Tidal Turbine Performance Prediction and Optimization using a BEMT-based Multi-layer Perceptron Surrogate Model

Dun Liu, Kamal Djidjeli, Zhiwei Hu, Abdessalem Bouferrouk
Corresponding author: dl3g20@soton.ac.uk
University of Southampton, UK.
University of the West of England, UK

Abstract: The growing global concern over climate change and the continued reliance on fossil fuels have intensified the demand for sustainable and clean energy solutions. Among various renewable energy sources, tidal energy is regarded as a promising alternative due to its high predictability, long-term sustainability, and relatively low environmental impact. Tidal turbines play a crucial role in harnessing tidal energy, and improving their performance remains a key research challenge.

This study aims to enhance the computational efficiency of high-fidelity turbine performance prediction by developing a surrogate modeling framework based on Blade Element Momentum Theory (BEMT). The proposed surrogate model is employed to improve the turbine performance by optimising the blade cross-section shape, with a particular focus on achieving an optimal balance between thrust and power output at the turbine level. The surrogate model is validated against BEMT calculations and demonstrates good agreement in predicting turbine performance metrics. The surrogate model presents a higher efficiency in turbine performance prediction and blade shape optimisation, which enables to screen 20,000 blade cross-section shapes in 6 minutes and is much faster than CFD simulations. The CFD results confirms the performance improvement and indicates that the selected cross-section shapes provide an average reduction in thrust of up to 29.062%, over a range of tip-speed ratio, while maintaining the power output.

Keywords: Computational Fluid Dynamics, Machine Learning, Turbulence Modelling.

1. Introduction

The growing urgency of climate change mitigation and the need to reduce reliance on fossil fuels have accelerated the deployment of sustainable energy technologies. Among marine renewables, tidal energy is particularly attractive due to its predictability, long-term sustainability, and comparatively low environmental impact [1]. Tidal stream turbines are the key technology in this sector, and enhancing the turbine performance remains a central challenge. However, in most cases, approaches like finite element methods and computational fluid dynamics simulations, which are supported by expensive hardware resources. Such methods also require relevant computational times and manually defined dense meshes, where equations are evaluated iteratively for each time step. Due to the explosive development in computational science, machine learning, and deep learning techniques have shown revolutionary growth in various areas, which has led to increased attention on neural networks in the field of fluid mechanics [2]. The surrogate model is utilised to reduce the requirement of dense meshing and iterative estimation over time steps and overcome the hardware limits.

In this work, a surrogate modelling framework is developed by integrating Blade Element Momentum Theory (BEMT) with a neural network. BEMT provides an efficient, physics-based tool for predicting turbine-level performance by coupling blade element aerodynamics with momentum theory, using sectional lift and drag characteristics from two-dimensional aerofoil data [3]. Leveraging this capability, BEMT is employed to generate performance data across a wide range of blade designs and operating conditions, which are then used to train an MLP neural network surrogate. The sectional aerodynamic inputs are obtained from low-fidelity polar prediction tools such as NeuralFoil and supplemented, where necessary, with high-fidelity CFD results to improve physical fidelity. This study focuses on improving the computational efficiency of simulating turbine performance with high accuracy, aiming to predict the effect of morphing trailing edge design operated at the turbine blade, and achieving a desirable trade-off between structural loading (thrust coefficient, C_T) and power output (power coefficient, C_P) of the entire tidal turbine.

2. Problem Statement

While traditional modelling approaches, such as BEMT, CFD, and experimental testing, provide valuable insight into tidal turbine performance, their application to large-scale optimisation problems is often limited by computational cost, modelling assumptions, and scalability in high-dimensional design spaces. In this context, machine learning and deep learning techniques have emerged as powerful tools for modelling highly nonlinear systems and constructing efficient predictive models. Multi-layer perceptrons (MLPs) constitute one of the most widely used neural network architectures in engineering applications, owing to their structural simplicity, strong nonlinear approximation capability, and computational efficiency, making them particularly suitable for surrogate modelling and performance prediction tasks driven by large volumes of simulation data. Li et al. [3] employed an MLP to replace computationally expensive CFD simulations in the optimisation of rotorcraft aerofoils under ultra-low Reynolds number conditions, enabling efficient prediction of lift-to-drag. Espinosa Bárcenas et al. [4] developed an MLP–based surrogate model to approximate aerodynamic wing performance, enabling rapid and accurate performance prediction for design and optimisation tasks. Zhou et al. [5] developed an MLP-based model to directly predict aerofoil lift coefficients from geometric coordinates and flow conditions, addressing the high computational cost and complex preprocessing requirements of CFD- and CNN-based aerodynamic evaluation. Li [6] proposed a recurrent MLP model to address short-term wind turbine power prediction, enabling one-step-ahead forecasting of highly fluctuating wind power outputs from wind speed/direction. Olayemi et al. [7] demonstrated that MLP can be effectively used to map flow parameters such as Mach number and angle of attack to aerodynamic lift coefficients of supercritical aerofoils, enabling fast surrogate-based evaluation for transonic aerodynamic design. Xie et al. [8] trained MLP surrogate models on CFD-based icing simulations to enable real-time prediction of wind turbine blade ice accretion mass and maximum thickness from meteorological parameters. Yilmaz and Özer [9] proposed an MLP-based adaptive pitch-angle controller to regulate wind turbine power above rated wind speed, addressing the overspeed and overloading problem by keeping generator power and rotational speed bounded under rapidly fluctuating high-wind conditions. Pelletier et al. [10] developed a multi-stage two-layer MLP power-curve model to achieve more accurate site-specific wind turbine performance assessment for underperformance detection. Abbaskhah et al. [11] used CNN- and MLP-based surrogate models to predict wind-turbine torque and thrust for blades with and without dimples across operating conditions, enabling efficient identification of the optimal pitch–wind-speed operating point.

Overall, these studies demonstrate that MLP provides a simple but highly effective data-driven modelling framework for a wide range of aerodynamic, hydrodynamic, and wind energy engineering problems. In the study, a MLP surrogate model is trained to predict the tidal turbine performance from the blade cross-section, which is assumed to be identical along the spanwise direction, with changing chord length and twist angle.

3. Methods

Building on the motivation outlined in the Introduction, this study develops a surrogate-based optimisation framework that links aerofoil cross-section geometry to turbine-level performance, as shown in Figure 3.1. An MLP is trained to learn the mapping between the blade sectional shape and the resulting power and thrust coefficients, C_P and C_T. Once trained, the surrogate model enables rapid performance prediction for a large number of candidate designs. This provides an efficient alternative to repeated high-cost CFD simulations during the early design stage. To construct the training dataset, sectional aerofoils are generated using a Bezier-curve representation controlled by a set of control points. This parameterisation allows both conventional profiles, such as NACA 4-digit aerofoils, and more flexible non-standard geometries to be produced efficiently. For each generated aerofoil, 2D aerodynamic polars are required as inputs to the BEMT model. These polars include lift and drag coefficients over a range of angles of attack and directly affect the accuracy of the predicted turbine performance. In this work, NeuralFoil is used to generate the sectional polars due to its high computational efficiency and its accuracy compared with established low-fidelity tools [12]. The polar data are then coupled with BEMT to calculate the corresponding C_P and C_T, which are used as supervised labels for surrogate model training.

After training, the MLP surrogate model is used to evaluate an extended design space that includes irregular aerofoil geometries and trailing-edge morphing configurations. The optimisation is performed through an iterative screening process. In the first iteration, candidate aerofoils are generated through random sampling to explore the design space. In the following iterations, the sampling strategy combines random exploration with local refinement. A proportion of the candidates is still generated randomly to maintain diversity, while the remaining candidates are produced by applying small mutations to the champion aerofoil selected from the previous iteration. This allows the framework to explore new regions while gradually improving promising designs. For each candidate, the aerofoil coordinates are converted into the same Bezier control-point vector used by the surrogate model. The trained MLP then predicts C_P and C_T over the selected TSR range. A hard filtering step is applied to remove candidates with unrealistic predicted values, which controls C_P between 0.300 and 0.593, C_T between 0.200 and 1.200. The remaining valid candidates are assessed using a multi-objective selection process, where C_P is maximised and C_T is minimised. Pareto-front analysis is used to identify non-dominated solutions, and knee points are selected to represent balanced trade-offs between power output and thrust loading at each TSR. In addition, a champion aerofoil is selected based on the average score C_P-0.2C_T across the full TSR range. This champion is then used as the parent geometry for the next iteration. The process continues until the improvement becomes negligible or the maximum number of iterations is reached. The final best-performing aerofoil is then exported for further aerodynamic validation using higher-fidelity simulations.

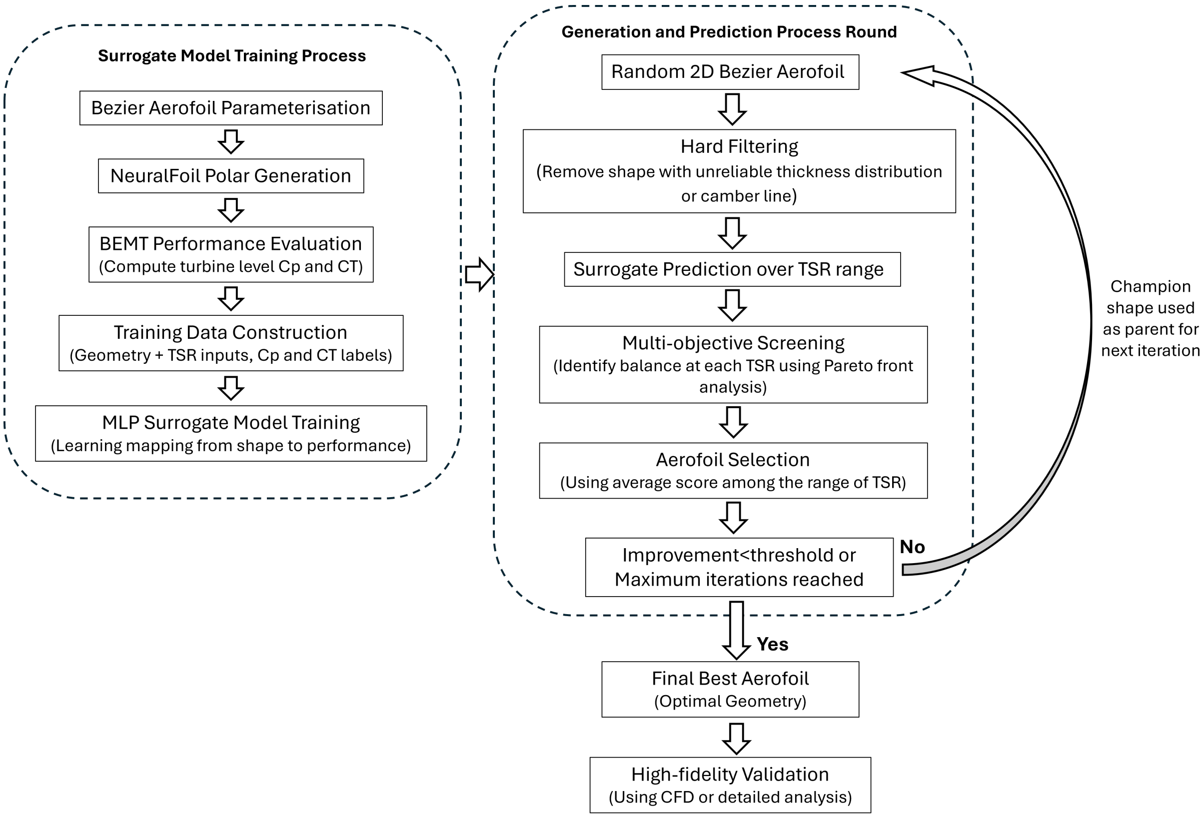

Figure 3.1: Surrogate-based iterative optimization framework, including the surrogate model training and aerofoil selection process.

3.1 Aerofoil Parameterization

To support the MLP training framework introduced in the previous section, a large population of aerofoil cross-sections is required for dataset generation and neural network training. In this study, aerofoil geometries are parameterised using a Bezier-curve formulation, which enables efficient sampling of both conventional profiles and more flexible, non-standard shapes for subsequent performance evaluation. The aerofoil surface is represented through a set of control points that define the thickness and camber line, providing a compact and consistent description of geometry for learning and optimisation.

This decomposition provides a clear and physically meaningful representation of shape variation, allowing aerofoil families to be generated by systematically adjusting camber and thickness while maintaining smooth surface geometry. Compared to traditional point-wise aerofoil representations, the Bezier control point formulation offers superior smoothness, parametric differentiability, and geometric compactness, making it particularly well-suited for optimisation, morphing, and generative design applications.

3.2 Blade Element Momentum Theory (BEMT)

The Blade Element Momentum Theory (BEMT) offers a reliable approach for calculating dynamic performance parameters, such as thrust and power, while accounting for the rotational motion of 2D cross-sections.

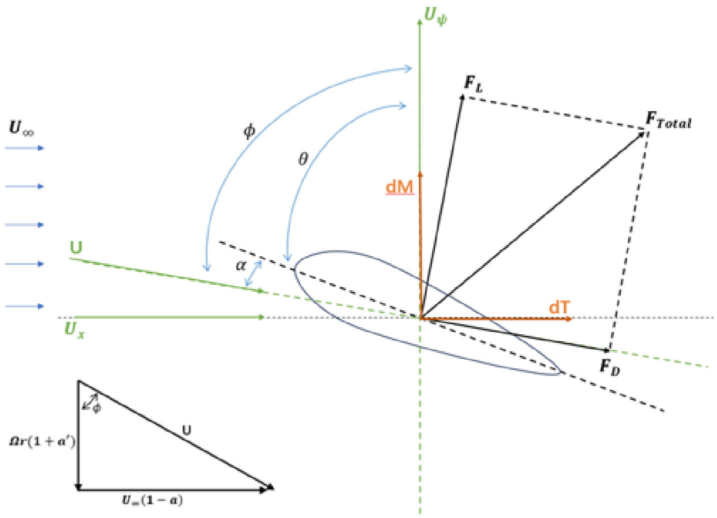

Figure 3.2: Velocity and forces components due to the free stream acting on a rotating blade cross-section.

The relationship between the velocities and forces can be described in Figure 3.2, and the relative velocity caused by the free-stream (U_∞, U_X) and turbine rotation (U_ϕ) can be presented in terms of the induced factors (a, a').

$$U_\phi = \omega ra' + \omega r = \omega r(1 + a') \tag{1}$$

$$U_X = U_\infty(1 - a) \tag{2}$$

The relative flow angle (ϕ) is defined between the incoming free-stream velocity (U_∞) and rotor motion, which can be presented as follows.

$$\phi = tan^{-1}(\frac{U_\infty(1 - a)}{\omega(1 + a')}) \tag{3}$$

The induced factors (a, a') present the relationship between the incoming flow and turbine rotation, hence can be determined by force in axial (C_X) and tangential directions (C_ψ) which can be found by lift and drag. The solidity (σ_r) presents the fraction of the swept area of the blade element in the annular stream.

$$\frac{a}{1 - a} = \frac{\sigma_r}{4sin^2\phi}, \quad \frac{a'}{1 + a'} = \frac{\sigma_r}{4sin\phi cos\phi} \tag{4}$$

$$\sigma_r = \frac{Nc(r)}{2\pi r} \tag{5}$$

The following loss factors for hubs (F_{hub}) and tips (F_{tip}) equations show the simple expression developed by Prandtl [13].

$$F_{tip} = \frac{2}{\pi}arccos(exp(-f_{tip})), \quad f_{tip} = \frac{N}{2}\frac{R_{tip} - r}{r|sin\phi|} \tag{6}$$

$$F_{hub} = \frac{2}{\pi}arccos(exp(-f_{hub})), \quad f_{hub} = \frac{N}{2}\frac{r - R_{hub}}{R_{hub}|sin\phi|} \tag{7}$$

$$F = F_{tip}F_{hub} \tag{8}$$

Hence, a and a' can be presented as follows:

$$a = [\frac{4Fsin^2\phi}{\sigma_r C_x} + 1]^{-1} \qquad a' = [\frac{4Fsin\phi cos\phi}{\sigma_r C_\psi} - 1]^{-1} \tag{9}$$

Finally, C_T and C_P can be found as

$$C_T = 4a(1-a), \qquad C_P = 4a(1-a)^2 \tag{10}$$

3.3 Machine Learning and Surrogate Model

Building on the BEMT procedure in Section 2.2, the present study constructs a supervised learning dataset in which turbine-level performance outputs (C_T, C_P) are generated efficiently for a large number of Bezier-parameterised aerofoils. For each aerofoil, sectional lift and drag polars (C_L, C_D) are first obtained using NeuralFoil, an open-source surrogate model for rapid two-dimensional aerodynamic prediction. These polars are then coupled with BEMT to compute the corresponding turbine performance coefficients, which serve as training labels for the neural network surrogate. To reduce the computational cost associated with repeated BEMT evaluations during design exploration and optimisation, a neural-network-based surrogate model is developed to directly predict C_T and C_P from the aerofoil geometry and operating condition. The input features consist of the Bezier control-point representation of the aerofoil cross-section, combined with the tip-speed ratio (TSR). In this work, the blade is assumed to employ an identical sectional aerofoil shape along the span, while the chord length and twist angle vary with radius as prescribed in the baseline blade design (i.e., sectional performance is evaluated consistently through the same BEMT formulation).

Aerofoil geometries are parameterised using a constrained Bezier representation of the camber line and thickness distribution (Section 2.1). Prior to dataset construction, each sampled aerofoil is represented by a fixed set of Bezier control points to ensure a consistent input dimension across all training samples. Specifically, the camber line is described by N_c control points and the thickness distribution by N_t control points, with both coordinates retained. The resulting geometric parameter vector, therefore, consists of $\left\{x_i^{(c)}, y_i^{(c)}\right\}_{i=1}^{N_c}$ for camber and $\left\{x_j^{(t)}, y_j^{(t)}\right\}_{j=1}^{N_t}$ for thickness. In the present study, a total of 48 geometric parameters is used, corresponding to the concatenation of all camber and thickness control-point coordinates. For each training sample, the surrogate input vector is constructed by concatenating the operating TSR with the geometric parameters in a fixed and deterministic order. The input feature vector (X) is defined as:

$$X = \left[\mathrm{TSR}, x_1^{(c)}, y_1^{(c)}, \ldots, x_{N_c}^{(c)}, y_{N_c}^{(c)}, x_1^{(t)}, y_1^{(t)}, \ldots, x_{N_t}^{(t)}, y_{N_t}^{(t)}\right] \tag{11}$$

Which ensuring a consistent mapping between the geometric representation and the surrogate model inputs. This ordering is strictly enforced during dataset construction and model training to guarantee reproducibility and interpretability. Since the surrogate model is designed to learn a smooth nonlinear mapping from aerofoil geometry and TSR to turbine performance coefficients. Gaussian error linear unit (GELU), as a non-linear activation function, is more suitable than the rectified linear unit (ReLU) for its smoother activation behaviour and improved ability to represent subtle variations in the design space [14].

The surrogate model is implemented as a feed-forward multi-layer perceptron (MLP) with a two-branch architecture, as shown in Figure 3.3. In total, the network consists of 3 fully connected layers after the initial branch-wise feature projection. The TSR branch maps the scalar TSR input into a 16-dimensional latent representation, while the geometry branch projects the 48-dimensional geometric parameter vector into a 64-dimensional latent representation. These latent features are then concatenated to form an 80-dimensional fused feature space. The combined embedding is subsequently passed through a shared regression head composed of 2 additional fully connected layers, in which the fused feature vector is first reduced to a 32-dimensional hidden representation and then mapped to the predicted thrust and power coefficients. This architecture enables the model to capture the coupled influence of operating condition and blade geometry on aerodynamic performance, while maintaining a moderate network depth that balances expressive capability and generalisation. Such a design is well-suited to the smooth, nonlinear relationships characteristic of BEMT-based aerodynamic responses. The surrogate neural network architecture mirrors the structure of the underlying physics model. The geometric parameters and the operating TSR are processed by separate fully connected sub-networks, allowing geometry-dependent and condition-dependent latent features to be extracted independently. These latent embeddings are subsequently concatenated and passed through a shared regression head to predict C_P and C_T. Such a modular design reflects the physical separation between blade shape and flow condition inherent in BEMT, while enabling efficient learning of their coupled influence on rotor-level aerodynamic performance.

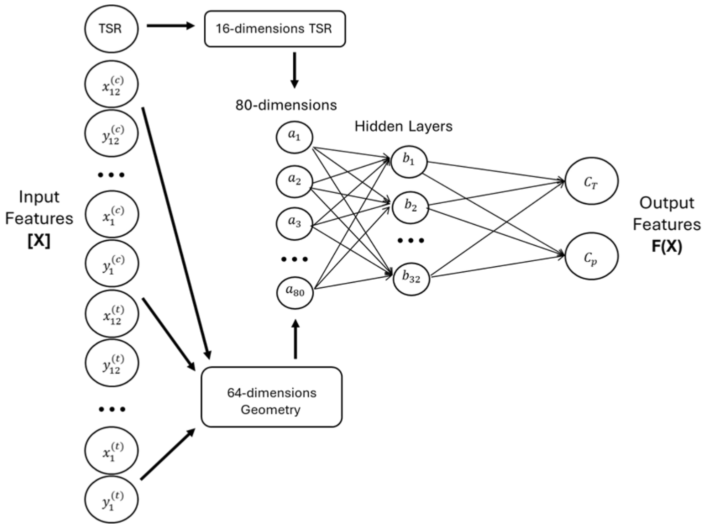

Figure 3.3: Architecture of the selected MLP Surrogate Model.

3.4 Computational Fluid Dynamics (CFD)

The 3D steady CFD simulations are conducted to validate the surrogate model predictions, which enable the investigation of blade deformation and bending behaviour under realistic aerodynamic loading, providing insight into the coupled fluid–structural response of the morphing blade. A 3D model without a support structure is built and presented below in Figure 3.4, with the rotation radius of 5 m and hub radius of 0.89 m.

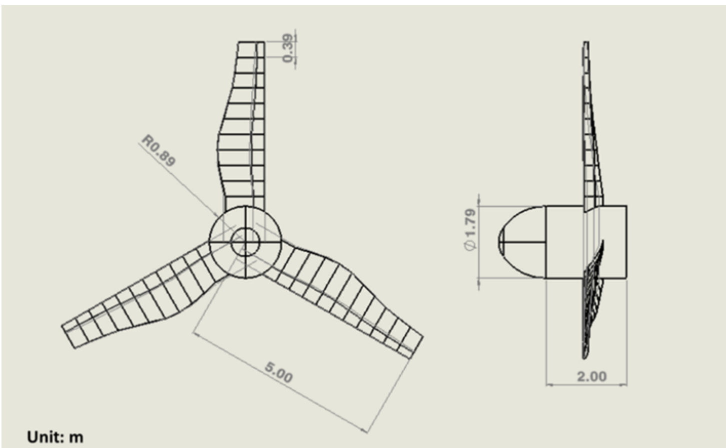

Figure 3.4: Dimensions of the provided 3-bladed tidal turbine model without a support structure, with the swirling area of 78.5 m^2.

The 3D steady CFD simulation was carried out under the Reynolds number of 500,000, using the blade tip chord length (0.658 meters) as the characteristic length, and TSR range over 2 to 6. The k-omega SST is selected as the turbulent model and the y^+ is limited to lower than 5. The 3D steady CFD model is validated by the exist CFD results and experimental data [15] as shown below in Figure 3.5.

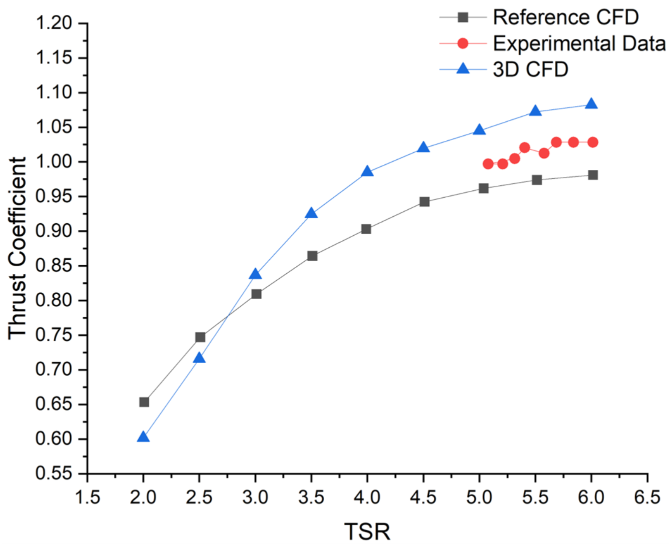

Figure 3.5: Comparison of the steady CFD simulation with the reference CFD and the experimental data [15] for validation.

At low to moderate TSR values (between 2 and 4), the 3D steady CFD results show good agreement with the reference CFD data [15], which is performed using the Reynolds Stress Model (RSM). The predicted C_T are close to the reference values, indicating that the present numerical setup can capture the main loading characteristics of the turbine. At higher TSR values (between 4.5 and 6), the present CFD results slightly overpredict the C_T compared with both the reference CFD and experimental data, as shown in Table 3.1. However, this marginal gain in precision (Table 3.1) for reference CFD results comes at the cost of significantly higher computational demand, and a more complex setup for the RSM model compared to the faster, more robust $k - \omega$ SST approach.

Table 3.1: Comparison of the relative error between the 3D steady CFD and the reference CFD against the experimental data at high TSR conditions.

TSR	Experimental Data [15]	Relative Error (3D Steady CFD)	Relative Error (Reference CFD [15])
5	0.9969	4.84%	3.52%
5.5	1.0205	5.09%	4.57%
6	1.0284	5.27%	4.59%

3.5 Multi-objective Optimisation (Pareto-front analysis)

Before final selection, both C_P and C_T are normalised within the Pareto set to allow a fair comparison between the two objectives. The normalised values are expressed as:

$$C_{P,\mathrm{nor}} = \frac{C_P - C_{P,\mathrm{min}}}{C_{P,\mathrm{max}} - C_{P,\mathrm{min}} + \epsilon}, \quad C_{T,\mathrm{nor}} = \frac{C_T - C_{T,\mathrm{min}}}{C_{T,\mathrm{max}} - C_{T,\mathrm{min}} + \epsilon} \tag{12}$$

where ϵ is a small numerical constant used to avoid division by zero. After normalisation, the knee point is selected as the solution closest to the ideal point. This ideal point represents the maximum normalised power coefficient and the minimum normalised thrust coefficient. Therefore, the selected knee point provides a balanced compromise between aerodynamic efficiency and load mitigation, rather than favouring only one objective. In addition to the Pareto-based selection, a scalar balance score is introduced to monitor the optimisation progress and define the stopping criterion. The score Sis calculated from the mean weighted performance across all TSR values:

$$S = \frac{1}{N_{\mathrm{TSR}}} \sum_{\mathrm{TSR}} (C_P - \lambda C_T) \tag{13}$$

where N_{TSR} is the number of TSR conditions. The weighting factor λ controls the penalty applied to C_T. In this study, λ is set to 0.2 to penalise excessive thrust loading while keeping power output as the primary optimisation objective, which is decided by processing a sensitivity test. This means that the optimisation favours designs with high C_P, but also gives preference to those that reduce C_T without causing a large power penalty.

The optimisation is performed in an iterative closed-loop manner. In each iteration, newly generated aerofoil geometries are evaluated by the trained surrogate model. Part of the new population is produced by perturbing the best aerofoil selected from the previous iteration. The remaining candidates are sampled randomly within the predefined geometric constraints. This strategy combines local refinement with global exploration. It allows the framework to improve promising designs while still searching for new regions of the design space. Overall, the method enables rapid screening of candidate aerofoils by combining Bezier-based parameterisation, surrogate-model prediction, Pareto-based selection, and iterative optimization.

4. Results and Discussion

This Chapter begins with the sensitivity test of the MLP surrogate model. The surrogate model predicted performance is verified using 3D CFD results and reference experimental data [15]. The limitation of the current surrogate model is discussed. Then the shapes selected by the optimisation framework, including the surrogate model and the multi-objective analysis, are presented with improved behaviour.

4.1 Sensitivity Test and Validation of the Surrogate Model

To examine how surrogate performance depends on network capacity, a set of singularity tests was conducted. The tests focus on two design factors: feature embedding size and regression-head depth. In the two-branch MLP, the TSR embedding dimension varied over 1, 4, 8, 16, and 32, while the geometry embedding dimension varied over 32, 48, 64, 128, and 256. These choices control how much information each branch can encode. Smaller embeddings restrict the feature representation and may lead to underfitting. Larger embeddings increase model capacity but can amplify overfitting and reduce training stability if the capacity is excessive. For each embedding combination, three regression-head configurations were evaluated: a shallow head with one hidden layer (32), a baseline head with two hidden layers (64, 32), and a deeper head with four hidden layers (128, 64, 64, 32). All other training settings were kept unchanged, and the same training–validation split was used in all runs to ensure a fair comparison. The trained models were then applied to predict the thrust and power coefficients of the baseline Wortmann FX 63-137 aerofoil, as the blade cross-section, across the investigated TSR range. Model accuracy in C_T and C_P, together with convergence behaviour during training, was compared across the tested architectures. This study clarifies the effects of embedding size and network depth on prediction quality and supports the selection of a practical surrogate architecture that balances accuracy and generalisation.

A. Effect of embedding dimension on surrogate model predictions.

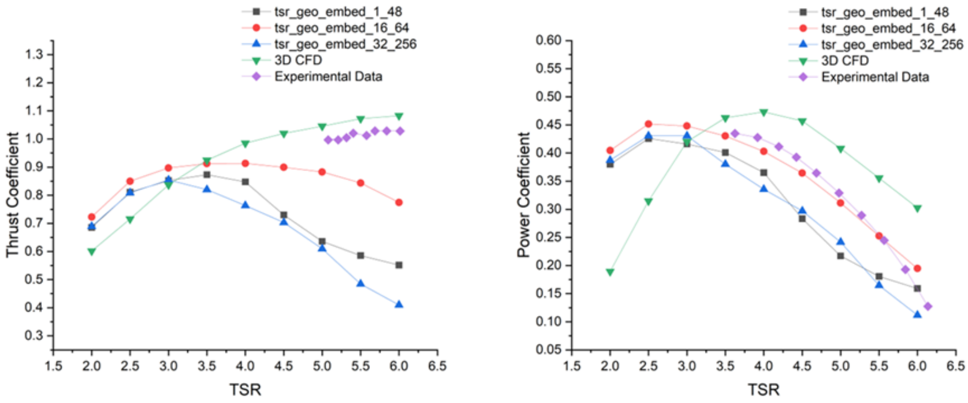

Figure 4.1: Effect of embedding dimensions on surrogate predictions of thrust (left) and power (right) coefficients within changing TSR for the Wortmann FX 63-137 aerofoil, compared with reference 3D CFD and experimental data [15].

Figure 4.1 compares the predicted turbine thrust and power coefficients for the baseline Wortmann FX 63-137 aerofoil using surrogate models trained with different embedding dimensions. Two extreme cases are included: a minimal embedding setting (`tsr_geo_embed_1_48`) and a highly expanded setting (`tsrgeo_embed_32_256`). These are compared with the selected intermediate configuration (`tsr_geo_embed_16_64`), together with the reference 3D CFD and experimental data.

Overall, the selected embedding size provides the most consistent agreement with the reference trends. It reproduces the peak region and the subsequent decay in C_P, and it gives a reasonable variation of C_T across the TSR range. The minimal embedding case shows limited representation capacity and deviates more strongly at moderate to high TSR. The very large embedding case does not improve the agreement and shows larger divergence at high TSR, suggesting reduced generalisation. These results indicate that moderate embedding dimensions offer a practical balance between feature representation and model robustness for this surrogate prediction task.

B. Effect of regression-head depth on surrogate model predictions.

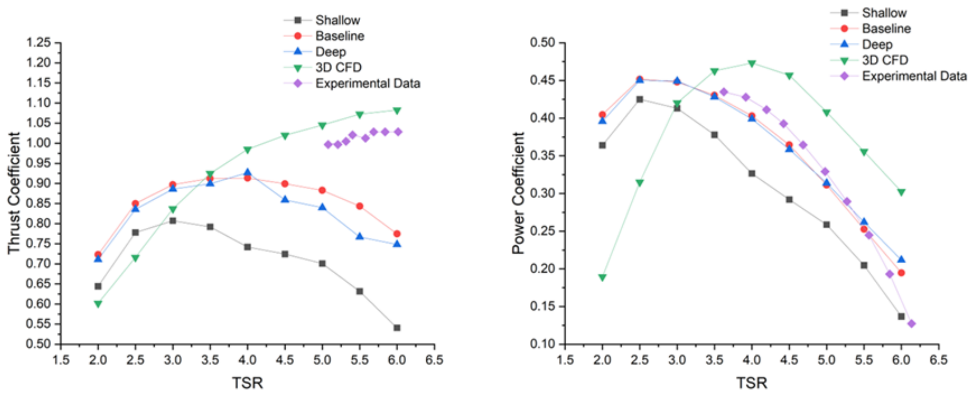

Figure 4.2: Effect of regression-head depth (shallow/baseline/deep) on surrogate predictions of thrust (left) and power (right) within changing TSR for the Wortmann FX 63-137 aerofoil, compared with reference 3D CFD and experimental data [15].

Figure 4.2 compares the predictions obtained with different regression-head depths. A shallow model, the selected baseline model, and a deeper model are shown together with the reference CFD and experimental results. The comparison shows that network depth has a clear influence on both prediction accuracy and robustness. The shallow model underpredicts the performance, especially at moderate to high TSR, and fails to follow the peak region accurately. This suggests that its representational capacity is insufficient to model the nonlinear aerodynamic response. Increasing the depth improves the overall trend prediction in both C_P and C_T. The baseline-depth model captures the peak region and the subsequent decline more consistently than the shallow case.

However, the deeper model does not provide a clear improvement over the baseline. Although it remains capable of following the overall trend, its deviations become slightly larger at higher TSR. This indicates that increasing depth beyond a moderate level does not necessarily improve predictive performance for the present task. Overall, the results suggest that a moderate head depth is sufficient and offers a better trade-off between accuracy and generalisation.

Based on the sensitivity study, the selected MLP architecture provides a practical balance between model capacity, predictive consistency, and computational efficiency. It is therefore adopted for the subsequent validation study.

The training history, Figure 4.3, shows a stable convergence behaviour for both training and validation datasets. The validation loss decreases rapidly in the early stage and remains relatively stable afterwards, indicating that the adopted network configuration provides a reasonable balance between fitting capability and generalisation performance. Although the C_T rediction remains more challenging than C_P, both targets are learned effectively without clear signs of severe overfitting.

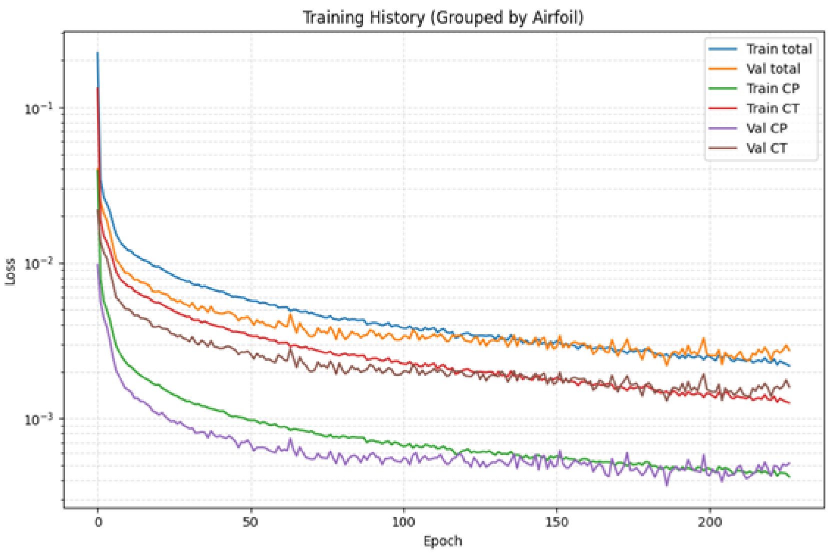

Figure 4.3: Training history of the MLP surrogate model under grouped-by-airfoil data splitting, including total loss and the individual losses for C_P and C_T on the training and validation sets.

After selecting suitable architecture through the sensitivity study, the surrogate model was validated against reference CFD and experimental data. The purpose of this validation is to assess whether the trained model can reproduce the main performance trends observed in higher-fidelity and experimental references. Figure 4.4 includes the turbine performance of the Wortmann FX 63137 cross-sectional model, from the BEMT-based surrogate prediction and reference data from both CFD and experiments.

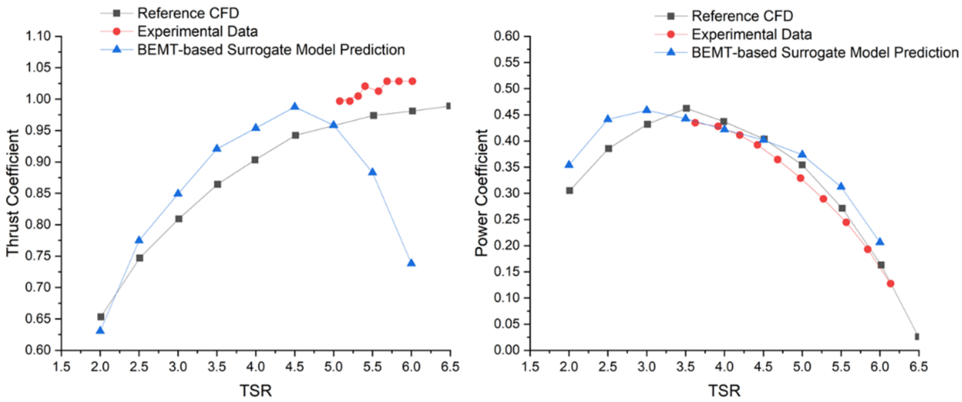

Figure 4.4: Comparison of thrust (left) and power (right) coefficients obtained from the proposed surrogate model, CFD and experimental data from the reference data [15].

For the C_P, the surrogate model follows the reference trend well across most of the TSR range, with particularly good agreement near the peak region. For the C_T, the overall trend is also captured, although larger deviations appear at high TSR. In this region, the model predictions become more conservative and deviate more clearly from reference data. These discrepancies are mainly linked to the assumptions embedded in the BEMT-based training data, rather than limitations of the regression model itself. Despite this, the surrogate model accurately represents the overall behaviour of both C_P and C_T across a wide range of operating conditions. The agreement with CFD and experimental data is strongest near the optimal operating region, TSR between 3.5 to 4.5. For higher TSR values (TSR > 5.0), the discrepancies with experimental data become pronounced due to the inherent limitations of Blade Element Momentum Theory (BEMT), as well as the underlying assumptions embedded in the BEMT-based training data. Overall, the surrogate model provides a good balance between predictive accuracy and computational efficiency, for TSR < 5, making it suitable for rapid performance evaluation and design space exploration.

4.2 Current Limitations of the Surrogate Model

Although the surrogate model shows favourable performance in terms of efficiency and robustness, several limitations remain.

First, the applicability of the model is constrained by the adopted geometric parameterisation. In this study, a fixed number of Bezier control points is used to represent the camber line and thickness distribution of the aerofoil. The number of control points is selected to reduce the input dimensionality of the model while still providing sufficient flexibility to accurately represent the aerofoil geometry. This approach is effective for capturing the dominant low-order shape features, but it may be less suitable for representing highly localised curvature changes or more complex leading-edge geometries. Increasing the number of control points could improve geometric flexibility, but this would also increase input dimensionality and data requirements.

Second, the reliability of surrogate predictions depends on whether the queried aerofoil lies within the distribution covered by the training dataset. If a new geometry differs significantly from the training set, prediction uncertainty may increase. In the present study, the Bezier fitting error is recorded as an auxiliary indicator of geometric consistency, which provides a practical way to identify possible inputs which are out of the distribution. However, this indicator has not yet been fully incorporated into the training process or into an explicit uncertainty quantification framework.

Third, the surrogate model is trained on aerodynamic data generated from 2D aerofoil polars coupled with BEMT. As a result, its predictive behaviour inherits the assumptions and limitations of the underlying BEMT framework. These include quasi-steady flow assumptions and the simplified treatment of wake behaviour. Consequently, the current surrogate model is mainly suited to steady-state performance prediction and does not explicitly account for unsteady aerodynamic effects, three-dimensional flow structures, or complex wake interactions.

Despite these limitations, the surrogate model remains a useful intermediate tool between low-fidelity and high-fidelity analysis. It does not aim to replace CFD or experimental testing. Instead, it provides an efficient means of rapid screening and preliminary optimisation. High-fidelity CFD simulations and experiments can then be applied to a smaller set of promising designs identified by the surrogate model. In this way, the surrogate supports a hierarchical and computationally efficient design framework.

4.3 Multi-objective Analysis

The multi-objective optimisation analysis is used to identify aerofoil candidates with improved performance. The results are represented by a Pareto front, which defines the set of optimal trade-off solutions within the feasible design space. Each point on the Pareto front corresponds to a non-dominated design, where no objective can be improved without degrading another. Rather than providing a single optimal solution, the Pareto front offers a range of design options. This allows informed selection based on additional considerations, such as structural performance or operational requirements.

The multi-objective optimisation framework was applied to 20,000 random aerofoil geometries, each parametrically generated and evaluated across 9 discrete TSRs ranging from TSR = 2.0 to TSR = 6.0. For each aerofoil, the C_T and C_P were computed, resulting in a comprehensive performance database capturing both aerodynamic efficiency and loading characteristics over a wide operational envelope. In the following Figures 4.5 to 4.7, the grey points denote all evaluated aerofoil geometries, while blue markers indicate non-dominated solutions forming the Pareto front. The red cross highlights the selected knee point for each TSR, representing a compromise solution where further improvements in C_P would result in a disproportionately large increase in C_T.

At low TSRs (Figures 4.5) corresponding to TSR = 2.0, 2.5, and 3.0, the Pareto fronts exhibit a steep slope, indicating that modest improvements in C_P require relatively large increases in C_T. In this regime, the flow is characterised by higher angles of attack and stronger loading, making aerodynamic efficiency gains increasingly expensive in terms of thrust. The distribution of feasible designs at low TSRs is also relatively compressed in CP, suggesting that the potential for efficiency improvement through airfoil shape variation is limited under these operating conditions. Consequently, knee points identified in this regime tend to lie closer to the mid-range of the Pareto front, reflecting a conservative balance between performance and load.

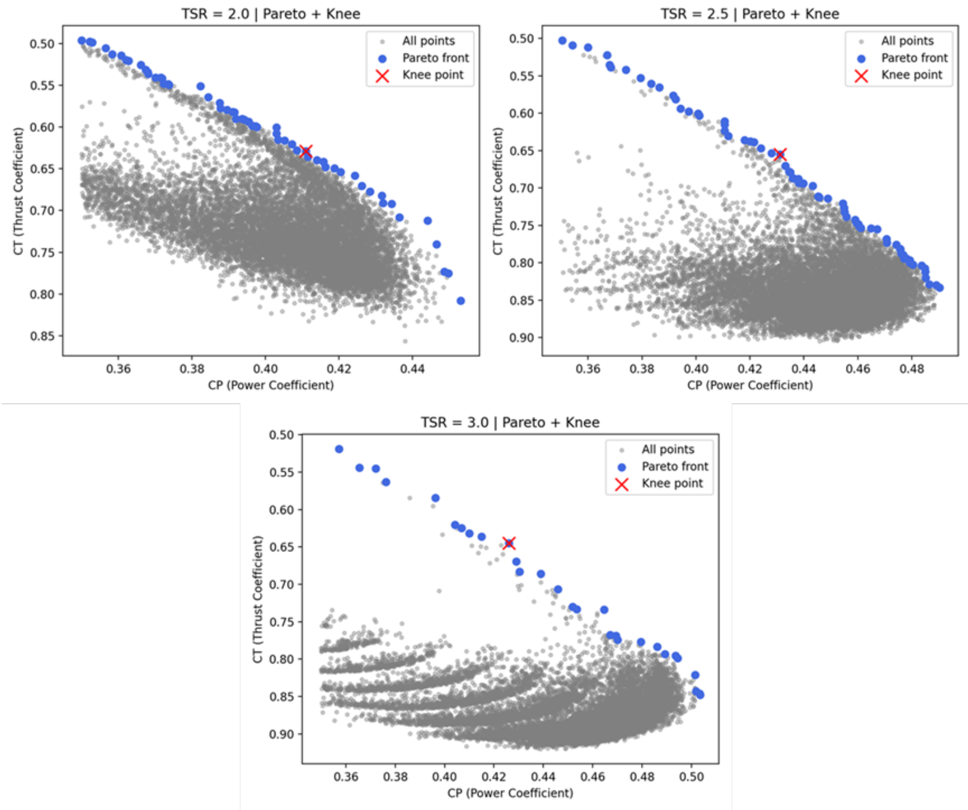

Figure 4.5: Pareto fronts and representative knee-point solutions in the CP–CT objective space for TSR = 2.0, 2.5, and 3.0.

In the intermediate TSR range (Figure 4.6), the Pareto fronts become smoother and more extended in C_P, indicating a more favourable trade-off between efficiency and thrust. The dispersion of feasible designs increases, particularly in the high C_P region, implying that airfoil geometry plays a stronger role in determining aerodynamic performance. This regime exhibits a pronounced "knee-like" curvature in the Pareto fronts, where a transition occurs from relatively efficient C_P gains at modest C_T penalties to rapidly increasing C_T for marginal C_P improvements. The knee points selected in this TSR range consistently align with this transition region, suggesting that these operating conditions offer the most meaningful compromise between energy capture and structural loading.

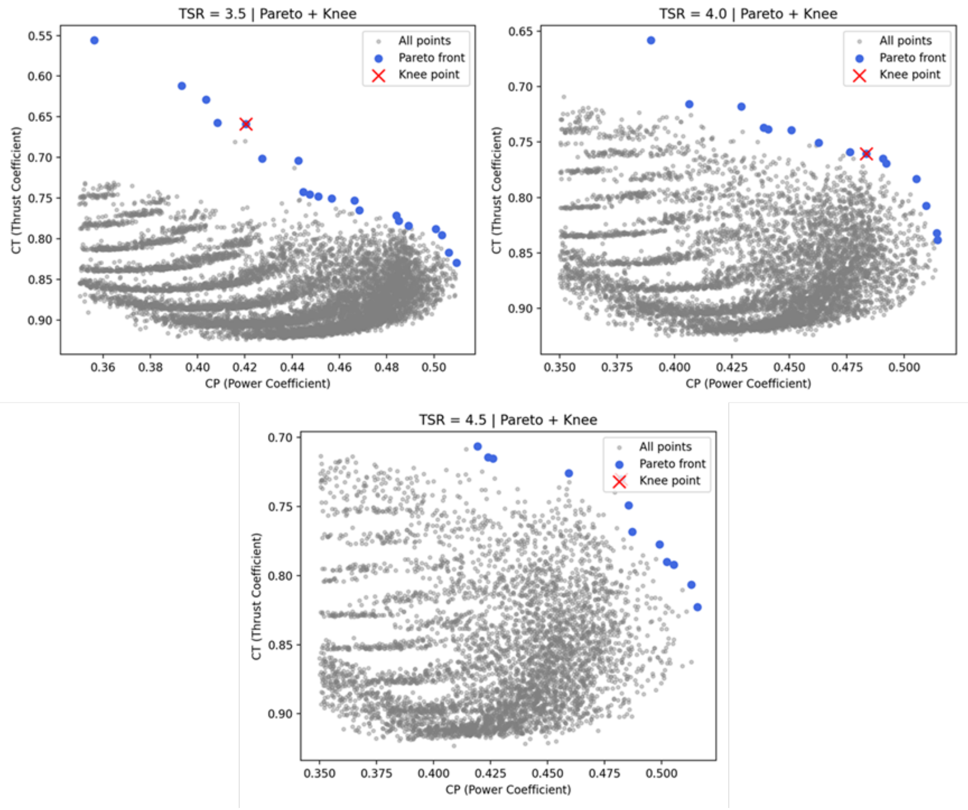

Figure 4.6: Pareto fronts and representative knee-point solutions in the CP–CT objective space for TSR = 3.5, 4.0, and 4.5.

At high TSRs (Figure 4.7), the Pareto fronts shift noticeably toward higher C_P values, reflecting improved aerodynamic efficiency due to reduced angles of attack and more favourable flow alignment. However, this efficiency gain is accompanied by a widening spread in C_T, particularly near the high C_P end of the front. In this regime, the Pareto fronts become flatter in C_P near their upper extremity, indicating diminishing returns: substantial increases in CT are required to achieve relatively small improvements in C_P. As a result, knee points tend to be located slightly away from the absolute maximum C_P solutions, favouring designs that avoid excessive thrust loading while still achieving near-optimal efficiency.

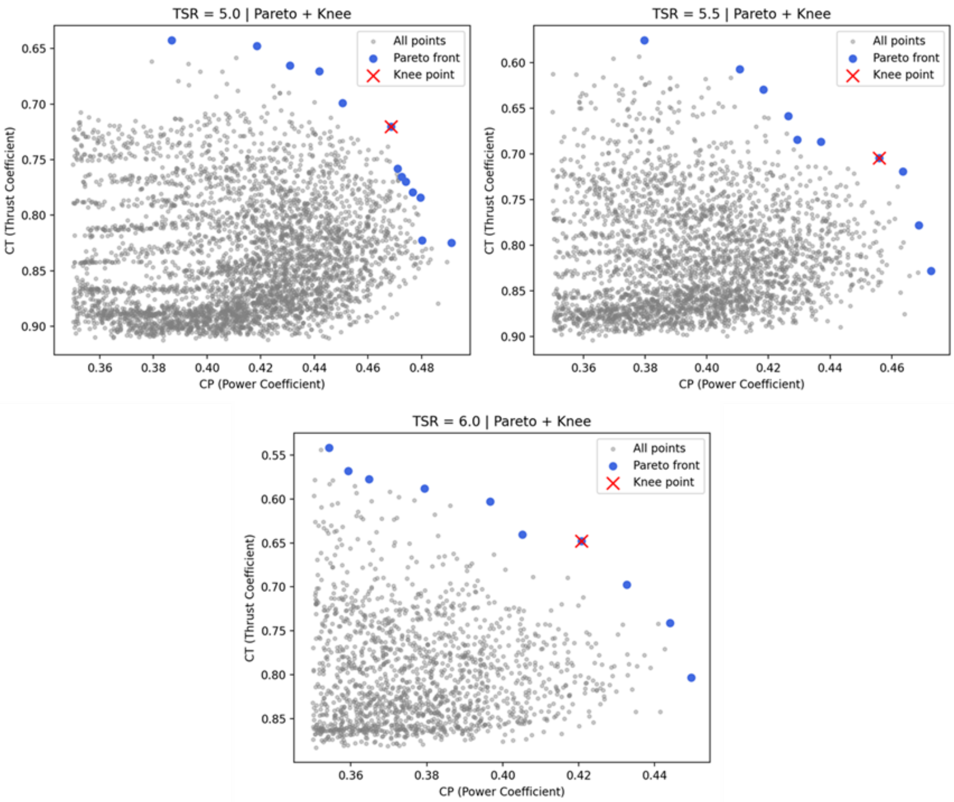

Figure 4.7: Pareto fronts and representative knee-point solutions in the CP–CT objective space for TSR = 5.0, 5.5, and 6.0.

The knee points identified in **Figures 4.5 to 4.7** represent balanced design solutions for each TSR. These points lie in regions where the trade-off between C_T and C_P changes rapidly. Beyond these points, further gains in power become increasingly costly in terms of thrust loading. It should be noted that the knee points do not correspond to the maximum C_P. Instead, they reflect a practical balance between aerodynamic efficiency and acceptable loading, which is important for structural integrity and long-term reliability. The consistent appearance of these points across all TSRs confirms the robustness of the multi-objective framework.

The results also show that no single aerofoil achieves optimal performance across all TSRs, mainly due to the limited coverage of the current design space. This suggests that improved designs may still be found by exploring a wider range of aerofoil geometries. Based on this observation, a set of candidate aerofoils is selected from the optimisation process. These shapes represent promising designs that balance power performance and loading and are further analysed in the following section.

4.4 Aerofoil Shape Optimisation

With the optimisation framework illustrated, the 3D turbine performance with the random aerofoil shapes is predicted through the preliminary trained surrogate model. The shapes are scored to search for the one that achieves the balance between loading and power output among the TSRs between 2 and 6.

Some of the aerofoil shapes selected (id43355, id10265, and id03808) through this optimisation process were then incorporated into the 3D turbine model and evaluated using CFD simulations. The following section presents the corresponding simulation results, which provide further insight into the aerodynamic and hydrodynamic effects of the selected optimised designs.

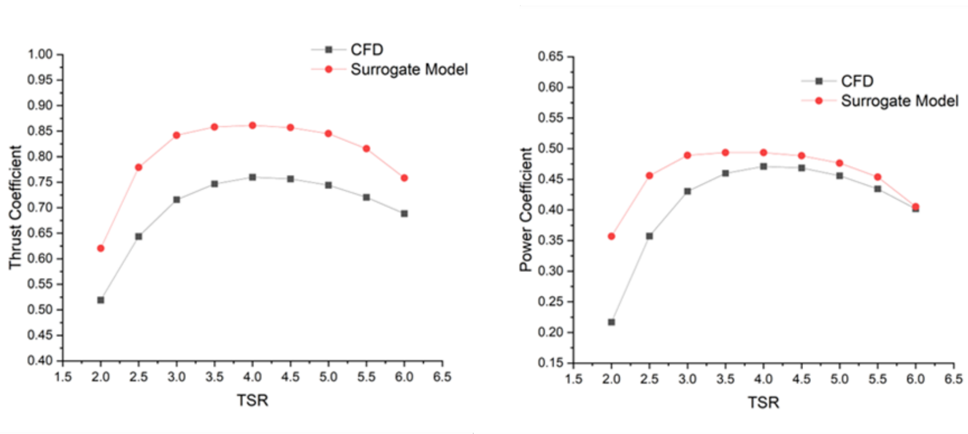

Figure 4.8: Comparison of C_T and C_P between the surrogate model and 3D CFD predictions for aerofoil id43355 over a range of (TSR).

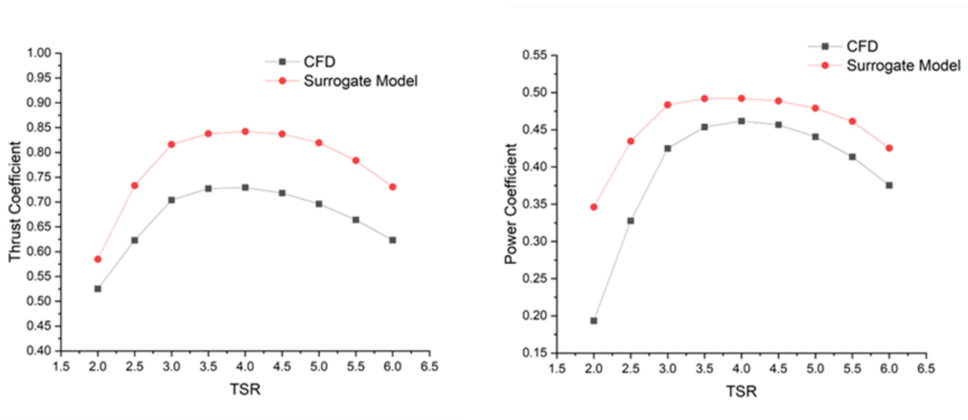

Figure 4.9: Comparison of C_T and C_P between the surrogate model and 3D CFD predictions for aerofoil id10265 over a range of (TSR).

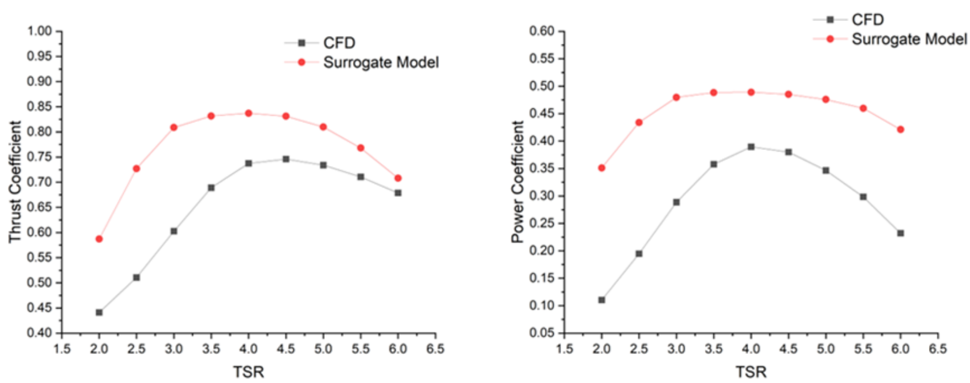

Figure 4.10: Comparison of C_T and C_P between the surrogate model and 3D CFD predictions for aerofoil id03808 over a range of (TSR).

Figures 4.8 to 4.10 compare the thrust coefficient (C_T) and power coefficient (C_P) predicted by the surrogate model with the 3D CFD results for three representative aerofoil geometries: id43355, id10265, and id03808. For all three cases, the surrogate model successfully captures the overall trend of both C_T

and C_P across the full range of TSR values. The predicted curves show the same increase at low TSR, the peak region around TSR of 3.5 to 4.5, and the subsequent decrease at higher TSR.

However, a consistent overprediction is observed across all geometries. The surrogate model predicts higher values of both C_T and C_P compared with the CFD results across most TSR values. This is caused by using BEMT-based datasets in the training process, which introduce inherent modelling errors. Nevertheless, the surrogate model remains effective for rapid screening, as it can efficiently identify aerofoil geometries with improved performance trends.

Overall, the surrogate model demonstrates good predictive capability for different aerofoil geometries. It accurately reproduces the main performance trends and peak locations, although it shows a tendency to overpredict the magnitude of both C_T and C_P, especially for more complex geometries and higher TSR values.

Table 4.1: Improvement in C_T and C_P for the turbine with different blade cross-section shapes (id10265, id14335, and id03808), compared to the original turbine with Wortmann FX 63-137.

	id10265		**id14335**		**id03808**	
TSR	C_T (%)	C_P (%)	C_T (%)	C_P (%)	C_T (%)	C_P (%)
2	-12.7%	2.1%	-13.8%	14.5%	-26.7%	-41.7%
2.5	-13.0%	4.0%	-10.1%	13.4%	-28.7%	-38.1%
3	-15.9%	1.1%	-14.4%	2.4%	-28.0%	-31.3%
3.5	-21.4%	-2.0%	-19.2%	-0.6%	-25.5%	-22.7%
4	-25.9%	-2.4%	-22.9%	-0.5%	-25.1%	-17.7%
4.5	-29.6%	-0.1%	-25.8%	2.5%	-26.8%	-16.8%
5	-33.4%	8.0%	-28.8%	11.7%	-29.8%	-15.1%
5.5	-38.1%	16.3%	-32.8%	22.1%	-33.7%	-16.1%
6	-42.4%	24.1%	-36.4%	32.8%	-37.3%	-23.3%

The percentage changes listed in Table 4.1 are consistent with the trends shown in Figure 4.11. The table provides a direct quantitative comparison with the baseline turbine, while the figure shows how these differences vary across the TSR range. For all three selected aerofoils, the reduction in C_T is clear over the full operating range, confirming that the new blade cross-sections can effectively reduce turbine loading. At the same time, the C_P results show that id14335 and id10265 maintain power output close to, or slightly above, the baseline in the main operating region, especially around TSR = 3 to 4. This agrees with the values in Table 4.1, where both aerofoils show clear improvements in C_P at higher TSR while still achieving substantial reductions in C_T.

Among the three candidates, id14335 and id10265 provide the best overall balance between power and loading. Their curves remain close to the baseline in C_P, but stay clearly below it in C_T. In contrast, id03808 achieves a larger reduction in C_T, but this comes with a clear loss in C_P across the full TSR range. This behaviour is also reflected in Table 4.1, where all C_P changes for id03808 are negative. Taken together, the table and the figure show that the surrogate-based screening process is able to identify blade cross-sections with improved overall performance. Although prediction errors remain in the surrogate model, the selected aerofoils still demonstrate clear benefits when evaluated using 3D CFD.

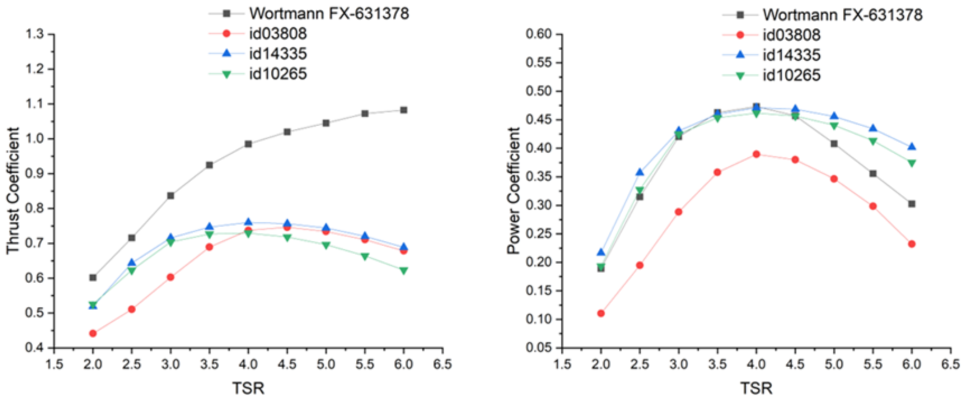

Figure 4.11: Comparison of C_T and C_P between the selected shapes with the original design.

Figure 4.11 compares the performance of the selected aerofoil geometries (id03808, id14335, and id10265) with the baseline Wortmann FX-63137 aerofoil in terms of C_T and C_P. For all three selected geometries, C_T remains lower than that of the baseline across the TSR range. This indicates a consistent reduction in hydrodynamic loading acting on the turbine. On average, a reduction of up to 29.062% in C_T is achieved while maintaining comparable power output.

The C_P results further highlight the difference between the selected designs. In the peak operating region around TSR = 3 to 4, id14335 and id10265 achieve similar peak C_P values to the baseline while maintaining reduced C_T. This demonstrates a better balance between power output and structural loading. Although id03808 also reduces C_T, it shows a lower overall C_P, making it less favourable in terms of combined performance. Overall, these results confirm that the selected aerofoil geometries provide a better trade-off between power output and loading. They also show that the surrogate model can be used effectively for preliminary design screening, even though some prediction error remains.

5. Conclusion and Future Work

In conclusion, this study presents an efficient optimisation framework for tidal turbine blade cross-section design. The framework uses an MLP surrogate model as the main prediction tool to rapidly estimate thrust and power performance, trained by low-fidelity BEMT data. By replacing repeated high-cost CFD evaluations with surrogate-based screening, the computational cost is greatly reduced. The comparison between surrogate predictions and 3D CFD results confirms the feasibility of the proposed framework for blade shape optimisation. The framework also demonstrates strong screening capability. It can iteratively identify promising blade cross-section shapes from 24 design groups, with 20,000 candidate geometries in each group, within 2 hours. The optimised shapes show improved hydrodynamic performance over the TSR range of 2–6. In particular, the selected designs achieve an average reduction of up to 29.062% in thrust coefficient while maintaining a comparable power coefficient. These results indicate that the proposed MLP-based optimisation framework can provide an effective and computationally efficient route for tidal turbine blade design.

The future work will focus on enhancing the accuracy of surrogate model prediction, especially at higher TSRs between 4.5 to 6, by ultilising high-fidelity CFD data and performing the transfer learning.

Acknowledgements

This study was funded by the University of Southampton, UK. The preliminary tidal turbine model was provided by the University of the West of England. The authors also acknowledge Kamal Djidjeli, Zhiwei Hu, and Abdessalem Bouferrouk, for their guidance.

References

[1] Simon P. Neill, Athanasios Angeloudis, Peter E. Robins, Ian Walkington, Sophie L. Ward, Ian Masters, Matt J. Lewis, Marco Piano, Alexandros Avdis, Matthew D. Piggott, George Aggidis, Paul Evans, Thomas A.A. Adcock, Audrius Židonis, Reza Ahmadian, and Roger Falconer. Tidal range energy resource and optimization – Past perspectives and future challenges. *Renewable Energy*, 127:763–778, November 2018.

[2] Fu-wei Zhu, Lan Ding, Bin Huang, Ming Bao, and Jin-Tao Liu. Blade design and optimization of a horizontal axis tidal turbine. *Ocean Engineering*, 195:106652, January 2020.

[3] Zhiqing Li, Yanbo Wang, Hongxing Zhou, and Xin Wang. Airfoil Design and Performance Optimization of Rotorcraft Robots through the Integration of LSTM-PSO and Multilayer Perceptron Technologies. In *2024 39th Youth Academic Annual Conference of Chinese Association of Automation (YAC)*, pages 1825–1832, Dalian, China, June 2024. IEEE.

[4] Oscar Ulises Espinosa Barcenas, Jose Gabriel Quijada Pioquinto, Ekaterina Kurkina, and Oleg Lukyanov. Surrogate Aerodynamic Wing Modeling Based on a Multilayer Perceptron. *Aerospace*, 10(2):149, February 2023.

[5] Jianbo Zhou, Rui Zhang, and Lyu Chen. Prediction of Airfoil Lift Coefficient Based on Multilayer Perceptron. *International Journal of Aeronautical and Space Sciences*, 26(1):13–20, January 2025.

[6] S. Li. Wind power prediction using recurrent multilayer perceptron neural networks. In *2003 IEEE Power Engineering Society General Meeting (IEEE Cat. No.03CH37491)*, pages 2325–2330, Toronto, Ont., Canada, 2003. IEEE.

[7] Department of Aeronautics and Astronautics, Faculty of Engineering and Technology, Kwara State University, Malete, Nigeria, Olalekan Olayemi, Oluwadolapo Salako, Department of Aeronautics and Astronautics, Faculty of Engineering and Technology, Kwara State University, Malete, Nigeria., Abdulbaqi Jinadu, Adebowale Obalalu, Department of Mathematical Science, Augustine University, Illara-Epe, Nigeria., Benjamin Anyaegbuna, and Department of Mechanical Engineering, Faculty of Engineering and Technology, University of Ilorin, Ilorin 240003, Nigeria. Aerodynamic lift coefficient prediction of supercritical airfoils at transonic flow regime using convolutional neural networks (CNNs) and multi-layer perceptions (MLPs). *Al-Qadisiyah Journal for Engineering Sciences*, 16(2):108–115, June 2023.

[8] Nan Xie, Qingqing Cao, Zhixiang Zeng, Kebo Ma, and Sizhun Zeng. Research on the Real-Time Prediction of Wind Turbine Blade Icing Process Based on the MLP Neural Network Model and Meteorological Parameters. *Processes*, 13(6):1910, June 2025.

[9] Ahmet Serdar Yilmaz and Zafer Özer. Pitch angle control in wind turbines above the rated wind speed by multi-layer perceptron and radial basis function neural networks. *Expert Systems with Applications*, 36(6):9767–9775, August 2009.

[10] Francis Pelletier, Christian Masson, and Antoine Tahan. Wind turbine power curve modelling using artificial neural network. *Renewable Energy*, 89:207–214, April 2016.

[11] Ahmad Abbaskhah, Hamed Sedighi, Pooria Akbarzadeh, and Ali Salavatipour. Optimization of horizontal axis wind turbine performance with the dimpled blades by using CNN and MLP models. *Ocean Engineering*, 276:114185, May 2023.

[12] Peter Sharpe and R. John Hansman. NeuralFoil: An Airfoil Aerodynamics Analysis Tool Using Physics-Informed Machine Learning, March 2025. arXiv:2503.16323 [physics].

[13] Y. El Khchine and M. Sriti. Tip Loss Factor Effects on Aerodynamic Performances of Horizontal Axis Wind Turbine. *Energy Procedia*, 118:136–140, August 2017.

[14] Dan Hendrycks and Kevin Gimpel. Gaussian Error Linear Units (GELUs), June 2023. arXiv:1606.08415 [cs].

[15] A. Mason-Jones, D.M. O'Doherty, C.E. Morris, T. O'Doherty, C.B. Byrne, P.W. Prickett, R.I. Grosvenor, I. Owen, S. Tedds, and R.J. Poole. Non-dimensional scaling of tidal stream turbines. *Energy*, 44(1):820–829, August 2012.