Multifidelity Active Learning for the Fast Identification of Incipient Damages in Aerospace Structures

FRANCESCO DI FIORE, MATTEO FENOGLIO
and LAURA MAININI

ABSTRACT

We propose MF-FREEDOM, a multifidelity active learning framework for efficient and accurate structural health monitoring of complex aerospace systems. The method combines (i) a two-stage informative signal compression that reduces data dimensionality while preserving critical features, and (ii) a multifidelity inference scheme that integrates high- and low-fidelity physics-based models to optimize diagnostic accuracy and computational efficiency. We apply MF-FREEDOM to detect incipient fractures in the carbon fiber skin of an aircraft wing achieving accurate damage detection performance with contained computational resources. By enhancing the efficiency and reliability of structural health monitoring, MF-FREEDOM supports the deployment of lightweight and maintainable structures playing a key role in advancing sustainable aviation.

INTRODUCTION

Structural Health Monitoring (SHM) is rapidly gaining prominence as a key enabler in the pursuit of sustainable aviation [1]. As the aerospace sector intensifies efforts to reduce emissions, extend aircraft life cycles, and minimize environmental impact, SHM technologies play a crucial role by enabling predictive maintenance and improving structural reliability. By supporting the adoption of lightweight structural designs, extending component service life, and reducing unscheduled maintenance, SHM contributes directly to resource efficiency and reduced lifecycle emissions. Moreover, SHM fosters the development of intelligent self-monitoring systems, which are essential for the deployment of autonomous and more-electric aircraft architectures. These capabilities are vital to building public and regulatory confidence in novel green aviation technologies. However, while the strategic role of SHM is widely recognized, its practical implementation in aerospace applications remains challenging [2].

Current SHM techniques often require large amounts of high-dimensional data—such as dense strain field measurements—making onboard data storage, real-time processing, and prompt diagnostics complex and inefficient. Moreover, these methods tend to exhibit low sensitivity to incipient damage, such as early-stage micro-fractures, which are criti-

Francesco Di Fiore, Research Associate, Email: f.difiore@imperial.ac.uk. Imperial College London and Brahmal Vasudevan Institute for Sustainable Aviation, London, United Kingdom

cal to detect for ensuring the safe operation of emerging aircraft structures. Additionally, the accuracy of conventional SHM is strongly dependent on high-fidelity models -– e.g. computational structural analysis – which imposes significant computational demands hindering the fast identification of damages.

To address these challenges, we introduce MF-FREEDOM, a multifidelity active learning framework for reliable and efficient diagnosis and monitoring tailored for complex aerospace systems. The framework integrates two key components: (i) a novel two-stage compression algorithm that extracts the most informative features from raw diagnostic signals, reducing data volume while retaining essential diagnostic content; and (ii) a multifidelity active learning strategy that leverages a hierarchy of physics-based models at different levels of fidelity to infer structural health, selectively incorporating high-fidelity simulations only where they most improve diagnostic accuracy. The framework is validated on a representative case study involving incipient fractures in a composite aircraft wing skin panel, an elusive but critical damage mode in carbon fiber laminates.

MF-FREEDOM

The proposed methodology [3] is structured around two core phases. The first phase focuses on optimizing the informative content of diagnostic signals while minimizing the computational cost of their management, leveraging a two-stage optimal compression strategy. The second phase enhances the accuracy of fault inference by efficiently integrating high-fidelity models through a multifidelity Bayesian inference scheme.

TWO-STAGE COMPRESSION STRATEGY

To reduce the dimensionality of diagnostic signals while preserving their informative content, MF-FREEDOM adopts a two-stage optimal compression strategy. The approach is inspired by previous works on sensor placement using dimensionality reduction techniques [4], and combines: (i) Dynamic Mode Decomposition (DMD) to identify dominant dynamical structures of the system, and (ii) a Self-Organizing Map (SOM) to project these features into a low-dimensional space while preserving topological information.

DMD [5] extracts coherent dynamical patterns of the nonlinear damaged system ψ from high-fidelity output signals $\tilde{\boldsymbol{y}}^{(L)}(\boldsymbol{k},\boldsymbol{x})$. Snapshot matrices $\tilde{\boldsymbol{Y}}$ and $\tilde{\boldsymbol{Y}}'$ are built from n_s simulations under varying incipient fault conditions $\boldsymbol{k}$ using a scaled Latin hypercube sampling strategy [4]. Due to the high dimensionality of $\tilde{\boldsymbol{Y}}$, we employ Singular Value Decomposition (SVD) to project snapshots onto a reduced number $n_w \ll n_s$ of informative modes $\boldsymbol{B}$, enabling the computation of the reduced DMD matrix:

$$\boldsymbol{A} = \boldsymbol{Y}'\boldsymbol{Y}^{+} \quad \text{with} \quad \boldsymbol{Y} = \boldsymbol{B}^{*}\tilde{\boldsymbol{Y}},\ \boldsymbol{Y}' = \boldsymbol{B}^{*}\tilde{\boldsymbol{Y}}' \tag{1}$$

The eigendecomposition of $\boldsymbol{A}$ yields the dominant dynamic modes $\boldsymbol{\Upsilon} = [\boldsymbol{\upsilon}_1, \ldots, \boldsymbol{\upsilon}_{n_w}]$ encoding the system's key behaviors.

The DMD modes $\boldsymbol{\Upsilon}$ are then compressed using a SOM neural network [6]. The SOM is trained with the combined dataset $\boldsymbol{T} = [\boldsymbol{x}, \boldsymbol{\upsilon}_1, \ldots, \boldsymbol{\upsilon}_{n_w}]$, and performs a non-linear projection onto a discrete neuron lattice. For each training sample $\boldsymbol{\tau}_i$, the best-matching unit (BMU) is found by minimizing the Euclidean distance:

$$l = \arg\min_{j} \|\boldsymbol{\tau}_i - \boldsymbol{w}_j\| \tag{2}$$

The SOM preserves topological similarities among inputs and clusters self-similar features. The result is an efficient low-dimensional encoding map $\hat{\boldsymbol{y}}(\boldsymbol{k}, \hat{\boldsymbol{x}})$ that retains only the most informative signal components. This compressed output is used in the online phase to reduce the computational complexity of the inference procedure.

MULTIFIDELITY ACTIVE LEARNING FOR DAMAGE INFERENCE

The informative map $\hat{\boldsymbol{y}}(\boldsymbol{k}, \hat{\boldsymbol{x}})$ is used to reduce the dimensionality of the SHM inverse problem:

$$\boldsymbol{k}^* = \arg\min_{\boldsymbol{k} \in \mathcal{K}} \gamma(\boldsymbol{k}, \hat{\boldsymbol{x}}), \tag{3}$$

where $\boldsymbol{k}^*$ is the actual health status of the system, $\gamma(\boldsymbol{k}, \hat{\boldsymbol{x}}) = ||\hat{\boldsymbol{y}}(\boldsymbol{k}^*, \hat{\boldsymbol{x}}) - \boldsymbol{y}_M^{\hat{}}(\boldsymbol{k}, \hat{\boldsymbol{x}})||$ measures the discrepancy in the compressed output space between the reference signal measured onboard $\hat{\boldsymbol{y}}$ and the monitoring signal $\boldsymbol{y}_M^{\hat{}}$ computed with the physics-based models available of the system. To solve this inverse problem efficiently, we adopt a multifidelity Bayesian optimization (MFBO) scheme [7–9], which accelerates damage inference by combining low- and high-fidelity model evaluations. This dynamic online procedure iteratively builds a multifidelity surrogate model of the discrepancy function and uses it to guide evaluations via a multifidelity acquisition function. At each iteration, the most informative damage configuration and fidelity level are selected based on expected utility. New observations update the surrogate model, repeating until convergence.

We construct a surrogate model using the multifidelity Gaussian Process (MFGP) regression [10] to approximate the discrepancy $\gamma^{(l)}(\boldsymbol{k})$ across fidelity levels. Based on the autoregressive model by Kennedy and O'Hagan [11], the lowest fidelity is modeled as a GP:

$$\gamma^{(1)} \sim GP(0, \kappa_1),$$

and higher fidelities are defined recursively as:

$$\gamma^{(l)} = \varrho^{(l-1)}(\boldsymbol{k})\gamma^{(l-1)}(\boldsymbol{k}) + \delta^{(l)}(\boldsymbol{k}), \tag{4}$$

where $\delta^{(l)}$ is an independent GP capturing discrepancies between fidelity levels.

Given a dataset of observed damaged configurations, the posterior distribution of the discrepancy is a GP characterized by its mean and variance:

$$\mu^{(l)}(\boldsymbol{k}) = \kappa_N^{(l)}(\boldsymbol{k})^T \left(\boldsymbol{K} + \sigma_\epsilon^2 \boldsymbol{I}\right)^{-1} \boldsymbol{f}, \tag{5}$$

$$\sigma^{2(l)}(\boldsymbol{k}) = \kappa((\boldsymbol{k}, l), (\boldsymbol{k}, l)) - \kappa_N^{(l)}(\boldsymbol{k})^T \left(\boldsymbol{K} + \sigma_\epsilon^2 \boldsymbol{I}\right)^{-1} \kappa_N^{(l)}(\boldsymbol{k}). \tag{6}$$

where the mean function μ represents the prediction of the discrepancy function over the entire damage space, and the standard deviation σ quantifies the associated uncertainty of the prediction.

We consider three acquisition functions for online damage inference: Multifidelity Expected Improvement (MFEI) [7], Multifidelity Probability of Improvement (MFPI) [8], and Multifidelity Max-value Entropy Search (MF-MES) [9]. Each function selects both a location $\boldsymbol{k}$ and fidelity level $l \in \{1, \dots, L\}$ to balance informativeness and cost.

Multifidelity Expected Improvement (MFEI) generalizes the Expected Improvement (EI) criterion to the multifidelity setting [7]. At each iteration, MFEI selects a fidelity-location pair $(\boldsymbol{k}, l)$ based on:

$$\text{MFEI}(\boldsymbol{k}, l) = \text{EI}(\boldsymbol{k}, L) \cdot \alpha_1(\boldsymbol{k}, l) \cdot \alpha_2(\boldsymbol{k}, l) \cdot \alpha_3(l), \tag{7}$$

where EI$(\boldsymbol{k}, L)$ denotes the EI computed using the high-fidelity model, and the utility terms encode:

$$\alpha_1(\boldsymbol{k}, l) = \text{corr}\left[\gamma^{(l)}(\boldsymbol{k}), \gamma^{(L)}(\boldsymbol{k})\right], \tag{8}$$

$$\alpha_2(\boldsymbol{k}, l) = 1 - \frac{\sigma_\epsilon}{\sqrt{\sigma^{2(l)}(\boldsymbol{k}) + \sigma_\epsilon^2}}, \tag{9}$$

$$\alpha_3(l) = \frac{\lambda^{(L)}}{\lambda^{(l)}}. \tag{10}$$

Here, α_1 captures statistical correlation between fidelities, α_2 accounts for the expected variance reduction due to the query, and α_3 reflects the cost-efficiency of fidelity l relative to the high-fidelity model.

Multifidelity Probability of Improvement (MFPI) extends the Probability of Improvement (PI) criterion in a similar fashion [8]. The acquisition score is computed as:

$$\text{MFPI}(\boldsymbol{k}, l) = \text{PI}(\boldsymbol{k}, L) \cdot \beta_1(\boldsymbol{k}, l) \cdot \beta_2(\boldsymbol{k}, l) \cdot \beta_3(l), \tag{11}$$

where PI$(\boldsymbol{k}, L)$ is the the high-fidelity PI, The terms β_1 and β_3 mirror their counterparts α_1 and α_2 in MFEI, and β_2 is defined as:

$$\beta_2(\boldsymbol{k}, l) = \prod_{i=1}^{n_l} \left[1 - R\left(\boldsymbol{k}, \boldsymbol{k}_i^{(l)}\right)\right]. \tag{12}$$

$$\tag{13}$$

The additional term β_2 penalizes regions with high sampling density via a radial kernel $R(\cdot, \cdot)$, promoting diverse exploration.

Multifidelity Max-value Entropy Search (MFMES) selects points that maximize information gain about the global minimum value γ^* per unit cost [9]:

$$\text{MFMES}(\boldsymbol{k}, l) = \frac{I(\gamma^*; \gamma^{(l)}(\boldsymbol{k}))}{\lambda^{(l)}}, \tag{14}$$

where $I(\gamma^*; \gamma^{(l)}(\boldsymbol{k}))$ is the mutual information between γ^* and a potential query at fidelity l. The information gain cannot be computed in closed form and is estimated via Monte Carlo sampling from the GP posterior.

COMPOSITE PLATE SHM PROBLEM

This SHM problem investigates the detection of incipient damages in a composite plate of an aircraft wing skin, where small cuts in the carbon fiber layers can significantly compromise mechanical integrity while remaining difficult to detect through standard non-destructive testing [12]. The problem is formalized as an inverse optimization task, where the objective is to identify fault parameters by minimizing the discrepancy between measured and simulated strain responses.

The composite plate consists of four layers of plain weave carbon prepreg (IM7/8552 AS4), laminated with a stacking sequence $[45°/0°/0°/45°]$. The plate dimensions are 102 mm in width, 456 mm in length, and 0.76 mm thickness per layer. A fiber cut is introduced in the third (0°) layer to simulate a critical structural fault. The system is loaded along the longitudinal axis, simulating the in-flight aerodynamic conditions of an aircraft wing. Damage is parameterized by $\boldsymbol{k} = [q_1, q_2, q_3, q_4]$, representing the horizontal and vertical position of the cut, its transverse extension, and the applied load, respectively. The goal is to solve:

$$\boldsymbol{k}^* = \min_{\boldsymbol{k} \in \mathcal{K}} \gamma(\boldsymbol{k}), \tag{15}$$

where the discrepancy function measures the root-mean-square error between reference and simulated strain fields:

$$\gamma(\boldsymbol{k}) = \sqrt{\frac{1}{N} \sum_{j=1}^{N} \frac{(S_{ref}^{j} - S_{mon}^{(l)j})^2}{S_{ref}^{j}}} \tag{16}$$

with N being the number of elements in the strain field. The search domain is $\mathcal{K} = [0, 102] \times [0, 456] \times [0, 30] \times [0, 20]$.

The strain field S is computed using the Reissner-Mindlin plate theory solved via Finite Element Method in MSC Patran/Nastran. Two models are developed at different levels of fidelity. The high-fidelity model is a 3D mesh composed of HEXA8 elements, refined near the cut to accurately resolve local strain concentrations. Each layer is discretized with three elements through thickness, and the cut is explicitly modeled as a rectangular notch in the third layer. This model captures subtle variations in strain caused by small damage or complex load interactions and serves both as a reference and as the highest available fidelity $S_{ref}(\boldsymbol{k})$ and $S_{mon}^{(2)}(\boldsymbol{k})$ and requires 40 minutes of CPU time. The low-fidelity model uses a 2D QUAD4 mesh with a coarser two-dimensional discretization approximating the same geometry in a simplified form and requiring 8 minutes of CPU time.

We define the relative computational costs as $\lambda^{(1)} = 0.2$ and $\lambda^{(2)} = 1$. Figures 1 and 2 show the computational mesh and resulting strain fields for both models. The example scenario is a cut in the third layer at $(x, y) = (40 \text{ mm}, 250 \text{ mm})$ with a 10 mm extension and 5 N load.

RESULTS AND DISCUSSION

This section presents and analyzes the results obtained using the MF-FREEDOM framework for the structural health monitoring problem of a composite plate. The performance of MF-FREEDOM is evaluated under three distinct multifidelity active learning

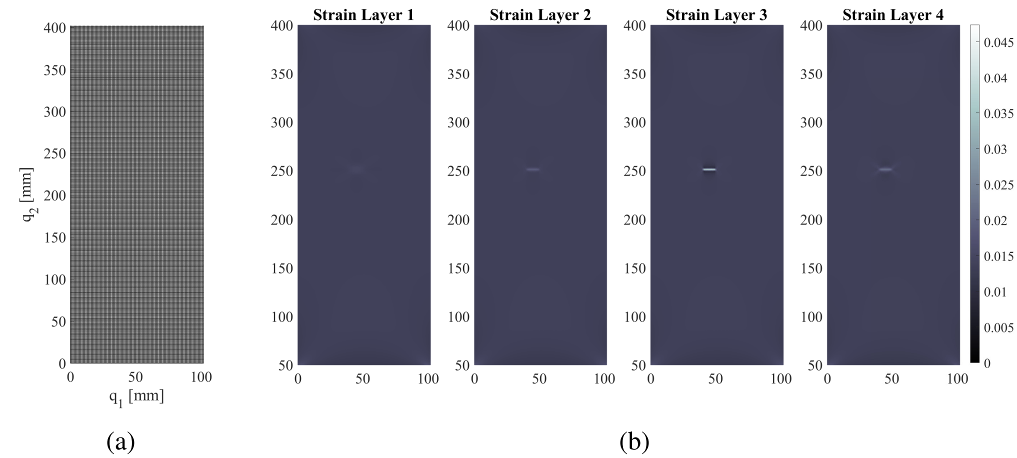

Figure 1. (a) High-fidelity discretization of the computational domain. (b) High-fidelity strain distribution across the four composite layers.

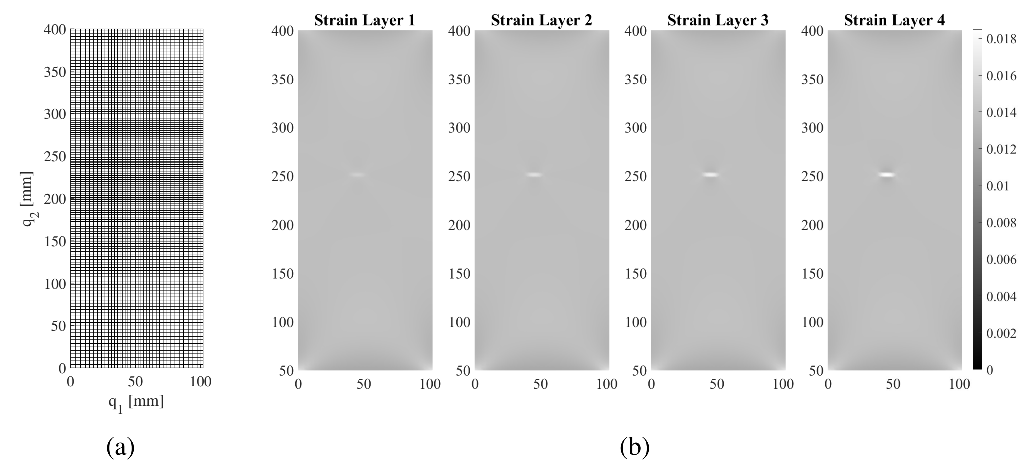

Figure 2. (a) Low-fidelity mesh discretization. (b) Corresponding strain field in the four layers of the damaged plate.

strategies, each employing a different multifidelity acquisition function namely MFEI, MFMES, and MFPI. These are compared against a baseline approach based on the Efficient Global Optimization (EGO) algorithm [13], which operates exclusively on high-fidelity simulations without any compression or multifidelity acceleration.

To assess the effectiveness of each optimization strategy, we consider the performance metric $\gamma^* = \min(\gamma(\boldsymbol{k}))$, which quantifies the minimum discrepancy between the predicted and actual damage parameters. A zero value of this metric indicates an exact identification of the damage configuration.

The evaluation is based on 25 different incipient damage scenarios generated via a scaled Latin Hypercube Sampling strategy. For each acquisition strategy, the SHM process is repeated across all 25 scenarios, and results are reported in terms of the median and interquartile range (25th–75th percentile) of the metric γ^*, plotted against the computational budget B. The budget is defined as the cumulative cost $\lambda_i^{(l)}$ associated with

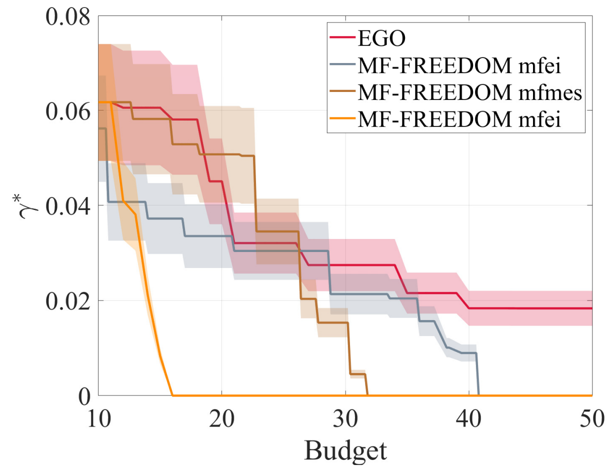

Figure 3. Statistics over 25 runs of the minimum discrepancy γ^* obtained with the MF-FREEDOM adopting the MFEI, MFMES and MFPI multifidelity active learning and the EGO algorithm.

querying the structural models at different fidelity levels throughout the optimization iterations.

Figure 3 reports the outcomes in terms of median and interval between the 25-th and 75-th percentiles for the assessment metric γ^* as function of the computational budget $B = \sum \lambda_i^{(l)}$ measured as the cumulative computational cost $\lambda_i^{(l)}$ used at each iteration i to evaluate the l-th structural level of fidelity. Overall, all the multifidelity active learning strategies implemented in the MF-FREEDOM allow the exact identification of the health status of the composite plate within the allocated budget, while the SHM strategy based on the EGO algorithm and relying uniquely on high-fidelity data fails in the identification task within the maximum budget. These outcomes show the clear advantage of combining an efficient compression stage of the strain signals with the multifidelity active learning for the identification of damages relying of physics-based models at different levels of fidelity: the MF-FREEDOM massively queries the low-fidelity structural model to accelerate the exploration of different damage configurations, and contains the evaluation of expensive high-fidelity analysis that are wisely queried to strategically refine the accuracy of the SHM process.

Among the multifidelity active learning strategies, MFPI exhibits the fastest convergence, achieving accurate damage identification with an average computational budget equivalent to just $B = 16.2$ high-fidelity evaluations. This is attributed to the inherently exploitative nature of the probability of improvement (PI) acquisition function, which prioritizes regions of the design space with high potential for improvement, thereby accelerating convergence when the damages space is low-dimensional. In contrast, MFEI and MFMES adopt more explorative behaviors, delaying exploitation and thus requiring a higher computational cost to achieve comparable accuracy. In particular, MFMES shows slower convergence due to its entropy-based acquisition strategy, which relies on a Monte Carlo approximation of the information gain. This approximation tends to underestimate the true information benefit, thus favoring cost-efficient but less informative low-fidelity evaluations.

CONCLUDING REMARKS

Overall, the MF-FREEDOM framework enables a significant acceleration of the SHM process while ensuring robust and reliable damage identification. This is achieved through the strategic integration of high-fidelity data and the systematic use of physics-based models across multiple levels of fidelity. By embedding the governing physical principles directly into the learning and SHM process, MF-FREEDOM maintains a strong connection with the underlying structural behavior which is critical for the trustworthiness of the results. At the same time, the two-stage compression strategy mitigates the computational burden associated with high-dimensional data significantly accelerating the online multifidelity active learning and allows the learner to efficiently navigate the damage space allocating high-fidelity evaluations more judiciously. In this way, MF-FREEDOM supports the development of lightweight and self-diagnosing structural systems contributing to the broader goal of sustainable aviation.

ACKNOWLEDGMENT

This work has been partially supported by the Brahmal Institute for Sustainable Aviation.

REFERENCES

1. Afonso, F., M. Sohst, C. M. Diogo, S. S. Rodrigues, A. Ferreira, I. Ribeiro, R. Marques, F. F. Rego, A. Sohouli, J. Portugal-Pereira, et al. 2023. “Strategies towards a more sustainable aviation: A systematic review,” *Progress in Aerospace Sciences*, 137:100878.
2. Diamanti, K. and C. Soutis. 2010. “Structural health monitoring techniques for aircraft composite structures,” *Progress in Aerospace Sciences*, 46(8):342–352.
3. Di Fiore, F., P. Berri, and L. Mainini. 2024. “Diagnosing Incipient Faults for a Faster Adoption of Sustainable Aerospace Technologies,” *AIAA Journal*, 62(8):3111–3126.
4. Mainini, L. and K. E. Willcox. 2017. “Sensor placement strategy to inform decisions,” in *18th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, p. 3820.
5. Schmid, P. J. 2010. “Dynamic mode decomposition of numerical and experimental data,” *Journal of fluid mechanics*, 656:5–28.
6. Kohonen, T. 1990. “The self-organizing map,” *Proceedings of the IEEE*, 78(9):1464–1480.
7. Huang, D., T. T. Allen, W. I. Notz, and R. A. Miller. 2006. “Sequential kriging optimization using multiple-fidelity evaluations,” *Structural and Multidisciplinary Optimization*, 32(5):369–382.
8. Ruan, X., P. Jiang, Q. Zhou, J. Hu, and L. Shu. 2020. “Variable-fidelity probability of improvement method for efficient global optimization of expensive black-box problems,” *Structural and Multidisciplinary Optimization*, 62(6):3021–3052.
9. Takeno, S., H. Fukuoka, Y. Tsukada, T. Koyama, M. Shiga, I. Takeuchi, and M. Karasuyama. 2020. “Multi-fidelity Bayesian optimization with max-value entropy search and its parallelization,” in *International Conference on Machine Learning*, PMLR, pp. 9334–9345.
10. Brevault, L., M. Balesdent, and A. Hebbal. 2020. “Overview of Gaussian process based multi-fidelity techniques with variable relationship between fidelities, application to aerospace systems,” *Aerospace Science and Technology*, 107:106339.
11. Kennedy, M. C. and A. O’Hagan. 2000. “Predicting the output from a complex computer code when fast approximations are available,” *Biometrika*, 87(1):1–13.
12. Giurgiutiu, V. 2015. “Structural health monitoring of aerospace composites,” :449–507.
13. Jones, D. R., M. Schonlau, and W. J. Welch. 1998. “Efficient global optimization of expensive black-box functions,” *Journal of Global optimization*, 13(4):455.