A study of Stall-Induced Vibrations using Surrogate-Based Optimization

Santhanam, Chandramouli; Riva, Riccardo; Knudsen, Torben

Published in:
Renewable Energy

DOI (link to publication from Publisher):
10.1016/j.renene.2023.05.054

Publication date:
2023

Document Version
Publisher's PDF, also known as Version of record

Renewable Energy 214 (2023) 201–215

Contents lists available at ScienceDirect

Renewable Energy

journal homepage: www.elsevier.com/locate/renene

A study of Stall-Induced Vibrations using Surrogate-Based Optimization

Chandramouli Santhanam [a,*], Riccardo Riva [b], Torben Knudsen [a]

[a] Department of Electronic Systems, Aalborg University, Fredrik Bajers Vej 7, Aalborg, Øst, 9220, Denmark
[b] Department of Wind Energy, Technical University of Denmark, Frederiksborgvej 399, Roskilde, 4000, Denmark

ARTICLE INFO

Keywords:
Aeroelastic stability
Stall Induced Vibrations
Surrogate based optimization
Delaunay triangulation
Wind turbine

ABSTRACT

Stall-Induced Vibrations (SIV) are an important design consideration for wind turbine blade design, especially for large, modern, wind turbines with highly flexible blades. Their severity depends on both the inflow and the structural characteristics of the blades. Studying SIV has a high computational cost, because it requires high-fidelity aeroelastic simulations, and potentially a large number of input variables. In an effort to reduce the computational cost of the domain exploration, in this work we have adopted a Surrogate-Based Optimization (SBO) framework. This way, the combination of input variables that lead to SIV can be explored with the minimum number of aeroelastic simulations. The proposed SBO framework can use any type of surrogate model, and leverages Delaunay triangulation to iteratively select samples to refine the surrogate model. The occurrence and severity of SIV on the IEA 10MW turbine is studied in a five variables space consisting of: wind speed, yaw angle, vertical wind shear, wind veer, and atmospheric temperature. A well-trained surrogate model is developed and used to predict the damping ratio of the first blade edgewise mode in the entire inflow space at a reduced computational cost. Sensitivity analysis of the predicted damping ratio shows that yaw angle is the most influential variable, while temperature is the least influential variable in terms of inflow conditions that can lead to the occurrence of SIV. Inflow conditions with a moderate yaw angle (around 10–25 deg), high wind speeds, and moderate to high negative veer are found to lead to severe SIV. This study should serve as a guiding tool to decide the scope of the more computationally expensive simulations such as high-fidelity CFD-based aeroelastic simulations which can provide a more accurate description of SIV.

1. Introduction

Stall-Induced Vibrations (SIV) are an aeroelastic instability that might happen when large portions of the wind turbine blade operate in moderate stall (defined in this work as an Angle of Attack (AoA) greater than 15 deg but lower than 40 deg), which leads to large internal loads. The typical negative lift — AoA gradient ($\partial C_L/\partial\alpha < 0$) in this region leads to a negative aerodynamic damping in the sections of the blade, which leads to the lift force being in phase with the velocity of the vibrating blade [1,2]. Although modern pitch-regulated turbines do not operate in stall like the older stall-regulated ones, computations show that SIV can lead to extreme edgewise vibrations in the turbine blades, especially during standstill conditions [3]. Standstill conditions exist for example during installation, when the turbine is stopped for maintenance, and due to extreme wind speeds. Apart from the risk that these vibrations can become vigorous, they also contribute to fatigue loading of the blades which leads to a reduction in the lifetime of the blades. With the need to lower the Levelised Cost of Energy driving the design towards larger and more flexible wind turbine blades, SIV is an important design consideration as the loads generated might severely damage the blades.

SIV have been studied using the Blade Element Momentum theory (BEM) based solvers in the limited yaw angle range around moderate stall regions. The Advanced Aerodynamic Tools for Large Rotors (AVATAR) project details a comprehensive SIV study of BEM-based aeroelastic solvers against a higher fidelity CFD based aeroelastic solver [4]. The results show that BEM-based solvers tend to over-predict standstill instabilities due to the utilization of static airfoil data. But since the basic assumption in BEM theory that the rotor can be modelled as a disc does not hold in standstill conditions, other models such as the Near Wake model (NW) have been proposed for aerodynamic modelling in standstill conditions [5]. The near wake model has been validated against analytical solutions for an elliptical wing and measurements against the NREL Phase VI rotor in standstill conditions. While the NW model deviates from the measurements in certain conditions, it is still a better choice than the Blade Element Momentum (BEM) model. It is also to be noted that the validity of such engineering models decreases as we move into the deep stall regions. Although engineering models suffer such problems in predicting

* Corresponding author.
E-mail address: chsa@es.aau.dk (C. Santhanam).

https://doi.org/10.1016/j.renene.2023.05.054
Received 16 December 2022; Received in revised form 24 March 2023; Accepted 11 May 2023
Available online 17 May 2023

stall behaviour in standstill conditions, they offer a huge reduction in the computational cost, while still having respectable validity in moderate stall regions. But, it can be seen that even with the cost reduction offered by the engineering models, the exploration of the inflow space is still a very computationally expensive process, and the cost increases drastically with the increasing number of dimensions to explore. For such systems with expensive target functions, Surrogate Based Optimization (SBO) is a technique that has been used to explore the input space effectively. In the context of wind turbine design, surrogate models have been used to study SIV [6], and in other similar applications involving costly evaluations, such as evaluation of site-specific loads [7], finding the optimal blade geometry parameters and control features [8], and study of the effect of multiple parameters on the blade design [9].

In SBO, an approximate but inexpensive alternative to the target function, called the surrogate model, is used to establish a quick-to-evaluate relationship between the input and output variables. The surrogate models can then be optimized via standard optimization techniques to identify points of interest in the domain. Popular choices of surrogate models include Polynomial Response Surfaces, Artificial Neural Networks (ANN), Kriging, Radial Basis Functions (RBF) etc. SBO techniques typically employ three key steps: (1) Selection of initial samples for simulation, (2) Construction of a surrogate model and (3) An adaptive sampling strategy to select samples for subsequent simulations. The process is repeated until the problem specific goals are reached. In many SBO algorithms, the adaptive sampling strategy is suited to a specific type of surrogate model, usually stochastic processes such as Kriging. The use of stochastic processes for exploiting an input space dates back to [10], and since then has gained attraction. A popular algorithm that has been successfully applied in many problems is the Efficient Global Optimization (EGO) algorithm proposed by [11], based on Kriging and a sampling strategy called Expected Improvement suited to Kriging. Other similar adaptive sampling techniques have been developed based on Kriging, such as Probability of Improvement, Mean Squared Error, and Confidence Bounds [12]. Some of these algorithms have been extended to RBF [13] and ANN [14], but as noted by [15], adaptive sampling strategies are predominantly based on Kriging because of the ready availability of the variance estimate along with the prediction. But depending on the nature of the unknown target function, it may be the case that a different choice of surrogate model works better [16,17]. Furthermore, most of these sampling methods have the primary goal of searching the parameter space for a minimum, while it is also often useful to obtain a well-trained surrogate that can be used for purposes other than finding the minima.

The focus of the current work is to propose methods and tools to study SIV in a given inflow space. To the authors' best knowledge, the problem of effectively exploring the behaviour of SIV using an SBO framework has not been attempted before and in this paper, we propose an SBO type framework to study SIV effectively. The proposed SBO framework is independent of the surrogate type, and in addition to identifying the most critical vibrations, has a wider scope to additionally identify regions in the domain where the target function can take critical values (below a certain threshold) with good confidence. The article is organized as follows. Section 2 describes the problem statement and scope of this work. Section 3 describes the simulation setup and methods to characterize SIV. Section 4 describes the proposed framework and illustrates the functionality of the framework on standard test functions and the considered SIV problem. Section 5 explains the effect of the considered variables on SIV.

2. Problem statement and scope

The aim of this work is to formulate an SBO framework that can, with the optimal number of simulations

1. Help study the effect of inflow and environmental conditions on SIV
2. Identify the conditions that lead to critical SIV in a wind turbine

The proposed SBO framework should not be surrogate-specific, i.e. it should be able to work with any choice of surrogate model. Studies of SIV in wind turbines [18,19] show that the occurrence and severity of SIV depend on inflow conditions, airfoil characteristics and structural characteristics of the blade. With this motivation, this framework is used to demonstrate the analysis of the influence of the following variables on SIV

- Wind Speed
- Yaw Angle
- Vertical Wind Shear
- Wind Veer
- Temperature

The parameter space, or inflow space, is defined as the five dimensional space made up of these variables. While there are other parameters that affect SIV, the number of variables is limited to five due to computational considerations.

Table 1
Ranges of the considered variables.

Variable	Range
Wind speed	[25, 60] ms^{-1}
Yaw Angle	[0, 40] deg
α	[−0.1, 0.4]
a_φ	[−1.2, 0.5]
Temperature	[−15, 20]°C

3. Simulation setup and SIV characterization

This section describes the turbine model, considerations in the simulation setup and the method used to quantify the effect of SIV. The IEA 10 MW turbine [20] is chosen for investigation. The occurrence and characteristics of SIV in this turbine are studied for a parked rotor with a 90 deg pitch angle, 0 deg azimuth angle (blade 1 pointing upwards), constant wind conditions and only the blades are considered flexible. Hence, only the blade modes are considered and not the turbine ones. The inflow conditions and simulation setup are similar to the considerations in the investigations of SIV in the AVATAR rotor [4]. The aeroservoelastic tool HAWC2 [21] is used to simulate the motion of the wind turbine. The Near Wake Model [5] is used for aerodynamic modelling, and the dynamic stall is modelled through the MHH Beddoes dynamic stall method [22]. A HAWC2 simulation is preferred over CFD simulations because of the computational time involved. While a typical HAWC2 simulation takes around 20 minutes, CFD simulations typically take much larger computational time. Additionally, the complexity of setting up the simulations is also higher than HAWC2 simulations. For an initial domain exploration, it is advantageous to use solvers like HAWC2, which is still costly for a 5 dimensional problem. The results of this initial exploration- (conceptual design phase) can help decide the focus of higher fidelity CFD simulations and lifting line methods that can be used for a detailed study of the instabilities.

The ranges of the considered variables are shown in Table 1.

The wind speed range above the cut-out wind speed is considered because it is expected to induce substantial vibrations. The yaw angle range is chosen so as to avoid sections of the airfoil operating in deep stall where the engineering models have limited validity. The shear exponent (α) is chosen to cover negative and positive shear. Wind veer is modelled through a coefficient (a_φ) as proposed in [23].

$$\begin{aligned}\Delta\varphi(z) &= \varphi(z) - \varphi(z_{hub}) \\ &\approx a_\varphi e^{-\sqrt{z_{hub}/h_{ME}}} \frac{z - z_{hub}}{\sqrt{z_{hub} h_{ME}}} \left(1 - \frac{z - z_{hub}}{2\sqrt{z_{hub} h_{ME}}} - \frac{z - z_{hub}}{4 z_{hub}}\right),\end{aligned} \quad (1)$$

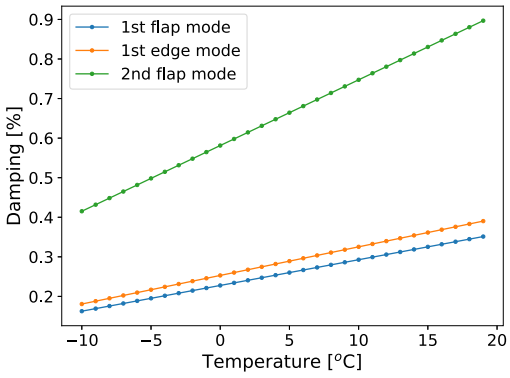

Fig. 1. Variation of damping of the first three modes with temperature.

where z represents the height of the point, $\Delta\varphi(z)$ is the veer angle at height z, z_{hub} is the hub height, h_{ME} is the modified Ekman atmospheric boundary layer depth, here set to 500 m. In Eq. (1), a_φ is a scaling factor to represent the veer, while the last two terms which depend upon the height z give the wind profile. a_φ depends on wind parameters and can be considered site-specific, and a recommended range given in [23] is [-1.2, 0.5]. As it can be seen from Eq. (1), positive values of a_φ correspond to positive veer angles above the hub and negative veer angles below the hub (veering), while negative values of a_φ correspond to negative veer angles above the hub and positive veer angles below the hub (backing). Also, $a_\varphi = 0$ corresponds to zero veer, and values further from zero imply higher veer in the respective direction.

We assume that temperature modifies the air density and blade structural damping, with the first modelled using the ideal gas equation and International Standard Atmosphere. With regards to the variation of structural damping, studies on the material properties of composites similar to the ones used in wind turbine blades show that in general, the structural damping increases with temperature [24,25]. In this work, a linear relationship between temperature and structural damping is assumed. The reference temperature for the normal damping values is considered as 15 °C, and a decrease of 50% in the structural damping is assumed at −10 °C. [26]. Thus, the variation of structural damping of mode i with temperature can be expressed as

$$\zeta_i(T) = \zeta_i^{ref}\left(0.7 + \frac{T}{50}\right) \tag{2}$$

where $\zeta_i(T)$ represents the structural damping of mode i at temperature T and ζ_i^{ref} represents the structural damping of mode i at 15 °C. While the exact variation of structural damping with temperature depends on the fibre material used in the blades and their orientation, experimental studies on glass fibre composites have reported the amount of variation considered in this study. The variation of the structural damping of the first three modes with temperature considering the variation described above is shown in Fig. 1.

In HAWC2, the structural damping is modelled via a stiffness-proportional damping model. As per this model, in the calculation of the damping matrix, factors k_x, k_y, k_z are multiplied to moments of inertia of the element about the x,y and z axes, I_x, I_y, and I_z, and inserted into the damping matrix. The details about the calculation of the damping matrix can be found in [27]. Thus, the structural damping of the blade is a function of the co-efficients. Typically the coefficients k_x, k_y, k_z are tuned until the required structural damping of the first three blade modes is achieved. For a given value of the coefficients and the structural model of the turbine blade, the structural damping can be quickly calculated using the program HAWCStab2 [28]. To calculate the value of the coefficients for a given temperature, the new values of damping of the first three modes are first calculated, and the value of $\mathbf{k} = [k_x, k_y, k_z]$ is tuned using HAWCStab2 until the new damping values for the first three modes are achieved. Thus $\mathbf{k}$ is found from the following optimization problem,

$$\mathbf{k} = \underset{\mathbf{k}}{\text{argmin}}\left(\sum_{i=1}^{3}(\zeta_i(\mathbf{k}) - \zeta_{i,req}(T))^2\right)$$

where $\zeta_i(\mathbf{k})$ represents the structural damping of the ith mode for a given value of $\mathbf{k}$, and $\zeta_{i,req}(T)$ represents the required structural damping of the ith mode for a given temperature T. The optimization problem is setup framework OpenMDAO [29], and solved using the Sequential Least SQuares Programming (SLSQP) algorithm [30].

SIV characterization

The edgewise bending moment at the root of blade 1 is chosen as the response to study SIV. A growth of edgewise bending moment with time implies the presence of SIV while a decay implies absence. The behaviour of the wind turbine for the different inflow conditions broadly falls into three categories: linearly stable, linearly unstable and non-linearly unstable (where the system enters a limit cycle after some time). The corresponding edgewise bending moments are shown in Fig. 2. It is generally observed that the responses where the system behaves as non-linearly unstable grow quicker than when the system behaves as linearly unstable. Hence, these nonlinearly unstable scenarios are more negatively damped than the linearly unstable ones, implying more critical vibrations.

The stability of the system with respect to SIV is characterized by identifying the damping ratio of the first blade edge mode. The damping ratio is chosen over other measures based on the amplitude such as peak-to-peak amplitude or max. amplitude because the BEM simulations have a tendency to predict amplitudes that are not physical, but nonetheless qualitatively comparable to CFD based simulations. The first 50 s are discarded as transient. The response is limited in time to be in the linear range to make it suitable for a damping identification method. One of the most widely used methods to identify the damping ratio is the logarithmic decrement method. However, the classic logarithmic decrement method is sensitive to the choice of peaks, and different peaks in the same response can lead to identification of significantly different damping ratios. Hence, we have used a variant of the logarithmic decrement method ratio which avoids this problem by considering many response peaks and hence is more robust.

The damping ratio is identified using this method as follows: The damped frequency of the dominant blade mode (first edge mode in the case of SIV) can be identified as the frequency corresponding to the peak of the Power Spectral Density (PSD) of the response. This is justified in this analysis because we consider an isolated blade (stiff tower, hub and nacelle) with no turbulent inflow, and hence the peaks of the PSD of the response is expected to be around the damped modes of the blade. Let ω_d represent the damped frequency of the first blade edge mode. The response, band-pass filtered around ω_d can be written as

$$y(t) = Ae^{-\zeta\omega_n t}\sin(\omega_d t + \phi), \tag{3}$$

where ζ represents the mode damping ratio, and ω_n the mode natural frequency. The term $Ae^{-\zeta\omega_n t}$ indicates the growth/decay of the peaks of the response, and so by indicating with (t_i, p_i) the time and value of the ith response peak, it can be written that

$$\log(p_i) = -\zeta\omega_n t_i + c, \tag{4}$$

where c is the line intercept. Thus, the mode damping ratio can be calculated from the slope of the straight line fitted to the logarithm of the bandpass filtered response peaks and the corresponding time instances. ω_n is calculated using HAWCStab2, and the straight line is fit with the least-squares method.

However, because of the finite frequency resolutions, and in some cases, the spectral leakage, the frequency at which the PSD peak happens varies slightly for different time series, and this poses a challenge

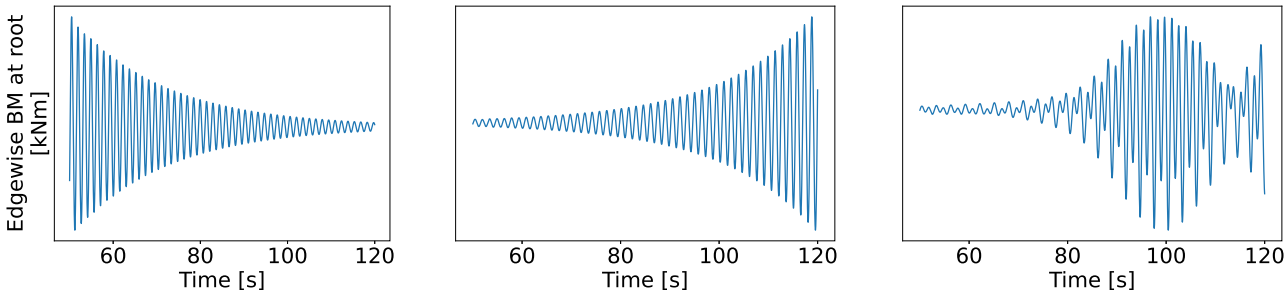

(a) *Linearly Stable* (b) *Linearly Unstable* (c) *Nonlinearly Unstable*

Fig. 2. Typical responses of the wind turbine blade under different yaw angles.

to the consistent identification of ω_d from the PSD peaks. So, to make the damping identification routine more robust and consistent, $y(t)$ in Eq. (3) is obtained by bandpass-filtering the response around ω_n. This is justified because ω_d and ω_n are related as $\omega_d = \omega_n\sqrt{1-\zeta^2}$, and for low values of ζ, $\omega_d \approx \omega_n$. This approximation is valid in the context of SIV, which are not expected to have very high values of ζ.

4. Surrogate based optimization setup

The proposed framework to study SIV includes an experiment design, a simulation environment, a surrogate model and an adaptive sample generation technique. A brief description of these components and the SBO process is given in this section.

4.1. Design of experiments

The initial samples for studying the parameter space are generated using a Latin Hypercube Sampling (LHS) method using the library `OpenTURNS` [31]. Such a sampling tends to distribute points all over the parameter space, avoiding potential clustering of points. To ensure that the initial samples contain points with a negative damping (presence of SIV), the following method is used to generate the samples. Roughly, the following conditions are guessed favourable for SIV — high wind speeds, moderate yaw angles, and low temperatures. Accordingly, trapezoidal distributions are considered for wind speed, yaw, and temperature with high probability density values at high wind speeds, moderate yaw angles, and low temperatures respectively. The effect of α and a_φ are not guessed beforehand, and hence a uniform distribution is considered for α and a_φ. A total of 100 initial samples are generated, of which 68 are generated using the LHS method, and 32 samples are generated by considering the combinations of the extreme values of the inflow variables (vertices of the hypercube formed by the inflow space).

4.2. Surrogate model

After the generation of samples, `HAWC2` simulations are performed, and for each simulation, the damping ratio is identified using the method described in Section 3. Then, a surrogate model is trained on the dataset comprising the initial samples and the corresponding damping ratios. The surrogate models establish an input–output relationship between the inflow variables and the damping ratio. This way, the damping ratio can now be expressed as a continuous variable across the domain, which helps in identifying a minimum value via standard optimization techniques. After investigating the performance of many surrogate models, Gaussian Process Regression (GPR) models (Kriging) which have been widely used in SBO problems are considered here. GPR models are interpolation models, and hence take the exact function value at the known sample locations, which makes them an attractive choice to be used in such applications where each evaluation of the target function is the result of a computationally expensive simulation. Also, in comparison to models which require a high training time like ANN, the GPR models take significantly lesser time to train, and this also makes them an attractive choice as the surrogates need to be trained multiple times in the proposed SBO framework. This part about training the surrogate multiple times is explained in detail in the forthcoming Section 4.3.1 which describes the exploration part of the adaptive sampling strategy. In this work, a GPR model with a squared exponential kernel is used which is briefly described below. The GPR model is implemented using the library `Scikit Learn` [32]. A brief description of the GPR model is as follows

Let $\mathbf{x}$ denote the independent variable in N dimensions (in this case the five inflow space variables) and y the dependent variable (in this case the damping ratio). Let $X = \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \ldots \mathbf{x}_n\}$ denote the set of sample points and let $y = \{y_1, y_2, \ldots y_n\}$ denote the value of the target function at the sample points. Let S denote the surrogate model that is obtained using X and y as training data. The GPR model for S uses a weighted sum of n functions with each function centred about the corresponding sample point. GPR models assume that the underlying target function is the result of a Gaussian Process (GP). The use of GP to model the underlying target function has its roots in geo-statistics where it was originally used to estimate surface heights of a given terrain based on heights sampled at a few locations.

$$S(\mathbf{x}) = \sum_{i=1}^{n} \lambda_i \psi(d(\mathbf{x}, \mathbf{x}_i)), \tag{5}$$

where λ_i are the weights (the model parameters), ψ is the squared exponential kernel that is centred around each point, and $d(\mathbf{x}, \mathbf{x}_i)$ denotes the distance between the points $\mathbf{x}$ and $\mathbf{x}_i$. The kernel is given by

$$\psi(d(\mathbf{x}_i, \mathbf{x}_j)) = \exp\left(-\frac{1}{2}\frac{d(\mathbf{x}_i, \mathbf{x}_j)^2}{l^2} \right) \tag{6}$$

where $l \geq 0$ is the length scale parameter. If l is a scalar, the kernel is called an isotropic kernel, and if $l = \{l_1, l_2, l \ldots l_N\}$ is a vector with a dimension N i.e. each independent variable has its own length scale, then the kernel is called an anisotropic kernel. In that case, the distance between the points is calculated after scaling the points with the length scale in each dimension. In this work, an anisotropic kernel is used. The length scale parameters are obtained by Maximum Likelihood Estimation.

The input variables are normalized to the range [0,1] while the target variable is standardized by scaling to unit mean and zero standard deviation, and the GPR model is trained on the normalized data. The reason for normalizing the inflow variables to the range [0,1] is to facilitate easier estimation of length scales and avoid numerical ill-conditioning when calculating the covariance matrix which is dependent on the distance between sample points. The reason for standardizing the output variable is because GPR models assume a zero mean prior on the target data. And similarly, whenever a prediction is to be made using the trained model, the data points to be predicted are normalized to [0,1], and the resulting output is then inverse scaled to obtain the predicted damping ratio.

4.3. Adaptive sampling

Adaptive sampling is a strategy to generate subsequent samples in the SBO process. The SBO process has two phases: exploration and exploitation, and each phase has its own adaptive sampling strategy. Exploration is the process of improvement of understanding of SIV behaviour in the domain while Exploitation is the process of understanding critical areas in the domain. The exploration phase aims at improving the accuracy of the surrogate model globally, while the exploitation phase aims at reaching the global minimum damping in the domain. In both phases, the samples are chosen iteratively. The adaptive sampling strategy for both phases is explained below.

4.3.1. Exploration

In the exploration phase, the aim is to improve the global performance of the surrogate model. One way to do this is to add more samples in the empty regions of the domain and in regions where the target function or the surrogate S is highly nonlinear. A way to identify empty regions in the input space is by the Delaunay triangulation of the domain. Delaunay triangulation is a way to divide the domain into triangles (simplexes in higher dimensions) with the sample points as vertices such that no sample point lies inside the circumcircle of any triangle thus being formed. The Delaunay triangulation maximizes the minimum angle of all the triangles and hence avoids the formation of skewed triangles. The simplexes with high volumes indicate large empty regions in the domain. Once the large regions are identified, the problem of complexity is addressed by placing the sample at the 'most non-linear' point within the large simplex. The samples in the exploration phase are selected in iterations or rounds, with k samples per round. A description of the selection of k new samples for one round is explained in the remainder of this subsection.

The Delaunay triangulation of the N dimensional space formed by X leads to T Delaunay simplices $\{D_1, D_2, \dots D_T\}$ with volumes $\{V_1, V_2, \dots V_T\}$, circumcentres $\{C_1, C_2, \dots C_T\}$, and circumradii $\{R_1, R_2, \dots R_T\}$ respectively. The volume V_t, circumcentre, C_t and circumradius R_t of a simplex D_t with known vertices can be calculated as

$$V_t = \sqrt{\left(\frac{(-1)^{N-1}}{2^N (N!)^2} |\mathcal{M}_t|\right)} \tag{7}$$

$$\mathcal{M}_t = \begin{bmatrix} 0 & 1 & 1 & \cdots & 1 \\ 1 & 0 & d_{12}^2 & \cdots & d_{1(N+1)}^2 \\ 1 & d_{21}^2 & 0 & \cdots & d_{2(N+1)}^2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & d_{(N+1)1}^2 & d_{(N+1)2}^2 & \cdots & 0 \end{bmatrix} \tag{8}$$

where d_{ij} denotes the distance between vertices $i\&j$, and $\mathcal{M}_t$ is called the Cayley–Menger matrix. R_t in cartesian coordinates, and C_t in Barycentric co-ordinates $\{\gamma_1, \gamma_2, \gamma_3, \dots, \gamma_{N+1}\}$ can be obtained by solving the following linear system

$$\mathcal{M}_t \begin{bmatrix} -2R_t^2 \\ \gamma_1 \\ \gamma_2 \\ \vdots \\ \gamma_{N+1} \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{9}$$

For the theorems related to calculating the volumes, circumcentres and circumradii, an interested reader is referred to [17]. Eq. (9) gives the circumcentre in barycentric co-ordinates, and it should be converted to cartesian coordinates for use later. The cartesian co-ordinate of C_t in the ith dimension can be calculated from its corresponding barycentric coordinate as

$$C_t^{(i)} = \sum_{j=1}^{N+1} \gamma_j x_j^{(i)} \tag{10}$$

where $x_j^{(i)}$ denotes the value in the ith dimension of the point $\mathbf{x}_j$.

From the Delaunay triangulation, the simplexes with the first k high volumes are chosen, and the samples are to be placed inside these simplexes. Let D^* denote one such simplex. Since D^* is a simplex in N dimensions, it has $N+1$ vertices $\{\mathbf{x}_{D1}, \mathbf{x}_{D2}, \dots \mathbf{x}_{DN}\}$. The location of the sample point within D^* is determined using a measure of complexity. To do this, first a linear approximation $L(\mathbf{x})$ is constructed using the $N+1$ points formed using the vertices of the simplex as independent variable values and the value of the target function at the vertices as the dependent variable value. Then the deviation function $\delta(\mathbf{x}) = |S(\mathbf{x}) - L(\mathbf{x})|$ gives a measure of the nonlinearity. However, maximizing δ within D^* may lead to the problem of the new sample point being too close to the vertices of the simplex, and hence may not help in effectively addressing the problem of empty regions. To address this, a distance function is defined as $\phi(\mathbf{x}) = \prod_{j=1}^{N+1} \|\mathbf{x}_{Dj} - \mathbf{x}\|$ as the product of the new sample point from all the vertices of D^*. Then, the product $\delta(\mathbf{x})\ x\ \phi(\mathbf{x})$ is optimized to balance the complexity and distance to existing points. Finally, to make the optimization problem simpler, the optimization is done within the circumsphere of D^* rather than D^* itself, as it simplifies the definition of the bounds of the problem. While this does increase the search region beyond D^* in some regions, this expansion can be argued as being better as we search a larger unexplored space. The final optimization problem that is solved is

$$\begin{aligned} &\text{max. } \delta(\mathbf{x})\phi(\mathbf{x}) \\ &\text{subject to } \|\mathbf{x} - C^*\|^2 \le R^{*2} \end{aligned} \tag{11}$$

where C^*, and R^* denote the circumcentre and circumradius of D^* respectively. At the end of the optimization, if the optimization fails, the next sample is located at the centroid of D^* whose coordinate in a given dimension is the average of the coordinate value of the vertices. In each round, the metric used to measure the accuracy of the surrogate model is the R^2 value calculated using the Leave One Out (LOO) approach. For each sample point in the domain $(\mathbf{x}_i, y_i)$, a surrogate model $S_i(\mathbf{x})$ is defined as the surrogate model trained by excluding $(\mathbf{x}_i, y_i)$ from the training data of S. Then the LOO accuracy is defined as

$$R^2 = 1 - \frac{\sum_{i=1}^{n} (y_i - \hat{y}_i)^2}{\sum_{i=1}^{n} (y_i - \bar{y}_i)^2}, \tag{12}$$

where $\hat{y}_i$ refers to the corresponding prediction made by S_i and $\bar{y}_i$ refers to the average value of the target function at all samples. This exploration metric provides a termination criterion for the exploration phase. The exploration phase is done until the R^2 value in the last n_{conv} consecutive iterations reaches an arbitrary threshold ϵ. Here n_{conv} and ϵ are set to 3 and 0.8 respectively. The threshold on R^2 is set with the consideration that the considered problem is in five dimensions and the number of points is of the order of 10^2. A higher threshold would require a higher number of simulations and hence incur a higher computational cost. The exploration process is summarized in Fig. 3.

4.3.2. Exploitation

In the exploitation phase, the aim is to find critical areas in the domain, i.e., areas of very low damping. As stated in the description of the framework, the goal of the framework is to identify critical areas in the domain in addition to the global minimum. To achieve this, the optimization of the surrogate model is carried out using a multi-start minimize method, and we identify m minima less than a specified threshold in each round. In the multi-start minimize method, the optimization of the surrogate model is performed for multiple initial values, and the minima are found each time. Many choices for the minimization algorithm exist, and in this work, the minimization was performed using the Sequential Least SQuares Programming (SLSQP) algorithm [30]. The minimization using SLSQP algorithm is implemented using the Python library `SciPy` [33]. The identified m minima are chosen as samples for the next iteration. As we try to study a larger dimensional space with only a few points, even after achieving satisfying accuracy of the surrogates in the exploration phase, the

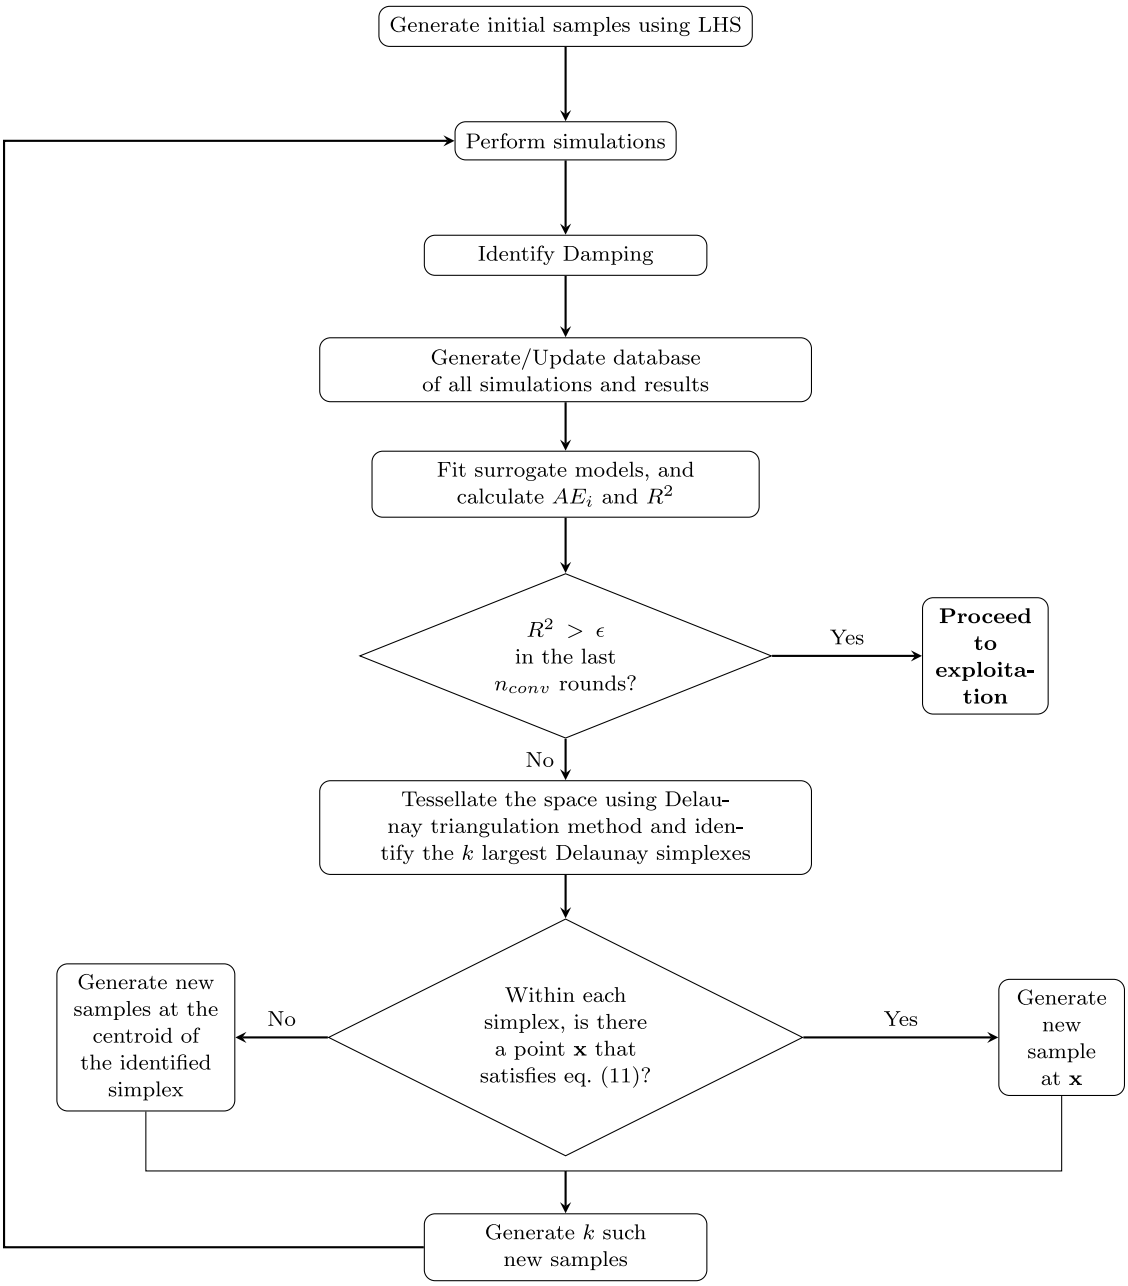

Fig. 3. The exploration process.

minima predicted by optimizing (minimizing) the surrogate model and the actual damping ratio from the HAWC2 simulation at the predicted minima may not be the same. So, the minimization process is repeated for a few iterations until the minima predicted by the surrogate model and the HAWC2 simulation are the same. Since we add samples at predicted local minima in addition to the global minima, the surrogate model is expected to be well trained around many points where the target function takes low values, and thus identify critical regions. If the global minima does not converge after $N_{exploit}$ rounds, to obtain an estimate of the global minima with the information available, the following scheme is employed.

4.3.3. Convergence scheme

Because the surrogate model is only an approximation to the true target function, the minimum value predicted by the surrogate model and the value of the target function at the predicted minima may not be the same, and thus the surrogate model predicts many 'false' minima. The predicted minima are false in the sense that the values of the target function at these points are not as low as predicted by the surrogate model. This problem of the actual and predicted minima not being the same has been observed in previous works with SBO [34]. In such cases, if an agreement is not achieved between the actual and the predicted minima after $N_{exploit}$ rounds, a new scheme is proposed to conclude convergence to a respectable minimum. The basic idea with the scheme is that, in the neighbourhood of the predicted minima, the observed target function values should be low, which is a justified assumption for a continuous target function. For every predicted minima, the corresponding Delaunay simplex with the closest centroid is identified, and the value of the target function is evaluated at the vertices of the simplex, and the average is calculated ($f_{D,\text{avg}}$). A threshold is then set on $f_{D,\text{avg}}$. The predicted minima whose $f_{D,\text{avg}}$ is higher than the threshold are regarded as possibly false, as they are surrounded by points with a high average value of the target function. Such predicted minima are not considered as samples for

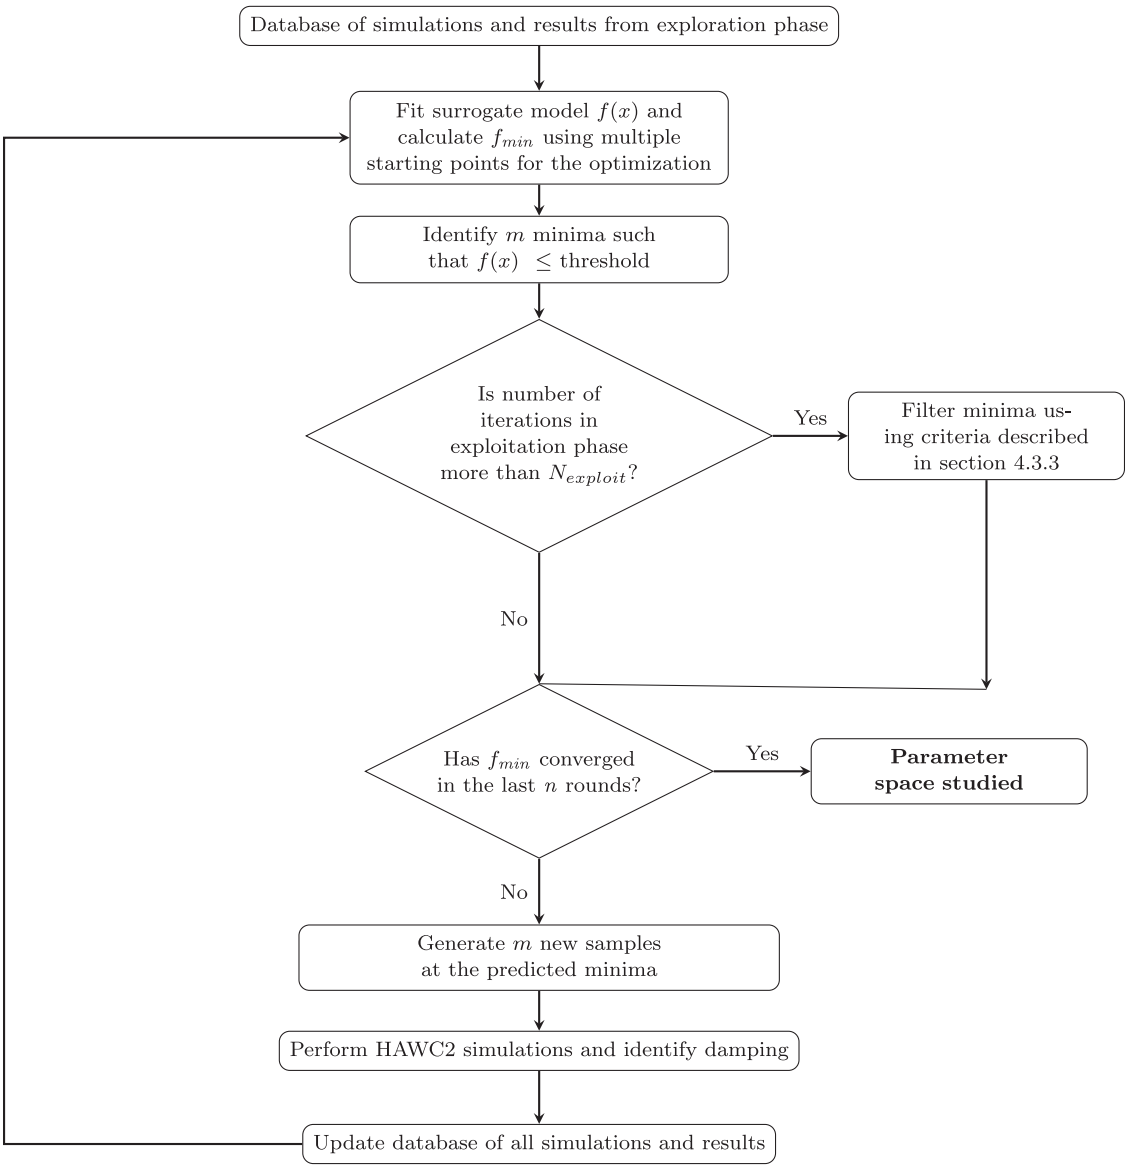

Fig. 4. The exploitation process.

the next round. The exploitation phase is repeated until the minima converges. In this work, $N_{exploit}$ is chosen as 20. The exploitation phase is summarized in Fig. 4.

The method of selecting samples using Delaunay triangulation and a measure of complexity is very similar to the algorithm proposed by [17], and the method of iteratively minimizing the surrogate model in the exploitation phase is similar to the method used in [35]. The key differences between these (and other similar works) and this work are:

1. In the exploration phase, we use a R^2 metric after every iteration as a measure of the quality of the surrogate and to decide when to stop the exploration phase, while this metric has not been proposed in the works that have used Delaunay triangulation for adaptive sampling. This metric serves two purposes - (a) Track the progress of the exploration phase and (b) To decide the termination of the exploration phase.
2. In the exploitation phase, if the minima cannot be found after a specified number of iterations, we propose a scheme based on the properties of the Delaunay simplexes to eliminate false predicted minima and conclude the process to a respectable minimum using the information available (Section 4.3.3). The previous studies that have used the method of iterative minimization of the surrogate model to identify the minimum have not used such a method to eliminate false minima in case convergence is not reached.

4.4. Framework validation

The proposed framework is validated against the six-hump camel back function and the Rosenbrock function. The camel back function is a 2 dimensional function given by

$$f(\mathbf{x}) = \left(4 - 2.1x_1^2 + \frac{x_1^4}{3}\right)x_1^2 + x_1x_2 + (-4 + 4x_2^2)x_2^2 \tag{13}$$

The function is usually defined in the domain $x_1 \in [-3,3]$, $x_2 \in [-2,2]$. The camel back function has six local minima and two global minima. It is visualized in Fig. 5(a). The global minima of $y = -1.0316$ occurs at $\mathbf{x}^* = (0.0898, -0.7126)$ and $(-0.0898, 0.7126)$.

The Rosenbrock function is a multidimensional function given by

$$f(\mathbf{x}) = \sum_{i=1}^{d-1}[100(x_{i+1} - x_i^2)^2 + (x_i - 1)^2] \tag{14}$$

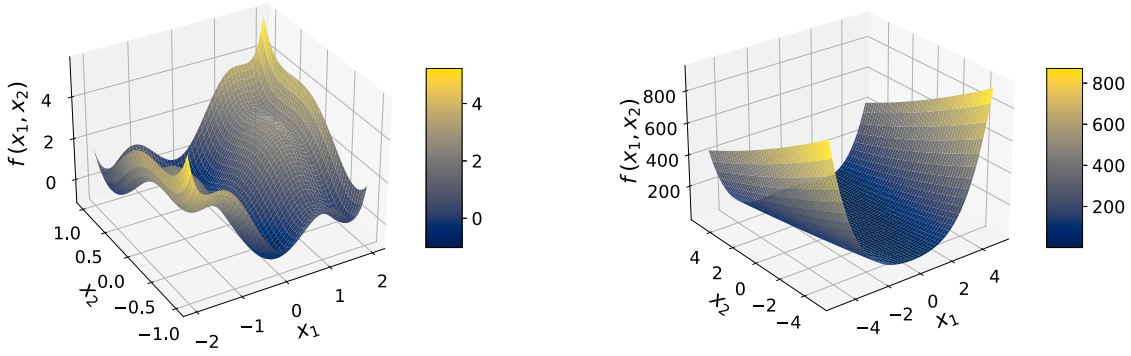

(a) *Six hump camel back function.* (b) *Rosenbrock function in two dimensions.*

Fig. 5. Test functions.

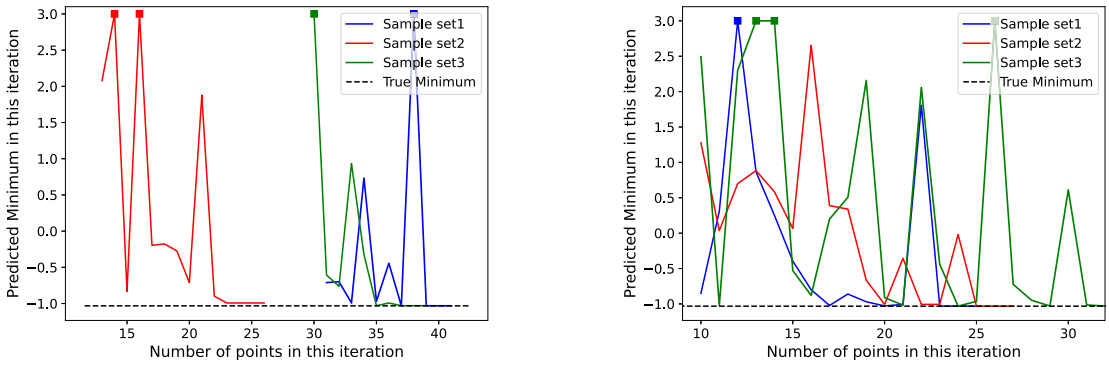

(a) *Exploitation phase of the proposed framework. Previous samples are a part of the exploration phase with 10 initial samples.*

(b) *EGO algorithm based on the EI adaptive sampling strategy and 10 initial samples.*

Fig. 6. Identification of minima of the six-hump camel back function with three different sets of initial samples. To make the plot clearer, values greater than 3 are saturated to 3 and marked in square in the plots.

The function is usually defined in the domain $\mathbf{x} \in [-5,10]$, and has a global minima of $y = 0$ at $\mathbf{x}^* = 1$. The minima lies in a narrow, parabolic, valley around $\mathbf{x}$=1. The Rosenbrock function in two dimensions is shown in Fig. 5(b).

The proposed framework is tested on the six-hump camel back function, and the 5-dimensional Rosenbrock function to see if the minima of these functions can be identified by the framework. To check for robustness against initial sampling, the process is done for three different sets of initial samples. The test functions are also evaluated on the popular SBO algorithm Efficient Global Optimization (EGO), with an adaptive sampling strategy based on the Expected Improvement (EI) function [11]. The results of the identification of the minima with both the algorithms are shown in Figs. 6 and 7 respectively. The EGO algorithm was implemented using the library SMT [36].

It can be seen that, for the 2-dimensional six-hump camel back function, the proposed framework requires more samples than the EGO algorithm to identify the minima, whereas for the 5-dimensional Rosenbrock function, for all the three different sets of initial sample sets considered, the EGO algorithm does not identify the true minima or achieve convergence, while the proposed framework converges consistently to the true minima of the function. The considered SIV problem involves 5 dimensions, and the performance of the framework in solving a similar 5 dimensional Rosenbrock problem is an indication of the suitability of the framework for the considered SIV problem.

4.5. Framework performance in analysing the SIV problem

The evolution of the surrogate accuracy and the evolution of minimum damping ratio is shown in Fig. 8. As expected, the accuracy of the surrogate model as measured by the R^2 metric calculated using a LOO approach increases with increase in number of points in the exploration phase. Once the required level of accuracy is reached consistently (R^2 =0.8 over three consecutive iterations), the exploitation phase is started to identify critical areas in the domain. Once the exploitation phase starts, the R^2 metric remains almost constant as the points added in the exploitation phase maybe closer to the existing points, and hence do not significantly affect the model performance.

In the exploitation phase, focus is on identifying critical conditions, and after a few iterations, the most critical condition is identified. The most critical SIV as predicted by the proposed framework is around −6%. The corresponding inflow conditions are around Wind speed =60 ms^{-1}, yaw angle = 18 deg, $a_\varphi = -1.2$, $\alpha = 0.4$, Temperature = 20 °C. It can be seen that during the convergence, the minimum does not keep decreasing sequentially with iterations. The reason for this non-monotonic trend is that, as mentioned in Section 4.3.3, the minimum damping ratio predicted by the surrogate model and the actual damping ratio from HAWC2 simulations do not always agree.

The convergence to the minima for the individual variables is shown in Fig. 9. Since there are five dimensions involved, the convergence is shown individually for each variable. The timeseries of the blade edgewise bending moment for some of the most critical conditions identified by the framework is shown in Fig. 10, and some of the inflow

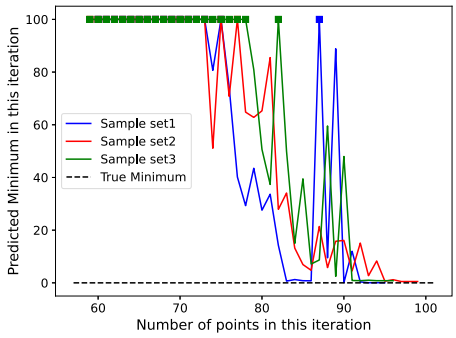

(a) *Exploitation phase of the proposed framework. Previous samples are a part of the exploration phase with 50 initial samples.*

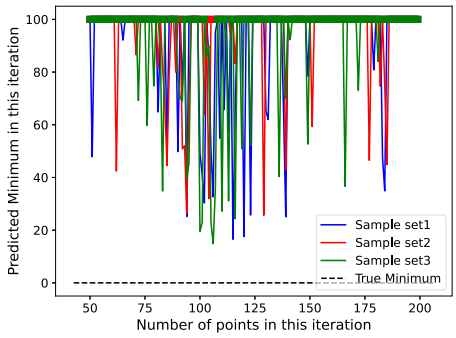

(b) *EGO algorithm based on the EI adaptive sampling strategy and 50 initial samples.*

Fig. 7. Identification of minima of the Rosenbrock function with three different sets of initial samples. To make the plot clearer, values greater than 100 are saturated to 100 and marked in square in the plots.

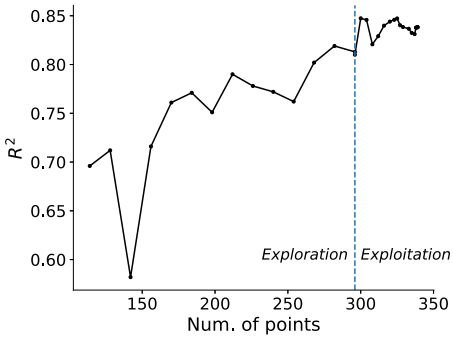

(a) *Evolution of Accuracy of the surrogate model.*

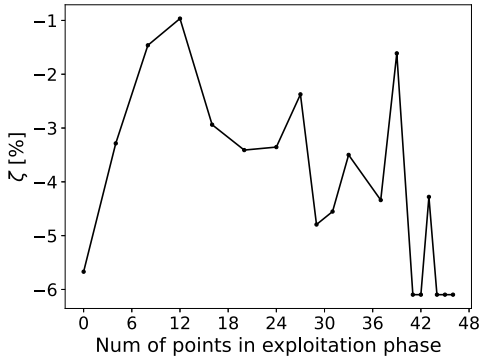

(b) *Evolution of Predicted minimum Damping in the exploitation phase.*

Fig. 8. Performance of framework.

Table 2
Some inflow conditions that lead to severe SIV.

Wind speed [ms^{-1}]	Yaw angle [deg]	a_φ [–]	α [–]	Temperature [°C]	ζ [–]
60	18	−1.12	0.4	20	−6.1%
60	18.6	−1	0.35	20	−5.7%
60	17.5	−0.81	0.12	3.01	−5.1%
60	18.5	−0.65	0.16	3.05	−4.9%
60	25.5	0.15	0.24	6.77	−4.8%

conditions that lead to severe SIV is shown in Table 2. It can be seen that the identified scenarios are spread over the domain, and that using the framework we can identify multiple critical areas in the domain as opposed to the classical SBO approach where the goal is to identify only the most critical point. The accuracy of the final trained surrogate model for different test data points is shown in the appendix.

For comparison, the EGO algorithm with the EI criterion is also used to get the minimum damping, as shown in Fig. 11.

The EGO algorithm has been run only for 20 iterations from the initial 100 samples but it is to be noted that the EGO algorithm identifies one new sample per iteration, and hence takes a much longer time compared to the framework proposed in this work which identifies multiple samples per iteration. hence, while the proposed framework and the EGO algorithm have been run for nearly the same amount of time, the EGO algorithm uses much fewer samples. Also, the focus of the EGO algorithm is to identify the global minimum, while the focus of the framework presented in this work is to also obtain a surrogate model that identifies multiple critical regions in the domain. Thus the feature that the proposed framework uses many samples is of great advantage since the surrogate model is trained better.

5. Influence of variables on SIV

The influence of the input variables on SIV is studied using sensitivity analysis. A global sensitivity analysis can help to understand the overall influence of the input variables, while a regional sensitivity analysis can help to identify the combination of variables that lead to certain types of damping behaviour. This can then be used to identify inflow conditions that lead to severe SIV.

Sobol Indices, which measure the fraction of variance of each input variable (and also combinations) in relation to the total variance of the output variable can be used to measure global sensitivity [37]. For an n-dimensional input space with $i, j \leq n$, the first order Sobol index S_i indicates the effect of individual variable i, while the second order index S_{ij} indicates the effect of the interaction of variables i and j, and so on for the higher indices, with a higher value indicating a higher influence. For an output variable $Y = f(X_i)$, the first order Sobol indices, S_i are given by

$$S_i = \frac{\mathrm{Var}[E(Y \mid X_i)]}{\mathrm{Var}(Y)} \tag{15}$$

The second order indices are given by

$$S_{ij} = \frac{\mathrm{Var}[E(Y \mid X_i, X_j)]}{\mathrm{Var}(Y)} - S_i - S_j \tag{16}$$

and so on.

Sobol indices are always in the range [0,1]. The calculation of Sobol Indices involves a large number of evaluations of the output with

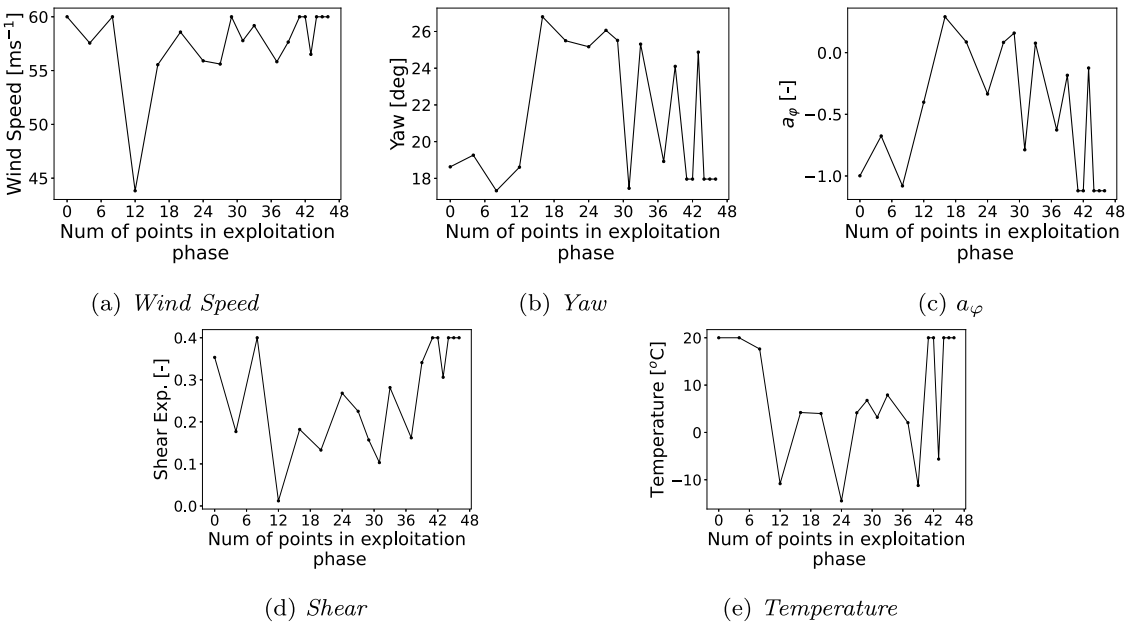

(a) *Wind Speed* (b) *Yaw* (c) a_φ

(d) *Shear* (e) *Temperature*

Fig. 9. Convergence to minima.

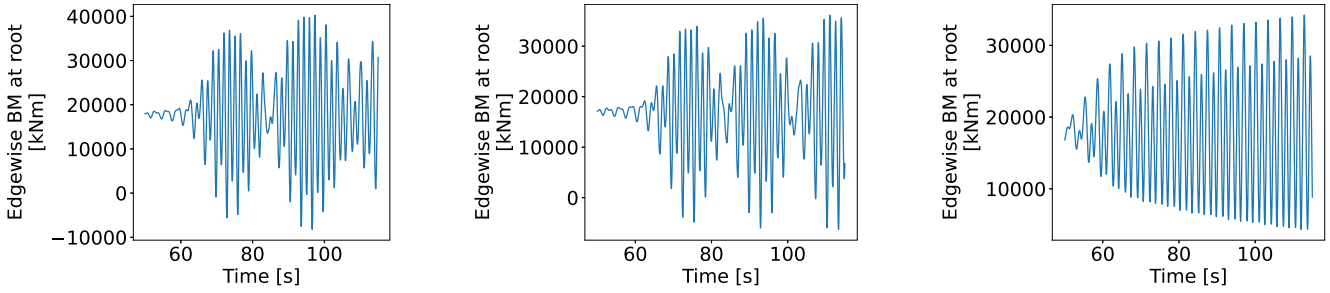

(a) $\zeta = -6.1\%$. (b) $\zeta = -5.7\%$. (c) $\zeta = -5\%$.

Fig. 10. Time series of the edgewise bending moment at the blade root for some of the severe conditions identified by the framework. The quick rise in the blade edgewise bending moments can be seen clearly.

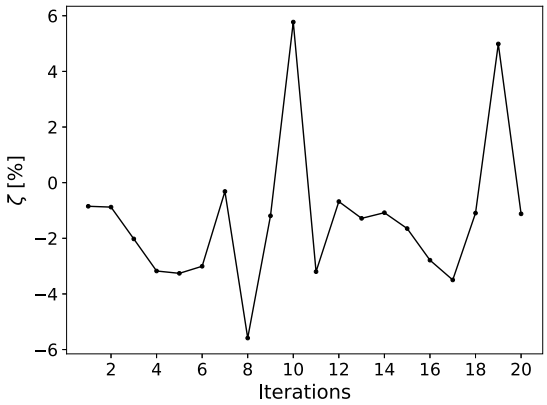

Fig. 11. Evolution of minimum damping using the EGO algorithm.

respect to the inputs, and the surrogate model is useful here. Multiple function evaluations are done quickly using the surrogate model, and the indices are calculated using the python library `SALib` [38]. The first order Sobol indices of the different variables are shown in Table 3.

It can be seen that the yaw angle is the most influential variable on SIV. Wind speed and veer have a mild influence, while shear and temperature have the least influence. The reason for the very low influence of temperature could be attributed to the fact that the effect of temperature is modelled through a change in structural damping, and for the turbine considered, the structural damping of the first edgewise mode itself is low (0.36%) as compared to the negative aerodynamic damping, which can be as high as −6%. Therefore, changes in temperature so not affect the overall damping significantly.

Table 3
First order Sobol indices.

Wind Speed	Yaw	Veer	Shear	Temperature
0.02	0.425	0.026	1.68E−03	4.4E−04

The quantity $1 - \sum S_i$ indicates the effect of interaction between the variables. Here, $\sum S_i = 0.473$, which implies a still significant interaction between the variables. To get the influence of the second order interaction among the variables, the second order Sobol indices are calculated. They are shown in Table 4.

It can be seen that there are significant interactions between wind speed and yaw, and yaw and veer. The interactions with respect to temperature and shear are very low, as for the first order Sobol index. Thus, it can be said that for this setup, yaw angle, wind speed and veer are the important variables in understanding the behaviour of SIV. This can also be seen from the values of the total Sobol indices shown in

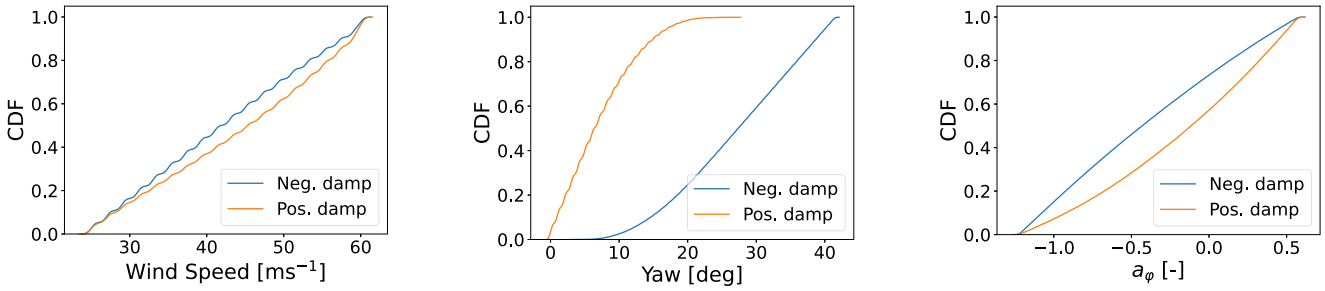

(a) *Wind Speed, D=0.0947* (b) *Yaw, D=0.7868* (c) a_φ, *D=0.1648*

Fig. 12. CDFs and KS test statistics of the different variables.

Table 4
Second order Sobol indices.

	Wsp	Yaw	Veer	Shear	Temperature
Wsp		0.146	0.006	0.02	0.001
Yaw	0.146		0.138	0.009	0.002
Veer	0.006	0.138		0	0
Shear	0.02	0.009	0.003		0
Temperature	0.001	0.002	0 .001	0	

Table 5
Total Sobol indices.

Wind Speed	Yaw	Veer	Shear	Temperature
0.38	0.924	0.335	0.11	0.007

Table 5 which include the combination of first, second and all the other higher order indices

5.1. Occurrence

Using the surrogate model, the inflow conditions that lead to SIV are also studied using techniques of regional sensitivity analysis [37]. The inflow space is divided into two regions based on the damping predicted using the surrogate model. In region B, the damping ratio is negative ($\zeta \leq 0$), and in the other region $\bar{B}$, the damping ratio is positive. The distribution of the variables in B and $\bar{B}$ are then analysed, and compared using a standard test such as the 2-sided Kolmogorov–Smirnov test (KS test) to analyse the similarity of the distributions. The test statistic for each test is $D_i = \sup\|F(X_i|B) - F(X_i|\bar{B})\|$ where F represents the Cumulative Density Function (CDF). The CDFs of X_i and the values of D_i are shown in Fig. 12.

For each X_i, the steeper regions of the CDF of B represent the values that lead to the output belonging to B. Also, a larger value of D_i indicates that the CDFs are significantly different, and hence if X_i is influential in deciding if the output lies in the region B. From Fig. 12, it can be seen that yaw angles with a large D are the most influential in the occurrence of a negative damping. The yaw angle range around [10, 40] deg is the most conducive for a negative damping to happen. Wind speed and a_φ have a mild influence in deciding the occurrence of SIV as the respective D values are low, and the distributions are still separated. As expected from the Sobol index values, the D values of shear and temperature are close to zero (not shown in Fig. 12), indicating that no specific values of these variables are favourable for SIV to occur.

5.2. Severity

The conditions that lead to severe SIV are analysed in a similar way, but now we define B as the region where ζ is severely negative. In this work, a scenario is severely negative, if $\zeta \leq -3\%$. The analysis of how the inflow space variables are distributed in B can help in understanding the conditions that lead to severe SIV. Fig. 13 shows the distribution of the inflow space variables in B.

It can be seen that the most severe SIV are predicted to occur at high wind speeds, and yaw angle in the range [10, 30] deg. It can also be seen that the distributions of the other variables span all over their defined ranges. So, it is predicted that severe negative damping can occur at any possible value in their range, although positive shear and negative veer are more probable to lead to severe SIV.

An analysis of the joint occurrence of the variables can provide additional information about the interaction between the variables in B. Since wind speed and yaw angle seem to be the most restrictive in terms of defining the region B, the combination with one of these variables is expected to provide the most information. The joint occurrence of the variables with yaw angle is shown in Fig. 14. It can be seen that the joint densities with wind speed and a_φ are concentrated in specific regions and also indicate some trends. For example, it can be seen that the yaw angle at which severe SIV occurs increases with increasing veer in the positive direction. Also, the probability of observing severe SIV is higher at negative veer than positive veer, indicating that situations with negative veer are more conducive than situations with positive veer for severe SIV to happen. Similarly, severe SIV are more probable at higher wind speeds. The joint densities with shear and temperature are relatively spread all over, indicating that the variables do not have a significant impact on the severity. In addition to identifying regions leading to severe damping ratio, the surrogate model can be used to get estimates of the damping ratio for severe conditions, as shown in [6].

The effectiveness of the surrogate based framework can be seen in the number of simulations necessary to study SIV in this domain with and without the surrogate based approach. Without the surrogate based framework, to generate the information presented in this section, simulations need to be conducted at at least 10 levels each for wind speed, yaw angle and temperature, and 5 levels for shear and veer each. Assuming a full factorial design to study SIV in the domain, this leads to a total number of $10^3 \times 5^2 = 25{,}000$ simulations. With the surrogate based framework, however, the behaviour of SIV in the domain is analysed with just a few hundred simulations. The trade-off is the uncertainty introduced by the surrogate model but as it can be seen, the computational benefit is huge.

It is to be noted that the severity of SIV as described in this work is measured using the damping ratio at low amplitudes of the edgewise bending moment. However, the simulations are performed using a Near Wake Model, which has been shown to have some deviations from actual measurements. Hence a more accurate picture of the instabilities is expected using a high cost, high-fidelity CFD based aeroelastic solver. But nonetheless, the methods presented here can help as a guiding tool to narrow the inflow conditions to be simulated using the CFD solvers, thereby optimizing the computational cost required to explore the inflow space.

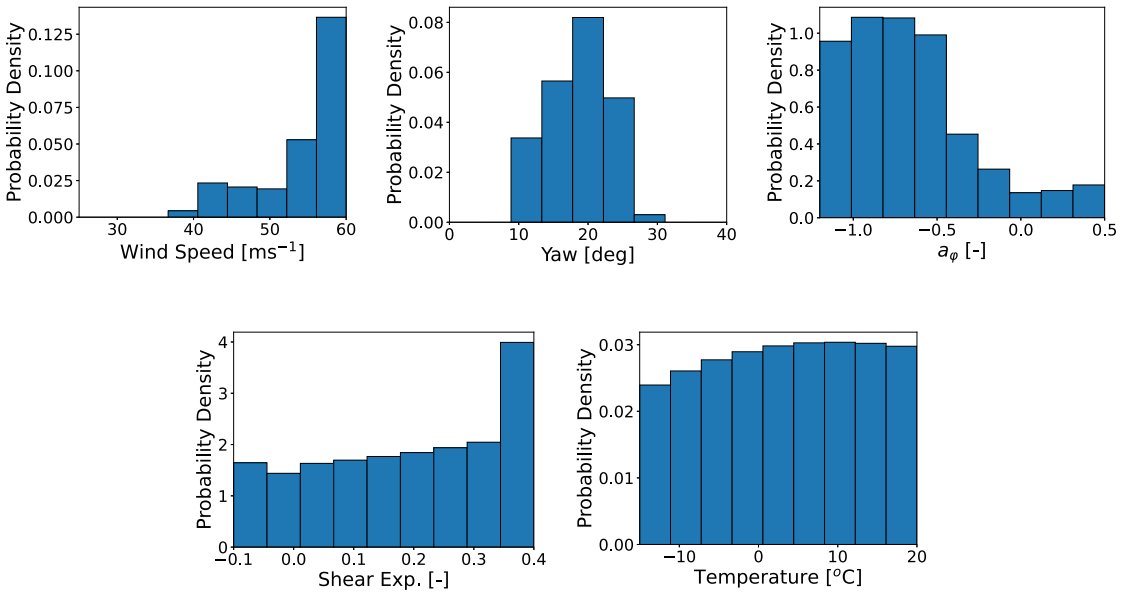

Fig. 13. Distribution of variables in regions that lead to severe negative damping.

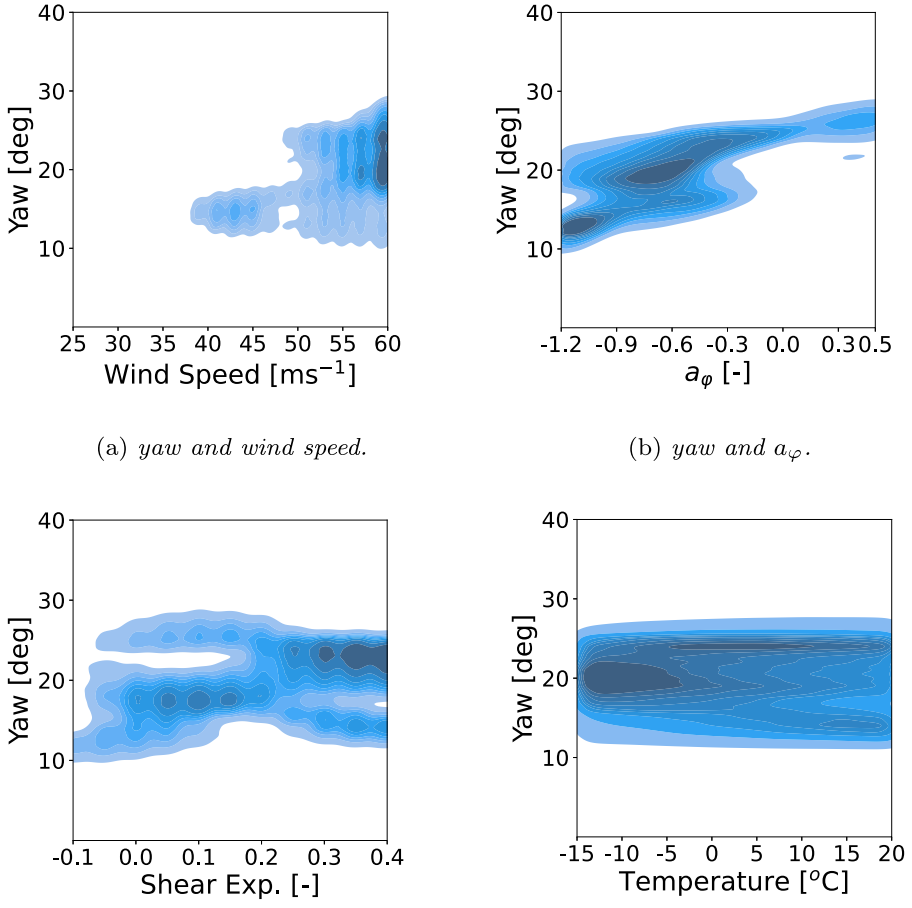

(a) *yaw and wind speed.* (b) *yaw and* a_φ.

(c) *yaw and shear.* (d) *yaw and temperature.*

Fig. 14. Joint densities of the variables in the severely negatively damped region.

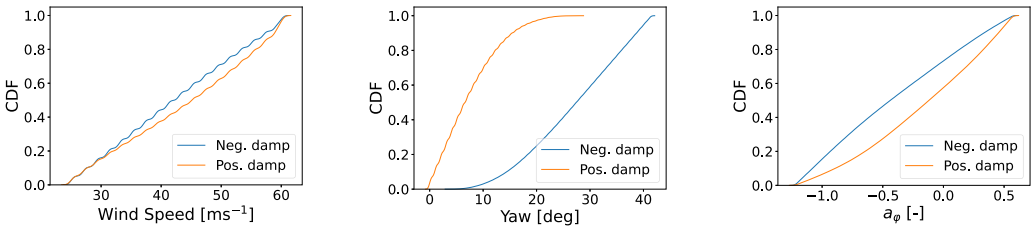

(a) *Wind Speed, D=0.0967* (b) *Yaw, D=0.7679* (c) a_φ, *D=0.1843*

Fig. A.15. CDFs and KS test statistics of the different variables for a different initial sample set.

6. Conclusion and future works

In this work, the effect of multiple inflow variables on SIV is studied using a surrogate-based optimization framework since a direct study at all the levels of the inflow variables would be computationally very expensive. The damping ratio of the first blade edgewise mode is considered as the stability measure and is identified by a method that is a variant of the logarithmic decrement method. The SBO process is split into two phases - the exploration phase where the samples for simulation are chosen iteratively to increase the accuracy of the surrogate model, and the exploitation phase where the samples are chosen iteratively to find the conditions leading to critical SIV. Using the developed SBO framework, it is found that the yaw angle has the highest influence in deciding the occurrence and the severity of SIV, while wind speed and wind veer have a mild influence. It is also found that the conditions leading to critical SIV in the IEA 10 MW turbine are in the moderate yaw angle region with negative veer and high wind speeds. The region of instability is large, spanning over a wide range of yaw angles in the presence of negative veer, and is small, spanning over a narrow yaw angle range in the presence of a positive veer. It is proposed that such a study can be useful in getting a quick estimate of critical conditions and the effect of turbine design changes.

Since the proposed framework is not limited to any particular type of surrogate model, the use of surrogate-specific features has not been explored to get an estimate of the uncertainty of the predictions. The development of methods to estimate uncertainties without too much exploitation of surrogate-specific features would be an improvement to the proposed framework. As future works, the following ideas are proposed.

1. The effect of other environmental and design variables related to the structural and aerodynamic characteristics of the blade on SIV could be studied using the proposed framework. Further, while the current study includes the two different sets of initial samples sampled using the same probability distributions, the effect of different initial samples could be studied in detail by considering other sampling methods and samples of different sizes to examine the convergence properties.
2. This study considers negative temperatures, where the blade might be covered with ice, which would affect the aerodynamic and structural properties of the blade. While HAWC2 does not have an ice accretion model, the effect of ice accretion on the aeroelastic behaviour could be studied by modelling the aerodynamic effects of ice accretion using CFD simulations and then using the resulting airfoil characteristics in an aeroelastic simulation, such as the study done in [39]. This would lead to better modelling of the effect of temperature on SIV.
3. Other types of aeroelastic instabilities such as Vortex-Induced Vibrations (VIV) which involve even costlier CFD simulations could be studied using similar approaches like the SBO framework suggested in this work.

CRediT authorship contribution statement

Chandramouli Santhanam: Conceptualization, Data curation, Methodology, Software, Writing – original draft. **Riccardo Riva:** Supervision, Writing – review & editing. **Torben Knudsen:** Supervision, Writing – review & editing.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Chandramouli Santhanam reports financial support was provided by Innovation Fund Denmark. Chandramouli Santhanam reports a relationship with Innovation Fund Denmark that includes: funding grants.

Data availability

Simulation data and code is available upon request.

Acknowledgements

This work has been supported by the PRESTIGE project (J.no. 9090-00025B), granted by Innovation Fund Denmark.

Appendix A. SIV analysis for a different initial sample set

The framework is tested for a different initial sample set. In this new sample set, 68 new samples are obtained assuming the same probability distributions as explained in Section 4.1, and 32 samples are placed at the vertices of the five dimensional hypercube as before. The key results obtained in Sections 5.1 and 5.2 are shown for this new set of initial samples in Figs. A.15 and A.16.

Comparing Fig. A.15 with Fig. 12, it can be seen that the curves and the test statistic values are also very similar. Thus, the information derived with respect to the most important variables is very similar to Section 5.1

Comparing Fig. A.16 with Fig. 14, it can be seen that the predicted regions of severe damping are very similar, and the conclusions derived about the interaction of the variables are also the same.

This shows the robustness of the framework to initial samples in terms of deriving key conclusions about SIV in the five dimensional domain.

Appendix B. Accuracy of the surrogate model

The trained surrogate model is validated by comparing the actual and predicted damping ratios for a few points from the total set of 343 points (points from both exploration and exploitation phase). This is shown in Fig. B.17 where different test data points are considered from the total set of points, and in each case the model is trained on the remaining points as training data points. The train-test split used is 80-20%.

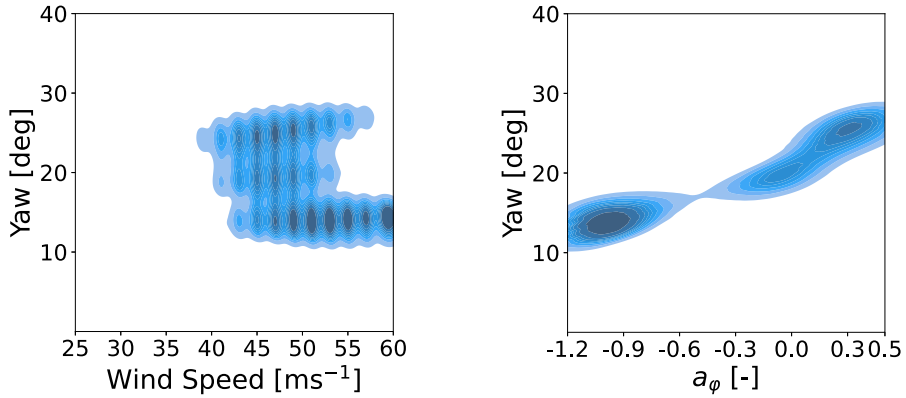

(a) *yaw and wind speed.* (b) *yaw and* a_φ.

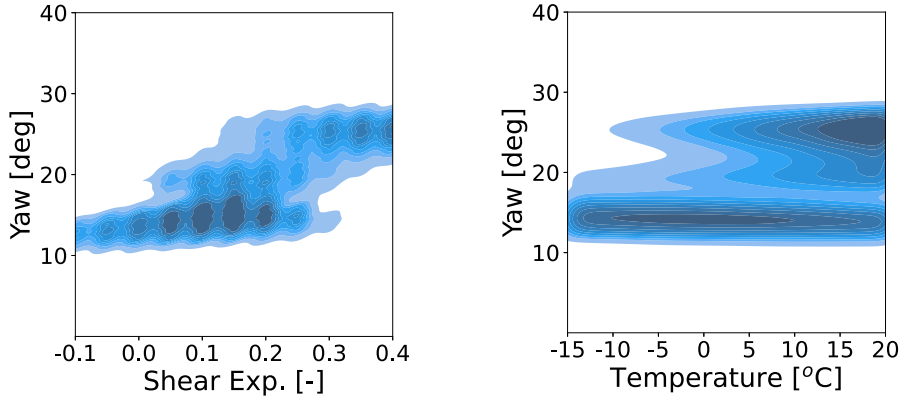

(c) *yaw and shear.* (d) *yaw and temperature.*

Fig. A.16. Joint densities of the variables in the severely negatively damped region for a different initial sample set.

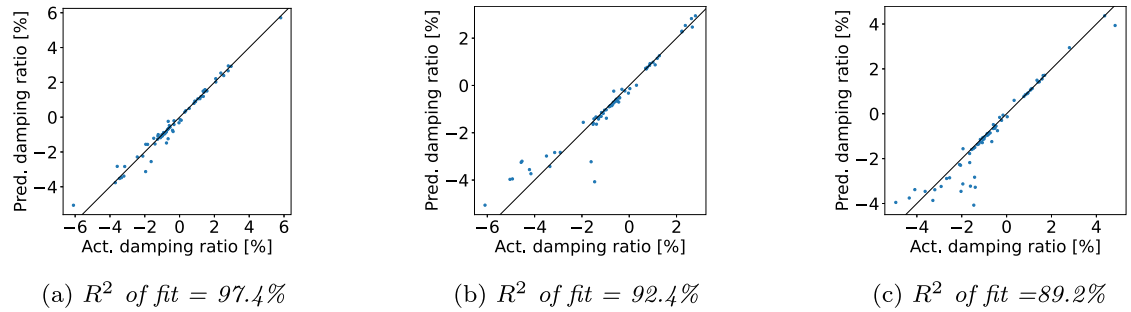

(a) R^2 *of fit* = *97.4%* (b) R^2 *of fit* = *92.4%* (c) R^2 *of fit* =*89.2%*

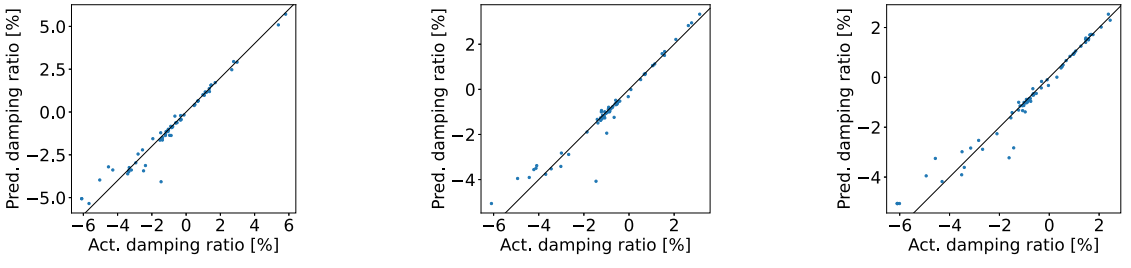

(d) R^2 *of fit* = *95.4%* (e) R^2 *of fit* =*94%* (f) R^2 *of fit* =*95.7%*

Fig. B.17. Comparison between actual damping ratio and damping ratio predicted by the final trained model for different training and test data points chosen as subsets of the total set of points.

It can be seen that the surrogate model has a slightly lower accuracy in the severe negative regions. This is because the length of the exploration phase is much longer than the length of the exploitation phase, and the number of training points in the severe SIV region is also less.

It can also be seen that variation of the model performance across different test datasets is not too much but still significant, which is a characteristic of surrogate models with limited datasets. For the five dimensional problem considered here, the total number of 343 points is still low to avoid this. This is also one of the main motivating reasons to use the LOO approach. In the LOO approach, this limitation is overcome as every point is included in the test database once, and the R^2 metric is calculated over all possible training and testing points, and a consistent measure of the model accuracy can be obtained easily.

References

[1] E.H. Dowell, A modern course in aeroelasticity, fourth ed., Springer, 2015, pp. 275–290.

[2] F. Zhao, Embedded Blade Row Flutter (Ph.D. thesis), Imperial College London, 2016.

[3] W. Skrzypiński, M. Gaunaa, Wind turbine blade vibration at standstill conditions-the effect of imposing lag on the aerodynamic response of an elastically mounted airfoil, Wind Energy 18 (3) (2014) 515–527, http://dx.doi.org/10.1002/we.1712.

[4] J. Heinz, N.N. Sørensen, V. Riziotis, M. Schwarz, S. Gomez-Iradi, M. Stettner, Aerodynamics of Large Rotors WP4 Deliverable 4.5, Tech. Rep. D4.5, ECN Wind Energy, Petten, The Netherlands, 2016.

[5] G.R. Pirrung, H.A. Madsen, S. Schreck, Trailed vorticity modeling for aeroelastic wind turbine simulations in standstill, Wind Energy Sci. 2 (2) (2017) 521–532.

[6] C. Santhanam, R. Riva, T. Knudsen, Surrogate models for predicting stall-induced vibrations on wind turbine blades, J. Phys. Conf. Ser. 2265 (2022) 032005.

[7] N. Dimitrov, M.C. Kelly, A. Vignaroli, J. Berg, From wind to loads: wind turbine site-specific load estimation with surrogate models trained on high-fidelity load databases, Wind Energy Sci. 3 (2) (2018) 767–790, http://dx.doi.org/10.5194/wes-3-767-2018.

[8] M. Thapa, S. Missoum, Surrogate-based stochastic optimization of horizontal-axis wind turbine composite blades, Struct. Multidiscip. Optim. 65 (2) (2022) 1–18.

[9] T. Barlas, N. Ramos-García, G.R. Pirrung, S. González Horcas, Surrogate-based aeroelastic design optimization of tip extensions on a modern 10 MW wind turbine, Wind Energy Sci. 6 (2) (2021) 491–504, http://dx.doi.org/10.5194/wes-6-491-2021.

[10] H.J. Kushner, A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise, J. Basic Eng. 86 (1964) 97–106, http://dx.doi.org/10.1115/1.3653121.

[11] D.R. Jones, M. Schonlau, W.J. Welch, Efficient global optimization of expensive black-box functions, J. Global Optim. 13 (4) (1998) 455–492, http://dx.doi.org/10.1023/A:1008306431147.

[12] Z.-H. Han, SurroOpt: A generic surrogate-based optimization code for aerodynamic and multidisciplinary design, in: 30th Proceedings of ICAS 2016, in: 2016–0281, International Council of the Aeronautical Sciences, Daejeon, South Korea, 2016.

[13] S. Bagheri, W. Konen, T. Bäck, Comparing kriging and radial basis function surrogates, in: Proc. 27. Workshop Computational Intelligence, Universitätsverlag Karlsruhe, Dortmund, 2017, pp. 243–259.

[14] N. Metta, R. Ramachandran, M. Ierapetritou, A novel adaptive sampling based methodology for feasible region identification of compute intensive models using artificial neural network, AIChE J. 67 (2) (2021) e17095, http://dx.doi.org/10.1002/aic.17095.

[15] S.S. Garud, N. Mariappan, I.A. Karimi, Surrogate-based black-box optimisation via domain exploration and smart placement, Comput. Chem. Eng. 108 (2019) 276–288, http://dx.doi.org/10.1016/j.compchemeng.2019.106567.

[16] A. Bhosekar, M. Ierapetritou, Advances in surrogate based modeling, feasibility analysis, and optimization: A review, Comput. Chem. Eng. 108 (2017) 250–267, http://dx.doi.org/10.1016/j.compchemeng.2017.09.017.

[17] S.S. Garud, N. Mariappan, I.A. Karimi, Surrogate-based black-box optimisation via domain exploration and smart placement, Comput. Chem. Eng. (2019) http://dx.doi.org/10.1016/j.compchemeng.2019.106567.

[18] M.H. Hansen, Aeroelastic instability problems for wind turbines, Wind Energy 10 (6) (2007) 551–577, http://dx.doi.org/10.1002/we.242.

[19] M. Stettner, M. Reijerkerk, A. Lünenschloß, V. Riziotis, A. Croce, L. Sartori, R. Riva, J. Peeringa, Stall-induced vibrations of the AVATAR rotor blade, J. Phys. Conf. Ser. 753 (2016) 042019, http://dx.doi.org/10.1088/1742-6596/753/4/042019.

[20] P. Bortolotti, H. Canet Tarrés, K. Dykes, K. Merz, L. Sethuraman, D. Verelst, F. Zahle, IEA Wind Task 37 on Systems Engineering in Wind Energy WP2.1 Reference Wind Turbines, Tech. Rep., International Energy Agency, 2019.

[21] T.J. Larsen, et al., How 2 HAWC2, The User's Manual, Tech. Rep., Risø National Laboratory, Denmark, 2007.

[22] M. Hansen, M. Gaunaa, H. Aagaard Madsen, A Beddoes-Leishman Type Dynamic Stall Model in State-Space and Indicial Formulations, Tech. Rep. 1354(EN), in: Denmark. Forskningscenter Risoe. Risoe-R, Risø National Laboratory, Denmark, 2004.

[23] A. Natarajan, N.K. Dimitrov, P.H. Madsen, J. Berg, M.C. Kelly, G.C. Larsen, J. Mann, D.R. Verelst, J. Dalsgaard Sørensen, H. Toft, I. Abdallah, N.-J. Tarp-Johansen, T. Krogh, J. Stærdahl, C. Eriksson, E. Jørgensen, F. Klintø, L. & Thesbjerg, Demonstration of a Basis for Tall Wind Turbine Design, EUDP Project Final Report, Chapter 2, Tech. Rep., DTU Wind Energy, 2016.

[24] G.T. Spirnak, J.R. Vinson, The effect of temperature on the material damping of graphite/epoxy composites in a simulated space environment, J. Eng. Mater. Technol. 112 (3) (1990) 277–279, http://dx.doi.org/10.1115/1.2903323.

[25] Y. Sefrani, J.-M. Berthelot, Temperature effect on the damping properties of unidirectional glass fibre composites, Composites B 37 (4–5) (2006) 346–355, http://dx.doi.org/10.1016/j.compositesb.2005.10.001.

[26] H. Zhen, W. Xiang, X. Xiao-lin, L. Zhi-Peng, W. Li-Ying, Study on glass fiber/epoxy gradient damping composites, Asian J. Chem. 25 (7) (2013) 3831.

[27] M.H. Hansen, Anisotrpic Damping of Timoshenko Beam Elements, Tech. Rep., Risø National Roskilde, Denmark, 2001.

[28] M. Hansen, Aeroelastic properties of backward swept blades, in: 49th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition, 2011, p. 260.

[29] J.S. Gray, J.T. Hwang, J.R.R.A. Martins, K.T. Moore, B.A. Naylor, OpenMDAO: An open-source framework for multidisciplinary design, analysis, and optimization, Struct. Multidiscip. Optim. 59 (4) (2019) 1075–1104, http://dx.doi.org/10.1007/s00158-019-02211-z.

[30] D. Kraft, A software package for sequential quadratic programming, Forsch.ber. Deutsche Forschungs Versuchsanstalt Luft Raumfahrt (1988).

[31] M. Baudin, A. Dutfoy, B. Iooss, A.-L. Popelin, OpenTURNS: An industrial software for uncertainty quantification in simulation, in: Handbook of Uncertainty Quantification, Springer International Publishing, Cham, 2017, pp. 2001–2038, http://dx.doi.org/10.1007/978-3-319-12385-1_64.

[32] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine learning in Python, J. Mach. Learn. Res. 12 (2011) 2825–2830.

[33] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S.J. van der Walt, M. Brett, J. Wilson, K.J. Millman, N. Mayorov, A.R.J. Nelson, E. Jones, R. Kern, E. Larson, C.J. Carey, İ. Polat, Y. Feng, E.W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E.A. Quintero, C.R. Harris, A.M. Archibald, A.H. Ribeiro, F. Pedregosa, P. van Mulbregt, SciPy 1.0 Contributors, SciPy 1.0: Fundamental algorithms for scientific computing in Python, Nature Methods 17 (2020) 261–272, http://dx.doi.org/10.1038/s41592-019-0686-2.

[34] M. Sessarego, N. Ramos-García, H. Yang, W.Z. Shen, Aerodynamic wind-turbine rotor design using surrogate modeling and three-dimensional viscous–inviscid interaction technique, Renew. Energy 93 (2016) 620–635.

[35] E. Davis, M. Ierapetritou, A centroid-based sampling strategy for kriging global modeling and optimization, AIChE J. 56 (1) (2010) 220–240.

[36] M.A. Bouhlel, J.T. Hwang, N. Bartoli, R. Lafage, J. Morlier, J.R.R.A. Martins, A Python surrogate modeling framework with derivatives, Adv. Eng. Softw. (2019) 102662., http://dx.doi.org/10.1016/j.advengsoft.2019.03.005.

[37] A. Saltelli, M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, S. Tarantola, Global Sensitivity Analysis. The Primer, John Wiley & Sons Ltd, 2008, pp. 184–200.

[38] J. Herman, W. Usher, SALib: An open-source python library for sensitivity analysis, J. Open Source Softw. 2 (9) (2017) http://dx.doi.org/10.21105/joss.00097.

[39] S. Gantasala, N. Tabatabaei, M. Cervantes, J.-O. Aidanpää, Numerical investigation of the aeroelastic behavior of a wind turbine with iced blades, Energies 12 (12) (2019) 2422.