


Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

## Comput. Methods Appl. Mech. Engrg.

journal homepage: [www.elsevier.com/locate/cma](http://www.elsevier.com/locate/cma)

# Finite elements for Matérn-type random fields: Uncertainty in computational mechanics and design optimization

Tobias Duswald<sup>b,c,\*</sup>, Brendan Keith<sup>d</sup>, Boyan Lazarov<sup>a</sup>, Socratis Petrides<sup>a</sup>, Barbara Wohlmuth<sup>c</sup>

<sup>a</sup> Lawrence Livermore National Laboratory, United States of America

<sup>b</sup> CERN, Geneva, Switzerland

<sup>c</sup> School for Computation, Information, and Technology, Technical University of Munich, Germany

<sup>d</sup> Division of Applied Mathematics, Brown University, United States of America

## ARTICLE INFO

### Keywords:

Random field generation  
Topology optimization  
Stochastic optimization  
Finite elements  
Manifolds

## ABSTRACT

This work highlights an approach for incorporating realistic uncertainties into scientific computing workflows based on finite elements, focusing on prevalent applications in computational mechanics and design optimization. We leverage Matérn-type Gaussian random fields (GRFs) generated using the SPDE method to model aleatoric uncertainties, including environmental influences, varying material properties, and geometric ambiguities. Our focus lies on delivering practical GRF realizations that accurately capture imperfections and variations and understanding how they impact the predictions of computational models as well as the shape and topology of optimized designs. We describe a numerical algorithm based on solving a generalized SPDE to sample GRFs on arbitrary meshed domains. The algorithm leverages established techniques and integrates seamlessly with the open-source finite element library MFEM and associated scientific computing workflows, like those found in industrial and national laboratory settings. Our solver scales efficiently for large-scale problems and supports various domain types, including surfaces and embedded manifolds. We showcase its versatility through biomechanics and topology optimization applications, emphasizing the potential to influence these domains. The flexibility and efficiency of SPDE-based GRF generation empowers us to run large-scale optimization problems on 2D and 3D domains, including finding optimized designs on embedded surfaces, and to generate design features and topologies beyond the reach of conventional techniques. Moreover, these capabilities allow us to model and quantify geometric uncertainties on reconstructed submanifolds, such as the interpolated surfaces of cerebral aneurysms provided by postprocessing CT scans. In addition to offering benefits in these specific domains, the proposed techniques transcend specific applications and generalize to arbitrary forward and backward problems in uncertainty quantification involving finite elements.

## 1. Introduction

This article demonstrates a scalable and expressive approach for generating *Gaussian random fields* with applications to modeling uncertainties in computational mechanics and design optimization. The goal is to illustrate the importance and the necessity of

\* Corresponding author at: CERN, Geneva, Switzerland.

E-mail addresses: [tobias.duswald@tum.de](mailto:tobias.duswald@tum.de) (T. Duswald), [brendan\\_keith@brown.edu](mailto:brendan_keith@brown.edu) (B. Keith), [lazarov2@llnl.gov](mailto:lazarov2@llnl.gov) (B. Lazarov), [petrides1@llnl.gov](mailto:petrides1@llnl.gov) (S. Petrides), [wohlmuth@cit.tum.de](mailto:wohlmuth@cit.tum.de) (B. Wohlmuth).

<https://doi.org/10.1016/j.cma.2024.117146>

Received 5 March 2024; Received in revised form 14 May 2024; Accepted 7 June 2024

Available online 24 June 2024

0045-7825/© 2024 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

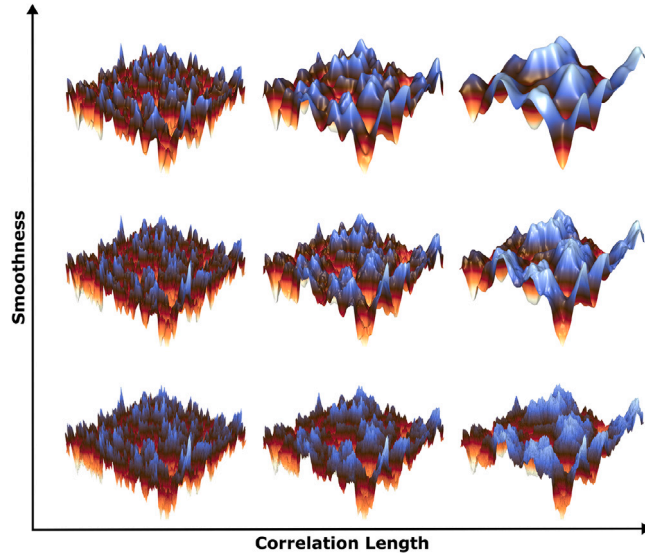


Fig. 1. Isotropic, homogeneous, Matérn-type Gaussian random fields of different smoothness and correlation length generated with the SPDE method. The smoothness and correlation-length increase along their respective axis. The examples are generated with homogeneous Neumann boundary conditions on  $D = (0, 1)^2$  using identical white noise realizations for all samples.

incorporating random fields with widely varying properties, such as correlation length, smoothness, and anisotropy, on arbitrary computational domains and embedded surfaces. The ability to easily control these properties and generate corresponding samples significantly reduces epistemic uncertainties in computational models. Therefore, we focus on examples with aleatoric uncertainties, describing the systems’ intrinsic randomness, assigned to the three categories described below. *Environmental and external excitations* describe external influences such as loads or other constraints. Random wind forces [1–4] acting on buildings [5] and wind turbines [6,7] serve as compelling examples in this category. *Material uncertainties* are associated with spatially varying material properties such as Young moduli, diffusion coefficients, or damping coefficients [8,9]. Such variations may be caused by material heterogeneity, changes occurring during manufacturing, or environmental degradation. Moreover, *geometric uncertainties* describe disturbances in the system’s geometry, for instance, the shape or the positioning of specific components. Examples arise in additive manufacturing processes where smooth surfaces may show artifacts and imperfections from manufacturing (cf. [10]). Computational models that fail to account for aleatoric uncertainties are prone to inaccurate predictions, poor risk assessment, and suboptimal designs [11].

The examples considered here leverage spatially correlated noise models. Mathematically, such noise is described by *random fields*, a generalization of a stochastic process modeling spatial variability. The fields are defined on a domain  $D \subset \mathbb{R}^d$ ,  $d = 1, 2, 3$ , and possess local notions of length scale, smoothness, and orientation (anisotropy). Practical system analyses and optimizations embed these fields into the physical model, e.g., as a random coefficient in a partial differential equation (PDE), and different field realizations appear in the inner loop of stochastic optimization algorithms and uncertainty quantification (UQ) workflows.

Among various types, *Gaussian random fields* (GRFs) are particularly notable as their mean value and covariance functions characterize them completely. The mean is a single-variable function  $m : D \rightarrow \mathbb{R}$ , and the covariance can be expressed as a two-variable function  $C : D \times D \rightarrow \mathbb{R}$ . As espoused below, we focus on *Matérn-type* random fields that can be defined on an arbitrary Lipschitz domain  $D \subset \mathbb{R}^d$ .

Classically, Matérn random fields are a special class of homogeneous GRFs with covariance defined over the domain  $D = \mathbb{R}^d$  as

$$C(\mathbf{x}, \mathbf{y}) = \frac{\sigma^2}{2^{\nu-1} \Gamma(\nu)} (\kappa r)^\nu \mathcal{K}_\nu(\kappa r), \quad r = \|\mathbf{x} - \mathbf{y}\|_2, \quad \kappa = \frac{\sqrt{2\nu}}{l}, \quad \mathbf{x}, \mathbf{y} \in \mathbb{R}^d, \tag{1}$$

where  $\sigma^2, \nu, l > 0$  are the variance, smoothness parameter, and correlation length, respectively. Moreover, in this definition, we have the gamma function  $\Gamma$  and the modified Bessel function of the second kind  $\mathcal{K}_\nu$ . The Matérn covariance (1), known for its flexibility and effectiveness, describes spatial uncertainties arising across diverse fields, including astronomy, health, engineering, imaging, environmental studies, geostatistics, and econometrics [12, Section 1.2]. Its versatility across numerous applications stems from its ability to interpolate between the exponential and squared-exponential covariance function. Moreover, it facilitates precise control over the field’s smoothness and its characteristic length scales via the parameters  $\nu$  and  $l$ , respectively. Matérn-type random fields generalize these attractive features to a larger space of covariance kernels on domains beyond  $D = \mathbb{R}^d$ . Fig. 1 illustrates random fields of Matérn-type covariance with different length scales and smoothness on the domain  $D = (0, 1)^2$ .

We consider problems requiring samples to be represented on the same discrete mesh utilized for modeling and analyzing the physics. The most popular techniques for generating such samples of GRFs employ *Karhunen–Loève* (KL) type series expansions or

filter-based Fourier methods. Both techniques represent the GRF as a linear combination of spatially varying basis functions with uncorrelated Gaussian random variables as coefficients. The size and type of basis determines the computational complexity of the random field generator. For correlation lengths that are relatively small compared to the characteristic size of the computational domain, the number of terms in KL-based expansions becomes large, and the associated computational cost is prohibitively expensive. On the other hand, Fourier-based methods are very computationally efficient for homogeneous GRFs but non-trivial to apply to the intricate design domains that are ubiquitous in computational mechanics.

The *Stochastic Partial Differential Equation* (SPDE) methodology proposed in [13] alleviates these limitations. The idea is to model a GRF as the solution to a fractional SPDE that smooths spatial Gaussian white noise  $\mathcal{W}$ , i.e., the inverse of the differential operator plays the role of a filter for convolving  $\mathcal{W}$ . In particular, Whittle [14,15] showed that the full-space ( $D = \mathbb{R}^d$ ) solution to

$$(\mathcal{I} - \kappa^{-2}\Delta)^k u = \eta\mathcal{W}, \quad k = \nu/2 + d/4, \quad (2)$$

is a GRF with covariance (1). Here,  $\mathcal{I}$  denotes the identity and  $\eta$  is a normalization constant depending on  $\sigma$ ,  $\nu$ , and  $l$ ; see, e.g., (11) below. Lindgren et al. [13] understood that one might exploit this link between (1) and (2) to efficiently generate random field realizations for downstream applications and coined the term *SPDE method*. The method does not impose any special limitations on the domain geometry, provides excellent flexibility for controlling the RF regularity, and allows for specifying particular behavior at the boundaries.

The method recently began to gain popularity in computational mechanics. For instance, Koh et al. [16] used the SPDE method in the context of Bayesian modeling and Gaussian process regression. Additionally and of significant relevance to this work, Guillemot et al. [17] and Ben-Yelun et al. [18] applied the methodology to optimize the topology of cantilevers and lattice structures, respectively. Here we demonstrate its efficacy in simulating complex uncertainties prevalent in scientific computing workflows for forward and inverse problems on modern distributed computing platforms. We present a numerical algorithm solving a generalized form of (2) to sample GRFs of Matérn covariance for computational workflows involving arbitrary meshed domains. The algorithm combines established techniques introduced in the context of the SPDE method; most notably, it promotes the general setup of [13], employs a rational approximation [19] to handle the fractional operator [20,21], and uses the efficient white noise sampling suggested in [22]. We provide a performant algorithm implementation, available as part of the open-source finite element library MFEM [23,24] (as of version 4.6). Standard domain decomposition techniques ensure perfect parallel scalability for large-scale optimization problems. The solver can be applied to 1D, 2D, or 3D domains, a surface part of the domain's boundary, or an arbitrary manifold embedded in the 3D space. The GRF discretization naturally utilizes the mesh of the physical problem and does not require additional discretization libraries and tools.

Highlighting the versatility of this methodology, we present its applications in two key areas we anticipate to benefit substantially: biomechanics and topology optimization. We show applications in modeling micro-structures with post-processed GRFs, illustrate the effects of uncertainties on complex geometries by considering GRFs on a cerebral aneurysm, and show comprehensive numerical experiments for topology optimization under uncertainty (thermal compliance and bridge design). In the latter set of examples, we observe new design features and topologies that are out of reach with established random field discretization methods. The paper proceeds with Section 2, which is devoted to methods generating GRFs, and presents our argument for adopting the SPDE method. Next, Section 3 contains an overview of the SPDE method and its numerical implementation. Sections 4 and 5 are devoted to biomechanics and topology optimization applications, respectively, demonstrating the potential of the method in these fields. The paper ends in Section 6 with a summary of our findings.

## 2. Numerical methods for generating Gaussian random fields

Gaussian random fields can be sampled through various techniques, each characterized by distinct theoretical underpinnings and practical applications. For a zero-mean GRF, a direct approach begins with specifying the covariance matrix  $\mathbf{C}$  between points at different locations and factorizing it with a Cholesky decomposition. Such matrix decomposition methods have been explored by numerous authors [25–28]. The so-called spectral method, utilizes the association between the covariance matrix and the power spectrum to perform this factorization in frequency space [29–35]. The centerpiece of this strategy is the Fourier transform, which facilitates the transformation to frequency space and the conversion of white noise into spatially-correlated random fields. Another approach is the *Karhunen–Loève* (KL) expansion [36–38], which leverages an optimal decomposition into eigenfunctions of the covariance kernel. Finally, PDE-based sampling techniques [12,13] have been developed, leveraging the connection between stochastic, fractional PDEs and certain GRFs [14,15]. Additional methods, such as the circulant embedding [39–41], also deserve recognition, as extensively reviewed in [42].

This section highlights some of the most prominent methodologies for generating GRFs. First, the *Expansion Optimal Linear Estimation* (EOLE) method [43], a variant of the KL expansion, is reviewed in Section 2.1. Second, the spectral method, employing the *fast Fourier transform* (FFT), is summarized in Section 2.2. Section 2.3 is dedicated to the SPDE method [13–15], the method of choice in this work. The selection of a particular method is contingent upon the specific objectives and constraints of the problem at hand, the desired characteristics of the random field (such as its correlation structure), and the computational resources available. While some methods are more suited to specific applications, others offer more flexibility regarding the attributes of the generated field. This section concludes by justifying and promoting the choice of the SPDE method for finite element-based applications.

### 2.1. The Karhunen–Loève expansion and the EOLE method

The EOLE method [43] is a KL-type series expansion method. In general, the KL expansion provides Gaussian random field samples by forming a linear combination of normalized eigenfunctions  $\phi_i(\mathbf{x})$  of the covariance kernel  $C(\mathbf{x}, \mathbf{y})$  with i.i.d. coefficients drawn from a normal distribution. In particular, for a GRF with mean  $m(\mathbf{x})$ , we have the expansion

$$u(\mathbf{x}) = m(\mathbf{x}) + \sum_{i=1}^{\infty} \alpha_i \sqrt{\lambda_i} \phi_i(\mathbf{x}), \tag{3}$$

where  $\int_D C(\mathbf{x}, \mathbf{y}) \phi_i(\mathbf{y}) d\mathbf{y} = \lambda_i \phi_i(\mathbf{x})$ ,  $\|\phi_i\|_{L^2(D)} = 1$ , and  $\alpha_i \sim \mathcal{N}(0, 1)$  for each  $i = 1, 2, \dots$  [44].

The EOLE method is a popular version of this expansion where the covariance kernel is sampled in a covariance matrix at a representative set of points  $\{\mathbf{x}^{(i)} \in D \mid i = 1, \dots, M\}$  in the domain  $D$ , and the eigenfunctions are approximated by their values at these points. For the selected set of points, the elements of the covariance matrix are computed as  $C_{ij} = C(\mathbf{x}^{(i)}, \mathbf{x}^{(j)})$ . The eigenvalues and the eigenvectors of the resulting matrix  $C \in \mathbb{R}^{M \times M}$  are denoted as  $\lambda_i$  and  $\mathbf{e}^{(i)}$ , respectively. The eigenvalues are ordered in non-increasing order, i.e.,  $\lambda_i \geq \lambda_j$  for all  $j > i$ . Following [43], the EOLE approximation of a GRF with mean  $m(\mathbf{x})$  and covariance  $C(\mathbf{x}, \mathbf{y})$  is given as

$$u(\mathbf{x}) \approx m(\mathbf{x}) + \sum_{i=1}^K \alpha_i \sum_{j=1}^M C(\mathbf{x}, \mathbf{x}^{(j)}) e_j^{(i)} / \sqrt{\lambda_i}, \quad K \leq M, \tag{4}$$

where  $\alpha_i \sim \mathcal{N}(0, 1)$  are again independent Gaussian random variables, and  $e_j^{(i)}$  denotes the  $j$ th component of the  $i$ th eigenvector.

In applications, the sum over the index  $i$  in (4) is truncated after  $K \ll M$  terms, i.e., the field is approximated using the  $K$  largest eigenvalues. For smooth random fields with correlation lengths comparable to the characteristic size of the domain  $D$ , the ordered sequence of eigenvalues decays fast [45], and the above expression provides a good approximation of the random field with only a relatively small number of terms  $K$ . Conveniently, this truncation reduces the computational complexity of evaluating  $u(\mathbf{x})$ . However, for rough oscillatory GRFs or computational domains much larger than the correlation lengths of the fields, an accurate representation requires a considerable number of terms, resulting in significantly increased computational costs.

### 2.2. Spectral methods

A *homogeneous* random field in  $\mathbb{R}^d$  ( $d > 1$ ) is the generalization of a stationary stochastic process in 1D. The statistical properties of such a field are invariant under parallel translations and, therefore, independent of the actual positions in space. In this case, the covariance simplifies to

$$C(\mathbf{x}, \mathbf{y}) = \hat{C}(\mathbf{x} - \mathbf{y}), \quad \mathbf{x}, \mathbf{y} \in \mathbb{R}^d \tag{5}$$

showing that the covariance between two points  $\mathbf{x}, \mathbf{y}$  may only depend on the length and direction of their connection. Further demanding invariance of the field’s statistical properties under all isometric transformations yields the important special case of an *isotropic*, homogeneous random field, i.e., the covariance function exclusively depends on the distance  $C(\mathbf{x}, \mathbf{y}) = \tilde{C}(|\mathbf{x} - \mathbf{y}|)$ .

Homogeneous random fields can be constructed by convolving spatial Gaussian white noise  $\mathcal{W}$  with a function  $T \in L^1(\mathbb{R}^d)$  [29–31,35,46,47]; namely,

$$u(\mathbf{x}) = (T * \mathcal{W})(\mathbf{x}). \tag{6}$$

Here, the function  $T$  is often referred to as a *filter* and the convolution is often interpreted as  $T$  smoothing (filtering) the white noise input. The Fourier transform  $\mathcal{F}$  is a computationally efficient way to evaluate the convolution in (6) because it becomes a pointwise multiplication in frequency space. Thus, (6) is computed as

$$u = \mathcal{F}^{-1}(\mathcal{F}(T * \mathcal{W})) = \mathcal{F}^{-1}(\mathcal{F}(T) \mathcal{F}(\mathcal{W})). \tag{7}$$

The filter  $T$  determines the covariance structure of the random field  $u$ . According to the Wiener–Khinchin theorem [48,49], the power spectrum  $\mathcal{P}(\mathbf{k})$ , which characterizes the distribution of variance and power across frequencies, is the Fourier transform of the covariance function  $\hat{C}(\mathbf{r})$ , where  $\mathbf{r} = \mathbf{x} - \mathbf{y}$ . Setting  $\mathcal{F}(T) = \sqrt{\mathcal{P}}$  results in a random field with a desired mean  $m$  and homogeneous covariance  $\hat{C}(\mathbf{r})$ ,

$$u = m + \mathcal{F}^{-1}(\sqrt{\mathcal{P}} \cdot \mathcal{F}(\mathcal{W})). \tag{8}$$

In practice, the FFT is utilized to evaluate (8), which requires accommodating artificial boundary conditions (e.g., periodic boundary conditions) on a truncated computational domain [50].

### 2.3. The SPDE method

The above approaches, i.e., the Karhunen–Loève expansion and the spectral method, require an explicit definition of the covariance function  $C$ . On the other hand, the SPDE method [13] defines  $C$  implicitly via the SPDE in (2). For an infinite computational domain, the Green’s function of the differential operator determines the covariance of the random field, thus replacing the filter function in the spectral method. As the method is based on the Matérn-type differential operator (2), we say that the

proposed method is limited to random fields of *Matérn-type covariance*. Generating GRFs with other types of covariance functions via a PDE-based methodology involves the derivation of alternative PDEs, which we do not consider here but refer the reader to [51, Section 1.3] for further details. However, notice that (2) can be solved on arbitrary domains and manifolds by replacing the Laplace operator with the Laplace–Beltrami operator. This helps generalize random fields of *Matérn-type covariance* to arbitrary domains. More details follow in Section 3.4. Additionally, the method allows for easy prescription of different boundary conditions and handling complicated domains, especially when the discretization process is performed using the Finite Element Method.

While the statistics literature typically uses the isotropic parameterization of the SPDE in (2), this work and the associated C++ implementation parametrize the SPDE following [52]. Recalling the smoothness parameter  $\nu > 0$  used to specify the random field’s regularity and introducing the matrix  $\Theta$  defining its correlation lengths in different directions, the SPDE reads

$$\left( I - \frac{1}{2\nu} \nabla \cdot (\Theta \nabla) \right)^{\frac{2\nu+d}{4}} u = \eta \mathcal{W}, \quad \Theta = \mathbf{R}^T \begin{pmatrix} l_1^2 & & \\ & \ddots & \\ & & l_d^2 \end{pmatrix} \mathbf{R}, \tag{9}$$

with  $\eta$  being a normalization constant. Here, the rotation matrix and the identity operator are denoted by  $\mathbf{R} \in \text{SO}(d)$  and  $I$ , respectively. Examples of such Matérn-type random fields with  $\Theta = l^2 I$  and homogeneous Neumann boundary conditions on the domain  $D = (0, 1)^2$  are depicted in Fig. 1, giving intuition for interpreting the regularity parameter  $\nu$  and correlation length  $l$ .

As demonstrated below, this parameterization also allows for a natural definition of anisotropy in engineering applications. Taking  $D = \mathbb{R}^d$ , the generalized Matérn covariance resulting from (9) is given as

$$C(\mathbf{x}, \mathbf{y}) = \sigma^2 M_\nu \left( \sqrt{2\nu} d(\mathbf{x}, \mathbf{y}) \right), \quad M_\nu(z) = \frac{2^{1-\nu}}{\Gamma(\nu)} z^\nu K_\nu(z), \tag{10}$$

with the anisotropic distance defined as  $d(\mathbf{x}, \mathbf{y}) = \sqrt{(\mathbf{x} - \mathbf{y})\Theta^{-1}(\mathbf{x} - \mathbf{y})}$ . Moreover, the normalization constant  $\eta$  leads to unit variance ( $C(\mathbf{x}, \mathbf{x}) = 1$ ) and is defined

$$\eta = \left( \frac{(2\pi)^{\frac{d}{2}} \sqrt{\det(\Theta)} \Gamma(\nu + \frac{d}{2})}{\nu^{\frac{d}{2}} \Gamma(\nu)} \right)^{\frac{1}{2}}, \tag{11}$$

with  $\Gamma$  denoting the gamma function. The corresponding GRF is homogeneous since (10) only depends on the anisotropic distance. However, taking  $D$  to be any Lipschitz domain in  $\mathbb{R}^d$ , (9) easily generalizes to a more general class of inhomogeneous Matérn-type GRFs upon replacing the matrix  $\Theta$  with the diffusion tensor  $\Theta(\mathbf{x})$ .

#### 2.4. Selection criteria and arguments for the SPDE method

Each of the random field generation methods summarized above has strengths and weaknesses that must be evaluated for the specific problem scenario. In this work, we are concerned with finite element workflows for computational mechanics and PDE-constrained optimization under uncertainty. The target problems are often defined on complex domains and can impose a wide range of requirements on the random field model. Typical workflows involve sampling multiple forward simulations of the physical model, with every sample corresponding to a different realization of the GRF and corresponding quantity of interest (QoI). The sampled QoIs may be used to estimate a posterior probability distribution (cf. Section 4) or an associated risk measure (cf. Section 5). For an optimization workflow (cf. Section 5), the gradients of the objective function and constraint functions can be estimated by solving an adjoint problem in addition to every forward problem.

We now match requirements of the random field generation method to the general workflows described above. Our underlying assumption is that the GRF model (9), with appropriate boundary conditions, is sufficiently general to capture the underlying uncertainty. First, the method must be able to generate random fields of Matérn-type covariance and handle all choices for its parameters  $\nu$  and  $\Theta = \Theta(\mathbf{x})$ . We emphasize the possible spatial dependence of  $\Theta$  because of the practical relevance of inhomogeneous GRFs [3]. Second, the method must be able to handle complicated domains  $D \subset \mathbb{R}^3$  and allow for defining general boundary conditions. Third, since the underlying physical processes will be simulated with finite elements, the random field discretizations should conveniently map to such data structures. Fourth, we demand (relative) scalability of the method. In particular, denoting the times to solve the physical problem and sample the random field as  $T_p$  and  $T_r$ , respectively, we demand that  $T_r \leq c T_p$  with  $c = \mathcal{O}(1)$ . Here, we pause to note that demanding a great difference between  $T_r$  and  $T_p$  (i.e.,  $c \ll 1$ ) is usually not necessary because there is typically no more than one random field generation step per physics solve. Thus, the dominant cost will always be determined by  $T_p$ . Lastly, we require the method to be memory efficient relative to the physics. In this case, letting  $M_r$  and  $M_p$  denote the memory footprint of the random field and the physical problem, respectively, we demand  $M_r \leq \tilde{c} M_p$  with  $\tilde{c} = \mathcal{O}(1)$ . In practice, limited computing resources may sometimes impose tighter upper bounds for  $c$  and  $\tilde{c}$  than  $\mathcal{O}(1)$ .

Karhunen–Loève-type expansions such as EOLE are algorithmically simple and effective for generating random fields and have been used successfully in engineering workflows, e.g., in topology optimization [53–55]. However, as mentioned in Section 2.1, methods in this class scale unfavorably as the correlation length diminishes. In (4), the number of points  $M$  sampling the space must be increased, yielding a large covariance matrix  $\mathbf{C} \in \mathbb{R}^{M \times M}$  whose decomposition becomes expensive. Additionally, its eigenvalues decay slowly, making evaluating the expansion increasingly challenging because the memory footprint grows linearly with the number of eigenvectors  $K$  and the overall computational cost scales roughly like  $\mathcal{O}(KM)$ . Decomposing the computational domain, evaluating the decomposition on each subdomain, and requiring continuity of the samples between the partitions reduces the size

of the covariance matrix and helps the method to scale [56]. However, such an approach can introduce errors on the interface between the subdomains and requires the utilization of regular geometries or advanced data structures not supported by general FEM libraries. Therefore, we rule out EOLE and similar approaches.

Spectral methods are excellent alternatives. FFT's  $\mathcal{O}(n \log n)$ -scaling enables fast generation of random fields with limited resources. Modern FFT libraries such as *heFFTe* [57] have proven that the algorithm scales on top-tier computing systems and complicated domains can be addressed with windowing techniques [50]. However, FFT-based methods become difficult, or infeasible, when approaching problems on two-dimensional manifolds embedded in  $\mathbb{R}^3$ . Moreover, the FFT is restricted to regular grids, and enforcing boundary conditions on anything but trivial domains requires expertise. Projecting the regular grid solution onto the finite element discretization of the physical problem is non-trivial, particularly in the context of distributed computing for which the data distribution of a graph-partitioned finite element problem does not coincide with the one used by the FFT libraries and extensive MPI-communication may be necessary for the mapping process. Finally, generating inhomogeneous fields requires space-dependent filter functions  $T$ . While such filters may be defined, the convolution theorem no longer holds, making the filtering of the white noise considerably more complex and resource intensive.

Here, we advocate for using the SPDE method. While we mainly restrict ourselves to examples from biomechanics and topology optimization under uncertainty, our arguments for employing the SPDE method extend to general application areas whose problems are encoded in PDEs and may be solved with the finite element method, provided that the Matérn-type covariance class accurately describes the uncertainties. FFT-based methods are arguably faster; however, even significantly shorter sampling times barely influence a PDE-constrained optimization workflow's overall runtime, and the SPDE method's convenience outweighs the FFT's computational advantages. The SPDE method allows arbitrary choices of regularity, orientation, and correlation lengths and samples random fields on arbitrary meshed domains. As the SPDE solution method here is based on finite elements, the random field can be computed on the same (distributed) mesh as the physical problem, and matching the random fields with the physical problem becomes a trivial task, even in setups involving thousands of CPU cores. We further note that the numerical solution of (9) is efficient and scalable (see Section 3) and, thus, our relative scalability demands,  $T_r \leq cT_p$  and  $M_r \leq \tilde{c}M_p$ , are satisfied. The SPDE method enables us to easily define boundary conditions for the GRF on complicated domains, and the method extends to a general class of inhomogeneous fields with  $\Theta = \Theta(\mathbf{x})$ . Hence, the SPDE method satisfies all requirements and carries the potential to become the method of choice in finite-element-driven application areas.

### 3. Implementing the SPDE method

The purpose of this section is to present an implementation-based overview of the SPDE approach to generating Matérn-type GRFs. For ease of presentation, we first choose to focus on the simplified PDE model (2) in the case where  $k$  is a positive integer. As shown in Section 3.1, the approach reduces to a Galerkin finite element method. Sections 3.2 and 3.3 discuss an efficient sampling of the spatial white noise vector and the handling of the fractional operator for real-valued  $k > d/4$ . From this point, the generalization to (9) is straightforward. Lastly, Section 3.4 explains how the SPDE generalizes to Matérn-type random fields on manifolds. We refer the reader to Appendix A for a small note on our MFEM implementation and its scalability.

#### 3.1. FEM discretization

Our description mainly follows [22, Sections 3.1]. Solving (2) with integer exponent  $k \geq 1$  on a domain  $D \subset \mathbb{R}^d$  reduces to a sequence of second-order PDEs:

$$\begin{cases} u_1 - \kappa^{-2} \Delta u_1 = \eta \mathcal{W} & \text{in } D, \\ u_{i+1} - \kappa^{-2} \Delta u_{i+1} = u_i & \text{in } D, \text{ for } i = 1, \dots, k-1, \end{cases} \quad (12)$$

with  $u := u_k$  and appropriate boundary conditions (e.g., Dirichlet or Neumann) on each  $u_i$ ,  $i = 1, \dots, k$ ; cf. [13].

The equations in (12) for  $u_i$ ,  $i = 2, \dots, k$ , lead to weak formulations that can be treated using standard finite element techniques; the interested reader is referred to [58] for further details. We focus on the case  $k = \eta = 1$  and, for simplicity, assume homogeneous Dirichlet boundary conditions. In this case, (12) reduces to

$$-\kappa^{-2} \Delta u + u = \mathcal{W} \text{ in } D, \quad u = 0 \text{ on } \partial D. \quad (13)$$

To solve (13) with finite elements, we proceed as follows. Let  $V_h = \text{span}(\phi_1, \phi_2, \dots, \phi_m) \subset H_0^1(D)$  be a finite element subspace spanned by the basis functions  $\phi_i$ ,  $i = 1, \dots, m$ . The Galerkin approximation  $u_h \approx u$  is then written

$$\text{Find } u_h \in V_h \text{ such that } \kappa^{-2} (\nabla u_h, \nabla v) + (u_h, v) = \langle \mathcal{W}, v \rangle, \text{ for all } v \in V_h, \quad (14)$$

where the notation  $(\cdot, \cdot)$  denotes the  $L^2(D)$ -inner product and  $\langle \mathcal{W}, \phi \rangle = \mathcal{W}(\phi)$  denotes the action of the generalized stochastic field  $\mathcal{W}$  onto the function  $\phi \in L^2(D)$ .

Our task is to rewrite (14) as a linear equation

$$\mathbf{Ax} = \mathbf{b}, \quad (15)$$

with a normally-distributed random right-hand side vector  $\mathbf{b}$ . We begin by referring to [59, Example 1.2 and Lemma 1.10], where it is shown that for the spatial white noise  $\mathcal{W}$  and any collection of functions  $\{\phi_i\}$  in  $L^2(D)$ , the actions

$$b_i = \langle \mathcal{W}, \phi_i \rangle \quad (16a)$$

are joint Gaussian random variables with zero mean and covariance given by

$$\mathbb{E}[b_i b_j] = (\phi_i, \phi_j). \quad (16b)$$

In turn, we find that expressing  $u_h = \sum_{i=1}^m x_i \phi_i(\mathbf{x})$  leads to the identities

$$A_{ij} = \kappa^{-2} (\nabla \phi_i, \nabla \phi_j) + (\phi_i, \phi_j) \quad \text{and} \quad b_i = \langle \mathcal{W}, \phi_i \rangle. \quad (17)$$

Moreover, by (16), we find that  $\mathbf{b}$  is a zero-mean Gaussian vector,

$$\mathbf{b} \sim \mathcal{N}(0, \mathbf{M}), \quad (18)$$

with covariance matrix equal to the standard finite element mass matrix  $M_{ij} = (\phi_i, \phi_j)$ .

Efficient procedures to assemble the finite element stiffness matrix  $\mathbf{A}$  are well-documented in the literature. It remains to state a technique to efficiently generate samples of the spatial white noise vector  $\mathbf{b}$ . This is done in Section 3.2 below. Once  $\mathbf{b}$  is sampled, the linear system (15) can be solved efficiently with standard FE system techniques, e.g., via various preconditioned linear equation solvers. The average cost can be further reduced if multiple samples of  $u_h$  are required simultaneously. Indeed, the matrix equation

$$\mathbf{A}\mathbf{X} = \mathbf{B}, \quad \mathbf{X} = \begin{bmatrix} | & & | \\ \mathbf{x}_1 & \cdots & \mathbf{x}_\ell \\ | & & | \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} | & & | \\ \mathbf{b}_1 & \cdots & \mathbf{b}_\ell \\ | & & | \end{bmatrix}, \quad (19)$$

is easily formed by first drawing  $\ell > 0$  samples of  $\mathbf{b} \sim \mathcal{N}(0, \mathbf{M})$  and then solved with the same preconditioner as if  $\ell = 1$ . Likewise, no new finite element solver is required if  $k > 1$  since the same elliptic operator appears in each of the equations in (12).

### 3.2. Sampling the spatial white noise vector

We understand from (18) that spatial white noise realizations  $\mathbf{b} \in \mathbb{R}^m$  can be generated by sampling a Gaussian vector whose covariance  $\mathbf{C}$  is given by the mass matrix  $\mathbf{M} \in \mathbb{R}^{m \times m}$ . Sampling such a vector can be thought of as executing a linear transformation that converts an  $n$ -dimensional Gaussian random vector  $\mathbf{z} = (z_1, \dots, z_n)^T \sim \mathcal{N}(0, \mathbf{I})$  with zero mean  $\mathbb{E}[\mathbf{z}] = \mathbf{0}$  and identity covariance  $\mathbb{E}[z_i z_j] = \delta_{ij}$ , to a new random vector

$$\mathbf{b} = \mathbf{H}\mathbf{z}, \quad (20)$$

where  $\mathbf{H} \in \mathbb{R}^{m \times n}$  decomposes the mass matrix  $\mathbf{M} = \mathbf{H}\mathbf{H}^T$ . Indeed, observe that for any such matrix  $\mathbf{H}$ , it holds that  $\mathbb{E}[\mathbf{b}] = \mathbf{H}\mathbb{E}[\mathbf{z}] = \mathbf{0}$  and, furthermore,

$$\mathbb{E}[\mathbf{b}\mathbf{b}^T] = \mathbb{E}[\mathbf{H}\mathbf{z}(\mathbf{H}\mathbf{z})^T] = \mathbf{H}\mathbb{E}[\mathbf{z}\mathbf{z}^T]\mathbf{H}^T = \mathbf{H}\mathbf{H}^T = \mathbf{M}. \quad (21)$$

Generating samples of  $\mathbf{z}$  is easy; pseudorandom number generators are available in all scientific software languages and may be used to independently sample each component  $z_i \sim \mathcal{N}(0, 1)$ ,  $i = 1, \dots, n$ . The main task is to construct  $\mathbf{H}$  efficiently. A tempting choice  $\mathbf{H} = \mathbf{L} \in \mathbb{R}^{m \times m}$  follows from the Cholesky factorization  $\mathbf{L}\mathbf{L}^T = \mathbf{M}$ . However, this choice may be prohibitively expensive if  $m$  is large since it involves an  $\mathcal{O}(m^3)$  factoring cost (from factoring  $\mathbf{L}\mathbf{L}^T = \mathbf{M}$ ) and  $\mathcal{O}(m^2)$  sampling cost (from matrix–vector multiplication  $\mathbf{b} = \mathbf{L}\mathbf{z}$ ). Instead, we promote the construction suggested in [22, Section 4.1] whose factoring and sampling costs both scale like  $\mathcal{O}(m)$ .

The idea is to leverage the assembly procedure that is used to construct global matrices in any typical finite element code (cf. Fig. 2) and perform the Cholesky decomposition in parallel at the element level. For the mass matrix, we have the factorization

$$\mathbf{M} = \mathbf{P}^T \mathbf{G}^T \mathbf{B}^T \mathbf{D} \mathbf{B} \mathbf{G} \mathbf{P}. \quad (22)$$

In this expression,  $\mathbf{P} \in \mathbb{R}^{m_{\text{local}} \times m}$  projects the (true) global degrees of freedom (DOFs) onto a space of local DOFs with redundancies at the intersections of local subdomains (usually restricted to separate processors) and  $\mathbf{G} \in \mathbb{R}^{m_{\text{elem}} \times m_{\text{local}}}$  further distributes these local DOFs to individual elements. Meanwhile, the matrix  $\mathbf{B} \in \mathbb{R}^{m_{\text{quad}} \times m_{\text{elem}}}$  and the diagonal matrix  $\mathbf{D} \in \mathbb{R}^{m_{\text{quad}} \times m_{\text{quad}}}$  involve function evaluation (at quadrature points) and perform multiplication with quadrature weights, respectively.

The block-diagonal matrix  $\tilde{\mathbf{M}} = \text{diag}(\mathbf{M}_T) = \mathbf{B}^T \mathbf{D} \mathbf{B}$  is comprised of the local mass matrices indexed by each individual element  $T$  in the mesh  $\mathcal{T}$ . Note that each  $\mathbf{M}_T$  is symmetric positive definite and, therefore, we can use a local Cholesky decomposition to factor each local mass matrix  $\mathbf{L}_T \mathbf{L}_T^T = \mathbf{M}_T$  independently. Moreover, defining  $\mathbf{H} = \mathbf{P}^T \mathbf{G}^T \text{diag}(\mathbf{L}_T) \in \mathbb{R}^{m \times m_{\text{elem}}}$ , we find that

$$\mathbf{M} = \mathbf{P}^T \mathbf{G}^T \tilde{\mathbf{M}} \mathbf{G} \mathbf{P} = \mathbf{P}^T \mathbf{G}^T \text{diag}(\mathbf{L}_T \mathbf{L}_T^T) \mathbf{G} \mathbf{P} = \mathbf{P}^T \mathbf{G}^T \text{diag}(\mathbf{L}_T) (\mathbf{P}^T \mathbf{G}^T \text{diag}(\mathbf{L}_T))^T = \mathbf{H}\mathbf{H}^T. \quad (23)$$

In summary, it is only necessary to factor the mass matrix locally, because of this, the cost of sampling the white noise vector  $\mathbf{b}$  scales linearly in the number of elements and DOFs. We encourage the interested reader to inspect our parallel and scalable implementation of the above sampling strategy in the latest release of MFEM [24,60].




Fig. 2. The finite element assembly procedure in a typical finite element code reduces to matrix multiplication. For the mass matrix, we multiply through the factorization  $M = P^T G^T B^T D B G P$  depicted above. Figure modified from [23,24].

### 3.3. Fractional exponents

The SPDE (2) includes a non-integer exponent, specifically  $k = \frac{2\nu+d}{4}$ . Solving boundary value problems with such fractional exponents, as they are known, presents unique challenges and demands substantial computational resources due to the non-local nature of the associated operator. However, the past decade has witnessed significant progress in developing efficient techniques to tackle these problems [61,62]. A notable approach is based on the Caffarelli–Silvestre extension, in which the fractional PDE solution is identified as a restriction of the solution to an integer-order boundary value problem set in a higher-dimensional semi-infinite cylinder [63]. This identification allows the higher-dimensional integer-order PDEs to be discretized using the Galerkin finite element method [64–66]. Another technique involves using an integral representation of the inverse fractional operator, combined with a specific quadrature rule, to break down the problem into a series of integer-order PDEs [67–69]. This technique has recently been applied to the SPDE [70]. An alternative strategy suggests modeling the fractional PDE as a transient pseudo-parabolic problem, where numerical integration over a ‘time’ parameter results in a sequence of integer-order PDEs [71,72]. In our implementation [60], we adopt a rational approximation method [20,21,73], which involves approximating the operator’s spectrum with a rational function. The multipole expansion of this function then leads to a set of independent integer-order PDEs that can be solved simultaneously. This method is summarized below.

The rational approximation method can be applied to a wide variety of elliptic differential operators  $\mathcal{A}$ , with spectrum  $\sigma(\mathcal{A}) \subset [\lambda_{\min}, \lambda_{\max}]$ ,  $0 < \lambda_{\min} < \lambda_{\max}$ . However, in (9), we are specifically interested in the fractional PDEs of the form

$$\mathcal{A}^k u = \eta \mathcal{W}, \quad k > d/4, \tag{24}$$

where  $\mathcal{A} := I - \kappa^{-2} \Delta$  with its associated boundary conditions. If the rational function  $\lambda \mapsto \sum_{n=1}^N \frac{c_n}{\lambda + d_n}$  approximates the function  $\lambda \mapsto \lambda^{-\alpha}$  on the interval  $[\lambda_{\min}, \lambda_{\max}]$ , i.e.,

$$\sum_{n=1}^N \frac{c_n}{\lambda + d_n} \approx \lambda^{-\alpha} \quad \text{for all } \lambda \in [\lambda_{\min}, \lambda_{\max}], \tag{25}$$

then the solution  $u$  can be approximated as the weighted average of the solutions of  $N$  problems written in terms of  $A$ ; namely,

$$u \approx \sum_{n=1}^N c_n u_n, \quad \text{where } (d_n I + \mathcal{A}) u_n = \eta \mathcal{W}. \tag{26}$$

Because  $A$  is an integer-order differential operator, each of the  $N$  problems above can be solved using the finite element method as described previously in this section. It remains to explain how to generate the weights  $c_n$  and the poles  $-d_n$  of the rational function (25). These can be obtained with one of the various rational approximation algorithms; see, e.g., [19–21]. In our MFEM implementation, we used the adaptive Antoulas–Anderson (AAA) algorithm proposed in [19]. From this point on, the generalization of the numerical scheme from (2) to (9) is straightforward and merely requires replacing the scalar number  $\kappa^{-2}$  with the matrix  $\Theta/2\nu$  in the weak form (14) and the corresponding matrix assembly.

The fractional exponent gives fine control over the smoothness of the GRF, which may not be achieved with integer-order exponents. Indeed, we remind the reader of Fig. 1, which depicts a variety of isotropic random fields that can be generated from the SPDE approach. In Fig. 3, we illustrate the impact of the regularity parameter  $\nu$  on GRFs generated via the generalized SPDE (9) with strong anisotropy. We use a pointwise non-linear transformation of the generated samples to demonstrate the effect. Such transformations are common in modeling material properties [74], loads, and inputs with marginal distributions [75] different from Gaussian. In the random field samples presented in Fig. 3, we threshold the generated GRF. All values of the GRF sample between zero and one are projected to one and all others to zero. We generated three GRF samples with enforced strong anisotropy and rotated them with different angles. The final image is obtained by pointwise summation and thresholding the result at one. The presented model can model the random distribution of fibers in composite materials. Close inspection of the images reveals that smoother fields, i.e., larger  $\nu$ , produce thicker fibers with well-defined spaces between them. The fibers’ size relative to the domain’s characteristic length can be scaled by modifying the correlation length  $l$ . On the other hand, less regular Gaussian fields,






Fig. 3. Combination of three, thresholded random fields to mimic material uncertainties for different smoothness (columns) and different white noise samples (rows). Thresholds are set to 0 and 1.

i.e., smaller values of the parameter  $\nu$ , result in well-separated clumps of thinner fibers. The clumping phenomena can be modeled only by increasing the roughness of the field. Thus, varying the parameter  $\nu$  combined with a non-linear transformation provides another control parameter for developing models of physically realistic distributions. Further examples for microstructure models based on the post-processing of Matérn-type random fields can be found in [10].

### 3.4. The SPDE method on manifolds

The SPDE approach is especially desirable to generate random fields on complex geometries and embedded surfaces. Indeed, the representation (12) provides a natural generalization to any sufficiently smooth manifold if the symbol  $\Delta$  is interpreted as the corresponding Laplace–Beltrami operator. In turn, instead of (1), the general formula for the covariance involves the eigenvalues  $\lambda_j$  and eigenfunctions  $\Psi_j$  of the Laplace–Beltrami operator,

$$C(\mathbf{x}, \mathbf{y}) = \sum_{j=1}^{\infty} \frac{c_k}{(k^2 + \lambda_j)^\alpha} \Psi_j(\mathbf{x})\Psi_j(\mathbf{y}). \tag{27}$$

Typically, a closed form for (27) cannot be derived because the eigenfunctions are unknown. Notable exceptions are intervals, rectangles, parallelepipeds, disks, sectors, spheres, spherical shells, ellipses, equilateral triangles, and circular and elliptical annuluses [76]. Truncated and spherical domains are particularly interesting in the context of the SPDE method and our examples. Khristenko and co-workers [50] thoroughly analyzed the effect of the boundaries on the SPDE covariance on  $D = (0, 1)$ . They derived closed from solutions for (27) with homogeneous Dirichlet and Neumann boundary conditions resulting in a sum over sine and cosine terms, respectively. They then found that the difference between the Matérn covariance (1) and the different expressions for (27) rapidly decays as the distance to the boundary increases and that the covariance is practically indistinguishable for points further than one to two correlation lengths away from the boundary. The authors additionally considered periodic boundaries, where they observe wrap-around effects due to the topology. Such effects are also present in the sphere. Here, the eigenfunctions are the *spherical harmonics* but (27) can be further simplified to a sum over the *Legendre* polynomials [51,77]. The Matérn-type covariance on the sphere is very similar to the Matérn covariance (1) if  $l/R \ll 1$ , where  $R$  denotes the radius of the sphere. Intuitively this can be understood from realizing that decreasing  $l$  on the sphere restricts the action of the filter defined by the Laplace–Beltrami operator to small local patches that are increasingly similar to  $\mathbb{R}^2$ . Finding the eigenfunctions and, therefore, the covariance of more intricate domains or manifolds requires numerical methods [78]. For further details regarding the SPDE on manifolds, we refer to Lindgren and co-worker’s original work [13] and the associated ten-year review [12]. Additional insights regarding the Matérn covariance on Riemannian manifolds can be found in the work of Borovitskiy and co-workers [79]. Lastly, we wish to highlight the recent work of Bonito et al. [70], providing a complete error analysis for GRFs on closed surfaces defined via the SPDE.

Fig. 4 displays a series of images depicting random field realizations on sphere, mobius strip, and Klein bottle geometries. These independent random fields were generated with our MFEM implementation of the SPDE method [60]. Additionally, we note that MFEM has long-standing native support for FEM discretizations of elliptic surface PDEs [23]. This feature was easy to leverage, and provided immediate support for the SPDE method on embedded surfaces without any code duplication.




Fig. 4. Matérn random fields on sphere, mobius strip, and Klein bottle geometries.




Fig. 5. The synthetic aneurysms are obtained from a reference geometry (grey mesh) [87] via the SPDE method and shifting the respective vertices in normal direction proportionally to the random field values (red-blue color bar) multiplied by a scaling parameter  $\alpha_S > 0$ . Regions with positive random field values are lifted while negative values are lowered. Matérn-type GRF parameters:  $l = 5.0$ ,  $\nu = 2.0$ , and  $\alpha_S = 0.2$ . We choose a larger scaling parameter  $\alpha_S$  than in Table 1 so that the geometric variations are easier to observe. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

#### 4. Simulating random fields on complex domains

To demonstrate the relevance of the SPDE method for geometric uncertainties, we borrow an example from biomedical engineering; namely, cerebral aneurysms. Cerebral aneurysms are weakened areas in the brain's blood vessel walls that bulge and sometimes rupture, potentially leading to fatality. Various factors influence cerebral aneurysm formation, including hemodynamics, genetics, and weakened vessel walls; see [80–82] and references therein. Once an aneurysm has formed, endovascular coiling has become a popular treatment option. In this procedure, medical professionals insert a thin wire into the aneurysm, forming a tiny coil to fill the space and prevent rupture [83–85].

Mathematical models for cerebral aneurysms and their treatment describe a complex interplay between fluid dynamics, biomechanics, and aneurysm geometry [82,86]. The computational domain, represented by a numerical mesh, is the foundation for these simulations. However, geometric uncertainties arise from the limitations of CT scans and subsequent post-processing techniques from which the meshes are extracted. These uncertainties propagate through the subsequent flow field and coiling simulations, potentially influencing predictions and conclusions that may be drawn from these analyses.

In the following, we assess uncertainties arising because of geometric ambiguities of the computational domain, i.e., the boundary of the synthetic aneurysm. Section 4.1 demonstrates the generation of synthetic aneurysms via the SPDE method followed by two forward problems for UQ presented in Section 4.2. We investigate how endovascular coiling deforms and stretches the aneurysm surface and complete the examples with blood flow simulations through perturbed geometries via the Stokes equation. The overarching goal of these forward problems is to illustrate the relationship between uncertainties in the aneurysm geometry and QoIs that influence the risk of rupture, such as the wall shear stress and pressure. This information can provide valuable medical insights and inform the development of more robust and reliable treatment strategies.

##### 4.1. Synthetic aneurysms

We generate synthetic aneurysms to quantify the influence of uncertainties in an aneurysm geometry. We proceed in three steps. First, we employ the SPDE method to sample an isotropic (i.e.,  $\Theta = I^2\mathbf{I}$ ) Matérn-type random field  $u$  on the surface of a reference

**Table 1**  
Parameters for the UQ aneurysm experiments.

Problem	SPDE			Membrane								
Parameter	$\nu$	$l_i$	$\alpha_S$	$T$	$h$	$\nu_M$	$g$	$a_k$	$b_k$	$t_1$	$t_2 = t_3$	$\alpha_M$
Value	2.0	5.0	0.1	0.1	0.05	0.5	10.0	4.0	1.0	20.0	0.05	0.05

aneurysm reconstructed from a CT scan [87]. Next, we estimate the outward-facing unit normal vector  $\mathbf{n}_k$  at each vertex  $\mathbf{v}_k$  of the triangulated aneurysm surface. To achieve this, we average the normals of the bordering elements on the triangulated surface. Lastly, we shift the vertices of the triangulated aneurysm surface proportionally to the values of the generated random fields; i.e., we translate the vertices  $\mathbf{v}_k \mapsto \mathbf{v}_k + \alpha_S u(\mathbf{v}_k) \mathbf{n}_k$ , where  $\alpha_S \in \mathbb{R}^+$ . This process introduces physical variations in the aneurysm geometry capable of modeling the uncertainty introduced by errors in the CT scans and associated post-processing techniques.

Fig. 5 shows an example of a synthetic aneurysm. It displays the grey mesh of the reference geometry [87] and the synthetic aneurysm colored in red and blue to indicate the values of the random field sample. The computational mesh comprises approximately  $71 \cdot 10^3$  nodes that define  $142 \cdot 10^3$  triangular elements. The shape of the random synthetic aneurysms depends in a natural way on the model parameters  $l$ ,  $\nu$ , and  $\alpha_S$ . Here, and throughout this section, we choose relatively long correlation lengths  $l$  and high smoothness parameters  $\nu$  to avoid creating unrealistic, rough, and highly-oscillatory aneurysm surfaces. The SPDE method allows us to generate as many synthetic aneurysm geometries as we wish in parallel. Computing a single random field realization as depicted in Fig. 5 ( $\nu = 2.0$ ,  $l = 5.0$ ) takes slightly less than one second on a standard notebook using 4 MPI ranks.

#### 4.2. Aneurysm deformation through coiling

We may now focus on uncertainty quantification with the synthetic aneurysm geometries. We consider an example motivated by endovascular coiling [88]. Here, we are interested in estimating how a medical coil inside the aneurysm stretches the surface. The mathematical model, introduced by Nobile and Vergara in [89], describes the displacement  $\eta$  of the aneurysm surface in the presence of a load  $f$  by a diffusion equation,

$$(\nabla \cdot T \nabla + \beta(\mathbf{x})) \eta(\mathbf{x}) = f(\mathbf{x}), \quad (28)$$

where the coefficient  $T$  describes the stiffness of the membrane and  $\beta$  is constructed from weighted curvature terms. Denoting the mean and Gaussian curvature as  $\rho_M$  and  $\rho_G$ , the latter coefficient reads

$$\beta(\mathbf{x}) = 2Th(2\rho_M(\mathbf{x})^2 - \rho_G(\mathbf{x})) + 2g(2\rho_M(\mathbf{x})^2 - (1 - \nu_M)\rho_G(\mathbf{x})), \quad (29)$$

with the parameters  $h$  and  $\nu_M$  denoting the membrane thickness and Poisson coefficient, respectively. The original authors derived a second-order time-dependent membrane equation from the Koiter shell model [90,91]. Our model further simplifies their pre-stressed model by assuming a stationary solution and an isotropic, spatially invariant coefficient  $T$ .

A detailed representation of the load  $f$  in (28) would require simulating the coiling [88] and coupling this to the membrane equation. Instead, we focus on the effect of geometric uncertainties on complicated manifolds and limit ourselves to an explicit model for the internal force  $f$  arising from the coiling. We sample a set of  $M$  points  $\mathbf{x}_{1,\dots,M}$  near the aneurysm center and model the force at a point  $\mathbf{x}$  on the aneurysm surface as

$$f(\mathbf{x}) = \max_{k \in \{1,\dots,M\}} (b_k \exp(-a_k^{-1}(\mathbf{x} - \mathbf{x}_k)^T \Theta_k (\mathbf{x} - \mathbf{x}_k))) \quad \text{with} \quad \Theta_k = \mathbf{R}_k^T \begin{pmatrix} t_1 & 0 & 0 \\ 0 & t_2 & 0 \\ 0 & 0 & t_3 \end{pmatrix} \mathbf{R}_k, \quad (30)$$

where  $a_k$ ,  $b_k$ , and  $t_i > 0$  parametrize the force. The matrices  $\mathbf{R}_k \in \text{SO}(3)$  are random rotation matrices obtained by sampling Euler angles.

Conveniently, the mathematical problem defined via (28) through (30) is similar to the SPDE problem. We discretize it using the finite element method and solve the arising linear system with the conjugate gradient method preconditioned with an algebraic multi-grid. For more details, we refer to Section 3. Computationally, the time for solving it is dominated by the approximate computation of the curvature terms in (29) [92], such that  $c = T_r/T_p \ll 1$  holds. Fig. 6 shows the solution of the forward problem for the parameter choices given in Table 1 and used throughout this section.

Based on the solution  $\eta$  of (28) depicted in Fig. 6(c), we define our quantity of interest (QoI) as the surface distortion resulting from shifting the vertices  $\mathbf{v}_k \mapsto \mathbf{v}_k + \alpha_M \eta(\mathbf{v}_k) \mathbf{n}_k$ , where  $\alpha_M \in \mathbb{R}^+$ . More specifically, we compute the *relative surface change* (RSC),

$$q_j = \frac{A_j^{\text{new}} - A_j^{\text{old}}}{A_j^{\text{old}}}, \quad (31)$$

an element-wise QoI, where  $A_j$  denotes the  $j$ th face surface area and  $j$  ranges over the number of faces in the mesh. Regions experiencing high stretches are prone to ruptures and, thus, are of significant interest for treatment planning [93]. Consequently, we are interested in the largest (or  $k$  largest) observed RSC values per sample.

Having defined the uncertainties, the mathematical forward model, and the QoIs, we now lay out the UQ workflow. We generate a set of 5000 synthetic aneurysms following the procedure described in Section 4.1 and using the parameters given in Table 1. The force field defined via (30) is generated once before generating the synthetic aneurysms and left unchanged for all samples.

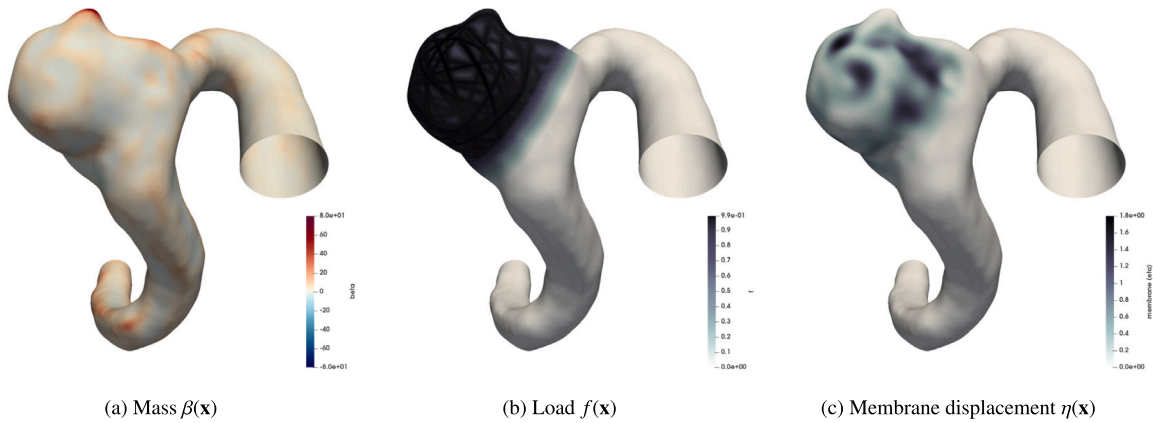


Fig. 6. The solution of Eq. (28) through Eq. (30) on the reference aneurysm. Images (a) and (b) show the functions defined in Eq. (29) and Eq. (30), respectively. (c) shows the solution  $\eta$  of Eq. (28) obtained with the finite element method. The model parameters are given in Table 1. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

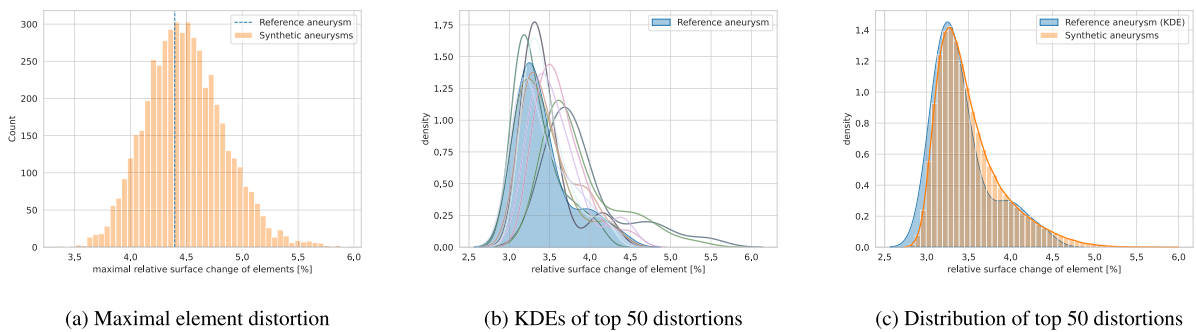


Fig. 7. Effects of uncertainties in the aneurysm geometry on the stretching of elements caused by a coil force. The *reference aneurysm* is obtained via a CT scan and depicted in Fig. 5. (a) Histogram of the maximal relative surface change of an element observed across 5000 synthetic samples. (b) KDE estimates of the largest 50 relative surface changes for the reference aneurysm (blue) and ten synthetic aneurysms. (c) Histogram of the largest 50 relative surface changes accumulated over 5000 samples. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Fig. 7 shows the results of this representative study. The histogram in Fig. 7(a) displays the distribution of the maximal RSC value in each sample. While the reference aneurysm obtained via a CT scan shows a maximal RSC of roughly 4.4% (blue), the synthetic aneurysms (orange) show a wide spread of the QoI, with some aneurysms showing RSCs of over 5.5%. This suggests that only simulating on the reference aneurysm may underestimate the risk of rupture during treatment. Fig. 7 (b,c) provides additional insights. Both figures show the distribution of the 50 largest RSC values for various samples. Fig. 7(b) shows the kernel density estimate (KDE) of the 50 largest RSC values for the reference aneurysm (blue, filled) and ten randomly chosen synthetic aneurysms. Fig. 7(c) presents a histogram of all of the top 50 RSC values across all synthetic RSC distributions. The heavy tails in both (b) and (c) provide additional support to the conclusion that only simulating on the reference aneurysm may underestimate the probability of a rupture.

The same GRF model can also be used for analyzing blood flow variations through different realizations of the aneurysm geometry. From basic fluid dynamics considerations, we expect the flow, particularly the pressure drop, to be very sensitive to geometric uncertainties. This can be understood from the Hagen–Poiseuille law, which states that the pressure drop is inversely proportional to the fourth power of the vessel’s radius [94]. However, because this law is derived for a cylindrical pipe instead of the complicated vessel geometry, we choose to numerically probe this hypothesis with a steady state Stokes flow model. We use a parabolic velocity profile prescribed at the inlet, free flow boundary conditions at the outlet, and no-slip boundaries at the vessel walls. The solutions for three different random vessels are shown in Fig. 8, and significant differences in the flow velocity and pressure distributions can be observed. Indeed, the first two vessels show similar pressure losses and velocity distributions. However, the third realization has a narrow diameter close to the outlet, which results in much more significant pressure loss and higher velocities towards the outlet. This change to the velocity profile will result in larger shear stresses on the aneurysm, increasing its risk of rupturing. The considered Stokes model is only utilized here for illustration purposes, and a more realistic blood flow model with a non-trivial Reynolds number and pulsating boundary conditions is outside of the scope of the current article, yet necessary for further medical analyses [82,86,95].




Fig. 8. Effects of geometric uncertainties for three realizations on a fluid flow throughout the aneurysm. The streamlines are colored by the pressure in the fluid using the color scale shown in the figure. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

## 5. Topology optimization under uncertainty

The ability to sample Gaussian random fields with a wide range of correlation lengths opens the possibility of investigating the impact of realistic operational, manufacturing, and material uncertainties on topology-optimized designs. Thus, the main goal of this section is to demonstrate these capabilities, discuss the computational complexity, and share details on the method's performance. In the following, Section 5.1 provides a concise introduction to the standard techniques and notations of topology optimization. The experienced practitioner may directly proceed to the selected examples, demonstrating the effect of such uncertainties. Section 5.2 considers a heat sink design problem subject to operational and material uncertainties. We treat the problem in 2D and on a surface embedded in 3D, and investigate how different correlation lengths, anisotropy, and symmetry affect the design. Section 5.3 presents the optimized design of bridges under stochastic excitation in 2D and 3D. If not explicitly stated to be otherwise, all GRFs generated via the SPDE (9) obey homogeneous Neumann boundaries. The vast collection of examples illustrates the flexibility and parallel scalability of the SPDE method, as well as the importance of incorporating uncertainties into the design problem.

### 5.1. Topology optimization

Topology optimization is an iterative design process aiming to find a material distribution by optimizing a specified performance measure and fulfilling a set of constraints [96]. Mathematically, the associated optimization problems can be written in the following form:

$$\text{minimize } j(\rho, \mathbf{u}) \text{ over } (\rho, \mathbf{u}) \in Z \times V \quad (32a)$$

$$\text{subject to } a(\mathbf{u}, \mathbf{v}; \rho) = \ell(\mathbf{v}) \text{ for all } \mathbf{v} \in V. \quad (32b)$$

Here,  $j: Z \times V \rightarrow \mathbb{R}$  is an integral functional subject to a set of constraints. The first constraint is a weak formulation of a partial differential equation (PDE). If the PDE is linear, then  $a: V \times V \rightarrow \mathbb{R}$  is a parametric bilinear form,  $\ell: V \rightarrow \mathbb{R}$  is a linear form,  $V$  is the admissible solution space of the weak formulation, and  $Z$  is the so-called design space defined as

$$Z = \left\{ \rho \in L^1(D) : 0 \leq \rho \leq 1 \text{ a.e.}, \int_D \rho \, dx \leq \gamma |D| \right\}, \quad (33)$$

for some fixed volume fraction parameter  $\gamma \in (0, 1)$ . Eq. (32b) is called the state equation, whose coefficients depend on a material distribution  $\rho \in Z$ , which in the topology optimization literature is referred to as the *density field*.

Assuming that (32b) is uniquely solvable for every  $\rho \in Z$ , the state variable  $\mathbf{u} = \mathbf{u}(\rho)$  in (32a) is often eliminated and the design problem is reduced to

$$\text{minimize } J(\rho) := j(\rho, \mathbf{u}(\rho)) \text{ over } \rho \in Z. \quad (34)$$

The above form provides the basis of the so-called nested design space approach in topology optimization [96]. It is often advantageous compared to the original formulation (32) because the state equations are always satisfied during the optimization process, often allowing the users to stop the iterative optimization process early.

For most practical problems, the density  $\rho \in Z$  should take on discrete values, zero or one, at every point in the design domain  $D \subset \mathbb{R}^d$ ,  $d = 2, 3$ . However, gradient-based optimization techniques require that the density vary continuously between zero and one. Such relaxation leads to 'grey' transition regions where  $0 < \rho(\mathbf{x}) < 1$ . The transition zone can be controlled by penalization [97],

projections [98], or by reformulating the optimization problem to include manufacturing uncertainties [53,99,100]. Applying a penalization, i.e., making all intermediate densities unfavorable, without incorporating additional restrictions on the density field leads to an ill-posed optimization problem, i.e., a mesh-dependent solution. A standard remedy is to employ filtering [101], where the material distribution is obtained by convolving the density field  $\rho$  with a filter kernel. Instead of an explicit convolution integral, the filtered density can be modeled as the solution of the following screened Poisson equation [102],

$$-\nabla \cdot r^2 \nabla \bar{\rho} + \bar{\rho} = \rho \quad \text{in } D, \quad (35)$$

where  $r > 0$  is often referred to as a length scale parameter for design. This PDE filter is equivalent to the standard convolution filter [102]. However, it offers additional benefits such as convenient enforcement of boundary conditions on the filtered density field  $\bar{\rho}$ , using the same finite element discretization machinery for solving the state PDE, and an efficient parallel implementation. Typically, the homogeneous Neumann boundary condition  $\nabla \bar{\rho} \cdot \mathbf{n} = 0$  is applied along the boundary  $\partial D$ , with  $\mathbf{n}$  denoting the outward-facing unit normal. However, mixed Dirichlet/Neumann boundary conditions enforcing the density to take values zero or one on specified parts of the boundary,  $\rho = 0 \vee 1$  on  $\Gamma_D \subset \partial D$ , are often employed in practical applications.

The most popular penalization approach in density-based compliance topology optimization is the so-called SIMP (solid isotropic material with penalization) interpolation [103]. Taking the modulus of the elasticity as an representative example, it is written as

$$E = E_{\min} + \bar{\rho}^p (E_{\max} - E_{\min}), \quad (36)$$

where  $\bar{\rho}$  is the filtered density field,  $p$  is a penalization parameter, and  $E_{\max}$  is the modulus of elasticity of the solid material.  $E_{\min}$  is a small positive constant ensuring that the state equation (32b) is well-defined. Depending on the application, analogous equations govern other physical coefficients in the state PDE, for instance, the material diffusion coefficient  $\kappa$  in (42) below. In [103],  $E_{\min}$  is set to zero, and in such cases,  $\rho$  and correspondingly  $\bar{\rho}$  has to be bounded from below to remain greater than zero. The penalization (36) is controlled by the power parameter  $p$ , which, in combination with an active volume constraint, makes intermediate densities  $0 < \rho(\mathbf{x}) < 1$  unfavorable. The penalization parameter is often set to be  $p \geq 3$ , with larger values resulting in a steeper transition between void and solid. The intermediate densities cannot be removed entirely in the final optimized solution, and if necessary, these parts of the design can be realized by using spatially varying microstructures [97].

Projections were introduced in [104] to sharpen the transition from void to solid. The original idea is to use a filter function with finite support combined with a pointwise Heaviside projection. Applying a zero threshold to the Heaviside projection – where all filtered densities above zero are projected to one – ensures that the final solid design corresponds to the support of the filtered density function  $\bar{\rho}$ , thereby guaranteeing the non-trivial size of all solid features in the design. Conversely, setting the threshold to one in the Heaviside projection – meaning all filtered densities below one are projected to zero – enforces a non-zero size upon all void features in the design. The Heaviside projection is replaced with a non-linear differentiable expression in practical implementations, with its transition sharpness is controlled by a single parameter.

The above approach does not work for the PDE filter regularization approach because the filter is a Green's function whose support coincides with the entire computational domain. Thus, the so-called robust formulation was introduced in [99] to alleviate the lack of a length scale in either the void or the solid part of the design for the classic convolution filter with finite support size, as well as to sharpen the transition regions for problems regularized with the screened Poisson Eqs. (35). The projection in [99] is performed point-wise using the following non-linear transformation:

$$\bar{\rho} = \frac{\tanh(\beta\eta) - \tanh(\beta(\bar{\rho} - \eta))}{\tanh(\beta\eta) - \tanh(\beta(1 - \eta))}, \quad (37)$$

where  $\eta \in (\eta_d, \eta_e)$  is a threshold parameter,  $0 < \eta_d \leq \eta_e < 1$ , and  $\beta$  controls the sharpness of the projection. The above expression approaches the Heaviside projection as  $\beta \rightarrow \infty$ . The projection alone does not enforce a length scale on the design. The length scale is enforced by minimizing the worst possible compliance from all possible realizations for  $\eta$  between  $\eta_d$  and  $\eta_e$ , with a bound on the worst possible volume. Thus, the optimization problem (32) is replaced by

$$\begin{aligned} & \text{minimize} && \max_{\eta_d \leq \eta \leq \eta_e} j(\bar{\rho}(\eta), \mathbf{u}) \quad \text{over } (\rho, \mathbf{u}) \in Z \times V \\ & \text{subject to} && a(\mathbf{u}, \mathbf{v}; \bar{\rho}(\eta)) = \ell(\mathbf{v}) \quad \text{for all } \mathbf{v} \in V \\ & && a_f(\bar{\rho}, w) = \ell_f(w; \rho) \quad \text{for all } w \in W, \end{aligned} \quad (38)$$

where the equation with the linear form  $\ell_f: W \times Z \rightarrow \mathbb{R}$  represents the PDE filter equation (35), and the design space is redefined to be

$$Z = \left\{ \rho \in L^1(D) : 0 \leq \rho \leq 1 \text{ a.e.}, \max_{\eta} \int_D \bar{\rho}(\rho) \, dx \leq \gamma |D| \right\}.$$

The compliance objective is the most popular choice of performance measure. It is monotonic with respect to the threshold parameter  $\eta$ . Thus, searching for the worst possible objective and material volume case can be avoided [98]. Threshold  $\eta = \eta_d$  represents the most dilated realization,  $\eta = \eta_e$  the most eroded one, with all realizations in between providing a simple model for the outcomes of a production process with uniform manufacturing errors along the perimeter of the topology-optimized design. Including these simplified production uncertainties in the optimization process avoids post-processing and engineering intervention between design and manufacturing [105].

Uniform erosion or dilation of the design is a crude approximation of real-life manufacturing uncertainties. These can vary spatially and, in many cases, introduce changes in the topology between different realizations. Therefore, more realistic models where the design's geometry or material stiffness varies spatially were proposed in [53,54]. For geometric uncertainties, the threshold  $\eta$  is modeled as a random field, and for material uncertainties, the random field represents the modulus of elasticity. In these models, finding the worst case is extremely expensive, and the maximum in (38) is replaced by the expectation or another risk measure.

In this work, we use several topology optimization problems to demonstrate various properties and effects of the SPDE approach to modeling random fields. The exact objectives and constraints are explained for every numerical example individually, ensuring a comprehensive understanding of the application and impact of the SPDE method in each case. It is important to note that the threshold interpretation of density-based topology optimization introduced in [99] shares similarities with many level-set approaches [106,107], and such formulations under geometric, material, or operational uncertainties can be found in [55,108–110]. Thus, the SPDE representation of random fields applies beyond the density-based formulations utilized here and could also be employed in other topology optimization approaches.

## 5.2. Heat sink

Optimizing thermal compliance is a textbook example of topology optimization [96]. The goal is to find a topology that extracts heat from a given domain  $D$  by minimizing the thermal compliance of the design, given a distributed heat source  $f$  on  $D$ . The problem may be considered with a constant (deterministic) or spatially-varying (stochastic) heat source. We first describe the physical model and the deterministic problem in 5.2.1. Afterward, we explore how GRFs influence the design in 5.2.2, explain efficient symmetrization techniques based on the SPDE method in 5.2.3, and demonstrate the effect of anisotropy in 5.2.4. Sections 5.2.5 and 5.2.6 present additional examples of optimized topologies on embedded manifolds and those subject to random material damage, respectively.

### 5.2.1. Deterministic thermal compliance optimization

We describe the topology with a zero/one bounded density field with intermediate grey values penalized using SIMP material interpolation with penalty parameter  $p = 3$ . We consider the domain  $D = (-0.5, 0.5)^2$  upon which we define a heat source represented by the function  $f$ . The domain's boundary  $\partial D$  is divided into two parts,  $\Gamma_N$  and  $\Gamma_D$ . The boundary  $\Gamma_N$  is insulated, i.e., the normal gradient of the heat flow along the boundary is zero. On  $\Gamma_D$ , with length  $1/7$  centered at the bottom of the domain, the temperature is fixed to zero. A diffusion equation describes the physics of the heat transfer, namely,

$$-\nabla \cdot \kappa(\bar{\rho}) \nabla u = f \quad \text{in } D, \quad \nabla u \cdot \mathbf{n} = 0 \quad \text{on } \Gamma_N, \quad u = 0 \quad \text{on } \Gamma_D, \quad \overline{\Gamma_N \cup \Gamma_D} = \partial D, \quad (39)$$

where  $u$  represents the temperature in the domain. The weak formulations of (39) is given as follows,

$$a_1(u, v) = \ell_1(v), \quad \forall v \in V_0 := \{v \in H^1(D) : v = 0 \text{ on } \Gamma_D\}. \quad (40)$$

with bilinear form  $a_1$  and linear form  $\ell_1$  defined by

$$a_1(u, v) = \int_D (\kappa(\bar{\rho}) \nabla u \cdot \nabla v) \, d\mathbf{x}, \quad \ell_1(v) = \int_D f v \, d\mathbf{x}. \quad (41)$$

At every point of the design domain, the diffusion coefficient is evaluated as

$$\kappa(\bar{\rho}) = \kappa_{\min} + \bar{\rho}^p (\kappa_{\max} - \kappa_{\min}), \quad (42)$$

with  $\bar{\rho}$  obtained by solving the finite element discretized PDE-filter equation (35) with weak bilinear and linear forms ( $W = H^1(D)$ ) defined as

$$a_f(\bar{\rho}, v) = \int_D (r^2 \nabla \bar{\rho} \cdot \nabla v + \bar{\rho} v) \, d\mathbf{x}, \quad \ell_f(v) = \int_D \rho v \, d\mathbf{x}. \quad (43)$$

In the deterministic setting, the heat source is a constant  $f = 1$ , and the objective is to minimize thermal compliance given as

$$j(\rho, u) = \int_D u f \, d\mathbf{x}, \quad (44)$$

with  $u$  defined via (39) and with  $\rho$  subject to a volume constraint

$$\int_D \bar{\rho}(\mathbf{x}) \, d\mathbf{x} \leq V. \quad (45)$$

We note that (44) depends indirectly on  $\rho$ , i.e.,  $u = u(\bar{\rho}(\rho))$ . The obtained optimized topology minimizing the thermal compliance for the uniform deterministic heat source is shown in Fig. 9. The deterministic problem is self-adjoint, and gradients of the thermal compliance with respect to the discrete representation of the original density field  $\rho$  are obtained using the chain rule. The domain is discretized using quad elements, with approximately 0.33M DOFs representing the temperature field.



Fig. 9. Optimized density  $\bar{\rho}$  for thermal compliance problem with deterministic heat source and volume constraint 50% (i.e.,  $\gamma = 0.5$ ). The length parameter in the PDE filter is set to  $r = 0.02$ .

### 5.2.2. Stochastic thermal compliance optimization via SPDE

In the stochastic setting, we model  $f$  as a zero-mean Matérn-type GRF with different correlation length parameters and smoothness  $\nu = 1$  obtained via the SPDE method, i.e., as the solution to (9) using the same discretization as the physical problem. The considered stochastic optimization problem is formulated akin to the deterministic one. The objective becomes the minimization of the *expected* thermal compliance defined as

$$j(\rho, u) = \mathbb{E} \left[ \int_D u f \, dx \right] \approx \frac{1}{N} \sum_{i=1}^N \int_D u_i f_i \, dx. \quad (46)$$

For the numerical examples to follow, we approximate the expectation operator with the fixed number of samples  $N = 300$ , and each optimization step uses the same 300 GRF samples. We determined the number of samples with numerical experiments systematically testing the effect of the sample size on the objective and resulting designs; from 300 samples onwards, the designs and objective showed only very minor differences. More details may be found in Appendix B. We remark that more sophisticated, adaptive algorithms can address the approximation of the expectation operator during stochastic optimization, such as the ones presented in [5,111–115]. They begin with a small number of samples and adaptively increase it as the solution approaches the minimum. This allows the algorithms to reduce the number of forward model evaluations needed to converge to the solution. However, due to the stochastic nature of the algorithm and the small number of initial samples, these algorithms are prone to falling into a local, asymmetric minimum. For the heat sink with a risk-neutral objective, i.e., mean compliance, an adaptive sampling algorithm is presented in the authors' earlier work [114] producing results similar to the ones presented in this section. As the presented SPDE approach applies to any optimization algorithm based on sampling, and our focus is on the effect of uncertainties on the design, the results presented here will use the most popular optimization algorithm in topology optimization, the Method of Moving Asymptotes (MMA) [116]. Additionally, one may use more suitable risk measures instead of the expectation in (46), for example, the *conditional value-at-risk* [5,117–119]. However, such risk measures often require specialized numerical optimization algorithms lying beyond the scope of this work.

Comparing the results of the deterministic and stochastic problems in Figs. 9 and 10, respectively, we observe significant differences. The deterministic optimization problem finds a tree-like structure trying to extract the heat through shortest paths to the support without violating the volume constraint. The stochastic formulation results in many closed loops of solid material decorated with more minor tree-like features. The reason for this difference can be understood by realizing that the only mechanism to extract heat from a point in a domain with uniform source is to connect conductive material to the fixed temperature support. On the other hand, when the source is a random field that provides both positive and negative values spatially varying throughout the domain, an area with negative heating can be locally neutralized by an area with positive heating by connecting them with conductive material. Therefore, the role of the closed loops in the obtained topologies is to establish local heat equilibrium, and only the excess heat is transferred to the support by the underlying tree-like skeleton. Decreasing the correlation length results in smaller loops reflecting the average size of the positively/negatively heated regions.

Intuitively, increasing the correlation length yields GRF samples that show fewer local fluctuations and become more locally uniform. This behavior can be observed in Fig. 1 along the correlation length axis. Here, the areas with similar field values become larger as  $l$  increases, and in the limit of  $l \rightarrow \infty$ , the SPDE solutions become constant in space. Thus, with increasing lengths, the stochastic case behaves more and more similar to the deterministic case with its uniform load. Our results support this reasoning; by lowering the correlation length from (a) to (d) in Fig. 10, we observe that the resulting designs develop structures which increasingly deviate from the deterministic case. While we observe two dominant main branches diagonally extending from the center to the top left and right corners in Panel (a) as for the deterministic problem, Panel (d) shows many small loops, and most branches have similar widths. From (a) to (d), the main branches become less pronounced as the correlation length decreases, resulting in significantly different designs compared to the deterministic case.

Another striking difference between the deterministic and stochastic designs is the lack of symmetry of the latter. The deterministic problem is symmetric about the  $y$ -axis, i.e., all operators and loads in the SPDE and physical problem are invariant under reflection. Introducing GRFs as loads breaks this symmetry since  $f(x, y) \neq f(-x, y)$  with probability equal to one for all sample






Fig. 10. Optimized density distributions  $\bar{\rho}$  (bottom) for thermal compliance problems with stochastic heat source of different correlation lengths (top, representative sample). The length parameter in the PDE filter is set to  $r = 0.02$ , and the volume constraint is 50% (i.e.,  $\gamma = 0.5$ ). The stochastic fields are isotropic, i.e.,  $\Theta = I^2 \mathbf{I}$ . The value of the objective function is denoted with  $j$ .

realizations. Consequently, using a finite number of samples inevitably leads to asymmetric designs. While considering larger sample sizes  $N$  yields increasingly symmetric designs (see Appendix B), this approach is computationally infeasible and we propose to enforce the natural symmetry via a modification to the algorithm.

### 5.2.3. Enforcing symmetries

Although the individual GRFs break the system’s symmetry, the expectation  $\mathbb{E}(f(x, y)) = \mathbb{E}(f(-x, y))$  preserves it. In fact, the reflection symmetry extends to the full probability density function of  $f$  at each  $(x, y) \in D$ . This, together with the same reflection symmetry in the state equation (39), boundary conditions, and other problem constraints ensures that at least a locally optimal symmetric design exists. While this is not sufficient to conclude that a globally optimal design must obey any symmetry, it shows that it is of substantial interest to investigate optimized designs exhibiting the same reflection symmetry found in the deterministic problem. The following changes to the algorithm ensure the design’s symmetry while simultaneously lowering the computational cost. We describe these changes in detail because their complexity is greater than in the deterministic optimal design setting.

As before, we begin with a symmetric initial guess for the density field satisfying the volume constraint (45). The gradients and the corresponding design updates preserve symmetry if we consider mirrored pairs of GRF realizations  $(f(x, y)$  and  $\bar{f}(x, y) = f(-x, y))$  instead of randomly chosen ones. Conveniently, the SPDE allows for constructing such pairs efficiently, only considering half of the domain  $D_{1/2} = (-0.5, 0) \times (-0.5, 0.5)$ . Moreover, the solution of the state equation (39) can also be computed on  $D_{1/2}$  since it exhibits the same symmetry.

We consider a GRF  $f$  generated via the SPDE method for a given, random white noise realization  $\mathcal{W}$  and decompose it into its symmetric and antisymmetric components, i.e.,  $f = (f_s + f_a)/\sqrt{2}$  and  $\bar{f} = (f_s - f_a)/\sqrt{2}$ . Linearity of the SPDE ensures that both  $f_s$  and  $f_a$  satisfy (9) with white noise realizations  $\mathcal{W}_s$  and  $\mathcal{W}_a$  that are proportional to the symmetric and antisymmetric components of  $\mathcal{W}$ . Symmetry and anti-symmetry imply  $\partial_x f_s(x = 0, y) = 0$  and  $f_a(x = 0, y) = 0$ . In other words,  $f_s$  and  $f_a$  obey homogeneous Neumann and Dirichlet boundary conditions, respectively, along the centerline of the domain  $D$ , defined  $\Gamma_C = \{(x, y) \in D \mid x = 0\}$ . Well-posedness of the SPDE (9) implies that we can solve the equation on  $D_{1/2}$  with homogeneous Dirichlet boundary conditions on  $\Gamma_C$  and white noise  $\mathcal{W}_a|_{D_{1/2}}$  to obtain  $f_a|_{D_{1/2}}$ . Likewise, using homogeneous Dirichlet boundary conditions and white noise  $\mathcal{W}_s|_{D_{1/2}}$  gives  $f_s|_{D_{1/2}}$ . Together,  $f_s|_{D_{1/2}}$  and  $f_a|_{D_{1/2}}$  allow us to construct  $f|_{D_{1/2}}$  and  $\bar{f}|_{D_{1/2}}$ , which can be reflected to construct  $f$  and  $\bar{f}$  at any point in  $D$ , if necessary. In practice, we do not derive  $\mathcal{W}_s$  and  $\mathcal{W}_a$  from  $\mathcal{W}$ , instead we simply generate two independent Gaussian white noise samples on  $D_{1/2}$  using the procedure explained in Section 3.2.

The GRFs  $f$  serve as loads in the state equation of this Eq. (39). Given that the physical problem is linear and also exhibits a reflection symmetry, we may restrict the entire optimization workflow to  $D_{1/2}$  by considering the symmetric and antisymmetric solutions of (39),  $u_s$  and  $u_a$ , respectively. The boundary conditions are analogous to the restricted SPDEs discussed in the previous paragraph, meanwhile,  $f_s$  and  $f_a$  become the loads for the symmetric and antisymmetric restricted state equations, respectively. Finally, we may recover the final temperature fields  $u$  and  $\bar{u}$  corresponding to  $f$  and  $\bar{f}$  from the equations

$$u = \frac{1}{\sqrt{2}} (u_s + u_a), \quad \bar{u} = \frac{1}{\sqrt{2}} (u_s - u_a). \tag{47}$$




Fig. 11. Optimized symmetric density distributions  $\bar{\rho}$  for thermal compliance problems with stochastic heat source, different correlation lengths, and volume constraint 50% (i.e.,  $\gamma = 0.5$ ). The length parameter in the PDE filter is set to  $r = 0.02$ . The stochastic fields are isotropic, i.e.,  $\Theta = I^2 \mathbf{I}$ . Symmetry is enforced across the vertical axis  $x = 0$  of the design domain. The optimized value of the objective function is denoted by  $j$ .

Although the above modification requires twice as many PDE solves, it reduces the overall computational cost because each solve requires only half of the original domain.

Fig. 11 shows the stochastic heat sink designs with enforced symmetry. The proposed modification produces symmetric designs with topological features similar to Fig. 10 for a lower computational cost. Comparing Figs. 10 and 11 shows that similar correlation lengths yield similar structures. For instance, the main branches in (a), (b), and (c) show strong similarities. Moreover, the characteristic loop sizes match. We observe that the point where the main branches split above  $x = 0$  moves closer to the bottom boundary with decreasing correlation length. This point seems to be consistently lower for the symmetrized designs. The most significant differences between symmetrized and non-symmetrized designs appear for short correlation lengths, i.e., in panels (d). Here, the main branches in the symmetrized design are completely separated, while the non-symmetrized branches are still connected. Overall, the latter branches appear thicker than the symmetrized ones.

The optimization of the symmetrized examples was carried out entirely on  $D_{1/2}$ . The images in Fig. 11 are post-processed such that the second half is added by reflection. While every optimization on  $D_{1/2}$  with subsequent reflection will produce symmetric results, we emphasize that the presented algorithm is equivalent to running the optimization on  $D$  with mirrored sample pairs. However, it bypasses the need for a symmetric mesh and reduces computational costs.

#### 5.2.4. Effects of anisotropy


Introducing different correlation lengths,  $l_i \neq l_j$  for  $i \neq j$ , when constructing the matrix  $\Theta$  in (9) allows the SPDE method to generate anisotropic random fields. The effect of such anisotropy on topology-optimized heat sinks is demonstrated in Fig. 12. Panel (a) shows the designs resulting from using long correlation lengths in the horizontal  $x$ -direction and short correlation lengths in the vertical  $y$ -direction, and vice-versa for panel (b). In both cases, the anisotropy does not break the mirror symmetry of the problem, and we employ the symmetrization technique proposed in Section 5.2.3 above. The two designs show significant differences that reflect the implied directionality of the different random heat sources. Additionally, we may rotate  $\Theta$  with a standard 2-dimensional rotation matrix; cf.  $\mathbf{R}$  in (9). Introducing this rotation breaks the symmetry, and we solve the problem without symmetrization. Panels (c) and (d) depict the designs given two different rotation angles  $\vartheta = \pm\pi/4$ . Overall, we find that anisotropy significantly biases the topologies, demonstrating the modeling freedom of the SPDE method. While not explicitly considered here, we remark that the method generalizes to spatially-varying coefficients  $\Theta = \Theta(\mathbf{x})$ , providing great flexibility in the construction of inhomogeneous fields.

#### 5.2.5. Optimization on embedded surfaces

We continue investigating stochastic thermal compliance optimization and consider a surface embedded in 3D. In this case,  $D$  is now a manifold equal to the surface of the unit-radius sphere intersected with a half-space passing through the point  $(0, 0, 4/5)$  and parallel to the  $x$ - $y$ -plane; see Fig. 13. More specifically,  $D$  is defined to be

$$D = \{(x, y, z) \in \mathbb{R}^3 \mid x^2 + y^2 + z^2 = 1 \text{ and } z < 4/5\}. \tag{48}$$

The problem formulation is identical to the one posed on the unit square; however, the boundary  $\partial D$  is now the circle  $\{(x, y, z) \in \mathbb{R}^3 \mid x^2 + y^2 + z^2 = 1 \text{ and } z = 4/5\}$ . We impose a homogeneous Dirichlet boundary condition along the entire boundary  $\Gamma_D = \partial D$ . Moreover, all Laplace-type operators in the problem formulation (39) are now interpreted as Laplace-Beltrami operators. For a detailed discussion of the covariance of GRFs on the sphere, including the Matérn-type covariance, we refer to [51,77,120]. We initialize the density to  $\rho_0 = (\sin(5\varphi) + 1)/2$  with  $\varphi \in (0, 2\pi)$  denoting the azimuthal angle. This prescription satisfies the volume constraint (45) with  $\gamma = 0.5$  and biases the topology towards a periodic pattern. Besides these minor differences, the optimization proceeds akin to the previous cases. Fig. 13 shows the density fields  $\bar{\rho}$  resulting from random heating modeled by isotropic Matérn-type GRFs of different correlation lengths. Similar to optimization on the unit square, decreasing the correlation length results in fast oscillatory fields, causing the optimization to form smaller closed loops of solid material to equilibrate the heat distribution locally. The tree-like skeleton transfers the excess heat to the support. Such capability extends to arbitrary surfaces embedded in  $\mathbb{R}^3$ , enabling heat sink designs on complex surfaces such as battery packs, tablets, and laptops [121].



(a)  $l_x = 0.2, l_y = 0.025, \vartheta = 0, j = 0.109$       (b)  $l_x = 0.025, l_y = 0.2, \vartheta = 0, j = 0.110$       (c)  $l_x = 0.2, l_y = 0.025, \vartheta = \frac{\pi}{4}, j = 0.114$       (d)  $l_x = 0.2, l_y = 0.025, \vartheta = \frac{-\pi}{4}, j = 0.114$

Fig. 12. Optimized density distributions  $\bar{\rho}$  (bottom) for thermal compliance problems with stochastic heat source of different correlation lengths (top, representative sample). The length parameter in the PDE filter is set to  $r = 0.02$ , the volume constraint is 50% (i.e.,  $\gamma = 0.5$ ) and the stochastic fields are anisotropic. For (a,c), symmetry is enforced across a vertical axis  $x = 0$ . The value of the objective function is denoted with  $j$ .

5.2.6. Effects of random material damage

Apart from operational uncertainties like external excitations or time variations in boundary conditions, material properties are another common source of randomness. Examples of such formulations can be found in [17,54,122]. The experience from these studies shows that differences between design topologies are barely noticeable for slowly varying uncertain material properties, modeled by random fields with long correlation lengths compared to the characteristic size of the design domain. On the other hand, [123] presents a very different outcome by modeling the localized damage caused by a small random projectile in the design domain. This form of uncertainty leads to significant differences in the optimized topologies. Indeed, redundant members are clearly observable in [123], providing numerous possibilities to transfer the load to the supports in case neighboring members fail.

These summarized examples from the literature represent two idealized corner cases that are relatively easy to model with limited computational resources but avoid the full range of spatial variations necessary for more realistic damage and material models. The SPDE approach allows for representing intermediate cases as demonstrated in Fig. 14. Instead of using SIMP interpolation (42) as before, we employ the robust formulation of topology optimization [99] with thresholds  $\eta_d = 0.3$  for the dilation and  $\eta = 0.7$  for the erosion parameters. The dilated volume is restricted to be less than 35% of the total volume of the design domain, and the filter parameter is selected to be  $r = 0.05$ . The projection parameter is chosen to be  $\beta = 8$ . In this example, the uncertainty is encompassed solely within the diffusion coefficient and the load is deterministic. More specifically, we use a uniform heat input  $f$  and model the random diffusion coefficient in (39) as

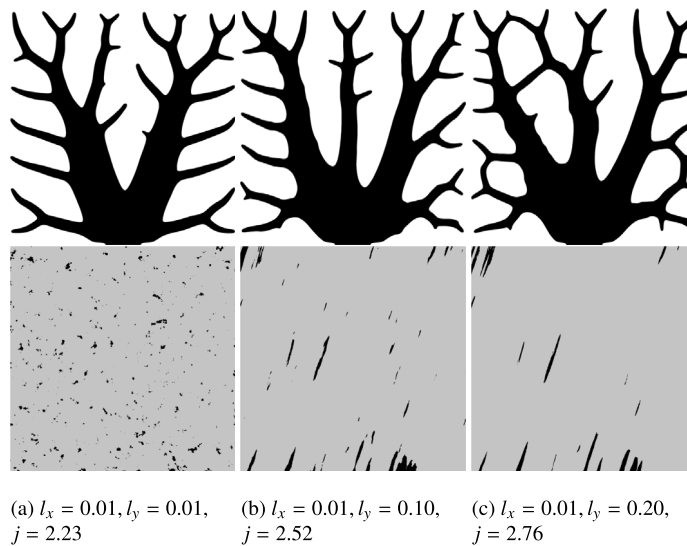
$$\kappa(\bar{\rho}) = \kappa_{\min} + \tau \bar{\rho} (\kappa_{\max} - \kappa_{\min}), \tag{49}$$

where  $\tau$  is a binary (zero/one) field obtained by a pointwise transformation of the SPDE solution. This transformation projects all values of above 2.5 to zero, and maps the remaining values to one. The zero values model local damage or material inclusions with very low conductivity.

Similar to the examples shown in Fig. 3, the shape of the damage regions, shown in the bottom row of Fig. 14, and their scale and relative positions can be controlled easily. The impact of such material variations can be seen on the top row. Transforming an isotropic Gaussian random field with a small correlation length cannot impact the optimized topology since it results in a model with small, well-distributed damage regions and characteristic sizes much smaller than the length scale enforced on the design. Increasing the correlation length in one of the directions and rotating the random field results in elongated damage regions following the prescribed orientation. Furthermore, the small damage regions are replaced by elongated damage zones with larger lengths than the length scale of the design. Therefore, the optimization provides alternative paths to transfer the heat to the support and creates closed loops of high-conductivity material. It should be pointed out that the selected material model visually affects the topology of an optimization formulation with mean value (risk-neutral) objective. This stands in contrast to all previously mentioned cases in the literature, where topology changes are only observed for risk-averse objective functions like the sum of mean and standard deviation or the min-max formulation.



**Fig. 13.** Optimized density distributions  $\bar{\rho}$  for thermal compliance problems with stochastic heat source, different correlation lengths, and volume constraint 50% on a cut sphere. The length parameter in the PDE filter is set to  $r = 0.02$ . The stochastic fields are isotropic. The values of the objective function are  $j_{(a)} = 39.74$ ,  $j_{(b)} = 17.75$ , and  $j_{(c)} = 7.61$ .



(a)  $l_x = 0.01, l_y = 0.01, j = 2.23$ , (b)  $l_x = 0.01, l_y = 0.10, j = 2.52$ , (c)  $l_x = 0.01, l_y = 0.20, j = 2.76$

**Fig. 14.** The figure shows optimized topologies (top) for heat sinks obtained using the robust formulation with materials uncertainties. Samples of the material distributions are shown in the bottom row and correspond to  $\tau$  in Eq. (49). In this row, the color grey indicates  $\tau = 1$ , while black indicates regions with  $\tau = 0$ . The random damage fields are rotated 0.3 radians counterclockwise with correlation lengths specified above for each optimization case.

### 5.3. Bridge design

We are now concerned with finding the optimal topology of a bridge to support a given load. We use the SPDE method to impose random fluctuations on the boundaries of a design problem, mimicking random horizontal loads. For all 2D and 3D solid mechanics examples, the system’s response is governed by the Navier–Cauchy PDE, given as

$$\begin{aligned}
 -\operatorname{div} \boldsymbol{\sigma}(\mathbf{u}) &= \mathbf{f}, \\
 \boldsymbol{\sigma}(\mathbf{u}) &= \mathbf{C} : \boldsymbol{\varepsilon}(\mathbf{u}) \quad \text{in } D,
 \end{aligned}
 \tag{50}$$

where  $\boldsymbol{\sigma}$  is the so-called stress tensor, and  $\boldsymbol{\varepsilon}$  is the linearized strain tensor defined as  $\boldsymbol{\varepsilon} = (\nabla \mathbf{u} + \nabla \mathbf{u}^T) / 2$ .  $\mathbf{C}$  is an elastic material properties tensor,  $\mathbf{u}$  denotes the displacement field, and  $\mathbf{f}$  is a distributed input supplied to the system. The boundary  $\partial D = \overline{\Gamma_{D_i} \cup \Gamma_{N_i}}$ ,  $i = 1, \dots, d$ , is decomposed into disjoint subsets for each component  $i = 1, \dots, d$ ,  $\Gamma_{D_i}$  with prescribed displacements  $u_i = 0$ , and  $\Gamma_{N_i}$  with prescribed traction  $t_i$ . The material properties tensor is isotropic and has the following form  $\mathbf{C}(\mathbf{x}) = E(\mathbf{x}) \mathbf{C}_0$ , where  $\mathbf{C}_0$  is a constant isotropic elasticity tensor obtained for a predefined Poisson ratio  $\nu$  smaller than 0.5 and modulus of elasticity one.  $E(\mathbf{x})$  is a spatially varying modulus of elasticity  $E(x) \in [E_{\min}, E_{\max}]$ ; cf. (36). Here, the Poisson’s ratio is kept constant; however, if necessary, it can also be modeled as a random field.

The variational formulation of (50) is given as

$$a(\mathbf{u}, \mathbf{v}) = \ell(\mathbf{v}) \quad \text{for all } \mathbf{v} \in V_0,
 \tag{51a}$$

with  $V_0 = \{ \mathbf{v} \in [H^1(D)]^d : v_i = 0 \text{ on } \Gamma_{D_i}, i = 1, \dots, d \} \subset V = [H^1(D)]^d$  and bilinear form  $a$  and linear form  $\ell$  defined by

$$a(\mathbf{u}, \mathbf{v}) = \int_D (\mathbf{C} : \boldsymbol{\varepsilon}(\mathbf{u})) : \boldsymbol{\varepsilon}(\mathbf{v}) \, dx,
 \tag{51b}$$

$$\ell(\mathbf{v}) = \int_D (\mathbf{f} \cdot \mathbf{v}) \, dx + \int_{\Gamma_N} (\mathbf{t} \cdot \mathbf{v}) \, ds.
 \tag{51c}$$

#### 5.3.1. Two-dimensional bridge design

The first set of examples is inspired by [110,124] and consists of 2D bridge design problems. The geometry of the 2D problem is shown in Fig. 15. We choose the design domain to be the unit square. Two distributed loads are applied on top of the design domain: a uniform vertical load  $f_v = 1$  and a spatially varying random horizontal distributed force  $f_h$ . The vertical and horizontal displacements at the bottom of the design domain are set to zero. We use the so-called robust formulation [98,99] discussed earlier to obtain black-and-white designs with a prescribed length scale. We optimize the mean compliance for the most eroded design obtained with  $\eta_e = 0.7$ , and the volume constraint is enforced on the most dilated design with  $\eta_d = 0.3$ . The maximum amount of solid material is limited to 30% of the total area of the design domain. The projection parameter in (37) is set to  $\beta = 8$ , and the length parameter in the PDE filter (35) is set to  $r = 0.02/\sqrt{12}$ . The interested readers are referred to [102] for more details on the length scale relation. For the PDE filter, we enforce the Dirichlet boundary condition  $\bar{\rho} = 1$  on the bridge deck, i.e., on top of the design domain,  $\bar{\rho} = 0$  on the vertical sides, and zero Neumann on the remaining boundary. The domain is discretized using second-order quad elements, with approximately 1M DOFs representing the displacement field and 0.5M DOFs representing the filtered density. The density field is optimized using the PETSc implementation [125] of the Method of Moving Asymptotes (MMA) [116], and the total number of iterations is set to 350. The optimization uses a fixed number of 400 realizations for  $f_h$ . The realizations do not change between the optimization iterations, resulting in a deterministic optimization problem similar to the one for thermal compliance. Symmetrized designs are obtained as in Section 5.2.3.

The results from the topology optimization for different initial density distributions and correlation lengths for the random spatially varying horizontal load  $f_h$  are presented in Figs. 16 and 17. The min–max optimization problem (38) is not convex, and the iterative optimization process will likely end up in a local minimum. Furthermore, the final topology is influenced significantly by the initial density distribution, which can be observed by comparing Figs. 16 and 17. The designs in the former figure are obtained from an initial density field proportional to a cosine function with a period of 1/4, and the designs presented in the latter are obtained from uniform initial density distribution. The non-uniform initial density distribution supports the appearance of five legs for transferring the vertical load to the support. On the other hand, the uniform initial guess results in two inclined main legs that branch to support the bridge deck.

As mentioned earlier, the optimization problem is not convex, and even though we have symmetry in the design domain with respect to the dashed line in the middle, we cannot expect a symmetric solution as the optimization process will likely end up in a local minimum. Furthermore, during the optimization, we use a limited number of realizations for the horizontal load  $f_h$  without guaranteeing the symmetry of the realized distribution concerning the dashed axis. Therefore, for optimization runs executed on the entire optimization domain, we clearly observe a lack of symmetry in the obtained topologies in Figs. 16 (f)–(j) and 17 (f)–(j). Since we enforce  $\bar{\rho} = 1$  for the boundary condition of the PDE filter on the deck, the solid material locally balances the negative and positive oscillations of the load, and only the unbalanced part of the horizontal force is transferred to the supports, forcing the topology optimization algorithm to develop diagonal support members.

The optimized topologies are significantly affected by the correlation length of the horizontal stochastic load  $f_h$ . Shorter correlation lengths result in a fast oscillatory load applied on the bridge’s deck. In these cases, the horizontal forces are balanced locally for the non-uniform initial distribution experiments, resulting in topologies without diagonal members required for stiffening the structure against global horizontal loads. The topologies obtained for correlation length parameter  $l = 0.01$  do not differ




Fig. 15. 2D bridge — design domain and boundary conditions. The vertical and horizontal loads are denoted by  $f_v$  and  $f_h$ , respectively.

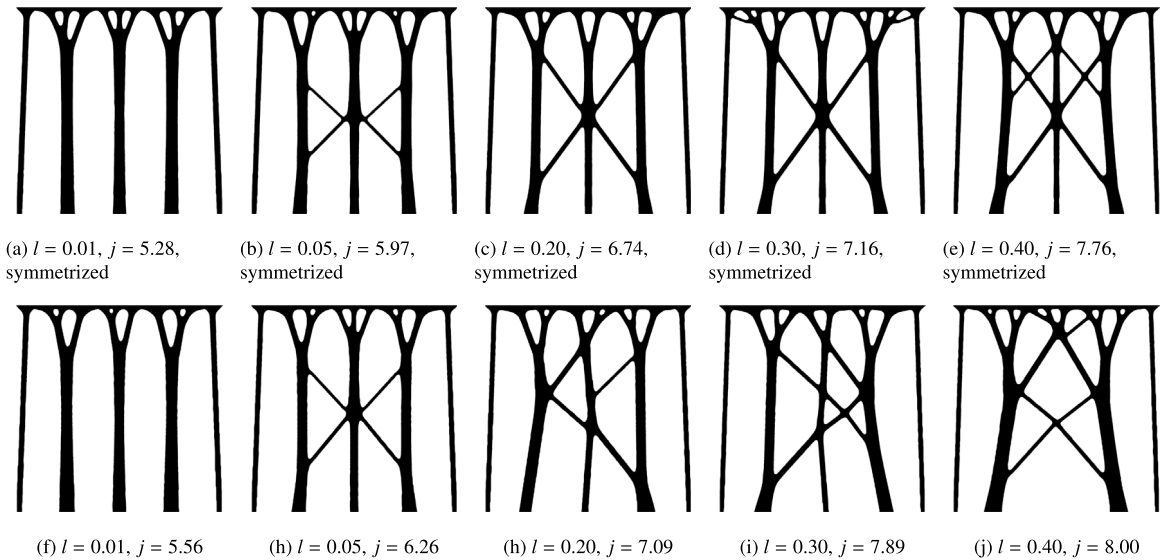


Fig. 16. Compliance-optimized design of a bridge structure with uniformly distributed unit vertical load and random horizontal Gaussian noise with a standard deviation 0.2 and different correlation lengths  $l$ . The correlation length  $l$  increases from left to right. The top (a–e) and bottom (f–j) rows display the results with and without enforced symmetry, respectively. The initial distribution, a scaled cosine wave, aids the appearance of five supporting legs in all designs.

visually from those obtained for the pure vertical load without horizontal disturbance. Therefore, to avoid redundancy, designs for deterministic cases with  $f_h = 0$  are omitted in Figs. 16 and 17.

On the other hand, for longer correlation lengths, the ergodic average along the bridge deck significantly differs from zero for most realizations, forcing the optimization to develop diagonal supporting structural members for transferring the unbalanced forces to the bottom of the design domain. For a uniform initial density distribution, the deck loads are transferred by only two legs to the bottom of the design domain. In this case, the optimization process develops small loops close to the deck for shorter correlation lengths versus larger loops for longer lengths. Like the multi-leg case above, the structural loops balance the forces as close as possible to the bridge deck, and the main two legs transfer only the excess, unbalanced load to the supports.

### 5.3.2. Three-dimensional bridge design

The final solid mechanics example demonstrates topology optimization of a 3D bridge structure with designs shown in Figs. 18, 19, 20. The bridge is loaded with a unit vertical distributed load. The optimization is performed only on a quarter of the domain shown in the figures. Periodic boundary conditions are enforced along vertical planes perpendicular to the deck axis and a symmetric boundary condition is placed on a vertical plane aligned with the long deck axis. The bridge is supported at the bottom of the domain. The computational domain  $D = (0, 1) \times (0, 0.5) \times (0, 1)$  is discretized with approximately 20M elements, resulting

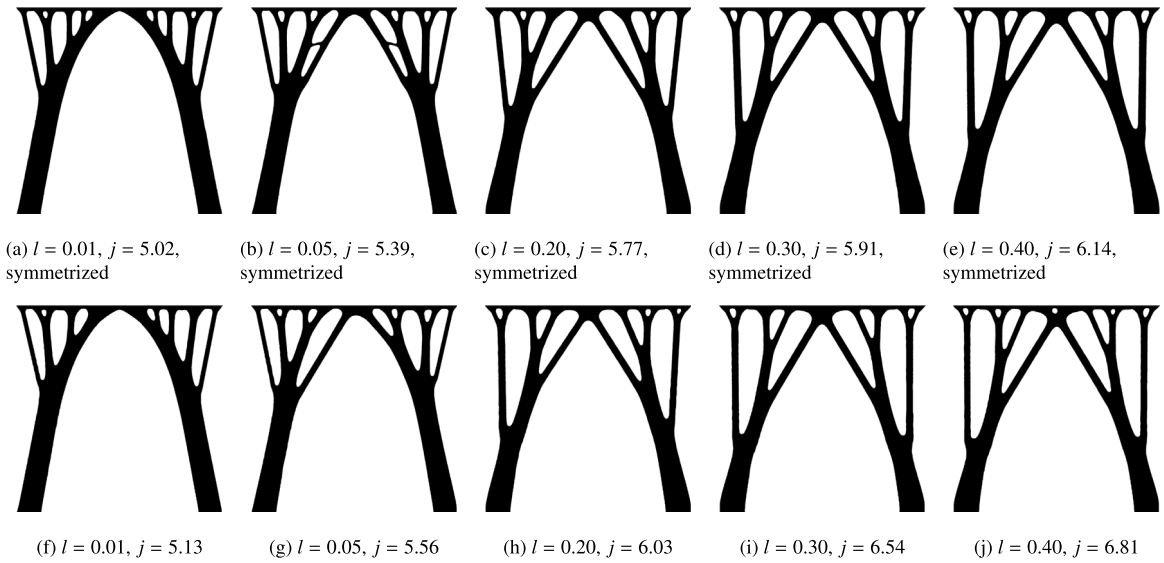


Fig. 17. Compliance-optimized design of a bridge structure with uniformly distributed unit vertical load and random horizontal Gaussian noise with a standard deviation 0.2 and different correlation lengths  $l$ . The correlation length  $l$  increases from left to right. The top (a-e) and bottom (f-j) rows display the results with and without enforced symmetry, respectively. The initial density distribution is uniform.

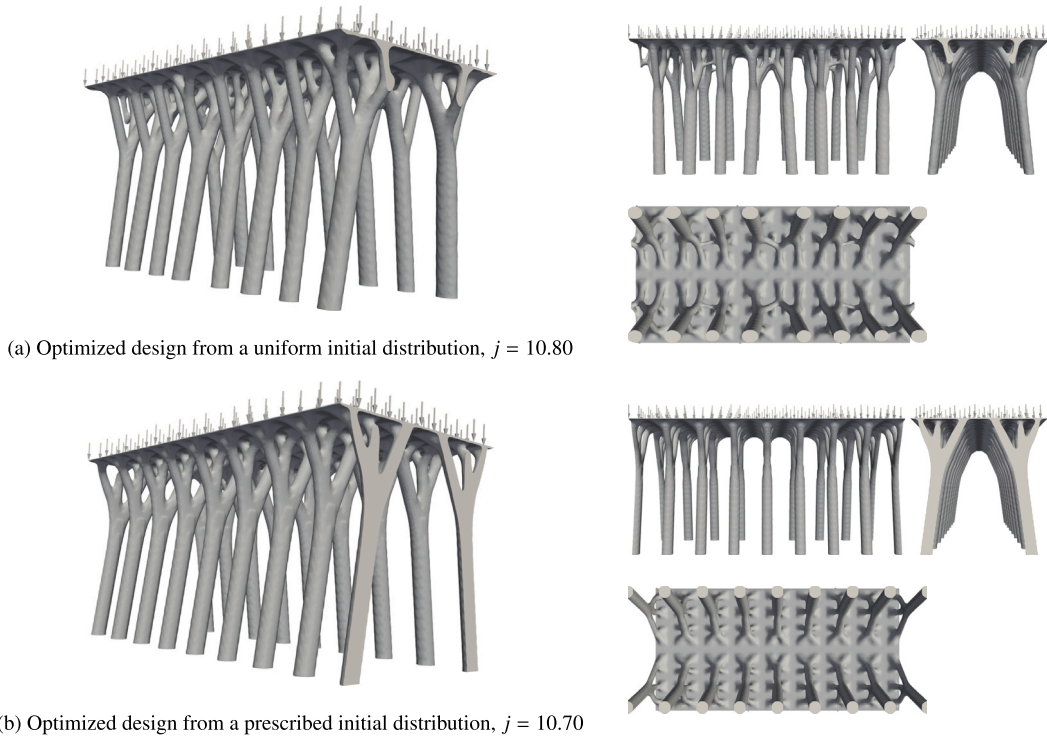


Fig. 18. Deterministic compliance optimization of a 3D bridge structure with uniform (top) and prescribed (bottom) initial density distribution.

in roughly 60M DOFs for the linear elasticity solver. We employ the robust formulation [99] with thresholds of  $\eta_d = 0.3$  and  $\eta_e = 0.7$  to enforce a strict length scale on the design. The filter parameter  $r$  is set to 0.04, and the volume constraint is set to 7% of the design volume.

The first two designs shown in Fig. 18 are executed with deterministic vertical load and demonstrate the impact of the initial density distribution on the final design. A predefined initial distribution affects the number of columns supporting the bridge and their locations. Similar to the 2D case, we add a Matérn-type random field, generated via the SPDE method, to represent the force

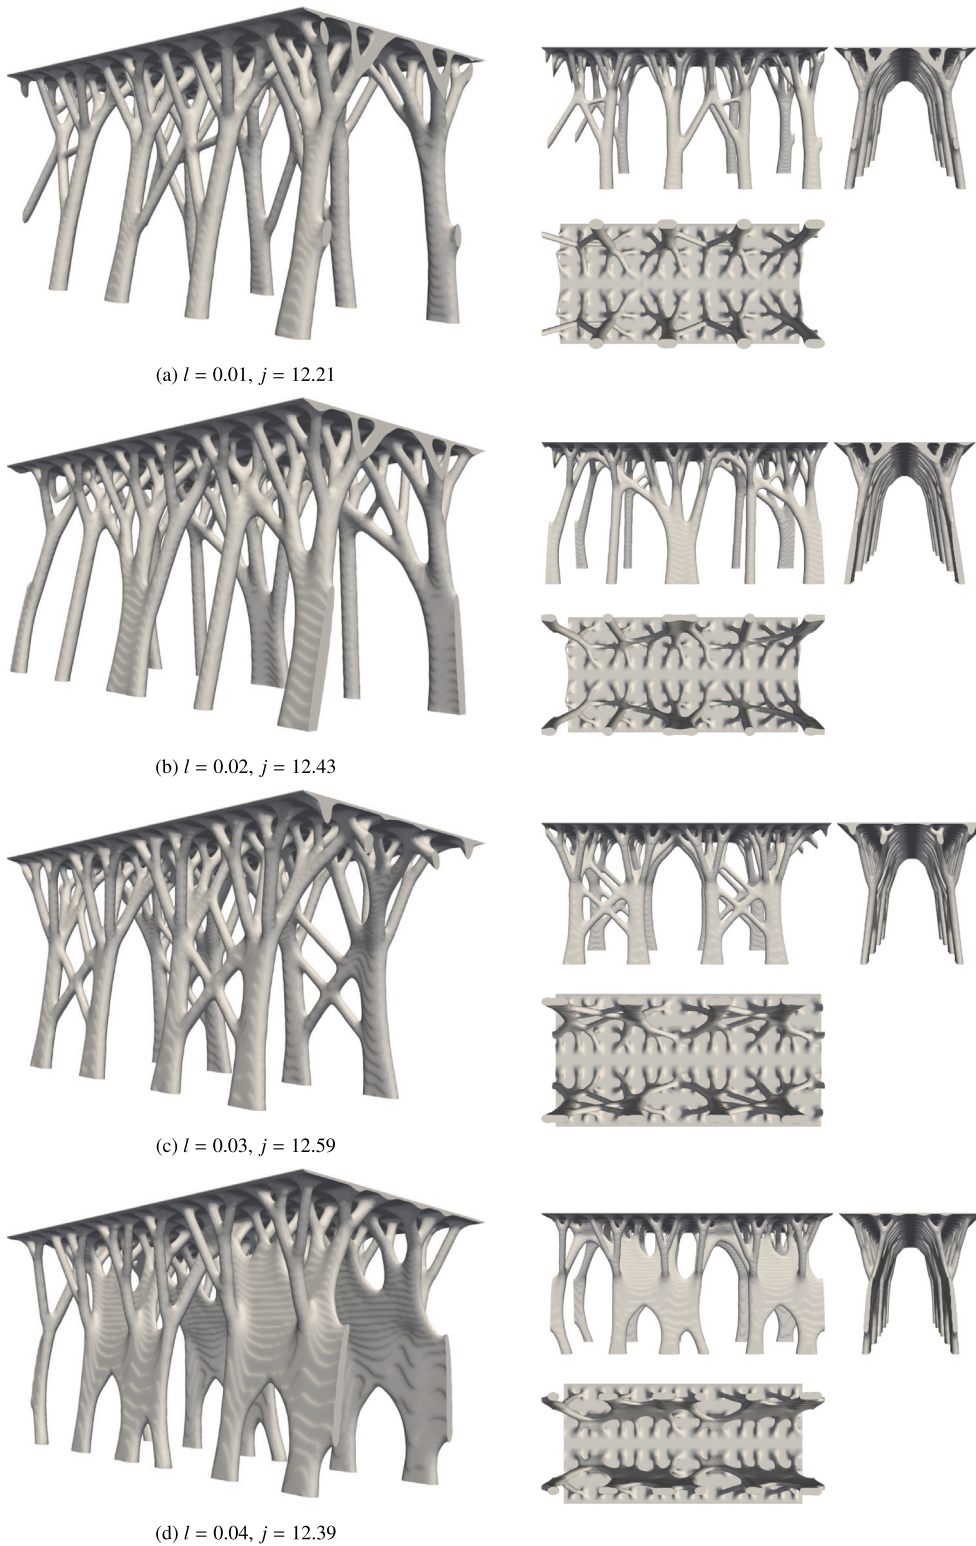


Fig. 19. Stochastic compliance optimization of a 3D bridge structure from a uniform initial density distribution.






Fig. 20. Stochastic compliance optimization of a 3D bridge structure from a prescribed initial density distribution.

on the bridge aligned with the deck axis. The experiments were executed with different correlation lengths for the random field and uniform and prescribed initial distributions. For small correlation lengths, the optimization tries to balance the forces close to the deck of the bridge by using inclined support legs connected to a well-defined column, which transfers the load to the support. Increasing the correlation length results in the development of diagonal elements and, for very long correlation lengths, shear walls. Since second-order effects, i.e., loss of stability, are not accounted for, the most effective way to transfer the excess horizontal force to the support would be through thin shear walls. However, the enforced length scale does not allow the formation of such thin features, and the optimization concentrates the material in column-like elements. For the nonuniform initial design density and short correlation lengths, the optimization process develops a single shear wall along the axis of the deck. The columns close to the two ends of the deck transfer only the vertical load to the support and do not affect the horizontal stiffness of the bridge. In

contrast to the 2D case, the unbalanced part of the nonuniform horizontal random excitation introduces an equivalent torsion force, which requires moving the vertical supports close to the two ends of the deck. This effect is observed for the long correlation length case shown in Fig. 20(c). We emphasize that the presented examples aim to demonstrate the scalability and versatility of the SPDE approach, and more realistic, physically-relevant bridge loads can be obtained using pointwise, memoryless GRF transformations or non-Gaussian random fields. The SPDE approach and the current MFEM implementation can be utilized to generate non-Gaussian random fields with Matérn-type covariances by replacing the Gaussian white noise with a non-Gaussian noise [126,127]. Discussions covering further generalizations to spatio-temporal processes can be found in [12].

## 6. Conclusions

This work aimed at incorporating realistic uncertainties into computational mechanics and design optimization workflows using Matérn-type Gaussian random fields generated via the SPDE method. We focused on aleatoric uncertainties arising from environmental influences, varying material properties, and geometric uncertainties. We consolidated an efficient numerical scheme, previously scattered in the literature, and provided an open-source implementation, which is now part of the widely used MFEM library. We further identified crucial requirements for effective random field generation within the modern scientific computing landscape. These requirements encompass the ability to handle diverse correlation lengths and field smoothness, address intricate domains and manifolds, align with finite element data structures, and computational costs on the order of the physical problem subject to uncertainties.

We find that the SPDE method effectively satisfies all the identified requirements. While alternative methods like FFT and EOLE offer advantages in specific niche areas, the SPDE method's overall versatility and scalability outweigh these isolated shortcomings. Its ability to extend naturally to manifolds enables uncertainty quantification of complex geometries and paves the way for design optimization under uncertainty on manifolds embedded in 3D space. Moreover, its integration into design problems bears the potential to discover novel topological features previously undetectable with established random field discretization methods. We demonstrated the method's capabilities through an extensive set of examples including the construction microstructures mimicking fiber materials, handling geometric uncertainties of cerebral aneurysms, and optimizing topologies for thermal and structural design, even on 2D surfaces embedded in 3D.

While the SPDE method demonstrates remarkable capabilities, it has limitations. The numerical scheme can become expensive for specific parameter regimes, particularly when the smoothness parameter  $\nu$  approaches infinity, leading to a long sequence of equations in (12). Solving the linear system associated with strongly anisotropic fields can also incur significant computational costs. Furthermore, the current implementation and theory are confined to Matérn-type covariances. Specific use cases might encounter covariance structures diverging from the Matérn kernel (1) or might even necessitate non-Gaussian fields. Various (local) transformations can be applied to construct fields with different statistical properties and may alleviate some of these issues. However, whether or not the Matérn covariance defines an appropriate spatial structure for a given problem must be assessed on a case-by-case basis. Ideally, such a decision should be supported by experimental data, aiding the identification of the covariance structure. While the general behavior of the SPDE covariance is well understood, obtaining a closed-form expression for (27) on embedded manifolds and other complex domains often proves challenging. However, even in such cases, the covariance (27) is typically well-approximated by (1) given a correlation length smaller than the characteristic features of the domain.

Our observations hold significant implications for a broad audience utilizing finite elements, including researchers in computational mechanics, design optimization, biomechanics, and inverse problems. The SPDE method simplifies incorporating spatial uncertainties into diverse scientific computing workflows. Additionally, the open-source implementation allows researchers to leverage this powerful tool with minimal effort within existing MFEM projects. These advancements are expected to drive the widespread adoption of the SPDE method.

Looking ahead, several promising avenues for future research present themselves. Extending the SPDE method beyond Matérn-type covariances would broaden its applicability to encompass a broader range of real-world uncertainty scenarios. Additionally, exploring the potential of neural operators for accelerating computations holds promise in enhancing the method's efficiency. Furthermore, developing dedicated numerical schemes tailored to simultaneously solving the equations arising from the rational approximation in (26) can significantly improve the method's computational performance. Collectively, these advancements aim to further refine and enhance the SPDE method for various applications across scientific computing domains.

In summary, this work establishes the SPDE method as a powerful and versatile tool for describing aleatoric uncertainties in modern scientific computing, ultimately leading to substantial advancements in uncertainty quantification, computational mechanics, and design optimization.

## CRedit authorship contribution statement

**Tobias Duswald:** Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Conceptualization. **Brendan Keith:** Writing – review & editing, Writing – original draft, Validation, Software, Methodology, Funding acquisition, Conceptualization. **Boyan Lazarov:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Conceptualization. **Socratis Petrides:** Writing – review & editing, Visualization, Validation, Software, Investigation. **Barbara Wohlmuth:** Writing – review & editing, Validation, Supervision, Methodology, Conceptualization.

## Declaration of competing interest

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

## Data availability

Our implementation of the SPDE solver is freely available as part of the open-source project MFEM (version 4.6)

## Acknowledgments

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344, the LLNL-LDRD Program under Project tracking No. 22-ERD-009, and Differentiating Large-Scale Finite Element Applications project supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research. Release number LLNL-JRNL-861284. This material is based upon work supported by the U.S. Department of Energy Office of Science, Early Career Research Program under Award Number DE-SC0024335. TD was sponsored by LLNL's Computing Scholar Program and the Wolfgang Gentner Programme of the German Federal Ministry of Education and Research (grant no. 13E18CHA). BW gratefully acknowledges the financial support provided by the German Science Foundation (DFG) under project number 465242983 within the priority programme *SPP 2311: Robust coupling of continuum-biomechanical in silico models to establish active biological system models for later use in clinical applications — Co-design of modeling, numerics and usability* (WO 671/20-1) — and the Federal Ministry of Education and Research (BMBF) as part of the *Multi-physics simulations for Geodynamics on heterogeneous Exascale Systems (CoMPS) project* (FKZ 16ME0651) inside the federal research program on *High-Performance and Supercomputing for the Digital Age 2021–2024 — Research and Investments in High-Performance Computing*. TD and BW would like to thank Markus Muhr, Stephan Lunowa, and Fabian Holzberger for insightful discussions regarding the cerebral aneurysm examples.

## Appendix A. Implementation and scalability

We implemented the SPDE method with MFEM [23,24], and the resulting MPI-parallel solver is part of MFEM's 4.6 release.<sup>1</sup> While the numerics are non-trivial, using the solver in MFEM-based projects only requires adding three lines of C++ code, see lines 11 to 13 in Algorithm 1. The statement in line 11 initializes the solver and assembles the mass and stiffness matrices in (14). Line 12 sets the seed for all MPI ranks. Lastly, line 13 assembles the white noise, constructs the linear systems arising from the rational approximation (26), and solves them with a preconditioned conjugate gradient algorithm [128,129].

---

**Algorithm 1** Application example of the MFEM based SPDE-solver (assumes pre-existing finite element space object "fespace"; see MFEM documentation for details).

---

```

1: ▷ Define the fractional PDE solution <
2: ParGridFunction u(&fespace);
3: u = 0.0;
4: ▷ Define GRF parameter <
5: const double nu = 2.0;
6: const double lengthscale = 0.05;
7: const int seed = 1;
8: ▷ Define the boundary conditions (defaults to Neumann). <
9: spde::Boundary bc;
10: ▷ Solve the SPDE problem <
11: spde::SPDESolver solver(nu, bc, &fespace, lengthscale, lengthscale, lengthscale);
12: solver.SetupRandomFieldGenerator(seed);
13: solver.GenerateRandomField(u);

```

---

We argue that the implementation inherits its scalability from MFEM [23,24] and *hypre* [130] because these libraries govern the computationally expensive aspects of the algorithm, i.e., the finite element assembly and solving the linear systems. Detailed scalability studies have been performed and presented by Fischer and co-workers [131]. They considered solving the equation  $(-\nabla \cdot \mu \nabla + \beta)u = f$  with MFEM and other numerical libraries via a diagonally preconditioned conjugate gradient method and measured the performance using (up to) 16,384 MPI ranks distributed among 512 nodes. Their results demonstrate MFEM's ability to scale on top-tier supercomputing systems for equations similar to (9). Additional performance analyses on GPU architectures may be found in [24]. Unlike [131], we use *hypre*'s BoomerAMG as a preconditioner. Henson and Yang [132] investigated and demonstrated the preconditioner's effectiveness and parallel scalability for a set of toy problems with different characteristics. These

<sup>1</sup> MFEM 4.6, [mfem.org/download/](https://mfem.org/download/), available in `miniapps/spde`

problems included the Laplace equation and related problems with anisotropy and the absence of symmetry. They found that the preconditioner works well across the set of example problems and scales to tens of millions of unknowns and more than a thousand processors.

During the development of our solver and the examples presented in Section 5, the solver routinely ran on hundreds of cores at Lawrence Livermore National Laboratory.

## Appendix B. Sample size choice

We optimize the topology with the Method of Moving Asymptotes [116] using a fixed set of samples. For the heat sink design in Section 5.2, we used a sample size of  $N = 300$ . For the bridge design in Section 5.3, we chose  $N = 400$ . To determine  $N$  in a problem-specific setup, we carefully investigate the effect of the sample size on the objective function and optimized design. Here, we show our numerical experiments justifying this choice for the heat sink optimization. Analogous experiments were carried out for the bridge examples.

The numerical experiments proceed akin to Section 5.2.2 with the notable difference that we use different sample sizes  $N \in \{50, 100, 200, 300, 400\}$ . We choose the correlation lengths  $l = 0.2$  and  $l = 0.05$  for our numerical experiments, representing the long and short correlation lengths considered in this work. We expect that larger correlation lengths require fewer samples than short correlation lengths. Fig. B.21 illustrates the result of this experiment, confirming this hypothesis. Comparing the cases  $N = 50$  and  $N = 400$ , we find that the longer correlation length yields more similar designs than the short correlation length. Overall, increasing the sample size increases the symmetry, i.e., the main branches resulting from the optimization become more evenly balanced. Regarding the objective function  $j$ , we find that its value converges once we pass 300 samples. This is also reflected in the design, which is very similar for the cases  $N = 300$  and  $N = 400$ . We conclude that computationally feasible increases beyond the sample size of  $N = 300$  barely affect the design and limit ourselves to 300 samples for each heat sink optimization experiment.

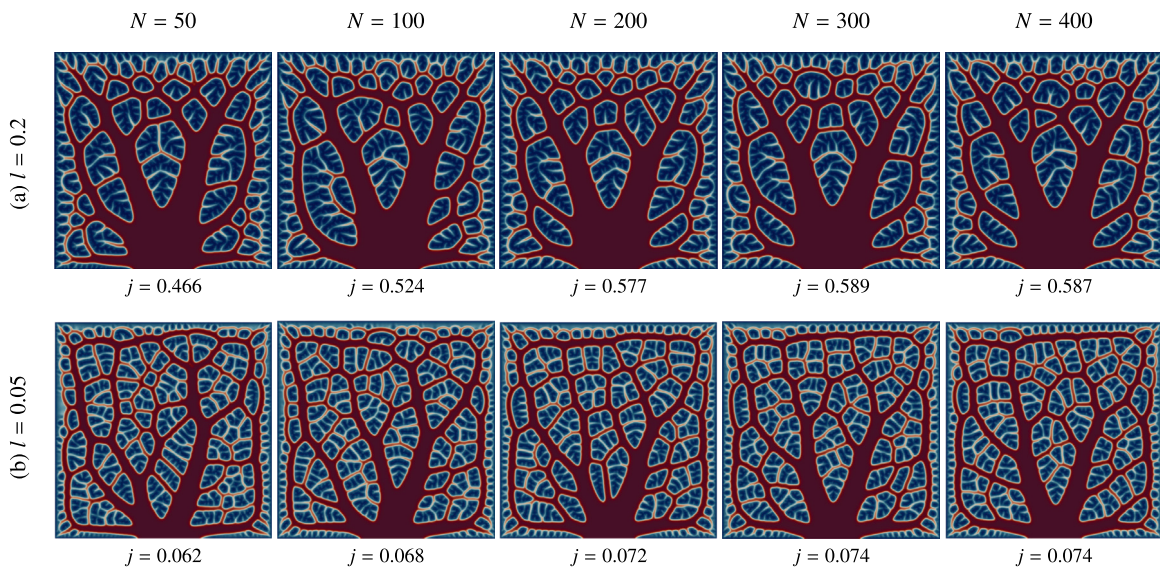


Fig. B.21. Effect of the sample size on the density  $\bar{\rho}$  and objective function  $j$ . The sample size increases from left to right.

## References

- [1] J. Mann, The spatial structure of neutral atmospheric surface-layer turbulence, *J. Fluid Mech.* 273 (1994) 141–168, <http://dx.doi.org/10.1017/S0022112094001886>.
- [2] J. Mann, Wind field simulation, *Probab. Eng. Mech.* 13 (1998) 269–282, [http://dx.doi.org/10.1016/S0266-8920\(97\)00036-2](http://dx.doi.org/10.1016/S0266-8920(97)00036-2).
- [3] B. Keith, U. Khristenko, B. Wohlmuth, A fractional PDE model for turbulent velocity fields near solid walls, *J. Fluid Mech.* 916 (2021) A21, <http://dx.doi.org/10.1017/jfm.2021.182>.
- [4] B. Keith, U. Khristenko, B. Wohlmuth, Learning the structure of wind: A data-driven nonlocal turbulence model for the atmospheric boundary layer, *Phys. Fluids* 33 (2021) <http://dx.doi.org/10.1063/5.0064394>.
- [5] A. Kodakkal, B. Keith, U. Khristenko, A. Apostolatos, K.-U. Bletzinger, B. Wohlmuth, R. Wüchner, Risk-averse design of tall buildings for uncertain wind conditions, *Comput. Methods Appl. Mech. Engrg.* 402 (2022) 115371, <http://dx.doi.org/10.1016/j.cma.2022.115371>.
- [6] A. Nybø, F. Nielsen, M. Godvik, Analysis of turbulence models fitted to site, and their impact on the response of a bottom-fixed wind turbine, in: *Journal of Physics: Conference Series*, vol. 2018, 2021, 012028, <http://dx.doi.org/10.1088/1742-6596/2018/1/012028>.
- [7] A. Nybø, F.G. Nielsen, J. Reuder, M.J. Churchfield, M. Godvik, Evaluation of different wind fields for the investigation of the dynamic response of offshore wind turbines, *Wind Energy* 23 (2020) 1810–1830, <http://dx.doi.org/10.1002/we.2518>.
- [8] S. Torquato, *Random Heterogeneous Materials*, vol. 16, Springer New York, New York, NY, 2002, <http://dx.doi.org/10.1007/978-1-4757-6355-3>.
- [9] A.E. Robertson, S.R. Kalidindi, Efficient generation of anisotropic N-field microstructures from 2-point statistics using multi-output Gaussian random fields, *Acta Mater.* 232 (2022) 117927, <http://dx.doi.org/10.1016/j.actamat.2022.117927>.

- [10] U. Khristenko, A. Constantinescu, P.L. Tallec, B. Wohlmuth, Statistically equivalent surrogate material models: Impact of random imperfections on the elasto-plastic response, *Comput. Methods Appl. Mech. Engrg.* 402 (2022) 115278, <http://dx.doi.org/10.1016/j.cma.2022.115278>.
- [11] A.D. Kiureghian, O. Ditlevsen, Aleatory or epistemic? Does it matter? *Struct. Saf.* 31 (2) (2009) 105–112, <http://dx.doi.org/10.1016/j.strusafe.2008.06.020>.
- [12] F. Lindgren, D. Bolin, H. Rue, The SPDE approach for Gaussian and non-Gaussian fields: 10 years and still running, *Spatial Stat.* 50 (2022) 100599, <http://dx.doi.org/10.1016/j.spasta.2022.100599>.
- [13] F. Lindgren, H. Rue, J. Lindström, An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic partial differential equation approach, *J. R. Stat. Soc. Ser. B Stat. Methodol.* 73 (2011) 423–498, <http://dx.doi.org/10.1111/j.1467-9868.2011.00777.x>.
- [14] P. Whittle, On stationary processes in the plane, *Biometrika* 41 (1954) 434, <http://dx.doi.org/10.2307/2332724>.
- [15] P. Whittle, *Stochastic processes in several dimensions*, *Bull. Inst. Intern. Stat.* 40 (1963) 974–994.
- [16] K.J. Koh, F. Cirak, Stochastic PDE representation of random fields for large-scale Gaussian process regression and statistical finite element analysis, *Comput. Methods Appl. Mech. Engrg.* 417 (2023) 116358, <http://dx.doi.org/10.1016/j.cma.2023.116358>.
- [17] J. Guillemot, A. Asadpoure, M. Tootkaboni, Topology optimization under topologically dependent material uncertainties, *Struct. Multidiscip. Optim.* 60 (2019) 1283–1287, <http://dx.doi.org/10.1007/s00158-019-02247-1>.
- [18] I. Ben-Yelun, A.O. Yuksel, F. Cirak, Robust topology optimisation of lattice structures with spatially correlated uncertainties, *Struct. Multidiscip. Optim.* 67 (2024) 16, <http://dx.doi.org/10.1007/s00158-023-03716-4>.
- [19] Y. Nakatsukasa, O. Sète, L.N. Trefethen, The AAA algorithm for rational approximation, *SIAM J. Sci. Comput.* 40 (2018) A1494–A1522, <http://dx.doi.org/10.1137/16M1106122>.
- [20] D. Bolin, K. Kirchner, The rational SPDE approach for Gaussian random fields with general smoothness, *J. Comput. Graph. Statist.* 29 (2020) 274–285, <http://dx.doi.org/10.1080/10618600.2019.1665537>.
- [21] S. Harizanov, S. Margenov, Positive approximations of the inverse of fractional powers of SPD M-matrices, in: G. Feichtinger, R.M. Kovacevic, G. Tragler (Eds.), *Control Systems and Mathematical Methods in Economics: Essays in Honor of Vladimir M. Veliov*, Springer International Publishing, Cham, 2018, pp. 147–163, [http://dx.doi.org/10.1007/978-3-319-75169-6\\_8](http://dx.doi.org/10.1007/978-3-319-75169-6_8).
- [22] M. Croci, M.B. Giles, M.E. Rognes, P.E. Farrell, Efficient white noise sampling and coupling for multilevel Monte Carlo with nonnested meshes, *SIAM/ASA J. Uncertain. Quantif.* 6 (2018) 1630–1655, <http://dx.doi.org/10.1137/18M1175239>.
- [23] R. Anderson, J. Andrej, A. Barker, J. Bramwell, J.-S. Camier, J. Cervený, V. Dobrev, Y. Doudout, A. Fisher, T. Kolev, W. Pazner, M. Stowell, V. Tomov, I. Akkerman, J. Dahm, D. Medina, S. Zampini, MFEM: A modular finite element methods library, *Comput. Math. Appl.* 81 (2021) <http://dx.doi.org/10.1016/j.camwa.2020.06.009>.
- [24] J. Andrej, N. Atallah, J.-P. Bäcker, J. Camier, D. Copeland, V. Dobrev, Y. Doudout, T. Duswald, B. Keith, D. Kim, T. Kolev, B. Lazarov, K. Mittal, W. Pazner, S. Petrides, S. Shiraiwa, M. Stowell, V. Tomov, High-performance finite elements with MFEM, 2024, [arXiv:2402.15940](https://arxiv.org/abs/2402.15940).
- [25] M.W. Davis, Production of conditional simulations via the LU triangular decomposition of the covariance matrix, *Math. Geol.* 19 (1987) 91–98, <http://dx.doi.org/10.1007/BF00898189>.
- [26] C.R. Dietrich, G.N. Newsam, Efficient generation of conditional simulations by Chebyshev matrix polynomial approximations to the symmetric square root of the covariance matrix, *Math. Geol.* 27 (1995) 207–228, <http://dx.doi.org/10.1007/BF02083211>.
- [27] P. Blanchard, O. Coulaud, E. Darve, *Fast Hierarchical Algorithms for Generating Gaussian Random Fields*, Inria Bordeaux Sud-Ouest, 2015.
- [28] M. Feischl, F.Y. Kuo, I.H. Sloan, Fast random field generation with H-matrices, *Numer. Math.* 140 (2018) 639–676, <http://dx.doi.org/10.1007/s00211-018-0974-2>.
- [29] J.M. Mejía, I. Rodríguez-Iturbe, On the synthesis of random field sampling from the spectrum: An application to the generation of hydrologic spatial processes, *Water Resour. Res.* 10 (1974) 705–711, <http://dx.doi.org/10.1029/WR010i004p00705>.
- [30] M. Shinozuka, Simulation of multivariate and multidimensional random processes, *J. Acoust. Soc. Am.* 49 (1971) 357–368, <http://dx.doi.org/10.1121/1.1912338>.
- [31] M. Shinozuka, C.-M. Jan, Digital simulation of random processes and its applications, *J. Sound Vib.* 25 (1972) 111–128, [http://dx.doi.org/10.1016/0022-460X\(72\)90600-1](http://dx.doi.org/10.1016/0022-460X(72)90600-1).
- [32] A. Gutjahr, B. Bullard, S. Hatch, General joint conditional simulations using a fast Fourier transform method, *Math. Geol.* 29 (1997) 361–389, <http://dx.doi.org/10.1007/BF02769641>.
- [33] F. Ruan, D. McLaughlin, An efficient multivariate random field generator using the fast Fourier transform, *Adv. Water Resour.* 21 (1998) 385–399, [http://dx.doi.org/10.1016/S0309-1708\(96\)00064-4](http://dx.doi.org/10.1016/S0309-1708(96)00064-4).
- [34] A. Lang, J. Potthoff, Fast simulation of Gaussian random fields, *Monte Carlo Methods Appl.* 17 (2011) <http://dx.doi.org/10.1515/mcma.2011.009>.
- [35] M.L. Ravalec, B. Noetinger, L.Y. Hu, The FFT moving average (FFT-MA) generator: An efficient numerical method for generating and conditioning Gaussian simulations, *Math. Geol.* 32 (2000) 701–723, <http://dx.doi.org/10.1023/A:1007542406333>.
- [36] M. Loève, *Probability Theory I*, vol. 45, Springer New York, New York, NY, 1977, <http://dx.doi.org/10.1007/978-1-4684-9464-8>.
- [37] C. Schwab, R.A. Todor, Karhunen–Loève approximation of random fields by generalized fast multipole methods, *J. Comput. Phys.* 217 (2006) 100–122, <http://dx.doi.org/10.1016/j.jcp.2006.01.048>.
- [38] D. Zhang, Z. Lu, An efficient, high-order perturbation approach for flow in random porous media via Karhunen–Loève and polynomial expansions, *J. Comput. Phys.* 194 (2004) 773–794, <http://dx.doi.org/10.1016/j.jcp.2003.09.015>.
- [39] C.R. Dietrich, G.N. Newsam, Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix, *SIAM J. Sci. Comput.* 18 (1997) 1088–1107, <http://dx.doi.org/10.1137/S1064827592240555>.
- [40] I.G. Graham, F.Y. Kuo, D. Nuyens, R. Scheichl, I.H. Sloan, Analysis of circulant embedding methods for sampling stationary random fields, *SIAM J. Numer. Anal.* 56 (2018) 1871–1895, <http://dx.doi.org/10.1137/17M1149730>.
- [41] M.H. Park, M. Tretyakov, A block circulant embedding method for simulation of stationary gaussian random fields on block-regular grids, *Int. J. Uncertain. Quantif.* 5 (2015) 527–544, <http://dx.doi.org/10.1615/Int.J.UncertaintyQuantification.2015013781>.
- [42] Y. Liu, J. Li, S. Sun, B. Yu, Advances in Gaussian random field generation: A review, *Comput. Geosci.* 23 (2019) 1011–1047, <http://dx.doi.org/10.1007/s10596-019-09867-y>.
- [43] C.-C. Li, A.D. Kiureghian, Optimal discretization of random fields, *J. Eng. Mech.* 119 (1993) 1136–1154, [http://dx.doi.org/10.1061/\(ASCE\)0733-9399\(1993\)119:6\(1136\)](http://dx.doi.org/10.1061/(ASCE)0733-9399(1993)119:6(1136)).
- [44] R.G. Ghanem, P.D. Spanos, *Stochastic Finite Elements: A Spectral Approach*, Springer New York, New York, NY, 1991, <http://dx.doi.org/10.1007/978-1-4612-3094-6>.
- [45] O.P.L. Maître, O.M. Knio, *Spectral Methods for Uncertainty Quantification: With Applications to Computational Fluid Dynamics*, Springer Netherlands, Dordrecht, 2010, <http://dx.doi.org/10.1007/978-90-481-3520-2>.
- [46] A.G. Journel, Geostatistics for conditional simulation of ore bodies, *Econ. Geol.* 69 (1974) 673–687, <http://dx.doi.org/10.2113/gsecongeo.69.5.673>.
- [47] D.S. Oliver, Moving averages for Gaussian simulation in two and three dimensions, *Math. Geol.* 27 (1995) 939–960, <http://dx.doi.org/10.1007/BF02091660>.
- [48] N. Wiener, Generalized harmonic analysis, *Acta Math.* 55 (1930) 117–258, <http://dx.doi.org/10.1007/BF02546511>.
- [49] A. Khintchine, Korrelationstheorie der stationären stochastischen Prozesse, *Math. Ann.* 109 (1934) 604–615, <http://dx.doi.org/10.1007/BF01449156>.
- [50] U. Khristenko, L. Scarabosio, P. Swierczynski, E. Ullmann, B. Wohlmuth, Analysis of boundary effects on PDE-based sampling of Whittle–Matérn random fields, *SIAM/ASA J. Uncertain. Quantif.* 7 (2019) 948–974, <http://dx.doi.org/10.1137/18M1215700>.

- [51] F. Lindgren, H. Bakka, D. Bolin, E. Krainski, H. Rue, A diffusion-based spatio-temporal extension of Gaussian Matérn fields, 2023, [arXiv:2006.04917](https://arxiv.org/abs/2006.04917).
- [52] U. Khristenko, A. Constantinescu, P.L. Tallec, J.T. Oden, B. Wohlmuth, A statistical framework for generating microstructures of two-phase random materials: Application to fatigue analysis, *Multiscale Model. Simul.* 18 (2020) 21–43, [http://dx.doi.org/10.1137/19M1259286](https://doi.org/10.1137/19M1259286).
- [53] M. Schevenels, B. Lazarov, O. Sigmund, Robust topology optimization accounting for spatially varying manufacturing errors, *Comput. Methods Appl. Mech. Engrg.* 200 (49–52) (2011) 3613–3627, [http://dx.doi.org/10.1016/j.cma.2011.08.006](https://doi.org/10.1016/j.cma.2011.08.006).
- [54] B.S. Lazarov, M. Schevenels, O. Sigmund, Topology optimization with geometric uncertainties by perturbation techniques, *Internat. J. Numer. Methods Engrg.* 90 (11) (2012) 1321–1336, [http://dx.doi.org/10.1002/nme.3361](https://doi.org/10.1002/nme.3361).
- [55] J. Martínez-Frutos, F.P. Esparza, *Optimal Control of PDEs under Uncertainty: An Introduction with Application to Optimal Shape Design of Structures*, Springer International Publishing, Cham, 2018, [http://dx.doi.org/10.1007/978-3-319-98210-6](https://doi.org/10.1007/978-3-319-98210-6).
- [56] L. de Carvalho Paludo, V. Bouvier, R. Cotteneau, Scalable parallel scheme for sampling of Gaussian random fields over very large domains, *Internat. J. Numer. Methods Engrg.* 117 (2019) 845–859, [http://dx.doi.org/10.1002/nme.5981](https://doi.org/10.1002/nme.5981).
- [57] A. Ayala, S. Tomov, A. Haidar, J. Dongarra, heFFTe: Highly efficient FFT for exascale, in: V.V. Krzhizhanovskaya, G. Závodszy, M.H. Lees, J.J. Dongarra, P.M.A. Sloot, S. Brissos, J.a. Teixeira (Eds.), *Computational Science - ICCS 2020*, Springer International Publishing, Cham, 2020, pp. 262–275, [http://dx.doi.org/10.1007/978-3-030-50371-0\\_19](https://doi.org/10.1007/978-3-030-50371-0_19).
- [58] T.J. Hughes, *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*, Courier Corporation, 2012.
- [59] T. Hida, H.-H. Kuo, J. Potthoff, L. Streit, *White Noise: An Infinite Dimensional Calculus*, vol. 253, Springer Science & Business Media, 2013.
- [60] T. Duswald, S. Petrides, B. Lazarov, B. Keith, MFEM Miniapp: Generating Gaussian Random Fields via the SPDE Method, 2023, GitHub repository, GitHub, <https://github.com/mfem/mfem/pull/3065>.
- [61] A. Bonito, J.P. Borthagaray, R.H. Nochetto, E. Otárola, A.J. Salgado, Numerical methods for fractional diffusion, *Comput. Vis. Sci.* 19 (2018) 19–46, [http://dx.doi.org/10.1007/s00791-018-0289-y](https://doi.org/10.1007/s00791-018-0289-y).
- [62] A. Lischke, G. Pang, M. Gulian, F. Song, C. Glusa, X. Zheng, Z. Mao, W. Cai, M.M. Meerschaert, M. Ainsworth, G.E. Karniadakis, What is the fractional Laplacian? A comparative review with new results, *J. Comput. Phys.* 404 (2020) 109009, [http://dx.doi.org/10.1016/j.jcp.2019.109009](https://doi.org/10.1016/j.jcp.2019.109009).
- [63] L. Caffarelli, L. Silvestre, An extension problem related to the fractional Laplacian, *Communications in partial differential equations* 32 (8) (2007) 1245–1260, [http://dx.doi.org/10.1080/03605300600987306](https://doi.org/10.1080/03605300600987306).
- [64] D. Meidner, J. Pfeifferer, K. Schürholz, B. Vexler,  $hp$ -Finite elements for fractional diffusion, *SIAM J. Numer. Anal.* 56 (4) (2018) 2345–2374, [http://dx.doi.org/10.1137/17M1135517](https://doi.org/10.1137/17M1135517).
- [65] L. Banjai, J.M. Melenk, R.H. Nochetto, E. Otárola, A.J. Salgado, C. Schwab, Tensor FEM for spectral fractional diffusion, *Found. Comput. Math.* 19 (4) (2019) 901–962, [http://dx.doi.org/10.1007/s12028-018-9402-3](https://doi.org/10.1007/s12028-018-9402-3).
- [66] L. Banjai, J.M. Melenk, C. Schwab, Exponential convergence of  $hp$  FEM for spectral fractional diffusion in polygons, *Numer. Math.* 153 (2023) 1–47, [http://dx.doi.org/10.1007/s00211-022-01329-5](https://doi.org/10.1007/s00211-022-01329-5).
- [67] A.V. Balakrishnan, et al., Fractional powers of closed operators and the semigroups generated by them, *Pacific J. Math.* 10 (2) (1960) 419–437.
- [68] A. Bonito, J. Pasciak, Numerical approximation of fractional powers of elliptic operators, *Math. Comp.* 84 (295) (2015) 2083–2110, [http://dx.doi.org/10.1090/S0025-5718-2015-02937-8](https://doi.org/10.1090/S0025-5718-2015-02937-8).
- [69] A. Bonito, W. Lei, J.E. Pasciak, On sinc quadrature approximations of fractional powers of regularly accretive operators, *J. Numer. Math.* 27 (2) (2019) 57–68, [http://dx.doi.org/10.1515/jnma-2017-0116](https://doi.org/10.1515/jnma-2017-0116).
- [70] A. Bonito, D. Guignard, W. Lei, Numerical approximation of Gaussian random fields on closed surfaces, *Comput. Methods Appl. Math.* (2024) [http://dx.doi.org/10.1515/cmam-2022-0237](https://doi.org/10.1515/cmam-2022-0237).
- [71] P.N. Vabishchevich, Numerically solving an equation for fractional powers of elliptic operators, *J. Comput. Phys.* 282 (2015) 289–302, [http://dx.doi.org/10.1016/j.jcp.2014.11.022](https://doi.org/10.1016/j.jcp.2014.11.022).
- [72] R. Lazarov, P. Vabishchevich, A numerical study of the homogeneous elliptic equation with fractional boundary conditions, *Fract. Calc. Appl. Anal.* 20 (2017) 337–351, [http://dx.doi.org/10.1515/fca-2017-0018](https://doi.org/10.1515/fca-2017-0018).
- [73] U. Khristenko, B. Wohlmuth, Solving time-fractional differential equations via rational approximation, *IMA J. Numer. Anal.* 43 (2023) 1263–1290, [http://dx.doi.org/10.1093/imanum/drac022](https://doi.org/10.1093/imanum/drac022).
- [74] J. Guillemot, Modeling non-Gaussian random fields of material properties in multiscale mechanics of materials, in: Y. Wang, D.L. McDowell (Eds.), *Uncertainty Quantification in Multiscale Materials Modeling*, in: Elsevier Series in Mechanics of Advanced Materials, Woodhead Publishing, 2020, pp. 385–420, [http://dx.doi.org/10.1016/B978-0-08-102941-1.00012-2](https://doi.org/10.1016/B978-0-08-102941-1.00012-2).
- [75] M. Grigoriu, *Stochastic Calculus: Applications in Science and Engineering*, first ed., Birkhäuser, 2002.
- [76] D.S. Grebenkov, B.-T. Nguyen, Geometrical structure of Laplacian eigenfunctions, *SIAM Rev.* 55 (2013) 601–667, [http://dx.doi.org/10.1137/120880173](https://doi.org/10.1137/120880173).
- [77] A. Lang, C. Schwab, Isotropic Gaussian random fields on the sphere: Regularity, fast simulation and stochastic partial differential equations, *Ann. Appl. Probab.* 25 (2015) [http://dx.doi.org/10.1214/14-AAP1067](https://doi.org/10.1214/14-AAP1067).
- [78] B. Levy, Laplace-beltrami eigenfunctions towards an algorithm that "understands" geometry, in: *IEEE International Conference on Shape Modeling and Applications 2006, SMI'06*, IEEE, 2006, p. 13, [http://dx.doi.org/10.1109/SMI.2006.21](https://doi.org/10.1109/SMI.2006.21).
- [79] V. Borovitskiy, A. Terenin, P. Mostowsky, M. Deisenroth, Matérn Gaussian processes on Riemannian manifolds, in: H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, H. Lin (Eds.), in: *Advances in Neural Information Processing Systems*, vol. 33, Curran Associates, Inc., 2020, pp. 12426–12437.
- [80] W.I. Schievink, Genetics of intracranial aneurysms, *Neurosurgery* 40 (1997) [http://dx.doi.org/10.1097/00006123-199704000-00001](https://doi.org/10.1097/00006123-199704000-00001).
- [81] J.L. Brisman, J.K. Song, D.W. Newell, Cerebral aneurysms, *N. Engl. J. Med.* 355 (2006) 928–939, [http://dx.doi.org/10.1056/NEJMra052760](https://doi.org/10.1056/NEJMra052760).
- [82] D.M. Sforza, C.M. Putman, J.R. Cebra, Hemodynamics of cerebral aneurysms, *Annu. Rev. Fluid Mech.* 41 (2009) 91–107, [http://dx.doi.org/10.1146/annurev.fluid.40.111406.102126](https://doi.org/10.1146/annurev.fluid.40.111406.102126).
- [83] F.K. Hui, D. Fiorella, T.J. Masaryk, P.A. Rasmussen, J.E. Dion, A history of detachable coils: 1987–2012, *J. NeuroInterventional Surg.* 6 (2014) 134–138, [http://dx.doi.org/10.1136/neurintsurg-2013-010670](https://doi.org/10.1136/neurintsurg-2013-010670).
- [84] L. Pierot, A.K. Wakhloo, Endovascular treatment of intracranial aneurysms, *Stroke* 44 (2013) 2046–2054, [http://dx.doi.org/10.1161/STROKEAHA.113.000733](https://doi.org/10.1161/STROKEAHA.113.000733).
- [85] J. Zhao, H. Lin, R. Summers, M. Yang, B.G. Cousins, J. Tsui, Current treatment strategies for intracranial aneurysms: An overview, *Angiology* 69 (2018) 17–30, [http://dx.doi.org/10.1177/0003319717700503](https://doi.org/10.1177/0003319717700503).
- [86] M. Frank, F. Holzberger, M. Horvat, J. Kirschke, M. Mayr, M. Muhr, N. Nebulishvili, A. Popp, J. Schwarting, B. Wohlmuth, Numerical simulation of endovascular treatment options for cerebral aneurysms, 2024, [arXiv:2402.00550](https://arxiv.org/abs/2402.00550).
- [87] X. Yang, D. Xia, T. Kin, T. Igarashi, Intra: 3D intracranial aneurysm dataset for deep learning, in: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR, 2020, pp. 2653–2663, [http://dx.doi.org/10.1109/CVPR42600.2020.00273](https://doi.org/10.1109/CVPR42600.2020.00273).
- [88] F. Holzberger, M. Muhr, B. Wohlmuth, A comprehensive numerical approach to coil placement in cerebral aneurysms: Mathematical modeling and in silico occlusion classification, 2024, [arXiv:2402.02798](https://arxiv.org/abs/2402.02798).
- [89] F. Nobile, C. Vergara, An effective fluid-structure interaction formulation for vascular dynamics by generalized robin conditions, *SIAM J. Sci. Comput.* 30 (2008) 731–763, [http://dx.doi.org/10.1137/060678439](https://doi.org/10.1137/060678439).
- [90] W.T. Koiter, On the foundations of the linear theory of thin elastic shells, in: *Proceedings of the Koninklijke Nederlandse Akademie van Wetenschappen*, Vol. B, 1970, pp. 169–195.

- [91] W.T. Koiter, On the non-linear theory of thin elastic shells, in: Proceedings of the Koninklijke Nederlandse Akademie van Wetenschappen, Vol. B, 1996, pp. 1–54.
- [92] D. Cohen-Steiner, J.-M. Morvan, Restricted delaunay triangulations and normal cycle, in: Proceedings of the nineteenth annual symposium on Computational geometry, ACM, New York, NY, USA, 2003, pp. 312–321, <http://dx.doi.org/10.1145/777792.777839>.
- [93] E. Levy, C.J. Koebe, M.B. Horowitz, C.A. Jungreis, G.L. Pride, K. Dutton, A. Kassam, P.D. Purdy, Rupture of intracranial aneurysms during endovascular coiling: Management and outcomes, *Neurosurgery* 49 (2001) <http://dx.doi.org/10.1097/00006123-200110000-00005>.
- [94] S.P. Suter, R. Skalak, The history of Poiseuille's law, *Annu. Rev. Fluid Mech.* 25 (1993) 1–20, <http://dx.doi.org/10.1146/annurev.fl.25.010193.000245>.
- [95] M. Horvat, S.B. Lunowa, D. Sytnyk, B. Wohlmuth, A lattice Boltzmann method for non-Newtonian blood flow in coiled intracranial aneurysms, 2024, [arXiv:2402.10809](https://arxiv.org/abs/2402.10809).
- [96] M.P. Bendsoe, O. Sigmund, *Topology Optimization - Theory, Methods and Applications*, Springer Verlag, Berlin Heidelberg, 2003.
- [97] M.P. Bendsoe, O. Sigmund, Material interpolation schemes in topology optimization, *Arch. Appl. Mech.* 69 (1999) 635–654, <http://dx.doi.org/10.1007/s004190050248>.
- [98] B.S. Lazarov, F. Wang, O. Sigmund, Length scale and manufacturability in density-based topology optimization, *Arch. Appl. Mech.* 86 (1) (2016) 189–218, <http://dx.doi.org/10.1007/s00419-015-1106-4>.
- [99] F. Wang, B. Lazarov, O. Sigmund, On projection methods, convergence and robust formulations in topology optimization, *Struct. Multidiscip. Optim.* 43 (6) (2011) 767–784, <http://dx.doi.org/10.1007/s00158-010-0602-y>.
- [100] B.S. Lazarov, F. Wang, Maximum length scale in density based topology optimization, *Comput. Methods Appl. Mech. Engrg.* 318 (2017) 826–844, <http://dx.doi.org/10.1016/j.cma.2017.02.018>.
- [101] B. Bourdin, Filters in topology optimization, *Internat. J. Numer. Methods Engrg.* 50 (2001) 2143–2158, <http://dx.doi.org/10.1002/nme.116>.
- [102] B.S. Lazarov, O. Sigmund, Filters in topology optimization based on Helmholtz-type differential equations, *Internat. J. Numer. Methods Engrg.* 86 (6) (2011) 765–781, <http://dx.doi.org/10.1002/nme.3072>.
- [103] M.P. Bendsoe, Optimal shape design as a material distribution problem, *Struct. Optim.* 1 (4) (1989) 193–202, <http://dx.doi.org/10.1007/BF01650949>.
- [104] J. Guest, J. Prevost, T. Belytschko, Achieving minimum length scale in topology optimization using nodal design variables and projection functions, *Internat. J. Numer. Methods Engrg.* 61 (2) (2004) 238–254, <http://dx.doi.org/10.1002/nme.1064>.
- [105] E. Andreassen, B.S. Lazarov, O. Sigmund, Design of manufacturable 3D extremal elastic microstructure, *Mech. Mater.* 69 (1) (2014) 1–10, <http://dx.doi.org/10.1016/j.mechmat.2013.09.018>.
- [106] O. Sigmund, K. Maute, Topology optimization approaches, *Struct. Multidiscip. Optim.* 48 (6) (2013) 1031–1055, <http://dx.doi.org/10.1007/s00158-013-0978-6>.
- [107] N. van Dijk, K. Maute, M. Langelaar, F. van Keulen, Level-set methods for structural topology optimization: a review, *Struct. Multidiscip. Optim.* 48 (3) (2013) 437–472, <http://dx.doi.org/10.1007/s00158-013-0912-y>.
- [108] S. Chen, W. Chen, A new level-set based approach to shape and topology optimization under geometric uncertainty, *Struct. Multidiscip. Optim.* 44 (1) (2011) 1–18, <http://dx.doi.org/10.1007/s00158-011-0660-9>.
- [109] J. Martínez-Frutos, D. Herrero-Pérez, M. Kessler, F. Periago, Robust shape optimization of continuous structures via the level set method, *Comput. Methods Appl. Mech. Engrg.* 305 (2016) 271–291, <http://dx.doi.org/10.1016/j.cma.2016.03.003>.
- [110] J. Martínez-Frutos, D. Herrero-Pérez, M. Kessler, F. Periago, Risk-averse structural topology optimization under random fields using stochastic expansion methods, *Comput. Methods Appl. Mech. Engrg.* 330 (2018) 180–206, <http://dx.doi.org/10.1016/j.cma.2017.10.026>.
- [111] R. Bollapragada, R. Byrd, J. Nocedal, Adaptive sampling strategies for stochastic optimization, *SIAM J. Optim.* 28 (4) (2018) 3312–3343, <http://dx.doi.org/10.1137/17M1154679>.
- [112] Y. Xie, R. Bollapragada, R. Byrd, J. Nocedal, Constrained and composite optimization via adaptive sampling methods, *IMA J. Numer. Anal.* (2023) drad020, <http://dx.doi.org/10.1093/imanum/drad020>.
- [113] F. Beiser, B. Keith, S. Urbainczyk, B. Wohlmuth, Adaptive sampling strategies for risk-averse stochastic optimization with constraints, *IMA J. Numer. Anal.* 43 (2023) 3729–3765, <http://dx.doi.org/10.1093/imanum/drac083>.
- [114] R. Bollapragada, C. Karamanli, B. Keith, B. Lazarov, S. Petrides, J. Wang, An adaptive sampling augmented Lagrangian method for stochastic optimization with deterministic constraints, *Comput. Math. Appl.* 149 (2023) 239–258, <http://dx.doi.org/10.1016/j.camwa.2023.09.014>.
- [115] T.T.-K. Lau, H. Liu, M. Kolar, AdAdaGrad: Adaptive batch size schemes for adaptive gradient methods, 2024, [arXiv:2402.11215](https://arxiv.org/abs/2402.11215).
- [116] K. Svanberg, The method of moving asymptotes - a new method for structural optimization, *Internat. J. Numer. Methods Engrg.* 24 (1987) 359–373, <http://dx.doi.org/10.1002/nme.1620240207>.
- [117] R.T. Rockafellar, S. Uryasev, Optimization of conditional value-at-risk, *J. Risk* 2 (2000) 21–42, <http://dx.doi.org/10.21314/JOR.2000.038>.
- [118] D.P. Kouri, T.M. Surowiec, Risk-averse PDE-constrained optimization using the conditional value-at-risk, *SIAM J. Optim.* 26 (1) (2016) 365–396, <http://dx.doi.org/10.1137/140954556>.
- [119] D.P. Kouri, T.M. Surowiec, A primal-dual algorithm for risk minimization, *Math. Program.* 193 (1) (2022) 337–363, <http://dx.doi.org/10.1007/s10107-020-01608-9>.
- [120] E. Jansson, M. Kovács, A. Lang, Surface finite element approximation of spherical Whittle-Matérn Gaussian random fields, *SIAM J. Sci. Comput.* 44 (2022) A825–A842, <http://dx.doi.org/10.1137/21M1400717>.
- [121] D. Martínez-Maradiaga, A. Damonte, A. Manzo, J.H. Haertel, K. Engelbrecht, Design and testing of topology optimized heat sinks for a tablet, *Int. J. Heat Mass Transfer* 142 (2019) 118429, <http://dx.doi.org/10.1016/j.ijheatmasstransfer.2019.07.079>.
- [122] B.S. Lazarov, M. Schevenels, O. Sigmund, Topology optimization considering material and geometric uncertainties using stochastic collocation methods, *Struct. Multidiscip. Optim.* 46 (2012) 597–612, <http://dx.doi.org/10.1007/s00158-012-0791-7>.
- [123] M. Jansen, G. Lombaert, M. Schevenels, O. Sigmund, Topology optimization of fail-safe structures using a simplified local damage model, *Struct. Multidiscip. Optim.* 49 (2014) 657–666, <http://dx.doi.org/10.1007/s00158-013-1001-y>.
- [124] J. Zhao, C. Wang, Robust topology optimization under loading uncertainty based on linear elastic theory and orthogonal diagonalization of symmetric matrices, *Comput. Methods Appl. Mech. Engrg.* 273 (2014) 204–218, <http://dx.doi.org/10.1016/j.cma.2014.01.018>.
- [125] N. Aage, B. Lazarov, Parallel framework for topology optimization using the method of moving asymptotes, *Struct. Multidiscip. Optim.* 47 (4) (2013) 493–505, <http://dx.doi.org/10.1007/s00158-012-0869-2>.
- [126] D. Bolin, Spatial Matérn fields driven by non-Gaussian noise, *Scand. J. Stat.* 41 (3) (2014) 557–579, <http://dx.doi.org/10.1111/sjos.12046>.
- [127] J. Wallin, D. Bolin, Geostatistical modelling using non-Gaussian Matérn fields, *Scand. J. Stat.* 42 (3) (2015) 872–890, <http://dx.doi.org/10.1111/sjos.12141>.
- [128] A.J. Wathen, Preconditioning, *Acta Numer.* 24 (2015) <http://dx.doi.org/10.1017/S0962492915000021>.
- [129] J. Xu, L. Zikatanov, Algebraic multigrid methods, *Acta Numer.* 26 (2017) 591–721, <http://dx.doi.org/10.1017/S0962492917000083>.
- [130] hypre: High Performance Preconditioners, <https://llnl.gov/casc/hypre>, <https://github.com/hypre-space/hypre>.
- [131] P. Fischer, M. Min, T. Rathnayake, S. Dutta, T. Kolev, V. Dobrev, J.-S. Camier, M. Kronbichler, T. Warburton, K. Świrydowicz, J. Brown, Scalability of high-performance PDE solvers, *Int. J. High Perform. Comput. Appl.* 34 (2020) 562–586, <http://dx.doi.org/10.1177/1094342020915762>.
- [132] V.E. Henson, U.M. Yang, BoomerAMG: A parallel algebraic multigrid solver and preconditioner, *Appl. Numer. Math.* 41 (2002) 155–177, [http://dx.doi.org/10.1016/S0168-9274\(01\)00115-5](http://dx.doi.org/10.1016/S0168-9274(01)00115-5).