A Variable-Importance-Aware Universal Kriging Framework For Surrogate-Based Structural Optimization

Muhammad Hassan Yaqoob and Huang Hai*

Hangzhou International Innovation Institute, Regional Centre for Space Science and Technology Education in Asia and the Pacific, School of Astronautics, Beihang University, Beijing, 100191, China

* *Corresponding author. E-mail: hhuang@buaa.edu.cn*

Received: Mar. 01, 2026; Accepted: Apr. 06, 2026

Surrogate modelling is essential for simulation-based engineering optimization when high-fidelity simulations are computationally expensive. This paper presents Fuzzy-Based Weighted Universal Kriging (FBW-UK), a surrogate modelling framework designed to improve Kriging performance by explicitly incorporating variable importance information into the modelling pipeline. FBW-UK integrates three components: a data-driven fuzzy weighting mechanism based on Pearson-correlation importance scores, which amplifies influential variables and suppresses weakly informative ones through a sigmoidal membership function; a universal linear trend that captures global drift before Gaussian process regression; and an ARD Matérn kernel that provides per-dimension length-scale adaptation for residual modelling. The framework is validated on a CubeSat structural optimization problem using high-fidelity finite element simulations, with separate surrogates constructed for mass, stress, and deflection. Comparative results against Response Surface Methodology and Simple Kriging show that FBW-UK delivers consistently better predictive performance across responses and evaluation metrics, with the largest gains observed for the non-linear, constraint-critical stress response. The method also demonstrates stronger data efficiency, achieving competitive accuracy under limited-sample conditions and reducing the number of expensive simulations required for reliable surrogate construction. When integrated into a surrogate based optimization workflow, FBW-UK yields a lighter feasible structural design while satisfying stress and deflection constraints. The resulting optimum is physically interpretable, and the learned fuzzy weights align with engineering intuition by identifying the most influential design variables. Overall, the study shows that importance-informed input-space rescaling can improve surrogate accuracy, robustness, and optimization effectiveness in computationally intensive engineering design.

Keywords: Fuzzy-Based Weighted Universal Kriging (FBW-UK); Surrogate modeling; Gaussian process regression; Structural optimization; CubeSat; Finite element analysis; Differential Evolution

http://dx.doi.org/10.6180/jase.202609_32.046

1. Introduction

Aerospace structural design increasingly depends on high-fidelity finite element analysis (FEA) to verify performance under launch and operational loads [1]. However, each simulation can be computationally expensive, making direct optimization impractical when hundreds or thousands of evaluations are required [2]. Surrogate modeling addresses this bottleneck by replacing expensive simulations with fast approximations inside optimization loops [3, 4].

Among surrogate methods, Gaussian process regression (GPR), commonly termed Kriging in engineering design, is widely used because it provides both predictions and uncertainty estimates [5]. These uncertainty estimates support

Bayesian optimization, constrained acquisition strategies, and risk-aware design decisions [6, 7]. As a result, Kriging has been applied successfully in structural, aerodynamic, and multidisciplinary optimization [8].

Despite these advantages, Kriging performance is sensitive to numerical conditioning, hyperparameter estimation quality, and limited data [9]. These issues intensify in high-dimensional design spaces, where covariance conditioning worsens and overfitting risk increases [10]. In constrained optimization, surrogate error near active constraints can misclassify feasibility and degrade search quality [11].

Recent work has addressed these limitations through multiple directions. Dimensionality-reduction methods project high-dimensional inputs into lower-dimensional latent spaces before Kriging [12]. Multi-fidelity and hierarchical Kriging variants improve efficiency in data-scarce settings [13]. ARD length scales can provide post-hoc feature-importance signals useful for variable pruning [14]. Soft-computing studies in surrogate-assisted optimization also show that fuzzy mappings can transform scalar indicators into bounded, optimization-relevant weights [15]. However, these approaches generally treat variable importance as a reduction tool, efficiency aid, or interpretive byproduct, rather than embedding it directly into preprocessing to reshape the input space before covariance construction [16]. In particular, ARD Universal Kriging—which combines an OLS linear trend with an ARD Matérn kernel fitted to residuals—already provides per-dimension length-scale adaptation through marginal-likelihood optimisation alone, without any explicit variable-importance signal. ARD therefore serves as the natural ablation baseline for FBW-UK: any accuracy difference between the two isolates the contribution of the fuzzy preprocessing stage, independent of all other architectural choices shared by both models.

This limitation is especially relevant in structural design problems where all variables remain physically meaningful but contribute differently across responses. CubeSat structural optimization provides a representative case. The rapid growth and diversification of CubeSat missions has increased demand for efficient and reliable design methodologies [17], and recent conceptual design tools emphasize faster, better-informed early-stage decisions to reduce downstream risk [18]. At the subsystem level, FEA is routinely used to assess structural safety under launch environments [19]; experimental studies show that modeling assumptions strongly affect prediction quality [20]; and mission-relevant structural/thermal analyses continue to rely on integrated numerical modeling and validation workflows [21]. Recent studies on operational CubeSat hardware combine modal, quasi-static, buckling, and random-vibration analyses with test validation to verify launchworthiness [22], while new structural concepts aimed at mass efficiency further increase the need for accurate response prediction [23]. In such settings, mass may be dominated by panel thickness, while stress and deflection depend on different multi-variable interactions. A surrogate that does not distinguish these sensitivities can over-smooth critical directions and under-smooth less important ones, reducing accuracy near the regions most important for design decisions.

To address this gap, we propose Fuzzy-Based Weighted Universal Kriging (FBW-UK), a surrogate framework that embeds variable-importance information directly into the Kriging preprocessing stage. FBW-UK combines three components: (i) a data-driven fuzzy weighting mechanism that maps Pearson correlation scores through a sigmoidal transformation to amplify influential variables and suppress weakly informative ones; (ii) a universal linear trend that removes global drift before GP regression; and (iii) an ARD Matérn Gaussian process fitted to residuals in the weighted input space, providing a second layer of adaptation through per-dimension length scales. The framework is trained independently for each response, allowing different sensitivity structures to be modeled for mass, stress, and deflection.

The contributions of this work are threefold. First, we introduce a fuzzy weighting mechanism that embeds response-specific variable importance directly into Kriging preprocessing, enabling input-space rescaling before covariance construction. Second, we show that this two-level adaptation strategy (importance-informed rescaling followed by ARD length-scale learning) improves predictive accuracy, with FBW-UK achieving $R^2 = 0.9349$ for stress prediction and a 7.3 percentage point improvement over Simple Kriging. Scalability is further demonstrated on an independent 800-sample dataset under two physically motivated stress-filtering regimes, where fuzzy preprocessing maintains strong stress prediction ($R^2 = 0.856$) while ARD without fuzzy weighting degrades to $R^2 = 0.685$ when training-set size is small relative to input dimensionality. Third, we demonstrate stronger data efficiency, with FBW-UK reaching accuracy at 100 samples that competing methods attain only at 200–400 samples, corresponding to estimated computational savings of 10–35 hours per optimization study.

The remainder of this paper is organized as follows. Section 2 presents the FBW-UK methodology, training algorithm, and software implementation. Section 2.5 describes the CubeSat structural optimization problem, finite

element model, dataset generation, and evaluation protocols. Section 3 reports comparative prediction results, data-efficiency behavior, and optimization outcomes. Section 4 concludes with key findings, limitations, and directions for future research.

2. Methods

This section presents the mathematical formulation of the proposed Fuzzy-Based Weighted Universal Kriging (FBW-UK) surrogate. FBW-UK follows the universal Kriging formulation in which the response is decomposed into (i) an explicit parametric trend representing global drift and (ii) a zero-mean Gaussian process (GP) representing local deviations. The distinguishing step in FBW-UK is an *importance-informed input preconditioning* procedure: a response-specific weight vector is computed from inexpensive association statistics and used to deterministically rescale the standardised inputs before GP fitting. This rescaling is *not* introduced as a fundamentally new covariance family beyond ARD; rather, it provides (a) an interpretable summary of variable influence and (b) a practical preconditioning signal for length-scale estimation under limited data and nonconvex marginal-likelihood optimisation.

The model is trained independently for each response (e.g., mass, stress, and deflection), so the fuzzy weights and kernel hyperparameters are response-specific.

2.1. Gaussian Process Regression and Universal Kriging

Let $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^d$ denote the input vector and $y(\mathbf{x}) \in \mathbb{R}$ the scalar response. In universal Kriging, the latent response is expressed as

$$y(\mathbf{x}) = m(\mathbf{x}) + \eta(\mathbf{x}), \tag{1}$$

where $m(\mathbf{x})$ is a low-order trend and $\eta(\mathbf{x})$ is a zero-mean GP:

$$\eta(\mathbf{x}) \sim \mathcal{GP}(0, k(\mathbf{x}, \mathbf{x}')). \tag{2}$$

Assume observations

$$y_i = y(\mathbf{x}_i) + \epsilon_i, \qquad \epsilon_i \sim \mathcal{N}(0, \sigma_n^2),$$

with training data $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$. Let $\mathbf{y} = [y_1, \ldots, y_n]^\top$ and $\mathbf{K} \in \mathbb{R}^{n \times n}$ with entries $\mathbf{K}_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$. For a test point $\mathbf{x}_*$, define $\mathbf{k}_* = [k(\mathbf{x}_*, \mathbf{x}_1), \ldots, k(\mathbf{x}_*, \mathbf{x}_n)]^\top$.

Conditioning yields a Gaussian predictive distribution. We distinguish:

(i) Latent-function prediction for $y(\mathbf{x}_*)$:

$$\mu_y(\mathbf{x}_*) = m(\mathbf{x}_*) + \mathbf{k}_*^\top \left(\mathbf{K} + \sigma_n^2 \mathbf{I}\right)^{-1} (\mathbf{y} - \mathbf{m}), \tag{3}$$

$$\sigma_y^2(\mathbf{x}_*) = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\top \left(\mathbf{K} + \sigma_n^2 \mathbf{I}\right)^{-1} \mathbf{k}_*, \tag{4}$$

where $\mathbf{m} = [m(\mathbf{x}_1), \ldots, m(\mathbf{x}_n)]^\top$.

(ii) Noisy-observation prediction for a future observation $y_*^{\text{obs}} = y(\mathbf{x}_*) + \epsilon$:

$$\text{Var}[y_*^{\text{obs}} \mid \mathcal{D}] = \sigma_y^2(\mathbf{x}_*) + \sigma_n^2. \tag{5}$$

In implementation, linear solves of the form $(\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1}\mathbf{v}$ are computed using a Cholesky factorisation for numerical stability rather than explicit matrix inversion.

Kernel hyperparameters are estimated by maximising the log marginal likelihood. Because this optimisation is typically nonconvex, multi-start optimisation with random restarts is used.

2.1.1. Kernel Choice: Matérn with ARD

We use a Matérn covariance with automatic relevance determination (ARD):

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \, k_\nu\big(r(\mathbf{x}, \mathbf{x}')\big), \qquad r(\mathbf{x}, \mathbf{x}') = \sqrt{\sum_{j=1}^{d} \frac{(x_j - x_j')^2}{\ell_j^2}}, \tag{6}$$

where $\ell_j > 0$ are dimension-specific length scales, σ_f^2 is the signal variance, and ν controls smoothness.

We set $\nu = 5/2$, which yields sample paths that are twice mean-square differentiable, a reasonable smoothness assumption for many structural responses. The nugget/noise variance σ_n^2 is retained both to model observation noise (when present) and to improve numerical conditioning for deterministic simulators.

2.2. Motivation and Scope of the FBW-UK Weighting

FBW-UK is motivated by two practical issues in engineering surrogate modeling:

Heterogeneous variable influence. Engineering responses are often driven by a subset of inputs, while other inputs have weaker influence. An isotropic kernel is overly restrictive in such settings because it imposes a single correlation length across all dimensions. ARD relaxes this by assigning dimension-specific length scales, but ARD alone does not provide a simple, response-specific importance summary and can be sensitive to hyperparameter estimation under limited data.

Mean–covariance entanglement under limited data. When global drift is present, a constant-mean GP must simultaneously represent both global and local structure, which can reduce data efficiency. Universal Kriging mitigates this by modeling the global drift explicitly and reserving the GP for local correlated residuals.

FBW-UK retains universal Kriging and ARD, and adds a low-cost, interpretable weighting stage that provides response-specific preconditioning of the input geometry before GP fitting.

2.3. FBW-UK Formulation

Let $\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_n]^\top \in \mathbb{R}^{n \times d}$ and $\mathbf{y} \in \mathbb{R}^n$.

2.3.1. Step 1: Standardisation

Each input dimension is standardised to remove scale effects:

$$\tilde{x}_{ij} = \frac{x_{ij} - \mu_j}{\sigma_j}, \quad \mu_j = \frac{1}{n}\sum_{i=1}^{n} x_{ij}, \quad \sigma_j = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(x_{ij} - \mu_j)^2}. \tag{7}$$

The resulting standardised design matrix is denoted by $\tilde{\mathbf{X}}$.

2.3.2. Step 2: Correlation-Based Importance Scores (Low-Cost Screening)

For each standardised input $\tilde{x}_j$, an inexpensive main-effect association score with the response is computed using the absolute Pearson correlation:

$$r_j = \left|\mathrm{corr}(\tilde{x}_j, \mathbf{y})\right|. \tag{8}$$

These values are normalised into simplex-constrained importance scores:

$$f_j = \frac{r_j}{\sum_{k=1}^{d} r_k}, \qquad \sum_{j=1}^{d} f_j = 1. \tag{9}$$

Because Pearson correlation captures linear association, f_j is used only as a low-cost screening signal; nonlinear effects and interactions may not be fully represented. Remaining nonlinear structure is captured by the residual GP with ARD in Step 6.

2.3.3. Step 3: Fuzzy Membership Mapping to Weights

To obtain smooth, bounded importance weights, the importance scores are mapped through a sigmoidal fuzzy membership function:

$$\tilde{w}_j = \frac{1}{1 + \exp\{-\alpha(f_j - f_0)\}}, \tag{10}$$

where $\alpha > 0$ controls the steepness and f_0 defines the transition point. The sigmoid form is chosen because it maps any real-valued importance score to the unit interval with a smooth, monotonic transition, ensuring that weakly relevant variables are down-weighted but never entirely discarded. A natural default is $f_0 = 1/d$ (uniform-importance baseline), while α controls the sharpness of the transition between low- and high-importance regimes.

In the experiments reported in this study, α and f_0 were selected by grid-search cross-validation on the 414-sample dataset. A grid of $f_0 \in \{0.125, 0.25, 0.375, 0.5, 0.625\}$ and $\alpha \in \{1, 3, 5, 7, 10, 15\}$ (30 candidate pairs) was evaluated using 5-fold CV with stress R^2 as the selection criterion. The combination ($\alpha = 5,\ f_0 = 0.5$) achieved the highest cross-validated stress R^2. Performance varied by only ± 0.027 across the full grid, confirming robustness to the specific parameter choice.

The fuzzy weights are then normalised to sum to one:

$$w_j = \frac{\tilde{w}_j}{\sum_{k=1}^{d} \tilde{w}_k}, \qquad \sum_{j=1}^{d} w_j = 1. \tag{11}$$

The resulting vector $\mathbf{w} = [w_1, \ldots, w_d]^\top$ is the response-specific fuzzy importance vector.

2.3.4. Step 4: Linear Trend Fit and Residualisation

A first-order trend is fitted on the standardised inputs:

$$m(\tilde{\mathbf{x}}) = \beta_0 + \sum_{j=1}^{d} \beta_j \tilde{x}_j. \tag{12}$$

Let $\mathbf{F} = [\mathbf{1},\ \tilde{\mathbf{X}}] \in \mathbb{R}^{n \times (d+1)}$ and $\boldsymbol{\beta} = [\beta_0, \beta_1, \ldots, \beta_d]^\top$. The ordinary least-squares estimate is

$$\hat{\boldsymbol{\beta}} = (\mathbf{F}^\top \mathbf{F})^{-1} \mathbf{F}^\top \mathbf{y}. \tag{13}$$

Residuals are then computed as

$$\boldsymbol{\varepsilon} = \mathbf{y} - \mathbf{F}\hat{\boldsymbol{\beta}}. \tag{14}$$

This trend-plus-residual decomposition follows the standard regression–Kriging interpretation: the linear trend captures global drift, and the GP models local correlated deviations.

2.3.5. Step 5: Importance-Informed Input Rescaling

The standardised inputs are rescaled using the fuzzy weights:

$$\mathbf{z}_i = \mathrm{diag}(\mathbf{w})\, \tilde{\mathbf{x}}_i, \qquad i = 1, \ldots, n, \tag{15}$$

and the weighted design matrix is $\mathbf{Z} = [\mathbf{z}_1, \ldots, \mathbf{z}_n]^\top$.

Interpretation and identifiability. For an ARD kernel, the weighted distance becomes

$$r(\mathbf{z}, \mathbf{z}') = \sqrt{\sum_{j=1}^{d} \frac{(w_j(\tilde{x}_j - \tilde{x}'_j))^2}{\ell_j^2}} = \sqrt{\sum_{j=1}^{d} \frac{(\tilde{x}_j - \tilde{x}'_j)^2}{(\ell_j / w_j)^2}}. \tag{16}$$

Hence, the rescaling induces effective length scales $\bar{\ell}_j = \ell_j / w_j$. In principle, this does not enlarge the ARD kernel family; rather, it provides an explicit importance representation and a practical preconditioning effect for hyperparameter optimisation.

2.3.6. Step 6: GP on Residuals in Weighted Space

A zero-mean GP is fitted to $(\mathbf{Z}, \boldsymbol{\varepsilon})$:

$$\varepsilon(\mathbf{z}) \sim \mathcal{GP}\Big(0,\ k_{\text{M52-ARD}}(\mathbf{z}, \mathbf{z}') + \sigma_n^2 \delta(\mathbf{z}, \mathbf{z}')\Big). \tag{17}$$

For $\nu = 5/2$, the ARD Matérn kernel is written as

$$k_{\text{M52-ARD}}(\mathbf{z}, \mathbf{z}') = \sigma_f^2 \prod_{j=1}^{d} \left(1 + \frac{\sqrt{5}|z_j - z_j'|}{\ell_j} + \frac{5(z_j - z_j')^2}{3\ell_j^2}\right) \times \exp\left(-\frac{\sqrt{5}|z_j - z_j'|}{\ell_j}\right). \tag{18}$$

Hyperparameters $\boldsymbol{\theta} = \{\sigma_f^2, \ell_1, \ldots, \ell_d, \sigma_n^2\}$ are estimated by maximising the log marginal likelihood using multi-start optimisation with R random restarts (here $R = 15$).

2.3.7. *Step 7: Prediction*

For a new input $\mathbf{x}^*$, the same transformations are applied:

$$\tilde{\mathbf{x}}^* = \frac{\mathbf{x}^* - \boldsymbol{\mu}}{\boldsymbol{\sigma}}, \qquad \mathbf{z}^* = \text{diag}(\mathbf{w})\, \tilde{\mathbf{x}}^*. \tag{19}$$

The final predictor is the sum of the trend and GP correction:

$$\hat{y}(\mathbf{x}^*) = m(\tilde{\mathbf{x}}^*) + \hat{\varepsilon}(\mathbf{z}^*), \tag{20}$$

with predictive uncertainty given by the GP posterior in weighted space (latent variance as in Eq. (4), and observation variance as in Eq. (5) when required).

Algorithm 1 summarises the training and prediction workflow.

Algorithm 1. FBW-UK Training

1: **Input:** Training data $\{\mathbf{X}, \mathbf{y}\}$
2: **Output:** Trained surrogate $\hat{y}(\mathbf{x}^*)$
3: Standardize $\mathbf{X}$ to obtain $\tilde{\mathbf{X}}$ (Eq. 7)
4: Compute correlations r_j (Eq. 8) and importance scores f_j (Eq. 9)
5: Compute fuzzy weights $\mathbf{w}$ (Eqs. 10–11)
6: Form $\mathbf{F} = [\mathbf{1}, \tilde{\mathbf{X}}]$ and fit the OLS trend: $\hat{\boldsymbol{\beta}} = (\mathbf{F}^\top \mathbf{F})^{-1} \mathbf{F}^\top \mathbf{y}$ (Eq. 13)
7: Compute residuals: $\boldsymbol{\varepsilon} = \mathbf{y} - \mathbf{F}\hat{\boldsymbol{\beta}}$ (Eq. 14)
8: Form the fuzzy-weighted matrix $\mathbf{Z} = \tilde{\mathbf{X}} \cdot \text{diag}(\mathbf{w})$ (Eq. 15)
9: Fit the GP model: $\mathcal{GP}(0, k_{\text{M52-ARD}} + \sigma_n^2 \delta)$ to $(\mathbf{Z}, \boldsymbol{\varepsilon})$ via log-marginal-likelihood maximization (Eqs. 17, 18)
10: Store the learned parameters: $\{\boldsymbol{\mu}, \boldsymbol{\sigma}, \hat{\boldsymbol{\beta}}, \mathbf{w}, \mathbf{L}, \boldsymbol{\alpha}_{\text{GP}}, \boldsymbol{\theta}^*\}$
11: **For Prediction at $\mathbf{x}^*$:**
12: Transform the input: $\tilde{\mathbf{x}}^* \leftarrow (\mathbf{x}^* - \boldsymbol{\mu})/\boldsymbol{\sigma}$ (Eq. 19)
13: Apply fuzzy weights: $\mathbf{z}^* \leftarrow \tilde{\mathbf{x}}^* \odot \mathbf{w}$ (Eq. 19)
14: Predict the output: $\hat{y}(\mathbf{x}^*) \leftarrow m(\tilde{\mathbf{x}}^*) + \mathbf{k}_*^\top \boldsymbol{\alpha}_{\text{GP}}$ (Eq. 20)

2.4. Computational Complexity

For each restart, training is dominated by a Cholesky factorisation of the $n \times n$ covariance matrix, giving $\mathcal{O}(n^3)$ time complexity and $\mathcal{O}(n^2)$ memory complexity. With R random restarts, the total training complexity is $\mathcal{O}(R\, n^3)$.

For prediction at a new point, kernel evaluation costs $\mathcal{O}(nd)$, and posterior mean/variance evaluation is $\mathcal{O}(n)$ once the Cholesky factor is available (via triangular solves).

The fuzzy preprocessing stage unique to FBW-UK adds $\mathcal{O}(nd)$ for computing Pearson correlations and $\mathcal{O}(d)$ for the sigmoid mapping, both of which are negligible relative to the $\mathcal{O}(Rn^3)$ GP fitting cost. In practice, FBW-UK training is approximately 5–8% slower than Simple Kriging due to the additional Pearson-correlation computation and fuzzy-weight mapping. Both GP-based models are approximately 600× slower than RSM, but the total surrogate training time for all three responses (∼41 s) constitutes less than 0.03% of the finite element data-generation cost (∼44.8 h for 414 evaluations when considering only valid solutions). For the current problem ($d = 8$, $n = 414$), training is tractable on a standard desktop in seconds.

2.5. Application to CubeSat Structural Optimization

This section presents the CubeSat structural frame optimization problem used to demonstrate the FBW-UK surrogate framework introduced in Section 2.3. The objective is to minimize structural mass subject to stress and deflection constraints under quasi-static launch loading. The section summarizes the design variables, high-fidelity finite element data generation, surrogate-model application and evaluation protocol, and the surrogate-based optimization workflow.

2.5.1. *Optimization Problem Formulation*

The CubeSat frame geometry is parametrized by eight independent continuous design variables. Due to structural symmetry, $x_2 = x_1$ and $x_4 = x_3$, so the independent design vector is

$$\mathbf{x} = [x_1, x_3, x_5, x_6, x_7, x_8, x_9, x_{10}]^\top.$$

The variables represent plate thicknesses and column cross-sectional dimensions, with bounds determined by manufacturability and geometric constraints. Table 1 summarizes the design variables and their ranges.

For each design, the required responses are total mass $m(\mathbf{x})$, maximum von Mises stress $\sigma(\mathbf{x})$, and maximum deflection $\delta(\mathbf{x})$, obtained from high-fidelity finite element analysis (FEA). The constrained optimization problem is defined as

$$\min_{\mathbf{x} \in \mathcal{X}} m(\mathbf{x}) \quad \text{subject to} \quad \sigma(\mathbf{x}) \le \sigma_{\text{allow}},\ \delta(\mathbf{x}) \le \delta_{\text{allow}}, \tag{21}$$

where $\sigma_{\text{allow}} = 2.66 \times 10^8$ Pa and $\delta_{\text{allow}} = 5 \times 10^{-4}$ m. Direct optimization is computationally expensive because each FEA evaluation requires several minutes.

Table 1. Design variables for the CubeSat frame optimization problem.

Variable	Description	Lower bound	Upper bound
x_1	Sideplate thickness	8.0 mm	13.0 mm
x_3	Internal divider thickness	2.0 mm	7.0 mm
$x_5, \ldots, x_{10}$	Column cross-sectional dimensions (6 variables)	30.0 mm	70.0 mm

2.5.2. *Finite Element Model and Dataset Generation*

High-fidelity simulations were performed in ANSYS Mechanical for the CubeSat frame under quasi-static launch loading conditions. The parametric 3U CubeSat frame ($340 \times 100 \times 100$ mm) comprises side plates, internal divider panels, and rectangular cross-section columns. The structural material is Al 6061-T6 with elastic modulus $E = 68{,}900$ MPa, Poisson's ratio $\nu = 0.33$, density $\rho = 2{,}770\,\mathrm{kg\,m^{-3}}$, yield strength $\sigma_y = 276$ MPa, and ultimate tensile strength $\sigma_u = 310$ MPa. The worst-case quasi-static launch load of $10.8\,g$ was applied along the longitudinal (Z) axis, meanwhile along (X) axis and (Y) axis a launch load of $8.5\,g$ was applied with the base constrained to represent the deployer support interface. The rails along X and Y faces are modelled as flexible supports. The design criterion is an allowable stress of $\sigma_{\text{allow}} = 266$ MPa (factor of safety 1.04 on yield). A mesh-convergence study, requiring less than 1% change in peak stress and deflection over five successive refinements, was performed to ensure solution independence and a mesh size comprising of 259,123 nodes and 140,305 element has been finalized for the study.

A Latin Hypercube Sampling (LHS) design was generated with a fixed random seed, producing 993 initial sample points over the eight-dimensional design space. Non-converged and non-physical simulation results were screened out, yielding a final dataset of **414 valid designs**. Each FEA evaluation required 360–420 s (∼97 h total for the full campaign). This dataset includes both feasible and infeasible configurations, allowing the surrogates to learn response behavior in the interior region and near constraint boundaries.

2.5.3. *Data Preprocessing and Partitioning*

The dataset was partitioned using a fixed 70/30 split (random seed 42), yielding 289 training samples and 125 test samples. This split is used consistently for model comparison in Section 3. Input preprocessing within FBW-UK follows the standardisation procedure defined in Eq. 7; predictions are reported in physical units for interpretation.

2.5.4. *Surrogate Model Application and Evaluation*

The FBW-UK model was trained independently for each response (mass, stress, and deflection) following the formulation in Section 2.3. For each response, the procedure consists of: input standardisation (Eq. 7), Pearson-correlation-based importance scoring (Eqs. 8–9), fuzzy weight generation (Eqs. 10–11), linear trend fitting and residualisation (Eqs. 12–14), and GP regression on the weighted input space using the ARD Matérn-5/2 kernel (Eqs. 17–18). In the reported experiments, the fuzzy-membership parameters are selected via grid-search cross-validation ($\alpha = 5$, $f_0 = 0.5$), and GP hyperparameters are estimated using log-marginal-likelihood maximisation with 15 random restarts.

Model predictions are produced through the FBW-UK predictor in Eq. 20, with predictive uncertainty obtained from the GP posterior (Eqs. 4–5).

FBW-UK is evaluated against three benchmark surrogates trained on identical data splits, input preprocessing, and hyperparameter optimisation settings:

- **Response Surface Methodology (RSM).** A second-order polynomial in all inputs, fitted by ordinary least squares. No covariance matrix is formed and no uncertainty estimate is produced.
- **Simple Kriging (SK).** An isotropic Matérn-5/2 Gaussian process with a single shared length scale, a white-noise nugget, and normalised output. Ten random restarts are used for hyperparameter optimisation.
- **ARD Universal Kriging (ARD).** Shares the same two-stage architecture as FBW-UK—an OLS linear trend removal followed by an ARD Matérn-5/2 GP fitted to residuals—but *omits the fuzzy preprocessing stage* (Steps 2–3 and 5 in Algorithm 1). ARD operates on standardised inputs directly, relying solely on marginal-likelihood-based length-scale optimisation for per-dimension adaptation. It serves as the *ablation baseline*: any accuracy difference between ARD and FBW-UK isolates the contribution of fuzzy importance weighting, with all other modelling choices held constant.

All GP-based models (SK, ARD, FBW-UK) use 15 random restarts, a noise-level nugget σ_n^2, and L-BFGS-B log-marginal-likelihood maximisation.

2.5.5. *Model Evaluation Protocol*

Predictive accuracy is evaluated on withheld test data using standard regression metrics, including the coefficient of

Table 2. Test-set predictive accuracy of RSM, Simple Kriging (SK), and FBW-UK on the full 414-sample dataset (70/30 split; $n_{\text{test}} = 125$). Bold values indicate the best result for each response–metric combination.

Response	Metric	RSM	SK	FBW-UK
Deflection (mm)	R^2	0.9974	0.9983	**0.9986**
	RMSE (mm)	2.35×10^{-3}	1.87×10^{-3}	$\mathbf{1.78 \times 10^{-3}}$
	MAE (mm)	1.66×10^{-3}	1.38×10^{-3}	$\mathbf{1.23 \times 10^{-3}}$
	MAPE (%)	0.416	0.337	**0.296**
Stress (MPa)	R^2	0.8729	0.8615	**0.9349**
	RMSE (MPa)	11.26	12.22	**7.51**
	MAE (MPa)	8.72	9.15	**4.96**
	MAPE (%)	3.42	3.60	**1.98**
Mass (kg)	R^2	0.9982	0.9983	**0.9986**
	RMSE (kg)	1.02×10^{-2}	1.01×10^{-2}	$\mathbf{8.10 \times 10^{-3}}$
	MAE (kg)	7.69×10^{-3}	7.39×10^{-3}	$\mathbf{4.71 \times 10^{-3}}$
	MAPE (%)	0.635	0.605	**0.383**

determination (R^2), root-mean-square error (RMSE), mean absolute error (MAE), and mean absolute percentage error (MAPE). These metrics are used consistently across all three responses to assess both global accuracy and error magnitude.

To evaluate data efficiency, the models are also retrained on progressively larger subsets of the dataset ($n = 100$, 200, 300, and 414), using the same split protocol and random seed. This allows direct comparison of convergence behavior under data-scarce conditions representative of expensive FEA campaigns.

2.5.6. *Surrogate-Based Optimization*

After validation, the trained FBW-UK surrogates replace direct FEA evaluations in the optimization loop. The surrogate-based problem is solved using predicted responses $\hat{m}(\mathbf{x})$, $\hat{\sigma}(\mathbf{x})$, and $\hat{\delta}(\mathbf{x})$ under the same design bounds and constraints as Eq. 21. Symmetry conditions remain enforced through $x_2 = x_1$ and $x_4 = x_3$.

A hybrid global–local strategy is used: Differential Evolution (DE) performs global exploration of the design space, and the best candidate is subsequently refined using Sequential Least Squares Programming (SLSQP). Because surrogate evaluation is orders of magnitude faster than FEA, this strategy enables extensive search at negligible marginal cost. The optimization outcome, including convergence behavior, the final design, and constraint activity, is reported in Section 3.

3. Results and discussion

This section compares Response Surface Methodology (RSM), Simple Kriging (SK), and the proposed Fuzzy-Based Weighted Universal Kriging (FBW-UK) on the 414-sample FEA-generated structural dataset, and additionally evaluates ARD Universal Kriging (ARD) alongside these models on an independently generated 800-sample dataset. The analysis covers (i) test-set predictive accuracy, (ii) error distribution and residual behaviour, (iii) dataset-size sensitivity, (iv) scalability on the 800-sample dataset under two physically motivated stress-filtering regimes, and (v) surrogate-based structural mass optimisation. Unless otherwise stated, results use the fixed 70/30 train–test split (289 training, 125 test samples, random seed 42).

3.1. Overall Predictive Accuracy

Table 2 reports test-set performance for deflection (mm), von Mises stress (MPa), and mass (kg) using R^2, RMSE, MAE, and MAPE. Figure 1 shows parity plots for all model–response combinations.

3.2. Deflection Prediction

All three models predict deflection with high accuracy ($R^2 \geq 0.9974$). FBW-UK is best across all four metrics, but the margins over SK and RSM are small (Table 2), consistent with the relatively smooth response behavior. Figure 1 (top row) shows only minor visual separation, with FBW-UK slightly tighter around the 45° line.

3.3. Stress Prediction

Stress provides the clearest separation among models. FBW-UK achieves $R^2 = 0.9349$, compared with 0.8729 (RSM) and 0.8615 (SK), corresponding to improvements of 6.2 and 7.3 percentage points, respectively. RMSE decreases from 11.26 MPa (RSM) and 12.22 MPa (SK) to 7.51 MPa (FBW-UK), and MAPE decreases from 3.42–3.60% to 1.98%.

Figure 1 (middle row) shows reduced systematic deviation for FBW-UK relative to the benchmarks. The stronger

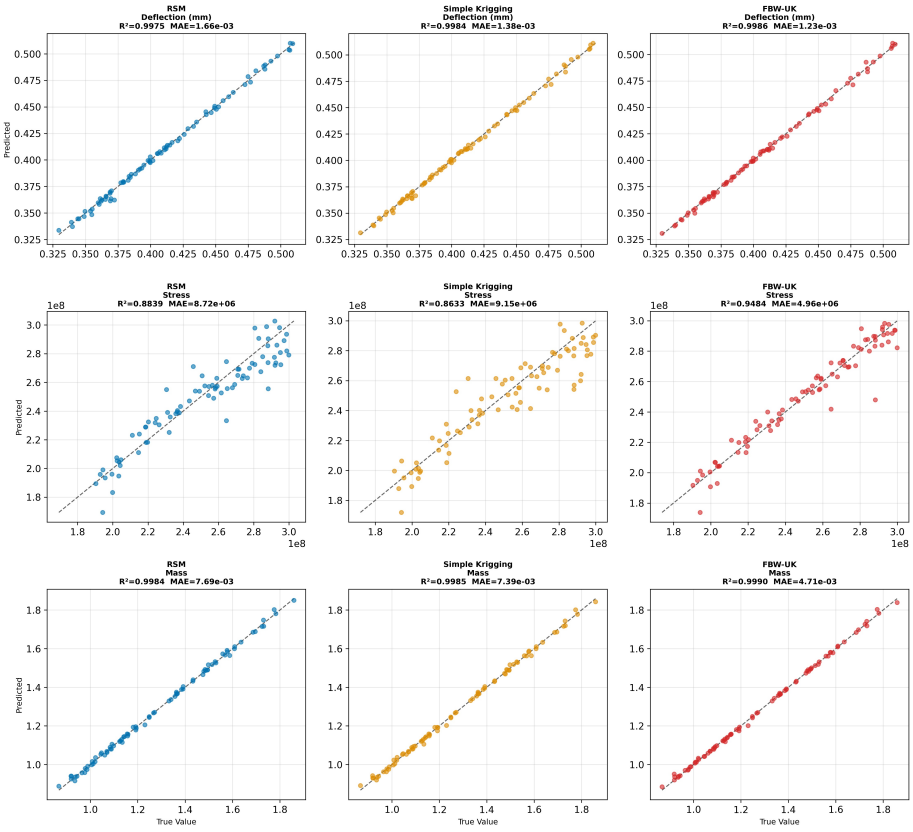

Fig. 1. Parity plots (predicted vs. actual) for all three models (columns) across the three structural responses (rows). Perfect predictions lie on the 45° line. FBW-UK points cluster tighter to the diagonal, particularly for stress where systematic bias in RSM and SK is reduced.

gains for stress are consistent with the more nonlinear and anisotropic response behavior. Furthermore, when absolute stress errors are examined as a function of input variables, RSM and SK errors concentrate at extreme values of panel thickness x_3 (the variable with the strongest nonlinear coupling to stress), whereas FBW-UK distributes errors more uniformly across the variable range, indicating that the fuzzy weighting mechanism reduces region-specific prediction bias.

3.4. Mass Prediction

Mass is predicted with very high accuracy by all models ($R^2 \geq 0.9982$). FBW-UK again gives the best performance across all metrics. Visual differences in Figure 1 (bottom row) are small, but Table 2 shows consistently lower error for FBW-UK (e.g., RMSE 0.00810 kg vs. $\approx$ 0.0101–0.0102 kg for SK/RSM).

3.5. Aggregate Metric Ranking

Across all 12 response–metric combinations in Table 2, FBW-UK attains the best value. The strongest separation appears for stress, while differences are smaller for deflection and mass.

3.6. Error Distribution and Residual Characteristics

To complement scalar metrics, error distributions were examined using absolute-error boxplots, residual violin plots, and cumulative distribution functions (supplementary material). The main observations are:

- **Median errors:** The median-error ranking matches RMSE/MAE. FBW-UK has the lowest medians for all responses, with the largest separation for stress.
- **Tail behavior:** For stress, FBW-UK shows smaller upper-tail absolute errors (75th–95th percentiles), indicating fewer large prediction errors in locally sensitive regions.
- **Residual symmetry:** For deflection and mass, residuals are approximately symmetric and centered near zero for all models. For stress, RSM and SK show mild skewness, whereas FBW-UK residuals are more symmetric.

- **Multi-metric consistency:** Radar-chart summaries (supplementary material) place FBW-UK at or near the best normalized performance across response–metric axes.

- **Spatial error patterns:** When absolute errors are plotted against each input variable, heteroscedastic patterns emerge. RSM and SK stress errors concentrate at extreme values of panel thickness x_3, whereas FBW-UK distributes errors more uniformly, consistent with the fuzzy weighting assigning higher importance to x_3.

Figure 2 shows sample-wise absolute errors for all three outputs. FBW-UK exhibits tighter clustering and fewer large outliers than RSM and SK, consistent with the aggregate metrics.

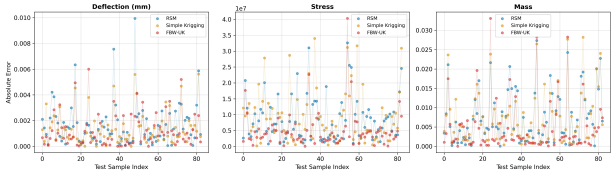

Fig. 2. Absolute error distribution across test samples for each model. Individual dots represent prediction errors for test samples, organized by output variable (deflection, stress, and mass).

3.7. Dataset Size Sensitivity and Model Convergence

To assess performance under limited-data conditions, the models were retrained on random subsets of size $n = 100$, 200, 300, and 414. For each subset, the same 70/30 train–test split (fixed random seed) was used, yielding training sets of 70, 140, 210, and 289 samples, respectively. Figure 3 shows convergence trends across all four metrics, and Table 3 reports the corresponding R^2 values.

The convergence trends in Figure 3 and the R^2 values in Table 3 show improvement with increasing dataset size for all models, with FBW-UK performing best in most cases, especially for stress.

3.8. Stress Convergence — The Critical Discriminator

The strongest separation occurs for stress at small sample sizes. At $n = 100$, FBW-UK achieves $R^2 = 0.8755$, compared with 0.7759 (SK) and 0.3894 (RSM), corresponding to gains of 11.5 and 48.7 percentage points, respectively. FBW-UK remains the best model for stress at all tested sizes (Table 3), indicating higher sample efficiency for the most nonlinear and constraint-relevant response.

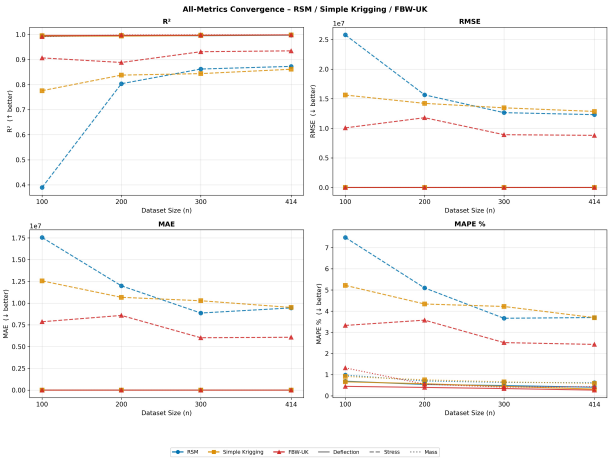

Fig. 3. Convergence curves showing all four accuracy metrics versus training-set size for all models and responses. FBW-UK consistently achieves the best (or tied-best) values across most sizes, with the largest advantage for stress.

3.9. Deflection and Mass Convergence

Deflection and mass converge rapidly for all models, with $R^2 > 0.99$ already at $n = 100$. FBW-UK is best at most sizes, with one exception for mass at $n = 100$ (SK: 0.9956 vs. FBW-UK: 0.9936). This difference disappears by $n = 200$, after which FBW-UK becomes best or tied-best for both responses.

3.10. Convergence Rate and Practical Implications

For stress, the FBW-UK curve increases rapidly from $n = 100$ to $n = 300$ and then shows smaller gains from $n = 300$ to $n = 414$, indicating diminishing returns beyond approximately 300 samples for this dataset. RSM and SK continue improving with added data but remain below FBW-UK.

Given an FEA runtime of 360–420 s per sample, these sample-size differences translate directly into additional computational cost. Concretely, FBW-UK at $n = 100$ already exceeds the stress-prediction accuracy that SK achieves at $n = 200$ (Table 3), eliminating the need for 100 additional FEA evaluations and saving approximately 10–12 hours of wall-clock time. The results therefore indicate that FBW-UK can reach a given stress-prediction accuracy with fewer simulations than RSM or SK.

3.11. Scalability on an 800-Sample Dataset

To assess model performance beyond the curated 414-sample dataset and to evaluate the contribution of fuzzy preprocessing relative to ARD-based length-scale learning alone, all four models (RSM, SK, ARD, and FBW-UK)

Table 3. R^2 on the 30% withheld test set as a function of dataset size (with a fixed 70/30 split). Bold values indicate the best model at each size–response combination.

Response	Model	$n = 100$	$n = 200$	$n = 300$	$n = 414$
Deflection (mm)	RSM	0.9926	0.9953	0.9954	0.9974
	SK	0.9938	0.9939	0.9961	0.9983
	FBW-UK	**0.9977**	**0.9966**	**0.9971**	**0.9986**
Stress (MPa)	RSM	0.3894	0.8038	0.8623	0.8729
	SK	0.7759	0.8381	0.8440	0.8615
	FBW-UK	**0.8755**	**0.8860**	**0.9279**	**0.9349**
Mass (kg)	RSM	0.9950	0.9973	0.9981	0.9982
	SK	0.9956	0.9969	0.9980	0.9983
	FBW-UK	0.9936	**0.9981**	**0.9987**	**0.9986**

were evaluated on an independently generated 800-sample Latin Hypercube design spanning the full ten-variable design space (including x_2 and x_4, which were eliminated by structural symmetry in the curated set). Two physically motivated filtering regimes were applied based on the Al 6061-T6 material properties:

- **Yield filter ($\sigma < 276$ MPa):** retains 312 samples (218 train / 94 test), corresponding to strictly elastic behaviour below the 0.2% proof stress.
- **UTS (Ultimate Tensile Strength) filter ($\sigma < 310$ MPa):** retains 447 samples (312 train / 135 test), permitting limited post-yield plasticity up to the tensile strength.

Table 4 reports test-set R^2 for all four models under both filtering regimes.

Several observations emerge from Table 4.

First, FBW-UK leads for stress prediction under yield filtering ($R^2 = 0.856$), exceeding ARD ($R^2 = 0.685$), SK ($R^2 = 0.845$), and RSM ($R^2 = 0.831$). Under UTS filtering, FBW-UK and ARD are effectively tied ($R^2 = 0.924$), both outperforming RSM and SK.

Second, ARD degrades markedly for stress under the yield filter ($R^2 = 0.685$), falling below all competitors including RSM and SK. This degradation arises from two compounding factors: (i) the 800-sample set uses all ten design variables, whereas the curated set uses eight (structural symmetry removes x_2 and x_4); and (ii) yield filtering retains only 218 training samples, yielding approximately 21.8 samples per input dimension. In this data-scarce, high-dimensional regime, marginal-likelihood-based ARD length-scale estimation becomes unreliable without the preconditioning signal provided by the fuzzy weighting stage. Under the UTS filter—where 312 training samples ($\approx$31 per dimension) are available—ARD recovers to $R^2 = 0.924$, matching FBW-UK.

Third, for deflection and mass, RSM and SK are competitive with the GP-based methods. Deflection is dominated by panel thickness x_3 with a near-linear sensitivity, and mass is a smooth algebraic function of geometry; neither response requires the nonlinear per-dimension adaptation that ARD and FBW-UK provide.

These results identify a practical boundary condition for fuzzy preprocessing: *the improvement over ARD is greatest when the target response is nonlinear, input dimensionality is high relative to training-set size, and reliable Pearson correlations can be estimated from the available data.* When these conditions are not met—as for deflection and mass with ample training data—simpler surrogates are competitive.

3.12. Surrogate-Based Structural Mass Optimization

Figure 4 shows the best predicted mass over 1,000 DE generations. The search improves rapidly in early generations, then transitions to slower refinement and a plateau near convergence. The best predicted mass falls below the best feasible design in the 414-sample dataset (0.8674 kg) early in the run, indicating that surrogate-based optimization identifies designs superior to those present in the initial space-filling sample. Representative milestones are visible in Figure 4: the best mass drops below the baseline by roughly 50 iterations, reaches 0.7692 kg by 200 iterations, and converges near 0.7672 kg by the end of the run.

3.12.1. Optimal Design and Feature Importance

The final SLSQP-refined optimum is listed in Table 5. The predicted minimum mass is 0.7672 kg. The deflection constraint is active at the optimum ($\hat{\delta} = 0.5000$ mm), whereas the stress constraint remains inactive with an 88.9 MPa margin, indicating a stiffness-governed optimum.

Table 6 reports the fuzzy feature weights for the mass surrogate. Panel thickness x_3 has the largest weight ($w = 0.431$), followed by x_1 ($w = 0.163$); the remaining variables contribute smaller shares. This ranking is consistent with

Table 4. Test-set R^2 on the 800-sample dataset under two stress-filtering regimes (70/30 split, seed 42). Bold indicates the best model for each response–filter combination.

Filter	Response	RSM	SK	ARD	FBW-UK
Yield ($\sigma < 276$ MPa)	Deflection	**0.9963**	0.9956	0.9893	0.9585
	Stress	0.8310	0.8448	0.6847	**0.8563**
	Mass	0.9979	0.9980	**0.9983**	**0.9983**
UTS ($\sigma < 310$ MPa)	Deflection	0.9965	**0.9976**	0.9900	0.9574
	Stress	0.8872	0.8670	0.9239	**0.9240**
	Mass	0.9984	0.9979	**0.9989**	**0.9989**

Table 5. Optimal design variables, predicted responses, and FEA verification at the optimum.

Variable / Response	Value	FEA	Bound or limit	Status
x_1	10.743 mm	10.743 mm	[8, 13]	Interior
x_3	**2.000 mm**	2.000 mm	[2, 7]	**At lower bound (LB)**
x_5	70.000 mm	70.000	[30, 70]	At upper bound (UB)
x_6	70.000 mm	70.000	[30, 70]	At UB
x_7	60.485 mm	60.485	[30, 70]	Interior
x_8	70.000 mm	70.000	[30, 70]	At UB
x_9	70.000 mm	70.000	[30, 70]	At UB
x_{10}	70.000 mm	70.000	[30, 70]	At UB
Mass $\hat{m}$	**0.7672 kg**	**0.7601 kg**	—	Objective
Stress $\hat{\sigma}$	177.1 MPa	168.2 MPa	$\leq$ 266 MPa	Margin: +88.9 MPa
Deflection $\hat{\delta}$	**0.5000 mm**	**0.485 mm**	$\leq$ 0.500 mm	**Binding**

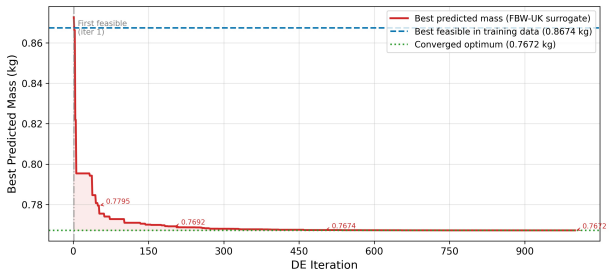

Fig. 4. Mass convergence history during Differential Evolution. Solid red line: best predicted mass (kg) at each of 1,000 iterations. Blue dashed line: best feasible mass in the 414-sample LHS training set (0.8674 kg). Green dotted line: converged surrogate optimum (0.7672 kg). Shaded region: mass improvement attributable to surrogate-guided global search.

the optimized design, where x_3 reaches its lower bound and several section dimensions move to upper bounds.

Tables 5 and 6 jointly summarize the optimized design and the corresponding fuzzy feature-weight distribution for the mass surrogate.

3.12.2. Design Space Topology

Surrogate contour analysis in the x_1–x_3 plane (with x_5–x_{10} fixed at their optimal values) shows that mass decreases primarily with decreasing x_3, while the deflection boundary governs the optimum. Stress varies with both x_1 and x_3 but remains below the allowable limit at the optimum, consistent with the inactive stress constraint in Table 5. These trends indicate a stiffness-controlled rather than strength-controlled optimum.

3.12.3. Optimization Outcome

The surrogate-based workflow (FBW-UK trained on 414 LHS samples, followed by DE and SLSQP polishing) achieves a minimum structural mass of **0.7672 kg**, an improvement of **0.1002 kg** (**11.55%**) relative to the best feasible design in the 414-sample dataset (0.8674 kg). The optimum is deflection-controlled ($\hat{\delta} = 0.500$ mm) with an inactive stress constraint and a margin of 88.9 MPa (33.4% below the allowable limit).

3.12.4. FEA Verification of the Surrogate Optimum

To independently verify the surrogate-predicted optimum, the ANSYS finite element model was re-evaluated at the exact optimal design point listed in Table 5. Figure 5 presents the mesh discretisation, von Mises equivalent stress distribution (maximum 168.2 MPa), and total deflection field (maximum 0.485 mm) at the optimal configuration.

The FEA-verified responses are compared with FBW-UK surrogate predictions in Table 5. The surrogate overestimates all three responses: mass by 0.93%, stress by

Table 6. Fuzzy feature weights for the mass surrogate.

Variable	Correlation $\lvert r\rvert$	Feature importance	Fuzzy weight w_j
x_1	0.416	0.241	0.163
x_3	**0.958**	**0.554**	**0.431**
x_5	0.088	0.051	0.073
x_6	0.035	0.020	0.063
x_7	0.004	0.003	0.058
x_8	0.066	0.038	0.069
x_9	0.075	0.044	0.070
x_{10}	0.088	0.051	0.073

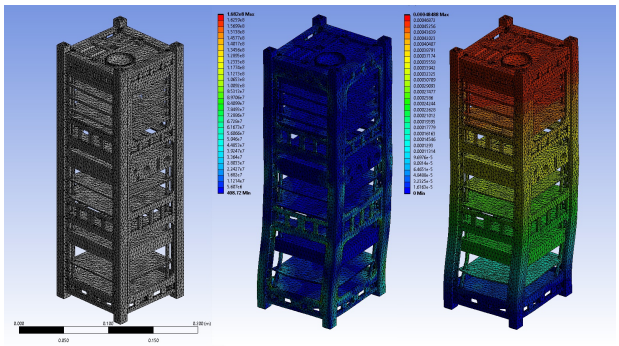

Fig. 5. FEA verification of the surrogate-predicted optimum in ANSYS Workbench. Left: mesh discretisation of the 3U CubeSat structure. Centre: von Mises equivalent stress contour. Right: total deflection contour

5.3%, and deflection by 3.1%. These deviations are consistent with the conservative bias commonly observed in Kriging-based surrogates near active constraint boundaries, where the posterior mean tends to smooth local response variations [24]. Notably, the FEA-verified deflection (0.485 mm) lies 3.0% below the 0.500 mm constraint limit, indicating that the deflection constraint is not truly binding at the optimal design point. This residual margin confirms that additional mass reduction beyond the current surrogate-predicted optimum is physically attainable—the FEA-confirmed mass of 0.7601 kg (0.93% lower than the surrogate prediction) corroborates this observation. The verified stress (168.2 MPa) remains 36.8% below the 266 MPa yield allowable, confirming that the design is stiffness-governed even after high-fidelity verification.

3.13. Software Implementation

A standalone Windows desktop application was developed to implement the full FBW-UK workflow (data import, model training, prediction, uncertainty quantification, and result visualization) through a graphical user interface requiring no programming expertise or prior Python installation. The software is written in Python (`tkinter`, `numpy`, `scipy`, `pandas`, `openpyxl`, `matplotlib`), packages the FBW-UK algorithm in an executable format, and supports model diagnostics, interactive plotting, batch prediction, and export of results.

4. Conclusion

The following key findings summarize the conclusions of this study:

- **FBW-UK Framework:** Fuzzy-Based Weighted Universal Kriging (FBW-UK) enhances surrogate modeling by combining correlation-based fuzzy feature weighting, a universal linear trend to remove global drift, and an ARD Matérn Gaussian process for residuals.

- **Evaluation and Comparison:** The framework was evaluated on CubeSat structural optimisation, predicting mass, stress, and deflection from high-fidelity finite element simulations. On the curated 414-sample dataset, FBW-UK outperformed Response Surface Methodology (RSM) and Simple Kriging (SK) in terms of predictive accuracy, particularly for stress prediction. On a separately generated 800-sample dataset spanning the full ten-variable design space, ARD Universal Kriging (ARD)—sharing all FBW-UK components except fuzzy preprocessing—was additionally benchmarked as an ablation baseline. Under yield-stress filtering, FBW-UK maintains stress $R^2 = 0.856$ while ARD drops to $R^2 = 0.685$, demonstrating that likelihood-based length-scale optimisation alone becomes unreliable in data-scarce, high-dimensional regimes without fuzzy preconditioning. Under UTS filtering, ARD recovers to $R^2 = 0.924$, matching FBW-UK, confirming that fuzzy preprocessing is most beneficial when training-set size is small relative to input dimensionality. FBW-UK also demonstrated stronger data efficiency by reducing the need for extensive simulations to achieve high accuracy.

- **Optimization Integration:** When integrated into a surrogate-based optimization workflow, FBW-UK produced a lighter structural design while satisfying stress and deflection constraints. The optimized design was physically interpretable, with panel thickness identified as the dominant mass driver, aligning with engineering expectations.

- **Generalizability:** The current validation is based on a single structural application (CubeSat frame optimization). Future work should extend the evaluation to additional engineering applications with diverse response landscapes to strengthen generality claims.

- **Limitations:** The study has several limitations that should be acknowledged. The fuzzy weighting stage relies on Pearson correlation, which captures only linear associations; strongly nonlinear or interaction-dominated importance structures may not be fully represented by this screening step. The preconditioning advantage over ARD diminishes as training-set size increases relative to input dimensionality, as observed in the UTS-filtered 800-sample results. The current validation is limited to a single structural application and standard test functions; broader engineering case studies are needed to confirm generality. Furthermore, surrogate prediction errors near active constraint boundaries can exclude feasible designs from the solution space. As demonstrated by the FEA verification (Table 5), the surrogate overestimates deflection at the optimum (0.500 mm predicted vs. 0.485 mm verified), causing the optimizer to treat the deflection constraint as binding when a 3.0% margin actually exists. Candidate designs with true responses marginally below constraint limits may thus be incorrectly classified as infeasible, effectively shrinking the explorable feasible region and preventing the optimizer from reaching the true global minimum [24]. While this conservatism is favorable from a structural safety perspective, infill strategies or constraint-relaxation techniques that account for surrogate prediction uncertainty could recover some of these excluded candidates.

- **Future Directions:** Future work will focus on developing nonlinear and interaction-aware importance measures (e.g., distance correlation, mutual information) to extend the fuzzy weighting beyond linear screening. Per-response adaptive tuning of the sigmoid parameters (α, f_0), potentially via Bayesian optimisation or response-specific CV, would allow the membership function to adapt to the sensitivity landscape of each output. Validation across additional engineering applications with diverse response landscapes will strengthen generality claims. Integration of adaptive sampling or active learning techniques with the FBW-UK framework may further enhance data efficiency. Finally, extension to multi-fidelity settings, where low-cost simulations provide initial importance estimates for high-fidelity surrogate construction, represents a promising direction.

Data availability

The compiled executable is publicly available at https://github.com/muhammadhassanyaqoob - commits/FBW-UK. The repository includes the FBW-UK core library, the desktop GUI application, extended dataset, and documentation for reproducing the surrogate results.

Acknowledgment

The authors would like to thank the Regional Centre for Space Science and Technology Education in Asia and the Pacific (China), affiliated with the United Nations, for providing technical resources for this research.

References

[1] R. Yondo, E. Andrés, and E. Valero, (2018) *"A review on design of experiments and surrogate models in aircraft real-time and many-query aerodynamic analyses"* **Progress in Aerospace Sciences *96***: 23–61. DOI: 10.1016/j.paerosci.2017.11.003.

[2] S. G. Kontogiannis and M. A. Savill, (2020) *"A generalized methodology for multidisciplinary design optimization using surrogate modelling and multifidelity analysis"* **Optimization and Engineering *21***: 723–759. DOI: 10.1007/s11081-020-09504-z.

[3] A. J. Keane and I. I. Voutchkov, (2020) *"Robust design optimization using surrogate models"* **Journal of Computational Design and Engineering *7***(1): 44–55. DOI: 10.1093/jcde/qwaa005.

[4] U. Kizhakkinan, P. L. T. Duong, R. Laskowski, G. Vastola, D. W. Rosen, et al., (2023) *"Development of a surrogate model for high-fidelity laser powder-bed fusion using tensor train and gaussian process regression"* **Journal of Intelligent Manufacturing *34***: 369–385. DOI: 10.1007/s10845-022-02038-4.

[5] D. Zhan and H. Xing, (2020) *"Expected improvement for expensive optimization: A review"* **Journal of Global Optimization *78***(3): 507–544. DOI: 10.1007/s10898-020-00923-x.

[6] B. Lei, T. Q. Kirk, A. Bhattacharya, D. Pati, X. Qian, R. Arroyave, B. K. Mallick, et al., (2021) *"Bayesian optimization with adaptive surrogate models for automated experimental design"* **npj Computational Materials** *7*: 194. DOI: 10.1038/s41524-021-00662-x.

[7] H. Røstum, S. Gros, and K. Aas-Jakobsen, (2025) *"Constrained Bayesian optimization for engineering bridge design"* **Structural and Multidisciplinary Optimization** *68*: 20. DOI: 10.1007/s00158-024-03951-3.

[8] A. Keane and I. Voutchkov, (2020) *"Surrogate approaches for aerodynamic section performance modelling"* **AIAA Journal** *58*(1): 16–24. DOI: 10.2514/1.J058687.

[9] J. Ollar, C. Mortished, R. Jones, J. Sienz, and V. Toropov, (2017) *"Gradient based hyper-parameter optimisation for well conditioned kriging metamodels"* **Structural and Multidisciplinary Optimization** *55*: 2029–2044. DOI: 10.1007/s00158-016-1626-8.

[10] M. A. Bouhlel, N. Bartoli, A. Otsmane, and J. Morlier, (2016) *"Improving Kriging surrogates of high-dimensional design models by partial least squares dimension reduction"* **Structural and Multidisciplinary Optimization** *53*(5): 935–952. DOI: 10.1007/s00158-015-1395-9.

[11] M. A. Bouhlel, N. Bartoli, R. G. Regis, A. Otsmane, and J. Morlier, (2018) *"Efficient global optimization for high-dimensional constrained problems by using the Kriging models combined with the partial least squares method"* **Engineering Optimization** *50*(12): 2038–2053. DOI: 10.1080/0305215X.2017.1419344.

[12] Y. Ge, J. Shi, Y. Li, and J. Shen, (2022) *"An efficient Kriging modeling method based on multidimensional scaling for high-dimensional problems"* **Algorithms** *15*(1): 3. DOI: 10.3390/a15010003.

[13] Y. He and J. Luo, (2024) *"Efficient Hierarchical Kriging Modeling Method for High-dimension Multi-fidelity Problems"* **Chinese Journal of Mechanical Engineering** *37*: 151. DOI: 10.1186/s10033-024-01136-z.

[14] Y. Deng, M. R. Eden, and S. Cremaschi, (2024) *"A Gaussian process embedded feature selection method based on automatic relevance determination"* **Computers & Chemical Engineering** *191*: 108852. DOI: 10.1016/j.compchemeng.2024.108852.

[15] I. Cruz-Vega, C. A. Reyes-García, H. J. Escalante, J. d. J. Rangel-Magdaleno, and J. M. Ramírez-Cortés, (2018) *"Surrogate modeling based on granular models and fuzzy aptitude functions"* **Applied Soft Computing** *65*: 21–32. DOI: 10.1016/j.asoc.2017.12.016.

[16] S. Li, S. Yuan, S. Liu, J. Wen, and Q. Huang, (2022) *"Research on an Accuracy Optimization Algorithm of Kriging Model Based on a Multipoint Filling Criterion"* **Mathematics** *10*(9): 1548. DOI: 10.3390/math10091548.

[17] T. Villela, C. A. Costa, A. M. Brandão, F. T. Bueno, and R. Leonardi, (2019) *"Towards the thousandth CubeSat: A statistical overview"* **International Journal of Aerospace Engineering** *2019*: 5063145. DOI: 10.1155/2019/5063145.

[18] C. Girardello, M. Tajmar, and C. Scharlemann, (2024) *"GREATCUBE+: conceptual design tool for CubeSat's design"* **CEAS Space Journal** *16*: 375–392. DOI: 10.1007/s12567-023-00509-9.

[19] F. T. Al-Maliky and M. J. AlBermani, (2018) *"Structural Analysis of Kufasat Using Ansys Program"* **Artificial Satellites** *53*(1): 29–35. DOI: 10.2478/arsa-2018-0003.

[20] A. N. Alhammadi, F. Jarrar, M. Al-Shaibah, A. Almesmari, T. Vu, A. Tsoupos, and P. Marpu, (2021) *"Effect of finite element model details in structural analysis of CubeSats"* **CEAS Space Journal** *13*(2): 231–246. DOI: 10.1007/s12567-020-00339-z.

[21] Y.-K. Park, G.-N. Kim, and S.-Y. Park, (2021) *"Novel Structure and Thermal Design and Analysis for CubeSats in Formation Flying"* **Aerospace** *8*(6): 150. DOI: 10.3390/aerospace8060150.

[22] A. Elshaal, M. Okasha, E. Sulaeman, A. H. Jallad, W. F. Aizat, and A. B. Alzubaidi, (2024) *"Structural Analysis of AlAinSat-1 CubeSat"* **The Egyptian Journal of Remote Sensing and Space Sciences** *27*(3): 532–546. DOI: 10.1016/j.ejrs.2024.06.006.

[23] G. Capovilla, E. Cestino, and L. Reyneri, (2023) *"Modular Multifunctional Composite Structure for CubeSat Applications: Embedded Battery Prototype Modal Analysis"* **Aerospace** *10*(12): 1009. DOI: 10.3390/aerospace10121009.

[24] B. Gaspar, A. P. Teixeira, and C. Guedes Soares, (2014) *"Assessment of the efficiency of Kriging surrogate models for structural reliability analysis"* **Probabilistic Engineering Mechanics** *37*: 24–34. DOI: 10.1016/j.probengmech.2014.03.011.