Bayesian meta-optimisation of variable stiffness composite cylinders for mass minimisation with manufacturing constraints

Almeida, José Humberto S.; Ashok, Aravind; Uzair, Muhammad; Castro, Saullo G.P.

DOI
10.1016/j.compstruc.2025.107868

Publication date
2025

Document Version
Final published version

Published in
Computers and Structures

Citation (APA)
Almeida, J. H. S., Ashok, A., Uzair, M., & Castro, S. G. P. (2025). Bayesian meta-optimisation of variable stiffness composite cylinders for mass minimisation with manufacturing constraints. *Computers and Structures*, *316*, Article 107868. https://doi.org/10.1016/j.compstruc.2025.107868

Computers and Structures 316 (2025) 107868

Contents lists available at ScienceDirect

Computers and Structures

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

Bayesian meta-optimisation of variable stiffness composite cylinders for mass minimisation with manufacturing constraints

José Humberto S. Almeida Jr. [a], Aravind Ashok [b], Muhammad Uzair [a], Saullo G.P. Castro [b,*]

[a] Department of Mechanical Engineering, LUT University, Lappeenranta, Finland
[b] Faculty of Aerospace Engineering, Delft University of Technology, Delft, The Netherlands

HIGHLIGHTS

- Bayesian optimisation for lightweight variable-stiffness composite design.
- Fibre paths, layer count, and thickness optimised under winding constraints.
- Bayesian optimisation uses up to 70 % fewer simulations than genetic algorithms.
- Results show smoother thickness and tailored stiffness for buckling control.
- The method enables efficient and adaptive design of advanced composite structures.

GRAPHICAL ABSTRACT

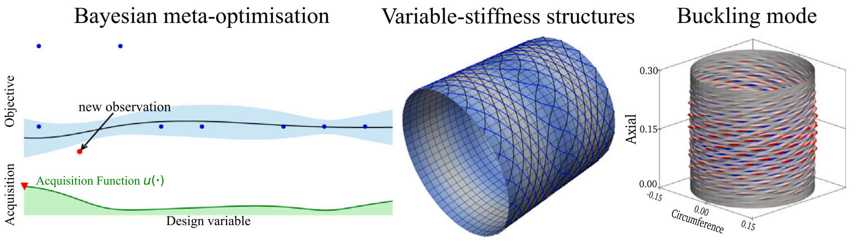

ARTICLE INFO

Keywords:
Mass minimisation
Bayesian metamodelling
Optimisation
Variable stiffness
Buckling
Manufacturing constraints

ABSTRACT

This study presents a Bayesian Optimisation (BO) framework for the mass minimisation of variable-stiffness (VS) composite cylinders under multiple buckling constraints, incorporating manufacturing limitations derived from filament winding processes. A computationally efficient single-curvature finite element model is used to evaluate the linear buckling response of multilayered shells. BO simultaneously optimises fibre paths, number of layers, and thickness distribution, achieving comparable or improved performance relative to a Genetic Algorithm (GA) while reducing simulation time by up to 70 %. Across most design loads, BO delivers structurally efficient solutions with smooth thickness transitions and local stiffness tailoring. Although GA outperformed BO in the highest load case in terms of weight and buckling capacity, BO retained competitive performance and demonstrated higher modal richness. Buckling mode analyses revealed that BO designs support mixed-mode instabilities with greater circumferential complexity, enhancing structural adaptability. In contrast, GA designs exhibited more uniform fibre paths and axial-dominated modes, reflecting conservative reinforcement strategies. These findings highlight the capability of BO to exploit complex design spaces more effectively, offering a scalable and data-efficient alternative to traditional optimisation methods. The proposed framework is particularly well suited for high-fidelity, simulation-driven design of advanced composite structures where computational cost and manufacturability are critical constraints.

---

[*] Corresponding author.
Email addresses: humberto.almeida@lut.fi (J.H.S. Almeida), aravindashok1422@gmail.com (A. Ashok), muhammad.uzair@student.lut.fi (M. Uzair), S.G.P.Castro@tudelft.nl (S.G.P. Castro).

https://doi.org/10.1016/j.compstruc.2025.107868
Received 28 March 2025; Accepted 12 June 2025

1. Introduction

The pursuit of high-performance and lightweight structures has intensified across various sectors, including aerospace, automotive, energy, and marine [1] industries. Among the promising solutions, variable stiffness (VS) composite structures [2] have emerged as a cutting-edge design paradigm due to their ability to spatially tailor fibre orientations, thereby enabling superior load-bearing and buckling resistance compared to conventional straight-fibre laminates [3,4].

This increased design freedom is primarily driven by recent advances in advanced manufacturing techniques, including automated fibre placement (AFP) [5], continuous tow shearing (CTS) [6], tailored fibre placement (TFP) [7], and more recently, additive manufacturing of continuous fibre composites [8,9] and filament winding [3,10]. These technologies allow for precise control over fibre paths, enabling engineers to manufacture increasingly complex and efficient structures [11]. However, the potential of VS composites comes at the cost of substantially larger design spaces, leading to high-dimensional and computationally expensive optimisation problems [12,13].

Traditional optimisation techniques, such as gradient-based and gradient-free methods (e.g., evolutionary algorithms), notably genetic algorithms (GA), have been extensively used in composite design [7,14,15]. While effective, these methods often require thousands of structural evaluations, making them impractical for iterative design involving high-fidelity finite element (FE) models and manufacturing constraints [16].

Bayesian optimisation (BO) is a sample-efficient, surrogate-based global optimisation technique particularly well-suited for expensive-to-evaluate black-box functions [17,18]. It constructs a probabilistic model, typically a Gaussian Process (GP), to approximate the objective function and iteratively selects new sampling points by maximising an acquisition function that balances exploration and exploitation [19].

In contrast to gradient-based methods, which require explicit derivatives and are often trapped in local minima when dealing with non-convex or noisy functions [20], BO operates without gradient information, making it robust for non-smooth or multi-modal problems. Compared to traditional gradient-free techniques such as Genetic Algorithms (GAs), BO offers faster convergence and significantly fewer function evaluations by leveraging probabilistic modelling and adaptive sampling. These features make BO particularly attractive for engineering design problems where function evaluations involve expensive simulations, such as the buckling analysis of composite structures with complex architectures [21].

To address these challenges, this study introduces a novel Bayesian optimisation (BO) framework for the mass minimisation of variable stiffness composite cylinders under multiple buckling load constraints and manufacturing feasibility limits. BO is a powerful global optimisation strategy that combines Gaussian process (GP) surrogate modelling with probabilistic acquisition functions, enabling efficient exploration of complex, high-cost design spaces with a limited number of evaluations [22–24]. Though widely adopted in machine learning, hyperparameter tuning [25] and materials informatics [26,27], its application to structural optimisation of VS composites remains largely underexplored.

This paper makes the following key contributions:

1. We propose the first Bayesian meta-optimisation framework for weight minimisation of variable stiffness composite cylinders under eigenvalue-based buckling loads and realistic manufacturing constraints;
2. We incorporate the recently developed single-curvature Bogner–Fox–Schmit–Castro (SC-BFSC) finite element formulation for accurate and efficient eigenvalue and mode-shape computations; and
3. We compare the performance of the proposed BO framework against a benchmark genetic algorithm, assessing not only the quality of the optimised solutions but also the computational efficiency.

The structure of the paper is as follows: Section 2 details the Bayesian optimisation strategy and the meta-optimisation framework; Section 3 describes the SC-BFSC finite element formulation used for buckling analysis; Section 4 presents numerical results and comparative analysis with the GA approach; and Section 5 concludes with key findings and future directions for research in data-driven optimisation of advanced composites.

2. Optimisation problem

Fig. 1 shows two adjacent filaments revolved around a cylinder with a total length of L. The blue and grey colours indicate whether a filament is frontal or dorsal concerning the back of the cylindrical shell. The filaments are always separated by the circumferential direction Δc, with $\Delta c_{min} \leq \Delta c \leq \Delta c_{max}$. The filament angle distribution $\theta(x)$ is interpolated using a second-order function interpolating three values of the filament winding angle θ_1, θ_2 and θ_3, respectively defined at $x = \{0, L/4, L/2\}$, for $0 \leq x \leq L/2$. Moreover, when $L/2 < x \leq L$ the second-order interpolation is done using θ_3, θ_4 and θ_5, defined at $x = \{L/2, 3\ L/4, L\}$, respectively. Note that $\theta_4 = \theta_2$ and $\theta_5 = \theta_1$ are used to impose a symmetric $\theta(x)$ with respect to the middle cross section of the cylinder. An overlap design is produced when using Δc_{min}, where no gaps are present because the filaments touch at the minimum filament angle; whereas Δc_{max} produces a gap design, without overlaps, where the filaments touch at the maximum filament angle. The present study focuses on overlap designs, with the premise that they are required to produce pressure-tight shells.

The main objective of the present optimisation is to minimise weight under buckling constraints by deciding on the number of layers and the independent variables defining each layer, i.e. θ_1, θ_2 and θ_3. A layer produced with overlapping variable-angle filaments of thickness h_{tow} will have a variable thickness $h(x)$ given by Eq. (1), where $\Delta\theta(x)$ represents the change in filament winding angle, as derived by Castro et al. [29] and further discussed by Wang et al. [28].

$$h(x) = \frac{h_{tow}}{\cos \Delta\theta(x)} \tag{1}$$

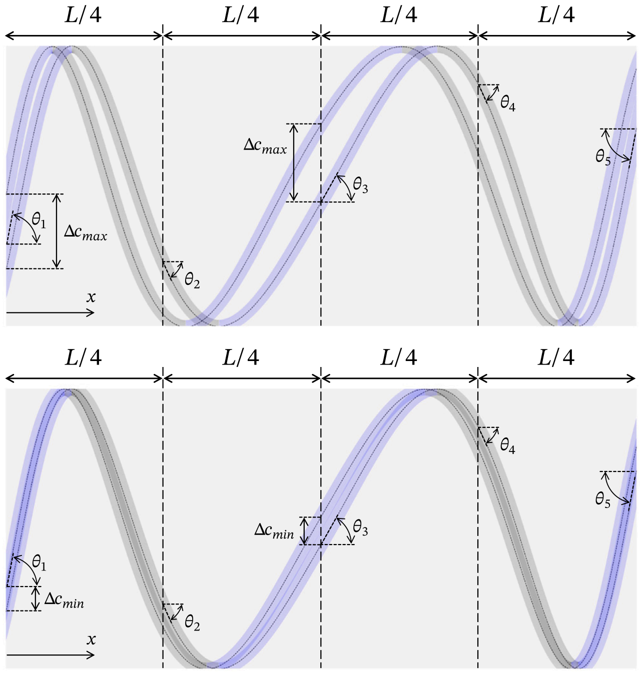

Fig. 1. Design space per layer, modified from Wang et al. [28]. Top: gap design. Bottom: overlap design.

Hence, in a variable-angle filament-wound (VAFW) design strategy, the total volume depends on the geometry of the cylinder, the number of layers, and the filament angle distribution for each layer. In the present study, the actual volume is divided by the volume of a constant-thickness single-layer shell, producing a relative volume that is more sensitive to the variables being optimised.

2.1. Constraints

With weight minimisation as the main objective, the buckling design load is introduced as the limiting constraint. Instead of introducing a constraint function in the objective function, a penalty function (**PF**) is used to keep the optimisation unconstrained, allowing an easier comparison amongst different algorithms. The penalty functions are based on the design load (**DL**) and the calculated critical buckling load of the current design (P_{cr}). The penalty function chosen is based on the square law, which is continuous over the frontier of the feasible region, keeping the information of boundary cases within the optimisation loop [30]. The penalty function based on the square-based law is given by

$$\mathbf{PF} = \max\left(1, \frac{1}{\lambda^2}\right) \tag{2}$$

$$\lambda = \frac{0.95 \times P_{cr}}{\mathbf{DL}} \tag{3}$$

Another constraint concerns manufacturability limits, where the designs are considered unfeasible when $\theta_{min} \leq \theta(x) \leq \theta_{max}$, with the following values $\theta_{min} = 3.3°$ and $\theta_{max} = 87.7°$ [28].

2.2. Design space

The optimisation problem is defined in such a way that it has the ability to decide the number of desired layers, with only a maximum number of layers given to keep the number of design variables. A conscious choice of 4 layers allows for comparisons of results from the work of Wang et al. [31] (geometrical dimensions and material properties presented in Table 1). It also allows checking the integrity of the problem without overloading CPU resources.

A maximum of 4 layers will be allowed in the investigated designs, with each layer having 4 independent design variables, being the three control points (θ_1, θ_2 and θ_3); and a Boolean variable **S** used to activate (**S** = True) or deactivate (**S** = False) each layer individually. Therefore, the problem has 16 design variables, wherein 12 variables define the composite steering path, and 4 Boolean variables control the topology, by defining the number of layers. The produced layups are balanced and the filament winding angle limits are $[3.3°, 87.7°]$ [28], here adopting with a discretisation of 4.68°, resulting in a design space of 2.21×10^{15} unique combinations.

The design space is illustrated in the buckling load versus mass graph of Fig. 2. Each number of layers is colour-categorised, resulting in a good representation of the buckling load limits for each case. It is an interesting observation that there is a gap between the design sub-spaces using 1 and 2 layers, which was also observed for other cylinder aspect ratios.

Mathematically, the optimisation problem can be written as:

$$\min_{\theta, \mathbf{S}} \quad \mathbf{Vol}([\boldsymbol{\theta}, \mathbf{S}]) \cdot \mathbf{PF}(\boldsymbol{\theta}, \mathbf{S}) \tag{4}$$

$$\text{where} \quad [\boldsymbol{\theta}, \mathbf{S}] = [\theta_1, \theta_2, \theta_3, S] \quad \text{for each independent layer set}$$

$$\text{subject to} \quad \theta_{min} \leq \theta(x) \leq \theta_{\max} \tag{5}$$

with:

- $\boldsymbol{\theta}$ representing the set of control points for the filament-winding angle distribution, for each independent layer.
- **S** used to activate or deactivate a layer.
- $\mathbf{PF}(\boldsymbol{\theta}, \mathbf{S})$ ensuring that designs not meeting the buckling load requirement are penalised.
- The constraint $\theta_{min} \leq \theta(x) \leq \theta_{\max}$ enforces manufacturing limitations on minimum and maximum filament winding angles.

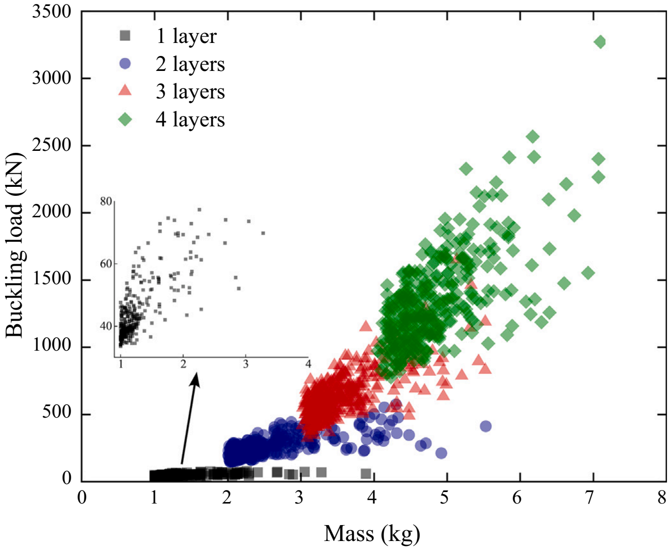

Fig. 2. Illustration of the design space.

Table 1
Geometric dimensions of the cylinders and material properties used in the optimisations.

Geometric dimensions	
L	0.3 m
D	0.3 m
Material properties	
E_{11}	90 GPa
E_{22}	7 GPa
ν_{12}	0.32
$G_{12} = G_{13}$	4.4 GPa
G_{23}	1.8 GPa
h_{tow}	0.4×10^{-3}m

3. The Bayesian optimisation framework

Fig. 3 schematically illustrates the working mechanism of the proposed BO framework [32]. The method utilises a surrogate model, typically a Gaussian Process, to approximate the true objective function based on a limited number of evaluations.

The horizontal axis represents a design variable, while the vertical axes show the objective function (top) and the acquisition function (bottom). The black curve in the upper plot represents the predicted surrogate model $\hat{f}(x)$, constructed from previously evaluated data points, which are shown as blue dots. The shaded region around the curve denotes the prediction uncertainty, capturing the confidence interval of the model.

The lower plot displays the acquisition function $u(x)$, which quantifies the utility of evaluating the objective function at each point in the design space. It guides the selection of the next evaluation point by balancing exploration (sampling regions with high uncertainty) and exploitation (sampling regions with low predicted objective values). The red point highlights the location of the next sample, selected at the maximum of the acquisition function.

This iterative approach allows the optimiser to efficiently explore the design space while minimising the number of expensive function evaluations. In the context of this study, the BO framework enables the discovery of lightweight, buckling-resistant composite configurations with significantly reduced computational effort.

BO uses acquisition functions for evaluating the likelihood of the points. The acquisition function uses the mean and variance values

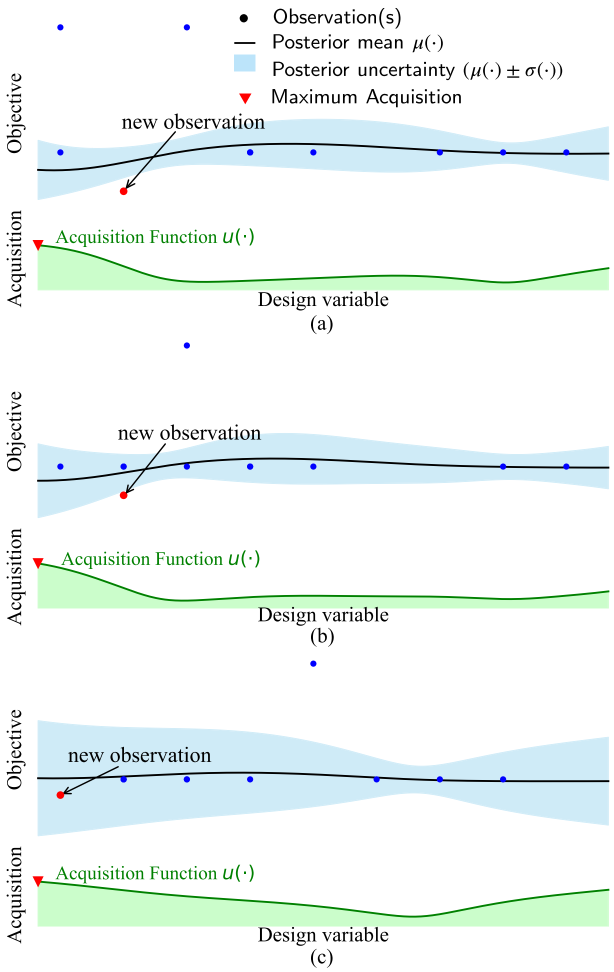

Fig. 3. The design space of the Bayesian optimisation framework over iterations (a) $t = 15$, (b) $t = 23$, and (c) $t = 33$.

provided by the Gaussian process regression model and calculates scores based on its Acquisition type. If $\mathbf{X}$ are the points evaluated and fit to a Gaussian model, and $\mathbf{X}^*$ are the new points to be assessed, then the acquisition function scores can be given as:

$$\text{Scores} = \text{Acq. func}\left(\mathcal{GP}\left(\mathbf{X}^*, \mathbf{X}\right)\right) \tag{6}$$

where $\mathcal{G}.\mathcal{P}$ is the Gaussian Process (GP) model, which outputs the mean and variance for the prediction of the output of new points $\mathbf{X}^*$ concerning $\mathbf{X}$.

With the goal of achieving an optimal surrogate model fit, an exhaustive design of experiments (DoE) is performed to understand how the BO framework is influenced by different hyperparameters, different kernels for the GP model, initial sample size and population size. Two test datasets of sizes 1000 and 5000 are used. The selection of kernels is based on the maximum log-likelihood value and the total computational time. Upon investigation, the Matern32 kernel is found to be best suited for the problem at hand. The selection of the initial sample size is based on the accuracy required by the GP model. The accuracy and error calculation are examined with the help of R^2 variance score and Mean Squared Error (MSE). The sample size is defined as a factor of input variables that would help with a more generic representation of the parameters for future use. With both the test data, it is seen that 10 times the input provides a good prediction accuracy.

After investigating the GP model and initial sample size, three acquisition functions are assessed, again aiming for an optimum BO framework: Probability of improvement (PI); Expected Improvement (EI) and Lower Confidence Bound (LCB). The characteristics of each function and exploration–exploitation weights are evaluated. Individually, LCB is found to be the most efficient. However, an overall improvement is observed when the acquisition functions are used on a rotational basis with each iteration. This is because although the selection of sample points has a good distribution over the design space, it is random in nature.

For consistency with the GA setup, the maximum number of iterations in the BO framework is fixed at 2500. To minimise redundant evaluations once the design space becomes saturated, two early-stopping criteria are implemented. The first criterion monitors the optimised weights and terminates the optimisation if 25 non-penalised designs with identical volume (rounded to ten decimal places), are encountered, indicating stagnation. The second criterion activates after 1000 iterations and terminates the run if the current optimum remains unchanged for 500 successive iterations, implying a lack of further improvement. These criteria ensure that the optimisation process remains efficient while maintaining the potential to identify high-quality solutions.

3.1. Surrogate model verification

The surrogate verification is achieved by predicting for test points when trained on a separate set of training points, wherein the deviation between the prediction and actual values is compared. The performance of different training hyperparameters is compared using the coefficient of determination/variance score $-\infty \leq R^2 \leq 1$, defined as the ratio between the variance explained by the model to the total variance [33]:

$$R^2 = 1 - \frac{\sum_{i=1}^{N_{test}} \left(y_i - \hat{y}_i\right)^2}{\sum_{i=1}^{N_{test}} \left(y_i - \bar{y}_i\right)^2}, \tag{7}$$

where y_i is the true response, $\hat{y}_i$ is the model prediction for the point, N_{test} is the total number of test samples and $\bar{y}_i$ is the mean of true test values. A good model can be characterised through an R^2 greater than 0.8; when $0.5 \leq R^2 \leq 0.8$ there is the presence of observable error in the model; and $R^2 < 0.5$ implies a weak model prediction and the presence of significant error. Values $R^2 < 0$ signify that the prediction is worse than the average model and should not be used.

The MSE is used to understand the error in the model. MSE is the mean of the overall squared prediction error, which is shown in Eq. (8):

$$MSE = \frac{1}{N_{test}} \sum_{i=1}^{N_{test}} \left(y_i - \hat{y}_i\right)^2 \tag{8}$$

A perfect model should have an MSE of 0, while any good model should have it close to zero.

3.1.1. Testing methods

There are several methods to test the surrogate model [33], such as Jackknife error, bootstrap error and K-fold cross-validation error. The K-fold cross-validation error splits the testing and training points. The model is trained with all subsets, and the process is repeated until each subset becomes the test set. If the samples can render a good representation of the design space, then the K-fold test will provide a good average R^2 score with minimum variance. Here, the verification is performed with a 5-fold cross-validation, as shown in Fig. 4.

3.2. Inputs and outputs

For the surrogate modelling, 2 cases are considered: 1st case is the use of input variables as orientation angles at control points and output as the objective; In the 2nd case, the inputs remain the same but the output used for modelling is changed to buckling load from the objective function. The reason is that this will give a direct correlation between the design variables and the buckling load, and the complexity of using penalty functions (PF) and manufacturing constraints is removed. The

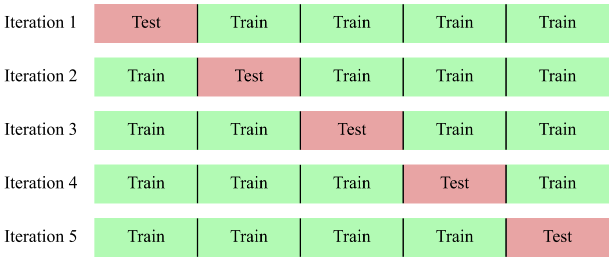

Iteration 1	Test	Train	Train	Train	Train
Iteration 2	Train	Test	Train	Train	Train
Iteration 3	Train	Train	Test	Train	Train
Iteration 4	Train	Train	Train	Test	Train
Iteration 5	Train	Train	Train	Train	Test

Fig. 4. The proposed 5-fold cross-validation process.

predicted buckling load is used to calculate the PF, which is multiplied by the relative volume to obtain the output.

With the help of K-fold cross-validation with R^2 and MSE scores, it is observed that in the 1st case, even though a more complex objective is used, it is very efficient, and the GP model is able to predict with good accuracy. The prediction is similar between the reduced switch variable and normal (15×16, refer to Section 4), even though there is a reduction in the number of input variables. The GP model can interpret the Boolean switch variables effectively. Interestingly, the 2nd case is the slowest in the optimisation process and the GP model prediction is lower than Case 1. This is attributed to post-processing of the predicted buckling loads, leading to larger errors, and the slower nature can be attributed to the second calculation of relative volume and objective function, which even though minor, added up when the data points increased.

3.3. Acquisition functions

The acquisition function utilises the information from the GP model about the input variables and translates them into scores. The highest score points are chosen for running the actual test, as this has the highest chance of being the optimal point or a point from an unexplored region. Three acquisition functions are focused on here: Probability of Improvement (PI), Expected Improvement (EI) and Lower Confidence Bound (LCB). The focus is on the static state of the acquisition function.

Probability of improvement

Probability of Improvement (PI) was first introduced by Kushner [34]. The function measures the maximum probability of a point improvement over its current best-known value. The function is purely exploitative in nature, which is a drawback since this may cause it to ignore points that have good value but high uncertainty. To reduce this effect, a trade-off parameter ξ is introduced in the function. It is formulated as Eq. (9). $\Phi()$ is the normal cumulative distribution function (CDF). In simple terms, CDF provides the probability of a function yielding a result lower than or equal result than that provided.

$$\begin{aligned} \mathrm{PI}(x,\xi) &= P\left[F_x \geq \mu_{\max} + \xi\right] \\ &= 1 - \Phi\left(\frac{\mu(x) - \left(\mu_{\max} + \xi\right)}{\sigma(x)}\right) \\ &\equiv \left(\frac{\mu(x) - \left(\mu_{\max} + \xi\right)}{\sigma(x)}\right) \end{aligned} \tag{9}$$

Expected improvement:

Although PI introduced the parameter ξ, it is still limited in selecting points likely to improve the objective. Expected Improvement (EI) is proposed to address this shortcoming [35,36].

Define:

$$Z = \begin{cases} \frac{\mu(\mathbf{x}) - \mu_{\max} - \xi}{\sigma(\mathbf{x})}, & \text{if } \sigma(\mathbf{x}) > 0 \\ 0, & \text{if } \sigma(\mathbf{x}) = 0 \end{cases} \tag{10}$$

Then, the acquisition function for EI is given by:

$$\mathrm{EI}(\mathbf{x}) = \begin{cases} \left(\mu(\mathbf{x}) - \mu_{\max} - \xi\right)\Phi(Z) + \sigma(\mathbf{x})\phi(Z), & \text{if } \sigma(\mathbf{x}) > 0 \\ 0, & \text{if } \sigma(\mathbf{x}) = 0 \end{cases} \tag{11}$$

In Eq. (11), the first term promotes exploitation, and the second encourages exploration. Φ and ϕ denote the cumulative distribution function (CDF) and the probability density function (PDF) of the standard normal distribution, respectively. The parameter ξ controls the trade-off between exploration and exploitation; a typical value is $\xi = 0.01$ [37,38], which can be increased to encourage more exploration. For large ξ, the algorithm approaches random sampling.

Lower/upper confidence bound:

The lower/upper confidence bound (LCB or UCB) is used for minimisation/maximisation problems. It tries to manage exploration–exploitation by being optimistic when faced with uncertainty. Eq. (12) gives the formulation for LCB, which is quite simple in nature. LCBs tend to require fewer iterations to reach the global optima [39,40].

$$\mathrm{LCB}(x) = \mu_x - \xi\sigma(x) \tag{12}$$

3.4. The BO framework

The BO begins with an initial population size, which is selected using the LHS sampling criterion. The maximum criterion is chosen in the Latin Hypercube Sampling (LHS), maximising the minimum Euclidean distance between points. These initial data points and their corresponding calculated objectives are used to fit the Gaussian regression model. The acquisition functions take the points and, with the help of the GP model, assess the points from the predicted mean, while the variance provides scores for each point. The best-scored point is evaluated, and

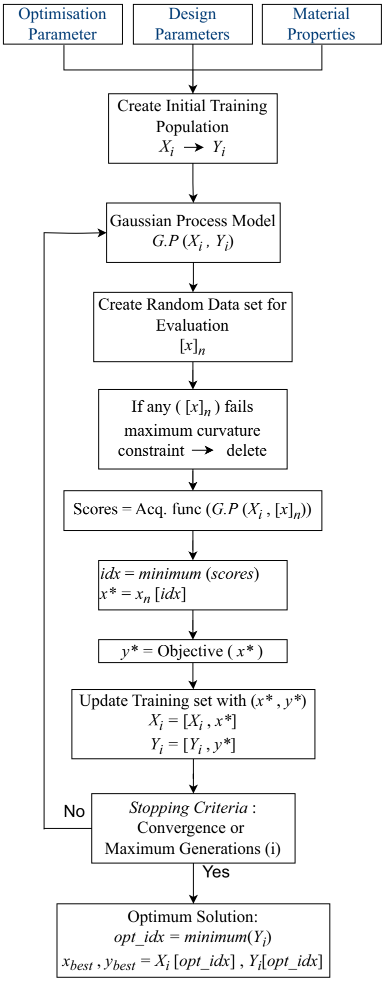

Fig. 5. The flowchart for the Bayesian optimisation framework.

Table 2
Performance of acquisition functions with varying weight factors.

Acquisition function	Weight	Obj. (50 kN)	Conv. (50 kN)	Obj. (1000 kN)	Conv. (1000 kN)
MPI	–	0.16622	240	0.00397	150
LCB	0.1	0.15602	360	0.00372	155
	0.2	0.15880	325	0.00397	148
	0.3	0.15581	320	0.00389	350
	0.4	0.16325	310	0.00384	290
	0.6	0.16112	310	0.00389	240
	0.8	0.16720	340	0.00389	240
EI	0.01	0.16290	355	0.00406	295
	0.025	0.15601	360	0.00410	310
	0.05	0.16278	380	0.00383	105
	0.1	0.16286	365	0.00372	155
Combo (LCB, PI, EI)	–	0.15410	245	0.00395	380

Table 3
The BO parameters.

Parameter	Value
Initial sampling size	10 times the input variables
Angle increment	0.5
Number of iterations	2500
Tolerance	1e-3
Population size	50
Acquisition function	Combination of EI, LCB and PI
Exploration weight	EI = 0.025,LCB = 0.3, PI = NA

its actual objective value is then added to the training population of inputs and outputs to be fit. The buckling constraint is part of the objective function, and the manufacturing constraints are applied externally. If constraints are not met, then the corresponding training point is not considered, avoiding spending time on unfeasible sets. This process is repeated until a maximum number of generations is reached or when the predicted solution has converged over a certain number of generations. The framework is illustrated in Fig. 5.

For finding the best parameters, the use of pseudo-random sets is considered. This helps replicate results that are used to compare the performance when parameters are changed. This is useful for the comparison of the acquisition functions. The performance of the acquisition functions is measured by monitoring their convergence to the minimum and their ability to find the optimum solution among the selected options.

Since each acquisition function prioritises different points, the idea proposed by Brochu et al. [41] could have been used to make use of all 3 sets of acquisition functions, selecting the best point from all sets. The disadvantage is that this would have increased the actual calculation by a factor of 3 (if all 3 MPI, EI and LCB are considered). Instead, the use of the acquisition function on a rotation-based approach is proposed. This uses a combination of LCB, EI and PI, which change at each iteration. After preliminary investigations, a rotation strategy with acquisition functions using EI (weight = 0.025), LCB (weight = 0.3) and PI, is chosen. The performance of the acquisitions function with varying weight factors and the parameters in the BO framework are shown in Tables 2 and 3.

3.5. Bayesian optimisation setup

The current framework is implemented in the Python language. Other than the standard libraries, “Scikit-Learn (sklearn)” module is used for the GP model fit [42] and “Scikit-Optimise (skopt)” is employed for LHS sampling [43].

4. Linear buckling in VAFW cylinders

4.1. Discretisation

The Single-Curvature Bögner-Fox-Schmit-Castro (SC-BFSC) is adopted to accurately calculate the linear buckling of variable-stiffness filament wound cylinders. As first proposed by Castro and Jansen [44], the BFSC introduces in-plane displacement (u, v) and its first derivatives as shown in Eq. (13), creating 10 DOFs per node, with $i = 1, 2, 3, 4$:

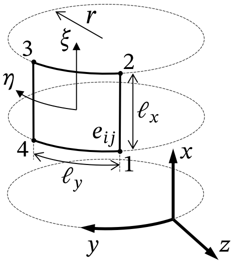

Fig. 6. The SC-BFSC element and coordinate system used to construct the cylindrical shell.

$$u_i^e = \left\{ u^i, u_x^i, u_y^i, v^i, v_x^i, v_y^i, w^i, w_x^i, w_y^i, w_{xy}^i \right\}^\top \tag{13}$$

The SC-BSFC has been successfully applied to composite shells using Sanders’ kinematics [31,45]. The cylinder is modelled with the x axis representing the axial (from $0 \rightarrow L$), y the circumferential (from $0 \rightarrow 2\pi r$) and z the radial normal direction, as illustrated in Fig. 6. For the finite element discretisation, the cylinder is divided into n sections along its axis and m along its circumference, creating $m \times n$ elements, which are kept with an aspect ratio of approximately one. The displacement fields u, v, w, are approximated as:

$$u_e, v_e, w_e = \sum_{i=1}^{4} S_i^{u,v,w} u_i^e \tag{14}$$

where $S_i^{u,v,w}$ are Hermite shape function matrices [46] that allow the in-plane and out-of-plane displacements to be approximated using 3rd-order polynomials:

$$\begin{aligned}
\boldsymbol{S}_i^u &= \begin{bmatrix} H_i & H_i^x & H_i^y & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \\
\boldsymbol{S}_i^v &= \begin{bmatrix} 0 & 0 & 0 & H_i & H_i^x & H_i^y & 0 & 0 & 0 & 0 \end{bmatrix} \\
\boldsymbol{S}_i^w &= \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & H_i & H_i^x & H_i^y & H_i^{xy} \end{bmatrix}
\end{aligned} \tag{15}$$

with the cubic Hermite functions H_i, H_i^x, H_i^y, H_i^{xy} calculated using natural coordinates as per [47]:

$$\begin{aligned}
H_i &= \frac{1}{16}(\xi + \xi_i)^2(\xi\xi_i - 2)(\eta + \eta_i)^2(\eta\eta_i - 2) \\
H_i^x &= -\frac{\ell_x}{32}\xi_i(\xi + \xi_i)^2(\xi\xi_i - 1)(\eta + \eta_i)^2(\eta\eta_i - 2) \\
H_i^y &= -\frac{\ell_y}{32}(\xi + \xi_i)^2(\xi\xi_i - 2)\eta_i(\eta + \eta_i)^2(\eta\eta_i - 1) \\
H_i^{xy} &= \frac{\ell_x \ell_y}{64}\xi_i(\xi + \xi_i)^2(\xi\xi_i - 1)\eta_i(\eta + \eta_i)^2(\eta\eta_i - 1)
\end{aligned} \tag{16}$$

where ℓ_x, ℓ_y are respectively the finite element dimensions along x, y, as represented in Fig. 6; and ξ, η are natural coordinates:

$$\begin{aligned} \xi &= \frac{2(x - x_c)}{\ell_x} \qquad (-1 \le \xi \le 1) \\ \eta &= \frac{2(y - y_c)}{\ell_y} \qquad (-1 \le \eta \le 1) \end{aligned} \tag{17}$$

such that ξ_i, η_i at each node are $(-1, -1)$, $(+1, -1)$, $(+1, +1)$ and $(-1, +1)$. With the proposed nodal connectivity for the SC-BFSC element, the nodal degrees-of-freedom $\boldsymbol{u}_{\boldsymbol{e}i}$ and the respective shape functions $\boldsymbol{S}_i^{u,v,w}$ are concatenated as:

$$\begin{aligned} \boldsymbol{u_e} &= \left\{ \boldsymbol{u_{e}}_1 \quad \boldsymbol{u_{e}}_2 \quad \boldsymbol{u_{e}}_3 \quad \boldsymbol{u_{e}}_4 \right\}^\top \\ \boldsymbol{S}^u &= \left[\boldsymbol{S}_1^u \quad \boldsymbol{S}_2^u \quad \boldsymbol{S}_3^u \quad \boldsymbol{S}_4^u \right] \\ \boldsymbol{S}^v &= \left[\boldsymbol{S}_1^v \quad \boldsymbol{S}_2^v \quad \boldsymbol{S}_3^v \quad \boldsymbol{S}_4^v \right] \\ \boldsymbol{S}^w &= \left[\boldsymbol{S}_1^w \quad \boldsymbol{S}_2^w \quad \boldsymbol{S}_3^w \quad \boldsymbol{S}_4^w \right] \end{aligned} \tag{18}$$

with $\boldsymbol{S}^u$, $\boldsymbol{S}^v$ and $\boldsymbol{S}^w$ being matrices of shape 1×40.

4.2. Total potential energy

Applying equivalent single-layer (ESL) theory, which assumes the heterogeneous laminate to be statistically equivalent to a single layer, the total potential energy for an SC-BFSC finite element is defined as Eq. (19):

$$\Phi = \frac{1}{2} \int_{y=y_1}^{y_4} \int_{x=x_1}^{x_2} (\boldsymbol{N}\boldsymbol{\epsilon} + \boldsymbol{M}\boldsymbol{\kappa})\, dx\, dy \tag{19}$$

where $\boldsymbol{N} = \left\{ N_{xx}, N_{yy}, N_{xy} \right\}^T$ represents the resultant membrane force and $\boldsymbol{M} = \left\{ M_{xx}, M_{yy}, M_{xy} \right\}^T$ are its distributed moments. The subscripts used in the integration limits indicate the node connectivity. Sanders' kinematics for cylindrical shell plates is adopted, as per Eq. (20):

$$\begin{aligned} \boldsymbol{\varepsilon} &= \begin{Bmatrix} \varepsilon_{xx} \\ \varepsilon_{yy} \\ \gamma_{xy} \end{Bmatrix} = \begin{Bmatrix} u_{,x} + \frac{1}{2} w_{,x}^2 \\ v_{,y} + \frac{1}{r} w + \frac{1}{2} w_{,y}^2 + \frac{1}{2} \frac{1}{r^2} v^2 - \frac{1}{r} v w_{,y} \\ u_{,y} + v_{,x} + w_{,x} w_{,y} - \frac{1}{r} v w_{,x} \end{Bmatrix} \\ \boldsymbol{\kappa} &= \begin{Bmatrix} \kappa_{xx} \\ \kappa_{yy} \\ \kappa_{xy} \end{Bmatrix} = \begin{Bmatrix} -w_{,xx} \\ -w_{,yy} + \frac{1}{r} v_{,y} \\ -2w_{,xy} + \frac{1}{r} v_{,x} \end{Bmatrix} \end{aligned} \tag{20}$$

with $(\cdot)_{,x} = \partial(\cdot)/\partial x$ used as a compact notation for partial derivatives. Here, ϵ represents the extensional strain and κ the rotational strain.

At the bifurcation point, there exists an equilibrium for the domain of the structure or assembly of finite elements, as shown in Eq. (21).

$$\delta\Phi = \sum_{e=1}^{n_e} \delta\Phi_e = \sum_{e=1}^{n_e} \int_{\Omega_e} \left(\boldsymbol{N}^T \delta\epsilon + \boldsymbol{M}^T \delta\kappa \right) d\Omega_e = 0 \tag{21}$$

The strain variation can be calculated by taking partial differentiation of the displacements and their shape functions as shown in Eq. (22). The curvature of the cylinder is factored into the circumferential strain from out-of-plane displacements.

$$\begin{aligned} \delta\boldsymbol{\varepsilon} &= \boldsymbol{B_m} \delta \boldsymbol{u_e} \\ \delta\boldsymbol{\kappa} &= \boldsymbol{B_b} \delta \boldsymbol{u_e} \end{aligned} \tag{22}$$

where $\boldsymbol{B_m}$ and $\boldsymbol{B_b}$ are defined as:

$$\boldsymbol{B_m} = \begin{bmatrix} \boldsymbol{S}_{,x}^u + w_{,x} \boldsymbol{S}_{,x}^w \\ \boldsymbol{S}_{,y}^v + \frac{1}{r} \boldsymbol{S}^w + w_{,y} \boldsymbol{S}_{,y}^w + \frac{1}{r^2} v \boldsymbol{S}^v - \frac{1}{r} v \boldsymbol{S}_{,y}^w - \frac{1}{r} w_{,y} \boldsymbol{S}^v \\ \boldsymbol{S}_{,y}^u + \boldsymbol{S}_{,x}^v + w_{,x} \boldsymbol{S}_{,y}^w + w_{,y} \boldsymbol{S}_{,x}^w - \frac{1}{r} v \boldsymbol{S}_{,x}^w - \frac{1}{r} w_{,x} \boldsymbol{S}^v \end{bmatrix} \tag{23}$$

$$\boldsymbol{B_b} = \begin{bmatrix} -\boldsymbol{S}_{,xx}^w \\ -\boldsymbol{S}_{,yy}^w + \frac{1}{r} \boldsymbol{S}_{,y}^v \\ -2\boldsymbol{S}_{,xy}^w + \frac{1}{r} \boldsymbol{S}_{,x}^v \end{bmatrix} \tag{24}$$

Noting that the partial derivatives of $\boldsymbol{S}^{u,v,w}$ are directly calculated from the shape functions of Eq. (18) in terms of the natural coordinates ξ, η. These derivatives are obtained using the Jacobian relations.

$$\begin{aligned} \frac{\partial}{\partial x} &= \frac{\ell_x}{2} \frac{\partial}{\partial \xi}, \\ \frac{\partial}{\partial y} &= \frac{\ell_y}{2} \frac{\partial}{\partial \eta}. \end{aligned} \tag{25}$$

4.3. Neutral equilibrium criterion

The neutral equilibrium criterion implies $\delta^2 \Phi = 0$, which is shown in Eq. (26).

$$\begin{aligned} \delta^2 \phi = \sum_{e=1}^{n_e} \delta^2 \phi_e = \sum_{e=1}^{n_e} \Bigg[&\int_{\Omega_e} \left(\delta \boldsymbol{N}^\top \delta \boldsymbol{\varepsilon} + \delta \boldsymbol{M}^\top \delta \boldsymbol{\kappa} \right) d\Omega_e \\ &+ \int_{\Omega_e} \left(\boldsymbol{N}^\top \delta^2 \boldsymbol{\varepsilon} + \boldsymbol{M}^\top \delta^2 \boldsymbol{\kappa} \right) d\Omega_e \Bigg] = 0 \end{aligned} \tag{26}$$

where the first part of the equation forms the constitutive stiffness matrix, and the second part represents the geometric stiffness, both being matrices of size 40×40. In the constitutive stiffness matrix, for linear analysis, the non-linear terms, when expanded into displacement terms containing w_x, w_y and w_{xy}, are ignored. The final expression is shown in Eq. (27).

$$\boldsymbol{K_e} = \iint_{xy} (\boldsymbol{B_m}^\top \boldsymbol{A} \boldsymbol{B_m} + \boldsymbol{B_m}^\top \boldsymbol{B} \boldsymbol{B_b} + \boldsymbol{B_b}^\top \boldsymbol{B} \boldsymbol{B_m} + \boldsymbol{B_b}^\top \boldsymbol{D} \boldsymbol{B_b}) dx\, dy \tag{27}$$

with $\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{D}$ being the laminated composite constitutive matrices, which are calculated based on the correct offset $d(x)$ of each filament-wound layer, rendering the correct positioning of each layer concerning the filament winding mandrel radius, as illustrated in Fig. 7.

The second integral is the geometric stiffness matrix $\boldsymbol{K_{G0}}$, and provides information on the non-linear effects of pre-buckling membrane stress $\boldsymbol{N_0}^\top = \{N_{0xx}, N_{0yy}, N_{0xy}\}$. Here it is considered that $\delta^2 << \delta$ [48], leading to Eq. (28):

$$\boldsymbol{K_{G0e}} = \iint_{xy} \begin{bmatrix} \boldsymbol{S}_{,x}^{wT} N_{0xx} \boldsymbol{S}_{,x}^w \\ \boldsymbol{S}_{,y}^{wT} N_{0yy} \boldsymbol{S}_{,y}^w \\ \boldsymbol{S}_{,y}^{wT} N_{0xy} \boldsymbol{S}_{,x}^w + \boldsymbol{S}_{,x}^{wT} N_{0xy} \boldsymbol{S}_{,y}^w \end{bmatrix} dx\, dy \tag{28}$$

4.4. Eigenvalue problem

The pre-buckled stress $\boldsymbol{N_0}$ of Eq. (28) is calculated from the nodal displacement of the corresponding element $\boldsymbol{u_{0e}}$. For a symmetric laminate, it can be written as per Eq. (29):

$$\boldsymbol{N_0} = \begin{Bmatrix} N_{0xx} \\ N_{0yy} \\ N_{0xy} \end{Bmatrix} = \boldsymbol{A} \begin{bmatrix} \boldsymbol{S^u}_{,x} \\ \boldsymbol{S^v}_{,y} + \frac{1}{r} \boldsymbol{S^u}_{,x} \boldsymbol{S^w} \\ \boldsymbol{S^u}_{,y} + \boldsymbol{S^v}_{,x} \end{bmatrix} \boldsymbol{u_{0e}} \tag{29}$$

where $\boldsymbol{u_{0e}}$ is calculated from the global displacement vector, solved using the static analysis given by:

$$\boldsymbol{u_0} = \boldsymbol{K}^{-1} \boldsymbol{f_0} \tag{30}$$

with $\boldsymbol{f_0}$ being the pre-buckling load. It is assumed that at the bifurcation point, there is a value of internal membrane stresses $\boldsymbol{N} = \lambda \boldsymbol{N_0}$ that satisfies the neutral equilibrium:

$$\delta \boldsymbol{u}^T \left(\boldsymbol{K} + \lambda \boldsymbol{K_{G0}} \right) = 0 \tag{31}$$

where $\boldsymbol{K}$ is the constitutive stiffness matrix and $\boldsymbol{K_{G0}}$ is the geometric stiffness matrix. For Eq. (31) to be zero, the determinant of the stiffness part should be zero:

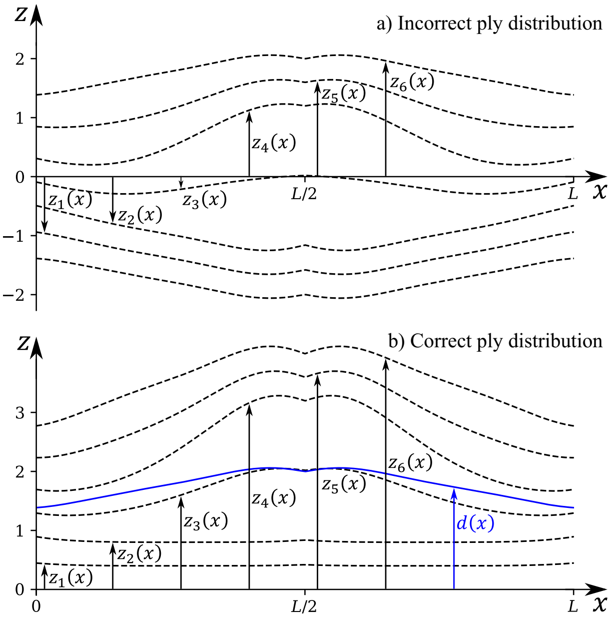

Fig. 7. (a) Incorrect ply distribution with a constant mid-surface at $z = 0$; and (b) correct ply distribution with a varying offset $d(x)$ applied to the mid-surface. The outward face position of each ply $z_k(x)$ for $k = 1, 2, \dots, 6$ is shown. Reproduced from Wang et al. [31].

$$\det\left(\boldsymbol{K} + \lambda \boldsymbol{K_{G0}}\right) = 0 \tag{32}$$

leading to an n^{th} order characteristic polynomial, whose roots λ represent the linear buckling eigenvalues.

5. Results and discussion

The GA-based optimisation is implemented [32] to serve as a benchmark for the BO framework, using the Python module OpenMDAO [49]. In GA, the *discretisation* of the design space of each variable is performed by defining the bits. In this algorithm, the design space is divided into the order of binary numbers (2^n). Here, the bit size is taken as 6 for all the steering angle variables. Therefore, the design space is divided into 128 (2^6) parts with increments of 1.32 deg.

Population Size: The recommended value for population size is 4 times the bit size. After trial and error with different sizes, 4 times the bit size is found to perform the best.

Mutation rate: The default value for mutation is 0.01. This was increased to 0.02, which resulted in better results. The possible reason for the improvement is the higher number of mutated children that are available, which results in better exploration.

Maximum generations: The convergence is tested with two loading cases: 50 kN, which is a very safe value for layer 1, and 500 kN, which is on the upper limit of using 2 layers and lower limit of employing 3 layers. This means that the minimum weight option is the 2-layer case, which is a tricky case for the optimiser. This helps to show the versatility of the solver in finding both a straightforward design load and a tricky one. The convergence in the GA occurs at about 95 generations, which is found to be similar for the rest of the load cases. Therefore, a maximum generation of 100 can be taken as the limit. The set of parameters used in GA is given in Table 4.

5.1. Verifying BO with GA

Table 4
The parameters used in the GA optimisations.

Parameters	Value
Maximum generation	100
bits size	6
Mutation rate P_m	0.02
Crossover rate P_c	0.1
Elitism	True

To verify the performance of the BO framework, a comparative analysis with GA is conducted. Fig. 8 displays the convergence histories for both methods under two loading conditions: 50 kN and 500 kN. The results demonstrate that BO outperformed GA in terms of convergence speed and solution quality. In the 50 kN case, BO converged to a near-optimal mass below 0.4 kg within approximately 300 iterations, whereas GA required over 1000 iterations to reach a similar mass. For the more demanding 500 kN case, BO achieved a converged solution near 1.1 kg in under 1200 iterations, while GA required up to 1800 iterations and still converged to a slightly heavier structure.

An interesting trend observed in the convergence behaviour of BO is its tendency to identify near-optimal low-mass solutions early in the optimisation cycle, often within the first 150 to 200 iterations. However, in later stages, the optimiser shifts towards exploring structurally more robust configurations. This may result in a slight increase in mass but better satisfy buckling and manufacturability constraints. This behaviour is particularly evident in the 50 kN case, where convergence is not solely driven by early-stage gains but also by subsequent refinements that align the solution with stricter feasibility limits. This pattern highlights the capacity of BO in balancing rapid early improvements with continued exploration for resilient, constraint-satisfying designs.

BO convergence curves are smoother and more stable, reflecting its data-efficient search strategy and the effectiveness of its surrogate modelling in navigating the design space. In contrast, GA exhibited a more erratic convergence pattern, especially for the 500 kN load case, suggesting a broader but less directed exploration of the design space.

These findings reinforce the advantages of BO in handling computationally expensive optimisation tasks. Its ability to efficiently balance exploration and exploitation makes it particularly well-suited for the design of advanced composites with complex design spaces. While GA remains a viable and widely used baseline method, BO offers significant computational savings and improved performance.

Optimisation results are presented for five representative design load cases: 50 kN, 100 kN, 200 kN, 500 kN, and 1000 kN. These load levels are selected to enable direct comparison with the benchmark solutions reported by Wang et al. [31], allowing for a more rigorous evaluation of the proposed BO strategy. Among the acquisition strategies explored, a rotational scheme, alternating between Probability of Improvement (PI), Expected Improvement (EI), and Lower Confidence Bound (LCB), offered the most favourable trade-off between convergence speed and computational cost. Exploration weights are set to 0.025 for EI and 0.3 for LCB to balance local exploitation and global exploration during the search effectively.

The optimised design variables and corresponding structural responses for both BO and GA are summarised in Table 5. Across all load cases, both methods achieved structurally sound solutions that meet buckling performance requirements while minimising structural weight. For example, in the 50 kN case, BO identified a configuration weighing 0.4135 kg that sustains a buckling load of 53.00 kN. As the applied load increases, both the structural complexity and number of design variables grow, as reflected in the 1000 kN case, where BO proposed a multi-layered configuration weighing 1.4102 kg and resisting up to 1038.45 kN of compressive load.

These results highlight the robustness of the BO framework in exploring moderately sized design spaces and converging toward lightweight, high-performance solutions. The diversity in optimal configurations across load levels illustrates the adaptability of the method to different structural demands.

A quantitative comparison of the optimal solutions obtained using BO and GA, summarised in Table 5, shows that BO consistently achieves results closely aligned with those of GA while significantly reducing

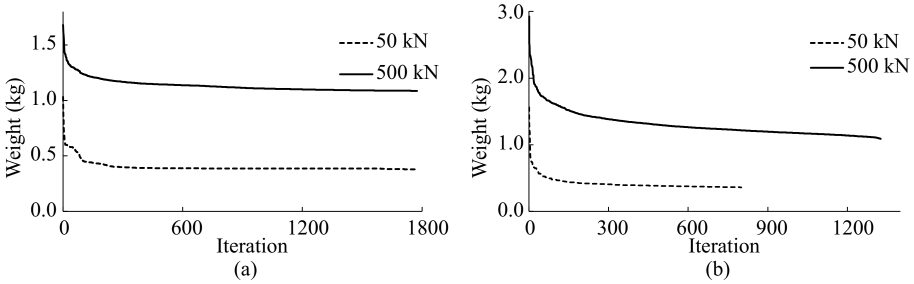

Fig. 8. Convergence plots for both (a) GA and (b) BO frameworks for two design loads.

Table 5
Comparison of BO and GA results for different design loads.

Design load	Method	Optimised design variables	Weight (kg)	Buckling load (kN)	Weight Diff. (%)*
50 kN	BO	[59.5, 87.5, 40.7]	0.4135	53.00	9.34
	GA	[42.3, 63.3, 53.3]	0.3782	52.80	–
100 kN	BO	[22.0, 22.0, 31.4] / [12.7, 12.7, 12.7]	0.7248	176.48	0.14
	GA	[7.3, 7.3, 7.3] / [27.3, 26.3, 27.3]	0.7238	204.57	–
200 kN	BO	[22.0, 31.4, 22.0] / [59.5, 59.5, 59.5]	0.7264	270.72	0.36
	GA	[19.3, 18.3, 19.3] / [29.3, 29.3, 29.3]	0.7238	248.47	–
500 kN	BO	[59.5, 40.7, 50.1] / [50.1, 40.7, 50.1] / [59.5, 59.5, 68.8]	1.0905	535.49	0.40
	GA	[35.3, 35.3, 34.3] / [25.3, 22.3, 24.3] / [64.3, 67.3, 64.3]	1.0861	581.13	–
1000 kN	BO	[3.3, 40.7, 78.2] / [40.7, 68.8, 31.2] / [22.0, 59.5, 59.5]	1.4102	1038.45	15.30
	GA	[18.3, 37.3, 33.3] / [34.3, 75.3, 64.3] / [9.3, 39.3, 35.3]	1.2229	1051.54	–

[*] Diff. (%) $= ((BO - GA)/GA) \times 100$

computational cost. In the 50 kN case, GA produced a slightly lighter design (0.3782 kg) compared to BO (0.4135 kg), with nearly identical buckling capacities (52.80 kN vs. 53.00 kN). For the 100 kN and 200 kN cases, the weight differences remained below 0.2 %, and buckling loads are also within a comparable range.

The most notable discrepancy arises at 1000 kN, where BO produced a design approximately 15.3 % heavier (1.4102 kg vs. 1.2229 kg), with a slightly lower buckling load (1038.45 kN vs. 1051.54 kN). This suggests that, in this instance, GA identified a more optimal configuration. However, during the optimisation process, BO discovered some feasible solutions around 1.33 kg that satisfied the buckling constraint (i.e PCR > 95 % of the design load), indicating that further tuning of acquisition parameters, such as the exploitation–exploration trade-off or adaptive sampling strategies, could enhance performance.

From a computational standpoint, BO demonstrates clear efficiency. While GA required 2401 function evaluations and approximately 8 h to converge, BO achieved its optimal solutions in just 800 evaluations and 2.5 h. This threefold reduction in computational effort, especially evident in the 50 kN case, underscores the suitability of BO for simulation-intensive optimisation of advanced composites.

Although the current BO implementation delivers excellent performance in most scenarios, convergence in complex or multimodal spaces could be further improved through modifications such as incorporating rotational acquisition functions, adjusting Pareto Front weighting strategies, or integrating a trust-region approach. Nonetheless, its data efficiency and consistent results make BO a compelling framework for design optimisation problems involving high-fidelity simulations.

5.2. Fibre paths and thickness distribution

Figs. 9–13 illustrate the optimised fibre orientations and thickness distributions obtained via BO (left) and GA (right) for increasing load levels. The visualisations highlight differences in design strategies between the two approaches as structural complexity increases.

For the 50 kN case (Fig. 9), both BO and GA generate relatively simple, symmetric fibre paths using a single layer. The helical winding

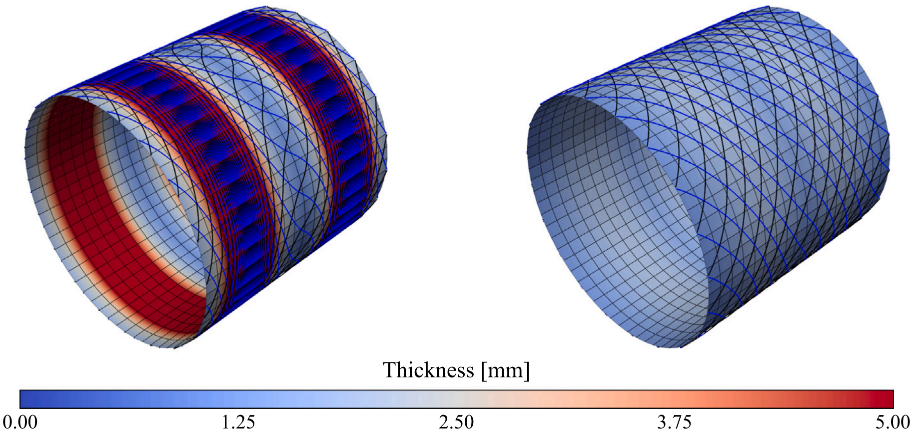

Fig. 9. Bayesian (left) and GA (right) results, 50 kN, single layer.

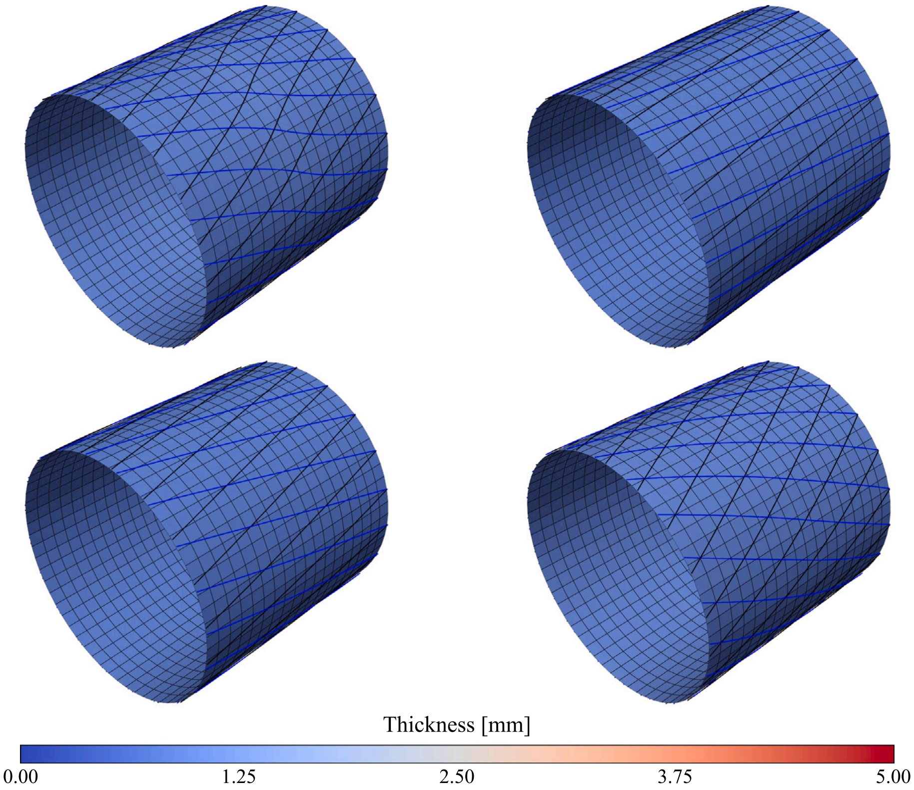

Fig. 10. Bayesian (left) and GA (right) results, 100 kN, single layer, layers 1 (top) and 2 (bottom).

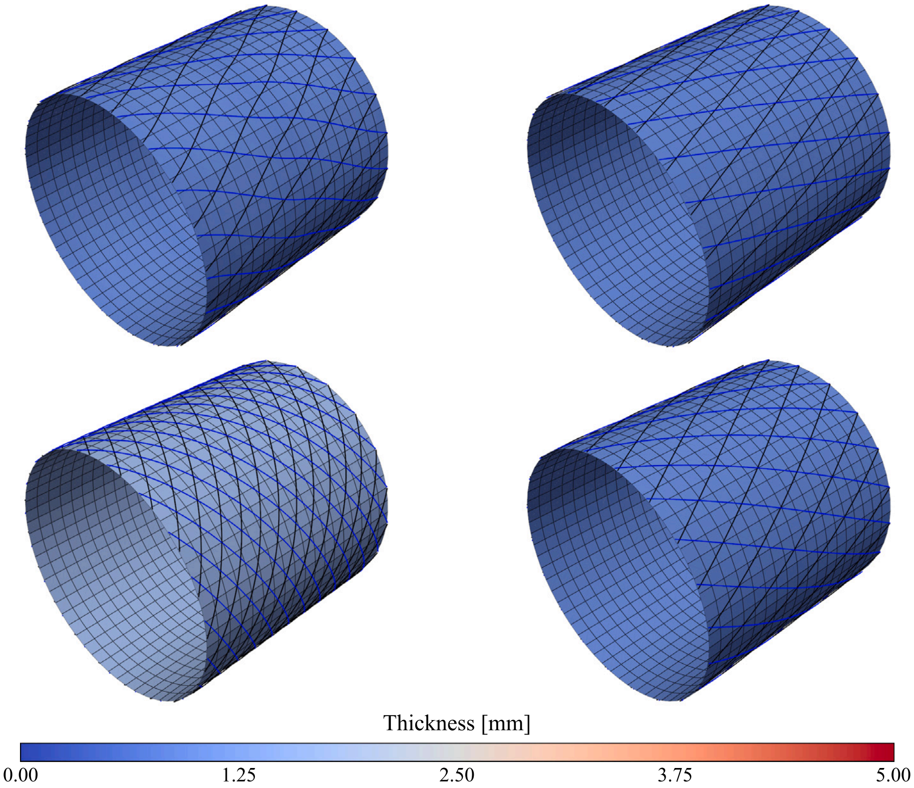

Fig. 11. Bayesian (left) and GA (right) results, 200 kN, layers 1 (top) and 2 (bottom).

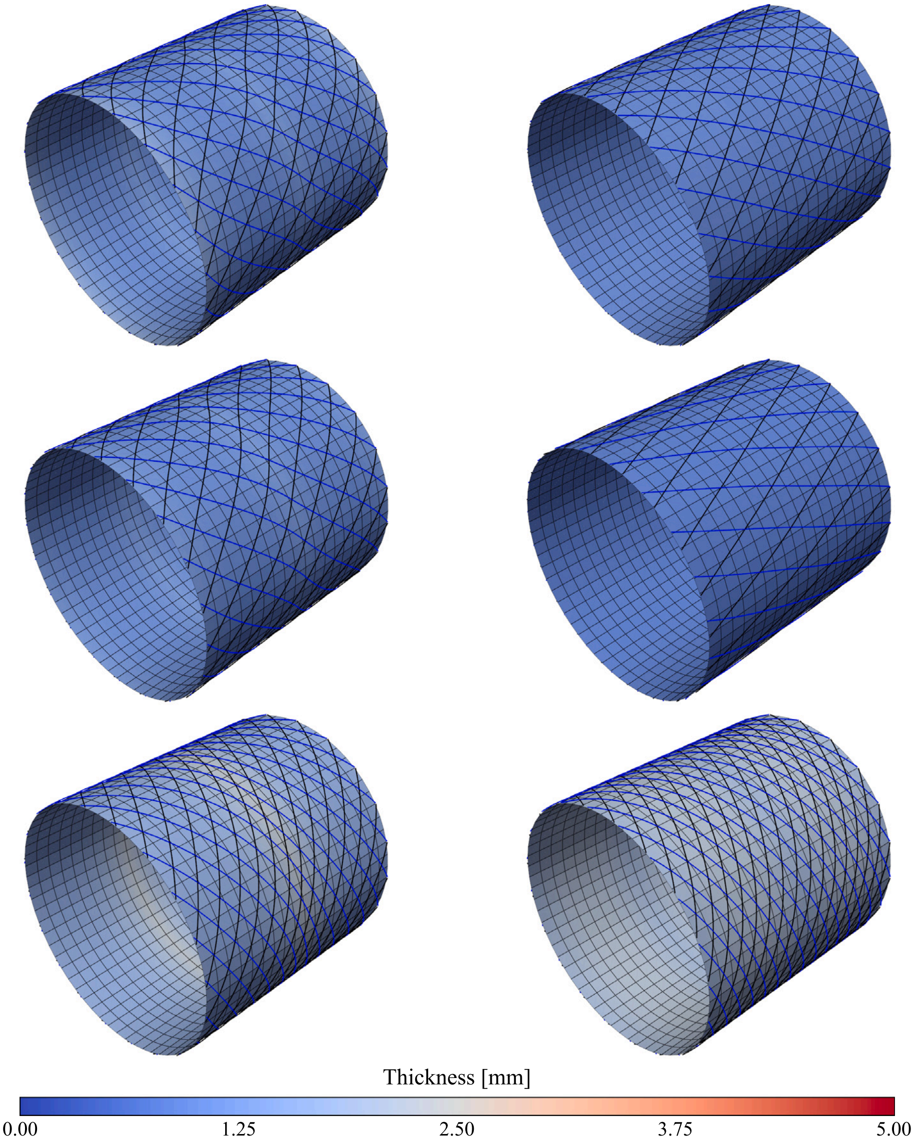

Fig. 12. Bayesian (left) and GA (right) results, 500 kN, layers 1 (top) to 3 (bottom).

patterns are nearly uniform, but the BO design introduces subtle spatial variation in thickness, likely responding to local stress concentrations. This is visible in the smooth gradient of red and blue tones, in contrast to the more uniform thickness of the GA solution, which maintains a consistent but less tailored configuration.

In the 100 kN case (Fig. 10), two layers are activated, and the benefits of BO's local adaptation begin to emerge. BO varies fibre angles across layers, particularly in the second layer, where slight asymmetries and refined orientation shifts are observed. Thickness is distributed more progressively, indicating material efficiency and structural intent. GA, on the other hand, employs near-identical helical paths across both layers with minimal variation, resulting in a less adaptive design that prioritises symmetry over local stress accommodation.

At 200 kN (Fig. 11), the difference in design strategy becomes more evident. BO further tunes the fibre orientations across both layers, introducing diagonal winding angles and controlled thickness gradients that reflect anticipated buckling-critical regions. The thickness distribution remains smooth and functionally graded. In contrast, the GA result exhibits more abrupt transitions and localised thickness increases, particularly in the second layer, possibly reflecting compensatory material placement due to a lack of local orientation refinement.

In the 500 kN case (Fig. 12), with three active layers, BO leverages its surrogate-assisted search to generate distinct fibre orientations in each layer, along with non-uniform thickness modulations that follow likely stress paths. This gradient-based reinforcement offers a more sophisticated material placement strategy. GA, by comparison, maintains globally repetitive fibre patterns across all layers with fewer local deviations. While both designs meet performance requirements, the BO solution exhibits superior material targeting and structural customisation.

The distinction becomes most pronounced at 1000 kN (Fig. 13). The solution provided by BO exhibits highly non-uniform fibre orientations and a smoothly varying thickness profile, particularly evident in the intermediate and bottom layers. The design suggests clear spatial alignment with load distribution and likely buckling deformation modes. Meanwhile, GA again employs more repetitive fibre architecture with noticeably thicker regions concentrated in the middle layer, indicating overcompensation rather than optimisation. The BO result shows better alignment of material usage with structural demand, reinforcing critical areas while avoiding excess.

Overall, the progression of results highlights the strength of BO in producing structurally efficient and spatially adaptive designs across

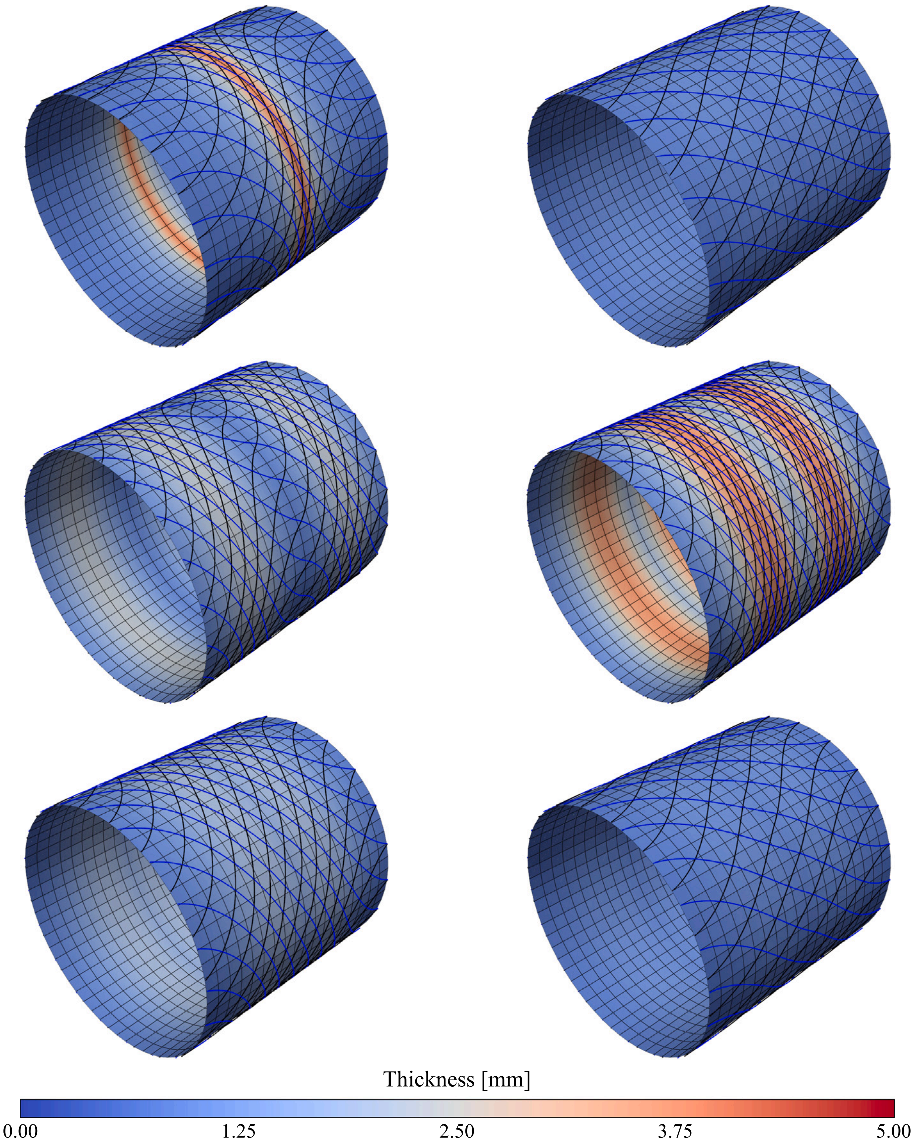

Fig. 13. Bayesian (left) and GA (right) results, 1000 kN, layers 1 (top) to 3 (bottom).

a wide range of loading conditions. BO systematically varies fibre orientations and thickness distributions to respond to local demands, enabling lightweight yet robust structures. GA, although capable of generating feasible solutions, tends to adopt more conservative and uniform strategies, often relying on increased thickness rather than refined orientation control. This comparison underscores the capability of BO to exploit the design space more intelligently, particularly for complex composite structures governed by buckling performance.

5.2.1. *Manufacturing considerations*

While the optimisation results demonstrate the structural efficiency of BO-generated designs, their manufacturability must be critically assessed, particularly with respect to automated composite fabrication processes such as filament winding (FW) and automated fibre placement (AFP). The following considerations are essential when evaluating the feasibility of implementing the optimised fibre paths.

Fibre path curvature and continuity. FW typically supports symmetric helical or hoop winding with continuous and relatively simple fibre trajectories [50]. Sharp changes in fibre angle, abrupt shifts between regions, or asymmetries across the shell, characteristics frequently seen in BO, optimised designs under higher load cases—are difficult to realise using standard FW techniques. Such features may require multi-pass programming, segmental winding, or complex mandrel adjustments.

In contrast, AFP technologies offer more flexibility. Fibre paths can be steered along variable angles, enabling the implementation of spatially adaptive BO strategies. However, AFP is not without constraints; tight curvatures and abrupt orientation changes can introduce gaps, overlaps, or local wrinkling, especially when the curvature radius falls below allowable thresholds for a given tow width [51].

These observations suggest that the more complex and refined fibre paths generated by BO are naturally better suited to AFP, whereas the regular, symmetric helical patterns of GA are more compatible with FW.

Layer transition and coverage. The BO framework frequently introduces inter-layer variation, where each layer exhibits distinct fibre angles and thickness distributions tailored to the local stress field. While advantageous structurally, these features introduce manufacturing complexity.

In FW, changing fibre orientations between layers is possible but requires reprogramming of the machine and may negatively impact layer compaction or resin flow consistency [3]. In AFP, such inter-layer variations are more manageable through selective ply stacking, but they

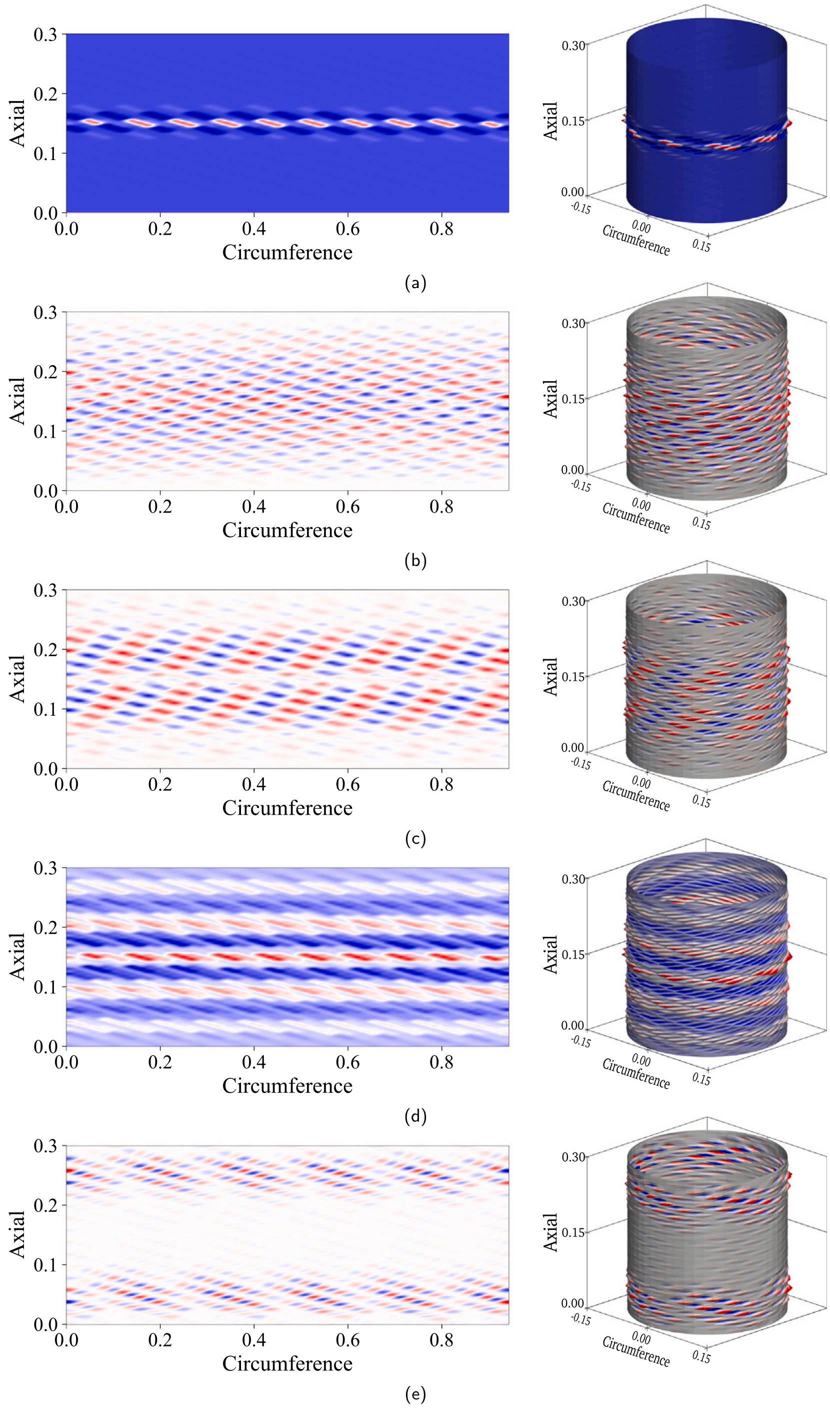

Fig. 14. Buckling modes for the optimum cylinders using BO for the design loads of (a) 50 kN, (b) 100 kN, (c) 200 kN, (d) 500 kN, and (e) 1000 kN. The buckling shapes are shown in 2D (left) and in 3D (right).

increase toolpath planning complexity and may still result in practical issues like tow gaps or overlaps during add/drop transitions.

Thickness gradient implementation. Thickness modulation is another critical design feature that distinguishes BO from GA results. The ability to vary thickness locally and continuously, as seen in BO outcomes, is challenging to implement using FW, which generally builds uniform layers. Achieving such gradients would require careful control of overlapping passes or local winding speeds, which are limited in practice.

AFP, however, is more capable of achieving thickness variation through tow-by-tow control, allowing selective material build-up or reduction [4]. GA solutions, with their inherently more uniform and symmetric thickness profiles, are simpler to manufacture via FW and thus offer a practical baseline for production readiness.

Overall, the non-uniform thickness and variable fibre orientations observed in BO designs align more naturally with AFP capabilities. However, when FW is the target manufacturing process, incorporating these advanced designs may require additional filtering, adaptation, or manufacturing-aware constraints within the optimisation loop.

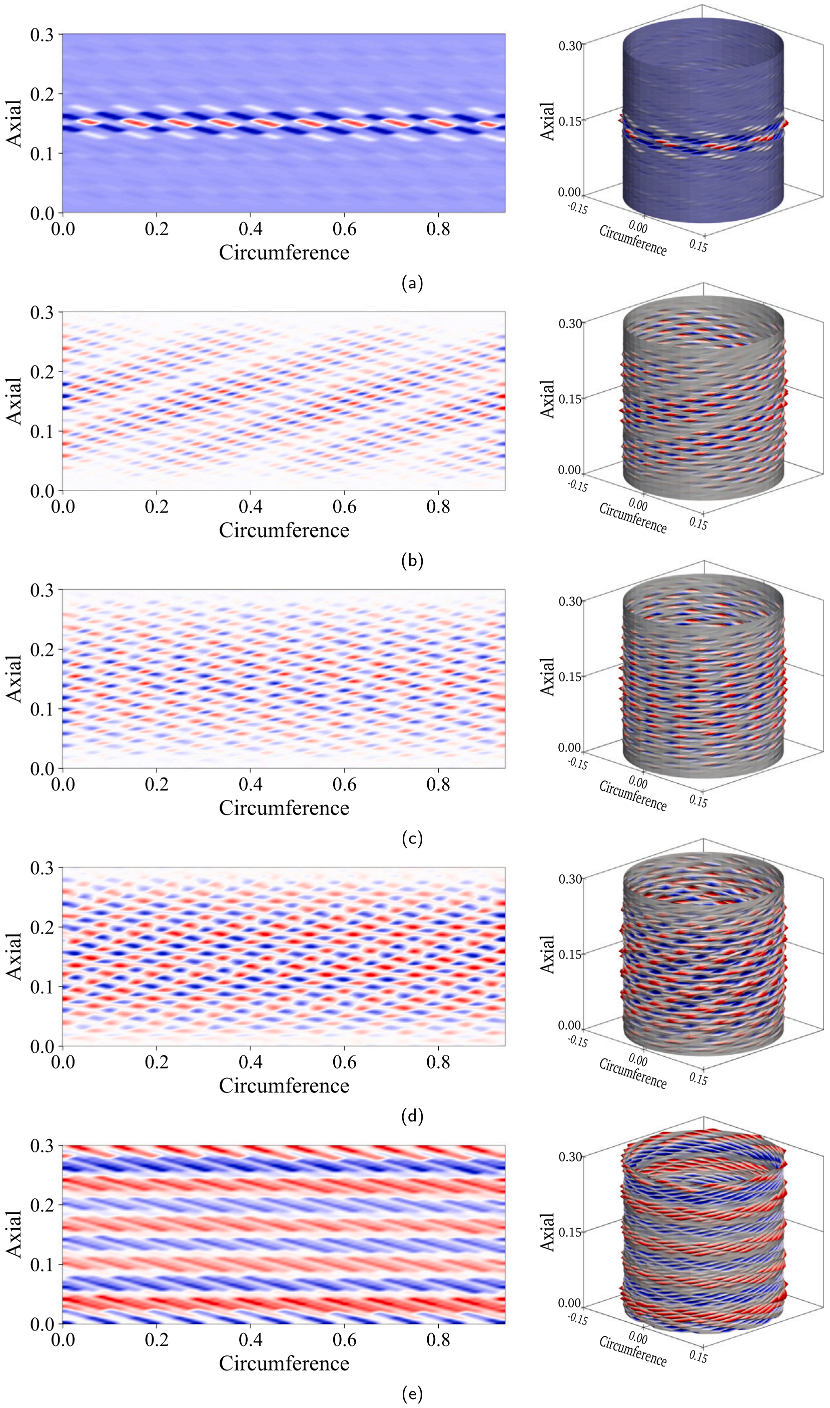

Fig. 15. Buckling modes for the optimum cylinders using GA for the design loads of (a) 50 kN, (b) 100 kN, (c) 200 kN, (d) 500 kN, and (e) 1000 kN. The buckling shapes are shown in 2D (left) and in 3D (right).

It is important to note that the manufacturing constraints embedded in the optimisation framework are derived based on filament winding limitations [31,52]. These included bounds on allowable fibre angles, continuity of fibre paths, and restrictions on abrupt inter-layer deviations. As such, the BO-generated designs remain within the feasible envelope of FW capabilities, despite appearing more locally refined than their GA counterparts.

However, some of the BO solutions, particularly under high load cases, approach the practical limits of FW, such as the minimum winding angle and layer asymmetry tolerances. While still manufacturable, these configurations may demand more advanced winding strategies, such as segmented or adaptive winding paths. Alternatively, their implementation may be facilitated through hybrid methods or AFP, which offer additional flexibility in fibre steering and layer-by-layer tailoring.

Therefore, while the current results demonstrate the ability of BO to operate within FW-compatible constraints, further enhancement of manufacturability could be achieved by explicitly incorporating process-aware penalisation or reward mechanisms into the optimisation objectives, particularly if broader production strategies (e.g., AFP or hybrid deposition) are to be considered.

5.3. Buckling mode shape analysis

Figs. 14 and 15 show the first-mode buckling shapes of the optimised variable-stiffness composite cylinders obtained using BO and GA, respectively, for design loads ranging from 50 kN to 1000 kN. These mode shapes reveal important characteristics regarding the distribution of stiffness and the structural response of the optimised configurations.

Figs. 14 and 15 show the first buckling modes for the optimised cylinders obtained using BO and GA, respectively, across all load levels. The results are presented in both 2D (axial vs. circumferential coordinate) and 3D views, allowing for a detailed comparison of modal characteristics.

At lower load levels (50 kN and 100 kN), both BO and GA yield diagonal buckling modes with evident torsional features. However, the BO solutions exhibit more localised and sharper lobe formations (Fig. 14(a)–(b)), suggesting enhanced circumferential stiffness modulation. In contrast, the GA designs (Fig. 15(a)–(b)) show more uniform and broader deformation bands, with lower spatial resolution in the lobes. These differences indicate that BO achieves finer tailoring of stiffness distribution, promoting the emergence of higher-order or mixed-mode deformations.

At intermediate load levels (200 kN and 500 kN), the distinction becomes more pronounced. The BO buckling modes Fig. 14(c)–(d) reveal a superposition of axial and circumferential waves, forming oblique and spatially heterogeneous patterns. The lobes are non-uniform yet structured, pointing to locally responsive reinforcement strategies shaped by the optimisation process. Conversely, the GA modes Fig. 15(c)–(d) display increased axial wave dominance, particularly at 500 kN, with relatively symmetric, repetitive lobes and reduced circumferential richness. This suggests that GA solutions prioritise axial stiffness, possibly sacrificing modal adaptability.

At the highest load case (1000 kN), both optimisation strategies converge to globally stable and symmetric buckling patterns. Nevertheless, the BO result Fig. 14(e) continues to exhibit finer-scale circumferential features and a slightly higher apparent mode order, reflecting sustained tailoring across layers. In contrast, the GA solution Fig. 15(e) appears more globally uniform, with longer axial half-wavelengths and reduced modal richness. This may indicate a stiffer but less responsive design, relying more on global thickness accumulation than localised reinforcement.

Overall, the updated buckling mode analysis reinforces the structural benefits of BO. By enabling adaptive local variations in stiffness and thickness, BO promotes the activation of complex, multi-directional buckling shapes that better align with the structural demands. While GA yields feasible and robust configurations, its tendency toward symmetric and conservative patterns may limit the activation of beneficial higher-order modes under complex loading conditions.

6. Conclusion

We proposed a Bayesian Optimisation (BO) framework for the mass minimisation of variable-stiffness composite cylinders subjected to a wide range of buckling-driven load cases (50 kN to 1000 kN). The formulation incorporated manufacturing constraints compatible with filament winding, and structural evaluations were performed using a computationally efficient single-curvature Bögner–Fox–Schmit–Castro (BFSC) finite element formulation tailored for thin-walled shells.

The BO framework demonstrated two key advances: by optimising local fibre orientations (through three control-point angles per layer), and by achieving convergence to near-optimal solutions with significantly fewer function evaluations than a baseline Genetic Algorithm (GA). Across most design loads, BO produced designs with comparable or improved buckling resistance while maintaining similar or slightly higher mass. Notably, for the 1000 kN case, GA resulted in a lighter design with marginally better structural performance, although BO solutions around 1.326 were also identified, suggesting scope for further tuning of hyperparameters.

The analysis of fibre paths and thickness distributions revealed that BO generates spatially adaptive, non-uniform designs with finer control over material allocation. These designs showed better alignment between stiffness variation and stress demand, particularly in higher load cases. While this enhances structural responsiveness, it also introduces complexity from a manufacturing standpoint. GA, in contrast, produced more symmetric and uniform configurations, which are inherently more compatible with filament winding.

Buckling mode analyses confirmed the benefits of BO in capturing rich, mixed-mode deformations, with higher circumferential lobe density and increased mode order compared to GA. These features reflect the ability of BO to locally tailor stiffness and delay global instabilities. In contrast, GA solutions tended towards lower-order axial-dominated buckling modes with less structural adaptability.

Importantly, BO achieved high-quality solutions with far fewer simulations—typically requiring only 800 evaluations and 2.5 h of computational time, compared to over 2400 evaluations and 8 h for GA. This data efficiency underscores the value of surrogate-assisted optimisation in simulation-driven design workflows.

In summary, the proposed BO framework offers a powerful and scalable approach for the design of lightweight composite structures under buckling constraints. Its ability to generate structurally adaptive solutions with significant computational savings makes it particularly well suited for complex, high-fidelity applications. Future work will explore the integration of experimentally measured geometric imperfections, enabling the optimisation of imperfect structures.

CRediT authorship contribution statement

José Humberto S. Almeida: Formal analysis, Investigation, Software, Validation, Writing – original draft. **Aravind Ashok:** Software, Methodology, Formal analysis. **Muhammad Uzair:** Visualization, Software. **Saullo G.P. Castro:** Conceptualization, Methodology, Supervision, Resources, Software, Writing – original draft.

Declaration of competing interest

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

Acknowledgements

HA has received funding from the Research Council of Finland under grant number 364167.

Data availability

All models used in the study are published on Zenodo, as indicated in the revised manuscript.

References

[1] Ma J, Zhang S, Wriggers P, Gao W, De Lorenzis L. Stochastic homogenized effective properties of three-dimensional composite material with full randomness and correlation in the microstructure. Comput Struct 2014;144:62–74. https://doi.org/10.1016/j.compstruc.2014.06.003

[2] Almeida JHS Jr., Bittrich L, Jansen E, Tita V, Spickenheuer A. Buckling optimization of composite cylinders for axial compression: a design methodology considering a variable-axial fiber layout. Compos Struct 2019;222:110928. https://doi.org/10.1016/j.compstruct.2019.110928

[3] Almeida JHS, St-Pierre L, Wang Z, Ribeiro ML, Tita V, Amico SC, et al. Design, modeling, optimization, manufacturing and testing of variable-angle filament-wound cylinders. Compos Part B Eng 2021;225:109224. https://doi.org/10.1016/j.compositesb.2021.109224

[4] Jing Z, Wang S, Duan L, Li B, Fan X. Buckling analysis and optimization of variable angle tow composite plates via ritz method and variable stiffness optimization. Comput Struct 2023;289:107156. https://doi.org/10.1016/j.compstruc.2023.107156

[5] Adriana W. Blom PBS, Gürdal Z. Optimization of a composite cylinder under bending by tailoring stiffness properties in circumferential direction. (Mar 2010). https://doi.org/10.1016/j.compositesb.2009.10.004

[6] Kim BC, Weaver PM, Potter K. Manufacturing characteristics of the continuous tow shearing method for manufacturing of variable angle tow composites. Compos Part

A Appl Sci Manuf 2014;61:141–51. https://doi.org/10.1016/j.compositesa.2014.02.019

[7] Almeida JHS Jr., Bittrich L, Nomura T, Spickenheuer A. Cross-section optimization of topologically-optimized variable-axial anisotropic composite structures. Compos Struct 2019 10;225:111150. https://doi.org/10.1016/j.compstruct.2019.111150

[8] Peeters DMJ, Lozano GG, Abdalla MM. Effect of steering limit constraints on the performance of variable stiffness laminates. Comput Struct 2018;196:94–111. https://doi.org/10.1016/j.compstruc.2017.11.002

[9] Almeida JHS Jr., Christoff BG, Tita V, St-Pierre L. A concurrent fibre orientation and topology optimisation framework for 3D-printed fibre-reinforced composites. Compos Sci Technol 2023;232:109872. https://doi.org/10.1016/j.compscitech.2022.109872

[10] Azevedo CB, Almeida JHS Jr, Flores HF, Eggers F, Amico SC. Influence of mosaic pattern on hygrothermally-aged filament wound composite cylinders under axial compression. J Compos Mater 2020;54(19):2651–59. https://doi.org/10.1177/0021998319899144

[11] Valvano S, Alaimo A, Orlando C. Higher-order models for the passive damping analysis of variable-angle-tow composite plates. Comput Struct 2023;280:106992. https://doi.org/10.1016/j.compstruc.2023.106992

[12] Albazzan M, Harik R, Tatting B, Gurdal Z, Blom-Schieber A, Rassaian M, et al. Optimization of cylinders with holes under bending using nonconventional laminates, In: 2018 AIAA/ASCE/AHS/ASC structures, structural dynamics, and materials conference; 2018. https://doi.org/10.2514/6.2018-1377.

[13] Fina M, Bisagni C. Buckling design optimization of tow-steered composite panels and cylindrical shells considering aleatory and epistemic uncertainties. Comput Mech 2025;23:1–34. https://doi.org/10.1007/s00466-024-02589-8

[14] Pathan MV, Patsias S, Tagarielli VL. A real-coded genetic algorithm for optimizing the damping response of composite laminates. Comput Struct 2018;198:51–60. https://doi.org/10.1016/j.compstruc.2018.01.005

[15] Almeida JHS Jr., Lisbôa TV, Spickenheuer A, St-Pierre L. A sequential finite element model updating routine to identify creep parameters for filament wound composite cylinders in aggressive environments. Comput Struct 2023;276:106939. https://doi.org/10.1016/j.compstruc.2022.106939

[16] Pitton SF, Ricci S, Bisagni C. Buckling optimization of variable stiffness cylindrical shells through artificial intelligence techniques. Compos Struct 2019;230:111513. https://doi.org/10.1016/j.compstruct.2019.111513

[17] Li Q, Ni P, Du X, Han Q, Xu K, Bai Y. Bayesian updating using accelerated hamiltonian Monte Carlo with gradient-enhanced kriging model. Comput Struct 2025;307:107598. https://doi.org/10.1016/j.compstruc.2024.107598

[18] Huang B, Sun M, Chen H, Wu Z. Stochastic static finite element model updating using the bayesian method integrating homotopy surrogate model. Comput Struct 2025;315:107769. https://doi.org/10.1016/j.compstruc.2025.107769

[19] Di Fiore F, Mainini L. Physics-aware multifidelity bayesian optimization: a generalized formulation. Comput Struct 2024;296:107302. https://doi.org/10.1016/j.compstruc.2024.107302

[20] Rao SS. Recent applications of evolutionary algorithms in mechanical engineering design. Eng Appl Artif Intel 2009;22(6):972–80. https://doi.org/10.1016/j.engappai.2008.09.002

[21] Dang C, Wei P, Faes MGR, Beer M. Bayesian probabilistic propagation of hybrid uncertainties: estimation of response expectation function, its variable importance and bounds. Comput Struct 2022;270:106860. https://doi.org/10.1016/j.compstruc.2022.106860

[22] Yamaguchi K, Phenisee SE, Chen Z, Salviato M, Yang J. Ply-drop design of non-conventional laminated composites using bayesian optimization. Compos Part A Appl Sci Manuf 2020;139:106136. https://doi.org/10.1016/j.compositesa.2020.106136

[23] Chuaqui TRC, Rhead AT, Butler R, Scarth C. A data-driven bayesian optimisation framework for the design and stacking sequence selection of increased notched strength laminates. Compos Part B Eng 2021;226:109347. https://doi.org/10.1016/j.compositesb.2021.109347

[24] Mathern A, Steinholtz OS, Sjöberg A, Önnheim M, Ek K, Rempling R, et al. Multi-objective constrained bayesian optimization for structural design. Struct Multidiscip Optim 2021;63(2):689–701. https://doi.org/10.1007/s00158-020-02720-2

[25] Raiaan MAK, Sakib S, Fahad NM, Al Mamun A, Rahman MA, Shatabda S, et al. A systematic review of hyperparameter optimization techniques in convolutional neural networks. Decis Analytics J 2024;11:100470. https://doi.org/10.1016/j.dajour.2024.100470

[26] Bessa MA, Pellegrino S. Design of ultra-thin shell structures in the stochastic post-buckling range using bayesian machine learning and optimization. Int J Solids Struct 2018;139-140:174–88. https://doi.org/10.1016/j.ijsolstr.2018.01.035

[27] Bessa MA, Glowacki P, Houlder M. Bayesian machine learning in metamaterial design: fragile becomes supercompressible. Adv Mater 2019;31(48):1904845. https://doi.org/10.1002/adma.201904845

[28] Wang Z, Almeida JHS Jr., St-Pierre L, Wang Z, Castro SGP. Reliability-based buckling optimization with an accelerated kriging metamodel for filament-wound variable angle tow composite cylinders. Compos Struct 2020;254:112821. https://doi.org/10.1016/j.compstruct.2020.112821

[29] Castro SGP, Donadon MV, Guimarães TAM. ES-PIM applied to buckling of variable angle tow laminates. Compos Struct 2019;209(June 2018):67–78. https://doi.org/10.1016/j.compstruct.2018.10.058

[30] Groh RM, Weaver P. Mass optimisation of variable angle tow, variable thickness panels with static failure and buckling constraints, In: 56th AIAA/ASCE/AHS/ASC structures, structural dynamics, and materials conference; 2015. https://doi.org/10.2514/6.2015-0452

[31] Wang Z, Almeida Jr JHS, Ashok A, Wang Z, Castro SGP. Lightweight design of variable-angle filament-wound cylinders combining kriging-based metamodels with particle swarm optimization. Struct Multidiscip Optim 2022;65(5). https://doi.org/10.1007/s00158-022-03227-8

[32] Aravind A, Uzair M, Almeida JHS Jr, Castro SGP. Bayesian optimization for VAFW CyLinders (bovafwcyl). 2025. https://doi.org/10.5281/zenodo.15488104

[33] Jiang P, Zhou Q, Shao X. Surrogate-model-based design and optimization. Springer Nature; 2020. https://doi.org/10.1007/978-981-15-0731-1_7

[34] Kushner H. A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. J Basic Eng 1964;86(1):97–106.

[35] Mockus J, Tiesis V, Zilinskas A. The application of bayesian methods for seeking the extremum. Towards Glob Optim 1978;2(117–129):2.

[36] Jones WDR, Schonlau M. Efficient global optimization of expensive black-box functions. J Global Optim 1998;13(4):455–92.

[37] Kuhn J, Spitz J, Sonnweber-Ribic P, Schneider M, Böhlke T. Identifying material parameters in crystal plasticity by bayesian optimization. Optim Eng 2021;23(3):1489–523. https://doi.org/10.1007/s11081-021-09663-7

[38] Xu H, Nakayama R, Kimura T, Shimizu R, Ando Y, Kobayashi S, et al. Tuning bayesian optimization for materials synthesis: simulating two- and three-dimensional cases. Sci Technol Adv Mater: Methods 2023 May;3(1). https://doi.org/10.1080/27660400.2023.2210251

[39] Auer P. Using confidence bounds for exploitation-exploration trade-offs. J Mach Learn Res 2003;3:397–422.

[40] Archetti F, Candelieri A. Bayesian optimization and data Science. Springer International Publishing; 2019. https://doi.org/10.1007/978-3-030-24494-1

[41] Brochu E, Hoffman MW, de Freitas N. Portfolio allocation for bayesian optimization. arXiv:arXiv:1009.5419 [Preprint]. 2011.

[42] Buitinck L, Louppe G, Blondel M, Pedregosa F, Mueller A, Grisel O, et al. API design for machine learning software: experiences from the scikit-learn project. In: ECML PKDD workshop: languages for data mining and machine learning; 2013. p. 108–22.

[43] Head T, Kumar M, Nahrstaedt H, Louppe G, Shcherbatyi I. Scikit-optimize/scikit-optimize. (Oct 2021). https://doi.org/10.5281/zenodo.5565057

[44] Castro SGP, Jansen EL. Displacement-based formulation of koiter's method: application to multi-modal post-buckling finite element analysis of plates. Thin-Walled Struct 2021;159:107217. https://doi.org/10.1016/j.tws.2020.107217

[45] Castro SG, Jansen EL. Displacement-based multi-modal formulation of koiter's method applied to cylindrical shells. In: AIAA scitech 2022 forum; 2022. https://doi.org/10.2514/6.2022-0256

[46] Ferreira LJF, Bittencourt ML. Hierarchical high-order conforming c1 bases for quadrangular and triangular finite elements. Int J Numer Methods Eng 2017;109(7):936–64. https://doi.org/10.1002/nme.5308

[47] Tsunematsu DQ, Donadon MV. Aeroelastic behavior of composite panels undergoing progressive damage. Compos Struct 2019;210:458–72. https://doi.org/10.1016/J.COMPSTRUCT.2018.11.065

[48] Castro SGP, Mittelstedt C, Monteiro FAC, Arbelo MA, Ziegmann G, Degenhardt R. Linear buckling predictions of unstiffened laminated composite cylinders and cones under various loading and boundary conditions using semi-analytical models. Compos Struct 2014;118:303–15. https://doi.org/10.1016/j.compstruct.2014.07.037

[49] Gray JS, Hwang JT, Martins JRRA, Moore KT, Naylor BA. OpenMDAO: an open-source framework for multidisciplinary design, analysis, and optimization. Struct Multidiscip Optim 2019;59(4):1075–104. https://doi.org/10.1007/s00158-019-02211-z

[50] Antunes MB, Almeida JHS Jr, Amico SC. Curing and seawater aging effects on Mechanical and physical properties of glass/epoxy filament wound cylinders. Compos Commun 2020;22:100517. https://doi.org/10.1016/j.coco.2020.100517

[51] Hong Z, Peeters D, Turteltaub S. An enhanced curvature-constrained design method for manufacturable variable stiffness composite laminates. Comput Struct 2020;238:106284. https://doi.org/10.1016/j.compstruc.2020.106284

[52] Castro SGP, Almeida JHS Jr., St-Pierre L, Wang Z. Measuring geometric imperfections of variable–angle filament–wound cylinders with a simple digital image correlation setup. Compos Struct 2021;276:114497. https://doi.org/10.1016/j.compstruct.2021.114497