

Document Version

Final published version

Licence

CC BY

Citation (APA)

Uzair, M., Castro, S. G. P., & Almeida, J. H. S. (2026). Multi-objective optimisation of variable-stiffness composite cylinders with geometric imperfections: minimising mass while maximising buckling capacity and knockdown factor. *Computers and Structures*, 323, Article 108141. <https://doi.org/10.1016/j.compstruc.2026.108141>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.

Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Multi-objective optimisation of variable-stiffness composite cylinders with geometric imperfections: minimising mass while maximising buckling capacity and knockdown factor

Muhammad Uzair^a, Saullo G.P. Castro^{b,*}, José Humberto S. Almeida Jr.^{a,*}


^a Department of Mechanical Engineering, LUT University, Lappeenranta, Finland

^b Faculty of Aerospace Engineering, Delft University of Technology, Delft, the Netherlands

HIGHLIGHTS

- Uncertainty-quantified optimisation of variable-stiffness composite cylinders.
- Principal component analysis with Latin hypercube sampling for imperfections.
- Gaussian process regression surrogates predict mass and collapse behaviour.
- Multi-objective design balances weight, buckling strength, and robustness.
- Bayesian and genetic optimisers show complementary efficiency and accuracy.

GRAPHICAL ABSTRACT



ARTICLE INFO

Dataset link: [Imperfection-tolerant multi-objective optimisation framework for variable-stiffness composite structures \(Original data\)](#)

Keywords:
Variable-stiffness
Geometric imperfections
Bayesian optimisation
Surrogate modelling
Buckling

ABSTRACT

This study presents an imperfection-tolerant, surrogate-assisted framework for the multi-objective optimisation of variable-stiffness (VS) composite cylinders that explicitly incorporates experimentally measured geometric imperfections. Principal component analysis (PCA) is applied to extract dominant imperfection modes from experimental data, and Latin hypercube sampling (LHS) is used to generate statistically consistent synthetic fields, which are subsequently mapped onto nonlinear finite element (FE) models. Linear buckling and geometrically nonlinear collapse analyses are performed under axial compression to determine the ideal and actual load-carrying capacities, from which the knockdown factor (KDF), quantifying imperfection sensitivity, is derived. Gaussian Process Regression (GPR) surrogates are trained to predict the mass and collapse loads of perfect and imperfect geometries with high cross-validated accuracy, while KDF is computed as their ratio. The framework enables simultaneous optimisation of three objectives: mass minimisation, collapse-load maximisation, and KDF maximisation by using Bayesian Optimisation (BO) and the Non-dominated Sorting Genetic Algorithm II (NSGA-II) independently. Results demonstrate that integrating experimentally informed imperfections with surrogate-based optimisation captures the key physical trends governing buckling and imperfection sensitivity, while achieving substantial computational savings relative to direct nonlinear analyses, and that both optimisers yield consistent Pareto fronts featuring smooth, manufacturable fibre trajectories that balance lightweight efficiency, strength, and robustness.

* Corresponding authors.

Email addresses: muhammad.uzair@lut.fi (M. Uzair), S.G.P.Castro@tudelft.nl (S.G.P. Castro), humberto.almeida@lut.fi (J.H.S. Almeida Jr.).

1. Introduction

Thin-walled composite cylinders are fundamental load-bearing components in aerospace, energy, and transportation applications due to their high specific stiffness, strength-to-weight ratios, and ability to be tailored through fibre architecture [1,2]. They are widely employed in aircraft fuselage barrels, launch vehicle cryogenic tanks and spacecraft pressure vessels, where reducing structural mass directly improves payload efficiency and operational cost. In these applications, axial buckling is often the governing failure mode. The challenge is that buckling strength is highly sensitive to imperfections: even small deviations such as ovality, fibre misalignments, or surface waviness can reduce the collapse load by 20 – 40% relative to idealised predictions [3–5].

Historically, this imperfection sensitivity has been addressed through empirical knockdown factors (KDFs), most notably those formalised in NASA SP-8007 [6]. While conservative and safe, such margins lead to significant weight penalties when applied to advanced composite shells. More recent guidelines, such as NASA/SP-8007-2020/REV-2 [7], and European initiatives, such as DESICOS [8], have shifted toward physics-based knockdown approaches, where geometrically nonlinear finite element (FE) analyses incorporating measured imperfections are used to derive less conservative yet reliable design margins. These developments emphasise the importance of accurately measuring and modelling imperfections for modern composite structures. However, no prior study has combined measured imperfection data with multi-objective optimisation in a reliability context, and this is the gap we address.

Castro et al. [9] implemented a simple digital image correlation (DIC) methodology to measure the full-field geometric imperfections of filament-wound composite cylinders. For the first time, such a measurement was conducted using only a single stereoscopic camera pair, with each acquisition capturing approximately 60° of the circumferential perimeter. This work provided the first experimental evidence of realistic imperfection fields in filament-wound cylinders. Building directly on this foundation, Almeida Jr. et al. [10] integrated these measured imperfection fields into nonlinear FE models, including geometric and material nonlinearities, enabling the numerical prediction of the collapse loads, validated against axial compression tests monitored by DIC. This combined metrology–simulation–experiment workflow confirmed that realistic imperfections are indispensable for accurate buckling assessment.

In parallel, the introduction of variable-stiffness (VS) laminates has fundamentally expanded the design space of composite shells by enabling spatially varying fibre orientations, redistributing stiffness in ways not possible with conventional constant-stiffness designs. Enabled by advanced manufacturing processes such as automated fibre placement (AFP) [11–13], rapid tow shearing (RTS) [14,15], tailored fibre placement (TFP) [16,17], and variable-angle filament-winding (VAFW) [10,18,19], VS laminates allow spatially varying fibre orientation that redistributes stiffness beyond the limits of conventional constant-stiffness (CS) designs [16,20]. VAFW is particularly attractive for cylindrical shells because it achieves segmented fibre steering using standard robotic winding equipment, offering a practical route to aerospace-scale VS structures [21]. Efforts on filament-wound cylinders have also explored time-dependent material behaviour using FE model updating approaches [22]. Recent advances have also addressed manufacturability aspects such as tow steering limits, curvature control, and fibre-path smoothness in VS laminates [23]. Optimisation studies have demonstrated the potential of VS laminates to increase buckling loads, delay instability, and reduce weight [24–26].

The design of such structures, however, involves multiple competing objectives. Multi-objective optimisation frameworks are therefore essential. Beyond buckling and mass objectives, strength-based formulations employing lamination parameters and global failure indices have demonstrated computational efficiency in laminate design [27]. Therefore, a designer must simultaneously minimise mass for efficiency, maximise collapse load for strength, and maximise KDF for

robustness against imperfections. These objectives are inherently conflicting, producing a Pareto front rather than a single optimal solution [28,29]. Capturing these trade-offs is particularly critical in aerospace design, where small weight increases may be justified by disproportionate gains in reliability.

Solving such problems requires meta-optimisation methods, which are global search strategies tailored for high-dimensional, nonlinear design spaces. Classical gradient-based optimisers often fail in such settings, especially when FE models must be evaluated under imperfection sensitivity. Recent Ritz-based formulations have been shown to be effective for the direct stiffness optimisation of variable-angle tow (VAT) composite plates under multiple loads and boundary conditions [30]. However, these approaches do not explicitly address the computational scalability challenges in multi-objective optimisation of complex design spaces, such as variable-angle filament-wound composite structures. Meta-optimisation strategies, including the Non-dominated Sorting Genetic Algorithm II (NSGA-II) and Bayesian Optimisation (BO), have also been proven effective in composite design. For example, Almeida Jr. et al. [31] presented the first application of BO to VS cylinders under manufacturing constraints, demonstrating its ability to achieve rapid convergence with relatively few FE evaluations. However, that study was restricted to idealised geometries and did not incorporate imperfection realism or explicit reliability metrics.

Previous work by Wang et al. [32] used a surrogate-based approach for buckling optimisation, but without experimental imperfection data. Moreover, Lincoln et al. [15] demonstrated that stiffness tailoring can improve cylinder buckling reliability, highlighting the benefit of incorporating imperfection tolerance into design. Such challenges motivate the use of meta-optimisation methods, where classical gradient-based optimisers often fail, particularly when FE models must capture imperfection sensitivity.

Taken together, these contributions, including the first full-field imperfection measurements using DIC [9], the integration of imperfections into nonlinear FE simulations validated experimentally [10], and the first application of BO to VS cylinders under manufacturing constraints [31], have progressively laid the foundation for the present work. Yet, a critical gap remains: optimisation frameworks and imperfection-sensitive analyses have thus far evolved largely in isolation. Although the literature has advanced in geometric imperfection characterisation, knockdown modelling, and optimisation strategies, a comprehensive framework that explicitly couples measured imperfections, reliability-aware metrics, and meta-optimisation for VS cylinders is still lacking.

In this work, we establish a multi-objective optimisation framework for VS composite cylinders that explicitly incorporates measured geometric imperfections. The experimentally acquired imperfection fields are reduced using principal component analysis (PCA), statistically expanded through Latin Hypercube Sampling (LHS), mapped onto FE models, and used to train Gaussian process regression (GPR) surrogates. Based on these surrogates, we optimise fibre angle distributions for three objectives: minimising mass, maximising nonlinear buckling load, and maximising the KDF, through both Bayesian optimisation (BO) and the non-dominated sorting genetic algorithm II (NSGA-II).

Hybrid frameworks combining Bayesian optimisation and evolutionary algorithms (EAs) have been explored in certain engineering and machine-learning domains [33,34], leveraging the sample-efficient, model-based search of BO with the population-based exploration of EAs. In contrast, the present work applies BO and NSGA-II independently on identical GPR surrogates. BO algorithm exploits GPR uncertainty to achieve data-efficient convergence, while NSGA-II performs global exploration via evolutionary operators. This separation enables a direct and systematic comparison between uncertainty-guided sampling (BO) and population-based exploration (NSGA-II) under the same surrogate modelling assumptions.

We hypothesise that incorporating realistic imperfection data into the optimisation process yields composite cylinder designs with enhanced buckling robustness, characterised by increased KDF values.

The remainder of the paper is structured as follows: Section 2 describes the design variables, FE modelling, surrogate construction, and optimisation strategies; Section 3 presents and discusses the optimisation results, while Section 4 summarises the key findings and implications for imperfection-sensitive composite design. The proposed framework bridges metrology, modelling, and optimisation, providing a pathway towards lightweight, imperfection-resistant composite structures with quantifiable reliability, an essential step for next-generation aerospace design.

2. Methodology

The overall framework, summarised in Fig. 1, outlines the integrated computational process developed in this study. It presents the FE modelling, imperfection reconstruction, surrogate-assisted response prediction, and multi-objective optimisation within a unified reliability-centred design workflow.

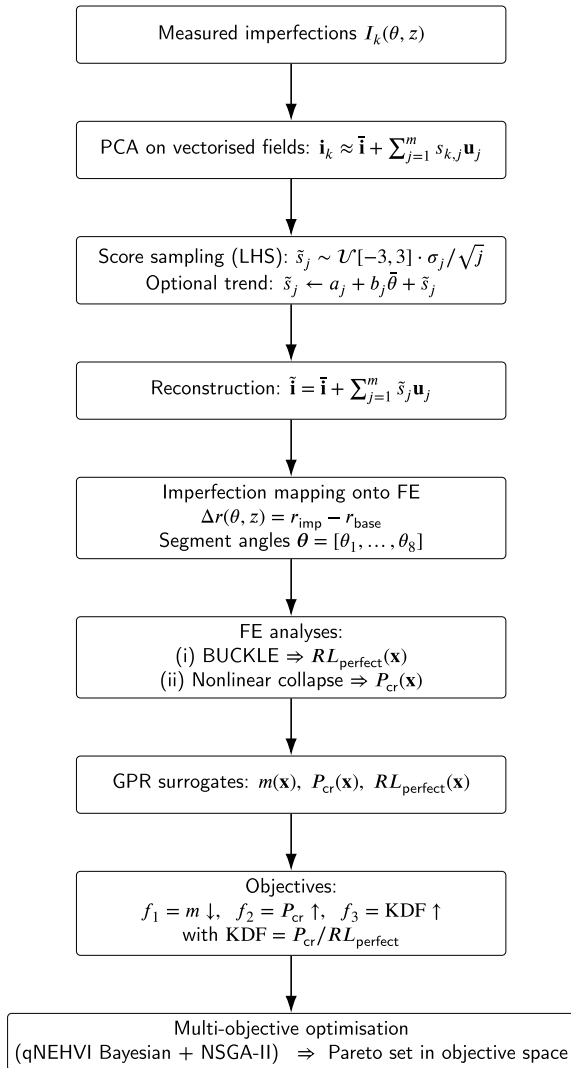


Fig. 1. Methodological flowchart linking PCA-based imperfection modelling, FE evaluation, surrogate construction, and multi-objective optimisation.

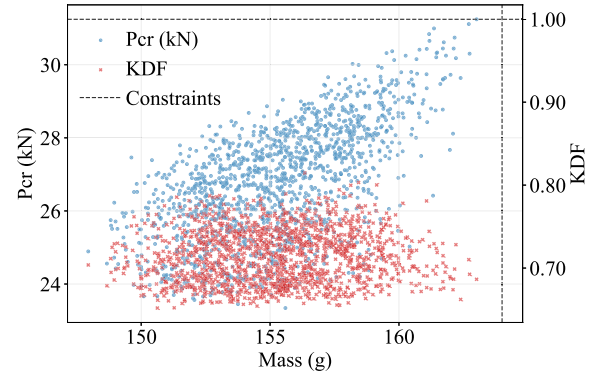


Fig. 2. Surrogate-predicted feasible design space illustrating objective trade-offs.

2.1. Design space parameterisation

The design space in this study is parameterised by the eight winding angles $\theta = [\theta_1, \dots, \theta_8]$, chosen to match the available experimental data of filament-wound cylinders reported in Castro et al. [35]. Every sampled configuration is mapped through FE analysis to three performance measures: structural mass m , nonlinear collapse load P_{cr} , and knock-down factor KDF. Fig. 2 illustrates the feasible region of this space by plotting the trade-off among the three performance measures, with the admissible domain restricted by the imposed constraints.

2.2. Finite element modelling

The evaluation of candidate designs relies on detailed FE models developed in Abaqus, combining Standard and Explicit solvers for linear and nonlinear analyses, respectively. Cylinders of length $L = 300 \text{ mm}$ and inner radius $R = 68 \text{ mm}$ were discretised with approximately 10^4 reduced-integration shell elements (S4R), arranged into 95 circumferential and 100 axial divisions to maintain near-square aspect ratios. Each axial frame was modelled with two plies of nominal 0.2 mm thickness, assigned winding angles θ_i according to the design vector \mathbf{x} . To account for fibre spreading at steeper angles, the effective ply thickness t_e was reduced according to Eq. (1) [32], ensuring constant fibre volume even when the tow is steered [36].

$$t_e = \frac{t_{\text{tow}}}{\cos(\Delta\theta)}, \quad (1)$$

where t_{tow} is the nominal tow thickness and $\Delta\theta(x)$ is the local change in winding angle at position x along the cylinder. In the present FE implementation, $\Delta\theta(x)$ is taken as the segment helical angle magnitude $|\theta_i|$ at the corresponding axial segment, so that the ply thickness written to the Abaqus input is $t_e = t_{\text{tow}} / \cos|\theta_i|$ (the sign of the angle affects orientation but not the thickness scaling). This update is applied consistently to all plies in each segment when angles are assigned. This correction prevents artificial thickening of the laminate that would otherwise occur due to mid-surface drift under fibre steering. By enforcing the tow-shear rule, laminate volume is conserved, and the stacking sequence remains physically consistent with manufacturing practice. The effect of this adjustment on the ply distribution is illustrated in Fig. 3, which contrasts the constant-thickness assumption with the corrected thickness distribution.

The geometric parameterisation used in all FE models corresponds to the eight-segment variable-angle filament-winding (VAFW8) configuration originally introduced by Almeida Jr. et al. [10]. As shown in Fig. 4, the cylinder is partitioned into eight equal axial frames, each defined by a constant winding angle θ_i applied uniformly through the laminate thickness. This discretised layout enables controlled axial stiffness variation while remaining compatible with filament-winding (FW) kinematics and forms the baseline for all subsequent analyses.




Fig. 3. Ply distribution in fibre-steered laminates (a) constant-thickness assumption; (b) corrected distribution from the tow-shear rule in Eq. (1) [37].




Fig. 4. The VAFW8 design space and control points (adapted from Almeida Jr. et al. [10]).

The lamina was modelled as a linearly elastic orthotropic T700/epoxy system with perfect interlaminar bonding, consistent with the experimentally observed buckling-dominated behaviour. The corresponding elastic constants and density are listed in Table 1.

The FE model, presented in Fig. 5, adopts a simplified variant of the experimental boundary conditions. The bottom edge is fixed in all translational degrees of freedom, while the top edge is subjected to an imposed axial displacement in the negative z -direction. Rotational degrees of freedom at both ends are left free. This choice captures the dominant loading mechanism of axial compression while avoiding over-constraint of the shell kinematics.

The overall geometry of the analysed cylinder was kept constant, with an inner radius of $R = 68$ mm and a length of $L = 300$ mm, matching the experimentally tested VAFW8 configuration by Almeida Jr. et al. [10]. This constraint on R/L ensures that all simulated designs share identical global slenderness and boundary conditions, thereby isolating the effect of variable fibre angles on the buckling response. Similarly, the adopted angular limits $55^\circ \leq \theta_i \leq 61^\circ$ are derived directly from the DIC-measured imperfection dataset of the same cylinder family reported by Castro and Almeida Jr. [35]. Restricting the design space to this experimentally supported range preserves statistical consistency between the surrogate models, the PCA-based imperfection fields, and the validated

Table 1
Material properties of T700 carbon/epoxy lamina [10,38,39].

Property	Value
Longitudinal modulus, E_1	141 GPa
Transverse modulus, E_2	9.5 GPa
Poisson's ratio, ν_{12}	0.32
Shear modulus, G_{12}	4.0 GPa
Shear modulus, G_{13}	4.0 GPa
Shear modulus, G_{23}	2.1 GPa
Density, ρ	1600 kg/m ³




Fig. 5. Finite element model and boundary conditions (adapted from Almeida Jr. et al. [10]).

specimens. Exploring wider angles would decouple the numerical surrogates from the available imperfection data, thereby compromising the physical representativeness of the optimisation framework.

Two analyses were performed for each configuration: (i) a linear buckling analysis of the perfect geometry,

$$(\mathbf{K} + \lambda \mathbf{K}_G) \phi = 0, \tag{2}$$

where \mathbf{K} is the elastic stiffness matrix, \mathbf{K}_G is the geometric stiffness matrix assembled for a reference load, λ denotes the eigenvalues (load multipliers), and ϕ are the corresponding buckling modes. The lowest eigenvalue λ_1 defines the linear buckling load as:

$$RL_{\text{perfect}} = \lambda_1 P_{\text{ref}}. \tag{3}$$

(ii) a geometrically nonlinear collapse analysis of the imperfect geometry using the dynamic explicit analysis under quasi-static loading. The reason to use dynamic explicit is to capture the explosive shift from elastic to kinetic energy that happens at the buckling collapse. Quasi-static behaviour was enforced by prescribing a sufficiently long step time, combined with mild mass scaling and bulk viscosity damping to suppress spurious dynamics. Energy histories were monitored to confirm quasi-static conditions, with the kinetic and artificial strain energies remaining below 1% of the internal energy throughout loading. The collapse load was defined as the peak of the load-shortening curve,

$$P_{cr}(\mathbf{x}) = \max_{\Delta} P(\Delta; \mathbf{x}). \tag{4}$$

The mass was calculated consistently from the FE mesh as

$$m(\mathbf{x}) = \rho \int_A t(\theta, z; \mathbf{x}) dA. \tag{5}$$

Validation against experimental load-shortening curves, Fig. 14, confirmed that the FE models reproduce both peak load and post-peak




Fig. 6. Example of a measured imperfection field on the unwrapped cylinder surface, highlighting boundary effects near the ends.

softening, making them suitable as high-fidelity ground truth for optimisation.

Since the collapse capacity of thin-walled shells is highly sensitive to geometric imperfections, these deviations from the ideal geometry must be properly characterised and incorporated into the structural model.

2.3. Imperfection modelling and mapping

Realistic imperfection fields were obtained directly from full-field DIC measurements of filament-wound cylinders [35]. Fig. 6 highlights the resin potting effects at the specimen ends, which motivated discarding the top 5% of the measured fields to eliminate this artefact. Each measured surface map was interpolated onto a uniform (θ, z) grid and vectorised into samples $i_k \in \mathbb{R}^N$. PCA reduced the dimensionality of this dataset,

$$i_k \approx \bar{i} + \sum_{j=1}^m s_{k,j} u_j, \quad (6)$$

where \bar{i} is the mean imperfection field, u_j are the PCA modes, and $s_{k,j}$ are modal coefficients. Modes with explained variance greater than 1% were retained until cumulative variance exceeded 95%, as illustrated in Fig. 7, yielding a compact, orthogonal basis that captures the dominant geometric irregularities.

To generate synthetic yet statistically consistent imperfection fields, the following steps were followed:

- (i) **Score sampling:** Synthetic mode coefficients \tilde{s}_j were sampled from a uniform $\pm 3\sigma_j$ range, assuming normally distributed amplitudes, with a $1/\sqrt{j}$ damping factor applied to suppress high-frequency content:

$$\tilde{s}_j \sim \mathcal{U}[-3, 3] \cdot \sigma_j / \sqrt{j}. \quad (7)$$
- (ii) **Optional trend:** A linear trend ($a_j + b_j \bar{\theta}$) could bias the sampling by mean winding angle. In this study, this option was disabled by setting $a_j = b_j = 0$, ensuring unbiased reconstruction.
- (iii) **Reconstruction:** Each synthetic imperfection field was reconstructed using the retained PCA modes as

$$\tilde{i} = \bar{i} + \sum_{j=1}^m \tilde{s}_j u_j, \quad (8)$$

yielding ensembles consistent with measured variability.

- (iv) **Mapping to FE mesh:** The reconstructed imperfection field $\tilde{I}(\theta, z)$ was mapped onto the FE model by perturbing the nodal coordinates:

$$r_{i,j}^* = r_{i,j} + \Delta r(\theta_i, z_j), \quad (9)$$

where $r_{i,j}^*$ and $r_{i,j}$ are the perturbed and perfect nodal coordinates, respectively.


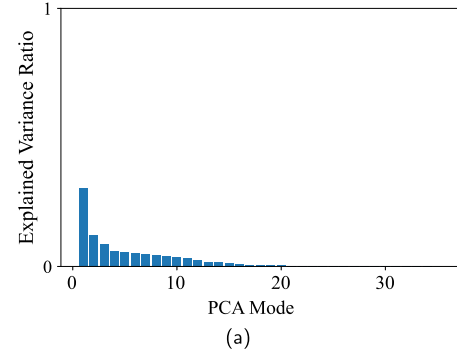


Fig. 7. PCA mode significance (a) explained variance ratio; (b) cumulative variance with the 95% threshold (dashed).

A diagnostic field was also constructed to visualise the mapped imperfections:

$$\Delta r(\theta, z) = r_{\text{imperfect}} - r_{\text{baseline}}, \quad (10)$$

confirming that the imposed geometries remained physically realistic, as shown in Fig. 12.

Since the number of measured imperfection fields available for VAFW8 family was limited ($n = 9$), a formal train/test statistical validation of the synthetic generator could not be performed. With a larger dataset, the conventional approach would be to fit the PCA model on a training subset, generate synthetic imperfections from the training scores, and compare the mean and standard deviation of imperfection amplitudes against a withheld test subset. Given the present constraint, we instead verified that the VAFW8 measured specimens from the dataset [35] display nearly identical imperfection amplitudes and minimal mass variability ($\leq 0.1\%$), indicating that the measured subset is effectively uniform. The synthetic reconstructions produced amplitudes of similar magnitude, supporting the realism of the PCA-LHS imperfections within the available evidence. The corresponding statistics are summarised in Table 2.

Although direct statistical verification of the reconstructed fields was not possible due to the limited number of experimental scans, their realism is indirectly supported by the close agreement between FE load-shortening predictions and the measured buckling behaviour. In

Table 2

Comparison of measured and synthetic VAFW8 imperfection statistics.

Dataset	RMS		Amplitude	
	Mean (mm)	SD (mm)	Mean (mm)	SD (mm)
Measured	0.16	0.00	1.12	0.00
Synthetic	0.10	0.01	0.85	0.12

summary, while we cannot statistically prove that the synthetic imperfection fields reproduce all real-world distributions (owing to the limited test dataset), the fact that they accurately replicate the experimental collapse response, displayed in Fig. 14, builds strong confidence in their realism.

This mapping ensures that each FE model reflects measured imperfection statistics while remaining computationally tractable. Having established both the modelling framework for perfect and imperfect cylinders, we now formalise the optimisation problem in terms of its design variables, objectives, and constraints.

2.4. Optimisation problem

The optimisation problem involves a trade-off among objectives. The first objective seeks to minimise the structural mass,

$$f_1(\mathbf{x}) = m(\mathbf{x}), \tag{11}$$

while the second aims to maximise the nonlinear collapse load of the imperfect geometry,

$$f_2(\mathbf{x}) = P_{cr}(\mathbf{x}). \tag{12}$$

To capture robustness against imperfections, a third objective is introduced in terms of the knockdown factor,

$$f_3(\mathbf{x}) = \text{KDF}(\mathbf{x}) = \frac{P_{cr}(\mathbf{x})}{RL_{\text{perfect}}(\mathbf{x})}. \tag{13}$$

The resulting tri-objective optimisation problem is expressed compactly as

$$\begin{aligned} \min_{\mathbf{x}} f_1(\mathbf{x}), \quad \max_{\mathbf{x}} f_2(\mathbf{x}), \quad \max_{\mathbf{x}} f_3(\mathbf{x}), \\ \text{s.t. } 55^\circ \leq \theta_i \leq 61^\circ \quad (i = 1, \dots, 8). \end{aligned} \tag{14}$$

The optimisation problem is treated as a true multi-objective problem, where all three objectives are considered simultaneously without scalarisation. Since the direct evaluation of these objectives for large design sets is computationally prohibitive, surrogate models are used to approximate the FE responses.

2.5. Surrogate modelling

The Gaussian Process Regression (GPR) surrogate, also known as Kriging, provides a probabilistic prediction of an unknown function $f(\mathbf{x})$ given training observations $D = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$. A Gaussian process assumes that function values follow a multivariate normal distribution:

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')), \tag{15}$$

where $m(\mathbf{x})$ is the mean function (taken as zero after data centring) and $k(\mathbf{x}, \mathbf{x}')$ is the covariance kernel encoding prior assumptions about smoothness and periodicity [40,41]. The present formulation corresponds to standard ordinary Kriging with a constant mean, which is mathematically equivalent to Gaussian process regression as commonly used in surrogate-based engineering design.

In this work, we employ a hybrid kernel combining a squared-exponential (RBF) term with a periodic component to capture both global smooth trends and the cyclic symmetry inherent to cylindrical coordinates:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{\|\mathbf{x}-\mathbf{x}'\|^2}{2\ell^2}\right) + \sigma_p^2 \exp\left(-\frac{2\sin^2(\pi\|\mathbf{x}-\mathbf{x}'\|/p)}{\ell_p^2}\right), \tag{16}$$

where σ_f^2, ℓ are the variance and length-scale of the RBF term, and σ_p^2, ℓ_p, p are the variance, length-scale, and period of the periodic term.

The hyperparameters $\theta = \{\sigma_f, \ell, \sigma_p, \ell_p, p\}$ are optimised by maximising the log-marginal likelihood:

$$\log p(\mathbf{y} | \mathbf{X}, \theta) = -\frac{1}{2} \mathbf{y}^\top \mathbf{K}^{-1} \mathbf{y} - \frac{1}{2} \log |\mathbf{K}| - \frac{n}{2} \log 2\pi, \tag{17}$$

with $\mathbf{K} = k(\mathbf{X}, \mathbf{X}) + \sigma_n^2 \mathbf{I}$ and σ_n^2 representing the noise variance [42].

After training, the predictive distribution for a new point \mathbf{x}_* , is the Gaussian with mean and variance:

$$\mu(\mathbf{x}_*) = \mathbf{k}_*^\top \mathbf{K}^{-1} \mathbf{y}, \quad \sigma^2(\mathbf{x}_*) = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\top \mathbf{K}^{-1} \mathbf{k}_*, \tag{18}$$

where $\mathbf{k}_* = [k(\mathbf{x}_*, \mathbf{x}_1), \dots, k(\mathbf{x}_*, \mathbf{x}_n)]^\top$. The lower confidence bound (LCB) used in optimisation is:

$$\hat{y}_{\text{LCB}}(\mathbf{x}_*) = \mu(\mathbf{x}_*) - 1.96 \sigma(\mathbf{x}_*). \tag{19}$$

This formulation provides not only a point prediction but also a quantitative measure of epistemic uncertainties, which is exploited by the Bayesian optimisation acquisition function to guide the search toward regions of great potential improvement [29,32].

An adaptive Kriging refinement was employed before surrogate training to pre-screen the full simulation dataset and identify the most informative samples [43]. The procedure iteratively evaluated prediction variance across the design space, adding points in regions of high model uncertainty until the surrogate error stabilised. In total, 2000 FE simulations were generated, from which ≈ 1000 were selected as the most informative subset for GPR training. This refinement ensured that the training data captured the dominant nonlinear trends of the responses while maintaining a compact, computationally efficient representation of the design space. The evolution of the Kriging model in Fig. 8 illustrates how the enrichment progressively densifies sampling in critical regions, yielding balanced coverage of the design domain without redundant evaluations.

GPR surrogates were subsequently trained to emulate the mapping

$$\mathbf{x} \mapsto \{m(\mathbf{x}), RL_{\text{perfect}}(\mathbf{x}), P_{cr}(\mathbf{x})\},$$

with KDF derived as the ratio between the corresponding predictions. All input variables were normalised and outputs standardised before training.

The GPR posteriors provided both mean predictions and local uncertainty estimates,

$$\hat{y}(\mathbf{x}) \pm 1.96 \sigma(\mathbf{x}),$$

enabling reliability-aware interpretation of surrogate responses. At all optimisation stages, the objective values were computed from the conservative lower confidence bounds of the GPR predictions (Eq. 19), ensuring that the search favoured designs that performed robustly even under model uncertainty. This use of GPR uncertainty effectively

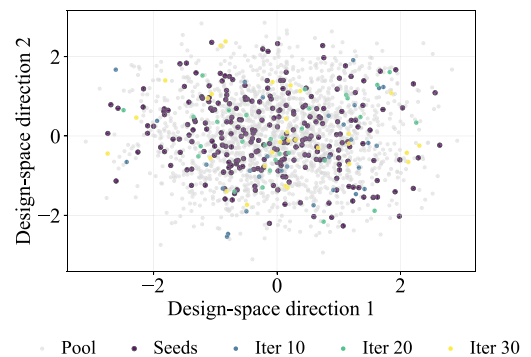


Fig. 8. Kriging-based pre-training refinement showing progressive selection of informative samples from the full simulation dataset.

Iteration 1	Test	Train	Train	Train	Train	Train	Train
Iteration 2	Train	Test	Train	Train	Train	Train	Train
Iteration 3	Train	Train	Test	Train	Train	Train	Train
Iteration 4	Train	Train	Train	Test	Train	Train	Train
Iteration 5	Train	Train	Train	Train	Test	Train	Train
Iteration 6	Train	Train	Train	Train	Train	Test	Train
Iteration 7	Train	Train	Train	Train	Train	Train	Test

Fig. 9. Schematic representation of the seven-fold cross-validation protocol, modified from Almeida Jr. et al. [31].

acts as a reliability-oriented design filter, encouraging exploration of safe regions of the design space while still allowing discovery of high-performance solutions through the q-Noisy Expected Hypervolume Improvement (qNEHVI) multi-objective acquisition.

Initial training samples were drawn from a Latin hypercube design of the eight-dimensional angle space. Predictive accuracy was assessed using a seven-fold cross-validation scheme, in which six subsets were used for training and one for testing in each iteration, ensuring that each data point was validated exactly once. The procedure is schematically illustrated in Fig. 9.

Model performance was quantified using the root-mean-square error (RMSE) and the coefficient of determination (R^2). The detailed cross-validation outcomes are discussed in Section 3.1, where surrogate predictions are compared directly with FE evaluations.

2.6. Optimisation strategies

With accurate GPR surrogates in place, the optimisation problem can be solved efficiently using advanced global search algorithms. Two complementary strategies were investigated: BO with qNEHVI and the population-based NSGA-II. The parameter settings for both approaches are summarised in Tables 3 and 4.

Bayesian optimisation: Surrogates from Section 2.5 provide predictive distributions of the objective vector

$$\mathbf{g}(\mathbf{x}) = [-m(\mathbf{x}), P_{cr}(\mathbf{x}), KDF(\mathbf{x})].$$

Pareto-optimal solutions are identified based on dominance relations in this three-dimensional objective space; a candidate design is considered optimal if no other feasible design improves one objective without

Table 3
Bayesian optimisation parameters.

Parameter	Value
Initial design points (Sobol)	50
Iterations (qNEHVI)	50
Batch size (q)	6
Trials	350
Acquisition function	qNEHVI
Surrogate model	GPR (RBF + ExpSineSquared)
Hyperparameter optimisation	Maximum marginal likelihood, 10 restarts
Constraints	$m \leq 164$ g, $KDF \leq 0.9999$

Table 4
Parameter settings for the NSGA-II.

Parameter	Value
Population size	160
Generations	250
Trials	40,000
Selection operator	Non-dominated sorting + crowding distance
Crossover operator	SBX
Crossover probability p_c	0.9
Mutation probability p_m	0.1

degrading at least one of the others. At each iteration, a batch of candidates

$$\mathbf{X} = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(q)}\},$$

was chosen by maximising the qNEHVI acquisition function, implemented through the BoTorch framework [44]. The hypervolume improvement is first defined as

$$\Delta HV(\mathbf{X}) = HV(\mathcal{P} \cup \{\mathbf{G}(\mathbf{X})\}) - HV(\mathcal{P}), \quad (20)$$

and the qNEHVI acquisition corresponds to its expected value,

$$\alpha_{\text{qNEHVI}}(\mathbf{X}) = \mathbb{E}[\Delta HV(\mathbf{X})], \quad (21)$$

where HV is the dominated hypervolume relative to a nadir point and \mathcal{P} is the current Pareto set.

The optimisation was initialised by a Sobol sequence design, after which the qNEHVI acquisition iteratively guided sampling toward regions with the highest expected hypervolume gain. Objective evaluations were performed using the conservative *lower confidence bounds* of the surrogate predictions, Eq. (19), favouring designs that remain reliable under predictive uncertainty. This uncertainty-aware formulation promotes conservative optimisation, preventing overly optimistic solutions. Feasibility constraints on mass and knockdown factor, Eq. (22), were imposed directly within the acquisition model to restrict sampling to feasible regions. Batches were evaluated synchronously according to the available FE cores, and the optimisation was terminated upon convergence of the hypervolume improvement, corresponding to 350 surrogate evaluations.

$$m \leq 164 \text{ g} \quad \text{and} \quad KDF \leq 0.9999. \quad (22)$$

NSGA-II: The optimisation problem was also solved independently using the Non-dominated Sorting Genetic Algorithm II (NSGA-II). A population size of 160 and 250 generations was used, corresponding to 40,000 surrogate-based objective evaluations, with no additional FE simulations required. Each individual encoded eight winding angles as real-valued genes. Offspring were generated by simulated binary crossover (SBX, $p_c = 0.9$) and polynomial mutation ($p_m = 0.1$), while parent selection employed binary tournaments. Elitist non-dominated sorting determined survivors, and crowding distance preserved diversity within each front.

To maintain computational efficiency, all objective evaluations within the evolutionary loop were performed using the lower confidence bounds of the surrogates, providing conservative estimates of the structural responses. This follows the same *lower confidence bound* surrogate logic used in BO runs, ensuring consistent and conservative evaluation across both optimisation strategies.

The resulting non-dominated front thus represents Pareto-optimal trade-offs inferred from uncertainty-aware surrogate predictions, whose fidelity was independently verified through cross-validation and variance analyses.

The two optimisation strategies were executed independently; BO applied uncertainty-guided acquisition on the surrogate models, while NSGA-II explored the design space using population-based evolutionary operators. Each method produced its own Pareto-optimal solutions, and their resulting sets and convergence histories are presented in Section 3.3.

However, not all designs on the preliminary Pareto sets are equally robust to geometric variability. In practice, imperfection sensitivity can undermine apparent performance advantages if the variability of knockdown response is excessive. To address this, the optimisation framework incorporates an explicit reliability assessment step, in which the ensemble statistics of knockdown factors are evaluated for each candidate. Only those designs that satisfy robustness requirements are retained, ensuring that the final Pareto fronts represent solutions that are not only light and strong but also consistently reliable.

2.7. Reliability and sensitivity metrics

Reliability and interpretability were assessed jointly to ensure that the reported designs are not only efficient but also robust and explainable. Reliability was quantified through ensemble evaluation of imperfection tolerance. For each candidate \mathbf{x} , N_s imperfection realisations were drawn from the PCA-LHS space and mapped onto the FE mesh. Each configuration was analysed using the Dynamic Explicit quasi-static procedure, yielding an ensemble of knockdown factors $\{\text{KDF}_j(\mathbf{x})\}_{j=1}^{N_s}$. The ensemble statistics were computed as

$$\begin{aligned} \mu_{\text{KDF}}(\mathbf{x}) &= \frac{1}{N_s} \sum_{j=1}^{N_s} \text{KDF}_j(\mathbf{x}), \\ \sigma_{\text{KDF}}(\mathbf{x}) &= \sqrt{\frac{1}{N_s - 1} \sum_{j=1}^{N_s} (\text{KDF}_j(\mathbf{x}) - \mu_{\text{KDF}}(\mathbf{x}))^2}, \end{aligned} \quad (23)$$

where μ_{KDF} measures average robustness, while σ_{KDF} quantifies sensitivity to imperfection variability. Only designs with acceptable σ_{KDF} were retained on the final Pareto fronts.

Complementarily, a global sensitivity analysis was conducted to quantify the influence of each design variable on the surrogate-predicted responses. Pearson correlation coefficients were computed between each winding angle θ_i and the surrogate posterior means of $m(\mathbf{x})$, $P_{\text{cr}}(\mathbf{x})$, and $RL_{\text{perfect}}(\mathbf{x})$ across a broad Latin hypercube sample of the admissible design space. The resulting global sensitivity indices are summarised in Table 5 and discussed in Section 3.2.

3. Results & discussion

The credibility of any optimisation framework depends on how faithfully its numerical foundations reflect physical reality. In this study, the integration of measured imperfection statistics, nonlinear collapse behaviour, and surrogate-based prediction must together reproduce the mechanisms governing buckling in filament-wound cylinders. Establishing this consistency is essential before interpreting the optimisation outcomes, as the significance of any observed design trend ultimately rests on the accuracy with which the computational models capture the imperfection-sensitive structural response.

3.1. Validation of imperfection modelling

Representative principal component modes obtained from the measured imperfection fields are illustrated in Figs. 10 and 11, shown on both the unwrapped cylinder surface and the reconstructed three-dimensional geometry. The PCA basis captures a wide spectrum of imperfection patterns, ranging from long-wavelength, global ovalisation-type distortions to short-wavelength, fibre-scale surface undulations. Based on a cumulative explained-variance criterion, the first fifteen modes were retained, capturing 95% of the total variance of the measured imperfection dataset. Each mode represents a statistically independent geometric fluctuation rather than a deterministic physical feature, highlighting that the experimental imperfections can be

Table 5

Global sensitivity indices (Pearson correlation coefficients) between winding angles and surrogate-predicted responses.

Angle	Mass Corr	P_{cr} Corr	RL_{perfect} Corr
θ_1	0.3803	0.1948	0.3880
θ_2	0.3773	0.1510	0.2134
θ_3	0.3415	0.0877	0.2003
θ_4	0.2750	-0.0089	0.1612
θ_5	0.3697	0.1711	0.2336
θ_6	0.3303	0.0500	0.1548
θ_7	0.3517	0.0263	0.1931
θ_8	0.3490	0.7362	0.3181

expressed as a combination of orthogonal spatial components. The truncation to these dominant modes preserves the characteristic imperfection wavelength content relevant to buckling behaviour while filtering out high-frequency components primarily associated with measurement noise.

This compact, low-dimensional representation retains the essential spatial characteristics of the measured data while filtering out random measurement noise and redundant correlations. Physically, the retained low-order PCA modes correspond to the dominant long-wavelength imperfection patterns observed in the measurements, while higher-order modes predominantly capture short-wavelength fluctuations associated with experimental noise and minor local irregularities. Truncating the expansion to the dominant modes therefore acts as an effective spatial low-pass filter, preserving the characteristic imperfection wavelengths governing buckling behaviour while suppressing non-structural measurement artefacts. The randomised linear combination of these orthogonal modes, using statistically consistent modal coefficients, ensures that the reconstructed imperfection fields retain the measured variability in both amplitude and spatial pattern without reintroducing measurement-specific noise.

Consequently, the PCA space provides an efficient and physically interpretable foundation for synthesising statistically consistent imperfection fields. These reconstructed fields are later employed in the nonlinear FE analyses to quantify sensitivity to imperfection amplitude and spatial pattern variability.

The realism of the imperfection mapping was verified through evaluation of the radial deviation field, defined in Eq. (10), as shown in Fig. 12. The contour plots display the spatial distribution of surface deviations along the cylinder, with the superimposed helical traces indicating the prescribed winding angles θ_i . The mapped fields demonstrate that the imperfections are coherently embedded within the finite element mesh, preserving both the winding geometry and the segmentation boundaries. The reconstructed surfaces exhibit continuous and physically plausible variations along the axial direction, confirming that the imposed imperfection fields accurately follow the designed fibre paths and are suitable for subsequent quantitative buckling assessment.

Fig. 13 presents the knockdown factors (KDFs) obtained from the nonlinear FE analyses as a function of the slenderness ratio r/t , alongside the NASA/SP-8007-2020/REV-2 [7] lower-bound guideline for isotropic shells. Unlike the empirical NASA curve, which represents metallic cylinders under idealised assumptions, each data point in the present dataset corresponds to a unique composite lay-up with segment-specific fibre orientations θ_i , leading to distinct stiffness matrices \mathbf{A} , \mathbf{B} , and \mathbf{D} . The KDF for each configuration was computed using Eq. (13), with the effective thickness inferred from total mass of the model to ensure consistency with the analytical definition of slenderness. This comparison highlights the broader variability and higher fidelity of the present database, which captures anisotropy- and imperfection-driven deviations beyond the scope of the NASA lower-bound correlation. The effective thickness of each configuration was determined from its computed mass, which is defined as

$$t_{\text{eff}} = \frac{m}{2\pi RL\rho}, \quad \frac{r}{t_{\text{eff}}} = \frac{R}{t_{\text{eff}}}, \quad (24)$$

with $R = 68$ mm, $L = 300$ mm, and $\rho = 1.6 \times 10^{-3}$ g/mm³. For a general laminate, the extensional, coupling, and bending stiffness matrices are assembled as

$$\mathbf{K}_{\text{ABD}} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B} & \mathbf{D} \end{bmatrix}, \quad \mathbf{D} = \sum_{k=1}^N \bar{\mathbf{Q}}^{(k)} \frac{(z_k^3 - z_{k-1}^3)}{3}, \quad (25)$$

where $\bar{\mathbf{Q}}^{(k)}$ is the transformed reduced stiffness matrix of ply k , and z_k denotes its top and bottom coordinates in the laminate. Since $\bar{\mathbf{Q}}^{(k)}$ depends on the fibre orientation θ_k , any variation in θ_i modifies both the bending stiffness \mathbf{D} and, through the tow-shear relation $t_e = t_{\text{tow}} / \cos |\theta_i|$, the local laminate thickness. Consequently, although all samples exhibit a




Fig. 10. Retained PCA imperfection modes shown on the unwrapped cylindrical surface.

similar geometric slenderness ($r/t \approx 90$), their effective flexural rigidity D_{11} and coupling terms differ, producing the observed spread in KDF values. This dispersion reflects the combined influence of stiffness tailoring and thickness variation, demonstrating that the FE dataset, while lying slightly above the NASA lower-bound envelope, captures the anisotropic behaviour and imperfection sensitivity typical of filament-wound composites.

Further evidence of fidelity is provided by the load-shortening response in Fig. 14. The comparison between the experimental curves [10], labelled *EX-VAFW8-1,2,3*, and the numerically simulated cylinders containing the reconstructed imperfection fields, labelled *S- i,j,k* , shows excellent correspondence in stiffness, peak load and post-buckling region. The numerical models capture the rapid load drop following the critical point and reproduce the characteristic asymmetric dimple observed experimentally. This agreement is particularly important because the emergence and location of the asymmetric dimple are governed by imperfection-triggered localised buckling mode selection, which cannot

be inferred from global stiffness or peak load alone. These results demonstrate that the PCA-based reconstruction and subsequent FE mapping accurately preserve the imperfection sensitivity inherent to real specimens. Consequently, the generated imperfect models can be considered reliable computational analogues for subsequent surrogate development and optimisation.

3.2. Surrogate behaviour and predictive accuracy

The predictive performance of the GPR surrogates was evaluated through seven-fold cross-validation, and the results are summarised in Fig. 15. The mass surrogate, Fig. 15(a), shows an almost perfect match with the FE data, reproducing the deterministic relationship between winding angle and structural weight with negligible error. Predictions of the linear buckling load for the perfect geometry, Fig. 15(c), remain highly correlated with the FE results, whereas the nonlinear collapse load, Fig. 15(b), shows slightly larger dispersion consistent with the




Fig. 11. Retained PCA imperfection modes represented on the wrapped surface.

intrinsic variability of imperfection-sensitive collapse behaviour. This spread does not indicate model bias but reflects the natural variability of nonlinear responses under distinct imperfection realisations. This behaviour arises because, in contrast to mass and linear buckling load, the nonlinear collapse load is governed by a strongly imperfection-sensitive post-buckling path, in which both the spatial pattern and amplitude of the imperfections influence the initiation and evolution of localised failure. As a result, the mapping between design variables and collapse load is inherently more scattered, since similar winding configurations can lead to different collapse responses under different imperfection realisations. This intrinsic multi-valued behaviour manifests as increased dispersion in cross-validation metrics and is therefore a physical characteristic of the problem rather than a limitation of the surrogate model.

Quantitatively, the seven-fold cross-validation yielded $R^2 = 0.99$ and $RMSE = 0.21$ g (1.3%) for the mass surrogate, $R^2 = 0.75$ and $RMSE = 0.89$ kN (8.6%) for the nonlinear collapse load, and $R^2 = 0.87$ and $RMSE = 0.55$ kN (5.1%) for the linear buckling load. As expected, the nonlinear collapse load exhibits a lower R^2 and higher RMSE than the mass and linear buckling load surrogates, reflecting the greater intrinsic variability introduced by imperfection-sensitive post-buckling behaviour rather than deficiencies in the surrogate modelling. These metrics confirm the high predictive accuracy of the surrogates and their suitability for subsequent multi-objective optimisation.

Across all three responses, the regression slopes remain close to unity, with no systematic deviation observed, indicating that the surrogates accurately capture the physical scaling and gradients of each response without underfitting or excessive smoothing. The maintained correlation




Fig. 12. Radial imperfection maps $\Delta r(\theta, z)$ after FE mapping.

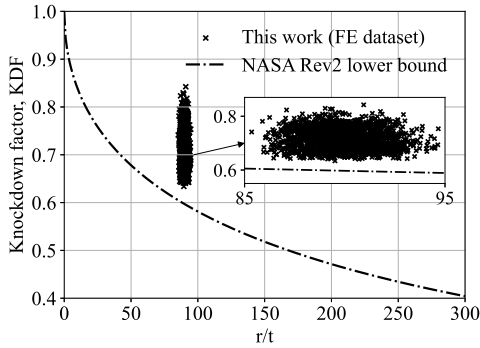


Fig. 13. Knockdown factor as a function of r/t , comparing NASA lower bound guide with present FE samples, modified from NASA/SP-8007-2020/REV-2 [7].

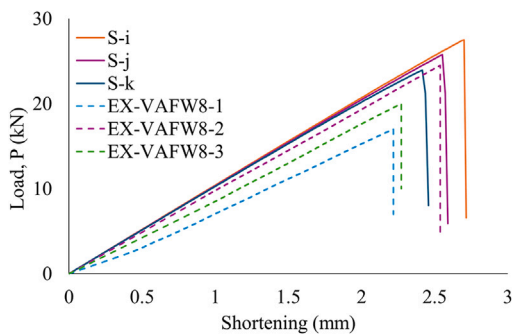
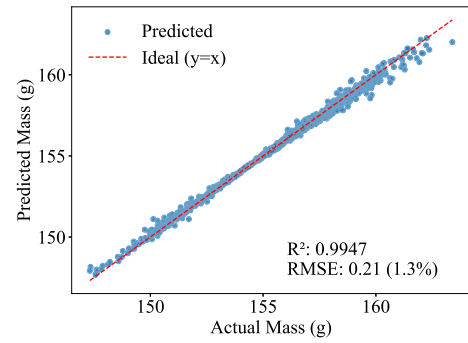


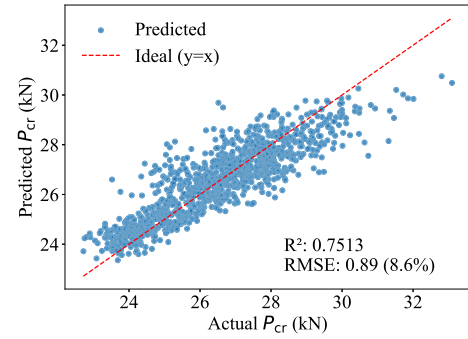
Fig. 14. Load-shortening model fidelity: experimental curves [10] and predicted responses of cylinders with synthetic imperfections.

for P_{cr} , even under pronounced geometric nonlinearity, demonstrates that the surrogate captures a genuine physical trend between the winding configuration and collapse strength rather than a spurious random statistical correlation. Overall, the GPR models provide a faithful and computationally efficient approximation of the FE responses, ensuring suitable use in the optimisation framework.


To further interpret the surrogate behaviour, a global sensitivity analysis was performed to quantify the influence of each winding angle θ_i on the predicted responses. Pearson correlation coefficients were



(a)



(b)



(c)

Fig. 15. Cross-validated surrogate predictions against true FE results for (a) mass, (b) nonlinear collapse load, and (c) linear buckling load of the geometrically perfect model.

evaluated between each θ_i and the posterior mean predictions of the mass $m(\mathbf{x})$, nonlinear collapse load $P_{cr}(\mathbf{x})$, and the linear buckling load of the perfect geometry $RL_{perfect}(\mathbf{x})$ over a Latin hypercube sample of the admissible design space. The results, summarised in Table 5, indicate a uniformly positive correlation between all θ_i and $m(\mathbf{x})$, reflecting the coupling between fibre angle and local laminate thickness. The collapse load P_{cr} shows a dominant dependence on θ_8 , corresponding to the constrained end segment, consistent with boundary-stiffening effects. The $RL_{perfect}$ response exhibits more distributed, moderate correlations across the shell length, suggesting that the global stiffness pattern, rather than local variations, governs linear stability and that multiple winding-angle segments contribute to the linear buckling response. These sensitivity indicators provide a compact quantitative interpretation of the relationships later visualised in the surrogate response surfaces.

The behaviour of the trained Kriging surrogates across the design domain is illustrated in Fig. 16 through two-dimensional contour slices extracted from the eight-dimensional response surfaces. Each plot

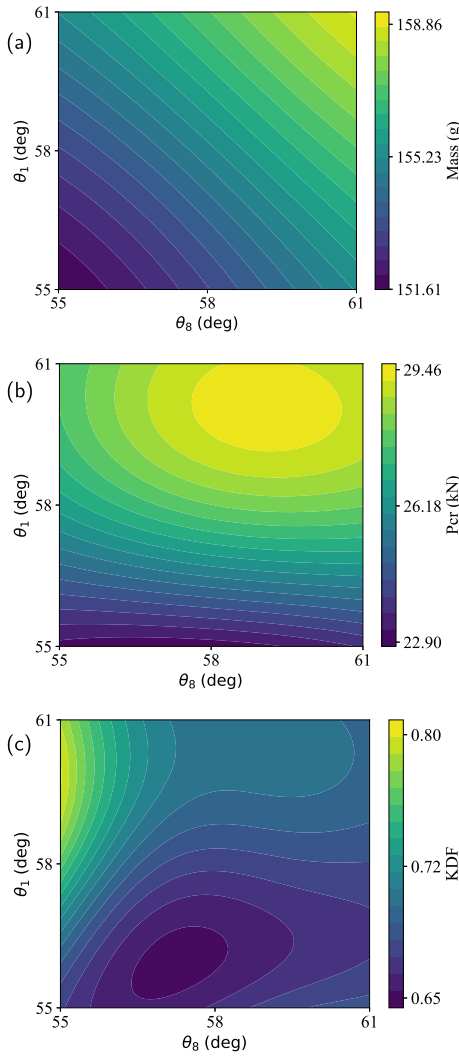


Fig. 16. Surrogate-predicted response surfaces with respect to segment angles (θ_1, θ_8) for (a) mass, (b) nonlinear collapse load, and (c) knockdown factor.

depicts the surrogate-predicted response for a given pair of winding angles $(\theta_i, \theta_j) \in [55^\circ, 61^\circ]^2$, while the remaining angles were fixed at the mean of their admissible range. This projection allows visualisation of local curvature and variable interaction, elucidating how segment-wise fibre adjustments influence the overall performance metrics.

The resulting response surfaces are smooth and physically consistent. The structural mass increases monotonically with the mean fibre angle due to the tow-shear thickness correction. This trend directly reflects the kinematic tow-shear correction, whereby the effective laminate thickness scales approximately with $t_e \propto 1/\cos|\theta|$ as the fibre steering angle increases, leading to a deterministic increase in mass. The monotonic mass response, therefore, indicates that the surrogate model is correctly learning the underlying manufacturing constraint associated with volume conservation during tow steering, rather than a spurious numerical correlation.

In contrast, P_{cr} exhibits a single broad optimum at mid-range angles where axial and circumferential stiffness are balanced. The KDF field displays gradual gradients without irregular oscillations, indicating stable surrogate behaviour. The continuity and monotonicity of these contours confirm that the Kriging interpolation captures genuine mechanical relationships among the winding angles rather than artefacts of numerical fitting, thereby reinforcing confidence in the reliability of the surrogate for multi-objective optimisation.

3.2.1. Extrapolation analysis and design-space justification

To assess whether expanding the fibre-angle range beyond the experimentally supported interval $55^\circ \leq \theta_i \leq 61^\circ$ would yield meaningfully improved and physically reliable Pareto-optimal solutions, we analysed the predictive uncertainty of the GPR surrogates when extrapolating outside the training bounds. The Gaussian process formulation provides not only a mean prediction but also a well-calibrated estimate of epistemic uncertainty, expressed as the posterior standard deviation $\sigma(x)$. This uncertainty grows when the model is queried at points far from the training data, indicating regions where the surrogate predictions become unreliable. The analysis does not aim to prove that superior designs cannot exist outside the investigated range; rather, it evaluates whether such designs could be identified with sufficient predictive confidence using the present surrogate and imperfection-based framework.

We performed a one-dimensional extrapolation analysis by varying all eight winding angles *simultaneously* along a line where $\theta_1 = \theta_2 = \dots = \theta_8$. This corresponds to a uniform fibre orientation along the entire cylinder, representing a simple but physically interpretable path through the design space that avoids introducing artificial angle gradients not present in the training data. For each common angle value θ in a sweep from 45° to 70° , we evaluated the posterior standard deviation of the surrogate-predicted nonlinear collapse load P_{cr} . Fig. 17 shows the resulting uncertainty profile.

As expected, the uncertainty remains low and nearly constant within the training bounds but rises markedly beyond them. This behaviour is intrinsic to the GPR kernel (Eq. 16), which provides confident interpolations where data are available but cannot guarantee accurate extrapolation. Consequently, any optimisation performed outside the 55° – 61° range would rely on predictions with high epistemic uncertainty, risking the production of physically unrealistic or non-manufacturable designs.

Moreover, the chosen angular limits are not arbitrary; they correspond to the *experimentally measured imperfection dataset* of the same cylinder family [9]. Using angles outside this range would decouple the PCA-LHS synthetic imperfections from the structural response, compromising the physical representativeness of the knockdown-factor evaluation. Since the KDF quantifies imperfection sensitivity, extrapolating beyond the measured imperfection range would invalidate the robustness metrics that are central to this reliability-aware framework.

Therefore, while a wider angle range might, in theory, alter stiffness-mass trade-offs (e.g., via steeper or shallower fibre orientations), the present study intentionally restricts the design space to the experimentally supported interval. This restriction ensures the following critical properties:

1. **Surrogate fidelity:** predictions remain reliable within the training domain;
2. **Imperfection consistency:** the PCA-LHS synthetic fields are statistically representative of real manufacturing variability;




Fig. 17. Posterior standard deviation of the nonlinear collapse load surrogate as a function of a representative winding angle θ_i .

The narrow angle range thus does not reflect a computational limitation but a deliberate methodological choice to maintain physical realism and predictive confidence. As shown in Section 3.3, even within this restricted domain, the optimisation reveals rich trade-offs among mass, collapse load and knockdown factor, yielding Pareto-optimal designs that balance lightweight efficiency with imperfection robustness under experimentally validated manufacturing variability.

Future work could expand the admissible angle range by acquiring additional full-field imperfection measurements for a broader set of winding configurations, thereby extending the reliability-aware optimisation framework to a wider design space without sacrificing predictive accuracy or physical consistency.

3.3. Optimisation results

The multi-objective optimisation produced coherent and physically interpretable Pareto fronts, confirming the internal consistency of the surrogate-predicted responses. Fig. 18 illustrates the trade-off among structural mass, nonlinear collapse load, and KDF obtained using BO and NSGA-II independently. Both solvers identified a Pareto surface, demonstrating that the surrogate models successfully captured the nonlinear couplings among mass, stiffness, and imperfection sensitivity. The dense clustering of points along the frontier indicates a smooth and continuous response landscape, where marginal improvements in collapse load or KDF inevitably require proportional sacrifices in mass. Despite relying on identical surrogate models, the Pareto fronts obtained by BO and NSGA-II differ in their distributions and densities. BO produces a more compact and locally refined Pareto front by adaptively concentrating evaluations in regions of maximal expected hypervolume improvement, whereas NSGA-II maintains population diversity across the objective space, resulting in a more dispersed set of non-dominated solutions. These differences reflect algorithmic exploration strategies rather than inconsistencies in the surrogate predictions, as both methods converge towards the same underlying trade-off trends.

Although both methods converged toward nearly identical envelopes, distinct algorithmic behaviours were observed. BO reached convergence within approximately 350 surrogate evaluations, a fraction of the computational cost of NSGA-II, which needed around 40,000 trials to achieve a similar Pareto front. The efficiency of BO originates from its qNEHVI acquisition strategy, which accounts for predictive noise while directing sampling toward regions of high uncertainty and expected hypervolume improvement, thereby concentrating computational effort where it yields the greatest gain in solution quality. In the present optimisation landscape, the objective responses vary smoothly with the winding-angle parameters and the Pareto-optimal set occupies a narrow, well-structured manifold within the high-dimensional design space; under such conditions, uncertainty-guided sampling enables BO to rapidly

localise the Pareto front. In contrast, NSGA-II requires many additional evaluations to discover and maintain this narrow envelope through stochastic population-based exploration, resulting in a more dispersed yet globally exhaustive set of solutions. While applied independently, the two approaches exhibit complementary strengths: BO achieves high precision through uncertainty-guided exploitation, whereas NSGA-II enhances diversity through evolutionary exploration. Their comparison demonstrates how different search philosophies converge toward consistent Pareto-optimal trends, collectively ensuring a balanced assessment of precision and diversity, an essential aspect of multi-objective problems dominated by competing physical mechanisms.

The convergence histories presented in Fig. 19 provide a quantitative view of this behaviour. The dominated hypervolume, a metric of Pareto front advancement, grew monotonically for both algorithms, plateauing once additional iterations produced less than 1% improvement. BO achieved 95% of its final hypervolume after roughly forty iterations, whereas NSGA-II reached comparable stability after about 180 generations. The absence of oscillations or regressions in the hypervolume trajectory evidences the numerical robustness of both solvers and confirms that the surrogate models capture the essential nonlinearities governing buckling and imperfection sensitivity.

To further ensure physical credibility, the Pareto sets were filtered using ensemble statistics of the knockdown factor, eliminating solutions that exhibited excessive sensitivity to random imperfection realisations. Each retained design, therefore, represents not only an optimal balance of objectives but also a configuration resilient to the stochastic variability inherent in manufacturing-induced geometry deviations. This reliability-aware filtering step effectively narrows the solution space to designs that balance lightweight efficiency with consistent nonlinear collapse performance, establishing a foundation for interpreting the design trade-offs and the representative configurations.

3.3.1. Representative designs

Representative Pareto-optimal designs obtained from both optimisation strategies are summarised in Table 6 and visualised in Fig. 20. Each pair of configurations corresponds to a distinct region of the Pareto front, with solutions optimised for minimum mass, maximum nonlinear collapse load, and maximum KDF. Despite originating from two fundamentally different optimisation strategies, the resulting designs display strong qualitative similarity, confirming the consistency and reliability of the surrogate-assisted optimisation framework.

Both algorithms converged towards fibre architectures exhibiting a gradual variation of winding angles along the cylinder length. In all cases, the end regions adopt shallower orientations, while the mid-sections become progressively steeper. This axial gradient in fibre angle redistributes stiffness, reducing local bending susceptibility near the

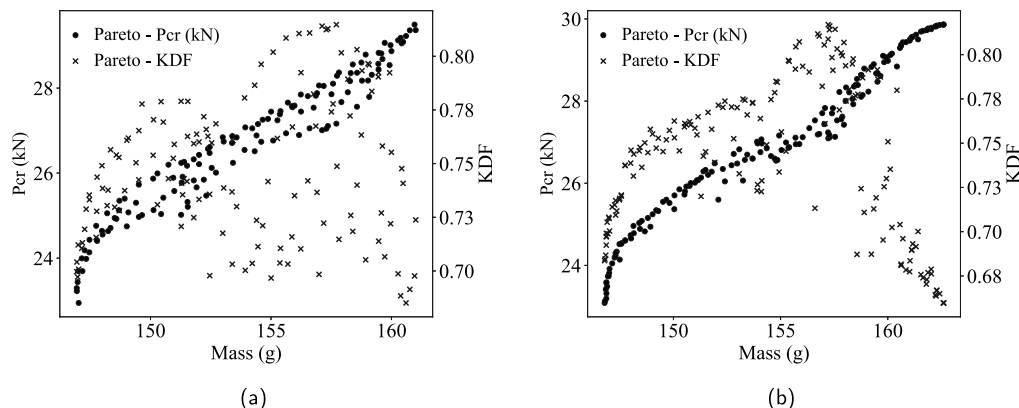


Fig. 18. Pareto-optimal fronts for tri-objective problem obtained using (a) BO and (b) NSGA-II.

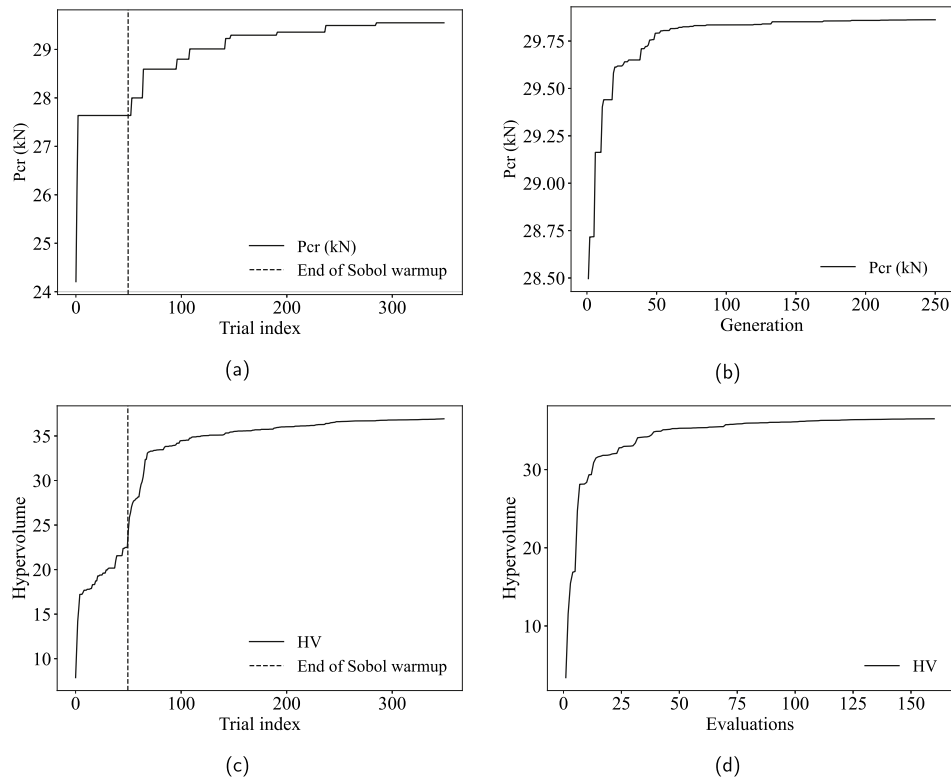


Fig. 19. Convergence behaviour of the optimisation algorithms. (a,b) evolution of the nonlinear collapse load P_{cr} and (c,d) evolution of the dominated hypervolume for BO and NSGA-II, respectively.

Table 6
Representative designs of BO and NSGA-II.

Design	Method	Optimised design variables	Mass (g)	P_{cr} (kN)	KDF
Min. Mass	BO	[55.0, 55.3, 55.7, 55.3, 55.9, 55.1, 56.3, 56.1]	146.94	23.29	0.70
	NSGA-II	[55.4, 55.6, 55.6, 55.7, 55.8, 55.7, 55.7, 55.7]	146.79	23.09	0.68
Max. P_{cr}	BO	[59.1, 58.9, 59.6, 60.2, 60.1, 59.0, 56.9, 60.3]	160.96	29.49	0.70
	NSGA-II	[59.3, 59.9, 59.9, 60.2, 60.1, 59.6, 58.2, 60.5]	162.60	29.86	0.66
Max. KDF	BO	[59.5, 59.7, 59.8, 59.5, 59.9, 55.0, 55.0, 60.6]	157.71	27.16	0.81
	NSGA-II	[59.4, 59.8, 59.7, 58.9, 59.6, 55.0, 55.0, 60.4]	157.23	27.09	0.82

loaded edges and delaying the onset of circumferential dimpling. The winding angle distributions obtained from BO and NSGA-II exhibit consistent, manufacturable profiles across all objectives. BO designs feature smoothly modulated segment transitions driven by its uncertainty-guided refinement, yielding physically coherent stiffness tailoring along the cylinder. NSGA-II produces trends that are comparable to those identified by BO, corroborating the robustness of the BO-identified design landscape. Both outcomes remain within manufacturability limits of FW, reinforcing the capability of BO to deliver manufacturing-ready solutions.

Comparison across objective extremes highlights the physical mechanisms underlying the trade-offs. In the minimum-mass configuration, the dominance of shallower angles results in lower axial stiffness and, hence, reduced collapse capacity but provides the most material-efficient structural form. At the opposite extreme, maximum P_{cr} designs rely on globally steeper orientations that increase the axial load-carrying capacity at the expense of added mass. Intermediate solutions, achieving maximum KDF, exhibit a more balanced stiffness distribution, with the end regions remaining shallow enough to maintain circumferential stability, while central segments adopt slightly steeper fibres to enhance overall rigidity. This combination mitigates imperfection sensitivity and promotes a stable nonlinear collapse response, consistent

with the imperfection-tolerant mechanisms observed experimentally in filament-wound shells.

These patterns confirm that the mechanical performance of variable-angle cylinders is not governed by a single dominant segment but by the collective smoothness of the stiffness trajectory along the shell. Excessively abrupt steering changes tend to localise stresses and trigger premature local buckling, whereas continuous fibre curvature sustains load redistribution. The correspondence between BO and NSGA-II solutions therefore emphasises the robustness of the optimisation landscape and the reproducibility of design trends across independent search strategies.

From a broader perspective, the observed trade-offs depict a clear physical hierarchy among the objectives where mass reduction inherently competes with both collapse strength and imperfection tolerance, while improvements in P_{cr} and KDF reinforce each other within a narrow region of the design space. This explains the nearly co-incident high-performance zone on the Pareto surface, where lightweight efficiency, strength, and robustness reach an optimal compromise. The representative designs thus capture the essential design philosophy of VAFW structures in which stiffness is tailored not only for strength maximisation but also for stability and reliability under imperfection sensitivity.

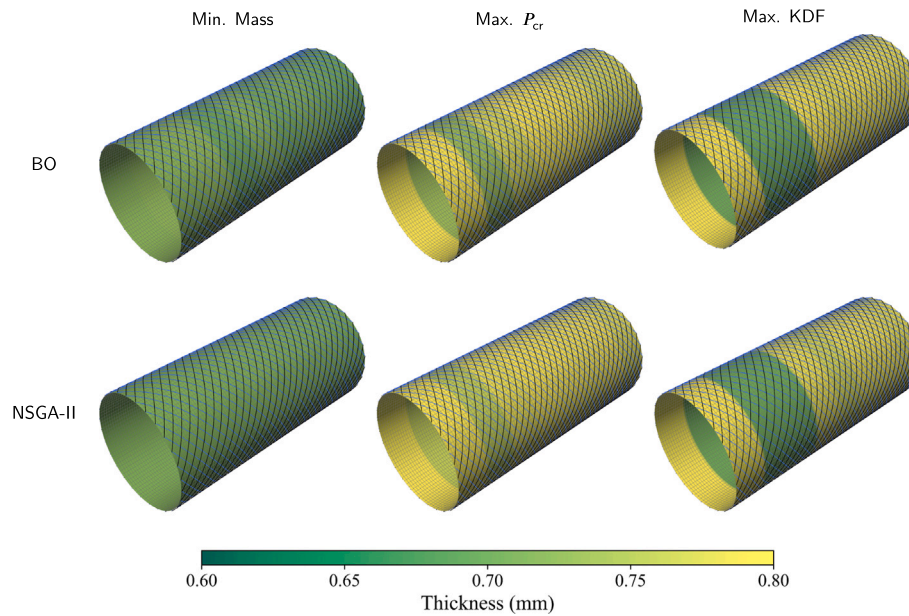


Fig. 20. Representative fibre architectures obtained using BO and NSGA-II, illustrating the respective designs associated with minimum mass, maximum collapse load (P_{cr}), and maximum knockdown factor (KDF).

3.3.2. Computational cost

Before comparing the two optimisation strategies, it is worth emphasising that the extensive number of design evaluations required to obtain these results would have been impractical without the surrogate-assisted framework. Each batch of ten parallel nonlinear FE analyses, executed simultaneously across available computational cores, required approximately two hours of wall-clock time. This computational demand underscores the prohibitive cost of direct design-space exploration and justifies the use of GPR surrogates, which replace high-fidelity simulations with rapid, uncertainty-aware predictions. The surrogate-based substitution enables the systematic evaluation of thousands of candidate designs and renders the multi-objective optimisation of imperfection-sensitive composite cylinders computationally tractable.

The computational effort of the two optimisers reflects their fundamentally distinct search paradigms. BO repeatedly refits GPR surrogates and solves acquisition-optimisation problems using multi-start stochastic search combined with Monte Carlo integration. These additional operations deliberately increase the per-iteration cost while providing rigorous uncertainty quantification and enabling sample-efficient exploration of the design space. In contrast, NSGA-II employs simpler evolutionary operators, such as selection, crossover and mutation, that introduce negligible computational overhead while allowing large populations of candidate designs to be evaluated in parallel.

These differences lead to distinct runtime profiles. Although both methods rely on identical surrogate models and therefore incur negligible cost per function evaluation, BO requires model retraining and acquisition optimisation at every iteration. In practice, BO performed approximately 350 surrogate evaluations compared with about 40,000 for NSGA-II. Despite this drastic reduction in function calls, the wall-clock time of BO was longer, which is a direct consequence of its iterative model fitting and probabilistic sampling. This additional cost is not a limitation but an intentional trade-off, whereby BO exchanges computational intensity for informed, uncertainty-guided decision-making.

Although more time-consuming, BO achieved Pareto-optimal solutions of comparable quality using two orders of magnitude fewer evaluations. This contrast highlights their complementary strengths: BO offers a principled and data-efficient search process, whereas NSGA-II favours rapid execution at the expense of a substantially larger evaluation budget. Together, these results illustrate the value of combining

exploration-oriented and uncertainty-aware optimisation approaches for efficient design of imperfection-sensitive composite structures.

3.3.3. Manufacturability and design insights

The optimisation outcomes reveal clear relations between fibre steering strategies, manufacturability limits and structural performance. Configurations obtained through BO exhibit stiffness gradients and gradual fibre transitions comparable to those from NSGA-II, both yielding architectures that are easier to realise through the FW process. These designs avoid abrupt curvature changes and excessive tow steering, which otherwise cause fibre waviness and resin-rich zones that degrade both buckling capacity and reproducibility. By contrast, extreme solutions located at the edges of the Pareto front often involve sharp angle variations near the mid-span or end regions, exceeding the practical winding curvature limits and leading to reduced geometric fidelity.

From a manufacturing standpoint, the optimised angle distributions indicate that feasible variable-angle trajectories should maintain continuous steering below approximately 10° per segment for the present mandrel geometry and tow width [10]. This constraint preserves tension uniformity during winding and promotes repeatable consolidation without inducing fibre wrinkling. Maintaining such smooth transitions also ensures a more uniform thickness distribution and dimensional stability after curing, both essential for minimising imperfection amplitude and sustaining high knockdown factors.

The design interpretation further indicates that the performance improvements stem from redistributing stiffness rather than simply increasing material. Intermediate Pareto configurations, particularly those maximising KDF, achieve higher collapse loads at nearly constant mass, demonstrating that moderate end angles and balanced axial tailoring can enhance buckling strength by up to 20–25% within current manufacturing capabilities. These results underline that optimisation-driven fibre paths, when constrained by physically meaningful manufacturing limits, can be directly translated into feasible winding patterns that balance structural efficiency with process robustness.

4. Conclusions

This study established a reliability-aware, multi-objective optimisation framework for variable-angle filament-wound (VAFW) composite

cylinders, explicitly embedding experimentally measured geometric imperfections into the design process. The framework couples experimentally informed imperfection characterisation with surrogate-assisted meta-optimisation, providing a scalable route for the design of imperfection-sensitive composite shells. By integrating principal component analysis (PCA) and Latin hypercube sampling (LHS) of DIC-measured imperfections with nonlinear FE modelling, Kriging surrogates, and multi-objective optimisation, the proposed approach achieves a robust synthesis of physical realism and computational efficiency.

The results confirm that the PCA–LHS reconstruction accurately reproduces the statistical features of measured imperfections, yielding knockdown factors consistent with the NASA/SP-8007-2020/REV-2 lower-bound trends. The surrogate models demonstrated high predictive fidelity under cross-validation, enabling rapid exploration of the tri-objective space defined by structural mass, nonlinear collapse load, and knockdown factor. Sensitivity analysis revealed that end-segment winding angles exert the dominant influence on buckling and collapse behaviour, reflecting the boundary-governed nature of filament-wound structures.

Comparative optimisation showed that Bayesian Optimisation (BO) achieved higher sampling efficiency and smoother convergence within a limited FE budget, whereas NSGA-II provided broader coverage of the design space through extensive population evolution. These results underline the complementarity of both strategies for navigating the high-dimensional, non-convex landscape of VAFW composite optimisation.

Collectively, the findings validate the proposed framework as an effective bridge between experimental imperfection data, surrogate modelling, and multi-objective meta-optimisation. The methodology enables informed trade-offs among weight efficiency, load-carrying capacity, and imperfection robustness, offering a practical pathway toward reliability-based design of advanced composite shells. Future work should aim to couple this approach with probabilistic modelling of manufacturing variability, material degradation, and damage evolution, thereby extending the reliability-centred paradigm toward qualification-ready aerospace structures. In addition, targeted experimental testing of selected Pareto-optimal designs manufactured under identical process constraints would provide a critical validation of the optimisation-predicted performance gains and further strengthen the link between computational design and practical deployment. Furthermore, extending the optimisation framework to alternative boundary conditions and loading scenarios (e.g., distributed surface loads or combined loading states) would allow exploration of a broader design space, as the optimal winding-angle ranges are inherently dependent on the assumed loading and support conditions.

Beyond methodological contributions, the framework delivers tangible value for engineering practice. The optimised VAFW cylinders achieve an effective balance between low mass, high buckling strength, and imperfection tolerance, which are key attributes for aerospace and pressure-vessel applications where stability margins are critical. By replacing thousands of costly nonlinear FE simulations with surrogate-assisted optimisation, the approach enables exploration of performance-robustness trade-offs that were previously computationally prohibitive. This capability positions the framework as a powerful design tool for integrating imperfection sensitivity into early-stage structural optimisation and certification workflows.

CRedit authorship contribution statement

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

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

SGP Castro's contributions to this work are funded by the innovation programme Luchtvaart in Transitie, which is co-funded by the Netherlands National Growth Fund. It is part of the 'H2Crash' project.

Data availability statement

The raw data required to reproduce the findings of this work are available for download from Zenodo through the following DOI: <https://doi.org/10.5281/zenodo.17577585>.

References

- [1] Pan H, Qu W, Yang D, Huang Q, Li J, Ke Y. Design and optimization of variable stiffness composite cylinders with the consideration of manufacturing interaction. *Appl Compos Mater* 2022;29(3):1249–73. <https://doi.org/10.1007/s10443-021-10007-y>
- [2] Coskun O, Turkmen HS. Pre-trained design optimization of variable stiffness composite cylinders modeled using bézier curves. *Struct Multidiscip Optim* 2023;66(2):41. <https://doi.org/10.1007/s00158-022-03480-x>
- [3] Kepple J, Herath M, Pearce G, Prusty G, Thomson R, Degenhardt R. Improved stochastic methods for modelling imperfections for buckling analysis of composite cylindrical shells. *Eng Struct* 2015;100:385–98. <https://doi.org/10.1016/j.engstruct.2015.06.013>
- [4] Hou Y, Fan H, Gu W. Buckling analysis of variable stiffness composite cylindrical shells under axial compression considering imperfection sensitivity. *Int J Struct Stab Dyn* 2024;2550226. <https://doi.org/10.1142/S0219455425502268>
- [5] Paudel A, Gupta S, Rowshan N, Thapa M, Mulani SB, Walters RW. Stochastic buckling of composite cylinder with geometric imperfection and global sensitivity analysis. In: AIAA SCITECH 2023 forum; 2023. p. 1093. <https://doi.org/10.2514/6.2023-1093>
- [6] Weingarten VI, Seide P, Peterson JP. Buckling of thin-walled circular cylinders, Special publication (sp), Hampton: NASA; 1968. <https://ntrs.nasa.gov/citations/19690013955>.
- [7] Hilburger MW. Buckling of thin-walled circular cylinders, Special publication (sp), Hampton: NASA; 2020. <https://ntrs.nasa.gov/citations/20205011530>.
- [8] Degenhardt R. New achievements in stability of composite aerospace structures. In: *Int. Conf. On vibrations and buckling*; 2016. <https://elib.dlr.de/114128/>.
- [9] Castro SGP, Almeida Jr. JHS, 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>
- [10] Almeida Jr. JHS, St-Pierre L, Wang Z, Ribeiro ML, Tita V, Amico SC, Castro SGP. 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>
- [11] Rouhi M, Ghayoor H, Fortin-Simpson J, Zacchia TT, Hoa SV, Hojjati M. Design, manufacturing, and testing of a variable stiffness composite cylinder. *Compos Struct* 2018;184:146–52. <https://doi.org/10.1016/j.compstruct.2017.09.090>
- [12] Clancy G, Peeters D, Oliveri V, Jones D, O'Higgins RM, Weaver PM. A study of the influence of processing parameters on steering of carbon fibre/peek tapes using laser-assisted tape placement. *Compos Part B Eng* 2019;163:243–51. <https://doi.org/10.1016/j.compositesb.2018.11.033>
- [13] Jamora VC, Wu KC, Kravchenko OG. Residual deformation analysis in composite shell structures manufactured using automated fiber placement. *Compos Struct* 2020;248:112482. <https://doi.org/10.1016/j.compstruct.2020.112482>
- [14] Lincoln R, Weaver P, Pirrera A, Groh RM. Manufacture and buckling test of a variable-stiffness, variable-thickness composite cylinder under axial compression. *American Institute of Aeronautics and Astronautics (AIAA)*; 2022. <https://doi.org/10.2514/6.2022-0664>
- [15] Lincoln R, Weaver PM, Pirrera A, Groh RM. Increasing reliability of axially compressed cylinders through stiffness tailoring and optimization. *Philos Trans R Soc A Math Phys Eng Sci* 2023;381(2244):20220034. <https://doi.org/10.1098/rsta.2022.0034>
- [16] Punera D, Mukherjee P. Recent developments in manufacturing, mechanics, and design optimization of variable stiffness composites. *J Reinforc Plast Compos* 2022;41(23–24):917–45. <https://doi.org/10.1177/07316844221082999>
- [17] Uhlig K, Tosch M, Bittrich L, Leipprand A, Dey S, Spickenheuer A, Heinrich G. Meso-scaled finite element analysis of fiber reinforced plastics made by tailored fiber placement. *Compos Struct* 2016;143:53–62. <https://doi.org/10.1016/j.compstruct.2016.01.049>
- [18] Souza GSC, dos Santos AL, Ávila Dias Maciel MM, Guedes RM, Viriato N, Amico SC, Tita V. Radial compression of thin-walled variable-angle tow filament-wound composite cylinders, Part I: manufacture, experimental, and computational analyses in the linear-elastic regime. *Compos Struct* 2025;366:119178. <https://doi.org/10.1016/j.compstruct.2025.119178>

- [19] Souza GSC, dos Santos AL, Ávila Dias Maciel MM, Guedes RM, Tavares P, Amico SC, Tita V. Radial compression of thin-walled variable-angle tow filament-wound composite cylinders, Part II: progressive crushing with and without stress concentrators. *Compos Struct* 2025;370:119364. <https://doi.org/10.1016/j.compstruct.2025.119364>
- [20] Rouhi M, Ghayoor H, Hoa SV, Hojjati M, Weaver PM. Stiffness tailoring of elliptical composite cylinders for axial buckling performance. *Compos Struct* 2016;150:115–23. <https://doi.org/10.1016/j.compstruct.2016.05.007>
- [21] Almeida Jr. JHS, 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>
- [22] Almeida Jr. JHS, Lisboa 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.compstruct.2022.106939>
- [23] 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.compstruct.2017.11.002>
- [24] Zhong J, Zheng Y, Chen J, Jing Z. Variable-stiffness composite cylinder design under combined loadings by using the improved kriging model. *Acta Mech Sinica* 2019;35(1):201–11. <https://doi.org/10.1007/s10409-018-0791-y>
- [25] Ma X, Wang F, Wang H, Hao P, Wang B. Post-buckling optimization of bending-induced variable stiffness composite cylinders considering worst geometric imperfections. *Thin-Walled Struct* 2021;169:108489. <https://doi.org/10.1016/j.tws.2021.108489>
- [26] Pelletier JL, Vel SS. Multi-objective optimization of fiber reinforced composite laminates for strength, stiffness and minimal mass. *Comput Struct* 2006;84(29):2065–80. <https://doi.org/10.1016/j.compstruct.2006.06.001>
- [27] Hong Z, Peeters D, Guo Y. Efficient strength optimization of variable stiffness laminates using lamination parameters with global failure index. *Comput Struct* 2022;271:106856. <https://doi.org/10.1016/j.compstruct.2022.106856>
- [28] Díaz J, Cid Montoya M, Hernández S. Efficient methodologies for reliability-based design optimization of composite panels. *Adv Eng Softw* 2016;93:9–21. <https://doi.org/10.1016/j.advengsoft.2015.12.001>
- [29] Guo Q, Hang J, Wang S, Hui W, Xie Z. Buckling optimization of variable stiffness composite cylinders by using multi-fidelity surrogate models. *Thin-Walled Struct* 2020;156:107014. <https://doi.org/10.1016/j.tws.2020.107014>
- [30] 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.compstruct.2023.107156>
- [31] Almeida Jr. JHS, Ashok A, Uzair M, Castro SGP. Bayesian meta-optimisation of variable stiffness composite cylinders for mass minimisation with manufacturing constraints. *Comput Struct* 2025;316:107868. <https://doi.org/10.1016/j.compstruct.2025.107868>
- [32] Wang Z, Almeida Jr. JHS, 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>
- [33] de Pison FM, Gonzalez-Sendino R, Aldama A, Ferreiro-Cabello J, Fraile-Garcia E. Hybrid methodology based on Bayesian optimization and ga-parsimony to search for parsimony models by combining hyperparameter optimization and feature selection. *Neurocomputing* 2019;354:20–6, recent Advancements in Hybrid Artificial Intelligence Systems. doi:<https://doi.org/10.1016/j.neucom.2018.05.136>
- [34] Yamashita T, Kino H, Tsuda K, Miyake T, Oguchi T. Hybrid algorithm of Bayesian optimization and evolutionary algorithm in crystal structure prediction. *Sci Technol Adv Mater Methods* 2022;2(1):67–74. <https://doi.org/10.1080/27660400.2022.2055987>
- [35] Castro SGP, Almeida Jr. JHS. VAFW cylinders 2020, S1, S2, S4, S8, stitched imperfections. Zenodo, 2021. <https://doi.org/10.5281/zenodo.4581164>
- [36] Castro SGP, Donadon MV, Guimarães TAM. ES-PIM applied to buckling of variable angle tow laminates. *Compos Struct* 2019;209:67–78. <https://doi.org/10.1016/j.compstruct.2018.10.058>
- [37] 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):140. <https://doi.org/10.1007/s00158-022-03227-8>
- [38] Composites TCR. Technical data sheet: Uf3357 TCR resin system, [Technical Data Sheet], [Cited: 2025-09-25], 2021. <https://www.tcrcomposites.com/lib/TDS-RD-0107-R001-UF3357.pdf>
- [39] Lu J, Zheng C, Wang L, Dai Y, Wang Z, Song Z. T700 carbon fiber/epoxy resin composite material hygrothermal aging model. *Materials* 2025;18(2). <https://doi.org/10.3390/ma18020369>
- [40] Santner TJ, Williams BJ, Notz WI. The design and analysis of computer experiments. 2nd ed. New York, NY: Springer; 2018. <https://doi.org/10.1007/978-1-4939-8847-1>
- [41] Gramacy RB. Surrogates: Gaussian process modeling, design, and optimization for the applied sciences. Chapman and Hall/CRC; 2020.
- [42] Rasmussen CE, Williams CKI. Gaussian processes for machine learning. The MIT Press; 2005. <https://doi.org/10.7551/mitpress/3206.001.0001>
- [43] Ulaganathan S, Couckuyt I, Ferranti F, Laermans E, Dhaene T. Performance study of multi-fidelity gradient enhanced kriging. *Struct Multidiscip Optim* 2015;51:1017–33. <https://doi.org/10.1007/s00158-014-1192-x>
- [44] Balandat M, Karrer B, Jiang DR, Daulton S, Letham B, Wilson AG, Bakshy E. BoTorch: a framework for efficient Monte-Carlo Bayesian optimization. *arXiv*, 2020. <https://doi.org/10.48550/arXiv.1910.06403>.