the propagation of the uncertainty through input variables or constraints is not included, the probabilistic structure of the framework is designed to integrate directly such an extension. Hence, in our framework the uncertainty becomes a core element in the design strategy, and this expands the scope of standard BO applications, which often treat uncertainty as a byproduct of the surrogate model.

A preliminary version of the proposed framework was introduced by the authors in an earlier study (Aristondo et al., 2025), where it was applied to the design of a simplified spar-type buoy mooring system. However, that initial implementation was limited to deterministic objective functions, treating their outputs as ground truth during the surrogate model fitting process. The present study extends this framework to accommodate objective functions with inherent uncertainty, incorporating the estimation of such uncertainty directly within the surrogate model fitting process. Furthermore, the earlier approach (Aristondo et al., 2025) addressed inequality constraints by assigning prohibitively high penalty values to infeasible design points. In contrast, the current implementation adopts a more general treatment of nonlinear constraints, which avoids introducing discontinuities or sharp gradients in the objective function, thereby preserving its smoothness and continuity.

To illustrate the methodology, three case studies are presented: (i) a closed-form multi-modal benchmark function, (ii) the optimisation of mooring system mass without penalty terms, and (iii) the same optimisation problem incorporating penalty terms. The uncertainty considered in this study focuses exclusively on the surrogate model's predictive uncertainty and does not account for uncertainty in the input variables, which is left for future work.

The paper has been structured in the following sections. Section 2 provides a formal definition of the design optimisation problem, de-

scribes the optimisation framework proposed in this study and provides detail on the implementation of the BO algorithm. Section 3 describes the case studies selected to demonstrate the methodology, providing details on the different objective functions and constraints considered in each case. Finally, Sections 4 and 5 contain the results of the optimisation problem and the main conclusions drawn from this work.

2. The optimisation problem

Our proposed optimisation framework is designed to provide a vector of optimised design parameters together with valuable insight about the actual distribution of the objective function across the design space. The objective function is assumed to be of an a-priori unknown shape, but can be evaluated at any point within the design space. Each evaluation is computationally expensive as it usually requires to run a high number of simulations. The evaluated samples are used to train a probabilistic surrogate model that estimates the objective function, which is also able to estimate the uncertainty of its own predictions. The trained surrogate model is then used to query the next sample to be evaluated until the stopping criterion is reached. A visual representation of the proposed framework is given in Fig. 1.

2.1. Problem formulation

The design optimisation of a floating wind turbine is a constrained nonlinear optimisation problem. Generally, it can be formulated as in Eq. (1).

$$\begin{aligned} & \min_{\mathbf{x}} f(\mathbf{x}) \\ \text{subject to : } & g_j(\mathbf{x}) \leq 0 \quad 0 \leq j \leq m \\ & l_i \leq x_i \leq u_i \quad 1 \leq i \leq n \end{aligned} \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^n$ is a vector containing all the design variables, $f(\mathbf{x})$ the objective function to be minimised and $g(\mathbf{x})$ the inequality constraints that must be satisfied. These inequality constraints define a feasible region F within the search space S , such that $F \subseteq S$. The search space S is defined by the lower (l_i) and upper (u_i) bounds for each variable in \mathbf{x} .

2.2. Optimisation framework

2.2.1. Bayesian optimisation

Bayesian optimisation is a global optimisation procedure, whose main goal is to approximate the real (unknown) distribution of an objective function $f(\mathbf{x})$, which is assumed to be defined on some bounded set S . BO creates a probabilistic model of $f(\mathbf{x})$ to guide decisions on where to further evaluate $f(\mathbf{x})$ within S . This also allows to effectively account for uncertainty and incorporate it into the optimisation problem (Snoek et al., 2012).

The central idea behind BO is to build a model that can be updated and queried to drive optimisation decisions (Shahriari et al., 2016). Let C be the value of $f(\mathbf{x})$, evaluated at any arbitrary point, and let D be the available data, representing the values of $f(\mathbf{x})$ that have already been evaluated. Since C has not been observed, it is treated as a random variable with a *prior* distribution $p(C)$ that captures the *prior* beliefs about the value of C before evaluating $f(\mathbf{x})$. Given the known data D and a *likelihood* model $p(D|C)$, a *posterior* distribution $p(C|D)$ can be inferred applying Bayes' rule shown in Eq. (2). This *posterior* distribution represents the updated beliefs about C after observing the data D . The value in the denominator, $p(D)$, is known as the marginal likelihood or evidence. The probabilistic model $p(C|D)$ is usually referred to as a surrogate model of $f(\mathbf{x})$.

$$p(C|D) = \frac{p(D|C) \cdot p(C)}{p(D)} \quad (2)$$

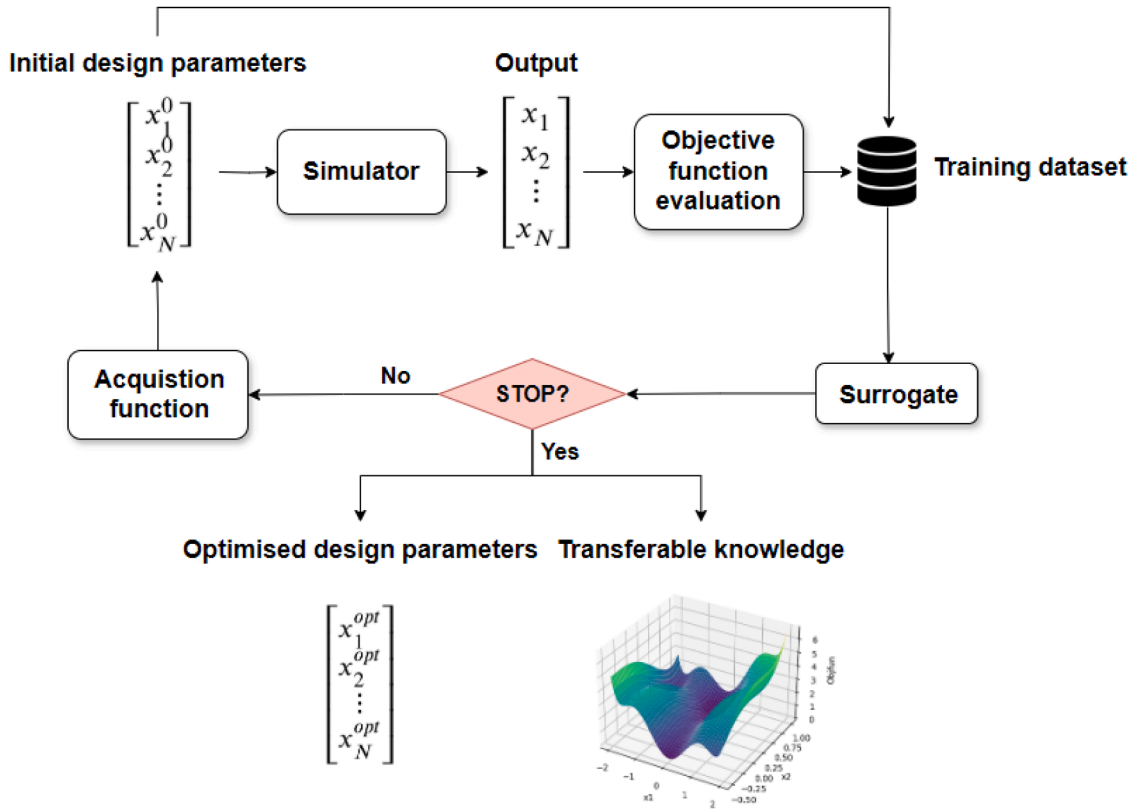


Fig. 1. Visual representation of the optimisation framework.

2.2.2. Surrogate model

The proposed BO framework utilises Gaussian Process (Rasmussen and Williams, 2008) (GP) regression as a surrogate model, which is widely adopted for BO due to its flexibility and capacity to quantify predictive uncertainty (Mulders et al., 2020). The objective function is assumed to be a realisation of a GP with mean function $\mu(\mathbf{x})$ and a covariance kernel matrix $k(\mathbf{x}, \mathbf{x}')$, as outlined in Eq. (3).

$$f(\mathbf{x}) \sim GP(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')) \quad (3)$$

where the mean function and the covariance matrix represent the expected value of the objective function and the variance of the prediction, as in Eq. (4).

$$\begin{aligned} \mu(\mathbf{x}) &= \mathbb{E}[f(\mathbf{x})] \\ k(\mathbf{x}, \mathbf{x}') &= \mathbb{E}[(f(\mathbf{x}) - \mu(\mathbf{x}))(f(\mathbf{x}') - \mu(\mathbf{x}'))] \end{aligned} \quad (4)$$

The covariance matrix $k(\mathbf{x}, \mathbf{x}')$ is computed using a set of kernel functions that encode the *prior* belief about the nature of $f(\mathbf{x})$. In a design problem, small variations of the input parameters are expected to yield similar values of $f(\mathbf{x})$. Therefore, $f(\mathbf{x})$ is expected to be continuous and smooth. This behaviour is encoded by using the squared exponential kernel, also known as RBF, which is defined in Eq. (5).

$$k(\mathbf{x}, \mathbf{x}') = a \cdot \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2l^2}\right) \quad (5)$$

where $\|\mathbf{x} - \mathbf{x}'\|$ is a vector of Euclidean distances, a is a scaling constant and l the length scale, which are computed as part of the fitting process of the GP to the available observations. Eq. (5) is infinitely differentiable and results in GPs that are smooth (Rasmussen and Williams, 2008).

2.2.3. Acquisition function

Once the surrogate model has been defined, the second key feature in BO is the acquisition function. Essentially, the acquisition function is used to determine the sequence of evaluations of $f(\mathbf{x})$. As each evaluation $f(\mathbf{x})$ is computationally expensive, the acquisition function must

determine the utility of candidate points for the next evaluation of $f(\mathbf{x})$, and be used to select the point that offers the greatest benefit. Usually, the acquisition function is carefully designed to trade off exploration of the search space and exploitation of current promising areas Shahriari et al. (2016).

Algorithm 1 Bayesian optimisation algorithm.

```

while  $\|\mathbf{x}_{n+1} - \mathbf{x}_n\| \geq 10^{-3}$  do
  select new  $\mathbf{x}_{n+1} = \arg \min LCB(\mathbf{x})$ 
  evaluate  $f_{cost}(\mathbf{x}_{n+1})$  to obtain  $y_{n+1}$ 
  augment data  $D_{n+1} = \{D_n, (\mathbf{x}_{n+1}, y_{n+1})\}$ 
  update surrogate model
end while

```

One of the most popular acquisition functions is the Expected Improvement (Jones et al., 1998) (EI), which intrinsically trades-off exploration and exploitation. The EI is designed for efficient global optimisation and is focused on solving the optimisation problem using the minimum number of evaluations. However, the optimisation framework presented in this study focuses on generating a surrogate model that is accurate over the whole design space, not only in optimising $f(\mathbf{x})$. Therefore, the Lower Confidence Bound (LCB) (Cox and John, 1992) has been used as the acquisition function (see Eq. (6)). The LCB function guides the search by minimising the cumulative regret bound (Zhang et al., 2022). The choice of LCB is based on the fact that it allows the designer to explicitly tune the exploration/exploitation ratio via a weight hyperparameter affecting the uncertainty term in Eq. (6). Minimising the prediction of the Eq. (3) has the benefit to exploit around the current best solution, while maximising the prediction's uncertainty has the benefit of exploring other under-sampled areas (Zhan and Xing, 2020). LCB is defined as

$$LCB(\mathbf{x}) = \mu(\mathbf{x}) - \beta \cdot k(\mathbf{x}, \mathbf{x}') \quad (6)$$

where $\beta \geq 0$ is a hyperparameter that allows to deal with the exploration/exploitation trade-off explicitly. This hyperparameter acts as a weight that modifies the relevance of $k(\mathbf{x}, \mathbf{x}')$ in Eq. (6). If $\beta \rightarrow \infty$ then the process explores the areas of higher uncertainty, completely neglecting the current promising areas. On the contrary, if $\beta \rightarrow 0$, the process exploits around the most promising areas, ignoring the high uncertainty regions. The minimisation of $LCB(\mathbf{x})$ as per Eq. (6) has been carried out using the COBYLA (Powell, 1994) algorithm with 20 random initializations. In this study, the BO process begins with a fixed value of β , which is then progressively reduced by a certain amount $\xi \in (0, 1]$ at each iteration, following Eq. (7).

$$\beta_{i+1} = \xi \cdot \beta_i \quad (7)$$

To run the BO process as described in Algorithm 1, an initial set of evaluations of the objective function is required in the first iteration of the loop. In this work, this initial grid search has been performed using equally spaced samples (see Section 3.5).

2.3. Inequality constraints

The optimisation of $f(\mathbf{x})$ must be restricted to solutions that are not only technically viable, but also that fulfil appropriate design criteria. The implementation of inequality constraints within the BO process is not straightforward, especially if the value of the constraint functions $g(\mathbf{x})$ cannot be computed before evaluating the objective function. In this study, a methodology similar to the one developed by Gardner et al. (2014) has been implemented, which is based on the concept of constrained acquisition function. At each iteration, each constraint function $g(\mathbf{x})$ is modeled using a GP trained on all previously evaluated samples. These surrogate models are subsequently employed as nonlinear constraints in the acquisition function optimisation, which determines the location of the next evaluation point, as outlined in Algorithm 1.

3. Case studies

The case study selected to demonstrate the optimisation methodology is the design optimisation of the mooring system of the public definition of the VoltornUS (Allen et al., 2020) semi-submersible platform, featuring the IEA 15MW Reference Wind Turbine (RWT) (Gaertner et al., 2020). To do so, three different objective functions are proposed. The first objective function only depends on the initial state of the mooring system and, therefore, can be evaluated analytically with no need of running any simulation. The intention of this first case is to serve as a verification of the implementation of the optimisation framework, as the shape and minima of such a function are known *a priori*. The second and third objective functions represent the mass of the mooring system. The third case includes a set of increasing penalisation terms closer to the boundaries of acceptable solutions to demonstrate the capabilities of the BO framework to solve more realistic problems with increasing complexity.

3.1. Numerical model

The numerical model representing the coupled aero-hydro-servo-elastic dynamics of the floating wind turbine has been implemented using OpenFAST v3.5.1 (NREL, 2024), with the publicly available data in the IEA 15MW RWT Github repository (IEA, 2024). OpenFAST offers a variety of modules to represent the dynamics and loading of different subsystems. In this study, ElastoDyn has been used to define the global mass and inertial properties of the coupled system, AeroDyn to compute the aerodynamic loading on both the rotor and the tower, HydroDyn for the hydrodynamic loading on the floating platform and MoorDyn (Hall, 2020) to represent the mooring system dynamics.

The aerodynamic loading on the rotor has been computed using AeroDyn's built-in Blade Element Moment code, which includes a series of empirical corrections for tip and hub losses and for skewed wake.




Fig. 2. Case study schematic and design variable bounds.

The hydrodynamic loading on the platform is computed using a hybrid model that includes potential flow matrices for radiation and diffraction (1st and 2nd order) wave loading, as well as the viscous term on Morison's (Morison et al., 1953) equation for viscous drag loading. The 2nd order wave loading has been computed using the full difference-frequency Quadratic Transfer Function.

3.2. Design variables

The mooring system considered in this study features three simple catenary lines of equal length, evenly distributed around the platform with an angle of 120 degrees. Under these conditions, the whole mooring system can be parametrised using only two design variables: the mass per unit length of each line m_L and the suspended length of the lines at rest L_{S0} . As the lines are assumed to be equal, having the platform in its undisplaced position and defining m_L and L_{S0} is enough to define the mooring system configuration. The lines have been modelled as almost infinitely long, and only the suspended part of the line has been considered when computing the system cost (see Section 3.5). More precisely, the laying length of the lines at rest has been defined as $L_{lay} = 10 \cdot L_{S0}$. The bounding values considered for the design variables can be found in Fig. 2.

3.3. Design load cases

3.3.1. Site selection

The metocean conditions used for the design correspond to the Biscay Marine Energy Platform (BiMEP), which is a test site for offshore renewable energy devices off the coast of Biscay, in the Cantabric sea. In BiMEP, the wave and current resources present a very strong directionality, North-West being the main direction for incoming waves (see Fig. 3a), and both East and West being the main heading directions for current (see Fig. 3b). Wind presents a less pronounced directionality, with strong West, North-North-West, South, South-South-West, East and North-North-East contributions (see Fig. 3c). The selection of extreme conditions has been carried out via the Inverse First Order Method (IFORM) contours available at IHCantabria (2024). The IFORM contours have been shown to perform well in the prediction of extreme sea states (Hauteclocque et al., 2022; Penalba et al., 2023), and are the approach recommended by design standards (DNV, 2021). The test site is divided into 4 berths, with depths ranging from 50 to 90m (IHCantabria, 2024).

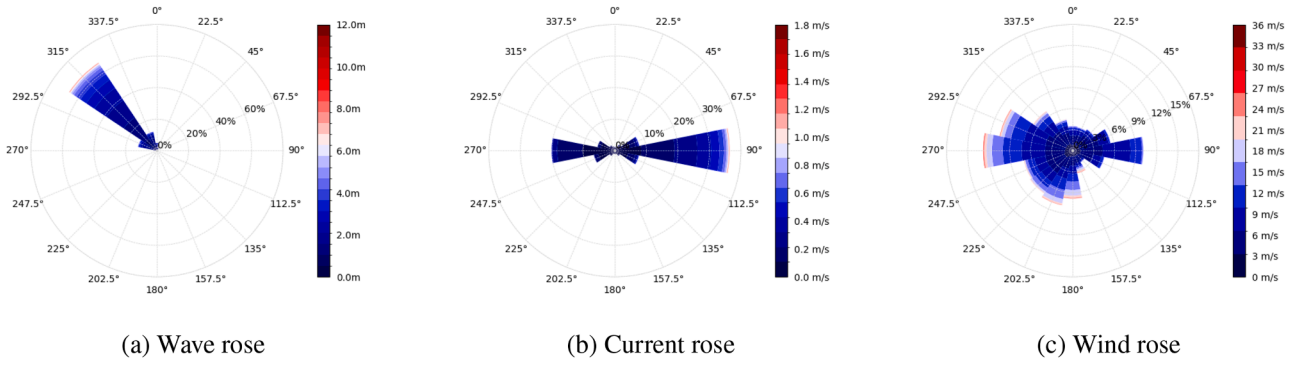


Fig. 3. Wave, current and wind roses at BiMEP (source IH Cantabria, 2024).

3.3.2. Load case selection

The selection of load cases has been carried out according to the DLC list required by the design standards defined by certification agencies such as DNV (DNV, 2021, 2016). This study focuses on the ultimate limit state (ULS) loading on mooring system. Fatigue Limit State (FLS) is as critical as ULS in mooring design. However, in this study we have limited our development to a proof-of-concept of our methodological approach rather than a full design of the FOWT, hence focusing only on ULS but without lack of generality. Indeed, the proposed approach constitutes a methodological framework fully applicable to FLS, if an appropriate selection of sea states is considered, and the cost function accounts for damage equivalent load and the associated constraints and boundary conditions.

Consequently, the complete set of DLCs has been whittled down to DLC1.6 and DLC6.1, following the guidance from DNV (2021). On the one hand, DLC1.6 corresponds to severe sea state conditions combined with rated or near-rated wind speeds, which result in the highest thrust forces being exerted on the rotor. On the other hand, DLC6.1 corresponds to extreme sea state conditions with extreme wind speed, with the wind turbine being either parked or idling. These two load cases result in the most demanding loading conditions for the mooring system, which was also observed in the analysis carried out for the original design (Allen et al., 2020).

About directionality, DLC1.6 imposes the assumption of co-directional wind, wave and current loading, whereas in DLC6.1 the effect of the combination of different directions should be analysed. To keep the number of simulations per evaluation from growing prohibitively, co-directionality has been assumed in both load cases in this study. Note that regarding the loads on the mooring system this assumption is conservative, as it leads to the highest displacement of the platform and, thus, the highest loading. In both cases, the rotor is assumed to be perfectly aligned with the incoming wind direction. A single incoming wind, wave and current direction has been considered per DLC: aligned with a mooring line, resulting in the maximum offset. A summary of the considered load cases can be found in Table 1.

The computational cost associated with the evaluation of each design candidate is directly proportional to the number of simulations considered per evaluation. The simulations were executed on a workstation equipped with a 13th Gen Intel® Core™ i9-13900K processor, utilizing 12 processing cores. Each individual simulation required approximately 1.5 h to complete, resulting in a total computational time of roughly 3 h per design candidate.

3.4. Constraints

In mooring design problems these criteria usually comprise a limit to the maximum load that a single mooring line is able to withstand and the maximum allowable excursion of the floating platform. Maintaining mooring line tension below the minimum breaking load guarantees the structural integrity of the mooring system and eventually the survivabil-

Table 1

Summary of DLCs considered for the evaluation of $f(\mathbf{x})$.

		DLC1.6	DLC6.1
Wind speed	[m s ⁻¹]	10.59	30
H_s	[m]	7.98	10.5
T_p	[s]	(10.59, 20.28)	(14.67, 18.12)
Current speed	[m s ⁻¹]	0.72	1.3
Direction	[deg]	0	0
Duration of seed	[s]	10,800	10,800
# of seeds	[-]	6	6

ity of the overall platform by preventing the failure of the lines due to sudden or extreme loads. The superimposition of a constrained excursion is essential for guaranteeing the overall functionality and safety of the FOWT. Indeed, excessive platform displacement can lead to damage in other components, such as the dynamic export cable, or even result in collisions with nearby structures or vessels.

3.4.1. Maximum tension

The guidelines of the floating wind design standards, such as DNV (2021), limit the maximum allowable value of the characteristic design load (T_d) of a mooring line. This constraint can be formally defined as outlined in Eq. (8).

$$g_1(\mathbf{x}) = T_d(\mathbf{x}) - C_s \cdot MBL(\mathbf{x}) \leq 0 \quad (8)$$

where MBL is the Minimum Breaking Load of the mooring line and $C_s = 0.95$ is a safety factor. Following the guidance from DNV (2021), T_d has been computed as in Eq. (9).

$$T_d = 1.3 \cdot T_{MPM} + 1.75 \cdot T_{mean} \quad (9)$$

where T_{MPM} is the most probable maximum tension and T_{mean} the mean tension on the line. To compute T_{MPM} , six three-hour simulations have been performed with different seeds. The peak values found in those six simulations have been adjusted to a Gumbel distribution, as required by design standards (DNV, 2021). After that, T_{MPM} has been computed as shown in Eq. (10).

$$T_{MPM} = \mu - 0.45 \cdot \sigma \quad (10)$$

where μ and σ are the mean and standard deviation of the Gumbel distribution. Following the guidance from Orcina (2023), the Minimum Breaking Load has been computed as $MBL = 1.96e7 \cdot D_L^2 \cdot (44 - 80D_L)$, which corresponds to grade 3 type of chains. From the same guide, $D_L = \sqrt{\frac{m_L}{21900}}$.

3.4.2. Maximum offset

The total platform offset is defined as the total horizontal displacement of the platform from its initial position. The maximum offset of the platform is limited to a certain value, generally to avoid damaging

some other subsystems such as the dynamic export cable. The maximum offset constraint is formulated as in Eq. (11).

$$g_2(\mathbf{x}) = \Delta(\mathbf{x}) - \Delta_{\max} \leq 0 \quad (11)$$

where $\Delta(\mathbf{x})$ is the maximum platform offset and Δ_{\max} is the maximum allowed offset.

3.5. Objective function

3.5.1. Validation based on a multi-modal objective function

As first objective function, we assumed a multi-modal two-variable function in Eq. (12) in order to validate the implementation and applicability of our methodology.

$$\begin{aligned} \min_{\mathbf{x}=(m_L, L_{S0})} \quad & f_1(\mathbf{x}) = 3 \cdot m_L \cdot L_{S0} + 10^5 \cdot [\sin(0.05 \cdot m_L) + \sin(0.075 \cdot L_{S0})] \\ \text{subject to:} \quad & 450 \leq m_L \leq 650 \\ & 120 \leq L_{S0} \leq 180 \end{aligned} \quad (12)$$

Eq. (12) represents a bounded, unconstrained minimisation problem which has a known minimum at $m_L = 469.51 \text{ kg/m}$, $L_{S0} = 144.10 \text{ m}$. Note that, although related to the mooring system properties, the value of $f_1(\mathbf{x})$ does not depend on any simulation result, but only on the properties of the mooring system at rest. Therefore, it can be evaluated analytically and has a known shape, making it ideal to validate the methodology.

3.5.2. Objective function based on mooring mass

The second objective function represents the required steel mass for the catenary mooring system.

$$\begin{aligned} \min_{\mathbf{x}=(m_L, L_{S0})} \quad & f_2(\mathbf{x}) = N_{\text{lines}} \cdot m_L \cdot \max(L_{Sj}(t)) \\ \text{subject to:} \quad & T_d(\mathbf{x}) - C_s \cdot MBL(\mathbf{x}) \leq 0 \\ & \Delta(\mathbf{x}) - \Delta_{\max} \leq 0 \\ & 450 \leq m_L \leq 650 \\ & 120 \leq L_{S0} \leq 180 \end{aligned} \quad (13)$$

where N_{lines} is the number of mooring lines, m_L the mass per unit length of each mooring line and $\max(L_{Sj}(t))$ the maximum value of suspended length of the j -th mooring line during the dynamic simulation. Eq. (13) represents the total suspended mass of the mooring system, which is used as a measure of the working chain portion. The mooring system spread is not included in the optimisation analysis, therefore the total mooring line length cannot be used explicitly as a measure of the system cost. During the analysis very long lay lengths have been considered and the maximum suspended length observed during each dynamic simulation has been considered as a measure of the working mooring line length.

3.5.3. Objective function based on mooring mass with penalisation

To add a bit of complexity to the objective function in Eq. (13), a penalisation term has been included to represent the proximity of the solution to the limit of the non-acceptable values determined by the constraints. These additional penalisation terms are intended to act as *soft* constraints to prevent optima that on the boundary of the acceptable solution space. The resulting objective function is shown in Eq. (14).

$$\begin{aligned} \min_{\mathbf{x}=(m_L, L_{S0})} \quad & f_3(\mathbf{x}) = \left(\omega_1 \cdot \frac{\Delta(\mathbf{x})}{\Delta_{\max}} + \omega_2 \cdot \frac{T_d(\mathbf{x})}{C_s \cdot MBL} \right) \cdot \\ & N_{\text{lines}} \cdot m_L \cdot \max(L_{Sj}(t)) \\ \text{subject to:} \quad & T_d(\mathbf{x}) - C_s \cdot MBL(\mathbf{x}) \leq 0 \\ & \Delta(\mathbf{x}) - \Delta_{\max} \leq 0 \\ & 450 \leq m_L \leq 650 \\ & 120 \leq L_{S0} \leq 180 \end{aligned} \quad (14)$$

Table 2

Initial grid search for the three case-studies. The values of $T_d(\mathbf{x})$ and $\Delta_{\max}(\mathbf{x})$ are used to compute the nonlinear constraints, and therefore independent of the considered case study.

m_L [kg m ⁻¹]	L_{S0} [m]	$T_d(\mathbf{x})$ [N]	$\Delta_{\max}(\mathbf{x})$ [m]	$f_1(\mathbf{x})$ [kg]	$f_2(\mathbf{x})$ [kg]	$f_3(\mathbf{x})$ [kg]
450	120	7.394e+06	31.31	1.545e+05	5.272e+05	3.885e+05
450	150	7.699e+06	22.96	0.570e+05	5.542e+05	3.573e+05
450	180	8.400e+06	17.85	2.747e+05	5.949e+05	3.595e+05
550	120	6.833e+06	29.74	3.091e+05	5.782e+05	3.613e+05
550	150	7.239e+06	21.40	2.206e+05	6.113e+05	3.287e+05
550	180	8.222e+06	16.87	4.473e+05	6.693e+05	3.430e+05
650	120	6.503e+06	28.30	3.636e+05	6.149e+05	3.469e+05
650	150	6.981e+06	20.10	2.841e+05	6.736e+05	3.175e+05
650	180	8.282e+06	16.24	5.198e+05	7.422e+05	3.469e+05

where $\frac{\Delta(\mathbf{x})}{\Delta_{\max}}$ represents the proximity of the solution to the maximum offset constraint and $\frac{T_d(\mathbf{x})}{C_s \cdot MBL}$ the proximity to the maximum tension constraint. ω_1 and ω_2 are weight values that define the relative relevance of each term. In this study, a value of $\omega_1 = 0.25$ and $\omega_2 = 0.75$ have been used.

3.5.4. Initial grid search

A uniform grid has been used to generate the samples on the initial grid search. The grid has been generated using three samples for each dimension, resulting in 9 evaluations of the objective function. The data on the grid search is given in Table 2.

4. Results


This section contains the results of the three optimisation case studies presented in Section 3. In all three cases, a value of $\beta_0 = 10$ and $\xi = 0.95$ have been employed. The selection of these values is based on a sensitivity analysis (see Appendix A) carried out comparing the performance of the BO framework with different combinations of β_0 and ξ . These parameters have shown to provide reasonable exploration while keeping the total iteration number below 50 in all cases, providing accurate estimations of $f(\mathbf{x})$ when converged.

4.1. Validation based on a multi-modal objective function


Fig. 4 shows the results of the validation. The points in the initial grid search have been marked using red squares, whereas the red crosses represent the points that have been evaluated during the optimisation process. The first thing to notice is that, after 26 iterations, the BO process has been able to converge to the true minimum of $f_1(\mathbf{x})$ with a negligible error, as shown in Fig. 4e. At the beginning, the exploration-exploitation trade-off presents a higher ratio of exploration, which explains the number of evaluations far away from the optimum region, allowing the BO process to detect the presence of a local minimum. After 26 iterations, exploitation is predominant, indicating that the BO has reached convergence. Notice that the acquisition function shows a clear and pronounced minimum in the region where the global minimum is located, as illustrated in Fig. 4d. Also notice that, in the vicinity of this region, the acquisition function and the prediction (see Fig. 4a) present very similar values.

The ability of the BO framework to estimate the distribution of $f_1(\mathbf{x})$ within the whole parametric space has been measured using the Mean Absolute Percentage Error (MAPE) (de Myttenaere et al., 2016). The MAPE represents the relative error between the true and predicted values in a regression problem. In this case, the estimated value of $f_1(\mathbf{x})$ depicted in Fig. 4a is compared against the true value of $f_1(\mathbf{x})$, shown in Fig. 4c. The MAPE is computed as shown in (15).


$$\text{MAPE}(y_t, y_p) = \frac{1}{n_s} \cdot \sum_{i=0}^{n_s-1} \frac{|y_{ti} - y_{pi}|}{\max(\epsilon, |y_{ti}|)} \quad (15)$$



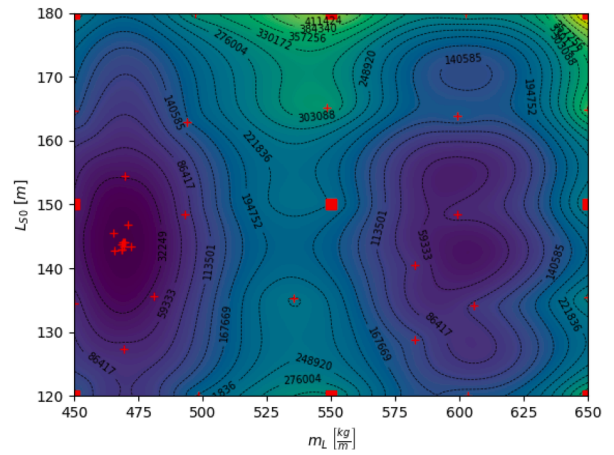
(a) Surrogate's estimation of $f_1(\mathbf{x})$




(b) Uncertainty




(c) Actual shape of $f_1(\mathbf{x})$



(d) Acquisition function



(e) Residual plot



(f) MAPE

Fig. 4. Results of case study 1 after 30 iterations of BO. The red squares indicate the samples from the grid search and the red crosses indicate the samples evaluated during the optimisation process. The error in (c) refers to the relative error between the estimated minimum at each iteration and the true minimum of $f_1(\mathbf{x})$. The MAPE compares the surrogate's estimation of $f_1(\mathbf{x})$ at each iteration with its actual shape (c).

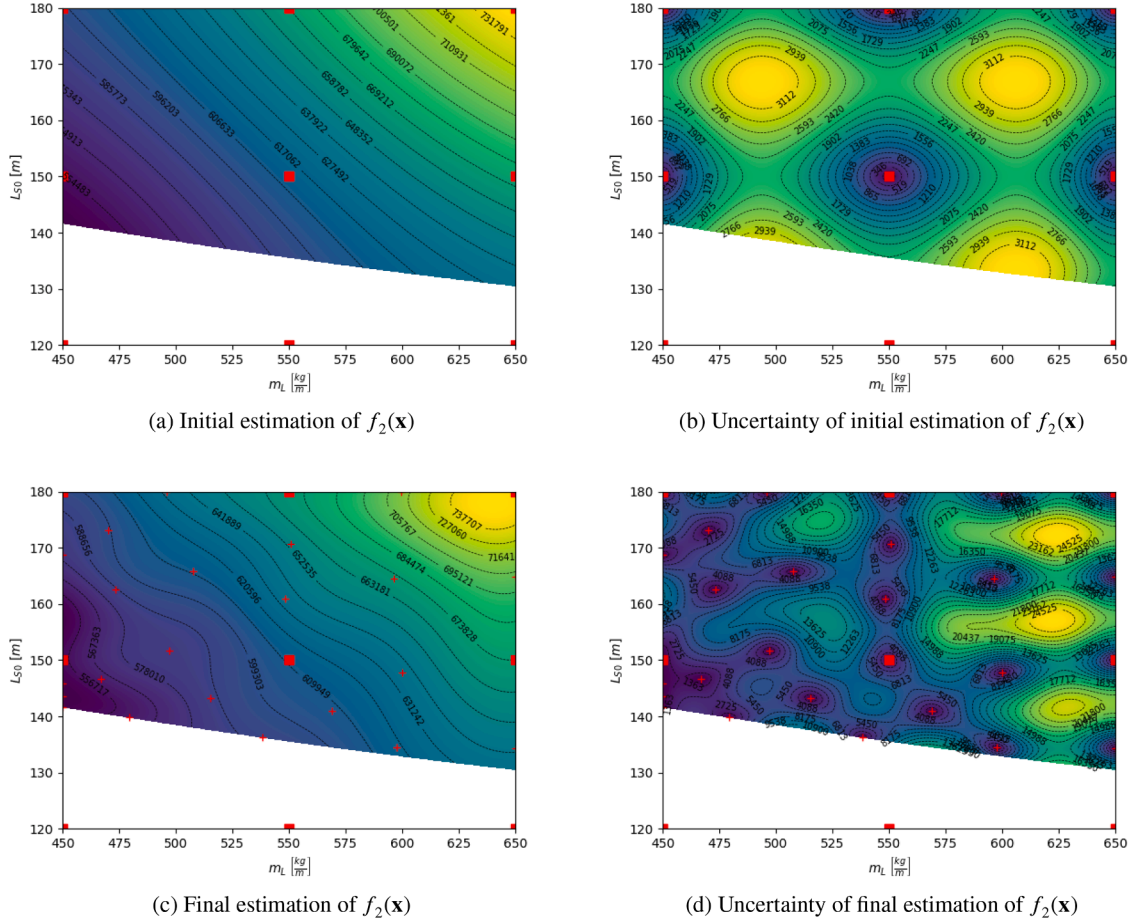


Fig. 5. Initial and final estimations of $f_2(\mathbf{x})$. The red squares indicate the samples from the grid search and the red crosses indicate the samples evaluated during the optimisation process. Initial and final estimations refer to surrogate's prediction before and after the BO has run, respectively.

where y_t is the true value, y_p the prediction, n_s the number of samples and ϵ a very small positive value to avoid a division by 0. As shown by Fig. 4f, the MAPE reduces as the BO progresses until the process converges and no more exploration is performed. At that stage, the MAPE stabilises at a value of 0.025, meaning that the mean prediction error is of 2.5% within the whole parametric space. Qualitatively, this can be seen comparing the prediction in Fig. 4a and the real value of $f_1(\mathbf{x})$ in Fig. 4c, which show good agreement, especially in the surroundings of the minimum, in which a higher number of samples of $f_1(\mathbf{x})$ have been evaluated. It is also consistent with the low values of prediction uncertainty, as shown in Fig. 4b.

4.2. Objective function based on mooring mass

The initial estimation of $f_2(\mathbf{x})$ as per Eq. (13) is shown in Fig. 5. The white shaded area shows the parametric region that violates any of the constraints described in Section 2.3. The initial grid search suggests a continuous and smooth decreasing tendency of $f_2(\mathbf{x})$ towards lower values of both m_L and L_{S0} . This is to be expected considering that shorter, lighter mooring lines have a lower mass. However, as the lines become lower and lighter the platform offset increases, which is shown as a widening of the shaded area representing the unfeasible region.

The uncertainty associated to this initial prediction, which can be found in Fig. 5b, is considered null for the points that have been included as part of the initial grid search, and grows as the distance to those points increase. The null value of uncertainty at the evaluated samples indicates that the GP prediction effectively interpolates the observed data. The regular increase of estimated uncertainty as distance

to the evaluated points increases is caused by the choice of Eq. (5) as a kernel function. However, notice that the overall uncertainty of the estimation is low ($\sim 10^3$) compared to the estimated value ($\sim 10^5$). The low values of estimated uncertainty indicate that the initial prediction is smooth, not showing ripples or steep gradients. Nevertheless, this initial estimation is completely biased by the sample distribution of the initial grid, which is somewhat arbitrary. In fact, the estimation of uncertainty in the last iteration (see Fig. 5d) shows generally higher uncertainty values. However, the highest uncertainty values at the end of the BO process are lower than 4% of the estimated value. This means that the nature of $f_1(\mathbf{x})$ is captured well even in the most under-sampled areas.

According to $f_2(\mathbf{x})$, the optimised mooring design corresponds to a configuration with $m_L = 450$ kg/m and $L_{S0} = 141.65$ m. At this stage, the BO process also ends up with an estimation of the real shape of $f_2(\mathbf{x})$, which is given in Fig. 5c together with its associated uncertainty in Fig. 5d. This estimation provides valuable insight from a decision-making perspective to analyse how the cost is distributed within the design space. Within the range of $450 \leq m_L \leq 525$, the cost increase is steeper for systems with a 10% increase of the initial suspended length than for systems with a 10% increase in the linear mass of the lines. However, as the lines become heavier ($m_L \geq 525$), increasing the initial suspended length seems more reasonable than choosing heavier mooring lines.

Note that the estimation of cost in the region comprising designs with lower m_L , which contains the optimum, presents a much lower uncertainty when compared to solutions with higher m_L (see Fig. 5d). A higher value of β_0 or ξ resulting in a higher exploration/exploitation ratio for more iterations would have resulted in a more densely evaluated design

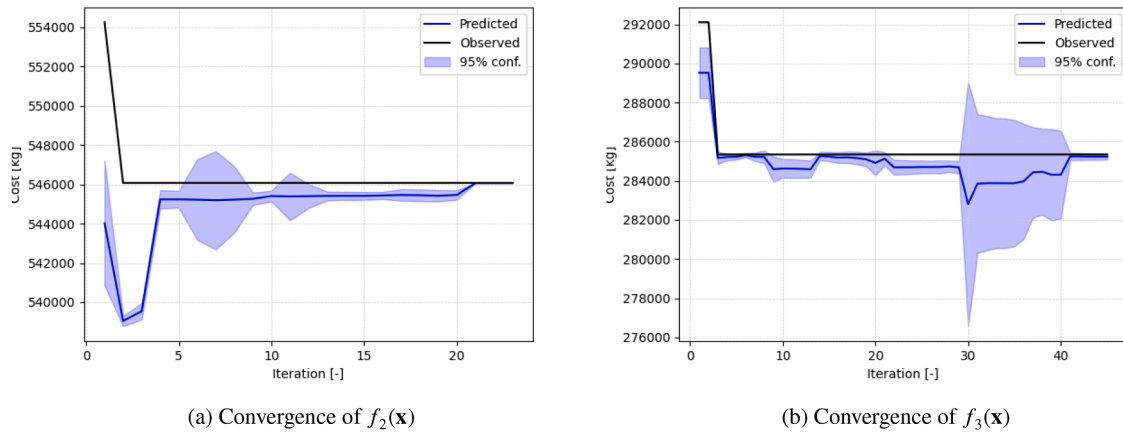


Fig. 6. Convergence plot of the cost values. The black line represents the minimum cost observed at each iteration. The blue line represents the minimum cost predicted by the surrogate model at each iteration, and the blue shaded area indicates the uncertainty of that estimation.

space with a lower overall uncertainty. However, $f_2(\mathbf{x})$ is not expected to have a very complex shape, given that the only nonlinearities are the ones introduced by the presence of peak values of the suspended length. Consequently, a higher exploration/exploitation ratio would result in a similar estimation of $f_2(\mathbf{x})$ at the expense of increasing evaluations to reach convergence.

The evolution of the minimum predicted value of cost is depicted in Fig. 6a. The black line represents the best observed point, which is the configuration among the ones evaluated so far at each iteration that has the minimum value of cost. The blue line represents the minimum cost value predicted by the surrogate model, with the 95% confidence interval shown as a blue shaded area. The best observed value (black line) drops after the 1st iteration and remains virtually constant during the rest of the optimisation. This indicates that optimum point has been sampled at the 1st iteration. However, note that during the sampling process the GP presents minima for other combinations of (m_L, L_{S0}) . This can be deduced from the predicted value (blue line), which must be always lower or equal to the best observed value. When the BO converges, these two lines must converge to the same value. The shaded region shows changes in the confidence interval of the predicted minimum. It represents the uncertainty associated to the predicted minimum. Note that in many iterations, the shaded area does not include the best observed point. This optimistic behaviour of the surrogate model is associated to the smoothness of $f_2(\mathbf{x})$. Moreover, the maximum discrepancy between the best observed and predicted cost values, which happens at the 1st iteration, is lower than 2%. The small discrepancy between values makes the estimation of uncertainty remain low, as no significant deviations from the estimated values are observed. This causes the surrogate model to under-predict uncertainty during exploration. Even though, as the number of iterations progresses, the BO is able to identify the global minimum in 21 iterations. Note that the number of objective function evaluations required by the proposed optimisation framework is significantly lower than other gradient-free optimisation frameworks. For instance, Faraggiana et al. (2022) solved an equivalent optimisation problem using GA, which also considered two design variables and the minimisation of a mass. The number of evaluations they required ranged between 140–210, in contrast to the 21 evaluations employed by the proposed framework to minimise $f_2(\mathbf{x})$.

4.3. Objective function based on mooring mass with penalisation

The initial estimation and uncertainty of $f_3(\mathbf{x})$ as per Eq. (14) are shown in Fig. 7a and 7b respectively. The initial estimation predicts a minimum towards the heavier side of m_L and the shorter side of L_{S0} (low pretension). Note that the addition of a penalisation term significantly affects the initial estimation of the optimal mooring line config-

uration, highlighting its high dependence on the selection of the optimisation objective. Also note that the shape of the initial estimation of uncertainty (Fig. 7b) is very similar to the initial uncertainty of $f_2(\mathbf{x})$. This is a consequence of using the same sample points to compute this initial grid search combined with the use of the same kernel function (see Eq. (5)).

Fig. 7c shows the estimation of the actual shape of $f_3(\mathbf{x})$ after the BO has converged. In this case, the surrogate model shows that a 10% variation on both m_L and L_{S0} lead to similar increases of $f_3(\mathbf{x})$, showing that the solution is not very sensitive to small variations around the optimum as long as the inequality constraints are not violated. For instance, if the linear mass needs to be reduced, the initial suspended length should be increased accordingly in order not to exceed the maximum allowed offset. Fig. 7c also shows that a dramatic decrease of the linear mass of the lines leads to a smaller cost increase than a dramatic increase of the initial suspended length. This behaviour can be explained by the fact that higher suspended lengths at the beginning lead to higher pretension values, which in turn lead to being closer to the limit established by Eq. (8).

The BO process converged after 41 iterations, almost doubling the number of iterations required for convergence in the optimisation of $f_2(\mathbf{x})$. The slower convergence leads to a higher sample density in the design space and a lower overall uncertainty, especially in the neighbourhood of the optimum. The cause for the slower convergence is the addition of the penalisation term, which introduces a higher level of variability in the shape of the objective function and, consequently, in the predictions of the surrogate model. It is worth noticing that the unfeasible regions, which are represented as white shaded areas in Figs. 5 and 7, are the same for both the optimisation of $f_2(\mathbf{x})$ and $f_3(\mathbf{x})$. The shape of the region where the nonlinear constraints are not fulfilled is also estimated using a GP, as explained in Section 2.3. Therefore, the fact that the estimated shape of this region is the same with two different sample densities and distributions reinforces the strength of this methodology to deal with nonlinear constraints.

Fig. 6b shows the evolution of the observed and predicted minimum values of $f_3(\mathbf{x})$, together with the 95% confidence interval of the predicted value. After 41 iterations, the process converges to an optimum value of $m_L = 650.0$ kg/m and $L_{S0} = 133.27$ m. Around iteration 30, a sudden decrease in the predicted minimum (blue line) with a sudden increase of the confidence interval can be observed. It is the expected response of the GP surrogate model when a certain evaluation results in a value which significantly deviates from the estimation. This happens when the objective function is evaluated in a point whose *prior* distribution shows a high level of uncertainty. As a consequence, the overall uncertainty increases for the next few iterations, showing a rapid decreasing tendency if the surrogate model remains stable. This decrease

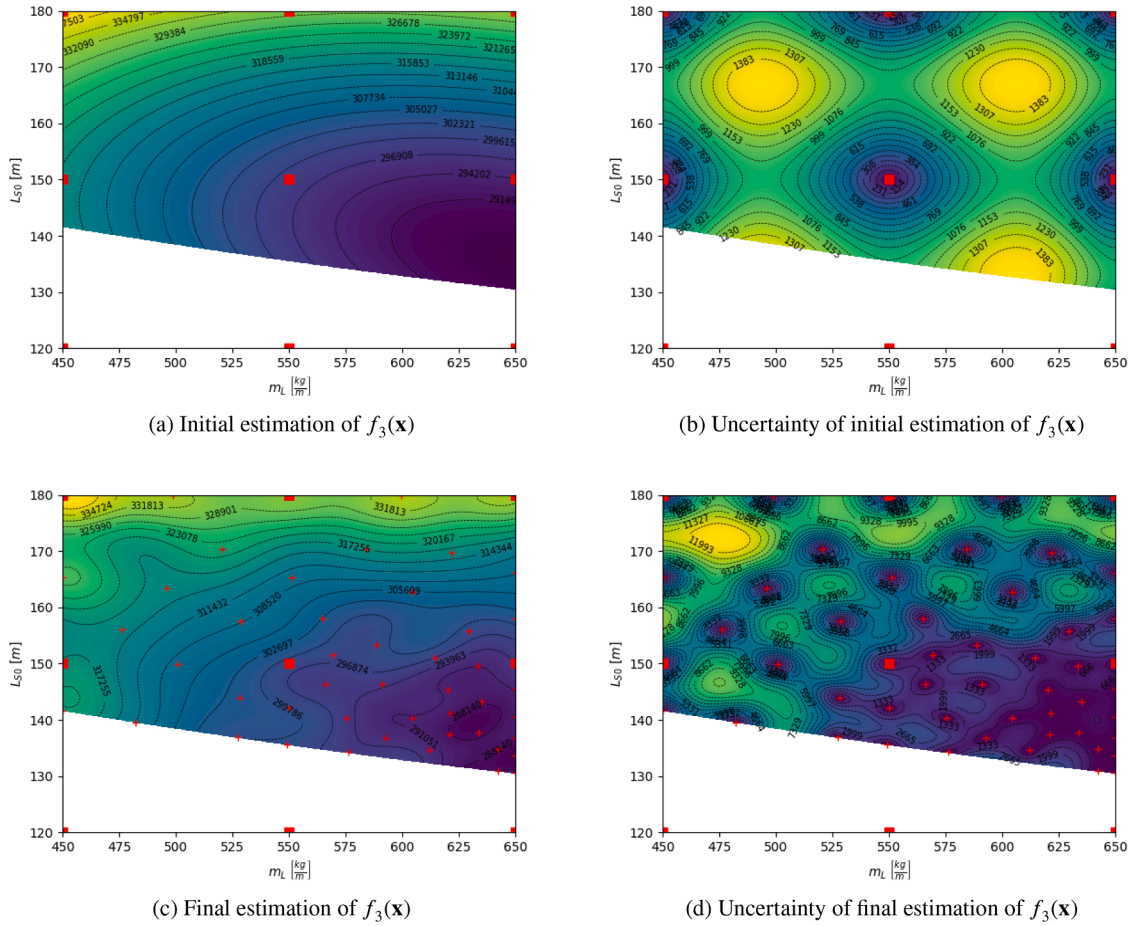


Fig. 7. Initial and final estimations of $f_3(\mathbf{x})$. The red squares indicate the samples from the grid search and the red crosses indicate the samples evaluated during the optimisation process. Initial and final estimations refer to surrogate's prediction before and after the BO has run, respectively.

and return to its previous tendency can be observed during iterations 29–40 in the predictions depicted in Fig. 6b.

4.4. Uncertainty quantification

A key advantage of the proposed framework lies in its ability to provide also a quantification of the uncertainty in surrogate model predictions, offering insights into the reliability of the estimated objective function. This quantification is needed in order that the uncertainty quantification can be integrated into the design process. In this study, uncertainty was assessed through the posterior variance of the GP surrogate model used to guide the optimisation of catenary mooring line designs.

At the conclusion of the optimisation process, the maximum uncertainty observed in the surrogate model predictions was found to be one to two orders of magnitude lower than the corresponding estimated values. This indicates a high level of confidence in the model's predictions across the design space. Two metrics have been employed to quantify the overall predictive uncertainty after the BO process: the peak uncertainty and the integrated uncertainty ratio.

The peak uncertainty ratio is defined as the ratio between the predictive uncertainty and the predictive mean at the location of maximum uncertainty. It provides an upper bound on the expected uncertainty in the surrogate's predictions. On the other hand, The integrated uncertainty ratio is defined as the ratio between the integral of the predictive uncertainty and the integral of the predictive mean across the design space. This metric offers an estimate of the overall expected uncertainty. Table 3 shows the values of these ratios for each case study.

Table 3

Peak uncertainty and integrated uncertainty of each case study at BO completion.

	Peak uncertainty	Integrated uncertainty ratio
Case study 1: Multi-modal objective function	12.19 %	4.43 %
Case study 2: Mooring system mass	6.67 %	1.84 %
Case study 3: Mooring system mass with penalisation	5.59 %	1.71 %

The results show that the overall uncertainty after BO is 4.43 % for Case Study 1, and approximately 1.8% for Case Studies 2 and 3. The lower uncertainty in the latter cases is attributed to two factors: a higher sample density across the design space and the nature of the objective function. Although the true shape of the objective function is unknown in Case Studies 2 and 3, the surrogate estimates shown in Figs. 5 and 7 suggest that the mooring mass (with and without penalisation) exhibits a single minimum within the design space and is nearly monotonic with respect to both m_L and L_{S0} . These characteristics facilitate more accurate surrogate modelling, resulting in lower predictive uncertainty.

5. Conclusions

This study presents an uncertainty-aware framework for the design of mooring lines of FOWTs which integrates Bayesian Optimisation into the design process. The proposed approach leverages a probabilistic surrogate model to approximate the distribution of the objective function

across the design space. The proposed BO framework is designed to offer enhanced decision-making capabilities compared to the optimisation methods often encountered in the literature, such as GA or PS. In addition to yielding an optimised design, the framework provides a probabilistic characterisation of the objective function and its associated uncertainty. This probabilistic representation enables the end user to assess how variations in design parameters influence the objective. In particular, it highlights regions of the design space where the model's predictions are associated with higher uncertainty. Based on this information, the user can make informed decisions—either accepting the current estimate with its associated uncertainty or choosing to perform additional evaluations to reduce uncertainty. It is important to note that increasing the number of evaluations implies a significant computational cost, and the decision to do so depends on the trade-off between computational effort and the potential value of improved accuracy.

When running the optimisation, the LCB acquisition function allows the user to explicitly determine the design space exploration strategy. Moreover, the proposed implementation allows to modify this strategy on the fly, allowing the designer to adapt the exploration/exploitation ratio at runtime. Finally, the inherent probabilistic nature of both the optimiser and the surrogate model enables the explicit incorporation of uncertainty into the optimisation process. The current implementation accounts for the inherent uncertainty of the objective function, which is estimated in the fitting process of the surrogate model.

Three case studies have been considered to demonstrate the aforementioned methodology: (i) a validation case with a multi-modal objective function that has a closed analytical representation, (ii) the design of the mooring system for a semi-submersible platform by minimising its useful mass and (iii) the design of the same mooring system considering a penalisation term together with the useful mass. The validation case demonstrates the ability of the optimisation framework to estimate the shape of a multi-modal function with a MAPE $\leq 2.5\%$ when compared to the real shape of the objective function. When applied to mooring design optimisation, both with and without applying a penalisation term to the system mass, the methodology is also able to estimate the shape of the objective function with the overall peak uncertainty below 4% of the estimated cost. However, the added complexity of the penalisation term in the objective function slows down the convergence of the optimiser, with the number of iterations required for convergence increasing from 21 to 41. It is noteworthy that the number of objective function evaluations required by the proposed optimisation framework is substantially lower than that of other gradient-free optimisation approaches. The results demonstrate that, when compared to the application of a Genetic Algorithm (GA) to an optimisation problem of equivalent dimensionality and with a similarly shaped objective function, the proposed methodology achieves convergence with approximately 80% fewer iterations. This low number of iterations required to achieve convergence is key for incorporating fully-coupled time-domain dynamic FOWT models. Although model complexity increases the evaluation cost, the efficient sampling strategy allows to reduce the number of evaluations when compared to other frameworks, thereby mitigating the need of reduced-order models.

Overall, this work constitutes a promising proof-of-concept for the proposed FOWT design optimisation framework. The design space for the mooring system was limited to two variables: linear weight and suspended length. This simplification does not compromise the generality of the approach and can be readily extended to include additional design variables in future work. Similarly, the objective function can be adapted to reflect broader design criteria, such as installation costs, reliability, and environmental impact. Future developments should also explore the integration of the methodology for the concurrent optimisation of multiple subsystems, such as mooring lines, platform geometry, and dynamic cables. Moreover, we have centred the current study on single-objective optimization. This initial proof-of-concept, despite its limitations in scope, has demonstrated the performance and robustness of the BO framework under a fully controlled setting, given that the minimisa-

tion of mooring system mass provides a relevant and interpretable result. However, the real-world floating offshore wind turbine design involves multiple, often conflicting objectives—such as minimising cost, maximising structural integrity, reducing platform motions and enhancing energy yield. The current methodology can be easily extended to cope with more realistic multi-objective optimisation. As a matter of fact, ongoing work is investigating the impact of multi-objective formulations on the final design, enabling a more comprehensive exploration of trade-offs and design synergies across different performance metrics. This extension will enhance the applicability of the framework to practical engineering scenarios and realistic decision-making processes. Finally, while this study accounts for uncertainty in the surrogate model predictions, future extensions should aim to incorporate additional sources of uncertainty, such as the variability in input parameters.

CRedit authorship contribution statement

Ander Aristondo: Writing – original draft, Visualization, Software, Methodology, Conceptualization; **Vincenzo Nava:** Writing – review & editing, Supervision, Methodology, Funding acquisition, Conceptualization; **Markel Penalba:** Writing – review & editing, Supervision, Conceptualization; **Miguel Esteras:** Writing – review & editing, Methodology, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This publication is part of the project RUL-ET, a research project funded by the Basque Government's ELKARTEK 2024 program under the grant No. KK-2024/00086. Additionally, Dr. Penalba is part of the research project PID2021-124245OA-I00 funded by MCIN/AEI/10.13039/501100011033 and by ERDF A way of making Europe, and the Basque Government's Research Group Program under the grant No. IT1505-22. Vincenzo Nava's work has been carried out within the RETURN Extended Partnership and received funding from the European Union Next-GenerationEU (National Recovery and Resilience Plan – NRRP, Mission 4, Component 2, Investment 1.3 – D.D. 1243 2/8/2022, PE0000005) – SPOKE TS 2.

Appendix A. Selection of β_0 and ξ

To select appropriate values for β_0 and ξ , a sensitivity analysis was conducted. The analysis considered $\beta_0 \in 5, 10$ and $\xi \in 0.9, 0.95, 0.99$, and evaluated the optimiser's performance on the minimisation of $f_1(\mathbf{x})$. Performance was assessed by tracking the evolution of three metrics across iterations: the mean absolute percentage error (MAPE), the location error, and the integrated uncertainty.

The MAPE quantifies the mean absolute percentage error between the true shape of $f_1(\mathbf{x})$ and its surrogate model estimation at each iteration. A perfect match yields a MAPE of 0, so lower values indicate better estimations. The MAPE is computed as shown in Eq. (15), and is replicated below for clarity:

$$\text{MAPE}(y_i, y_p) = \frac{1}{n_s} \cdot \sum_{i=0}^{n_s-1} \frac{|y_{ti} - y_{pi}|}{\max(\epsilon, |y_i|)} \quad (\text{A.1})$$

where y_i is the true value, y_p the predicted value, n_s the number of samples, and ϵ a small positive constant to avoid division by zero.

The location error measures the distance between the best observation at the current iteration (\mathbf{x}_{best}) and the true minimum (\mathbf{x}_{true}) of the objective function:

$$\text{Location error} = |\mathbf{x}_{\text{best}} - \mathbf{x}_{\text{true}}| \quad (\text{A.2})$$

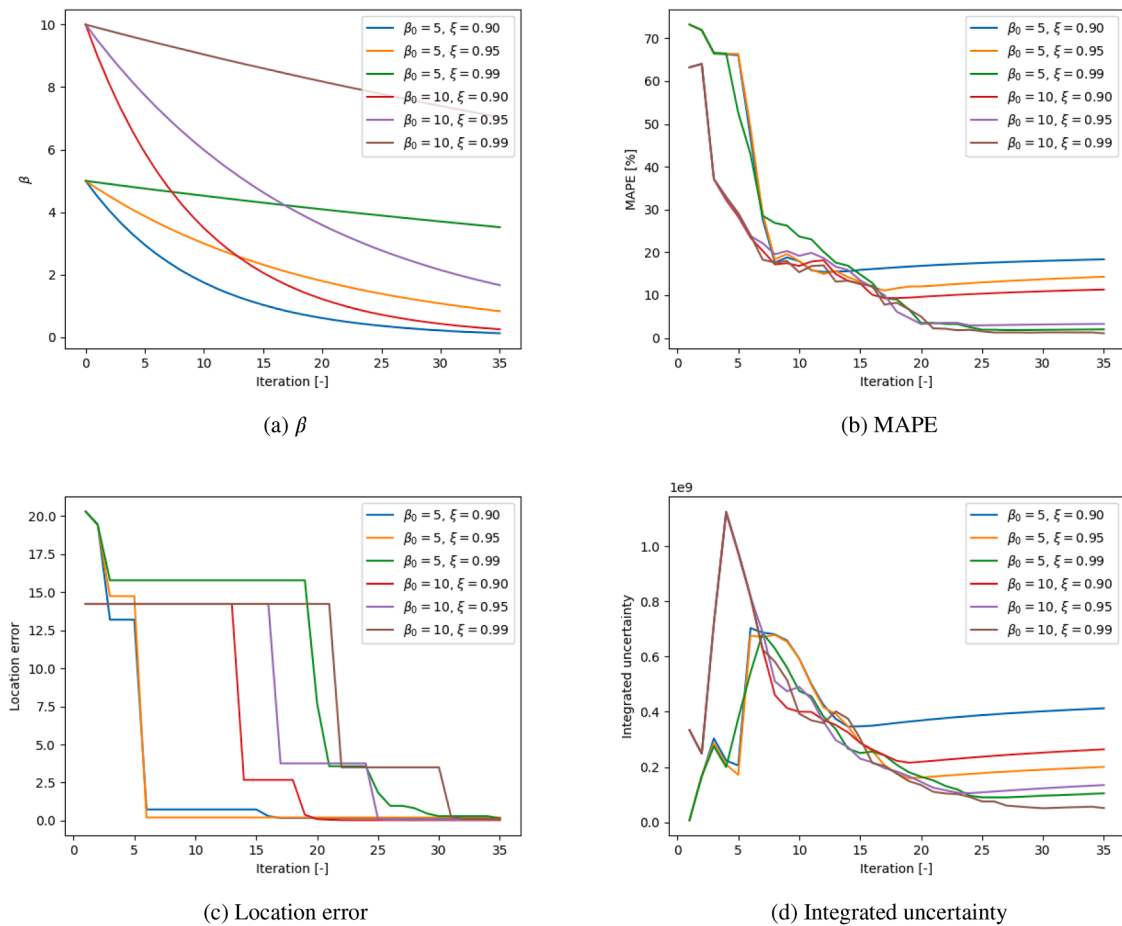


Fig. A.1. Sensitivity analysis results for selecting β_0 and ξ . (a) shows the evolution of β for each strategy. (b) presents the evolution of the MAPE between the predicted and true shape of $f_1(x)$ across iterations for each parameter combination. (c) and (d) illustrate the evolution of the location error and the integrated uncertainty, respectively, over iterations for each combination.

The integrated uncertainty quantifies the total predictive uncertainty remaining after each iteration. It is computed as:

$$\text{Integrated uncertainty} = \iint_{m_L, L_{S0}} \sigma(x) dx \tag{A.3}$$

where $\sigma(x)$ is the predictive standard deviation of the Gaussian Process (GP) model, and $\mathbf{x} = (m_L, L_{S0})$.

Fig. A.1 presents the results of the sensitivity analysis. Each optimisation run was performed for 35 iterations. The iteration at which the location error reaches zero indicates when the optimiser has found the true minimum, serving as estimation for convergence.

The results show that a value of $\xi = 0.99$ consistently yields a low MAPE, regardless of β_0 . Consequently, these configurations also result in the lowest overall uncertainty after convergence. However, they typically require more than 30 iterations to locate the global minimum. As shown in Fig. A.1a, this is due to the slow decay of β associated with $\xi = 0.99$, which promotes prolonged exploration. In this regime, the uncertainty term in the LCB acquisition function dominates the sampling process for a longer period.

In contrast, lower values of $\xi = 0.90$ lead to faster convergence but at the cost of higher MAPE and integrated uncertainty. This behaviour reflects a more exploitative search strategy, driven by the rapid decay of β and the resulting dominance of the exploitation term in the LCB.

For the intermediate value $\xi = 0.95$, the optimiser’s behaviour depends significantly on the choice of β_0 . The combination $\beta_0 = 5$ and $\xi = 0.95$ results in a highly exploitative strategy, with rapid convergence but higher final MAPE and uncertainty. On the other hand, the combination $\beta_0 = 10$ and $\xi = 0.95$ achieves a better balance between explo-

ration and exploitation. It explores the design space more thoroughly, achieves a final MAPE and uncertainty comparable to the more exploratory $\xi = 0.99$ cases, and converges in fewer than 30 iterations.

Based on these findings, the combination $\beta_0 = 10$ and $\xi = 0.95$ was selected for this study.

References

Allen, C., Viscelli, A., Dagher, H., Goupee, A., Gaertner, E., Abbas, N., Hall, M., Barter, G., 2020. Definition of the UMaine VoltturnUS-S Reference Platform Developed for the IEA Wind 15-Megawatt Offshore Reference Wind Turbine. Technical Report NREL/TP-5000-76773, 1660012, MainId:9434. NREL. <https://doi.org/10.2172/1660012>

Aristondo, A., Abanda, A., Esteras, M., Nava, V., Penalba, M., 2025. A methodology for expert knowledge imbrication in mooring system design using Bayesian based optimization. In: Innovations in Renewable Energies Offshore - Proceedings of the 6th International Conference on Renewable Energies Offshore, RENEW 2024, p. 361–367. <https://doi.org/10.1201/9781003558859-40>

Brochu, E., Cora, V.M., de Freitas, N., 2010. A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning. arXiv:1012.2599 [cs].

Centeno-Telleria, M., Yue, H., Carrol, J., Aizpuru, J.I., Penalba, M., 2024. O&M-aware techno-economic assessment for floating offshore wind farms: a geospatial evaluation off the North Sea and the Iberian Peninsula. Appl. Energy 371, 123684. <https://doi.org/10.1016/j.apenergy.2024.123684>

Cousin, A., Delépine, N., Guiton, M., Munoz Zuniga, M., Perdrizet, T., 2024. Optimal design of experiments for computing the fatigue life of an offshore wind turbine based on stepwise uncertainty reduction. Struct. Saf. 110, 102483. <https://doi.org/10.1016/j.strusafe.2024.102483>

Cox, D.D., John, S., 1992. A statistical method for global optimization. In: [Proceedings] 1992 IEEE International Conference on Systems, Man, and Cybernetics, pp. 1241–1246 vol.2. <https://doi.org/10.1109/ICSMC.1992.271617>

Di Fiore, F., Nardelli, M., Mainini, L., 2024. Active learning and Bayesian optimization: a unified perspective to learn with a goal. Arch. Comput. Methods Eng. 31 (5), 2985–3013. <https://doi.org/10.1007/s11831-024-10064-z>

- DNV, 2016. DNV-ST-0437 Loads and site conditions for wind turbines. DNV, 2021. DNV-ST-0119 Floating wind turbine structures.
- Dou, S., Pegalajar-Jurado, A., Wang, S., Bredmose, H., Stolpe, M., 2020. Optimization of floating wind turbine support structures using frequency-domain analysis and analytical gradients. *J. Phys. Conf. Ser.* 1618 (4), 042028. <https://doi.org/10.1088/1742-6596/1618/4/042028>
- Du, X., Liang, J., Qian, G., Yang, Y., Xie, P., Zhang, K., 2024. A computational framework for the co-optimization of platform hydrodynamic and passive structural control of floating offshore wind turbines. *Ocean Eng.* 293, 116609. <https://doi.org/10.1016/j.oceaneng.2023.116609>
- Díaz, H., Guedes Soares, C., 2023. Cost and financial evaluation model for the design of floating offshore wind farms. *Ocean Eng.* 287, 115841. <https://doi.org/10.1016/j.oceaneng.2023.115841>
- Echevarrieta, J., Arza, E., Pérez, A., 2024. Speeding-up evolutionary algorithms to solve black-box optimization problems. *IEEE Trans. Evol. Comput.* PP, 1. <https://doi.org/10.1109/TEVC.2024.3352450>
- Faraggiana, E., Sirigu, M., Ghigo, A., Bracco, G., Mattiazio, G., 2022. An efficient optimization tool for floating offshore wind support structures. *Energy Rep.* 8, 9104–9118. <https://doi.org/10.1016/j.egy.2022.07.036>
- Fu, A., Liu, Y., Liu, Z., Liu, H., Liu, X., Song, C., Zhang, Z., 2025. Gaussian process-enhanced predictive control for individual pitch regulation in floating offshore wind turbines. *Ocean Eng.* 332, 121394. <https://doi.org/10.1016/j.oceaneng.2025.121394>
- Fylling, I., Berthelsen, P.A., 2011. WINDOPT: an optimization tool for floating support structures for deep water wind turbines. In: Volume 5: Ocean Space Utilization: Ocean Renewable Energy. ASME/EDC, Rotterdam, The Netherlands, pp. 767–776. <https://doi.org/10.1115/OMAE2011-49985>
- Gaertner, E., Rinker, J., Sethuraman, L., Zahle, F., Anderson, B., Barter, G.E., Abbas, N.J., Meng, F., Bortolotti, P., Skrzypinski, W., Scott, G.N., Feil, R., Bredmose, H., Dykes, K., Shields, M., Allen, C., Viselli, A., 2020. IEA Wind TCP Task 37: Definition of the IEA 15-Megawatt Offshore Reference Wind Turbine. Technical Report. NREL. <https://doi.org/10.2172/1603478>
- Gardner, J.R., Kusner, M.J., Xu, Z., Weinberger, K.Q., Cunningham, J.P., 2014. Bayesian optimization with inequality constraints. In: Proceedings of the 31st International Conference on Machine Learning - Volume 32. JMLR.org, p. II-937–II-945.
- Hall, M., 2020. MoorDyn V2: new capabilities in mooring system components and load cases. In: International Conference on Offshore Mechanics and Arctic Engineering, Vol. 9, Ocean Renewable Energy, p. V009T09A078. <https://doi.org/10.1115/OMAE2020-19341>
- Hauteclouque, G.d., Mackay, E., Vanem, E., 2022. Quantitative comparison of environmental contour approaches. *Ocean Eng.* 245, 110374. <https://doi.org/10.1016/j.oceaneng.2021.110374>
- Hegseth, J.M., Bachynski, E.E., Martins, J. R.R.A., 2020. Integrated design optimization of spar floating wind turbines. *Mar. Struct.* 72, 102771. <https://doi.org/10.1016/j.marstruc.2020.102771>
- IEA, 2022. World energy outlook 2022. International Renewable Energy Agency. <https://doi.org/10.1787/3a469970-en>
- IEA, 2024. IEA-15-240-RWT. IH Cantabria, 2024. Site Condition Assessment - BiMEP. Technical Report. IH Cantabria.
- IRENA, 2024. Floating offshore wind outlook. International Renewable Energy Agency.
- IRENA, GWEC 2023. Enabling frameworks for offshore wind scaleup: innovations in permitting. International Renewable Energy Agency.
- Jones, D.R., Schonlau, M., Welch, W.J., 1998. Efficient global optimization of expensive black-box functions. *J. Global Optim.* 13 (4), 455–492. <https://doi.org/10.1023/A:1008306431147>
- Leimeister, M., Collu, M., Kolios, A., 2022. A fully integrated optimization framework for designing a complex geometry offshore wind turbine spar-type floating support structure. *Wind Energy Sci.* 7 (1), 259–281. <https://doi.org/10.5194/wes-7-259-2022>
- Lemmer, F., Yu, W., Müller, K., Cheng, P.W., 2020. Semi-submersible wind turbine hull shape design for a favorable system response behavior. *Mar. Struct.* 71, 102725. <https://doi.org/10.1016/j.marstruc.2020.102725>
- Lim, J., Choi, M., Lee, S., 2023. A bayesian optimization algorithm for the optimization of mooring system design using time-domain analysis. *J. Mar. Sci. Eng.* 11 (3). <https://doi.org/10.3390/jmse11030507>
- López-Queija, J., Tena, A., Jugo, J., Aristondo, A., Robles, E., 2024. Simultaneous design optimisation methodology for floating offshore wind turbine substructure and feedback-based control strategy. *Appl. Ocean Res.* 150, 104120. <https://doi.org/10.1016/j.apor.2024.104120>
- Morison, J.R., Johnson, J.W., O'Brien, M.P., 1953. Experimental studies of forces on piles. *Coast. Eng. Proc.* (4), 25. <https://doi.org/10.9753/icce.v4.25>
- Mulders, S., Pamososuryo, A., van Wingerden, J., 2020. Efficient tuning of individual pitch control: a Bayesian optimization machine learning approach. *J. Phys. Conf. Ser.* 1618 (2), 022039. <https://doi.org/10.1088/1742-6596/1618/2/022039>
- de Myttenaere, A., Golden, B., Le Grand, B., Rossi, F., 2016. Mean absolute percentage error for regression models. *Neurocomputing* 192, 38–48. Advances in artificial neural networks, machine learning and computational intelligence. <https://doi.org/10.1016/j.neucom.2015.12.114>
- Nava, V., Ruiz-Minguela, P., Perez-Moran, G., Rodriguez Arias, R., Lopez Mendia, J., Villate-Martinez, J.-L., 2019. Installation, Operation and Maintenance of Offshore Renewables. The Institution of Engineering and Technology. chapter 11. pp. 397–424. https://doi.org/10.1049/PBPO129E_ch11
- NREL, 2024. OpenFAST v3.5.1.
- Ojo, A., Collu, M., Coraddu, A., 2025. Geometric design parameterization and optimization of spar floating offshore wind turbine substructure. *Ocean Eng.* 332, 121378. <https://doi.org/10.1016/j.oceaneng.2025.121378>
- Orcina, 2023. OrcaFlex Documentation. Edition 11.4.
- Patryniak, K., Collu, M., Coraddu, A., 2022. Multidisciplinary design analysis and optimisation frameworks for floating offshore wind turbines: state of the art. *Ocean Eng.* 251, 111002. <https://doi.org/10.1016/j.oceaneng.2022.111002>
- Pegalajar-Jurado, A., Borg, M., Bredmose, H., 2018. An efficient frequency-domain model for quick load analysis of floating offshore wind turbines. *Wind Energy Sci.* 3 (2), 693–712. <https://doi.org/10.5194/wes-3-693-2018>
- Penalba, M., Zarketa-Astigarraga, A., Branson, P., Robertson, B., 2023. Impact of resource uncertainties on the design of wave energy converters. In: Proceedings of the European Wave and Tidal Energy Conference. Vol. 15. <https://doi.org/10.36688/ewtec-2023-529>
- Pollini, N., Pegalajar-Jurado, A., Bredmose, H., 2023. Design optimization of a TetraSpar-type floater and tower for the IEA wind 15MW reference wind turbine. *Mar. Struct.* 90, 103437. <https://doi.org/10.1016/j.marstruc.2023.103437>
- Powell, M.J.D., 1994. A Direct Search Optimization Method That Models the Objective and Constraint Functions by Linear Interpolation. Springer Netherlands, Dordrecht. pp. 51–67. https://doi.org/10.1007/978-94-015-8330-5_4
- Rafiee, A., Haase, M., Malcolm, A., 2022. Multi-objective Bayesian hull form optimisation for high-speed craft. *Ocean Eng.* 266, 112688. <https://doi.org/10.1016/j.oceaneng.2022.112688>
- Rafiee, A., Repalle, N., 2024. Harnessing uncertainty: a generalized polynomial chaos approach to mooring load analysis in offshore floating wind turbines. In: International Conference on Offshore Mechanics and Arctic Engineering. Vol. Volume 7: Ocean Renewable Energy, p. V007T09A051. <https://doi.org/10.1115/OMAE2024-128260>
- Rasmussen, C.E., Williams, C. K.I., 2008. Gaussian processes for machine learning. Adaptive Computation and Machine Learning, MIT Press, Cambridge, Mass., third ed..
- Rohrer, P.J., Bachynski-Polić, E.E., Hegseth, J.M., 2023. Gradient-based design optimization of fully-flexible floating wind turbines using modal analysis. In: International Conference on Offshore Mechanics and Arctic Engineering. Vol. 8: Ocean Renewable Energy, p. V008T09A030. <https://doi.org/10.1115/OMAE2023-101930>
- Shahriari, B., Swersky, K., Wang, Z., Adams, R.P., de Freitas, N., 2016. Taking the human out of the loop: a review of Bayesian optimization. *Proc. IEEE* 104 (1), 148–175. <https://doi.org/10.1109/JPROC.2015.2494218>
- Snoek, J., Larochelle, H., Adams, R.P., 2012. Practical Bayesian Optimization of Machine Learning Algorithms. Vol. 25. Curran Associates, Inc.
- Sykes, V., Collu, M., Coraddu, A., 2023. A review and analysis of optimisation techniques applied to floating offshore wind platforms. *Ocean Eng.* 285, 115247. <https://doi.org/10.1016/j.oceaneng.2023.115247>
- WFO, 2024. Global Offshore Wind Report.
- Yin, Q., Zhu, H., Yang, J., Ni, H., Gao, B., 2025. Research and application of optimization method for semi-submersible platform mooring system based on deep learning. *Ocean Eng.* 329, 121084. <https://doi.org/10.1016/j.oceaneng.2025.121084>
- Yuan, J., He, S., Cheng, Z., Hu, B., Ai, C., 2024. Optimized model-free adaptive load mitigation control for offshore floating wind turbines using Bayesian techniques. *Ocean Eng.* 314, 119672. <https://doi.org/10.1016/j.oceaneng.2024.119672>
- Zhan, D., Xing, H., 2020. Expected improvement for expensive optimization: a review. *J. Global Optim.* 78 (3), 507–544. <https://doi.org/10.1007/s10898-020-00923-x>
- Zhang, S., Yang, F., Yan, C., Zhou, D., Zeng, X., 2022. An Efficient Batch Constrained Bayesian Optimization Approach for Analog Circuit Synthesis via Multi-objective Acquisition Ensemble. *IEEE Trans. Computer-Aided Des. Integr. Circuits Syst.* 41 (1), 1–14. arXiv:2106.15412 [cs, eess]. <https://doi.org/10.1109/TCAD.2021.3054811>