A Multi-fidelity, Multi-start Approach for the Aerodynamic Design of Propellers

Tao Zhang, Mark Woodgate, George Barakos
CFD Laboratory, James Watt School of Engineering, University of Glasgow, UK.
Yu Luo
School of Mechanical Engineering, Guizhou University, China.

ABSTRACT

This paper presents a new approach for aerodynamic shape design and optimisation, combining high-fidelity gradient-based optimisation with multi-fidelity surrogate-based optimisation. It exploits the gradient-based optimisation history (both function values and gradients) to train gradient-enhanced multi-fidelity surrogate models, and uses the surrogate to indicate potentially better global solutions to restart the gradient-based local search. This multi-fidelity, multi-start (MFMS) approach retains the high efficiency of gradient-based methods in searching high-dimensional design spaces and helps evade suboptimal local solutions via surrogates. A key novelty here is the introduction of a multi-fidelity neural network (MFNN), which seamlessly fuses physical models, multi-fidelity data, and gradients, providing accurate surrogate predictions on very sparse samples. The proposed MFMS aerodynamic shape optimisation framework was verified and evaluated via benchmark supercritical aerofoil shape optimisation. The framework was further demonstrated using the HART II rotor optimisation case as part of the International Design Workshop on Rotor Blade Optimisation (InDeWo).

NOTATION

*

Symbols

AoA Angle of Attack, [deg]

C_d Drag coefficient, [-]

α Design variable, [-]

ρ Density, [kg/m^3]

$FoM = T\sqrt{T/(2\rho\pi R^2)}/P$ Figure of Merit, [-]

C_l Lift coefficient, [-]

C_m Moment coefficient, [-]

C_p Pressure coefficient, [-]

R Rotor radius, [m]

r/R Rotor radial position, [-]

T Rotor Thrust, [N]

Q Rotor Torque, [Nm]

V_{tip} Rotor tip velocity, [m]

$\theta_{0.75}$ Rotor collective pitch at 75% span, [deg]

$C_T = \frac{T}{0.5\rho V_{tip}^2 \pi R^2}$ Thrust coefficient, [-]

$C_Q = \frac{Q}{0.5\rho V_{tip}^2 \pi R^3}$ Torque coefficient, [-]

Θ Neural Network Hyperparameters, [-]

*

Acronyms

CFD Computational Fluid Dynamics

FFD Free Form Deformation

HMB3 Helicopter Multi-Block 3 CFD framework

HOPT HMB Hybrid Optimisation Toolbox

HF, LF High-fidelity, Low-fidelity data or models

L, NL Linear and Non-linear neural networks

MFMS Multi-fidelity, Multi-start

MFNN Multi-fidelity Neural Network

SLSQP Sequential Least-suqare Quadratic Programming

I. INTRODUCTION

With advancements in simulation methods and tools over the past decades, aerodynamic shape optimisation has become an integral part of aerospace design. Nevertheless, important challenges, e.g. robustness, scalability, complexity, and efficiency (Ref. 1), still restrict its wider adoption in practical applications. The aerodynamic shape problem is characterised by high-dimensional and potentially multi-modal design spaces dominated by complex flow physics. The high-dimensionality leads to expensive design explorations, while the physical complexity requires high-fidelity tools such as computational fluid dynamics(CFD). Together, these often render aerodynamic shape optimisation impractical. Multi-modality is also not uncommon given the dimensionality and physical complexity (Refs. 2, 3). This could be an opportunity to find families of good designs, but also challenge the optimisation algorithms when the quest is to find the global optima.

It is well established that gradient-based optimisation methods are particularly efficient for high-dimensional problems (Ref. 1). The gradient information effectively drives each design variable towards improvement. Compared to classic surrogate-based, gradient-free approaches, the gradient-based methods require significantly fewer function evaluations. However, the risk of being trapped in local optima is inherent due to this gradient-driven nature. Restarting the optimisation from different initial guesses is a popular solution, but where to restart and how to verify whether a solution is the true global optimum remain open questions.

Multi-fidelity surrogate-based optimisation (Ref. 4) is another popular option for high-dimensional problems. The idea is to leverage models of different fidelity levels to sketch the large design space. Low-fidelity models are often used to quickly scan the entire design space for potential optimal regions, while high-fidelity models are used for localised and focused search. Surrogate models that can assimilate information from different fidelity levels, e.g. co-Kriging (Ref. 5), multi-fidelity Gaussian regression (Ref. 6), and more recently, multi-fidelity neural networks (Refs. 7–9), are often at the core of such multi-fidelity strategies. Ideally, multi-fidelity searches can easily evade sub-optimal local solutions and locate the true global optima, given sufficient exploration of the design space. However, considering the surrogate-based and gradient-free nature, the algorithms may still struggle with high-dimensional problems if the underlying surrogate models' performance is poor.

This work proposes a new hybrid strategy combining gradient-based and multi-fidelity surrogate-based optimisation, i.e. the Multi-fidelity, Multi-start (MFMS) approach. The key idea is to exploit the gradient-based local optimisation history to update multi-fidelity surrogates, and use the multi-fidelity surrogates to scan the global design space and indicate new restart points. Gradient-based optimisation is at the core of this hybrid approach, and the multi-fidelity surrogates offer gradient-based approaches the capability of finding global optimal solutions via coordinated restarts based on quantitative evaluations. The overall optimisation cost scales linearly with classic gradient-based approaches; hence, the new approach is especially suitable for high-dimensional problems. It is also shown that this hybrid approach greatly improved the robustness and flexibility of gradient-based search in the case of unsuccessful optimisation processes.

Compared to previous attempts (Refs. 9, 10) combining gradient-based and multi-fidelity approaches, this new hybrid strategy stands out with efficiency, robustness, and a more symbiotic integration of models, data, gradient information, and optimisation algorithms. We adopted a novel Multi-fidelity Neural Network (MFNN) (Ref. 7) architecture and linked it directly with physical models rather than data, enabling multi-fidelity and hybrid physics-/data-based learning. This hybrid model demonstrated excellent regression capability on very sparse samples.

The paper is organised as follows: Section II elaborates on the proposed MFMS methodologies along with all key components, including the MFNN architecture and our in-house Helicopter Multi-Block(HMB)-Adjoint framework (Refs. 10, 11) and the Hybrid Optimisation Toolbox(HOPT). Section III presents demonstrations and evaluations of the proposed MFMS approach, via canonical optimisation problems, benchmark aerofoil optimisation cases (ADODG benchmark case II (Ref. 1)), and rotary wing design optimisation problems. Last but not least, conclusions and future work are presented in Section IV.

II. METHODS AND TOOLS

The Multi-fidelity, Multi-start (MFMS) Optimisation Approach

The proposed multi-fidelity, multi-start approach combines both methods, and draws on each method's advantages to offset their weaknesses. The idea is to exploit the gradient-based optimisation history (both objective/constraint function values and gradients) to build high-performance surrogates, and use the surrogates to indicate potentially better solutions to restart the gradient-based optimisation. This is illustrated in Figure 1.

For the gradient-based algorithms, global surrogate-based optimisation is used as a restart strategy to bypass local optima. This restart strategy is quantitative and targeted towards global optima as the surrogates evolve. Another key advantage is that the surrogates provide a global sketch of the design space, as opposed to the single-threaded, local exploration of the design space in classic gradient-based optimisation.

From the surrogate-based algorithms' perspective, gradient-based local searches are used as the infill strategy to explore the design space. Given the gradient-based nature, this infill strategy is always centred around local optima; hence, it is particularly suitable for optimisation purposes. Of course, the surrogate models are key in the proposed MFMS framework. They still need the capability to handle sparse samples

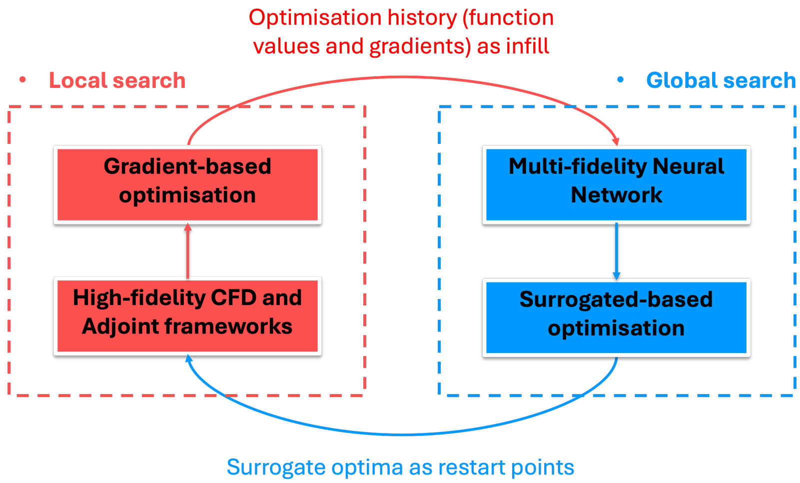

Figure 1. Flowchart of the multi-fidelity, multi-start optimisation approach.

and the particular ability to assimilate gradient information. This is why we introduced the Multi-Fidelity Neural Network (MFNN) as a key novelty.

In terms of computational cost, the proposed MFMS is much more efficient than classic global optimisation methods such as particle swarm or genetic algorithm (Ref.). The total computational cost is the cost of each gradient-based local search multiplied by the number of restarts. The cost in MFNN construction, update, and optimisation is negligible. Hence the MFMS cost scales linearly with the cost of HF local optimisation. It is of course more expensive than a single gradient-based optimisation, but it is the necessary cost to pay to bypass local optimal solutions. Nevertheless, the overall cost only linearly scaled, hence the proposed approach is suitable for high-dimensional optimisation such as aerodynamic shape optimisation.

The Multi-fidelity Neural Network (MFNN)

The performance of the surrogate models is key in the proposed MFMS framework. The models must be able to provide accurate predictions upon sparse samples from the gradient-based exploration, and must have the capability to assimilate gradients. In this light, we adopted the novel Multi-Fidelity Neural Network (MFNN) models. The MFNN architecture in this study is based on references (Refs. 7, 9) and is illustrated in Figure 2.

The idea is to approximate the correlation between high- and low-fidelity results as:

$$y_{HF} = F(y_{LF}(x), x) = F_l(y_{LF}(x), x) + F_{nl}(y_{LF}(x), x), \quad (1)$$

where F_l and F_{nl} represent linear and non-linear correlations, respectively. The MFNN is a composite network that approximates this correlation with several different neural networks, as shown in Figure 2. NN^l_{HF} and NN^{nl}_{HF} are two separate feed-forward networks representing F_l and F_{nl}, respectively. NN^l_{HF} is a linear network without any activation functions, while NN^{nl}_{HF} is a non-linear network regulated by the hyperbolic tangent function in this work. These two networks take as input the design variables x and the low-fidelity prediction y_{lf} at the same time. They are trained using a high-fidelity dataset. The final output combines the output of these networks as $y_{MF} = \omega F_l + (1-\omega)F_{nl}$, where $\omega \in [0,1]$ controls the linear and non-linear contributions. Note that ω is also a trainable parameter and is determined during the training process.

It should be highlighted that the low-fidelity input y_{lf} is also a function of the design variables x. The mapping from x to y_{lf} may use another neural network or use a physics-based low-fidelity model(LFM) directly.

In this work, a key difference from (Refs. 7, 9) is that we directly used LFMs to provide the low-fidelity data where possible. This is because the training of a neural network to fit LFMs is not at all easier than fitting the HF model directly, especially when the LFM is complex. The LFM provides much more accurate predictions and gradients than its neural network surrogates. This option substantially improves the MFNN modelling capability.

Nevertheless, where LFM is not available in this work, neural networks are a reasonable option to represent the LFM. We kept these LF neural networks frozen while training the linear and non-linear correlation networks. This does not deteriorate accuracy, as the NN_{LF} is supposed to be trained on low-fidelity data and then to supply information to other networks as intended. However, separating the training process simplifies the network architecture and reduces the training complexity.

MFNN training with data and gradients The MFNN is trained via either HF or LF data alone, or their mixture, and

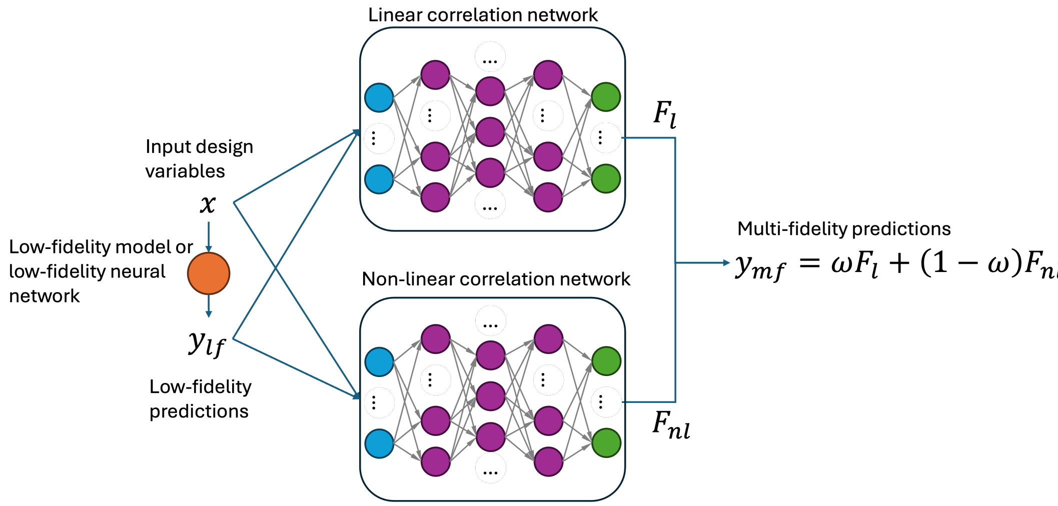

Figure 2. Architecture of the Multi-fidelity Neural Network (MFNN).

the training process can also be enhanced by gradients if available. In our work, when used for optimisation purposes, we often have an LF database with abundant LF samples and a HF database with limited HF samples. If no gradient information is available, the MFNN is trained purely with HF/LF data with the loss function:

$$L_{data}(\theta_{NL},\theta_L) = \frac{1}{N_{hd}}\sum_{i=1}^{N_{hd}}(y_i^{HF}-\hat{y}_i^{MF})^2 + \frac{w}{N_{ld}}\sum_{j=1}^{N_{ld}}(y_j^{LF}-\hat{y}_j^{MF})^2 \tag{2}$$

where N_{hd} and N_{ld} are the numbers of HF and LF data respectively. y^{HF} and y^{LF} are the high- and low-fidelity training datasets, respectively. $\hat{y}^{MF}$ is the MFNN prediction at the same LF/HF sampling points. w is a weighting factor that ensures the LF contribution is small. The LF contribution is necessary when there are no or a very limited number of HF samples. It enforces the MFNN to fit the LF data in the absence of HF information and avoids non-physical interpolations. As HF samples increase, the LF contribution becomes negligible, and HF loss takes priority. In our work, we used $w \approx 10^{-4}O(\frac{1}{N_{hd}}\sum_{i=1}^{N_{hd}}(y_i^{HF}-\hat{y}_i^{MF})^2)$.

When gradient information is available. The loss function of the MFNN is defined as:

$$L(\theta_{NL},\theta_L) = L_{data}(\theta_{NL},\theta_L) + \frac{1}{N_g}\sum_{k=1}^{N_g}(\nabla y_k^{HF}-\nabla\hat{y}_k^{MF})^2, \tag{3}$$

where θ_{NL} and θ_L are the hyper-parameters of the nonlinear and linear networks in Figure 2, respectively. This first term of the RHS of Equation (3) represents the mean squared error (MSE) between the MFNN prediction and the HF/LF dataset. The second term of the RHS represents the MSE between the MFNN gradients and the gradients of the HF dataset, and this is how the model was enhanced by gradients. ∇y^{HF} is the HF gradient in the gradient dataset with N_g samples. HF gradients can be computed via methods such as adjoint, complex steps, finite differences, or automatic differentiation while computing the HF dataset. Gradients of the LFM could be incorporated as well, following the same principle, but in our work, we considered only HF gradients, as LF data and models are not considered accurate anyway. Note that the gradient data sample points may not necessarily coincide with the dataset, hence N_{hd} and N_g might be different.

$\nabla\hat{y}^{MF}$ is the gradient of the MFNN predication. If the LF input is provided by another neural network, the gradients are readily available thanks to the built-in and streamlined automatic differentiation. However, it must be highlighted that, if the LF input is provided directly by a low-fidelity physical model (LFM), then following the chain rule, we should have:

$$\nabla\hat{y}_j^{MF}(x_k, y_j^{LFM}) = \frac{\partial\hat{y}_j^{MF}}{\partial x_k} + \frac{\partial\hat{y}_j^{MF}}{\partial y_j^{LFM}}\frac{\partial y_j^{LFM}}{\partial x_k}, \tag{4}$$

where x_k is the kth input variable. y_j^{LFM} is the LFM prediction and $\frac{\partial y_j^{LFM}}{\partial x_k}$ is its gradients. This means that the LFM must provide gradients too. These gradients may not be readily available and might need methods like adjoint or finite differences.

Regardless, this gradient SME loss part enforces the MFNN to have the same gradients as the HF function. This helps reduce the number of sampling points, especially for high-dimensional information. Assuming that HF data and gradients are taken at the same sampling points, and HF function evaluations and gradient evaluations are of similar costs (via e.g. adjoint methods), then with N function evaluations and N gradient evaluations, we can have $N+dN$ information, thanks to the gradients (d is the problem dimension). Compared to non-gradient-enhanced approaches, the ideal speedup s is linear as $s=\frac{N+dN}{2N}=\frac{1+d}{2}$. This means that the gradient-enhanced strategy is particularly efficient for high-dimensional problems. However, the embedding of gradient information of course introduces substantially more computations during the network training process. HF gradients are sometimes difficult to acquire and often blended with strong noise that hinders the training convergence. In practice, the gradient loss may be introduced only at limited sampling points and assigned with cautious weights.

The training of the LF network, if it is necessary, can follow the same loss function composition as in Equation (3), except that all data and gradients come from LF predictions.

High-fidelity HMB3-Adjoint CFD framework

For the high-fidelity computation of the aerodynamics and gradients, we used our in-house Helicopter Multi-Block (HMB3) CFD framework (Refs. 11, 12). The code has been widely used in simulations of rotorcraft flows. HMB3 solves the Unsteady Reynolds Averaged Navier-Stokes (URANS) equations in integral form using the Arbitrary Lagrangian Eulerian (ALE) formulation for time-dependent domains, which may include moving boundaries. The Navier-Stokes equations are discretised using a cell-centred finite volume approach on a multi-block, structured grid:

$$\frac{d}{dt}(\mathbf{W}_{i,j,k}V_{i,j,k}) = -\mathbf{R}_{i,j,k}(\mathbf{W}_{i,j,k}), \tag{5}$$

where i,j,k represent the cell index, $\mathbf{W}$ and $\mathbf{R}$ are the vector of conservative flow variables and residual, respectively, and $V_{i,j,k}$ is the volume of the cell i,j,k. In the present work, simulations of propellers in hover were performed in the non-intertia Rotating Reference Frame (RRF) method (Ref. 11) in HMB3. This converted the modelling to steady simulations with rotational periodicity and enabled the use of adjoint computations for gradients.

To evaluate the convective fluxes, the Osher-approximate Riemann solver is used, while the viscous terms are discretised using a second-order central difference scheme. The 3^{rd} order MUSCL (Monotone Upstream-centred Schemes for Conservation Laws) approach was used to provide high-order accuracy in space. The chimera/overset grid method was used in this work. In the present work, simulations were performed with the $k-\omega$ SST (Shear Stress Transport) (Ref. 13) turbulence model.

Along with high-fidelity aerodynamic solutions, the HMB3 framework also provides accurate gradient information following the adjoint formulation. The adjoint method (Ref. 12) is known for its efficient handling of hundreds of thousands of design variables with a handful of objective functions. This is especially suitable for the aerodynamic shape optimisation problem at hand, which involves dozens of design variables governing shapes with just a few objective/constraint functions.

In the context of CFD analysis, a typical objective function I can be explicitly defined as $I(\mathbf{W}(\alpha),\alpha)$, a function subject to the design variable vector α and the flow variables $\mathbf{W}$. The flow variables still satisfy the governing equation in the form $\mathbf{R}(\mathbf{W}(\alpha),\alpha=0$.

The sensitivity equation of I relative to the design variables α is as follows:

$$\frac{dI}{d\alpha} = \frac{\partial I}{\partial \alpha} + \frac{\partial I}{\partial \mathbf{W}}\frac{\partial \mathbf{W}}{\partial \alpha}. \tag{6}$$

Introducing an adjoint vector λ correlating the cost function and the governing equations, the sensitivity equation can be recast in the adjoint form as:

$$\begin{aligned} \frac{\partial \mathbf{R}}{\partial \mathbf{W}}^{\mathbf{T}}\lambda &= -(\frac{\partial I}{\partial \mathbf{W}^T}), \\ \frac{dI}{d\alpha} &= \frac{\partial I}{\partial \alpha} + \lambda^T\frac{\partial \mathbf{R}}{\partial \alpha}. \end{aligned} \tag{7}$$

The major effort in solving the sensitivity equation aligns with solving the adjoint vector λ coupled in the linear system in Equation (7). This linear system scales with the number of objective functions and is irrelevant to the design variables. This property makes the adjoint formulation especially suitable for aerodynamic shape optimisation. Details of the HMB3-Adjoint formulations can be found in previous optimisation studies of rotors, wings, and ducted propellers (Refs. 11,12).

The Hybrid Optimisation Toolbox(HOPT)

The proposed MFMS algorithm, along with the key components, e.g. the MFNN model, optimisers, data management, and utilities such as parameterisation tools and drivers, are all implemented and integrated in our in-house Hybrid Optimisation Toolbox(HOPT) framework. HOPT is implemented as a high-level Python library that provides interfacing, communication, and streamlined automation for various optimisation tools. It is interfaced with our HMB3-Adjoint framework to extract aerodynamic solutions and gradients. The proposed MFNN architecture, along with all necessary data management, automatic differentiation, and training procedures, was implemented within HOPT using PyTorch. For the gradient-based optimisation, the Sequential Least-Squares Quadratic Programming (SLSQP) algorithm (Ref. 14) as provided in the NLopt library (Ref. 15) was linked with HOPT to drive the optimisation process. For the surrogate-based optimisation using the MFNN models, a simple classic genetic algorithm was used.

Low-fidelity aerodynamic models

For the aerofoil and the rotary wing optimisation cases in this work, we used two different low-fidelity physical models as the LF input end to the MFNN. The aerofoil cases, the NeuralFoil tool (Ref. 16) was used. This is a neural network representation of the classic XFoil tool (Ref.). It was used for its capability to provide gradients that are needed for the MFNN training. As for the rotary wing case, a simple BEMT model was composed, and the gradients were computed via automatic differentiation.

III. NUMERICAL RESULTS

Evaluation of the Multi-fidelity, Multi-start (MFMS) optimisation algorithm

The proposed multi-fidelity, multi-start (MFMS) optimisation algorithm was first evaluated using the Styblinski-Tang function (Ref.). The function is defined as:

$$f(\mathbf{x}) = \frac{1}{2}\sum_{i=1}^{n}(x_i^4 - 16x_i^2 + 5x_i), \tag{8}$$

where $\mathbf{x} = (x_1, x_2, x_3, ..., x_i, ..., x_n)$ is the n-dimensional input and $n \geq 1$ is the dimension. Over the domain $x_i \in [-5, 5]$, this function is highly multi-modal and has multiple local optima and one global optima at $f(-2.903534, -2.903534, ..., -2.903534) \approx -39.166n$.

The LF approximation we have used in this case is simply :

$$f_{LF}(\mathbf{x}) = \frac{1}{8}\sum_{i=1}^{n} x_i^4, \tag{9}$$

which is uni-modal has only one global optima at the origin.

The MFMS algorithm was first demonstrated and assessed in 2D, where the Styblinski-Tang function in Equation (8) has 4 local optima. Figures 3(a) to 3(d) demonstrate the optimisation search and the MFNN updates during the optimisation.

Figure 3(a) shows the initial MFNN representation. This was trained using only 500 random LF samples. No initial HF data was used. It hence represented the LF model at this initial stage. The '+' sign denotes the local optimum of the MFNN found using exhaustive local optimisation starting from 50 random locations. The cost of this is negligible thanks to the rapid MFNN predictions. Note that this local optimum was note exactly at the origin due to errors in MFNN approximation. Regardless, this local optimum was used as the starting point for HF local optimisation. The HF search used a SLSQP algorithm and its searching history is represented by the black lines and red dots. Note that in this case, the HF search only managed to identify a local optimum. Regardless, the HF data and gradients generated in this search process were used to update the MFNN.

Figure 3(b) presents the updated MFNN. With very limited exploration of the HF design space, the MFNN was able to recover patterns of the HF function with the help of gradient information. All local optima on the updated MFNN were then identified, and HF local optimisation was initialised from these locations, as illustrated in Figure 3(c). Note that the local optimum in the first quadrant in Figure 3(c) did not initialise any HF search. This is because this point already existed in the previous HF database, hence there was no need to the search anymore. After exploring the other optima, all local optima including the global optima were identified. Overall, 4 HF local searches were carried out, which is exactly the minimal necessary to find all local optima using local optimisation algorithms. The updated MFNN is shown in Figure 3(d) with dots representing the final HF data points. This approximation is very close to the HF function in Equation 8, especially near the local optima.

However, it must be stressed that the proposed MFMS optimisation is not a global optimisation algorithm and cannot guarantee a global optimum. The purposes of the proposed algorithm are to improve local optima delivered by gradient-based optimisation, and to sketch the design space exploiting the optimisation history. In the ideal case, where the MFNN approximates the HF function well and the exploration of the HF design space is sufficient, the MFMS may identify all existing local optima with the minimal necessary amount of local searches. In practice, however, due to differences between the MFNN and the HF function, the MFMS may miss the global optimum region and terminates prematurely. The MFMS could not proceed further if all MFNN local optima already exist in the HF database. This problem of premature convergence also exists in many global optimisation algorithms such as particle swarm and simple surrogate-based optimisation. This is mainly due to insufficient exploration of the design space.

Table 1 presents an evaluation of this problem. It also presents comparisons against the simple random restart approach (case C1), i.e. restarting HF local search via random LHS samples. For case C1, the HF local search restarted from 1,000 different random initial positions over the domain $[-5, 5] \times [-5, 5]$. On average, each HF search required 10 HF function evaluations and 9 HF gradient evaluations. However, the possibility of finding the global optimum was only 26.1%, which actually corresponds to the area ratio of the valley region near the global optimum.

Case C2 used the MFNN to indicate new restart points following the MFMS process. The MFNN was initialised using 200 random LHS samples, and the MFMS optimisation was repeated 1,000 times. On average, each optimsiation trial required 34 HF function evaluations and 34 gradient evaluations. These are, of course, more costly than optimisation from a single initial guess as in case C1. However, these are necessary extra costs to search for better local optima. The number 34 corresponded to about 4 restarts of the HF local search, which were the minimal necessary to find all local optima. However, the possibility of finding the global optimum in this test was around 90.9%. Though this was substantially higher than the simple random search approach, it demonstrated that the proposed MFMS could not guarantee a global optimum as aforementioned, due to in sufficient exploration of the design space.

A possible solution to mitigate this issue is to initialise the MFNN with limited HF samples, as demonstrated with case C3 in Table 1. The MFNN of C3 was initialised with 200 LF samples plus 20 HF samples uniformed sampled in the design space with LHS. This improved the MFNN accuracy, thereby improving the chance of finding a global optimum. C3 was also repeated 1, 000 times with random LF and HF samples in each trial. The chance.

However, if we further increase the number of initial HF sam-

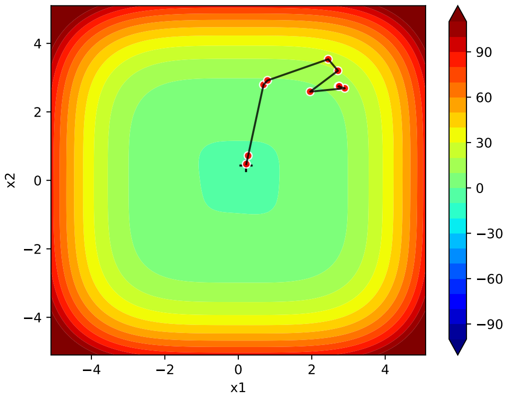

(a) Initial MFNN approximation with its local optimum denoted by '+'. The black lines and red dots represent the HF search history starting from this local optimum.

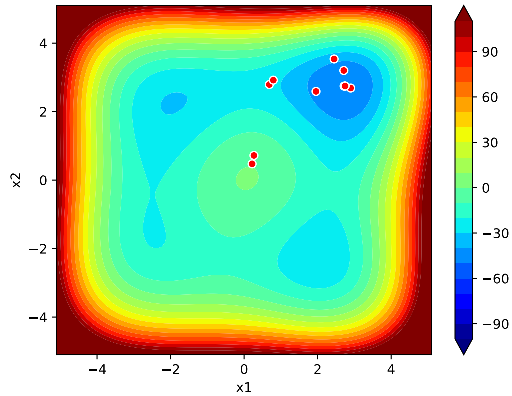

(b) Updated MFNN using HF data and gradients generated in the HF local search process. The dots represent the current HF database.

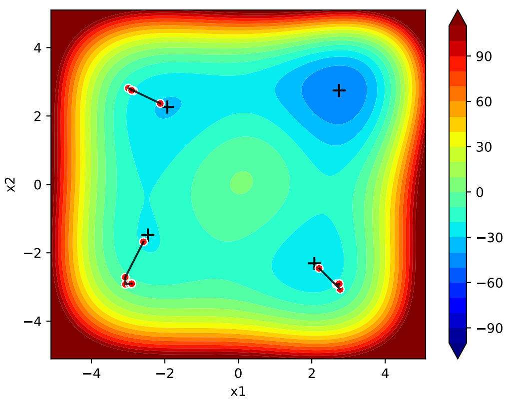

(c) Local optima on the updated MFNN ('+') and the HF search history starting from these local optima (lines and dots).

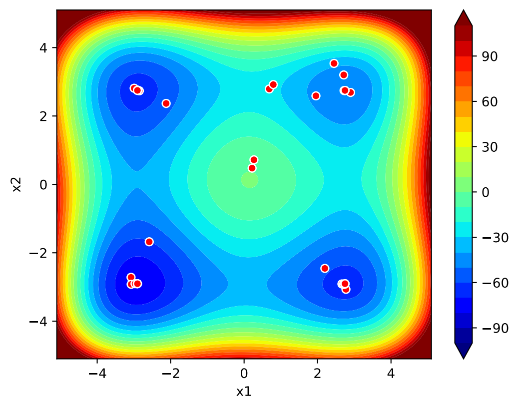

(d) Updated MFNN using the HF optimisation data and gradients. The dots represent the current HF database.

Figure 3. Optimisation of the 2D Styblinski-Tang function using the proposed MFMS approach.

Table 1. Comparisons between the simple random restart HF optimisation and the MFMS approach with and without initial HF samples. All cases were repeated 1, 000 times with random Latin Hypercube Sampling (LHS).

Cases	Initial LF samples	Initial HF samples	Restart methods	Average HF function evaluations	Average HF gradient evaluations	Global optima chance
C1	-	-	Random	10	9	26.1%
C2	200	0	MFNN	34	34	90.9%
C3	200	20	MFNN	26	26	99.7%

ples, the MFMS would devolve into a surrogate-based optimisation approaches. For high-dimensional design problems, this may not be feasible. There is tradeoff between final solution and costs.

Other solutions include analysing the distribution and topology of the final HF samples, and decide on where to add more. Inserting uncertainty analysis and borrow ideas from EGO may help as well, but these are for future work.

The ADODG benchmark case II: drag minimisation of the RAE2822 aerofoil

This section demonstrates and validates the proposed MFMO algorithm using the benchmark RAE2822 aerofoil case of ADODG (Refs. 1, 17), which aims to minimise drag at the condition of Ma = 0.734 and Re = 6.5×10^6. This is a critical condition where shock waves and shock-boundary-layer interactions are encountered. The complete optimisation problem is formulated as:

$$\begin{aligned} \text{minimise}: \; & C_d, \\ \text{by varying:} \; & (\alpha_i, AoA), \\ \text{subject to:} \; & C_l = 0.824, \\ & C_m \geq -0.092, \\ & A \geq A_0, \end{aligned} \tag{10}$$

where α_i are parameters controlling the aerofoil shape, and the AoA is also included as a design variable. A is the area of the aerofoil, and A_0 is the initial aerofoil area. This is a highly non-linear optimisation problem, with a stringent equality lift constraint and 2 inequality constraints on the moment and area.

Mesh convergence study of the original RAE2822 aerofoil RANS solutions of the initial RAE2822 shape were first computed using the HMB3 CFD framework with the $k-\omega$ turbulence model (Ref. 13). (This differs from most existing studies on this case, which use the S-A model.) We used multi-block, fully-structured meshes around the aerofoil with a typical C-type topology. The aerofoil was modelled with a sharp trailing edge. A mesh convergence study was first carried out as detailed in Table 2. The drag counts steadily decreased as we increased cell numbers in all directions. However, between meshes 3 and 4, the drag change was around only 1 count, hence mesh 3 was adopted for the optimisation study.

Figure 4(a) presents the pressure solution using mesh 3. A strong shock wave can be noted on the upper surface at this critical condition, which is prone to strong shock-boundary-layer interactions. Figure 4(b) presents all surface pressure comparisons between all meshes of Table 2. Differences between meshes 2, 3, and 4 are only minor, which adds to the confidence of using mesh 3 for optimisation studies in the following work. Note that this case is known as a highly sensitive case, where slight differences in turbulence modelling, meshing, and numerical settings can result in important differences in drag predictions. As summarised by Ref (Ref. 17), different studies reported diverse drag predictions. Our current drag predictions were slightly high but still within the range reported.

Optimisation setup, history, and convergence In this work, we adopted the typical Class Shape Transformation (CST) parametrisation for the aerofoil shape. A total of 16 parameters were used for the shape – 8 for the upper surface and 8 for the lower surface, respectively. These CST coefficients α_i were bounded within ± 0.1 with respect to the initial RAE2822 shape. Figure 5 illustrates the shape variation ranges relative to the original RAE2822 shape in this work. Shape changes were implemented as mesh deformations in the CFD modelling via an inverse distance weighted deformation procedure (Ref. 12). As aforementioned, the AoA was also considered as a design variable, and was bounded within ± 1 degrees relative to the initial AoA of 2.885 degrees. During the optimisation, all AoA values and their derivatives were converted to radians to normalise the design variables. No further normalisation or scaling of the design variables was applied. In total, there were 17 design variables considered in this case.

To initialise the MFNN in the MFMO process, we sampled the design space with 10,000 LF samples and 50 HF samples separately via LHS sampling. The LF data and model were constructed using the NeuralFoil (Ref. 16) tool as described in Section II, while the HF data was generated using the HMB3 solver. No gradient information was computed at this initial stage.

Three separate MFNNs were trained for the lift, drag, and moment, respectively. This is to reduce complexity in the MFNN architecture, and reduce difficulty in the MFNN normalisation and training. For the drag MFNN, the final MFNN output was further regulated with a square operation, which eliminates negative drag values that are non-physical.

After the MFNN initialisation, a first gradient-based optimisation was launched from the initial RAE2822 shape. The gradients were computed using the HMB3-Adjoint, and the optimisation problem was solved via an SLSQP algorithm with hard constraints. The optimisation history was then used to update the MFNNs. The global optimum of the MFNNs was then sought via a simple genetic algorithm, and this optimum was used as the restart point for the next gradient-based HF optimisation cycle. This MFMS process iterated until a convergence was achieved. Figure 6 and Table 3 presents the convergence history of the MFMS for this case. 4 restarts were invoked, and the MFNN and the HF optimisation eventually reached very close solutions, where further HF cycles would not produce any new optimum, hence the MFMS process was stopped.

Note that, as shown in Table 3, the 3rd restarted HF optimisation failed due to excessive shape deformation that resulted in non-physical solutions. This is a typical situation that challenges the error-handling capabilities of optimisation algorithms and frameworks. For the current HF gradient-based optimisation coupled with CFD, this could happen due to, e.g. excessive shape and mesh deformations or non-physical conditions during the searching. In such conditions, classic gradient-based optimisation would have to restart from a different initial guess, and the computational resources are wasted. However, it should be highlighted that in our MFMS framework, failed optimisation still provides valuable information to the MFNN surrogates and helps guide subsequent searches regardless. As shown in Figure 6, the MFNNs were still able to indicate restart points from the failed 3rd restart and quickly achieved convergence in the 4th restart.

Table 3 also presents the number of HF function and gradient evaluations during each HF optimisation cycle. Via adjoint approaches, the HF function and gradient evaluations are of similar costs. For this aerofoil optimisation case, the initial optimisation from RAE2822 and the first restart required the most costs, while the following second and fourth restarts needed only half the costs since their initial guesses were close to the final solutions. The third restart also took similar costs

Table 2. Mesh convergence study using 4 different sets of meshes at Cl=0.824.

Mesh	Cell numbers, [k cells]	Cd, (counts)
1	34.4	221.86
2	76.72	220.51
3	122.4	216.32
4	189.2	215.34

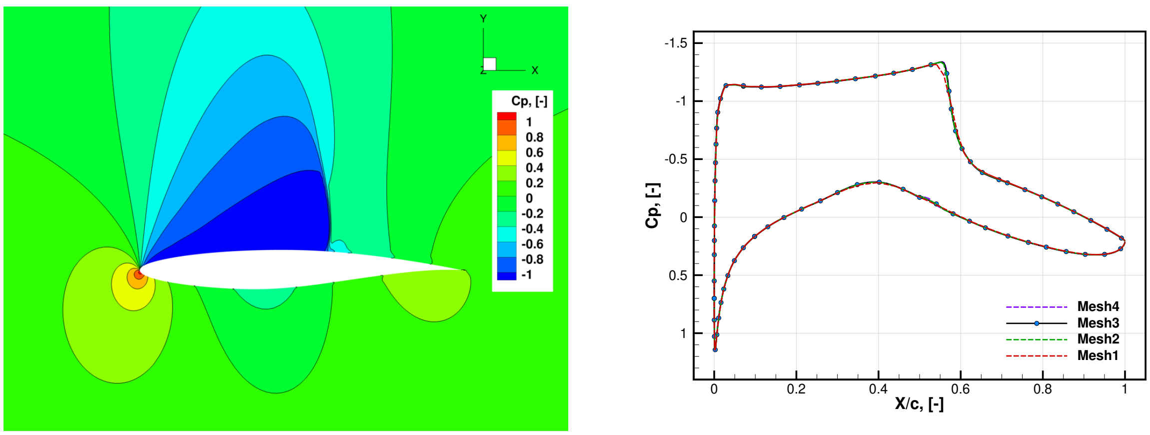

(a) Pressure coefficient contour using mesh 3 of Table 2.

(b) Surface pressure coefficient distribution variations using different meshes.

Figure 4. Pressure coefficient solutions using different meshes of Table 2.

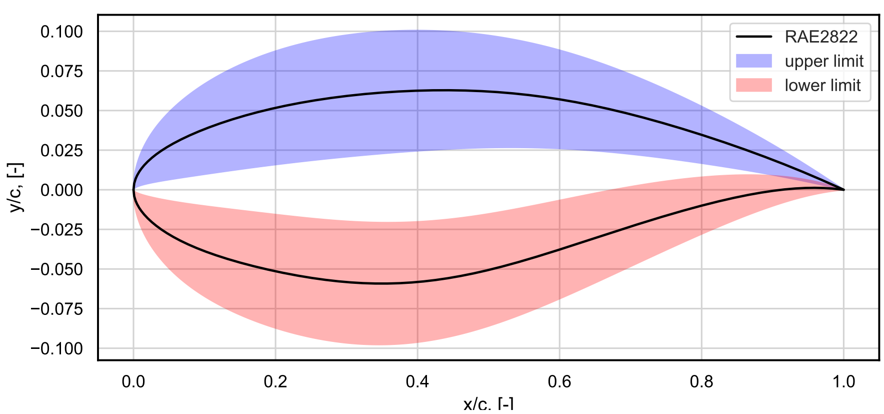

Figure 5. Shape variation ranges for the RAE2822 optimisation study.

as the initial due to failed search, and all this information was exploited in the MFNNs and was not wasted. Convergence history of the initial optimisation is presented in Figures 7(a) to 7(d), and all constraints were satisfied within minor tolerance.

The final drag values from all HF optimisations were also presented in Table 3. All successive optimisations reported similar optimised drag values near 136 counts, with differences lower than 0.7 counts. These minor differences were negligible and were mainly due to differences in mesh deformation, flow convergence, and optimisation convergence. Compared to the original drag value of 216.32 using the same mesh 3 in Table 2, the current optimisation reduced the drag by about 80 counts, which is also within the typical range reported by Ref (Ref. 17).

Optimised shapes and flow solutions Figures 8(a) to 8(d) present comparisons of optimised shapes from each success-

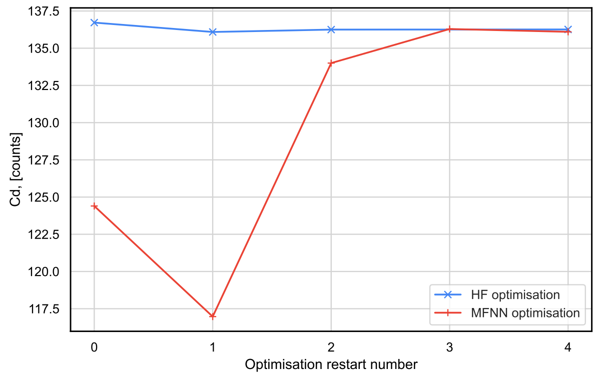

Figure 6. Convergence of the MFNN model and the restarted HF optimisation.

Table 3. Overall optimisation results and costs with multiple restarts.

Optimisation case	Cd, [counts]	HF function evaluations	HF gradient evaluations	Total costs
Initial optimisation	136.72	44	30	74
Restart 1	136.09	49	30	79
Restart 2	136.25	26	20	46
Restart 3	(failed)	50	29	79
Restart 4	136.26	20	15	35

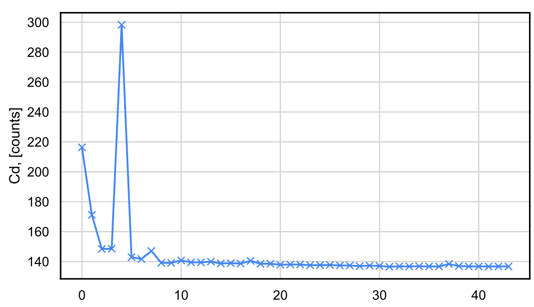

(a) Objective (drag) history.

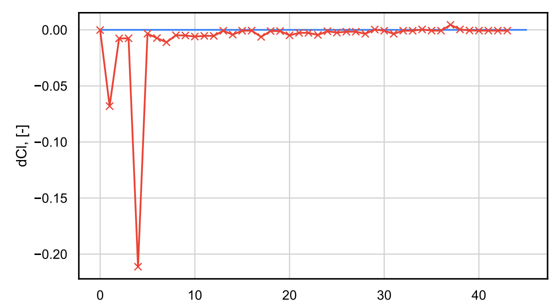

(b) Equality lift constraint history.

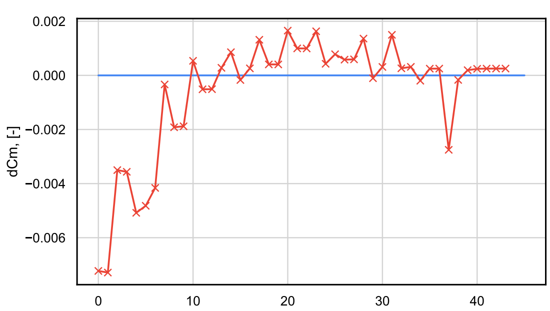

(c) Inequality moment constraint history.

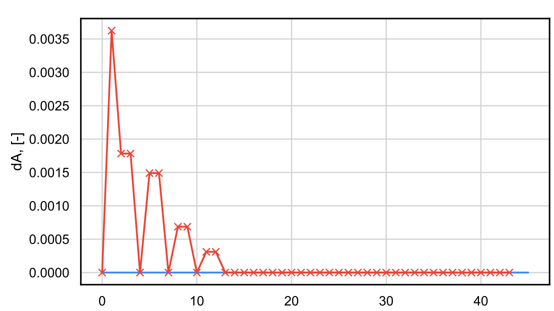

(d) Inequality area constraint history.

Figure 7. HF gradient-based optimisation convergence history of the initial optimisation case of Table 6 starting from the original RAE2822 shape.

ful optimisation of Table 3. Figure 8(a) shows that all optimal shapes have only minor differences. Combined with the similar drag values reported in Table 3, these suggest that only one optimal solution was sought by the MFMS. This agrees with findings reported by other studies (Ref. 17) on this case. For the first restart, as shown in Figure 8(b), the MFNN surrogates suggested slightly different initial shapes. However, for the 2nd and 4th restarts, the initial shapes were almost identical to the final solutions, and differences were mostly in the *AoA*. This agrees with the observation that restarts 2 and 4 required fewer iterations to converge.

Figures 9(a) and 9(b) present the flow solutions of the optimised shapes. It can be noted that the flow field was shock-free after the optimisation. This is the main reason for the drag reduction, and this finding agrees with existing optimisation studies (Ref. 17). As shown in Figure 8(a), this was achieved by slightly expanding the aft half of the aerofoil to create a super-critical shape that delays the formation of shocks. At the same time, the lower-surface leading edge was shrunk to maintain area, lift, and moment. Pressure distribution comparisons in Figure 9(b) highlight these aerodynamic changes. Note that there were still minor pressure differences near the lower leading edge among optimal solutions. It is difficult to consider these as evidence for multiple solutions, as the shape and drag differences were negligible. We consider these mainly due to minor differences in mesh deformation, flow convergence, and optimisation convergence.

Evaluation and exploitation of the MFNN surrogates Figures 10(a) and 10(b) present the distributions of all HF samples generated during the optimisation process. The 'x' denotes the initial 50 HF LHS samples used to establish the initial MFNNs. Dots connected with lines represent the complete search history during each restarted optimisation. In total, there were 239 HF samples and 124 sets of gradient data. These are certainly more costly than a single gradient-based optimisation, but are still much more efficient than typical surrogate-based optimisation, given the dimensionality of 17. From a practical point of view, it is difficult to know the modality of the design space a priori; hence, restarting and sampling are necessary costs to extract information, establish confidence, and avoid sub-optimal solutions.

For the current MFMS approach, it is shown in Figures 10(a) and 10(b) that the design space exploration was centred around the optimal solution, with occasional breaches into larger domains. This was due to the gradient-driven nature. With this concentration of samples, the local MFNN accuracy is certainly higher near the optimum. Hence, the MFNN surrogates can give confident performance predictions when there are small excursions near the optimum. However, it is not clear whether the global accuracy was improved.

To understand the evolution of the MFNN surrogates' global accuracy, we generated an extra 20 HF samples across the entire design space for validation purposes. Figures 11(a) to 11(c) present the correlation between the drag MFNN surrogate and the validation HF dataset during the MFMS optimisation process. For the initial MFNN in 11(a) with only 50 initial HF samples, it struggled to predict the validation or initial optimisation datasets. However, as the initial optimisation data and gradients were assimilated, the correlation was immediately improved as shown in Figure 11(b). When the optimisation finished the 3rd restart, as shown in Figure 11(c), the correlation between the MFNN and the HF dataset was greatly improved. It was able to predict most of the data points in the 4th restart accurately. This highlights that the global accuracy of the MFNN surrogates was also improved, although the sampling was centred around the optimum.

With this confidence, the MFNN surrogates could be exploited post-optimisation, to inquire performance or interpolate new designs. As demonstrated in Figure 12, the MFNN surrogates could offer quick suggestions of new designs when there are changes in the design objectives and constraints. For demonstration purposes, we relaxed the lift constraint and removed the area constraint. The MFNN quickly suggested a group of designs whose drag was below 0.014. We can also notice that in this case, the lower leading-edge surface enjoyed more freedom in shape, while the upper surface shape was more critical to delaying shock formations.

Applications to rotary wing shape optimisation

This section applies the MFMS approach to the optimisation of the HART II rotor (Ref. 18) in hover as part of the International Design Workshop on Rotor Blade Optimisation (InDeWo) (Refs. 19, 20). The HART II rotor is four-bladed. The rotor blade has a radius R of 2 metres, and a reference chord length C_0 of 0.121 metres. The blade has a constant chord, a linear twist, and a rectangular planform without any anhedral/dihedral or sweep. The sectional shape is the NACA23012 aerofoil fitted with a tail tab. More detailed blade definitions can be found in Ref (Ref. 18). In hover, the blade rotates at an *RPM* of 1042 rev/min. This results in a tip Mach number of $Ma = 0.64$. The Reynolds number based on the reference chord and the tip speed is $Re = 1.79 \times 10^6$.

Optimisation problem formulation and parametrisation In our study, the optimisation problem at hand is to find the optimal rotor shape that minimises the torque while maintaining a target thrust in hover. The optimisation problem is formulated as follows:

$$\begin{aligned} &\text{minimise}: C_Q, \\ &\text{by varying: } \{\alpha\}, \\ &\text{subject to: } C_T = 0.01522, \end{aligned} \tag{11}$$

$\{\alpha\}$ are the shape design variables. In this case, we adopted the Free Form Deformation (FFD) approach (Ref. 21) for the blade shape parametrisation. As shown in Figure 13, the rotor blade was wrapped in a box bounded by the initial FFD lattices. New blade shapes were introduced by offsetting these FFD lattices, which brings elastic deformation to the shape. The design variables were hence the offset relative to the initial lattice, and the initial design was represented as $\alpha = \mathbf{0}$.

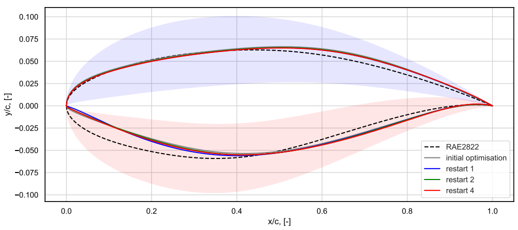

(a) Overall comparisons between the original RAE2822 and optimised shapes from each successful restart.

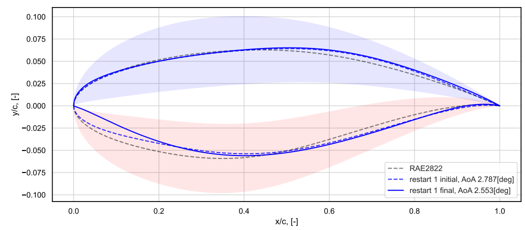

(b) Comparisons between the initial and final shapes in Restart 1.

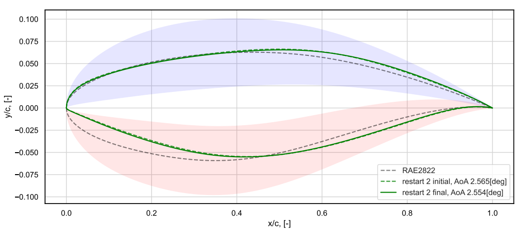

(c) Comparisons between the initial and final shapes in Restart 2.

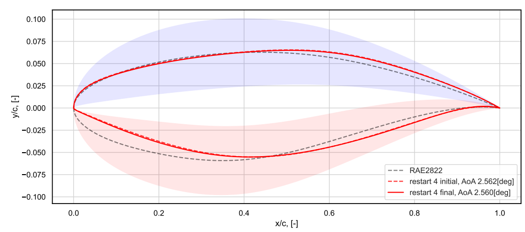

(d) Comparisons between the initial and final shapes in Restart 4.

Figure 8. Comparisons of final optimal aerofoil shapes from each restart of Table 3

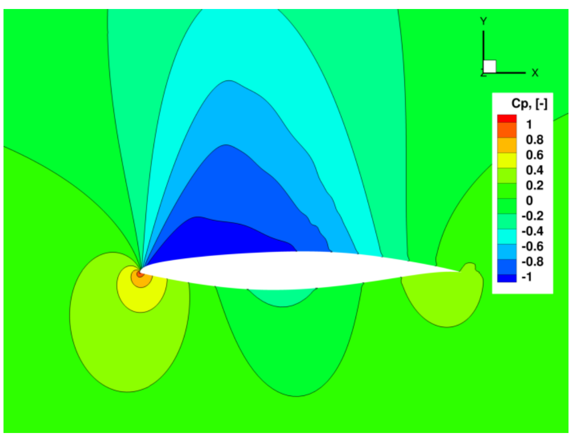

(a) Pressure coefficient contour of Restart 4.

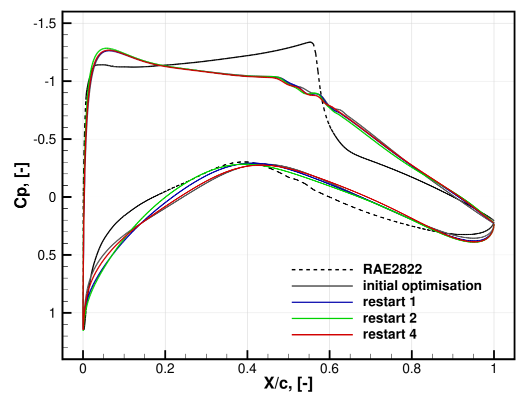

(b) Comparisons of surface pressure coefficient distributions from restarts.

Figure 9. Pressure solutions of the optimised aerofoil shapes.

In this work, we restricted the lattice offsets to only the y and z directions to avoid changing the blade radius. We used 7, 2, and 3 sections of FFD lattices along the blade spanwise, chordwise, and thickness directions, respectively. Moreover, on each spanwise section, as highlighted in Figure 13, the 2 lattices along the thickness were coupled to share the same offsets. This was to slightly restrict sectional shape deformation and focus more on the blade planform optimisation.

In total, we had 42 design variables for this case(7 sections, 3 free lattices per section, and 2 degrees of freedom per lattice). This parametrisation strategy included all key blade shape features, i.e. the twist, chord, anhedral/dihedral, and sweep. The sectional shape was also included, but restricted to mostly camber changes.

It should be highlighted that multiple local optimal solutions should be expected from this parameterisation strategy. For instance, given any optimal solution, a collective offset of all lattices along the z direction would lead to an equally good solution. The sectional camber and pitch changes may cancel out, too. Dimensionality reduction analysis, or simply more constraints, may help reduce the design space. Nevertheless, we allowed such multi-modality to evaluate whether the proposed MFMS approach is capable of finding multiple local solutions as expected. We were also interested in investigating the differences between different local solutions as different design families.

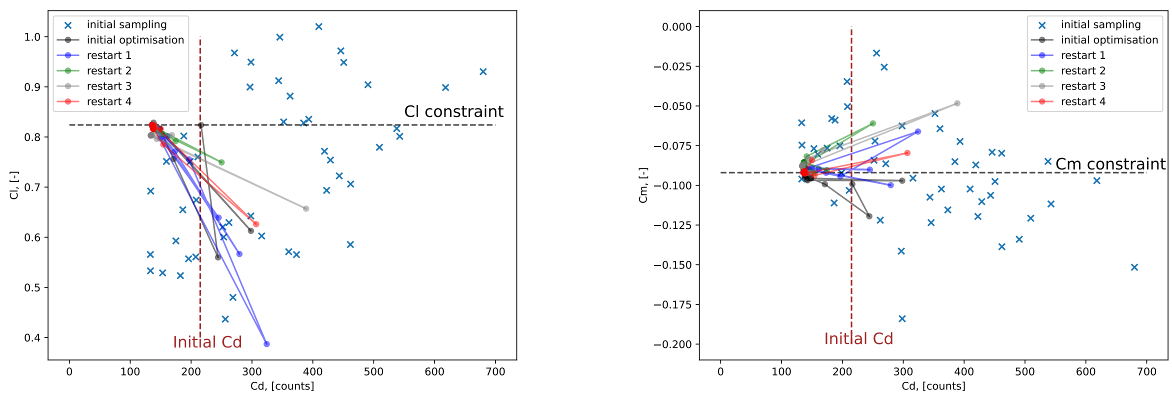

(a) Data distributions projected onto the Cl-Cd domain.

(b) Data distributions projected onto the Cm-Cd domain.

Figure 10. Overall distributions of high-fidelity data over the MFMO process. 'x' denotes the initial 50 LHS HF sampling to establish the initial MFNN. Lines with dots represent the search history in each restart. Data from the failed Restart 3 case was also exploited.

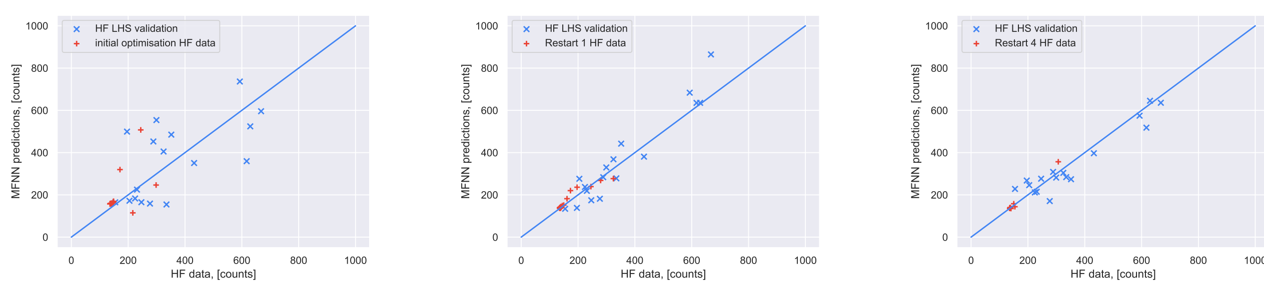

(a) Initial MFNN and its correlations with the HF validation dataset and the initial optimisation dataset.

(b) MFNN after assimilating the initial HF optimisation, and its correlations with the HF validation and the 1st restart datasets.

(c) MFNN after assimilating the 3rd HF restart, and its correlations with the HF validation and the 4th restart datasets

Figure 11. Evaluation of the drag MFNN accuracy during the optimisation. The blue 'x' denotes an extra 20 HF LHS samples for validation purposes. The red '+' denotes the next restart HF data before it was assimilated. The line represents the 1:1 correlation.

Meshing and initial solutions The high-fidelity HMB3 RANS simulations were carried out in a rotating reference frame. The fully structured, overset mesh topology is presented in Figure 14. For this four-bladed rotor, only 1/4th of the domain was modelled with periodic boundaries to save computational resources.

Table 4 presents the mesh convergence study. All meshes here have the same block structure and overset topology. The finest mesh consisted of 13.84 cells in total, and the medium and coarse meshes were derived from this by coarsening the blade and the background. In all simulations, the rotor thrust was trimmed to the target C_T of 0.01522 by adjusting the collective pitch angle $\theta_{0.75}$. As shown in Table 4, for all meshes, the final $\theta_{0.75}$ was about 10 degrees. From coarse to fine meshes, the convergence in torque and *FoM* is evident. The fine and medium meshes reported almost identical results. The *FoM* of the initial design at this thrust is around 0.68. However, it is worth noticing that the coarse mesh results had only very small discrepancies in this case.

Figures 15(a) and 15(b) present the flow solutions of the initial rotor design using the finest mesh from Table 4. Figure 15(a) shows the resolved tip vortices using the dimensionless Q criterion, coloured with pressure coefficients. The modelling strategy clearly resolved the delicate tip vortex system. Figure 15(b) presents the pressure solutions. We noticed that there was a small shock wave near the blade tip around 92.5% to 97.5% r/R. This is further illustrated in the pressure coefficient contours at the corresponding sections.

Blade optimisation results The optimisation study used the coarse grid from Table 4. This is to strike a balance between computational costs and accuracy. The coarse grid results were reasonably accurate, as shown in Table 4. The final optimisation solution will be validated using fine grids. For all FFD lattices, the z offsets were bounded within $\pm 0.125C_0$. Their y offsets have a linear bound within $\pm 1/64 i_s C_0$, where

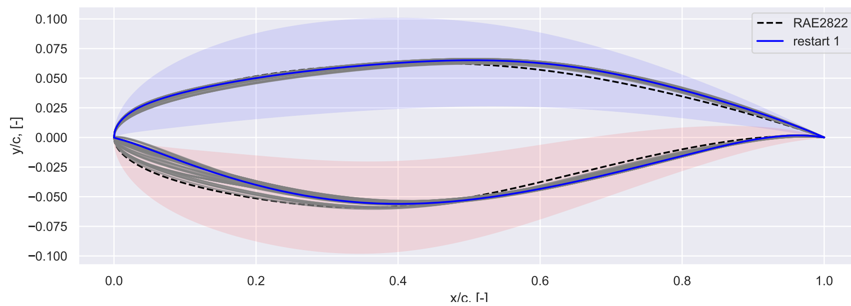

Figure 12. Feasible shapes (shades areas) for relaxed constraints of $Cl \geq 0.82$, $Cd \leq 0.014$, and $Cm \geq -0.092$.

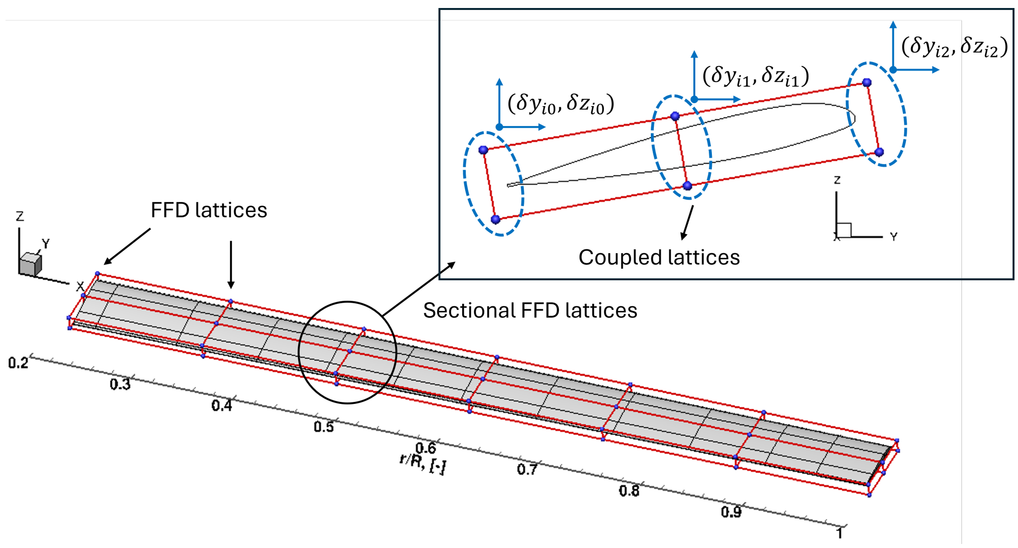

Figure 13. Free Form Deformation(FFD) parameterisation of the initial blade shape. The FFD lattices were only allowed to move along the y and z directions. Lattices along the thickness direction on each section were coupled and shared the same offsets.

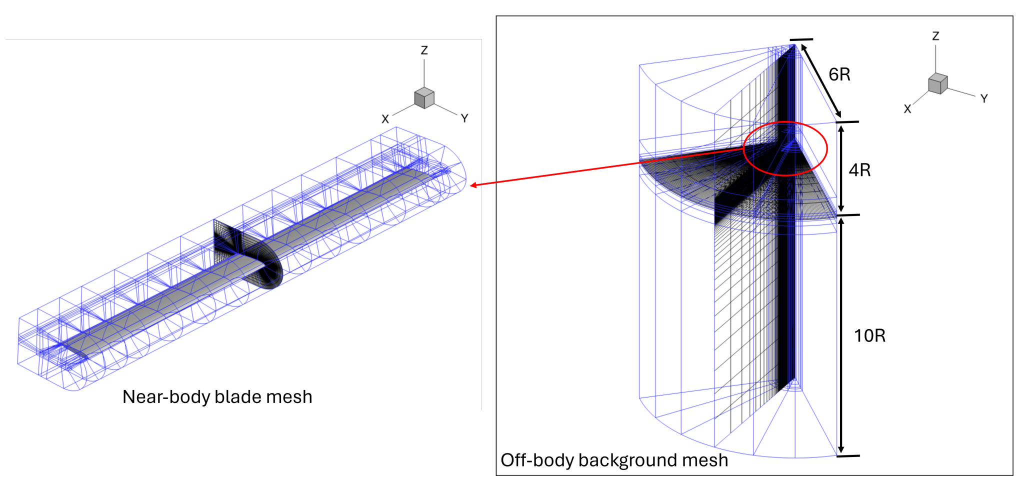

Figure 14. Fully structured, multi-block, overset mesh topology for hover simulation in the rotating reference frame.

Table 4. Mesh convergence study of the initial rotor blade design. All rotors were trimmed to the target C_T of 0.01522 by adjusting the collective pitch angle.

Mesh	Cell number, [million]			$\theta_{0.75}$, [deg]	C_Q, [-]	FoM, [-]
	Near-body	Off-body	Total			
Coarse	0.77	1.02	1.79	10.02	0.001476	0.665
Medium	5.70	1.02	6.72	10.05	0.001432	0.684
Fine	5.70	8.14	13.84	10.01	0.001434	0.682
Test (Ref. 18)	-	-	-	-	-	0.67

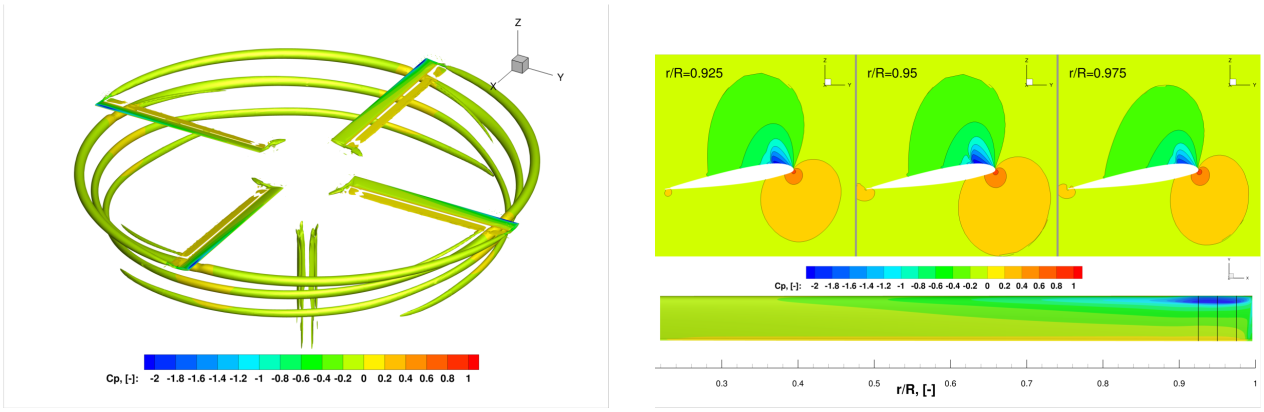

(a) Tip vortices visualised using the dimesionless Q criterion.

(b) Blade surface pressure solutions and tip shock waves.

Figure 15. Flow solutions of the initial blade using the fine mesh in Table 4

$i_s = 1, 2, ..., 7$ is the lattice section index. This was to avoid overlapping near the root and to have more design freedom near the blade tip.

Following the proposed MFMS approach, an initial MFNN was built, linking a simple BEMT model and the high-fidelity CFD results. An initial 100 designs were sampled across the entire design space and were evaluated using the high-fidelity CFD approach. Since there was minimal restrictions on the design variables, some of the designs had infeasible geometries (e.g. inverted/overlapping surfaces), and the high-fidelity evaluations were aborted. Eventually, 63 out of the 100 samples were successfully evaluated and were used to initialise the MFNN.

The first gradient-based optimisation batch started from the baseline HART II blade shape, and 3 successive restarts were carried out, each with a different initial guess indicated by the MFNN. The overall optimisation results are presented in Table 5, and the final optimisation solutions were validated using the medium grid from Table 4. The convergence history of the initial and the restarted optimisation is shown in Figures 16(a) to 16(d), respectively.

Starting from the baseline shape, the initial optimisation improved the *FoM* by about 8.1% compared to the baseline design. The 1st restarted optimisation required 75 iterations to converge, which was substantially increased compared to the initial optimisation. This was due to the inaccuracy of the MFNN estimation at this initial stage. The initial guess for the first restart was far from optimal. Nevertheless, it converged to a different local solution with 7.1% improvement. The 2nd restarted optimisation also converged to about 7.1% improvement in the *FoM*, but it will be shown in later sections that the solution was different from the first restart. The 3rd restart managed to bring the largest improvement of 9.2% in the *FoM*. These are already very similar numbers, and validation results using the medium grid suggest that all solutions have almost the same improvement of about 7.6% to 7.9% as shown in Table 5.

Overall, solutions from gradient-based local searches have very similar performance improvements; however, their shapes are distinct. This is illustrated in Figures 17(a) to 17(g) along with comparisons with the baseline HART II blade shape. The optimisation solutions showed important differences in terms of sectional shapes from root to tip. Compared to the baseline HART II design, the initial optimisation mostly adjusted the chord and twist distributions, with minor changes to the sectional shapes. The 1st restart introduced similar changes for inboard sections before $r/R = 0.85$. Towards the tip, the twist angle and the chord were slightly reduced. The 2nd restarted optimisation reduced the camber for sections before $r/R = 0.55$ along with chord and twist changes. The 3rd restart also reduced the camber near $r/R = 0.4$ and $r/R = 0.55$. Near the tip, the sectional camber and chord length were both increased.

Blade platform changes are detailed in Figures 18(a) to 18(d), which further highlight the shape differences. In Figures 18(c) and 18(d), the sweep and an-/dihedral were measured by the

Table 5. Performance comparisons between the baseline and optimised designs (pending fine grid validation).

	Coarse grid optimisation results		Medium grid verification	
Optimisation case	FoM, [-]	Relative improvement, [-]	FoM, [-]	Relative improvement, [-]
Baseline	0.665	-	0.684	-
Initial optimisation	0.719	8.1%	0.737	7.7%
Restart 1	0.712	7.1%	0.738	7.9%
Restart 2	0.712	7.1%	0.736	7.6%
Restart 3	0.726	9.2%	0.738	7.9%

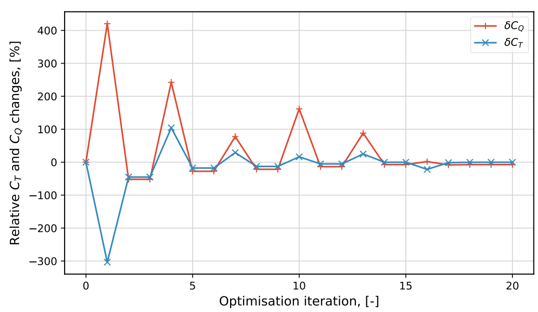

(a) Optimisation history of the initial optimisation case.

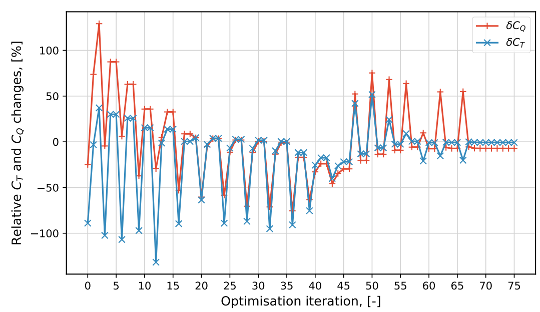

(b) Optimisation history of the 1st restart.

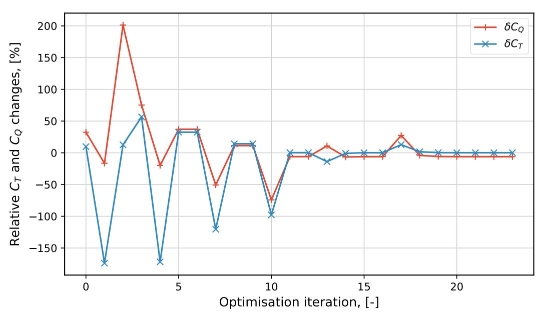

(c) Optimisation history of the 2nd restart.

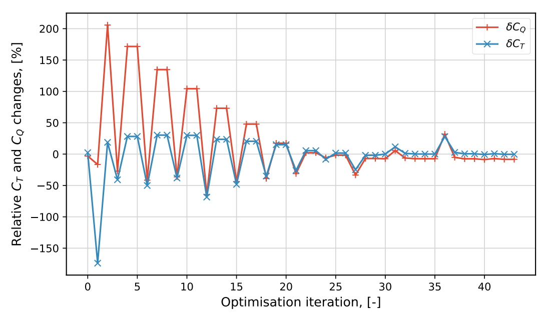

(d) Optimisation history of the 3rd restart.

Figure 16. Gradient-based optimisation convergence history.

quarter-chord point offset in the y and z directions, respectively. In Figures 18(a) to 18(c), the changes were normalised by the reference chord length of the original HART II blade $C_0 = 0.121m$.

In terms of chord changes in Figure 18(a), the initial optimisation resulted in minor changes compared to the baseline shape. The chord was slightly increased by about 2% C_0 for sections below $r/R = 0.7$, while the tip chord was reduced by about 1% C_0. For the 1st restart, the chord change was much more salient, especially near the tip. From root to about $r/R = 0.65$, the chord was increased by about 5% C_0. The chord was then quickly reduced from $r/R = 0.65$ towards the tip following a nonlinear trend. At the tip, the chord was reduced by about 10% C_0. The 2nd restart had a similar chord distribution, but near the tip, the chord was further reduced to about 15% C_0. The 3rd restart had almost identical chord distributions as the initial optimisation before $r/R = 0.6$. The chord was then slightly reduced till $r/R = 0.75$ and quickly increased to about 8% towards the blade tip.

Changes in the blade twist are drastic. All optimisation results derived nonlinear twist distributions, increasing twist near the root and reducing twist towards the tip. These twist changes, combined with the chord changes in Figure 18(a), indicate that the optimisation shifted loads towards the blade inner sections. The 3rd restart even introduced a negative local pitch angle of about -2.5 degree near the tip. However, the tip sectional chord and camber were both increased as shown in Figure 17(g), which may cancel out the twist change. This highlights the multi-modality of the design problem.

As for sweep, the changes were mild, but the differences were clear. The sweep changes in Figure 18(c) were less than 5% C_0 except for the 3rd restart. The initial optimisation intro-

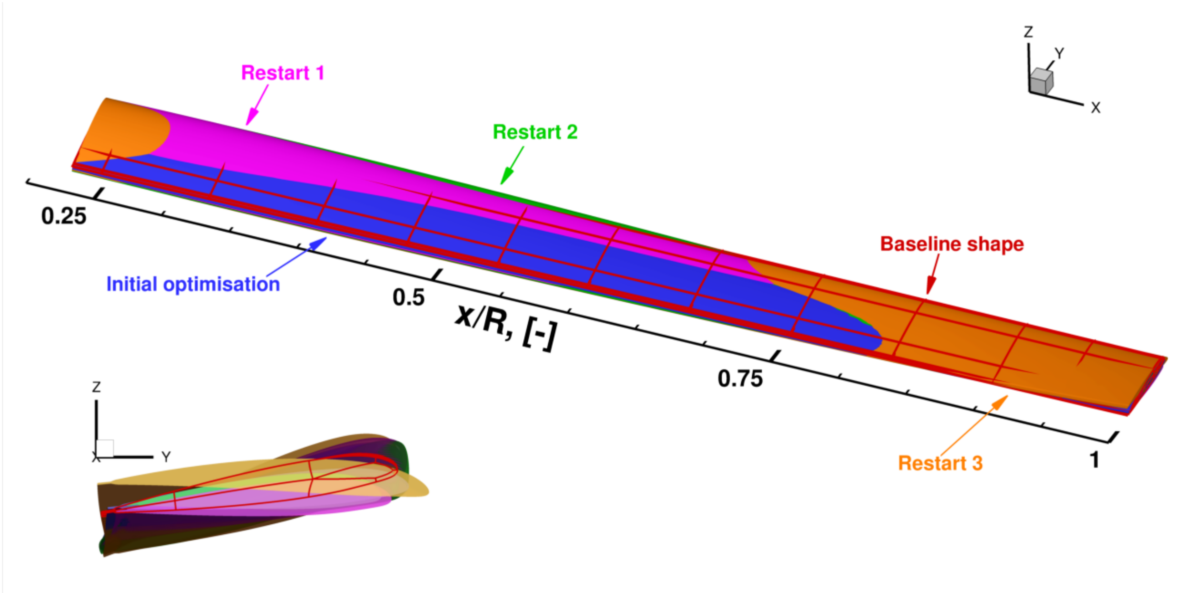

(a) Overall blade shape comparisons.

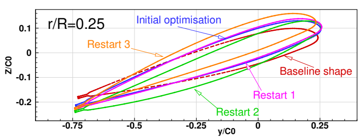

(b) Section $r/R = 0.25$.

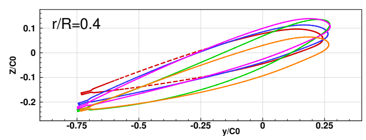

(c) Section $r/R = 0.40$.

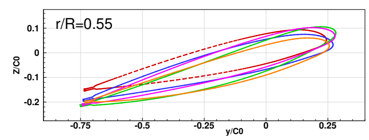

(d) $r/R = 0.55$.

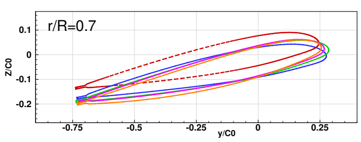

(e) $r/R = 0.70$.

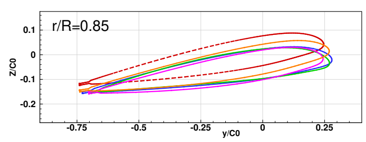

(f) $r/R = 0.85$.

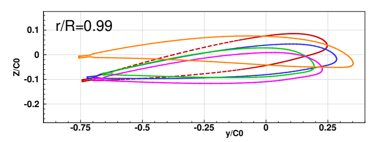

(g) $r/R = 0.99$.

Figure 17. Comparisons between the baseline and optimised blade shapes.

duced a linear forward sweep throughout the blade, reaching about 4% near the tip. This forward sweep should help delay the formation of shock waves. The 1st and 2nd restarts, on the other hand, introduced a similar forward sweep for the inner sections below $r/R = 0.7$, but quickly switched to a backwards sweep towards the tip. This is the typical blade platform change that delays shock waves. The 3rd restart had the highest non-linear forward sweep. Before $r/R = 0.7$, the forward sweep was increased then reduced, peaking at about 2.5% near $r/R = 0.5$. From $r/R = 0.7$ towards the blade tip, the forward sweep was drastically increased to about 8%.

The anhedral/dihedral changes were similar for the initial optimisation and the 1st and 2nd restarts, as shown in Figure 18(d). All these results lifted sections near the root by about 2.5% C_0 and moved down sections near the tip by about 5% to 7.5% C_0. The 3rd restart had a diffrent distribution, with higher lifting up at around 4% near the root, and more flattened sections from $r/R = 0.55$ towards the tip. These changes effectively introduced a non-linear anhedral towards the tip, which should help alleviate the tip vortex induction.

The blade loading distributions were extracted and are shown in Figures 19(a) and 19(b). For all optimisation cases, the tip loading was reduced and the loads were shifted towards the blade inner sections, as shown in Figure 19(a). This is also reflected in the tangential force distributions in Figure 19(b). The baseline HART II blade had a linearly increasing tangential force distribution towards the tip. All optimisation solutions unloaded the tip and shifted more loads towards the inner sections, though each had a different distribution. The sectional tangential forces are directly linked to the torque and power consumption in this case. More loading on inner sections means lower overall torque and power. The loading distributions were consistent with previous optimisation studies of the HART II blade (Ref. 19).

The loading changes are closely correlated with the shape changes discussed above. All optimisation solutions had an increase in twist that helped shift the load to the inner sections. The increased chord lengths near the inner sections had the same effect. For the 1st and 2nd restarts, the reduced chords near the tip helped reduce the tip loading too. However, due to

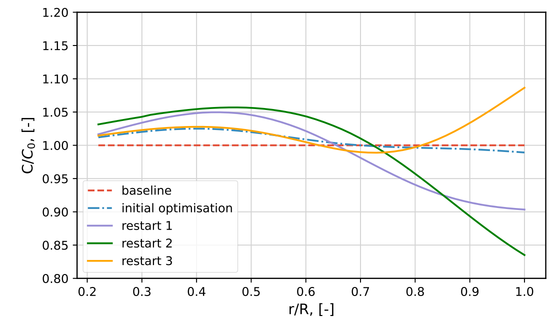

(a) Chord distribution.

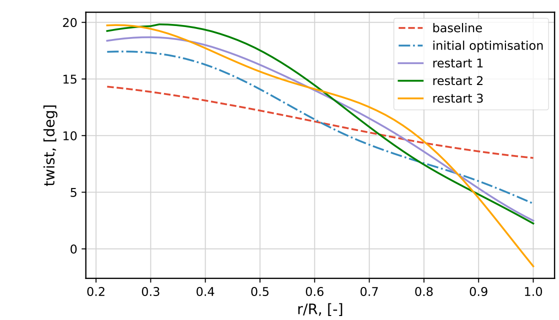

(b) Twist angle distribution (collective pitch included).

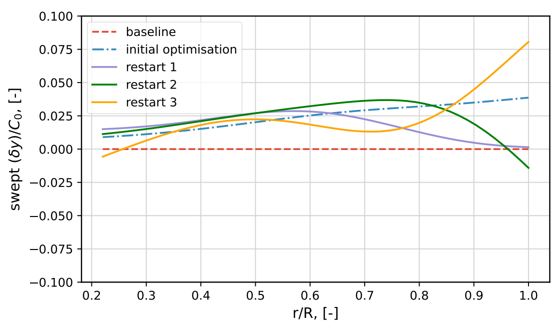

(c) Sweep distribution (measured using sectional quarter chord point offsets).

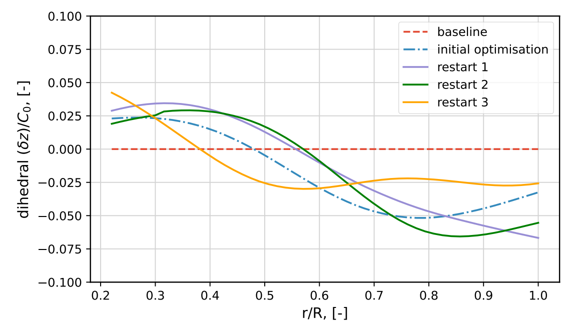

(d) Anhedral/dihedral distribution (measured using sectional quarter chord point offsets).

Figure 18. Comparisons of the baseline and optimised blade planforms.

the high freedom of the parameterisation, there are some conflicting geometric features. For instance, the chord length of the 3rd restart was substantially increased near the tip, and the sectional camber was increased too. However, the tip loading was still reduced due to negative twist angles at the tip. This highlights the multi-modality of the design space, as different design features may cancel out and lead to the same overall performance. More practical constraints, e.g. bending moments, may help eliminate some of the solutions.

Figures 20(a) to 20(d) present the flow solutions near the rotor blade tip. Compared to the baseline solution in Figure 15(b), it is clear that all optimisation cases managed to either eliminate or reduce the tip shock waves, via combined planform and subtle sectional shape changes. This effectively reduced the tip tangential forces and hence the overall torque and power consumption. The 3rd restart in Figure 20(d) completely eliminated the shock waves at the tip.

Last but not least, Figures 21(a) and 21(b) present the distribution of all rotor designs during the optimisation process. These included the initial 63 sampling to establish the MFNN, and samples from each gradient-based local search. Note that these results were computed using the coarse grid from Table 4 to save computational costs, hence the FoM results were slightly lower. Nevertheless, it is interesting to notice in 21(a) that most of the rotor designs formed an efficiency front near $FoM = 0.7$. Every few samples managed to reach $FoM = 0.75$ at higher thrust. Figure 21(b) provides a close-up view near the baseline design. The optimisation searches were again concentrated near the thrust constraints, and many samples offered practically similar performance. It can be noted that samples from the 3rd restart focused on regions with better performance thanks to the successively improved MFNN.

Overall, this practical optimisation case demonstrated the proposed MFMS approach's capability in exploring bumpy design spaces and finding families of improved designs. The optimisation results had distinct shapes, but many offered similar performance. This was achieved by combinations of changes in the twist, chord, sweep, an-/dihedral, and sectional shapes. This highlights both complexity and freedom in aerodynamic shape design optimisation, especially for such loosely constrained problems.

IV. CONCLUSIONS

This work presented a multi-fidelity, multi-start (MFMS) approach for aerodynamic shape optimisation. The details of the methodology were elaborated, and the implementations and tools were verified and demonstrated via canonical optimisation problems, benchmark aerofoil optimisation cases, and practical rotary wing design optimisation. From the current results, the following conclusions can be drawn:

1. With the help of the MFNN, linked with physical models, the proposed MFMS hybrid optimisation approach is

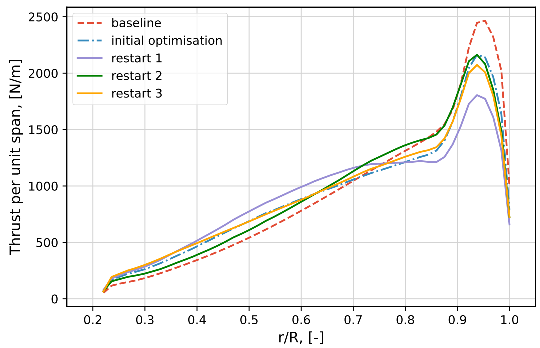

(a) Thrust distributions along the blade.

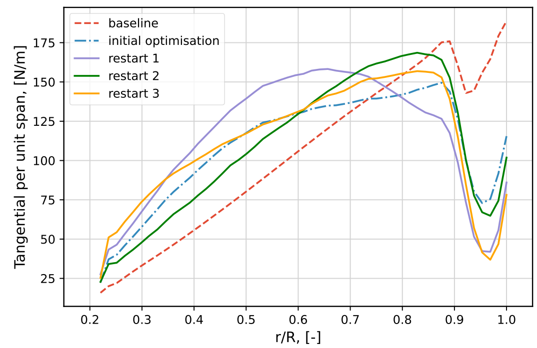

(b) Tangential force distribution along the blade.

Figure 19. Comparisons of blade loadings between the baseline and the optimised shapes.

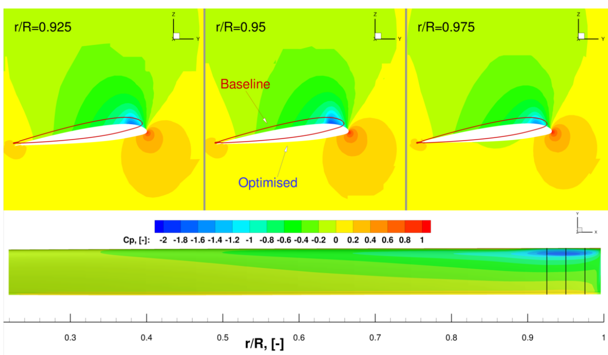

(a) Initial optimisation.

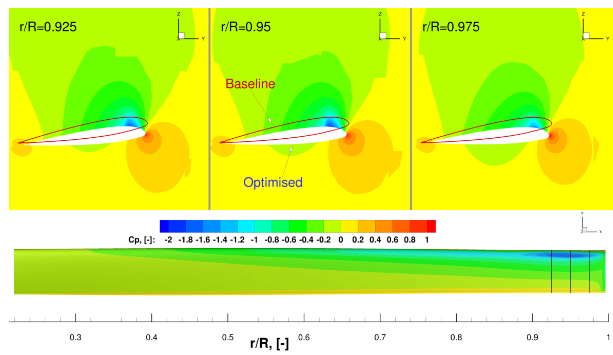

(b) Restart 1.

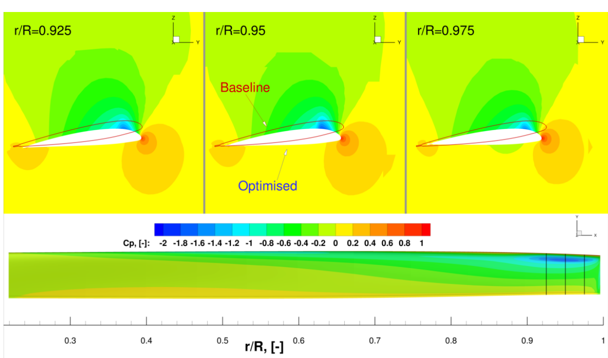

(c) Restart 2.

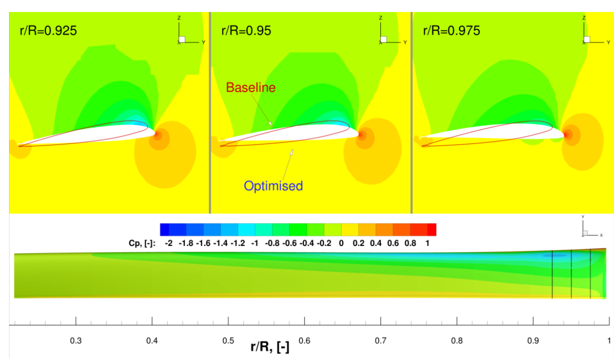

(d) Restart 3.

Figure 20. Flow solutions of the optimised blades on the surface and near the blade tip (Tip $Ma = 0.64$).

shown to be capable of evading sub-optimal local solutions found by gradient-based optimisation. Evaluations using the multi-modal Styblinsky-Tang function suggest that the overall cost was only linearly scaled with the number of restarts.

2. For the benchmark RAE2822 aerofoil optimisation case, the MFNN reduced the drag by 80 counts. This was achieved by delaying the shock formation. Four restarts were carried out, and all converged to the same solution in this tightly constrained problem. The MFMS also demonstrated capabilities to assimilate failed optimisations, highlighting the method's robustness and flexibility. Though trained on local optimisation history, the MFNN was capable of providing reasonable predictions for the entire design space as the restarts progressed.

3. The MFMS approach was successfully applied to the HART II rotor blade optimisation case. The shape parameterisation allowed changes in chord, twist, an-/dihedral, sweep, and sectional shapes. Multimodality was expected due to conflicting geometric features. The

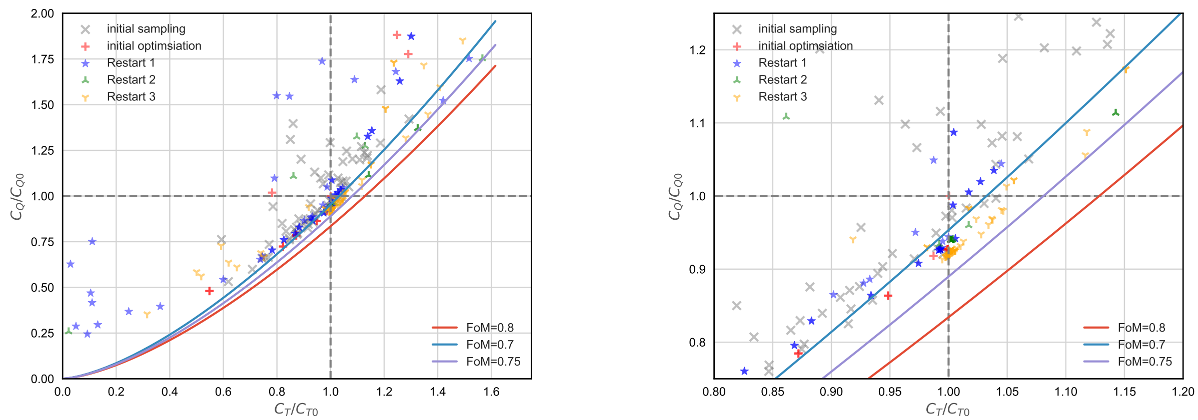

(a) Overall thrust and torque distributions.

(b) Close-up view near the baseline HART II design.

Figure 21. Thrust and torque distributions of all samples collected during the optimisation. Note that these samples were computed using the coarse grid of Table 4, which tends to give slightly lower *FoM* predictions.

MFMS approach managed to find 4 local solutions with very similar improvements in the *FoM* of about 8%, but with distinct geometry details. These 4 solutions shared common features, such as increased twist and non-linear anhedral, but design details were distinct. All solutions managed to upload the tip by weakening or eliminating the tip shock wave and shifting the loading further inboard. These results highlighted the proposed MFMS approach's capability in handling multi-modal design spaces.

Future work will expand the framework to include multidisciplinary optimisation, especially for aeroelastic considerations of rotor blades. For the InDeWo rotary wing design case, we will explore more practical constraints, larger design spaces, and further design options.

ACKNOWLEDGEMENT

The authors would like to acknowledge the MoHRSS Exchange Scheme S20240126 supporting Dr Yu Luo.

REFERENCES

1. Joaquim RRA Martins. Perspectives on aerodynamic design optimization. In *AIAA Scitech 2020 Forum*, page 0043, 2020.

2. Nicolas P Bons, Xiaolong He, Charles A Mader, and Joaquim RRA Martins. Multimodality in aerodynamic wing design optimization. *AIAA Journal*, 57(3):1004–1018, 2019.

3. Oleg Chernukhin and David W Zingg. Multimodality and global optimization in aerodynamic design. *AIAA journal*, 51(6):1342–1354, 2013.

4. Alexander IJ Forrester, András Sóbester, and Andy J Keane. Multi-fidelity optimization via surrogate modelling. *Proceedings of the royal society a: mathematical, physical and engineering sciences*, 463(2088):3251–3269, 2007.

5. Donald E Myers. Matrix formulation of co-kriging. *Journal of the International Association for Mathematical Geology*, 14(3):249–257, 1982.

6. Marc C Kennedy and Anthony O'Hagan. Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87(1):1–13, 2000.

7. Xuhui Meng and George Em Karniadakis. A composite neural network that learns from multi-fidelity data: Application to function approximation and inverse pde problems. *Journal of Computational Physics*, 401:109020, 2020.

8. Jun Tao and Gang Sun. Application of deep learning based multi-fidelity surrogate model to robust aerodynamic design optimization. *Aerospace Science and Technology*, 92:722–737, 2019.

9. Jethro R Nagawkar, Leifur T Leifsson, and Ping He. Aerodynamic shape optimization using gradient-enhanced multifidelity neural networks. In *AIAA SciTech 2022 Forum*, page 2350, 2022.

10. Tao Zhang, George N Barakos, Malcolm Foster, et al. Multi-fidelity aerodynamic design and analysis of propellers for a heavy-lift evtol. *Aerospace Science and Technology*, 135:108185, 2023.

11. Tao Zhang and George N. Barakos. High-fidelity numerical analysis and optimisation of ducted propeller aerodynamics and acoustics. *Aerospace Science and Technology*, 113:106708, 2021.

12. Massimo Biava and George N Barakos. Optimisation of ducted propellers for hybrid air vehicles using high-fidelity cfd. *The Aeronautical Journal*, 120(1232):1632–1657, 2016.

13. Florian R Menter. Two-equation eddy-viscosity turbulence models for engineering applications. *AIAA journal*, 32(8):1598–1605, 1994.

14. D. Kraft. Algorithm 733: Tomp-fortran modules for optimal control calculations. *ACM Transactions on Mathematical Software*, 20(3):262–281, 1994.

15. S.G. Johnson. The nlopt nonlinear-optimization package. https://nlopt.readthedocs.io/en/latest/, retrieved 3 Sep 2020.

16. Peter D. Sharpe. *Accelerating Practical Engineering Design Optimization with Computational Graph Transformations*. PhD thesis, Massachusetts Institute of Technology, 2024.

17. Xiaolong He, Jichao Li, Charles A Mader, Anil Yildirim, and Joaquim RRA Martins. Robust aerodynamic shape optimization—from a circle to an airfoil. *Aerospace Science and Technology*, 87:48–61, 2019.

18. Berend G van der Wall. A comprehensive rotary-wing data base for code validation: the hart ii international workshop. *Aeronautical Journal*, 115(1164):91, 2011.

19. Gunther Wilke, Joëlle Bailly, Keita Kimura, and Yasutada Tanabe. Jaxa-onera-dlr cooperation: results from rotor optimization in hover. *CEAS Aeronautical Journal*, 13(2):313–333, 2022.

20. Keita Kimura, Gunther Wilke, Joëlle Bailly, and Yasutada Tanabe. Jaxa-onera-dlr cooperation: results from rotor optimization in forward flight. In *ERF 2022*, 2022.

21. Thomas W Sederberg and Scott R Parry. Free-form deformation of solid geometric models. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, pages 151–160, 1986.

CLASSIFICATION & COPYRIGHT