Delft University of Technology

Document Version
Final published version

Licence
CC BY

Citation (APA)
Guo, L., Kansara, H., Khosroshahi, S. F., Zhang, G. Q., & Tan, W. (2026). Multi-fidelity Bayesian data-driven design of energy absorbing spinodoid cellular structures. *Computers and Structures*, *323*, Article 108155. https://doi.org/10.1016/j.compstruc.2026.108155

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.

This work is downloaded from Delft University of Technology.

Computers and Structures 323 (2026) 108155

Contents lists available at ScienceDirect

Computers and Structures

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

Multi-fidelity Bayesian data-driven design of energy absorbing spinodoid cellular structures

Leo Guo[a], Hirak Kansara[b], Siamak F. Khosroshahi[b], GuoQi Zhang[a,*], Wei Tan[b,*]

[a] Department of Electronic Components, Technology and Materials, Delft University of Technology, Mekelweg 4, Delft, 2628CM, Netherlands
[b] School of Engineering and Materials Science, Queen Mary University of London, Mile End Road, London, E1 4NS, United Kingdom

ARTICLE INFO

Keywords:
Data-driven design
Multi-fidelity
Bayesian optimization
Cellular structures
Energy absorption

ABSTRACT

Finite element (FE) simulations of structures and materials are becoming increasingly accurate, but also more computationally expensive as a collateral result. This development occurs in parallel with a growing demand for data-driven design. To reconcile the two, a robust and data-efficient optimization method called Bayesian optimization (BO) has been previously established as a technique to optimize expensive objective functions. The mesh width of an FE model can be exploited to evaluate an objective at a lower or higher fidelity (cost & accuracy) level, which is the domain of multi-fidelity BO (MFBO) applications. However, BO and MFBO are usually not directly compared in the literature. Moreover, sampling quality and assessing design parameter sensitivity are often underrepresented parts of data-driven design. This paper combines global sensitivity analysis and (MF) BO into a novel, efficient Bayesian data-driven framework. We compare the performance of BO with that of MFBO by maximizing the energy absorption (EA) problem of spinodoid cellular structures. The findings show that similar or better designs are suggested by MFBO with 16% fewer expensive objective evaluations compared to BO when maximizing the EA. The results, which are made open-source, serve to support the utility of multi-fidelity techniques across expensive data-driven design problems.

1. Introduction

Energy-absorbing structures, such as crumple zones, are essential in applications like vehicle safety, where their primary function is to convert kinetic energy into plastic deformation to protect occupants. By minimizing the forces experienced during collisions, these structures play a crucial role in the crashworthiness of vehicles, a key consideration in design and manufacturing. Traditionally, materials like steel and aluminum have been used in these applications. With advancements in 3D metal printing [1], this process of additive manufacturing allows for the creation of complex prototypes that enhance energy absorption efficiency. More recently, however, there has been a shift towards using advanced composite materials and architected metamaterials, such as cellular structures [2] and shell-based structures [3,4], that offer superior strength-to-weight ratios and improved energy absorption.

With the development of materials science, mechanical metamaterials have emerged as a promising solution for lightweight and efficient energy absorption (EA). Architected materials, including periodic structures like the aforementioned cellular solids, offer unique mechanical advantages but are often limited by stress concentrations that lead to premature failure. Innovations such as triply periodic minimal surfaces (TPMS) [5] introduce doubly curved, non-self-intersecting structures that distribute stress more evenly. However, while TPMS provides higher material stiffness, the aperiodic spinodoid metamaterials [6] are more robust against imperfections in fabrication.

Optimization of these complex energy-absorbing structures may require a strategic approach to balance multiple conflicting objectives, such as maximizing EA while minimizing peak forces (PF) experienced during impact. Conventional design methods often rely on trial-and-error or topology optimization [7–9]. Meanwhile, topology optimization iteratively redistributes material within a design domain but can produce geometries with sharp angles and irregularities that complicate fabrication. With both of these traditional methods having major drawbacks, data-driven design is a strong contender to address this type of optimization problem.

One of the cornerstones of the data-driven design process is design space sampling. Popular examples include grid sampling [10], uniform random sampling, and Latin hypercube sampling. While these methods have proven to be effective in some use cases, Sobol' sampling has proven to be more effective in gaining insight when data is scarce due to

---

* Corresponding authors.
Email addresses: g.q.zhang@tudelft.nl (G. Zhang), wei.tan@qmul.ac.uk (W. Tan).

https://doi.org/10.1016/j.compstruc.2026.108155
Received 12 September 2025; Accepted 12 February 2026

Available online 20 February 2026
0045-7949/© 2026 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

an expensive objective: by using Sobol' samples, it is possible to perform variance-based sensitivity analysis of the objective with respect to the design parameters.

However, sampling alone is not sufficient to provide a robust objective optimization scheme. In literature, data-driven optimization approaches leverage machine learning models as surrogates for costly simulations. Among them, Bayesian optimization (BO) with Gaussian process regression has gained traction for optimizing meta-material structure designs with high computational cost [11–13], including spinodoid structures [14]. BO is a popular data-efficient objective optimization methodology. It aims to leverage the proven regression capabilities of Bayesian machine learning methods to assess the likelihood of the optimum location. This leads to a delicate balance between exploring new regions of high uncertainty where the black box function may have a global (or at least a better local) optimum, and exploiting regions where an optimum is already expected to be present. The core of BO consists of first training a surrogate model on a set of training data, populated by evaluated design experiments, and subsequently optimizing a proxy belief model based on the surrogate regression model – also called an acquisition function. The design that is found through this process is then evaluated and augmented to the original training data, after which the cycle starts anew. See Fig. 1 for a schematic overview of BO.

BO can be flexibly adapted to settings with data that adhere to particular input or output structures. Multi-fidelity BO (MFBO), which relies on surrogate models that handle the cost-accuracy trade-off between low- and high-fidelity data, allows for accelerated convergence by leveraging inexpensive approximations alongside selective high-fidelity evaluations. This approach enables more comprehensive exploration of the design space while maintaining the accuracy needed for realistic applications. Numerous successful mechanical engineering design optimizations have been performed by using MFBO. These cases include mathematical and physical phenomena with complex behavior, such as inverted pendulums [15], fractional advection-dispersion equations [16], haemodynamics [17], and non-linear state-space models [18]. Additionally, engineering design processes have been developed with MFBO, including the design of antennas [19], fixed-wing drones [20], helicopter blades [21], aerofoils [22–24], flip-chip packages [25], horizontal road alignment [26], and analog circuits [27].

The multi-fidelity techniques have proven particularly useful when the objective functions are provided by finite element (FE) simulations, where the fidelity trade-off is intuitively encoded by the mesh size [28,29]. By incorporating elastic-plastic material behaviors into quasi-static FE compression simulations, complex structural responses under compression can be captured. The positive track record of MFBO applications renders the EA optimization problem a suitable problem to tackle with a multi-fidelity optimization methodology.

This work aims to present an investigation into two gaps within the current practice of data-driven design optimization. Firstly, a notable lack in the literature exists with regard to an integrated Bayesian framework to perform robust and efficient data-driven design. This is partly due to the under-explored role of design space sampling in sensitivity analysis. Secondly, to the authors' knowledge, there is no precedent towards any direct comparisons between MFBO and its single-fidelity counterpart, especially in the context of applications to solve expensive engineering problems. The core motivation to present such a direct comparison is the authors' belief that it is possible to make savings in computational expense made by single-fidelity BO by exploiting transfer learning in the multi-fidelity setting. Achieving this will underscore the potential of using MFBO as a key contender to perform robust and efficient data-driven design.

In order to address both issues, this paper introduces a novel two-stage approach: first, Saltelli sampling is used to perform global, variance-based sensitivity analysis of the EA objective across two fidelity levels. This offers key insight into the importance of each design parameter with respect to the objective. Subsequently, both BO and MFBO are applied to optimize spinodoid topologies for EA, enabling a rigorous, side-by-side evaluation of optimization performance.

Our work combines Sobol' sensitivity analysis (Section 2.1) and Bayesian optimization (Section 2.2) into one Bayesian data-driven framework. See Fig. 2 for a schematic overview of the framework.

Remark that the individual parts of the Bayesian framework have previously been described and applied in literature, while the framework as a whole is a novel development to promote robust and efficient data-driven design practice. Another novel element of this framework is the reuse of the Saltelli samples as the initial DoE for (MF)BO, removing additional computational expense to sample such a DoE separately.

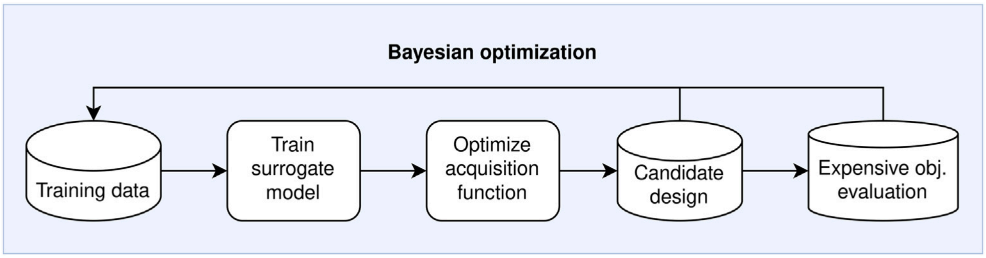

Fig. 1. Flowchart diagram of BO.

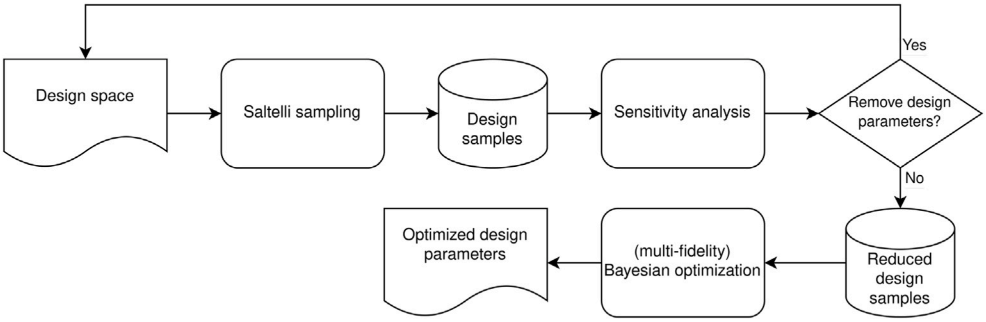

Fig. 2. Flowchart diagram of the Bayesian data-driven framework, which includes BO (Fig. 1).

This study is among the first to demonstrate the practical benefits and trade-offs of MFBO in complex, physics-driven design spaces, as an alternative to single-fidelity BO. To support transparency and reproducibility, all code and data are made openly available via GitHub. The manuscript concludes with a critical analysis of MFBO's applicability to safety-critical materials design, outlining future directions for advancing multi-fidelity optimization strategies.

2. Data-driven methodology

2.1. Sobol' sensitivity analysis

Let $f : [0,1]^D \to \mathbb{R}$ be an objective function of the design parameters $x_1, x_2, \ldots, x_D$. Our final goal is to find a global optimum (maximum) of f on the scaled design space, which is the unit cube. This can be written as:

$$\begin{aligned} &\text{Maximize} \quad f(x) \\ &\text{Subject to} \quad 0 \le x_i \le 1, \quad i = 1, \ldots, D \end{aligned} \tag{1}$$

To gain an initial insight into the response of f across the (scaled) design space, it is essential to know how much f varies by perturbing the design parameters. This is precisely the goal of sensitivity analysis, which is a common technique used in practice when phenomena involving expensive or scarce data are concerned, such as hydrology [30] and financial decision-making [31]. Indeed, a review of various methods and applications identifies sensitivity analysis as part of the best practices to ensure modeling quality in general [32].

Originally introduced by Sobol' [33], the Sobol' sensitivity analysis (SSA) approach treats the objective and design parameters $f(X_1, \ldots, X_D) = Y$ as fully stochastic, with $X_1, \ldots X_D, Y$ being modeled as random variables. Next, let $I = \{i_1, \ldots, i_d\} \subset \{1, \ldots, D\}$ be a subset of size d, also called a d-set. Define $\mathbf{X}_I := (X_{i_1}, \ldots, X_{i_d})^\top$. By the law of total variance,

$$\mathrm{Var}(Y) = \mathbb{E}(\mathrm{Var}(Y|\mathbf{X}_I)) + \mathrm{Var}(\mathbb{E}(Y|\mathbf{X}_I)), \tag{2}$$

which splits the total, unconditional variance of the dependent variable Y into its unexplained and explained components with respect to the selected design parameters according to the indices in I.

By dividing the explained variance component $\mathrm{Var}(\mathbb{E}(Y|\mathbf{X}_I))$ by the total variance of Y, one obtains a normalized variance metric, also called the Sobol' sensitivity index [34,35] with respect to $\mathbf{X}_I$:

$$S_I := \frac{\mathrm{Var}(\mathbb{E}(Y|\mathbf{X}_I))}{\mathrm{Var}(Y)}. \tag{3}$$

The special univariate case where $I = \{i\}$ for some $i \in \{1, \ldots, D\}$,

$$S_i := \frac{\mathrm{Var}(\mathbb{E}(Y|X_i))}{\mathrm{Var}(Y)}, \tag{4}$$

is called the first-order sensitivity index with respect to the design parameter X_i, and it is of particular interest.

Dual to the first order index S_i is the concept of total effect index with respect to X_i, defined as the ratio between the *unexplained* variance of all design variables *except* X_i, and the total variance:

$$S_{T_i} := \frac{\mathbb{E}(\mathrm{Var}(Y|\mathbf{X}_{\sim i}))}{\mathrm{Var}(Y)} = 1 - \frac{\mathrm{Var}(\mathbb{E}(Y|\mathbf{X}_{\sim i}))}{\mathrm{Var}(Y)}, \tag{5}$$

where $\sim i := \{1, \ldots, D\} \setminus \{i\}$.

Calculating the required statistical moments for S_i and S_{T_i} generally implies the calculation of multiple intractable integrals, and is therefore almost always approximated in literature. Several established methods exist to calculate S_i and S_{T_i}, for example, by using Gaussian quadrature in the context of Gaussian process regression [36]. This method uses Latin hypercube sampling to approximate the integrals needed to calculate the statistical moments corresponding to the sensitivity indices. However, a more commonly used approximation to evaluate the expressions in Eqs. (4) and (5) is given by Saltelli et al. [37], as briefly described in Appendix A. Importantly, Saltelli samples are quasi-random in nature. This means that they can be appended in a space-filling fashion without the risk of clustering, which is a potential issue when building accurate surrogate models. An existing open-source implementation of SSA with Saltelli samples in Python called SALib [38] is used throughout the manuscript.

2.2. Single- and multi-fidelity Bayesian optimization

The core of Bayesian optimization is to solve the optimization problem posed in Eq. (1) by transforming a surrogate model into knowledge about a candidate design that is believed to evaluate to a better objective value. The most popular type of surrogate model is a Gaussian process regression model [39], which gives rise to a regression-predictive normal distribution with probability density function $\hat{\phi}^{(i)}$ at Bayesian optimization iteration i. A more detailed exposition can be found in Appendix B.

It is possible to shape the information that $\hat{\phi}^{(i)}$ provides into a design selection policy that suggests a candidate $\mathbf{x}^{(i)} \in [0,1]^D$ for objective evaluation. In short, this is achieved by means of a so-called acquisition function or infill function $\alpha : [0,1]^D \to \mathbb{R}$, which intuitively assigns a measure of belief $\alpha(\mathbf{x}; \hat{\phi}^{(i)})$ to the design $\mathbf{x} \in [0,1]^D$ that $f(\mathbf{x})$ can achieve a better optimum than $f(\mathbf{z})$ for any design row $\mathbf{z}^\top$ in $\mathbf{X}^{(i-1)}$.

As α expresses a level of belief across the design parameter space, design selection takes place by locating such $\mathbf{x} \in [0,1]^D$ at which α is maximized: see Eq. (6).

$$\mathbf{x}^{(i)} := \mathrm{argmax}_{\mathbf{x} \in [0,1]^D} \alpha(\mathbf{x}; \hat{\phi}^{(i)}) \tag{6}$$

Popular examples of acquisition functions α are the expected improvement (EI) acquisition [40] and the upper/lower confidence bound (UCB) acquisition [41]. See Eqs. (7) and (8) respectively.

$$\alpha_{\mathrm{EI}}(\mathbf{x}; \hat{\phi}^{(i)}) := \hat{\sigma}^{(i)}(\mathbf{x})(z^{(i)}(\mathbf{x})\hat{\Phi}(z^{(i)}(\mathbf{x})) + \hat{\phi}(z^{(i)}(\mathbf{x}))) \tag{7}$$

$$\alpha_{\mathrm{UCB}}(\mathbf{x}; \hat{\phi}^{(i)}, \beta) := \hat{\mu}^{(i)}(\mathbf{x}) + \beta\hat{\sigma}^{(i)}(\mathbf{x}) \tag{8}$$

In the UCB acquisition function, β is a hyperparameter, usually set to 2 to refer to a 95% confidence interval.

In practice, a logarithmic variant of the EI acquisition function (LogEI) is used due to numerical stability during optimization [42]. The derivations and intuition behind these formulations can be referred to in Refs. [43,44].

The BO schematic introduced in Fig. 1 is summarized in Algorithm 1 when GPR is used as a surrogate modeling methodology.

In step 8 of Algorithm 1, the Rec function returns the recommended design-objective pair that is found throughout the optimization process. This corresponds to the design that achieves the optimized objective throughout all of the iterations.

The goal of (single-fidelity) Bayesian optimization is to optimize an objective function f. When there are M objectives $f_1, \ldots, f_M$, ordered in

Algorithm 1 Bayesian optimization with Gaussian process regression (BO).

```
Require: Design of training experiments $\mathbf{D}^{(0)}$, covariance function $\kappa$,
    acquisition function $\alpha$, number of iterations $I$
1: for $i = 1, \ldots, I$ do
2:     $\hat{\theta}^{(i)}_{\mathrm{MLE}} \leftarrow \mathrm{argmin}_{\theta \in \mathcal{T}_\kappa} \ln(\det(\mathbf{K}_\theta(\mathbf{X}^{(i-1)}))) + \mathbf{y}^{(i-1)\top}\mathbf{K}_\theta^{-1}(\mathbf{X}^{(i-1)})\mathbf{y}^{(i-1)}$
3:     $\hat{\phi}^{(i)} \leftarrow \hat{\theta}^{(i)}_{\mathrm{MLE}}$
4:     $\mathbf{x}^{(i)} \leftarrow \mathrm{argmax}_{\mathbf{x} \in [0,1]^D} \alpha(\mathbf{x}; \hat{\phi}^{(i)})$
5:     $y^{(i)} \leftarrow f(\mathbf{x}^{(i)})$
6:     $\mathbf{D}^{(i)} \leftarrow (\mathbf{D}^{(i-1)}, (\mathbf{x}^{(i)\top}, y^{(i)}))^\top$
7: end for
8: $(\mathbf{x}_{\mathrm{rec}}, y_{\mathrm{rec}}) \leftarrow \mathrm{Rec}(\mathbf{D}^{(I)})$
9: return $(\mathbf{x}_{\mathrm{rec}}, y_{\mathrm{rec}})$
```

increasing levels of fidelity, the commonly assumed goal is to optimize only f_M [45–48]. In particular, the interest does *not* lie with optimizing $f_1, \dots, f_{M-1}$, and data acquired from any of these $M-1$ objectives merely serve an auxiliary purpose to transfer knowledge to fidelity level M. This draws an important distinction between multi-output BO (MOBO) and BO that utilizes MOGP regression models, which is called multi-fidelity Bayesian optimization (MFBO): while MOBO tries to optimize the various objective functions simultaneously [49–51], MFBO leverages the knowledge transfer facilitated by MOGPR to optimize f_M, and is the focal point of this manuscript. In particular, multi-task Gaussian process regression is used as the MOGPR method of choice. More details can be found in Appendix B.

In the context of multi-fidelity Bayesian optimization, it is important to address the literature surrounding acquisition functions capable of handling multi-fidelity data sets. A large class of acquisition functions involves the arithmetic combination of functions that pertain to various factors involving the design space and/or the fidelity space. In particular, ideas from single-fidelity acquisition functions get exploited and augmented to handle multiple levels of fidelity. This is similar to how constrained EI [52] or local penalization of generic acquisitions [53] is constructed in the single-fidelity scenario.

An example of a multiplicative acquisition is an augmented version of the EI acquisition function based on multi-fidelity sequential kriging optimization [54], also called variable-fidelity EI (VF-EI). It is possible to modify VF-EI by using the logarithmic EI acquisition function instead [42], giving rise to VF-LogEI. This multi-fidelity acquisition function multiplies the EI function applied to the highest fidelity, with fidelity-dependent parameters such as the ratio of computational cost and correlation between the highest fidelity and the fidelity in question. Other similar multi-fidelity acquisition functions built on top of single-fidelity improvement-based acquisitions include extended expected improvement [55], variable-fidelity probability of improvement (VF-PI) [56], and variable-fidelity upper confidence bound (VF-UCB) [47]. For this manuscript, the main focus will be on the VF-(Log)EI and VF-UCB. See Eqs. (9) and (10).

$$\alpha_{\text{VF-LogEI}}(\mathbf{x}, m; \hat{\phi}^{(i)}) := \alpha_{\text{LogEI}}(\mathbf{x}; \hat{\phi}_m^{(i)}) \cdot \text{CR}(m) \cdot \rho(m); \tag{9}$$

$$\alpha_{\text{VF-UCB}}(\mathbf{x}, m; \hat{\phi}^{(i)}) := \omega_1 \cdot \hat{\mu}_m^{(i)}(\mathbf{x}) + \omega_2 \cdot \hat{\sigma}_m^{(i)}(\mathbf{x}) \cdot \text{CR}(m). \tag{10}$$

Here, CR(m) stands for the cost ratio, i.e., the ratio of (computational) expense of one evaluation of f_m compared to one evaluation of f_M, and $\hat{\phi}_m^{(i)}$ is the m-th fidelity regressive-predictive normal distribution at iteration i with mean $\hat{\mu}_m^{(i)}$ and standard deviation $\hat{\sigma}_m^{(i)}$ (Eq. (B.24)). Furthermore, in the case of $M = 2$, $\rho(m = 1)$ is the Pearson correlation coefficient between the low- and high-fidelity objectives, otherwise $\rho(m = 2) = 1$. Finally, in the case of VF-UCB, ω_1, ω_2 are weights that depend on the coefficients of variation related to f_m throughout the entire design space. The weights ω_1 and ω_2 are given by

$$\omega_1 := \frac{c_1}{c_1 + c_2}, \quad \omega_2 := \frac{c_2}{c_1 + c_2}, \tag{11}$$

with

$$c_1 := \frac{\left(\int_{[0,1]^D} \hat{\mu}_M^2(\mathbf{x}) \, \mathrm{d}x - \left(\int_{[0,1]^D} \hat{\mu}_M(\mathbf{x}) \, \mathrm{d}x \right)^2 \right)^{1/2}}{\int_{[0,1]^D} \hat{\mu}_M(\mathbf{x}) \, \mathrm{d}x} \tag{12}$$

$$c_2 := \frac{\left(\int_{[0,1]^D} \hat{\sigma}_M^2(\mathbf{x}) \, \mathrm{d}x - \left(\int_{[0,1]^D} \hat{\sigma}_M(\mathbf{x}) \, \mathrm{d}x \right)^2 \right)^{1/2}}{\int_{[0,1]^D} \hat{\sigma}_M(\mathbf{x}) \, \mathrm{d}x} \tag{13}$$

as the coefficients of variation of the posterior (or regressive-predictive) mean and standard deviation, respectively. These can be approximated numerically by means of sampling the design space.

The Bayesian optimization algorithm in the multi-fidelity scenario is largely similar to Algorithm 1, and is presented in Algorithm 2.

Algorithm 2 Bayesian optimization with MFGPR (MFBO).

```
Require: Design of train experiments $\mathcal{D}^{(0)}$, covariance function $\mathcal{K}$,
    model parameter vector domain $\mathcal{T}$, acquisition function $\alpha$, number
    of iterations $I$, cost function $c$, budget $B$
 1: $b \leftarrow 0$
 2: for $i = 1, \dots, I$ do
 3:     if $b > B$ then
 4:         break
 5:     end if
 6:     $\hat{\theta}_{\text{MLE}}^{(i)} \leftarrow \text{argmin}_{\theta \in \mathcal{T}} \ln(\det(\mathbf{K}_\theta(\mathcal{X}^{(i-1)}))) + \mathbf{y}^{(i-1)\top} \mathbf{K}_\theta^{-1}(\mathcal{X}^{(i-1)}) \mathbf{y}^{(i-1)}$
 7:     $\hat{\phi}^{(i)} \leftarrow \hat{\theta}_{\text{MLE}}^{(i)}$
 8:     $(\mathbf{x}^{(i)\top}, m^{(i)}) \leftarrow \text{argmax}_{\mathbf{x} \in [0,1]^D, m \in \{1,\dots,M\}} \alpha(\mathbf{x}, m; \hat{\phi}^{(i)})$
 9:     $y^{(i)} \leftarrow f_{m^{(i)}}(\mathbf{x}^{(i)})$
10:     $b \leftarrow b + c(m^{(i)})$
11:     if $m = m^{(i)}$ then
12:         $\mathbf{D}_m^{(i)} \leftarrow (\mathbf{D}_m^{(i-1)}, (\mathbf{x}^{(i)\top}, y^{(i)}))^\top$
13:     else
14:         $\mathbf{D}_m^{(i)} \leftarrow \mathbf{D}_m^{(i-1)}$
15:     end if
16: end for
17: $(\mathbf{x}_{\text{rec}}, y_{\text{rec}}) \leftarrow \text{Rec}(\mathcal{D}^{(I)})$
18: return $(\mathbf{x}_{\text{rec}}, y_{\text{rec}})$
```

Table 1
Design parameters for the 4D EA optimization problem.

Design parameter	Lower bound	Upper bound
ρ	0.3	0.6
θ_1	0°	90°
θ_2	0°	90°
θ_3	0°	90°

One main difference between single- and multi-fidelity BO is the augmentation of the data set at each fidelity, showcased by steps 11–15. If fidelity m is selected to represent the fidelity level to sample at for iteration i, only $\mathbf{D}_m^{(i)}$ is updated, while all the data sets at the other fidelities remain unchanged. Similar to single-fidelity BO, the multi-fidelity counterpart has a recommendation function that depends on the final multi-fidelity data set after the final iteration I. This recommender returns the largest value among $\{f_M(\mathbf{x}_1^*), f_M(\mathbf{x}_2^*), \dots, f_M(\mathbf{x}_M^*)\}$ and the corresponding design parameter vector, where $\mathbf{x}_m^*$ is the incumbent maximizer of f_m after iteration I.

3. Problem description and data analysis

3.1. Finite element analysis and objective description

To optimize the energy absorption objective of the spinodoid structure according to the data-driven framework in Fig. 2, four relevant design parameters have been identified, along with their domains. See Table 1 for an overview.

From a physical point of view, ρ represents the relative density of the structure as a material-to-air ratio, and the cone angles $\theta_1, \theta_2, \theta_3$ give rise to the configuration of the structure itself [57,58]. See Fig. 3 for an illustration.[1]

Varying the values of $\theta_1, \theta_2, \theta_3$ will lead to different material structures and interconnections. Some example structures are visualized in Fig. 4.

As can be observed from Fig. 4, different sets of angular parameters can yield qualitatively distinct structures, especially when the cone angles are at or near 0°. A small θ_3, for example, yields a columnar structure, which represents an interconnected network of pillars (Fig. 4(a)), while small values of θ_1 and θ_2 yield a lamellar material configuration, consisting of connected layers (Fig. 4(d)).

---

[1] It should be noted that the depicted structure serves as a visual aid and the values of ρ are kept constant throughout the structure for any data-related results.

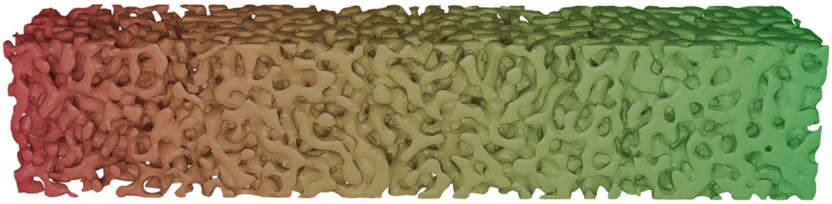

Fig. 3. Spindoid structure with varying levels of $\rho \in [0.3, 0.6]$, increasing from left to right, for fixed values $\theta_1 = 90°$, $\theta_2 = 0°$ and $\theta_3 = 0°$.

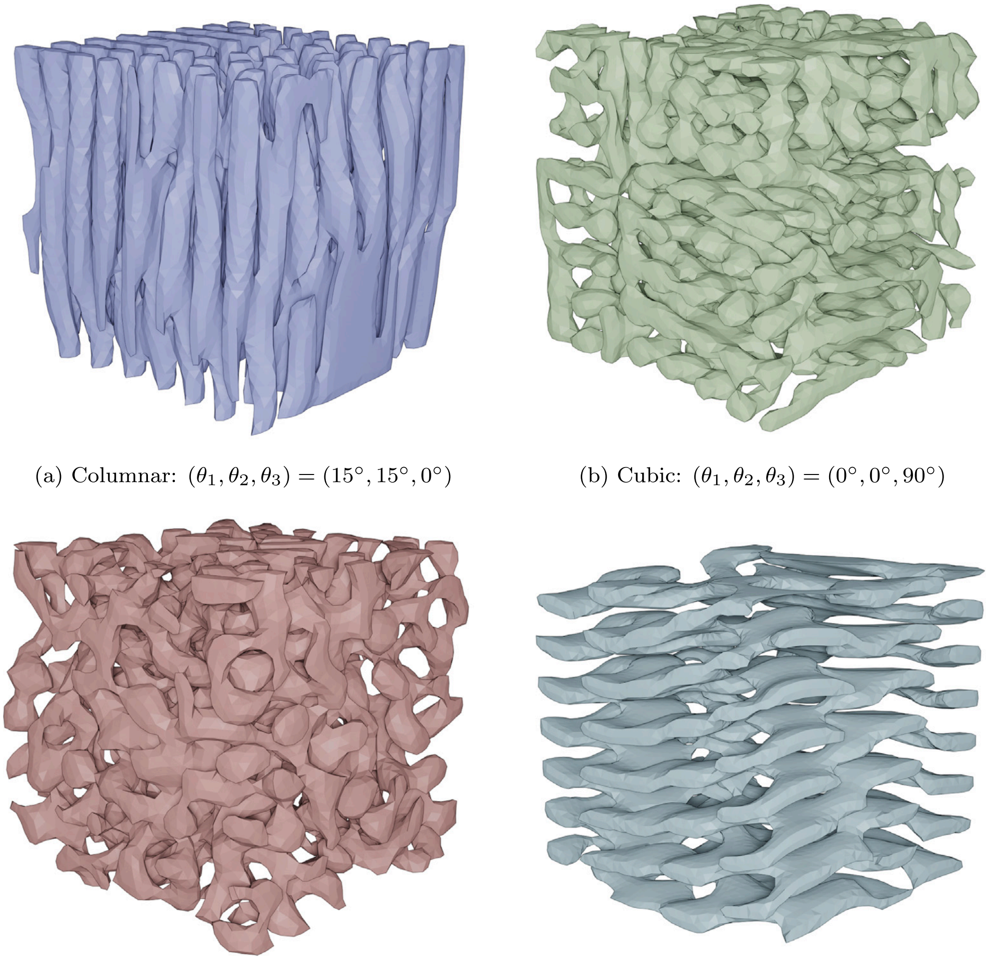

(a) Columnar: $(\theta_1, \theta_2, \theta_3) = (15°, 15°, 0°)$

(b) Cubic: $(\theta_1, \theta_2, \theta_3) = (0°, 0°, 90°)$

(c) Isotropic: $(\theta_1, \theta_2, \theta_3) = (20°, 20°, 20°)$

(d) Lamellar: $(\theta_1, \theta_2, \theta_3) = (0°, 0°, 30°)$

Fig. 4. Spinodoid structures with $\rho = 0.3$. Each caption includes the class name of the structure and the θ-angles that yield the structure.

In the extreme case where θ_3 is near the maximum value, the spinodoid structure becomes cubic in nature (Fig. 4(b)), while equal cone angles yield a more isotropic material configuration (Fig. 4(c)).

The spinodal structures, generated by means of a MATLAB script, are subjected to one-dimensional compressive loading simulations. These simulations are handled by the ABAQUS Explicit FE solver with tetrahedral (C3D4) elements. The simulation setup includes two rigid plates. One plate, acting as a loader, moves until a maximum strain of 50% is achieved, while the other serves as a stationary anvil. In addition, general contact is defined to simulate interactions between the plates and the structure as well as self-contact, enabling densification and preventing self-interaction during deformation. The constitutive model, which underlies the FE model, relating the Young's moduli, Poisson's ratios, and shear moduli, has been used before in literature. In practice, fused deposition modeling (FDM) is used to fabricate the structure. This type of additive manufacturing induces directional dependency into the material as the layers are printed. To capture this anisotropy, an orthotropic model is used, which allows the definition of Young's modulus along the printing (longitudinal) direction, as well as in the transverse direction. To capture the effect of plasticity, the Hill anisotropic criteria are used to determine the onset of plastic yielding. A comprehensive description is beyond the scope of this manuscript and can be found in the work by Kansara et al. [59].

As the compressive loading simulation is performed, a force P is applied across a displacement x until a threshold $\delta_{\max}$ is achieved. From this, a force-displacement behavior $P(x)$ can be observed. The relevant data-driven objective related to the scope of this problem is the normalized energy absorption value, which equals the integral of the force across the continuous displacement variable, divided by the EA value obtained from a cube made out of solid material with relative density ρ, denoted as EA_s.

$$\mathrm{EA} = \frac{1}{\mathrm{EA}_s} \int_0^{\delta_{\max}} P(x)\,\mathrm{d}x. \tag{14}$$

It has been experimentally observed that the amount of energy that the spinodoid material can absorb behaves non-trivially as a function of the θ-values [59]. For example, EA reaches a local maximum when θ_1 and θ_2 are around 45°, given $\theta_3 = 15°$. The spinodoid structure's energy absorption is therefore a suitable objective function f to consider as the target of an optimization (maximization) problem.

3.2. 4D data

The starting point of this data-driven analysis is the sampling of objective data. The sampling is performed such that it is compatible with performing SSA.

Naturally, the computational cost depends on the mesh resolution at which the spinodoid structure is generated; it is expected that the fine mesh (Fig. 5(b)) simulations will take longer than the coarse mesh (Fig. 5(a)) simulations. The mesh resolution is a dimensionless number that is used directly in the FE solver, where higher numbers indicate finer meshes. For this application, a bi-fidelity scenario is assumed, where the low-fidelity model corresponds to a mesh resolution of 20, and the high-fidelity one has a resolution of 30.

It should be noted that the simulation cost depends on the design parameters. To empirically estimate how much the computational cost differs between the objectives at the two fidelity levels, sets of Saltelli samples are used to generate a statistic of computational run times, presented in Fig. 6. For the high-fidelity objective, 96 design samples are used, whereas the low-fidelity objective is evaluated on the same design samples, along with an additional 288 samples, for a total of 386 samples.

Fig. 6 shows that the outliers in both fidelity cases are all located above the upper whisker. This underlines the high computational costs associated with simulating the loading experiment. Given the long simulation times compared to the duration of a BO iteration with the same computational power, it makes sense to consider Bayesian data-driven techniques to optimize the spinodal structural design for its energy absorption.

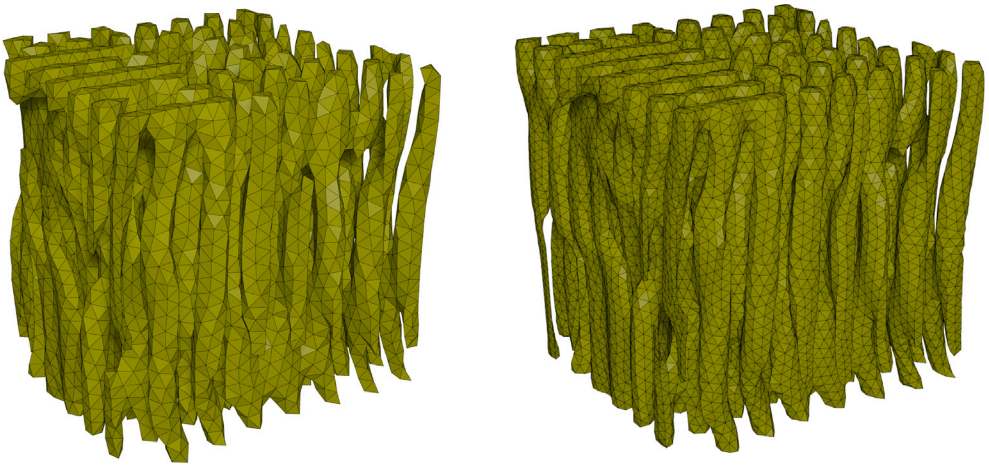

(a) Mesh resolution 20 (b) Mesh resolution 30

Fig. 5. Spinodal structure meshes at two different mesh resolutions for fixed values $\rho = 0.3$, $\theta_1 = 15°$, $\theta_2 = 15°$, and $\theta_3 = 0°$. The coarse mesh represents the low-fidelity model, while the fine mesh is considered high-fidelity. Intuitively, a low value for θ_3 arranges the structure into columns aligned in the loading direction.

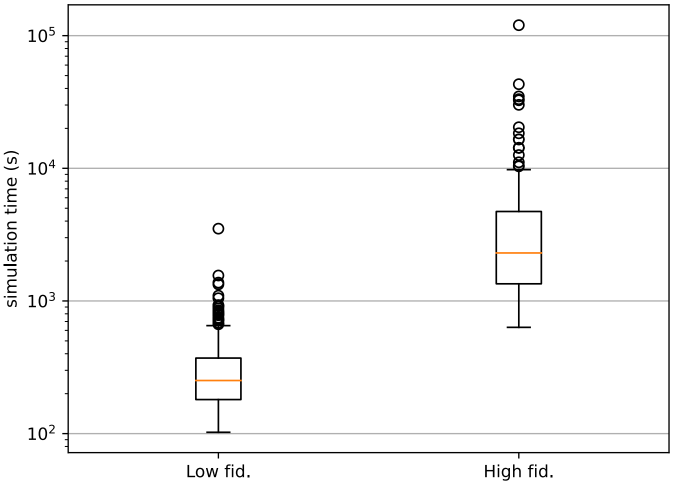

Fig. 6. Box plots of the distribution of simulation times for mesh resolution 20 (left) and 30 (right). The median simulation time when the coarser mesh is applied is 251 s, much lower than the median simulation time of the high-fidelity samples (2300 s). In each box plot, the blank dots correspond to the outliers, which do not fall within the 1.5-interquartile range with respect to the median.

Our data-driven workflow continues with SSA based on the previously obtained samples. It should be noted that the dimensionality of the design space (four) is low. Nonetheless, performing SSA is of vital importance because the relative density of the material ρ is a design parameter. Due to mass scaling effects [60] in finite element analysis, a significant reduction in computational time can be achieved. Combined with the fact that ρ can affect the energy absorption capabilities of the spinodoid structure, it is relevant to quantify the sensitivity of the objective towards this design parameter, also in comparison with the angular design parameters $\theta_1, \theta_2, \theta_3$.

To perform SSA thoroughly, sufficient samples are needed. To ascertain that this condition is met, the Saltelli samples are cumulatively applied to the analysis, giving rise to the sensitivity index convergence plot shown in Fig. 7.

It can be concluded from Fig. 7 that the sensitivity indices are stabilizing when more than 150 design samples are utilized. The sensitivity of both low- and high-fidelity EA objectives with respect to ρ is approaching 1, while the S_i and S_{T_i} values of θ_1, θ_2, θ_3 approach 0. This means that the material's energy absorption is highly dependent on its relative density, especially compared to the geometric parameters.

3.3. 3D data

Given the findings related to the 4D design space SSA shown in Fig. 7, the angular design parameters do not clearly affect the objective. Therefore, the next step is to select a fixed value for ρ to observe the residual variance-based sensitivity among $\theta_1, \theta_2, \theta_3$. From a physics-driven perspective and previously established numerical experiments [59], BO will quickly converge to designs with a value of $\rho = 0.6$, due to those allowing for high values of energy absorption. This is confirmed by additional data analysis in Fig. C.15. However, a higher value for ρ corresponds to computationally more expensive FE simulations due to mass scaling. For this study, therefore, the value $\rho = 0.3$ is chosen. After sampling the reduced 3D design space, it was found that the median high-fidelity simulation time equals 780 s, while the low-fidelity simulations take 150 s on average. This allows for the faster procurement of data-driven results, while acknowledging that the optimization of ρ is trivial.

The fixture of ρ reduces the 4D optimization problem to one with a 3D design space. This new design space is treated with the same sensitivity index convergence analysis method that gave rise to Fig. 7 in the 4D scenario. Due to the perceived lack of sensitivity of EA with respect to $\theta_1, \theta_2, \theta_3$ from Fig. 7, even more samples are expected to be required to achieve stabilizing sensitivity indices. The resulting observations are shown in Fig. 8.

From Fig. 8, it can be seen that the angular design parameters that give rise to the spinodal structure are all approximately equally sensitive towards energy absorption. Moreover, it is noted that the sensitivity indices take longer to stabilize. With both mesh resolutions, the indices only stabilize when more than 400 design samples are used, a significant increase compared to the 150 needed for the case with the 4D design space. This indicates the difficulty for the SSA procedure in determining the sensitivity indices due to less structure in the overall data.

Furthermore, for the sake of applying the VF-EI acquisition function (Eq. (9)) during MFBO, it is of interest to investigate the global level of correlation between the low- and high-fidelity data points. In order to ensure the absence of any bias, this inter-fidelity correlation is calculated once, before the optimization procedure. See Fig. 9.

As is visible from Fig. 9, a cluster of designs with low-fidelity EA values at around 0.018 exists due to a numerical artifact resulting from the FE meshing procedure, resulting in the same EA value for a range of different θ-values. This cluster was removed before the correlation coefficient was calculated. Furthermore, a moderate level of (linear) correlation between the low- and high-fidelity objectives can be observed. This motivates the application of multi-fidelity Bayesian techniques, due to the transfer learning potential from the low- to high-fidelity levels. The high computational cost, in addition to the observed non-triviality of the EA objective landscape with respect to the angular design parameter, makes the reduced three-dimensional problem interesting for Bayesian optimization.

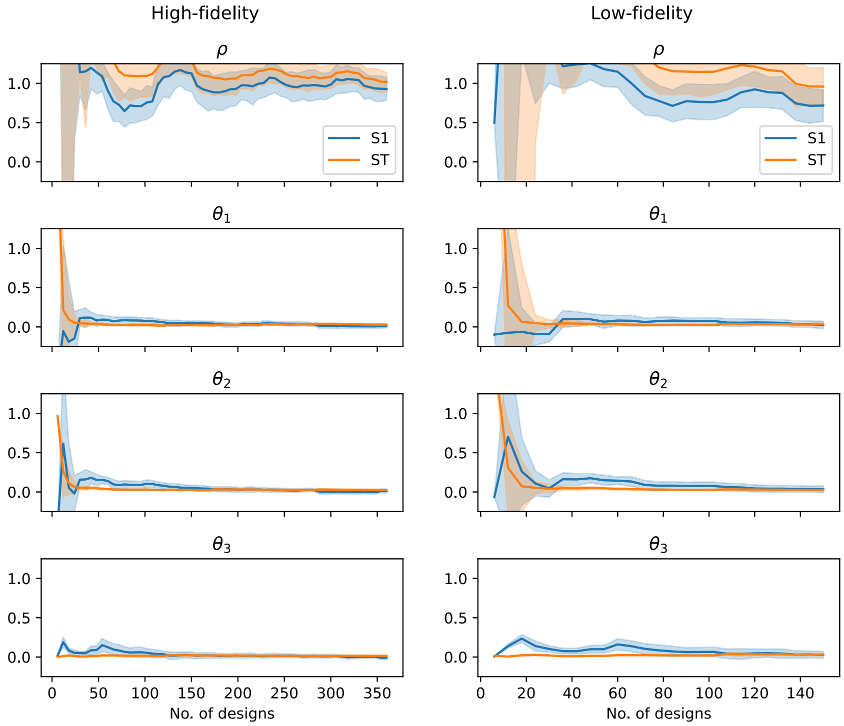

Fig. 7. First (S1) and total (ST) order Sobol' sensitivity indices and corresponding bootstrapped confidence interval per design parameter as a function of the number of Saltelli designs.

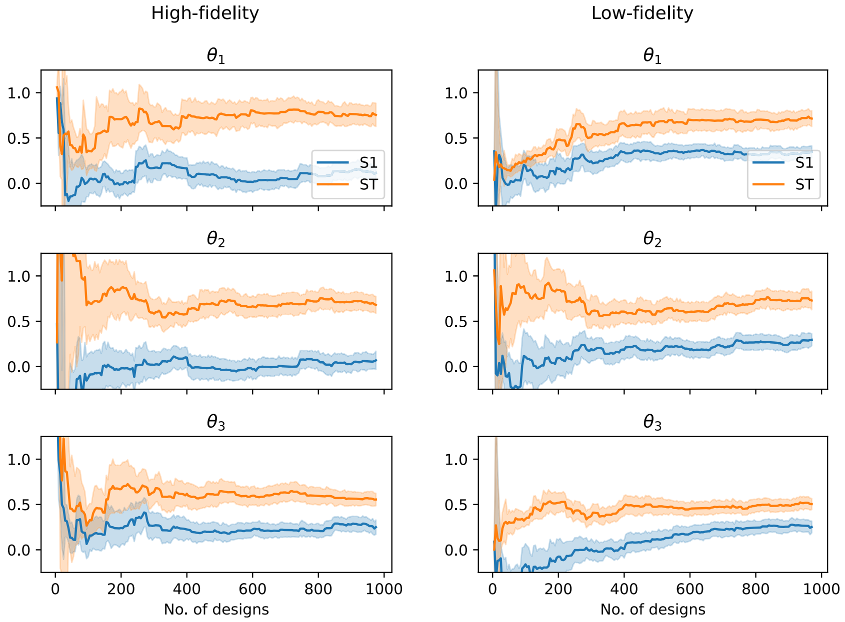

Fig. 8. First (S1) and total (ST) order Sobol' sensitivity indices of the 3D EA objective (fixed $\rho = 0.3$), and corresponding bootstrapped confidence interval per design parameter as a function of the number of Saltelli designs. The discrepancy between the first and total-order indices for all design parameters indicates that higher-order joint effects are statistically significant with respect to varying the objective.

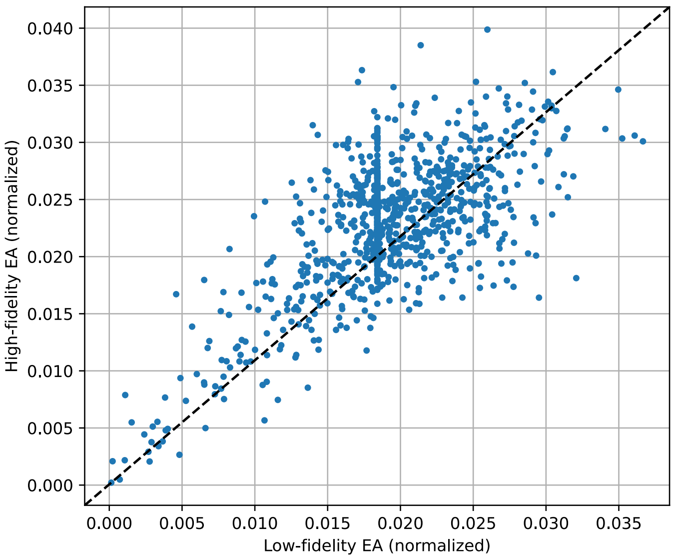

Fig. 9. Low- and high-fidelity EA values at the coincident 3D Saltelli design samples used in Fig. 8. The correlation coefficient between the two sets of data points is approximately 0.68.

4. Optimization results

BO (Algorithm 1) and MFBO (Algorithm 2) are applied to optimize the reduced, 3D spinodoid material design. A total of four combinations of (base) covariance functions (Matérn kernel, RBF kernel) and (single- or multi-fidelity) acquisition functions (Logarithmic EI, UCB) are applied to maximize the normalized (high-fidelity) EA. The Adam optimizer shall be used with 10 random restarts to resolve the inner optimization loops, i.e., the numerical optimization of the likelihood as well as optimization of the acquisition function. All design parameter and objective data are assumed to be normalized to conform to the zero mean assumption of any (MF) GPR model.

It should be noted that the selection of these hyperparameters is not, by any means, based on the recommendations of any hyperparameter optimization scheme. This reasoning stems from the primary nature of the comparison between BO and MFBO: similar kernel and acquisition types are used throughout the optimization experiment runs, so the single- and multi-fidelity optimization schemes are being compared on an equal footing.

Furthermore, the stopping criterion of (MF)BO as described in Algorithms 1 and 2, is solely based on a budget of high-fidelity equivalents. It is possible to define a stopping criterion that is based on the convergence of the incumbent optimum objective value. However, the focus of our work is placed on the comparison between single- and multi-fidelity optimization methods in a real-world, high-cost, and budget-constrained scenario. Furthermore, since the optimized value of the EA objective is unknown, it is infeasible to assess at which level this convergence threshold should be located. Hence, no explicit convergence criterion was used throughout the optimization runs.

For this experiment, the computational expenditure for the initial DoE $\mathbf{D}^{(0)}$ is equivalent to 160 high-fidelity evaluations, while the optimization budget is $I = 50$ high-fidelity equivalents. The samples that populate this DoE are reused from the 3D sensitivity index convergence analysis performed in Fig. 8 of Section 3.3. The optimization histories are shown in Fig. 10.

Fig. 10 shows the expended cost on the horizontal axis per unit of time needed to perform one high-fidelity FE simulation. While this unit is not a constant value, as Fig. 6 shows, an average (median) wall-clock simulation time of 2300 s can be assumed.

As explained in Section 3, it is possible to exploit the cost-accuracy trade-off that naturally arises with mesh resolution when FE simulations are concerned. This corresponds to a small low-to-high fidelity cost ratio of 0.19, which is obtained by dividing the average simulation time of a low-fidelity (150 s) and a high-fidelity simulation (780 s). It is therefore sensible to use multi-fidelity Bayesian optimization (Algorithm 2) for solving the energy absorption maximization problem.

The goal of this numerical experiment is to compare BO to MFBO, so a total computational cost equivalent to 160 high-fidelity evaluations is spent on the initial DoE $\mathcal{D}^{(0)}$. Eighty samples are inherited from high-fidelity Saltelli designs used in the initial DoE $\mathbf{D}^{(0)}$ of the single-fidelity BO scenario. Four hundred samples are reused from the previously obtained Saltelli samples that resulted in the low-fidelity EA sensitivity analysis of Fig. 8. The same optimization budget B of 50 high-fidelity evaluation equivalents is used, compared to the single-fidelity BO experiment in Fig. 10. Lastly, since B is in units of high-fidelity evaluations, the cost function c and the cost ratio CR are equal: i.e., $c\ (m = 1) = 0.19$ and $c\ (m = 2) = 1$. The multi-task Bayesian optimization (MTBO) objective histories are shown in Fig. 11.

Resolving the likelihood optimization problem to train each MTGPR at every MTBO iteration took, on average, 10 wall-clock seconds. This is insignificant compared to the expensive FE simulations, even in the low-fidelity scenario.

The resulting optimized designs by the BO and MFBO runs after the computational budget was expended are reported in Table 2.

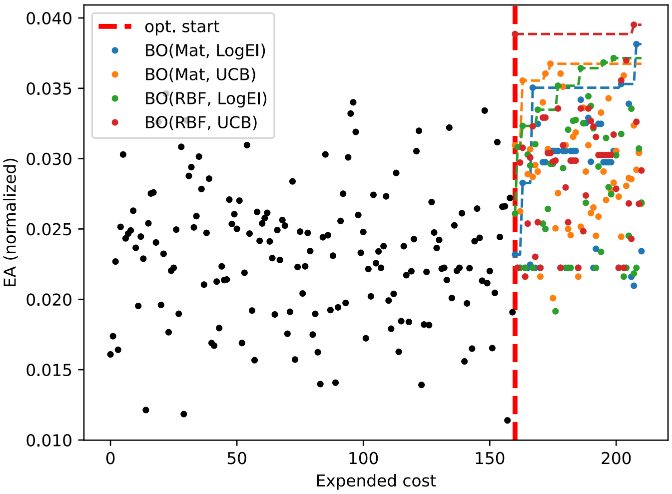

Fig. 10. Several standard BO histories (colored dots) and each cumulative maximum (colored dashed line). The black dots correspond to the common initial DoE, whereas the vertical thick red dashed line indicates the cost value at which BO was initiated. The optimization was allowed to run for a computational cost equivalent to 50 high-fidelity evaluations.

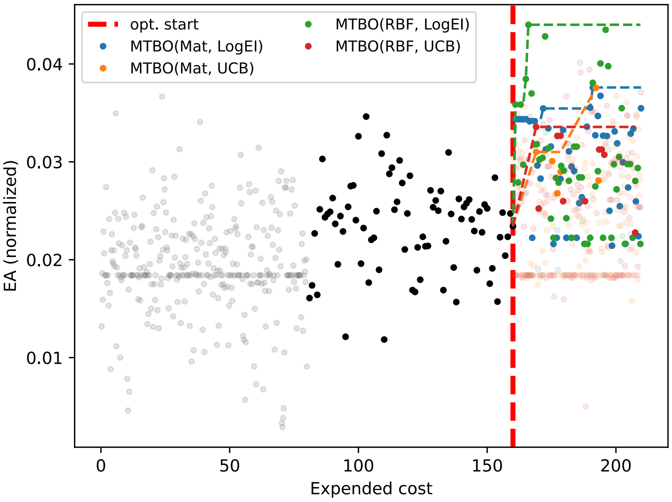

Fig. 11. Several EA optimization histories with multi-task Bayesian optimization. The solid dots indicate high-fidelity data points, while the faded dots correspond to low-fidelity data. The evaluated EA value is plotted across the expended computational cost in units of high-fidelity data points. The black dots represent the initial high-fidelity data, shared across all BO runs.

Table 2
Optima to the 3D EA optimization problem found by various single- and multi-fidelity BO configurations. The best single- and multi-fidelity designs are indicated in boldface.

Acquisition	Kernel	Optimized values			
		θ_1	θ_2	θ_3	EA (normalized)
α_{LogEI}	κ_{Mat}	46.6°	56.5°	10.0°	0.038139
	κ_{RBF}	41.7°	56.7°	0.0°	0.037146
α_{UCB}	κ_{Mat}	36.8°	62.9°	0.0°	0.037123
	κ_{RBF}	**37.6°**	**61.7°**	**0.0°**	**0.039517**
$\alpha_{\text{VF-LogEI}}$	κ_{Mat}	44.1°	46.2°	8.9°	0.040766
	κ_{RBF}	**44.4°**	**43.9°**	**2.4°**	**0.044007**
$\alpha_{\text{VF-UCB}}$	κ_{Mat}	41.6°	50.2°	3.9°	0.040872
	κ_{RBF}	44.2°	45.2°	6.0°	0.038878

From Table 2, it can be seen that the values that multi-task BO recommends are generally better than the single-fidelity BO counterparts. Importantly, it can be seen that there is an 11% overall increase in the optimally located normalized EA value across all kernels and acquisitions, while there is an average 8.6% increase among methods with the same covariance and acquisition function types. This latter fact can likely be attributed to the knowledge transfer that occurs from low- to high-fidelity. Barring the possibility that this improvement falls within the numerical uncertainty due to the FE simulation procedure, this is supported by the substantial correlation between the low- and high-fidelity data, as noted from Fig. 9. Moreover, it should be noted that three out of four values in Table 2 resemble columnar structures, with optimal values of $\theta_3 = 0°$. This stands in contrast to the slightly higher values of θ_3 found by all of the multi-task Bayesian optimizers ($2° \sim 9°$). These values for θ_3 could be indicative of some degree of horizontal reinforcement between the column-like structures. These structures are bending-dominated structures, which are known to be more compliant and dissipate energy more evenly in the loading direction [61]. In particular, the horizontally reinforced structure will buckle in a progressive, sequential manner. This is a desirable property for energy absorption, compared to the other spinodoid structure classes (such as a lamellar configuration) that exhibit strong deformations under similar load paths.

The previously presented results show a benefit in the energy absorption when MFBO is applied. To investigate whether the multi-fidelity optimization method converges faster than its single-fidelity counterpart, Fig. 12 shows a direct comparison between the histories of the best-performing BO algorithm and the best-performing MFBO algorithm.

Fig. 12 encapsulates the main motivation to use MFBO over BO. It takes 46 high-fidelity cost units for the best-performing BO scheme to achieve (negligible) improvement in the EA value with respect to the initial data. However, the best-performing MTBO scheme can suggest a similar (or superior) design within 5 cost units. Moreover, it can suggest four more designs with a better EA value within the subsequent 30 units of cost. On average, MFBO achieves these design parameters after 33 cost units, while it takes BO to reach its top five design recommendations after 41 cost units. This corresponds to a difference of 8 cost units on a budget of 50, i.e., an efficiency gain of 16%. In summary, MFBO

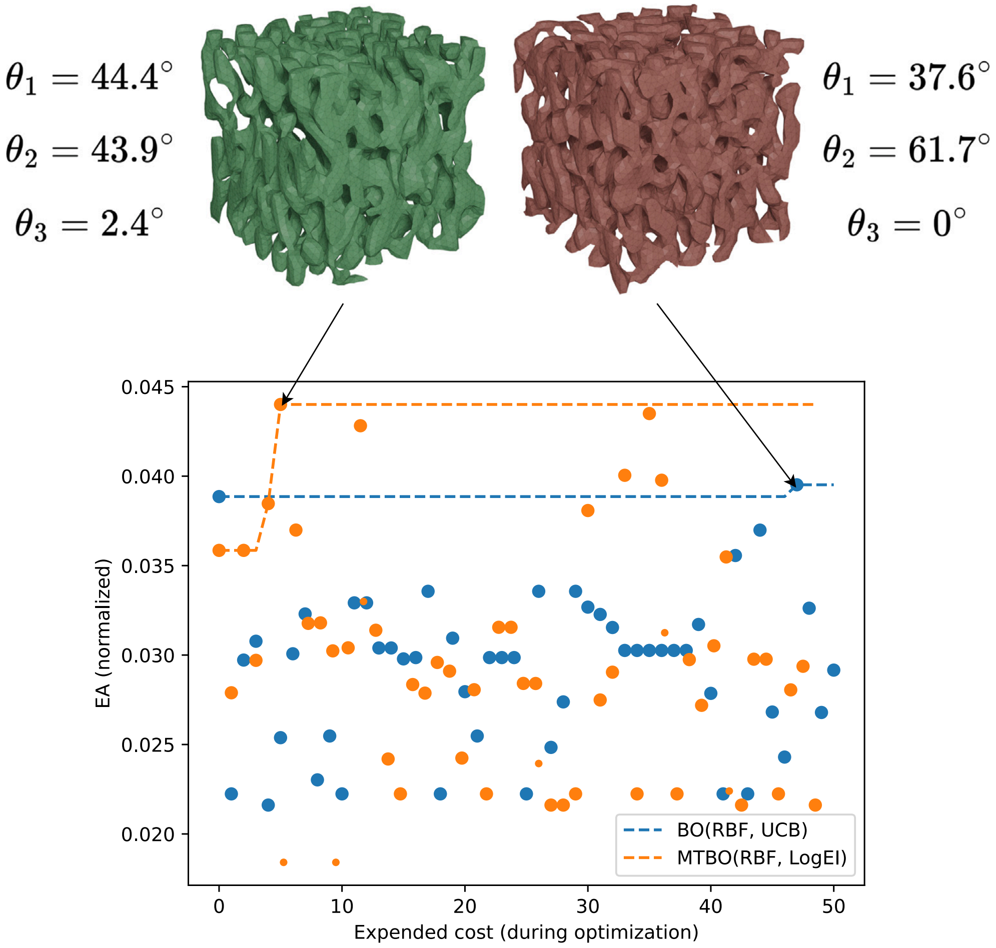

Fig. 12. The EA-optimal spinodoid structures and θ-values (Tables 2) found by the best-performing BO (top right) and the best MTBO (top left) schemes, superimposed with the optimization histories. Low-fidelity EA evaluations during the MTBO run are indicated with the smaller orange dots.

can outperform BO in terms of computational efficiency, given the same budget.

5. Future work

While the novel Bayesian data-driven framework has yielded promising results, we recognize various potential extensions to this research. These are suggested below.

- **Expanding the optimization benchmark**, This paper presents a comparison between single- and multi-fidelity BO. The results show the effectiveness of MFBO to solve a multi-fidelity design problem by leveraging transfer learning from a lower fidelity level. However, it is possible to expand this framework to more critically assess the benefit of multi-fidelity methods. Apart from single-fidelity BO, there exist several popular global, gradient-free optimization methods that can be used. Examples include schemes such as Nelder-Mead [62] or evolutionary algorithms. One of the most promising such heuristics, called the covariance matrix adaptation evolution strategy (CMA-ES) [63], has previously received accolades in competitive objective optimization settings [64].
- **Empirical-statistical validation**. The results presented in Figs. 10 and 11 are based on single optimization runs. Despite the observed increase in EA values found by the best-performing MFBO scheme, this does not warrant the conclusion that it outperforms the single-fidelity counterpart. This is partly due to the lack of multiple optimization runs per optimization type: a different initial DoE might show a more modest improvement by using multi-fidelity. Performing similar optimization procedures on several different initial DoE configurations will be able to provide more meaningful statistical insight into whether MFBO outperforms BO in terms of objective value or computational efficiency.

  The discretization of the low-fidelity FE model has led to data clustering, as shown in Fig. 9. This could potentially be a negative influence on the quality of the MOGPR models used during MFBO. The authors recommend the consideration of data filtering to alleviate such clustering effects.
- **Experimental validation**. The optimized results in this manuscript have been purely based on simulation results. Given the precedent concerning additively manufactured spinodoid structures [1], it is desirable to experimentally validate the optimized design parameters by means of a real-life loading test on a manufactured spinodoid sample.
- **Design and fidelity space considerations**. Following the most popular convention in multi-fidelity literature, a total of $M = 2$ fidelity levels were assumed to be part of the multi-fidelity data structure. However, the effect of an increase in M is yet to be studied in the context of the energy absorption maximization problem. For example, how do the MTBO results presented earlier compare to a scenario with an intermediate fidelity level ($M = 3$)?

  Furthermore, the Sobol' sensitivity analysis on the 4D case has showcased a significant skew towards reliance on the relative density of the material. While this allowed for the simplification of the design problem at hand, it will be especially interesting to apply MFBO to a design problem with a higher design space dimensionality. One possibility is to consider graded spinodoids [65].
- **Multi-objective considerations**. While BO and MFBO have proven to optimize the energy absorption fairly well across the design space, it is not the only objective of interest in this class of material design problems. Minimizing the peak force (PF) that the material bears during loading is another important objective. One way to handle multiple objectives by means of a single-objective method is by scalarizing the EA and PF values into one objective [66,67]. This scalarized objective can then be optimized by way of a single-fidelity or multi-fidelity BO method.

  Lastly, owing to the success of applying multi-objective Bayesian optimization to cooling duct design [68] and a similar spinodoid design problem [59], it has been of interest to combine multi-fidelity and multi-objective BO (MOBO) into one framework. Previously, literature has shown that it is possible to translate an MFBO problem into an MOBO problem by introducing a trust function [69]. Heuristically speaking, a trust function measures the level of validity to use a low-fidelity objective to model a high-fidelity objective, similar to the usage of Pearson's correlation coefficient in Fig. 9. In addition, the correlation between the objectives at different fidelity levels might be dependent on the individual design parameters, whereas this study makes use of a global measure of linear correlation to estimate the similarity between low- and high-fidelity EA.

6. Conclusions

This work introduces a practical, robust, and efficient Bayesian data-driven design framework. This novel framework combines proven and established Bayesian methods, variance-based sensitivity analysis, and BO into one unified pipeline. This framework is used to provide a direct comparison between the performance of BO and MFBO with similar kernel and acquisition types. The engineering use case to which this is applied pertains to optimizing the energy absorption of spinodoid structures. To this end, a potential set of design parameters is established, and variance-based sensitivity analysis shows that the dimensionality of the design space can be reduced due to the strong and monotonous dependency of EA on the density ρ. Subsequently, on the reduced design space by fixing $\rho = 0.3$, both BO with GPR as well as MTBO are applied to search for the design configuration with the highest energy absorption, given an initial DoE containing samples that have also been used for SSA. It is shown that, with the given computational budget constraints, MTBO is able to yield similar results or outperform single-fidelity BO across all the hyperparameter variations that have been probed. This raw improvement in terms of EA value cannot be decisively ascribed to transfer learning due to stochastic effects. However, even if the EA values found by the best performing BO and MFBO schemes are the same, the multi-fidelity scheme can achieve this with an average 16% cut in terms of high-fidelity cost units.

CRediT authorship contribution statement

Leo Guo: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Hirak Kansara:** Writing – review & editing, Visualization, Validation, Software, Resources, Investigation, Data curation, Conceptualization. **Siamak F. Khosroshahi:** Resources, Methodology, Conceptualization. **GuoQi Zhang:** Supervision, Funding acquisition. **Wei Tan:** Writing – review & editing, Supervision, Project administration, Investigation, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

W. Tan reports that financial support was provided by EPSRC New Investigator Award. G.Q. Zhang reports that financial support was provided by ECSEL Joint Undertaking. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

This work was supported by the ECSEL Joint Undertaking (JU) under Grant 826417. The JU receives support from the European Union's

Horizon 2020 research and innovation program and Germany, Austria, Spain, Finland, Hungary, Slovakia, the Netherlands, and Switzerland.

Wei Tan, Hirak Kansara, and Siamak F. Khosroshahi furthermore acknowledge the financial support from the EPSRC New Investigator Award (grant No. EP/V049259/1). This work was also selected by the ERC and funded by UK Research and Innovation (UKRI) under the UK government's Horizon Europe funding guarantee (No. EP/Y037103/1).

Lastly, the authors extend much gratitude to Dr. Miguel A. Bessa for the valuable discussions, insights, and support throughout the project leading to this manuscript.

Appendix A. Saltelli sampling

Let $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N \in [0,1]^{2D}$ be the first N terms of a $2D$-Sobol' sequence. To be explicit, let $\mathbf{x}_j := (x_{j,1}, x_{j,2}, \ldots, x_{j,2D})^\top$ be the components of $\mathbf{x}_j$ for any $j \in \{1, \ldots, N\}$. Next, define

$$\mathbf{A} := (\mathbf{a}_1^\top, \ldots, \mathbf{a}_N^\top)^\top \text{ where } \mathbf{a}_j := (x_{j,1}, \ldots, x_{j,D})^\top, \tag{A.1}$$

$$\mathbf{B} := (\mathbf{b}_1^\top, \ldots, \mathbf{b}_N^\top)^\top \text{ where } \mathbf{b}_j := (x_{j,D+1}, \ldots, x_{j,2D})^\top. \tag{A.2}$$

Next, for any $i \in \{1, \ldots, D\}$, define the matrix $\mathbf{A}_\mathbf{B}^{(i)}$ as the $N \times D$ matrix containing coordinate-wise cross-combinations of rows that populate $\mathbf{A}$ and $\mathbf{B}$:

$$(\mathbf{A}_\mathbf{B}^{(i)})_{j,k} := \begin{cases} a_{j,k}, & k \neq i, \\ b_{j,k}, & k = i, \end{cases} \tag{A.3}$$

where $j \in \{1, \ldots, N\}$, $k \in \{1, \ldots, D\}$ and the component notations are derived from $\mathbf{a}_j =: (a_{j,1}, \ldots, a_{j,D})^\top$ and $\mathbf{b}_j =: (b_{j,1}, \ldots, b_{j,D})^\top$. See Fig. A.13 for a graphic illustration on how to construct $\mathbf{A}_\mathbf{B}^{(i)}$.

Finally, let $\mathbf{A}_\mathbf{B}^{(i)} := (\mathbf{c}_1^{(i)\top}, \ldots, \mathbf{c}_N^{(i)\top})^\top$ be decomposed into rows. Then, the following approximations hold:

$$\mathrm{Var}(\mathbb{E}(Y|X_i)) \approx \frac{1}{N} \sum_{j=1}^{N} f(\mathbf{b}_j)(f(\mathbf{c}_j^{(i)}) - f(\mathbf{a}_j)), \tag{A.4}$$

$$\mathbb{E}(\mathrm{Var}(Y|\mathbf{X}_{\sim i})) \approx \frac{1}{2N} \sum_{j=1}^{N} (f(\mathbf{a}_j) - f(\mathbf{c}_j^{(i)}))^2. \tag{A.5}$$

Moreover, the statistical confidence of these approximations can be numerically tested by means of bootstrapping [35,70].

An aggregate of all design samples necessary to perform SSA, i.e., $\mathbf{A}, \mathbf{B}$ and $\mathbf{A}_\mathbf{B}^{(i)}$ for all $i \in \{1, \ldots D\}$, is called a collection of Saltelli samples, where it should be noted that the quantity of these samples is always a multiple of $D + 2$. An example of two-dimensional Saltelli sampling is given in Fig. A.14.

From Fig. A.14, it can be seen that the sequentially dependent distribution of the Saltelli samples fills the space. This makes Saltelli sampling the preferred method to access samples, not only to perform sensitivity analysis, but also to be used as a data set to build a regression model on, particularly as a central step in the Bayesian optimization algorithm.

Appendix B. Gaussian process regression

In literature, a wide variety of regression models exist to model the relationship between $\mathbf{X}$ and $\mathbf{y}$. In terms of complexity, models range from linear ones such as ordinary least squares, ridge [71] or lasso [72] regression to the highly non-linear artificial neural network (ANN) models. Among these, some regression models have been successfully employed from which posterior information can be extracted in order to suggest

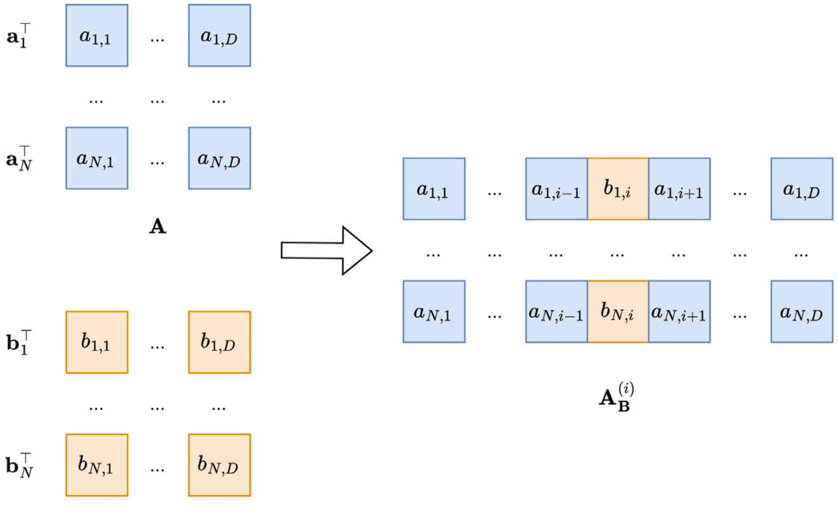

Fig. A.13. A schematic construction of $\mathbf{A}_\mathbf{B}^{(i)}$ given $\mathbf{A}$ and $\mathbf{B}$.

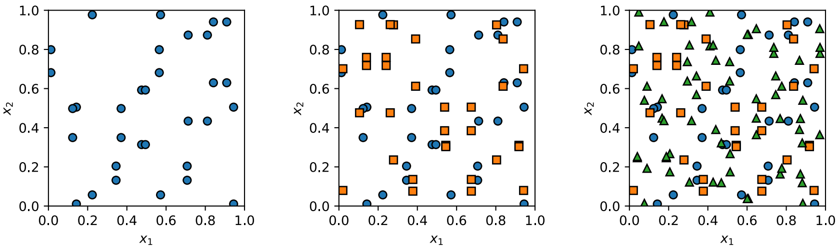

Fig. A.14. Saltelli sampling on a unit square with 32, 64, and 128 subsequent samples (different colors and shapes). These samples are respectively based on 8, 16, and 32 subsequent Sobol' samples in 4D space.

new data in the Bayesian optimization. Examples include random forest regression [73], tree-structured Parzen estimators [74], and Bayesian ANNs (BNNs) [75]. However, as the prominent reviewers of Bayesian optimization methods have iterated [43,44], the Gaussian process model is the most popular type of model used in practice to perform regression within the Bayesian optimization context.

B.1. Single-fidelity Gaussian process regression

Given a design of experiments $\mathbf{D}$ with N experiments, it is assumed that $\mathbf{X}$ is an independent variable with respect to an explanatory model, while $\mathbf{y}$ is a dependent variable.

In the context of Bayesian statistics, it is assumed that any given set of evaluations of f is the realization of a jointly distributed random variable. Specifically, f could be modeled as a Gaussian process (GP) with zero[2] prior mean. Concretely, for any $\mathbf{x} \in [0,1]^D$, it is assumed that $f(\mathbf{x})$ is a random variable with $\mathbb{E}(f(\mathbf{x})) = 0$ and prior covariance function or kernel κ; for any $\mathbf{u}, \mathbf{v} \in [0,1]^D$, the defining property is as follows:

$$\kappa(\mathbf{u},\mathbf{v}) := \mathrm{Cov}(f(\mathbf{u}), f(\mathbf{v})). \tag{B.1}$$

Let $\theta_1, \dots, \theta_{T_\kappa}$ be the GP model parameters, collected into a model parameter vector $\boldsymbol{\theta} = (\theta_1, \dots, \theta_{T_\kappa})$. Since $\boldsymbol{\theta}$ contains precisely the model parameter coordinates on which κ depends, the notation $\kappa(\mathbf{u},\mathbf{v}) = \kappa_{\boldsymbol{\theta}}(\mathbf{u},\mathbf{v})$ arises.

A popular type of prior covariance function is the radial basis kernel (RBF) kernel [39,77,78], given by:

$$\kappa_{\mathrm{RBF},\boldsymbol{\theta}_{\mathrm{RBF}}}(\mathbf{u},\mathbf{v}) := c \cdot \exp\left(-\frac{\|\mathbf{u}-\mathbf{v}\|^2}{2\lambda^2}\right) + s^2\delta(\mathbf{u}-\mathbf{v}), \tag{B.2}$$

where c is called the amplitude parameter, λ is the length scale parameter and s^2 is the aleatoric noise variance parameter. The noise variance, as the name suggests, only affects the scenario in which the covariance is calculated for $\mathbf{u} = \mathbf{v}$ as indicated by the Dirac-δ function. In summary, the GP model parameters associated with the RBF kernel are $\boldsymbol{\theta}_{\mathrm{RBF}} := (c, \lambda, s^2)^\top$.

As dependent variables in a Bayesian context, it is assumed that $\mathbf{y}$ is a realization of a random vector $\mathbf{Y}$, and that any model parameter vector $\boldsymbol{\theta}$ is a realization of a random vector $\boldsymbol{\Theta}$. The Gaussian process assumption implies that $\mathbf{Y}$ follows a multivariate normal distribution when a model parameter vector $\boldsymbol{\theta}$ is established. In statistical terms: the prior distribution of $\mathbf{Y}$ conditioned on $\boldsymbol{\Theta} = \boldsymbol{\theta}$ is given by

$$\mathbf{Y}|(\boldsymbol{\Theta} = \boldsymbol{\theta}) \sim \mathcal{N}(\mathbf{0}, \mathbf{K}_{\boldsymbol{\theta}}), \tag{B.3}$$

with covariance matrix $\mathbf{K}_{\boldsymbol{\theta}} = \mathbf{K}_{\boldsymbol{\theta}}(\mathbf{X}) := (\kappa_{\boldsymbol{\theta}}(\mathbf{x}_i, \mathbf{x}_j))_{i,j=1,\dots,N}$.

For any given $\mathbf{x} \in [0,1]^D$, it is now of interest to find the distribution of $f(\mathbf{x})$ given $\mathbf{y}$ and $\boldsymbol{\theta}$, which is also called the (parameterized) posterior distribution. Here lies the crux of assuming f to be a GP, since $(\mathbf{Y}^\top, f(\mathbf{x}))^\top$ is another multivariate normal random variable. Then, following some mathematical manipulations [39,79], it can be shown that the parameterized posterior is also multivariate normal:

$$f(\mathbf{x})|(\mathbf{Y} = \mathbf{y}, \boldsymbol{\Theta} = \boldsymbol{\theta}) \sim \mathcal{N}(\mu_{\boldsymbol{\theta}}(\mathbf{x}), \sigma^2_{\boldsymbol{\theta}}(\mathbf{x})), \tag{B.4}$$

where

$$\mu_{\boldsymbol{\theta}}(\mathbf{x}) := \kappa_{\boldsymbol{\theta}}(\mathbf{x},\mathbf{X})^\top \mathbf{K}_{\boldsymbol{\theta}}^{-1}\mathbf{y}, \tag{B.5}$$

$$\sigma^2_{\boldsymbol{\theta}}(\mathbf{x}) := \kappa_{\boldsymbol{\theta}}(\mathbf{x},\mathbf{x}) - \kappa_{\boldsymbol{\theta}}(\mathbf{x},\mathbf{X})^\top \mathbf{K}_{\boldsymbol{\theta}}^{-1}\kappa_{\boldsymbol{\theta}}(\mathbf{x},\mathbf{X}), \tag{B.6}$$

with $\kappa_{\boldsymbol{\theta}}(\mathbf{x},\mathbf{X}) := (\kappa_{\boldsymbol{\theta}}(\mathbf{x},\mathbf{x}_1), \dots, \kappa_{\boldsymbol{\theta}}(\mathbf{x},\mathbf{x}_N))^\top$ as a column vector of length N. It is also convenient to denote $\mu_{\boldsymbol{\theta}}(\mathbf{Z}) := (\mu_{\boldsymbol{\theta}}(\mathbf{z}_1), \dots, \mu_{\boldsymbol{\theta}}(\mathbf{z}_P))^\top$ as a column vector of length P whenever $\mathbf{Z} = (\mathbf{z}_1^\top, \dots, \mathbf{z}_P^\top)^\top$ is a column of P design rows. The same notation is introduced for $\sigma^2_{\boldsymbol{\theta}}$.

While Eq. (B.4) provides a closed form description of the (parameterized) posterior distribution of $f(\mathbf{x})|(\mathbf{Y} = \mathbf{y}, \boldsymbol{\Theta} = \boldsymbol{\theta})$, it is desirable to obtain the distribution of $f(\mathbf{x})|(\mathbf{Y} = \mathbf{y})$, also called the posterior predictive distribution (PPD), independent of the model parameter vector $\boldsymbol{\theta}$.

In practice, the PPD cannot be calculated analytically [39]. In order to handle this, the dependency on the model parameter vector $\boldsymbol{\theta}$ of the intermediate posterior predictive distribution is removed by approximating it with a point estimate.

One choice for such an estimate, in the context of Gaussian process modeling, is the maximum (log) likelihood estimate (MLE) [39]. By defining the model parameter likelihood function as

$$\ell(\boldsymbol{\theta};\mathbf{D}) := p(\mathbf{Y}|\boldsymbol{\Theta} = \boldsymbol{\theta})[\mathbf{y}], \tag{B.7}$$

the MLE is given by

$$\hat{\boldsymbol{\theta}}_{\mathrm{MLE}} = \mathrm{argmin}_{\boldsymbol{\theta}\in\mathcal{T}_\kappa} \ln(\det(\mathbf{K}_{\boldsymbol{\theta}}(\mathbf{X}))) + \mathbf{y}^\top \mathbf{K}_{\boldsymbol{\theta}}^{-1}(\mathbf{X})\mathbf{y}, \tag{B.8}$$

where $\mathcal{T}_\kappa := P_1 \times P_2 \times \cdots \times P_{T_\kappa} \subset \mathbb{R}^{T_\kappa}$ is the T_κ-dimensional domain of model parameter vectors $\boldsymbol{\theta}$ and where $P_t \subset \mathbb{R}$ is defined to be the domain of all permissible values of θ_t, $t \in \{1, \dots T_k\}$.

Because ℓ equals the probability density function of a normal distribution as a function of $\mathbf{y}$, this likelihood function can also be referred to as a Gaussian likelihood function.

Some calculations yield the following approximation for the PPD density:

$$\begin{aligned} p(f(\mathbf{x})|\mathbf{Y} = \mathbf{y})[y] &\approx p(f(\mathbf{x})|\mathbf{Y} = \mathbf{y}, \boldsymbol{\Theta} = \hat{\boldsymbol{\theta}}_{\mathrm{MLE}}) \\ [y] &= \phi_{\hat{\mu}(\mathbf{x}),\hat{\sigma}^2(\mathbf{x})}(y) =: \hat{\phi}(y; \mathbf{x}, \mathbf{y}, \kappa), \end{aligned} \tag{B.9}$$

where

$$\phi_{\mu,\sigma^2}(z) := \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2}\left(\frac{z-\mu}{\sigma}\right)^2\right)$$

denotes the normal probability density function with mean μ and variance σ^2. Furthermore, the notations

$$\hat{\mu} := \mu_{\hat{\boldsymbol{\theta}}_{\mathrm{MLE}}}, \quad \hat{\sigma}^2 := \sigma^2_{\hat{\boldsymbol{\theta}}_{\mathrm{MLE}}} \tag{B.10}$$

are used for the optimized parameterized posterior parameters in Eqs. (B.5) and (B.6).

Commonly, it is assumed that the model parameter likelihood ℓ, Eq. (B.7), is differentiable with respect to the model parameter vector $\boldsymbol{\theta}$. Then, due to the differentiable nature of the optimization problem in Eq. (B.8), it is most often solved by using gradient-based methods such as L-BFGS or Adam.

B.2. Multi-output Gaussian process regression

Given the scarcity of data to construct a design of experiments $\mathbf{D}$, a common solution is to acquire higher-throughput data, at a risk of such data having larger (aleatoric) uncertainty or inherent biases compared to $\mathbf{D}$. The trade-off between these levels of fidelity gives rise to a data structure that consists of a large volume of low-cost, low-accuracy data and a limited volume of high-cost, high-accuracy data [80]. Prioritizing efficiency over accuracy can be achieved by experimental [29,81], computational [82,83] or analytical strategies [84,85]. The assumption throughout this manuscript will be such that there are an M number of such trade-offs, each of which is associated with a single-output objective function $f_m : [0,1]^D \to \mathbb{R}$, where $m \in \{1, \dots, M\}$. The utilized convention is that f_1 has the lowest fidelity (i.e., cost and accuracy), and f_M has the highest fidelity. The data structure of a multi-fidelity design of experiments will be generically written as a vertical concatenation $\boldsymbol{\mathcal{D}} := (\mathbf{D}_1, \dots, \mathbf{D}_M)^\top$ of M DoEs, ordered by increasing

---

[2] This is a simplifying assumption in concordance with f being standardized. However, it is not strictly necessary [39,76].

levels of fidelity. Similar notation for the multi-fidelity input data $\mathcal{X} := (\mathbf{X}_1, \dots, \mathbf{X}_M)^\top$ and multi-fidelity output data $\mathbf{y} := (\mathbf{y}_1, \dots, \mathbf{y}_M)^\top$ are introduced, where $\mathbf{D_m} := (\mathbf{X}_m, \mathbf{y}_m)$ for $m \in \{1, \dots, M\}$. It is furthermore assumed that $\mathbf{X}_m$ consists of N_m designs with D design variables, and therefore has dimensions $N_m \times D$, while $\mathbf{y}_m$ is correspondingly a column vector of length N_m consisting of the (realized) values $f_m(\mathbf{X}_m)$.

A single-output objective f can be treated like a (single-output) GP (SOGP). This can be generalized to handling a multi-output objective $\mathbf{f} := (f_1, \dots, f_M)^\top$. In the case of an SOGP model, the covariance kernel κ is a real-valued function. A multi-output Gaussian process (MOGP) model with M outputs generalizes κ to an $M \times M$ matrix-valued function: $\boldsymbol{\mathcal{K}}(\cdot,\cdot) := (\kappa_{i,j}(\cdot,\cdot))_{i,j=1}^{M}$. As with the SOGP case, the covariance $\boldsymbol{\mathcal{K}}$ is dependent on a vector of model parameters θ, which for the sake of notational clarity will be suppressed for the remainder of this subsection. Each component $\kappa_{i,j}$ is to be interpreted as the (single-output) covariance function between fidelities i and j. For example, if $\mathbf{u}$ is a design to be evaluated at fidelity i and $\mathbf{v}$ at fidelity j,

$$\kappa_{i,j}(\mathbf{u}, \mathbf{v}) := \mathrm{Cov}(f_i(\mathbf{u}), f_j(\mathbf{v})). \tag{B.11}$$

Following similar reasoning as for SOGP models, the parameterized posterior distribution is an M-dimensional multivariate normal distribution,

$$\mathbf{f}(\mathbf{x})|(\mathbf{Y} = \mathbf{y}, \boldsymbol{\Theta} = \theta) \sim \mathcal{N}(\boldsymbol{\mu}(\mathbf{x}), \boldsymbol{\Sigma}(\mathbf{x})), \tag{B.12}$$

which models each fidelity simultaneously. The parameterized posterior distribution comprises $\boldsymbol{\mu}(\mathbf{x}) = (\mu_1(\mathbf{x}), \dots, \mu_M(\mathbf{x}))^\top$ as the posterior M-dimensional mean and $\boldsymbol{\Sigma}(\mathbf{x}) = (\sigma^2_{i,j}(\mathbf{x}))_{i,j=1,\dots,M}$ as the $M \times M$ posterior covariance matrix, whenever $\mathbf{x} \in [0,1]^D$. The components of these statistical parameters are readily interpretable as follows: $\mu_m(\mathbf{x})$ is the mean and $\sigma^2_m(\mathbf{x}) := \sigma^2_{m,m}(\mathbf{x})$ is the variance at fidelity $m \in \{1, \dots, M\}$.

Despite the introduced complexity compared to SOGP, their expressions in terms of the prior covariance function $\boldsymbol{\mathcal{K}}$ are similar, cf. Eqs. (B.5) and (B.6):

$$\boldsymbol{\mu}(\mathbf{x}) = \boldsymbol{\mathcal{K}}(\mathbf{x}, \mathcal{X})^\top \mathbf{K}^{-1} \mathbf{y}, \tag{B.13}$$

$$\boldsymbol{\Sigma}(\mathbf{x}) = \boldsymbol{\mathcal{K}}(\mathbf{x}, \mathbf{x}) - \boldsymbol{\mathcal{K}}(\mathbf{x}, \mathcal{X})^\top \mathbf{K}^{-1} \boldsymbol{\mathcal{K}}(\mathbf{x}, \mathcal{X}). \tag{B.14}$$

In the MOGP scenario, the covariance matrix

$$\mathbf{K} := \begin{pmatrix} \mathbf{K}_{1,1}(\mathbf{X}_1, \mathbf{X}_1) & \dots & \mathbf{K}_{1,M}(\mathbf{X}_1, \mathbf{X}_M) \\ \vdots & \ddots & \vdots \\ \mathbf{K}_{M,1}(\mathbf{X}_M, \mathbf{X}_1) & \dots & \mathbf{K}_{M,M}(\mathbf{X}_M, \mathbf{X}_M) \end{pmatrix} \tag{B.15}$$

has dimensions $(N_1 + \dots + N_M) \times (N_1 + \dots + N_M)$. This matrix is to be interpreted as having an $M \times M$ block structure with blocks of size $N_i \times N_j$ for any fidelity levels $i, j \in \{1, \dots, M\}$. Each block is defined as:

$$\mathbf{K}_{i,j}(\mathbf{X}_i, \mathbf{X}_j) := \begin{pmatrix} \kappa_{i,j}(\mathbf{x}_{i,1}, \mathbf{x}_{j,1}) & \dots & \kappa_{i,j}(\mathbf{x}_{i,1}, \mathbf{x}_{j,N_j}) \\ \vdots & \ddots & \vdots \\ \kappa_{i,j}(\mathbf{x}_{i,N_i}, \mathbf{x}_{j,1}) & \dots & \kappa_{i,j}(\mathbf{x}_{i,N_i}, \mathbf{x}_{j,N_j}) \end{pmatrix}. \tag{B.16}$$

Furthermore,

$$\boldsymbol{\mathcal{K}}(\mathbf{x}, \mathcal{X}) := \begin{pmatrix} \kappa_{1,1}(\mathbf{x}, \mathbf{X}_1) & \dots & \kappa_{1,M}(\mathbf{x}, \mathbf{X}_1) \\ \vdots & \ddots & \vdots \\ \kappa_{M,1}(\mathbf{x}, \mathbf{X}_M) & \dots & \kappa_{M,M}(\mathbf{x}, \mathbf{X}_M) \end{pmatrix} \tag{B.17}$$

is an $(N_1 + \dots + N_M) \times M$ matrix consisting of stacked column vectors of length N_m, $m \in \{1, \dots, M\}$, defined by

$$\kappa_{i,j}(\mathbf{x}, \mathbf{X}_i) := \begin{pmatrix} \kappa_{i,j}(\mathbf{x}, \mathbf{x}_{i,1}) \\ \vdots \\ \kappa_{i,j}(\mathbf{x}, \mathbf{x}_{i,N_i}) \end{pmatrix}. \tag{B.18}$$

There exist several methods in the literature that attempt to relate the covariance functions $\kappa_{i,j}$ to one another based on an underlying structure. One way is the linear auto-regressive model [86,87], which requires the discrepancy between the fidelities to be modeled as a Gaussian process. However, this method is not commonly used in combination with Bayesian optimization. This is because the subsequent design matrices need to be nested to obtain the model parameters of $\kappa_{i,j}$ accurately, i.e., $\mathbf{X}_0 \supset \mathbf{X}_1 \supset \dots \supset \mathbf{X}_M$. This is a highly restrictive condition to perform optimization, which requires appending the design matrix with a new candidate design at any arbitrary fidelity level.

A structural assumption that is compatible with Bayesian optimization is the multiplicative separation between design and fidelity [88]. Since multi-task Bayesian optimization [50] is a well-established method, which has also been implemented in Python [89], it was selected as the reference multi-fidelity Bayesian optimization method for our investigation.

Given a single-output covariance function κ and a positive semi-definite matrix $\mathbf{B} := (b_{i,j})_{i,j=1\dots,M}$, it can be postulated that $\boldsymbol{\mathcal{K}} \doteq \boldsymbol{\mathcal{K}}_{\mathrm{MT}}(\kappa) = (\kappa_{i,j})_{i,j=1,\dots,M}$ where

$$\kappa_{i,j}(\mathbf{u}, \mathbf{v}) := b_{i,j} \cdot \kappa(\mathbf{u}, \mathbf{v}) \tag{B.19}$$

for any designs $\mathbf{u}, \mathbf{v} \in [0,1]^D$. In practice, the positive semi-definiteness of $\mathbf{B}$ is ensured by defining it as a Gram matrix $\mathbf{B} = \mathbf{A}^\top \mathbf{A} + \mathrm{diag}(\mathbf{a})$, where the model parameters in the low-rank matrix $\mathbf{A}$ and the non-negative vector $\mathbf{a}$ can be learned by optimizing a likelihood function (see below, Eq. (B.22)).

The resulting type of MOGP is called a multi-task GP (MTGP). As a consequence, Eq. (B.15) can be rewritten as a block-wise Hadamard product,

$$\mathbf{K} = \begin{pmatrix} b_{1,1}\mathbf{K}(\mathbf{X}_1, \mathbf{X}_1) & \dots & b_{1,M}\mathbf{K}(\mathbf{X}_1, \mathbf{X}_M) \\ \vdots & \ddots & \vdots \\ b_{M,1}\mathbf{K}(\mathbf{X}_M, \mathbf{X}_1) & \dots & b_{M,M}\mathbf{K}(\mathbf{X}_M, \mathbf{X}_M) \end{pmatrix} = \mathbf{B} \odot \mathbf{K}_{\mathrm{Had}}, \tag{B.20}$$

where $\mathbf{K}_{\mathrm{Had}} := (\mathbf{K}(\mathbf{X}_i, \mathbf{X}_j))_{i,j=1\dots,M}$ is a block matrix with blocks defined in the sense of Eq. (B.3).

Furthermore, if all input designs for each fidelity are homotopic, i.e., $\mathbf{X}_1 = \dots = \mathbf{X}_M =: \mathbf{X}_{\mathrm{hom}}$, this implies that

$$\mathbf{K} = \begin{pmatrix} b_{1,1}\mathbf{K}_{\mathrm{Kro}} & \dots & b_{1,M}\mathbf{K}_{\mathrm{Kro}} \\ \vdots & \ddots & \vdots \\ b_{M,1}\mathbf{K}_{\mathrm{Kro}} & \dots & b_{M,M}\mathbf{K}_{\mathrm{Kro}} \end{pmatrix} = \mathbf{B} \otimes \mathbf{K}_{\mathrm{Kro}} \tag{B.21}$$

as a Kronecker product instead, where $\mathbf{K}_{\mathrm{Kro}} := \mathbf{K}(\mathbf{X}_{\mathrm{hom}}, \mathbf{X}_{\mathrm{hom}})$, again defined in the sense of Eq. (B.3).

The benefits, with regards to regression, are twofold: $\mathbf{K}_{\mathrm{Kro}}$ is much smaller than the block matrix $\mathbf{K}_{\mathrm{Had}}$, which speeds up the calculation of the covariance matrix inverse $\mathbf{K}^{-1}$ for the calculation of the maximum likelihood estimate (Eq. (B.22)), and also parameterized posterior mean (Eq. (B.13)) and variance (Eq. (B.14)). Meanwhile, the calculation of the inverse itself is also simplified by the fact that $\mathbf{K}^{-1} = \mathbf{B}^{-1} \otimes \mathbf{K}^{-1}_{\mathrm{Kro}}$. These advantages have been recognized by various authors, and MTGPs have been successfully used as regression models in practice, such as biomedical applications [90,91], assessing machine translation quality [92], and school exam data analysis [93].

However, in the context of data with varying levels of cost, it is not sensible to assume homotopical designs across all fidelity levels. This drawback reduces the gain in computational efficiency to a side note, and an MTGP will be constructed using the block-wise Hadamard product within the context of this manuscript.

Similar to single-fidelity Gaussian process regression, the goal of MOGP regression (MOGPR) is to find the maximum likelihood model parameter vector $\hat{\theta}_{\mathrm{MLE}}$ for a supposed covariance kernel. The expression for the maximum likelihood estimate of the model parameter vector is similar to the single-fidelity GP regression scenario in Eq. (B.8):

$$\hat{\theta}_{\text{MLE}} = \text{argmin}_{\theta\in\mathcal{T}} \ln(\det(\mathbf{K}_\theta(\mathcal{X}))) + \mathbf{y}^\top \mathbf{K}_\theta^{-1}(\mathcal{X})\mathbf{y}. \tag{B.22}$$

After locating $\hat{\theta}_{\text{MLE}}$, the posterior predictive distribution of $\mathbf{f}(\mathbf{x})|(\mathbf{Y} = \mathbf{y})$ can be approximated by a regression-predictive, multivariate normal distribution with probability density function $\hat{\phi}$, in a similar fashion to the single-fidelity scenario, cf. Eqs. (B.9) and (B.10):

$$p(\mathbf{f}(\mathbf{x})|\mathbf{Y}=\mathbf{y})[y] \approx p(\mathbf{f}(\mathbf{x})|\mathbf{Y}=\mathbf{y},\boldsymbol{\Theta}=\hat{\theta}_{\text{MLE}})[y] = \phi_{\hat{\boldsymbol{\mu}}(\mathbf{x}),\hat{\boldsymbol{\Sigma}}(\mathbf{x})}(y) =: \hat{\phi}(y;\mathbf{x},\mathcal{D},\mathcal{K}), \tag{B.23}$$

where

$$\phi_{\boldsymbol{\mu},\boldsymbol{\Sigma}}(\mathbf{z}) := \frac{1}{(2\pi)^{M/2}\sqrt{\det(\boldsymbol{\Sigma})}} \exp\left(-\frac{1}{2}(\mathbf{z}-\boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{z}-\boldsymbol{\mu})\right)$$

denotes the multivariate normal probability density function with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. Furthermore, $\hat{\boldsymbol{\mu}} := \boldsymbol{\mu}_{\hat{\theta}_{\text{MLE}}}$ (Eq. (B.13)) and $\hat{\boldsymbol{\Sigma}} := \boldsymbol{\Sigma}_{\hat{\theta}_{\text{MLE}}}$ (Eq. (B.14)) are the mean vector and covariance matrix of the posterior distribution, parameterized by the maximum likelihood model parameter vector. Finally, it is useful to define

$$\hat{\phi}_m := \phi_{\hat{\mu}_m(\mathbf{x}),\hat{\sigma}_m^2(\mathbf{x})}. \tag{B.24}$$

Appendix C. Spinodal design parameter sensitivity

Section 3 introduced the global, variance-based sensitivity analysis of the EA objective with respect to the 4D design space and the subsequently reduced 3D design space. These analyses were able to yield a comprehensive insight into the design parameter effects. As a confirmation of the conclusions drawn from those insights, one could consider visualizing the effect of each design with respect to the objective. Such an overview is obtained by projecting each EA value down onto the individual design parameter axes. This results in the scatter plots in Fig. C.15.

The leftmost scatter plot in both Fig. C.15(a) and Fig. C.15(b) appears to show a strong relationship between ρ and the normalized EA value, whereas the projections onto the θ-dimensions exhibit much less structure. This holds for both the low- and high-fidelity objectives. This reaffirms the physical intuition of the strong dependence between the energy absorption of a spinodoid material and its relative density. Moreover, it legitimizes Sobol' sensitivity analysis as an effective tool to discover dependency patterns between design parameters and the objective, as the conclusion from Fig. 7 shows.

Similar to the 4D case, a coordinate-wise projection scatter plot was produced for the reduced 3D design space, as shown in Fig. C.16.

This confirms the difficulty for the SSA procedure in determining the sensitivity indices due to less structure in the overall data in the reduced design space, as was concluded from Fig. 8.

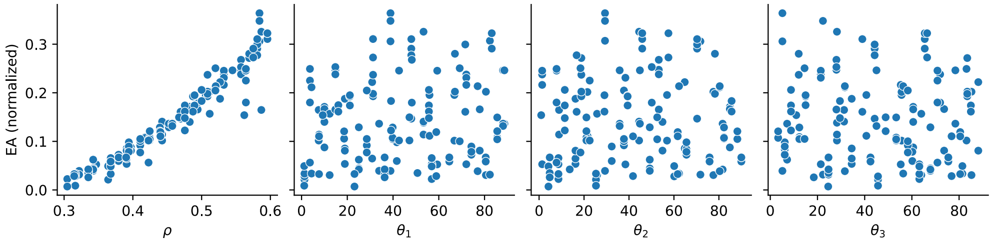

(a) High-fidelity EA values per design parameter

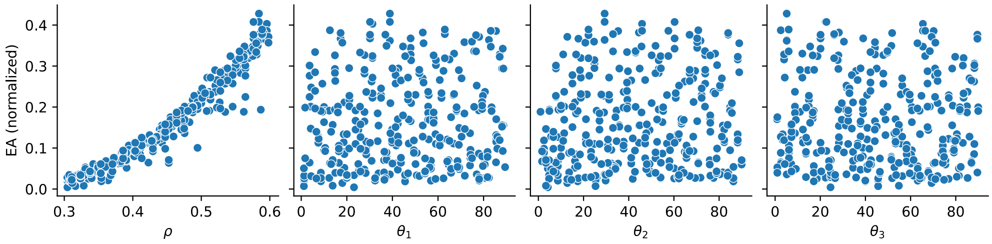

(b) Low-fidelity EA values per design parameter

Fig. C.15. The collection of Saltelli samples used in Fig. 6 and their normalized EA value projections onto the design parameter axes.

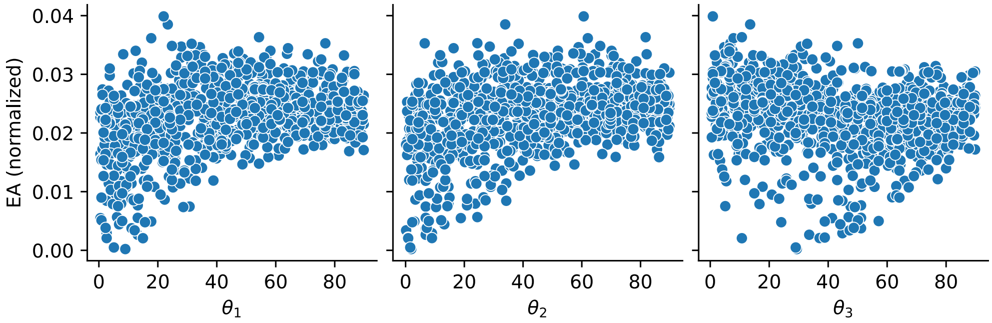

(a) High-fidelity EA values per design parameter

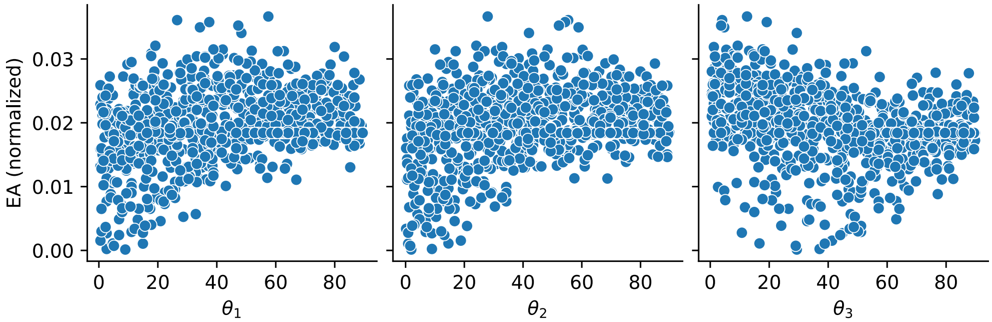

(b) Low-fidelity EA values per design parameter

Fig. C.16. The collection of 980 Saltelli samples (each fidelity) and their normalized EA value projections onto the remaining design parameter axes when $\rho = 0.3$ is fixed.

Data availability

All presented data and the implemented workflow in this manuscript are open-source and accessible via GitHub: https://github.com/llguo95/MFB.

References

[1] Wu Y, Fang J, Wu C, Li C, Sun G, Li Q. Additively manufactured materials and structures: a state-of-the-art review on their mechanical characteristics and energy absorption. Int J Mech Sci 2023;246:108102.

[2] Harris JA, McShane GJ. Impact response of metallic stacked origami cellular materials. Int J Impact Eng 2021;147:103730.

[3] Han SC, Lee JW, Kang K. A new type of low density material: shellular. Adv Mater (Deerfield Beach Fla) 2015;27(37):5506–11.

[4] Guell Izard A, Bauer J, Crook C, Turlo V, Valdevit L. Ultrahigh energy absorption multifunctional spinodal nanoarchitectures. Small 2019;15(45):1903834.

[5] Al-Ketan O, Abu Al-Rub RK. Multifunctional mechanical metamaterials based on triply periodic minimal surface lattices. Adv Eng Mater 2019;21(10):1900524.

[6] Cahn JW. On spinodal decomposition. Acta metallurgica 1961;9(9):795–801.

[7] Bendsoe MP, Sigmund O. Topology optimization: theory, methods, and applications. Springer Science & Business Media; 2013.

[8] Zhang Y, Hsieh M-T, Valdevit L. Mechanical performance of 3D printed interpenetrating phase composites with spinodal topologies. Composite Struct 2021;263:113693.

[9] Guo Y, Sharma S, Kumar S. Inverse designing surface curvatures by deep learning. Adv Intell Syst 2024:2300789.

[10] Yu X, Zhou J, Liang H, Jiang Z, Wu L. Mechanical metamaterials associated with stiffness, rigidity and compressibility: a brief review. Prog Mater Sci 2018;94:114–73.

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

[12] Shin D, Cupertino A, de Jong MHJ, Steeneken PG, Bessa MA, Norte RA. Spiderweb nanomechanical resonators via Bayesian optimization: inspired by nature and guided by machine learning. Adv Mater 2022;34(3):2106248.

[13] Kuszczak I, Azam FI, Bessa MA, Tan PJ, Bosi F. Bayesian optimisation of hexagonal honeycomb metamaterial. Extrem Mech Lett 2023;64:102078.

[14] Raßloff A, Seibert P, Kalina KA, Kästner M. Inverse design of spinodoid structures using Bayesian optimization. Comput Mech 2025:1–22.

[15] Marco A, Berkenkamp F, Hennig P, Schoellig AP, Krause A, Schaal S, Trimpe S. Virtual vs. real: trading off simulations and physical experiments in reinforcement learning with Bayesian optimization. In: 2017 IEEE international conference on Robotics and Automation (ICRA). IEEE; 2017. pp. 1557–63.

[16] Pang G, Perdikaris P, Cai W, Karniadakis GE. Discovering variable fractional orders of advection–dispersion equations from field data using multi-fidelity Bayesian optimization. J Comput Phys 2017;348:694–714.

[17] Perdikaris P, Karniadakis GE. Model inversion via multi-fidelity Bayesian optimization: a new paradigm for parameter estimation in haemodynamics, and beyond. J R Soc Interface 2016;13(118):20151107.

[18] Imani M, Ghoreishi SF, Allaire D, Braga-Neto UM. Mfbo-ssm: multi-fidelity Bayesian optimization for fast inference in state-space models. In: Proceedings of the AAAI conference on artificial intelligence, vol. 33. 2019. pp. 7858–65.

[19] Wu Q, Wang H, Hong W. Multistage collaborative machine learning and its application to antenna modeling and optimization. IEEE Trans Antennas Propag 2020;68(5):3397–409.

[20] Charayron R, Lefebvre T, Bartoli N, Morlier J. Towards a multi-fidelity & multi-objective Bayesian optimization efficient algorithm. Aerospace Sci Technol 2023;142:108673.

[21] Ariyarit A, Sugiura M, Tanabe Y, Kanazaki M. Hybrid surrogate-model-based multi-fidelity efficient global optimization applied to helicopter blade design. Eng Optim 2018;50(6):1016–40.

[22] Tyan M, Nguyen van N, Lee J-W. Improving variable-fidelity modelling by exploring global design space and radial basis function networks for aerofoil design. Eng Optim 2015;47(7):885–908.

[23] Kou J, Zhang W. Multi-fidelity modeling framework for nonlinear unsteady aerodynamics of airfoils. Appl Math Model 2019;76:832–55.

[24] Peng X, Kou J, Zhang W. Multi-fidelity nonlinear unsteady aerodynamic modeling and uncertainty estimation based on hierarchical kriging. Appl Math Model 2023;122:1–21.

[25] Tran A, Wildey T, McCann S. Smf-bo-2cogp: a sequential multi-fidelity constrained Bayesian optimization framework for design applications. J Comput Inf Sci Eng 2020;20(3).

[26] Aziz M, Hare W, Jaberipour M, Lucet Y. Multi-fidelity algorithms for the horizontal alignment problem in road design. Eng Optim 2020;52(11):1848–67.

[27] Zhang S, Lyu W, Yang F, Yan C, Zhou D, Zeng X, Hu X. An efficient multi-fidelity Bayesian optimization approach for analog circuit synthesis. In: 2019 56th ACM/IEEE design Automation conference (DAC). IEEE; 2019. pp. 1–6.

[28] Li S, Xing W, Kirby R, Zhe S. Multi-fidelity Bayesian optimization via deep neural networks. Adv Neural Inf Process Syst 2020;33.

[29] Cupertino A, Shin D, Guo L, Steeneken PG, Bessa MA, Norte RA. Centimeter-scale nanomechanical resonators with low dissipation. Nature Commun 2024;15(1):4255.

[30] Lenhart T, Eckhardt K, Fohrer N, Frede H-G. Comparison of two different approaches of sensitivity analysis. Phys Chem Earth Parts A/B/C 2002;27(9–10):645–54.

[31] Baucells M, Borgonovo E. Invariant probabilistic sensitivity analysis. Manage Sci 2013;59(11):2536–49.

[32] Borgonovo E, Plischke E. Sensitivity analysis: a review of recent advances. Eur J Oper Res 2016;248(3):869–87.

[33] Sobol IM. Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates. Math Comput Simul 2001;55(1–3):271–80.

[34] Opgenoord MMJ, Allaire DL, Willcox KE. Variance-based sensitivity analysis to support simulation-based design under uncertainty. J Mech Des 2016;138(11):111410.

[35] Marelli S, Lamas C, Konakli K, Mylonas C, Wiederkehr P, Sudret B. Uqlab user manual–sensitivity analysis. Report UQLab-V1 2019:2–106.

[36] Song J, Wei P, Valdebenito MA, Faes M, Beer M. Data-driven and active learning of variance-based sensitivity indices with Bayesian probabilistic integration. Mech Syst Signal Process 2022;163:108106.

[37] Saltelli A, Annoni P, Azzini I, Campolongo F, Ratto M, Tarantola S. Variance based sensitivity analysis of model output. design and estimator for the total sensitivity index. Comput Phys Commun 2010;181(2):259–70.

[38] Herman J, Usher W. Salib: an open-source Python library for sensitivity analysis. J Open Source Softw 2017;2(9):97.

[39] Williams CKI, Rasmussen CE. Gaussian processes for machine learning, vol. 2. MIT Press Cambridge, MA; 2006.

[40] Jones DR, Schonlau M, Welch WJ. Efficient global optimization of expensive blackbox functions. J Glob Optim 1998;13(4):455–92.

[41] Srinivas N, Krause A, Kakade SM, Seeger M. Gaussian process optimization in the bandit setting: no regret and experimental design. [arXiv preprint] arXiv:0912.3995. 2009.

[42] Ament S, Daulton S, Eriksson D, Balandat M, Bakshy E. Unexpected improvements to expected improvement for bayesian optimization. [arXiv preprint] arXiv:2310.20708. 2023.

[43] Garnett R. Bayesian optimization. Cambridge University Press; 2023.

[44] Shahriari B, Swersky K, Wang Z, Adams RP, De Freitas N. Taking the human out of the loop: a review of Bayesian optimization. Proc IEEE 2015;104(1):148–75.

[45] Kandasamy K, Dasarathy G, Schneider J, Póczos B. Multi-fidelity Bayesian optimisation with continuous approximations. In: International conference on machine learning. PMLR; 2017. pp. 1799–808.

[46] Zhang Y, Han Z-H, Zhang K-S. Variable-fidelity expected improvement method for efficient global optimization of expensive functions. Struct Multidiscip Optim 2018;58(4):1431–51.

[47] Jiang P, Cheng J, Zhou Q, Shu L, Hu J. Variable-fidelity lower confidence bounding approach for engineering optimization problems with expensive simulations. AIAA J 2019;57(12):5416–30.

[48] Do B, Zhang R. Multi-fidelity bayesian optimization in engineering design. [arXiv preprint] arXiv:2311.13050. 2023.

[49] Maddox WJ, Balandat M, Wilson AG, Bakshy E. Bayesian optimization with high-dimensional outputs. Adv Neural Inf Process Syst 2021;34:19274–87.

[50] Swersky K, Snoek J, Adams RP. Multi-task Bayesian optimization. Adv Neural Inf Process Syst 2013;26.

[51] Laumanns M, Ocenasek J. Bayesian optimization algorithms for multi-objective optimization. In: International conference on parallel problem solving from nature. Springer; 2002. pp. 298–307.

[52] Schonlau M. Computer experiments and global optimization. University of Waterloo; 1997. https://dspacemainprd01.lib.uwaterloo.ca/server/api/core/bitstreams/113a2333-5543-4b94-bfeb-52b448ab35ca/content.

[53] González J, Dai Z, Hennig P, Lawrence N. Batch Bayesian optimization via local penalization. In: Artificial intelligence and statistics. PMLR; 2016. pp. 648–57.

[54] Huang D, Allen TT, Notz WI, Miller RA. Sequential kriging optimization using multiple-fidelity evaluations. Struct Multidiscip Optim 2006;32(5):369–82.

[55] Liu Y, Chen S, Wang F, Xiong F. Sequential optimization using multi-level cokriging and extended expected improvement criterion. Struct Multidiscip Optim 2018;58(3):1155–73.

[56] Ruan X, Jiang P, Zhou Q, Hu J, Shu L. Variable-fidelity probability of improvement method for efficient global optimization of expensive black-box problems. Struct Multidiscip Optim 2020;62(6):3021–52.

[57] Cahn JW. Phase separation by spinodal decomposition in isotropic systems. J Chem Phys 1965;42(1):93–9.

[58] Zheng L, Kumar S, Kochmann DM. Data-driven topology optimization of spinodoid metamaterials with seamlessly tunable anisotropy. Comput Methods Appl Mech Eng 2021;383:113894.

[59] Kansara H, Khosroshahi SF, Guo L, Bessa MA, Tan W. Multi-objective Bayesian optimisation of spinodoid cellular structures for crush energy absorption. Comput Methods Appl Mech Eng 2025;440:117890.

[60] Voet Y, Sande E, Buffa A. A theoretical analysis of mass scaling techniques. Comput Mech 2025:1–24.

[61] Kumar S, Tan S, Zheng L, Kochmann DM. Inverse-designed spinodoid metamaterials. npj Comput Mater 2020;6(1):73.

[62] Olsson DM, Nelson LS. The nelder-mead simplex procedure for function minimization. Technometrics 1975;17(1):45–51.

[63] Hansen N, Müller SD, Koumoutsakos P. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (cma-es). Evol Comput 2003;11(1):1–18.

[64] Loshchilov I. Cma-es with restarts for solving CEC 2013 benchmark problems. In: 2013 IEEE congress on evolutionary computation. IEEE; 2013. pp. 369–76.

[65] Liu Y, Wang H, Yan L, Huang J, Liang Y. Mechanical properties of homogeneous and functionally graded spinodal structures. Int J Mech Sci 2024;269:109043.

[66] Marler RT, Arora JS. The weighted sum method for multi-objective optimization: new insights. Struct Multidiscip Optim 2010;41:853–62.

[67] Knowles J. Parego: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. IEEE Trans Evol Comput 2006;10(1):50–66.

[68] Ellison M, DiazDelaO FA, Ince NZ, Willetts M. Robust optimisation of computationally expensive models using adaptive multi-fidelity emulation. Appl Math Model 2021;100:92–106.

[69] Irshad F, Karsch S, Döpp A. Leveraging trust for joint multi-objective and multi-fidelity optimization. Mach Learn Sci Technol 2024;5(1):015056.

[70] Efron B. Bootstrap methods: another look at the jackknife. In: Breakthroughs in statistics: methodology and distribution. Springer; 1992. pp. 569–93.

[71] Hoerl AE, Kennard RW. Ridge regression: biased estimation for nonorthogonal problems. Technometrics 1970;12(1):55–67.

[72] Tibshirani R. Regression shrinkage and selection via the lasso. J R Stat Soc Ser B Stat Methodol 1996;58(1):267–88.

[73] Hutter F, Hoos HH, Leyton-Brown K. Sequential model-based optimization for general algorithm configuration. In: International conference on learning and intelligent optimization. Springer; 2011. pp. 507–23.

[74] Bergstra J, Bardenet R, Bengio Y, Kégl B. Algorithms for hyper-parameter optimization. Adv Neural Inf Process Syst 2011;24.

[75] Snoek J, Rippel O, Swersky K, Kiros R, Satish N, Sundaram N, Patwary M, Prabhat M, Adams R. Scalable Bayesian optimization using deep neural networks. In: International conference on machine learning; 2015. pp. 2171–80.

[76] Bishop CM, Nasrabadi NM. Pattern recognition and machine learning, vol. 4. Springer; 2006.

[77] Schulz E, Speekenbrink M, Krause A. A tutorial on Gaussian process regression: modelling, exploring, and exploiting functions. J Math Psychol 2018;85:1–16.

[78] Liu M, Chowdhary G, Da Silva BC, Liu S-Y, How JP. Gaussian processes for learning and control: a tutorial with examples. IEEE Control Syst Mag 2018;38(5):53–86.

[79] Holt W, Nguyen D. Essential aspects of Bayesian data imputation. In: An invitation to undergraduate research in risk management: actuarial science, mathematical finance, and sports analytics. Springer; 2025. pp. 207–236.

[80] Fernández-Godino MG, Park C, Kim N-H, Haftka RT. Review of multi-fidelity models. [arXiv preprint] arXiv:1609.07196. 2016.

[81] Forrester AIJ. Black-box calibration for complex-system simulation. Philos Trans R Soc A Math Phys Eng Sci 2010;368(1924):3567–79.

[82] Balabanov V, Grossman B, Watson L, Mason W, Haftka R. Multifidelity response surface model for hsct wing bending material weight. In: 7th AIAA/USAF/NASA/ISSMO symposium on multidisciplinary analysis and optimization; 1998. p. 4804.

[83] Fidkowski K. Quantifying uncertainties in radiation hydrodynamics models. Ann Arbor 2014;1001:48109.

[84] Forrester AIJ, Sóbester A, Keane AJ. Multi-fidelity optimization via surrogate modelling. Proc R Soc A Math Phys Eng Sci 2007;463(2088):3251–69.

[85] Le Gratiet L, Cannamela C, Iooss B. A Bayesian approach for global sensitivity analysis of (multifidelity) computer codes. SIAM ASA J Uncertain Quantif 2014;2(1):336–63.

[86] Kennedy MC, O'Hagan A. Predicting the output from a complex computer code when fast approximations are available. Biometrika 2000;87(1):1–13.

[87] Liu H, Cai J, Ong Y-S. Remarks on multi-output Gaussian process regression. Knowl Based Syst 2018;144:102–21.

[88] Williams C, Bonilla EV, Chai KM. Multi-task Gaussian process prediction. Adv Neural Inf Process Syst 2007:153–60.

[89] Gardner J, Pleiss G, Weinberger KQ, Bindel D, Wilson AG. Gpytorch: blackbox matrix-matrix Gaussian process inference with GPU acceleration. Adv Neural Inf Process Syst 2018;31.

[90] Alaa AM, van der Schaar M. Deep multi-task Gaussian processes for survival analysis with competing risks. In: Proceedings of the 31st international conference on neural information processing systems; 2017. pp. 2326–34.

[91] Dürichen R, Pimentel MAF, Clifton L, Schweikard A, Clifton DA. Multi-task Gaussian process models for biomedical applications. In: IEEE-EMBS international conference on biomedical and health informatics (BHI). IEEE; 2014. pp. 492–5.

[92] Cohn T, Specia L. Modelling annotator bias with multi-task Gaussian processes: an application to machine translation quality estimation. In: Proceedings of the 51st annual meeting of the association for computational linguistics (volume 1: long papers); 2013. pp. 32–42.

[93] Li P, Chen S. Hierarchical Gaussian processes model for multi-task learning. Pattern Recognit 2018;74:134–44.