

Towards efficient uncertainty quantification in complex and large-scale biomechanical problems based on a Bayesian multi-fidelity scheme

Jonas Biehler · Michael W. Gee · Wolfgang A. Wall

Received: 20 May 2014 / Accepted: 2 September 2014 / Published online: 23 September 2014
© Springer-Verlag Berlin Heidelberg 2014

Abstract In simulation of cardiovascular processes and diseases patient-specific model parameters, such as constitutive properties, are usually not easy to obtain. Instead of using population mean values to perform “patient-specific” simulations, thereby neglecting the inter- and intra-patient variations present in these parameters, these uncertainties have to be considered in the computational assessment. However, due to limited computational resources and several shortcomings of traditional uncertainty quantification approaches, parametric uncertainties, modeled as random fields, have not yet been considered in patient-specific, nonlinear, large-scale, and complex biomechanical applications. Hence, the purpose of this study is twofold. First, we present an uncertainty quantification framework based on multi-fidelity sampling and Bayesian formulations. The key feature of the presented method is the ability to rigorously exploit and incorporate information from an approximate, low fidelity model. Most importantly, response statistics of the corresponding high fidelity model can be computed accurately even if the low fidelity model provides only a very poor approximation. The approach merely requires that the low fidelity model and the corresponding high fidelity model share a similar stochastic structure, i.e., dependence on the random input. This results in a tremendous flexibility in choice of the approximate model. The flexibility and capabilities of the framework are demonstrated by performing uncertainty quantifi-

cation using two patient-specific, large-scale, nonlinear finite element models of abdominal aortic aneurysms. One constitutive parameter of the aneurysmatic arterial wall is modeled as a univariate three-dimensional, non-Gaussian random field, thereby taking into account inter-patient as well as intra-patient variations of this parameter. We use direct Monte Carlo to evaluate the proposed method and found excellent agreement with this reference solution. Additionally, the employed approach results in a tremendous reduction of computational costs, rendering uncertainty quantification with complex patient-specific nonlinear biomechanical models practical for the first time. Second, we also analyze the impact of the uncertainty in the input parameter on mechanical quantities typically related to abdominal aortic aneurysm rupture potential such as von Mises stress, von Mises strain and strain energy. Thus, providing first estimates on the variability of these mechanical quantities due to an uncertain constitutive parameter, and revealing the potential error made by assuming population averaged mean values in patient-specific simulations of abdominal aortic aneurysms. Moreover, the influence of correlation length of the random field is investigated in a parameter study using MC.

Keywords Abdominal aortic aneurysm · Uncertainty quantification · Bayes · Monte Carlo · Random fields

J. Biehler · W. A. Wall (✉)
Institute for Computational Mechanics, Technische Universität München, Boltzmannstr. 15, 85747 Garching b. München, Germany
e-mail: wall@lnm.mw.tum.de

M. W. Gee
Mechanics and High Performance Computing Group, Technische Universität München, Parkring 35, 85748 Garching b. München, Germany

1 Introduction

The assessment of biomechanical factors of cardiovascular and other diseases using computational tools has led to tremendous progress in the understanding of the underlying processes and causes leading to severe medical conditions and often death. Furthermore, these computational models can offer new, unprecedented diagnostic and predictive pos-

sibilities, presenting a promising avenue to further advance in medical health care. Hence, the improvements achieved in the last decade have fueled the desire for the application of these tools in clinical settings and medical research.

However, most computational models require a multitude of patient-specific data and input parameters to accurately predict the relevant patient-specific biomechanical factors. Many of these aforementioned input parameters are unavailable or can only be estimated in some way. Often when truly patient-specific data are inaccessible assumptions about these parameters are made, or population-averaged values are used. The difference between the actual patient-specific values and any assumed or averaged values in the input parameters in a deterministic model of course translates to uncertainty or potential error in the computed quantities.

For instance, in the context of computational models for biomechanical problems, already the process to get from medical image data to a discretized model is a potential source for uncertainty and error in the computational geometry. Additionally, this task is often also dependent on the operator performing the segmentation and meshing process and thus not objective and reproducible. The differences in model geometry of course translate into differences in numerical results (Moore et al. 1997). Furthermore, the resolution of state-of-the-art clinical CT or MR scanners often is not high enough to determine certain important geometrical features. An even more severe problem is that patient-specific physical parameters, such as constitutive properties, almost never can be measured accurately. Many of these parameters intrinsically vary within and between patients and are also significantly affected by lifestyle, age, and diseases, which lead to further substantial variations. Moreover, as pointed out previously by Chen et al. (2013), uncertainty is introduced also by the often ambiguous choice of mathematical models and applied boundary conditions, which will affect the simulation results.

Let us consider, e.g., 3D models of the human vasculature in the following. As the wall thickness of blood vessels can generally not be obtained from medical images, assumptions about the wall thickness and its regional distribution are made when creating structural models of the vessel wall. Especially in cases where the patient suffers from a cardiovascular disease like arteriosclerosis or abdominal aortic aneurysm (AAA), vessel walls are subject to pathological changes resulting in a wall thickness deviating strongly from population mean values and exhibiting large spatial intra-patient variations (Reeps et al. 2012; Raghavan et al. 2006, 2011).

Another source of uncertainty are the unknown patient-specific constitutive properties of the vessel wall. As in the case of the wall thickness, the constitutive properties of the vessel wall are severely affected by advancing cardiovascular diseases. As shown in many ex vivo studies, the non-linear material properties of an aneurysmatic arterial wall

are subject to large inter- and intra-patient variations, and constitutive parameters determined by tensile tests vary significantly within a patient, patient cohort, and experimental setup in the various published studies (Reeps et al. 2012; Marini et al. 2011; Raghavan et al. 2006, 2011; Vande Geest et al. 2006b,c; Raghavan et al. 1996; Vallabhaneni et al. 2004; Thubrikar et al. 2001). Even in non-arteriosclerotic human aorta, the stiffness parameters reported in the literature easily span one order of magnitude (Roccabianca et al. 2014). Despite the progress that has been made regarding in vivo characterization of constitutive properties based on pulse wave velocity or local changes in diameter of the aorta (Schulze-Bauer and Holzapfel 2003; Astrand et al. 2011; Vappou et al. 2010), current in vivo approaches are not yet able to detect short-scale spatial variations. Furthermore, none of the above-mentioned approaches has been successfully applied and validated for, e.g., AAA or human aorta by tensile tests.

While the exact knowledge of patient-specific values for wall thickness and constitutive parameters may be irrelevant for pure computational fluid models, this is not the case for solid or fluid–solid interaction models. Therefore, in absence of truly patient-specific parameters, it is extremely important to be aware of the uncertainty in population averaged or assumed patient-specific model parameters. In lack of true patient-specific values, these uncertainties should be incorporated into the computational model to quantify the impact of the uncertainty on the computed results and to obtain more reliable predictions or worst-case scenarios. The identification and quantification of the uncertainties in the computational results will help to access the margin of error of simulation results due to uncertain model parameters, thus allowing estimates on the accuracy of deterministic models. This is crucial if computational models are, ultimately, to be used for diagnosis of cardiovascular diseases, or to obtain recommendations for prevention and treatment of individual patients. A prominent example is the patient-specific rupture risk stratification of AAA using finite element models to determine whether the patient should undergo elective surgery. Additionally, the results could help to improve computational models and provide directions for future research, as the major contributors to the output uncertainty are identified.

In order to achieve such a quantification of uncertainties, two major steps are required. First, a suitable probabilistic description of the uncertain model parameters must be found, and second, an appropriate approach to propagate the uncertainties through the model is needed. As far as uncertain material properties and arterial wall thickness are considered, two options exist to model these uncertainties in a probabilistic fashion. The simpler approach models the constitutive parameters as a spatially constant random variable, which follows a prescribed probability distribution, while thus neglecting spatial variability. The other option,

which also accounts for spatial variations, are so-called random fields (Vanmarcke 2010; Shinozuka and Jan 1972). An essential feature of random fields is the ability to capture the spatial correlation, encoded in the autocorrelation function (ACF), of a varying property at different locations.

In either case, the resulting computational model is no longer deterministic, and hence, a stochastic problem must be solved. Various approaches to solve this kind of stochastic boundary value problem have been proposed in recent decades (Spanos and Zeldin 1998; Stefanou 2009; Panayirci and Schuëller 2011; Xiu 2010; Eldred 2009). However, only few approaches scale to large nonlinear problems with moderate to high stochastic dimension. The majority of the proposed approaches fall into one of the following two categories. The first is based on some form of Polynomial Chaos Expansion (PCE) (Xiu 2010; Ghanem and Spanos 2003; Xiu 2009). The major disadvantage of solution schemes based on PCE is the curse of dimensionality, which renders these approaches infeasible for problems with high stochastic dimension as is the case when considering random fields. The second broad class of algorithms is based on sampling approaches like Monte Carlo (MC) and all its variants (Metropolis and Ulam 1949; McKay et al. 1979; Bourgund and Bucher 1986; Cliffe et al. 2011; Pellissetti and Schuëller 2009; Pradlwarter et al. 2007; Olsson and Sandberg 2002; Neal 2001; Hurtado and Barbat 1998). As reported by Stefanou (2009), direct MC might still be the only universal approach for uncertainty quantification (UQ) in complex systems, because it only requires a deterministic solver and the ability to draw samples from the distributions of the uncertain input parameters. However, while the stochastic problem at hand can in principle be solved using direct MC, the computational costs associated with standard MC render this approach infeasible for realistic, nonlinear and hence quite complex biomechanical models, as MC requires typically several thousand model evaluations.

In this study, we present a framework which incorporates a probabilistic description of spatially varying model parameters, based on experimental data. Moreover, this framework allows for the efficient solution of the resulting stochastic boundary value problem using approximate models in combination with nonparametric Bayesian regression (Koutsourelakis 2009). We demonstrate that this approach is able to reproduce the MC reference solution at a fraction of the computational costs and delivers accurate results even if strongly simplified approximate models are used. Thus, the computational costs of UQ can be reduced to an acceptable level even for large-scale, nonlinear, and complex biomechanical problems with uncertainties modeled as random fields.

As a first step toward full UQ incorporating all relevant uncertain model parameters in patient-specific biomechanics, we exemplarily consider two patient-specific AAA geometries and the uncertainty in one of the constitutive para-

eters of the aneurysm wall. The constitutive parameter is modeled as random field, since our experimental data suggest that the chosen parameter exhibits high inter- and intra-patient variations. We present the method and demonstrate, through a proof of concept study, the impact of uncertainties present in the constitutive parameter in the arterial wall on mechanical quantities usually related to AAA rupture potential, such as von Mises stress/strain or strain energy density. In addition, we perform a parameter study to investigate the influence of the correlation length of the random field using direct MC.

Thus, the purpose of this work is twofold. On the one hand, a novel UQ approach is presented, and its efficiency and accuracy are shown for large-scale nonlinear models. On the other hand, in addition to the methodology presented in this paper, the uncertainties in the simulation results of patient-specific AAA models are quantified, and the ramifications for patient-specific computational assessment of AAA rupture potential are discussed.

To the knowledge of the authors, this is the first time that UQ with respect to a spatially correlated random quantity is performed employing nonlinear patient-specific finite element models. Hitherto, research regarding UQ in cardiovascular mechanics was limited to either idealized geometries and uncertainties modeled as independent random variables rather than correlated stochastic processes. For instance, Sankaran and Marsden (2011) investigated idealized geometries of AAAs considering parametric uncertainties such as the radius, while Chen et al. (2013) and Xiu and Sherwin (2007) studied the effect of uncertain parameters in one-dimensional models of the arterial network. Although we consider patient-specific models of AAA, we stress the fact that our approach is very general and can be applied to a wide range of problems and is by no means limited to this particular application.

The remainder of this paper is structured as follows. In Sect. 2, the stochastic boundary value problem and the probabilistic model used to describe the random and spatial variations of the constitutive parameter are provided. Section 3 summarizes the multi-fidelity MC approach utilized here and highlights important aspects. The described methodology is applied to two large, nonlinear, patient-specific finite element models of AAAs, and the results of these model problems are discussed in Sect. 4. By means of a simple proof of concept example, we show in Sect. 5 that geometric uncertainties can be considered in addition to the constitutive uncertainty as well. Conclusions are drawn in Sect. 6.

2 Stochastic problem formulation

The material properties of AAAs are subject to large inter- and intra-patient variations in all of the material properties (Reeps et al. 2012; Raghavan et al. 2011). Hence, we use the

random field approach to formulate a stochastic material law for the aneurysmatic arterial wall. These random fields are constructed based on our own experimental data from tensile tests. Reliable information about the spatial distribution and correlation of constitutive properties in AAA is scarce; hence, some auxiliary assumptions have to be made for the characterization of the field. Nevertheless, the assumptions regarding the probabilistic characteristics of the random field are far less restrictive than those typically made in a deterministic problem setting where the variations or uncertainties in the model parameters are commonly ignored entirely. In addition, the sensitivity of assumptions on the results can also be tested by the proposed approach.

2.1 Probabilistic constitutive model

A homogenous, univariate, lognormal three-dimensional random field is used to describe the inter- and intra-patient variations of one constitutive parameter of a hyperelastic constitutive law for the AAA wall. More specifically, we use the hyperelastic constitutive model proposed in (Raghavan and Vorp 2000) as a starting point for our probabilistic description of the arterial wall:

$$\Psi = \alpha(\bar{I}_1 - 3) + \beta(\bar{I}_1 - 3)^2 + \frac{\kappa}{\eta^2}(\eta \ln J + J^{-\eta} - 1) \quad (1)$$

where \bar{I}_1 is the first invariant of the isochoric Right-Cauchy-Green strain tensor $\bar{\mathbf{C}}$ and $J = \det \mathbf{F}$ is the determinant of the deformation gradient \mathbf{F} . The parameters of the isochoric part α and β can be determined by fitting the strain energy function given above to experimentally measured stress-stretch curves as we precisely described in (Reeps et al. 2012). A fixed bulk modulus of $\kappa = \frac{2\alpha}{1-2\nu}$ with $\nu = 0.49$ was used throughout this work and $\eta = -2$.

The parameters α and β were fitted for $n = 142$ tensile test specimens harvested in open surgery using the methodology described in (Reeps et al. 2012). Figure 1 depicts the histograms of the constitutive parameters α and β and reveals that the intra-patient variations are of the same order

of magnitude as the inter-patient variations. It was found that β exhibits a larger variance as well as a greater variation between individual patients as compared to α . Hence, as a first step towards a fully probabilistic constitutive framework, we choose to model only the parameter β as a random field leaving the remaining constitutive parameters $\alpha = 0.059$ and $\kappa = 5.9$ MPa deterministic population-averaged mean values.

2.1.1 Generation of non-Gaussian random fields using spectral representation and translation process theory

Solving stochastic finite element problems with uncertainties modeled as random fields using a sampling-based approach requires the computation of discretized sample functions, often referred to as realizations, of the random field. The spectral representation method in combination with translation process theory is a widely used approach to generate these sample functions based on a large number of trigonometric functions with independent random phase angles (Shinozuka and Jan 1972). First, a Gaussian random field is created with the spectral representation method, which is then mapped into a non-Gaussian field using translation process theory (Grigoriu 1995).

In order to use the spectral representation method, two probabilistic characteristics of the random field, the first-order probability distribution and the autocorrelation function (ACF), need to be specified. Experimental data from tensile tests was utilized to obtain an estimate for the probability distribution using 142 tested specimens from 49 patients, which provides empirical information on the range and relative frequency of β , see Fig. 1b. A lognormal probability distribution for the parameter β , which is visualized in Fig. 1b, was obtained by applying the *lognfit* function in matlab to this data.

$$p_{\beta}(\beta) = \frac{1}{\beta\sigma_0\sqrt{2\pi}} e^{-\frac{(\ln \beta - \mu)^2}{2\sigma_0^2}}; \quad (2)$$

$$\mu = 1.0857, \quad \sigma_0 = 0.9205.$$

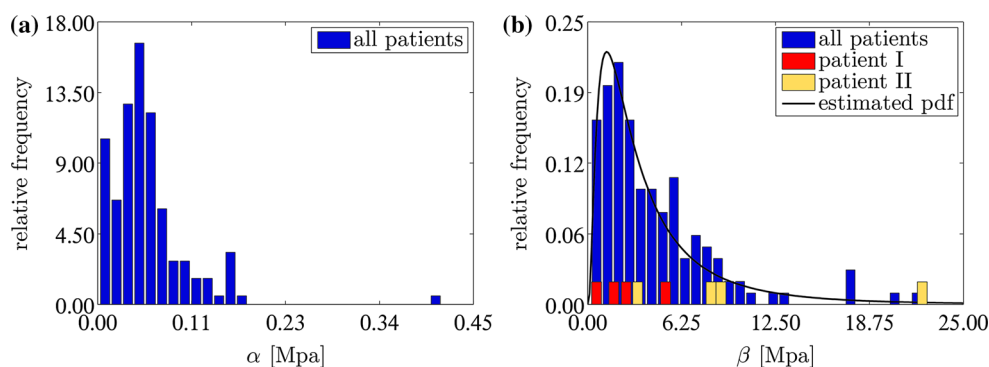


Fig. 1 Constitutive parameters obtained from tensile testing. **a** Histogram of fitted α parameter. **b** Histogram of fitted β parameter

As usual in stochastic mechanics (Charmis et al. 2007), our experimental data did not contain sufficient information to fully determine the second characteristic of the random field, the ACF. Based on the reasonable assumption that the constitutive parameters vary smoothly in space, we chose the squared exponential ACF

$$R_{NG}(\boldsymbol{\tau}) = \sigma_\beta^2 e^{-\left(\frac{|\boldsymbol{\tau}|}{d}\right)^2}. \tag{3}$$

The ACF describes the decay of covariance, which, for homogenous fields, is solely based on the distance between two points $\boldsymbol{\tau} = \mathbf{x} - \mathbf{x}'$ and a parameter called correlation length d . The correlation length controls the distance between two points above which the correlation between the values at these points approaches practically zero. The parameter σ_β^2 denotes the variance of the lognormal distribution in (2). The ACF can be transformed into the power spectral density (PSD) of the random field using the Wiener-Khinchine relationship

$$S_{NG}(\boldsymbol{\kappa}) = \frac{1}{(2\pi)^3} \int_{-\infty}^{\infty} R_{NG}(\boldsymbol{\tau}) e^{i\boldsymbol{\kappa}^T \boldsymbol{\tau}} d\boldsymbol{\tau}, \tag{4}$$

where $\boldsymbol{\kappa}$ denotes the frequency vector. The PSD is needed to create sample functions of the random field using the spectral representation method (Shinozuka and Jan 1972). The formula to generate sample functions of a univariate three-dimensional Gaussian random field is given in (Shinozuka and Deodatis 1996):

$$C(\mathbf{x}) = \sqrt{2} \sum_{n_1=0}^{N_1-1} \sum_{n_2=0}^{N_2-1} \sum_{n_3=0}^{N_3-1} \sqrt{2S_{NG}(\kappa_{1n_1}, \kappa_{2n_2}, \kappa_{3n_3}) \Delta\kappa_1 \Delta\kappa_2 \Delta\kappa_3} \left[\begin{aligned} &\cos(\kappa_{1n_1}x_1 + \kappa_{2n_2}x_2 + \kappa_{3n_3}x_3 + \xi_{n_1n_2n_3}^{(1)}) \\ &+ \cos(\kappa_{1n_1}x_1 - \kappa_{2n_2}x_2 + \kappa_{3n_3}x_3 + \xi_{n_1n_2n_3}^{(2)}) \\ &+ \cos(\kappa_{1n_1}x_1 + \kappa_{2n_2}x_2 - \kappa_{3n_3}x_3 + \xi_{n_1n_2n_3}^{(3)}) \\ &+ \cos(\kappa_{1n_1}x_1 - \kappa_{2n_2}x_2 - \kappa_{3n_3}x_3 + \xi_{n_1n_2n_3}^{(4)}) \end{aligned} \right], \tag{5}$$

with $\Delta\kappa_i = \frac{\kappa_{iu}}{N_i}$ and $\kappa_{in_i} = n_i \Delta\kappa_i$. Herein, κ_{iu} represents the cutoff frequency above which the PSD is of insignificant magnitude and can assumed to be zero. This threshold depends on the correlation length of the field and has to be adjusted accordingly. The $\xi_{n_1n_2n_3}^{(i)}$ denote the independent random phase angles which are uniformly distributed between $[0, 2\pi]$.

Sample functions generated with (5) are asymptotically Gaussian for $N \rightarrow \infty$ by virtue of the central limit theorem. However, a value of $N_i = 64$ yielded sufficiently accurate results and hence was used throughout this work for the generation of sample functions. Using translation process theory,

the sample of the Gaussian random field can be transformed into a sample of a non-Gaussian random field, the first-order probability density of which obeys the desired lognormal probability density function (PDF) (Grigoriu 1995)

$$G(\mathbf{x}) = F_{NG}^{-1}\{\Phi[C(\mathbf{x})]\}, \tag{6}$$

where F_{NG}^{-1} is the inverse cumulative distribution function (CDF) of the desired lognormal distribution as given in (2) and Φ is the normal CDF of the underlying Gaussian process. The nonlinear transform described above leads to a correlation distortion in the sense that (6) not only changes the first-order distribution of the stochastic process but also affects the ACF. In order to account for this distortion as well as any potentially arising incompatibilities between prescribed CDF and ACF (Grigoriu 1995), we implemented a framework based on (Shields et al. 2011), which allows to efficiently generate sample functions of random fields that match a prescribed PDF as well as the desired ACF. For a more detailed in-depth overview over algorithms for generating non-Gaussian random fields, the reader is referred to (Bocchini and Deodatis 2008) and (Stefanou and Papadrakakis 2007).

With the non-Gaussian random field described above, we can extend the constitutive law for the aneurysmatic arterial wall to incorporate the uncertainty in the parameter $\beta = \beta(\mathbf{x}, \boldsymbol{\xi})$. Thus, β becomes a function of the location and the random phase angles $\xi_{n_1n_2n_3}^{(i)}$, which are summarized in the vector $\boldsymbol{\xi}$ for notational convenience. Once a sample of random phase angles is drawn, a realization of the three-dimensional random field is computed using (5) and (6). The constitutive parameter $\beta(\mathbf{x}, \boldsymbol{\xi})$ can then be evaluated, e.g., at the midpoint of each element in the AAA wall. This value is then assigned to the element as local constitutive parameter. The full stochastic version of the strain energy function, which now depends on the location \mathbf{x} as well as on the vector of random phase angles $\boldsymbol{\xi}$, is

$$\Psi(I_1, J, \mathbf{x}, \boldsymbol{\xi}) = \alpha(\bar{I}_1 - 3) + \beta(\mathbf{x}, \boldsymbol{\xi})(\bar{I}_1 - 3)^2 + \frac{\kappa}{\eta^2}(\eta \ln J + J^{-\eta} - 1). \tag{7}$$

Note that the methodology to create sample functions of random fields described above can also be used to consider models with any other uncertain quantities, e.g., uncertain wall thickness as later also used in one numerical example. In case of multiple uncertain model parameters, each of those can be modeled individually as random field using the formulas given above, if the parameters are uncorrelated. However, if these parameters are correlated, cross-correlated random vector fields should be used as a probabilistic model. The generation of non-Gaussian cross-correlated random vector fields is more involved and an area of ongoing and vivid research. Several methods for the generation of such

cross-correlated sample functions have for example been proposed in (Popescu et al. 1998; Field and Grigoriu 2012; Vořechovský 2008; Cho et al. 2013). In this work, we restrict ourselves to random fields, which are not cross-correlated for ease of exposition, and because at this stage, we can only speculate about potential cross-correlations between different model parameters. But while sample generation will be more involved in the cross-correlated random vector field case, we do not see a reason why the proposed Bayesian multi-fidelity approach should not work as well.

2.2 Nonlinear stochastic boundary value problem

Having established a probabilistic description of the material properties, the boundary value problem that has to be solved becomes stochastic as well. For a formal description of the problem, a probability space $(\Omega, \mathcal{F}, \mathcal{P})$ is defined, where Ω is the event space, \mathcal{F} a σ -algebra and \mathcal{P} a probability measure. The random phase angles now define a mapping $\xi : \Omega \rightarrow [0, 2\pi]^d$, where d corresponds to the total number of phase angles used to generate sample functions of the random field. Additionally, we ask that the random phase angles obey a uniform distribution $\pi_{\xi_i} = \mathcal{U}(0, 2\pi)$. The stochastic boundary value problem can then be written in the most general form as

$$\mathcal{L}(\mathbf{u}(\mathbf{x}, \xi), \xi) = \mathbf{f}, \quad (8)$$

with the nonlinear stochastic differential operator \mathcal{L} depending on the vector $\xi \in \mathbb{R}^d$ of real random variables describing, e.g., the uncertainty in the material properties or the wall thickness. Only deterministic boundary conditions and body loads are considered in this work. The solution $\mathbf{u}(\mathbf{x}, \xi)$ of the problem, the displacement, thus becomes a stochastic process as well, so do all quantities derived from the displacement field or functions thereof, such as the von Mises stress σ_{VM} , von Mises strain e_{VM} or strain energy density Ψ . For the sake of brevity, we will refer to all quantities of interest as $y(\mathbf{x}, \xi)$ in the following, and in addition, we will omit the dependence on the location and random phase angles for notational convenience. Apart from the expected value $\mathbb{E}[y]$ or the coefficient of variation (COV) of the quantity of interest, we are interested in obtaining the density $\pi_y(y)$ or a response statistic like

$$Pr [y \in \mathcal{A}] = \int 1_{\mathcal{A}}(y(\xi)) \pi_{\xi}(\xi) d\xi, \quad (9)$$

where \mathcal{A} is a π_{ξ} -measurable subset and $1_{\mathcal{A}}$ is the corresponding indicator function. The case where $\mathcal{A} = \{y | y > y_0\}$ is of particular interest if we want to know the probability that, e.g., the von Mises stress exceeds a certain threshold y_0 ; thus, (9) can be used to define a global or local failure probability.

3 Solution of the nonlinear stochastic boundary value problem

Solving the stochastic boundary value problem using sampling-based approaches like Monte Carlo (MC) only requires a deterministic solver and the ability to draw samples from the distribution $\pi_{\xi}(\xi)$. The distribution of the quantity of interest $\pi_y(y)$ can then be approximated by computing a number (N_{SAM}) of random samples $y_i = y(\xi_i)$, with $\xi_i \sim \pi_{\xi}(\xi)$; which are sometimes also referred to as particles. These particles are equipped with a weight W^i according to their weight or probability, respectively. With slight abuse of notation, the approximation of $\pi_y(y)$ can be computed by:

$$\pi_y(y) \approx \sum_{i=1}^{N_{\text{SAM}}} W^i \delta_{y^{(i)}}(y(\xi)) \quad (10)$$

where $\delta_{y^{(i)}}$ is the delta-Dirac mass. In case of direct MC the weight of all samples is the same and equal to $1/N_{\text{SAM}}$. Hence:

$$\pi_y(y) \approx \frac{1}{N_{\text{SAM}}} \sum_{i=1}^{N_{\text{SAM}}} \delta_{y^{(i)}}(y(\xi)). \quad (11)$$

The accuracy of the MC approximation depends on the number of independent samples N_{SAM} with a probabilistic error of $O(1/\sqrt{N_{\text{SAM}}})$. In contrast to Polynomial Chaos Expansion (PCE)-based approaches, the accuracy does not depend on the stochastic dimension of the problem per se. However, for an accurate computation of the response statistics, the number of required samples can be prohibitively large, especially if one is interested in the tails of the distribution to evaluate failure probabilities. With a solution time of current state-of-the-art models in computational biomechanics of the order of several CPU hours, direct MC with several thousand samples is impractical for the application to patient-specific models and large patient cohorts. The use of advanced MC schemes like subset simulation, importance sampling or sequential Monte Carlo (SMC) methods can alleviate the computational burden (Neal 2001; Au and Beck 2001; Del Moral et al. 2006). However, the aforementioned advanced sampling schemes are not applicable if one is interested in uncertainty quantification with respect to multiple response quantities simultaneously, e.g., stresses and strains.

3.1 Bayesian multi-fidelity Monte Carlo

To reduce the computational effort, we implemented a framework based on a multi-fidelity sampling strategy employing the concepts proposed in (Koutsourelakis 2009; Kennedy and O'Hagan 2000). The general idea is to incorporate models with different levels of sophistication or fidelity in

a sampling-based UQ approach. Roughly speaking, the sampling is done on an inexpensive, approximate, low fidelity model. Then, to account for the discrepancy between the low fidelity and high fidelity solution, a probabilistic correction factor is used, for the calculation of which only very few evaluations of the expensive high fidelity model are required.

We start, given a particular system of interest, by assuming that we have two computational models for this system, an accurate, expensive high fidelity model and a cheap approximate version. The latter could be obtained, for instance, by creating a much coarser discretization. Suppose that we are interested in computing a particular quantity $y(\xi)$ or rather its probability density $\pi_y(y)$. Analogous to $y(\xi)$, the quantities of interest computed with the high fidelity model, we will refer to the corresponding quantity and its density in the approximate model as $x(\xi)$ and $\pi_x(x)$ from here on. For example, if y is the von Mises stress at a particular location in the accurate model, x denotes the von Mises stress evaluated at the same spatial location in the low fidelity model. Since the approximate model is cheap to evaluate, the density $\pi_x(x)$ can be readily approximated using any kind of sampling algorithm such as direct MC or SMC.

If the low fidelity model and the high fidelity model have the same stochastic structure, that is, they show a similar dependence on the random input parameters, then there is a statistical correlation between x and y , as shown in Fig. 2a. This correlation can be exploited to accurately predict $\pi_y(y)$ and does not at all depend on how good the low fidelity model approximates a value in a deterministic sense. Let us first consider the most extreme case, the one where we have a direct one-to-one relationship between the two model outputs x and y , e.g., as shown in Fig. 2b. It is obvious that it would be easy to come up with a regression model that can reproduce the interrelation between x and y . In combination with $\pi_x(x)$ obtained from sampling, the calculation of $\pi_y(y)$ is trivial. However, the relationship between x and y can also be exploited if no strict deterministic relationship exists. A noisy correlation as shown in Fig. 2a can also be exploited. If the interrelation between the two model outputs is captured as

conditional probability density $p(y|x)$, $\pi_y(y)$ can be readily computed using

$$\pi_y(y) = \int p(y|x)\pi_x(x)dx. \tag{12}$$

Hence, the methodology has two major requirements. First, the approach of course requires an approximate model that, while being computationally cheap, still provides sufficient stochastic information about the quantity of interest. Second, an efficient way to infer the conditional density $p(y|x)$ is needed. As proposed by Koutsourelakis (2009), we use a Bayesian regression model to estimate the conditional density $p(y|x)$. Therefore, we need a set of training samples $\{(x_j, y_j)\}_{j=1}^n$, i.e., solutions of the low and high fidelity model based on the same realizations of the random field or the same set of random phase angles ξ_j , respectively. The Bayesian regression model then establishes a probabilistic link between the low and high fidelity model, thus allowing us to compute $\pi_y(y)$ using (12) or $Pr[y \in \mathcal{A}]$ if we rewrite (9):

$$\begin{aligned} Pr[y \in \mathcal{A}] &= \mathbb{E}_x [Pr[y \in \mathcal{A}|x]] \\ &= \int Pr[y \in \mathcal{A}|x]\pi_x(x)dx. \end{aligned} \tag{13}$$

As a consequence, the regression model, in combination with the information from the low fidelity model, serves as a data fit surrogate for the high fidelity model. Another interpretation is that the regression model provides a probabilistic correction factor to account for the discrepancy between high and low fidelity solution. In order to facilitate the understanding for the reader, we want to provide a sketch of the basic steps of the proposed approach before explaining the details of the Bayesian regression and the computation of $\pi_y(y)$ in the following subsections. Furthermore, we want to highlight some advantages of the proposed approach at this point.

Given an existing accurate, high fidelity model of our system, an approximate low fidelity model has to be constructed. This approximate model should have a similar sto-




Fig. 2 Correlation between solution on low fidelity model x and high fidelity model y . **a** Noisy correlation. **b** One-to-one dependence

Table 1 Basic steps of the proposed multi-fidelity MC scheme

<i>Bayesian multi-fidelity MC</i>	
1.	Construct a low fidelity and a high fidelity model of the system of interest, e.g., by creating a coarse and a fine discretization
2.	Compute $\pi_x(x)$ by sampling of the cheap low fidelity model
3.	Select a very small subset of the samples $\{x_j\}_{j=1}^n$, fully covering the (empirical) support of $\pi_x(x)$
4.	Compute corresponding high fidelity solution $\{y_j\}_{j=1}^n$ for the selected subset of samples
5.	Use the samples $\{x_j, y_j\}_{j=1}^n$ from both models as training data for a Bayesian regression model to obtain a (probabilistic) relationship between x and y , i.e., $p(y x)$
6.	Approximate $\pi_y(y)$ and $Pr[y \in \mathcal{A}]$ from $p(y x)$ and $\pi_x(x)$ using (12) and (13), respectively

chastic structure as the high fidelity model. However, it is not necessary that the approximate model provides an accurate approximation in a deterministic sense, as needed in deterministic multi-level/grid schemes. The only prerequisite is a mere statistical correlation between the high fidelity and the approximate model, meaning that the approximate model and the high fidelity model have a similar dependence on the random input parameters. This very weak requirement offers a tremendous flexibility in the choice of approximate models. One way to create an approximate model is to simply use a coarser discretization. Additionally, model reduction techniques beyond coarsening the spatial and temporal discretization can be used. As shown in one of our numerical examples in Sect. 4, the multi-fidelity MC approach also works if a simpler physical model is used. In principle, also model reduction techniques such as proper orthogonal decomposition could provide suitable low fidelity models.

In the second step of the proposed approach, the density $\pi_x(x)$ of our quantity of interest x has to be determined using the low fidelity model. Thereby, the sampling can be carried out using any kind of sampling algorithm and is by no means limited to direct MC. Independent of the particular choice of the sampling algorithm computing $\pi_x(x)$ by sampling is much cheaper as compared to sampling of the high fidelity model. We used approximate models which are up to a factor of 120 cheaper to compute than the corresponding high fidelity model. In some cases, a factor of 1,000 can be achieved (Koutsourelakis 2009). Then, a small subset $\{x_j\}_{j=1}^n$ of all computed samples is selected. For an accurate computation of $\pi_y(y)$, these samples should fully cover the support of $\pi_x(x)$, i.e., cover the whole range of values that x can take. Since usually only a particulate approximation of $\pi_x(x)$ is available, this range is approximated by the samples with the smallest and largest x value, respectively. In our experience, 100–200 evenly distributed samples which cover the range of all possible x values yield excellent results. Now the high fidelity solution $\{y_j\}_{j=1}^n$ corresponding to the $\{x_j\}_{j=1}^n$ is computed. The dataset $\{(x_j, y_j)\}_{j=1}^n$ is then used to fit a Bayesian regression model, which provides probabilistic information on y given x . More specifically, the Bayesian regression provides the desired conditional probability dis-

tribution $p(y|x)$. Compared to the costs associated with the computation of the high fidelity solution for the n samples, the costs of fitting the regression model are negligible. While we adopt the regression model used in (Koutsourelakis 2009) with minor modifications, the multi-fidelity MC approach is, of course, not limited to this particular choice of regression model. Once the conditional density $p(y|x)$ is obtained, we can use standard probability theory to compute $\pi_y(y)$. Overall, the proposed multi-fidelity MC approach can be summarized in six steps, which are given in Table 1.

3.1.1 Bayesian regression model

Here, we will briefly outline the main concepts of the Bayesian regression model proposed by Koutsourelakis (2009), which we adopt and modify by the addition of a linear term. As mentioned before, one key aspect of the proposed procedure lies in the efficient determination of the conditional density $p(y|x)$ or conditional probability $Pr[y \in \mathcal{A}|x]$. We start by assuming an a priori unknown functional relationship between x and y plus some additive Gaussian noise:

$$y(\xi_i) = f(x(\xi_i); \theta) + \sigma Z_i, \quad Z_i \sim \mathcal{N}(0, 1) \quad (14)$$

with $f(x, \theta)$ being a combination of a first-order polynomial and a sum of Gaussian kernel functions

$$f(x, \theta) = a_0 + a_1 x + \sum_{j=1}^k a_j e^{-\tau_j |v_j - x|^2}, \quad (15)$$

where the vector θ contains all parameters of the model

$$\theta = \{k, \{a_j\}_{j=0}^k, \{\tau_j\}_{j=2}^k, \{v_j\}_{j=2}^k\}. \quad (16)$$

Thus, given the model parameters θ the conditional distribution $p(y|x, \theta, \sigma)$ is Gaussian:

$$p(y|x) \approx p(y|x, \theta, \sigma) = \mathcal{N}(f(x, \theta), \sigma) = \phi(y - f(x, \theta), \sigma), \quad (17)$$

where $\phi(y - f(x, \theta), \sigma)$ is the Gaussian PDF defined as

$$\begin{aligned} \phi(y - f(x, \theta), \sigma) &= \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left\{-\frac{(y - f(x, \theta))^2}{2\sigma^2}\right\}. \end{aligned} \tag{18}$$

Following a fully Bayesian approach, all unknown parameters are treated as random variables that are equipped with a probability distribution. Hence, instead of trying to obtain point estimates, the goal is to infer a probability distributions for each of the model parameters in θ and the variance of the noise σ^2 . This of course means that the conditional distribution $p(y|x)$, where the parameters θ are integrated out, is not necessarily Gaussian, and more complex distributions can be obtained. Exploiting the statistical correlation between the two quantities x and y , we can use the training samples $\{(x_j, y_j)\}_{j=1}^n$ to infer the parameters θ of the regression model as well as the variance of the Gaussian noise σ^2 using Bayes' rule

$$\begin{aligned} \pi_{\theta, \sigma^2}(\theta, \sigma^2) &= p(\theta, \sigma^2 | (x_{1:n}, y_{1:n})) \\ &= \frac{p((x_{1:n}, y_{1:n}) | \theta, \sigma^2) p(\theta) p(\sigma^2)}{p(x_{1:n}, y_{1:n})} \tag{19} \\ &\propto p((x_{1:n}, y_{1:n}) | \theta) p(\theta) p(\sigma^2). \end{aligned}$$

After specifying the prior distributions $p(\theta)$ and $p(\sigma^2)$, Bayes' rule modifies these prior distributions based on additional available information. In our case, this information is provided by the training samples $\{(x_j, y_j)\}_{j=1}^n$. Using these training samples, it is possible to evaluate how probable it is to observe the data given a particular choice of θ and σ^2 using the likelihood $p((x_{1:n}, y_{1:n}) | \theta, \sigma^2)$. The combination of the prior and the likelihood gives rise to the posterior distribution of the parameters $p(\theta, \sigma^2 | (x_{1:n}, y_{1:n}))$ given the data. The Bayes' rule adjusts the distribution of the model parameters such that they are both probable under the prior and compatible with the observed evidence or data.

3.1.2 Prior distribution

We adopted the choice of priors given in (Koutsourelakis 2009) for the model parameters. Algebraic rearrangement and integration yield the following prior for the model parameters:

$$\begin{aligned} p(\theta | s, a_\tau, a_\mu, a_0, b_0) &= \frac{1}{(s + k)^{k+1}} \\ &\times \prod_{j=1}^k \frac{\Gamma(a_\tau + 1)}{\Gamma(a_\tau)} \frac{a_\tau^{a_\tau}}{\tau_j^{(a_\tau-1)}} \frac{1}{a_\mu} \frac{1}{(a_\tau \tau_j + a_\mu^{-1})^{(a_\tau+1)}} \tag{20} \\ &\times \frac{b_0^{a_0}}{\Gamma(a_0)} \frac{1}{(2\pi)^{(k+2)/2}} \frac{\Gamma(a_0 + \frac{k}{2} + 1)}{\left(b_0 + \frac{1}{2} \sum_{j=0}^k a_j^2\right)^{a_0+k/2+1}}, \end{aligned}$$

where $s = 1.0, a_\tau = 1.0, a_\mu = 0.01, a_0 = 1.0, b_0 = 1.0$ were used throughout this work. Additionally, a *Inv - Gamma*(a, b) prior was chosen for the variance σ^2 of the Gaussian noise in (14). The hyperparameters $a = 2$ and $b = 10^{-6}$ were used throughout this work. As additive Gaussian noise was assumed, the likelihood of the parameters $p((x_{1:n}, y_{1:n}) | \theta, \sigma)$ for n training samples $\{(x_j, y_j)\}_{j=1}^n$ is given by

$$\begin{aligned} p((x_{1:n}, y_{1:n}) | \theta, \sigma) &= \prod_{i=1}^n p((x_i, y_i) | \theta, \sigma) \\ &= \frac{1}{(2\pi)^{n/2}} \frac{1}{\sigma^n} \exp\left\{-\frac{1}{2\sigma^2} \sum_{i=1}^n (y_i - f(x_i, \theta))^2\right\}. \end{aligned} \tag{21}$$

One major advantage of this approach is that the complexity of the Bayesian regression model, that is the number of Gaussian kernels k in (15), is also a model parameter that is inferred from the data. Hence, also highly irregular or non-linear interrelations between x and y can be captured. Moreover, as the chosen Poisson prior for k ensures that sparse regression models are favored, over-fitting is avoided and the simplest regression model explaining the data is chosen (Koutsourelakis 2009).

3.1.3 Computing the posterior

After generating the training samples, Bayes' rule is applied in an iterative manner inserting one pair of training samples (x_j, y_j) at a time, using the resulting posterior distribution as a prior for the next pair. Naturally, as more training samples are used, the influence of the prior diminishes. The posterior density $\pi_{\theta, \sigma^2}(\theta, \sigma^2)$ cannot be obtained as an analytic expression; hence, it is approximated using $N_{\text{particles}}$ discrete particles equipped with a weight proportional to their probability.

$$\pi_{\theta, \sigma^2}(\theta, \sigma^2) \approx \sum_{l=1}^{N_{\text{particles}}} W^l \delta_{(\theta, \sigma^2)^l}(\theta, \sigma^2) \tag{22}$$

The position and weights of these particles can be computed using , e.g., advanced SMC schemes (Doucet and de Freitas 2001; Kemp 2003). In this work, we use the SMC algorithm proposed in (Koutsourelakis 2009).

3.1.4 Computing solution statistics of the high fidelity model

Having computed a particulate approximation of the posterior density $\pi_{\theta, \sigma^2}(\theta, \sigma^2)$, we can evaluate the conditional distribution $p(y|x)$. With the established quantitative probabilistic relationship between x and y , the computation of

an estimate of $\pi_y(y)$ or $Pr[y \in \mathcal{A}]$ is now straightforward. A frequently used point estimate in Bayesian models is the posterior mean. The posterior mean $\hat{\pi}_y(y)$ can be computed as:

$$\begin{aligned} \hat{\pi}_y(y) &= \mathbb{E}_{\theta, \sigma, x}[\pi_y(y)] \\ &= \int p(y|x, \theta, \sigma) \pi_x(x) \pi_{\theta, \sigma^2}(\theta, \sigma^2) d\theta d\sigma^2 dx. \end{aligned} \tag{23}$$

Using particulate approximations of $\pi_x(x)$ and the joint posterior density $\pi_{\theta, \sigma^2}(\theta, \sigma^2)$

$$\begin{aligned} \pi_x(x) &\approx \sum_{i=1}^{N_{SAM}} W^i \delta_{x^{(i)}}(x), \\ \pi_{\theta}(\theta, \sigma^{-2}) &\approx \sum_{l=1}^{N_{particles}} W^l \delta_{(\theta, \sigma^l)}(\theta, \sigma^2), \end{aligned} \tag{24}$$

the posterior mean $\hat{\pi}_y(y)$ can be readily approximated with

$$\hat{\pi}_y(y) \approx \sum_l^{N_{particles}} \sum_i^{N_{SAM}} W^l W^i \phi(y - f(x_i; \theta_l), \sigma_l), \tag{25}$$

where ϕ denotes the Gaussian PDF as in (18). Obviously, we can also evaluate the probability of y exceeding a specific threshold y_0 . Given the model parameters θ and the variance σ^2 the probability of y exceeding a specific threshold y_0 can be computed by

$$Pr[y < y_0; \theta, \sigma] = \int q_{\mathcal{A}}(x; \theta, \sigma) \pi_x(x) dx \tag{26}$$

with $q_{\mathcal{A}}(x; \theta, \sigma)$ being defined as

$$\begin{aligned} q_{\mathcal{A}}(x; \theta, \sigma) &= \int_{y_0}^{\infty} p(y|x, \theta, \sigma) dy \\ &= \Phi\left(\frac{f(x; \theta) - y_0}{\sigma}\right), \end{aligned} \tag{27}$$

where $\Phi(x)$ is the standard normal CDF. The posterior mean approximation of $q_{\mathcal{A}}$, computed using the particulate approximation of $\pi_{\theta, \sigma^2}(\theta, \sigma^2)$, can be used to produce an estimate of the exceedance/failure probability for a specific value of x

$$\begin{aligned} \hat{q}_{\mathcal{A}}(x) &= \mathbb{E}_{\theta, \sigma}[q_{\mathcal{A}}] \\ &\approx \sum_{l=1}^{N_{particles}} W^l \Phi\left(\frac{f(x; \theta_l) - y_0}{\sigma_l}\right). \end{aligned} \tag{28}$$

Substituting into (26) and using (24) yields an estimate for the failure probability in the sense that on average, given the

training samples, the probability of y exceeding a threshold y_0 is given by:

$$\mathbb{E}_{\theta, \sigma}[Pr[y < y_0]] \approx \sum_{i=1}^{N_{SAM}} W^i \hat{q}_{\mathcal{A}}(x_i) \tag{29}$$

Another advantage of the Bayesian approach is that in addition to point estimates like the posterior mean, confidence intervals can be estimated from the posterior, quantifying the uncertainties in the inferred parameters and hence in the regression model. P-quantiles $q_{\mathcal{A}, p}(x)$ for the failure probability can be readily estimated using:

$$\begin{aligned} Pr[q_{\mathcal{A}}(x) \leq q_{\mathcal{A}, p}(x)] \\ \approx \sum_{i=1}^{N_{particles}} W^i H\left(q_{\mathcal{A}}^i(x) - q_{\mathcal{A}, p}(x)\right) = p, \end{aligned} \tag{30}$$

where H denotes the heavy side function. We use the $p = 1$ and $p = 99\%$ quantiles $q_{\mathcal{A}, 0.01}$ and $q_{\mathcal{A}, 0.99}$ to provide confidence intervals for exceedance probabilities in the following sections.

4 Patient-specific examples

It is important to emphasize that the proposed approach is very general and can be readily applied to a wide range of problems. To demonstrate the capabilities of our approach, we perform UQ using two realistic AAA models with patient-specific geometries from our database, referred to as male67 and male71, as example for a large-scale nonlinear solid mechanics problem with uncertain constitutive properties. The geometry of the lumen and, if present, the intraluminal thrombus (ILT) were reconstructed from CT data. The model male71 exhibits ILT while male67 does not. The wall thickness cannot not be determined from the CT data and was assumed to be a constant 1.57 mm throughout the model. While this assumption is made here to restrict ourselves to uncertain material parameters for ease of discussion, we stress that the presented method can also be applied to consider other uncertainties, like uncertain regional variations in wall thickness. To demonstrate this capability, we consider a simpler generic example in a proof of concept study in the next section. In this third example, we study an uncertain constitutive property together with an uncertain wall thickness, thereby both of those quantities are modeled as random fields.

For the patient-specific examples, we use our existing workflow, which is described in detail for the deterministic case in (Reeps et al. 2010; Maier et al. 2010) and obtain high fidelity finite element models with a mesh size of roughly

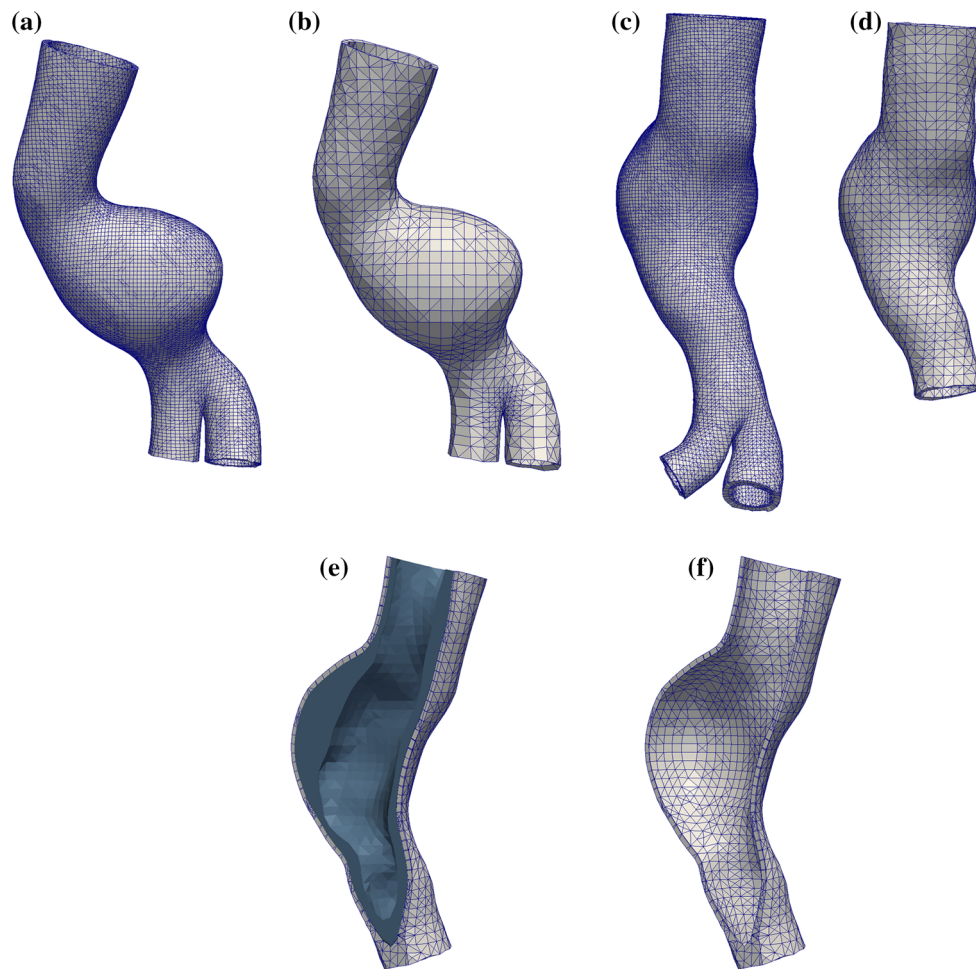


Fig. 3 AAA finite element models based on patient-specific geometries. **a** Male67 high fidelity model. **b** Male67 approximate model. **c** Male71 high fidelity model. **d** Male71 approximate model. **e** Cross-

sectional view of male71 approximate model with ILT. **f** Cross-sectional view of male71 approximate model without ILT

1 mm, which are shown in Fig. 3a, c. These hybrid finite element discretizations consist of 15,228 (male67) and 169,791 (male71) linear hexahedra-, wedge- and tetrahedra-shaped elements, resulting in a problem size of 61,674 and 292,044 degrees of freedom, respectively. For the aneurysm male67, an approximate model which is shown in Fig. 3b is constructed by coarsening the discretization. Using an element size of approximately 3 mm, the resulting coarser model has 2,670 elements and 8,679 degrees of freedom. For the aneurysm male71, two approximate models were created. For the first approximate model, a coarser discretization and a truncated geometry were used as depicted in Fig. 3d, e, respectively, which results in 19,579 elements and 25,875 degrees of freedom. Secondly, to further reduce the computational effort, another approximate model was created using model reduction in the sense that the ILT was omitted entirely in addition to a coarser discretization and geometric truncation, see Fig. 3d, f, respectively. Therein, the size reduces to

3,164 elements and 10,320 degrees of freedom. The reduction in model size and complexity yields a tremendous reduction in computational costs. For the aneurysms male67 and male71, coarsening of the discretization and model reduction yield an approximate finite element model that is between 10 and 50 times cheaper than the original high fidelity model, respectively. The CPU time required to compute one sample of the different models along with the numbers of degrees of freedom is summarized in Table 2. We used a 6 core Intel Xeon 3.2 GHz workstation with 12 GB memory to establish the solution times of the different models. The simulations were computed on one core, except the high fidelity model of the male71 aneurysm, for the solution of which four cores were used. We applied simple clamped boundary conditions at all in- and outlets of the aneurysms and imposed an orthogonal pressure follower load, to mimic the load imposed by the blood pressure on the luminal surface of the ILT. If no ILT is present, pressure is applied on the lumi-

Table 2 Size of studied finite element models and solution times for one sample

Model	Solution time (CPU seconds)	Degrees of freedom
Male67 high fidelity	2,000	61,674
Male67 approximate	195	8,679
Male71 high fidelity	10,800	292,044
Male71 approximate with ILT	400	25,875
Male71 approximate without ILT	215	10,320

nal surface of the vessel wall. The shear stresses, induced by the blood flow, are negligible as compared to the pressure load. Hence, a full fluid–structure interaction simulation is not required. Since the geometries of the aneurysms were obtained from CT images, the in vivo imaged configuration is not stress free, but represents a loaded spatial configuration. In order to obtain meaningful results, this “prestressed” state has to be accounted for. To imprint the stresses from an assumed diastolic pressure of 87 mmHg in the in vivo imaged state, we used a modified updated Lagrangian formulation (MULF) (Gee et al. 2009, 2010). The obtained prestrained/prestressed state implicitly defines a stress-free reference geometry, while neglecting growth and remodeling processes in the wall. This implicitly defined geometry of course depends on the constitutive properties of the arterial wall and the ILT and hence varies from patient to patient and from sample to sample in an UQ analysis.

After the prestressing phase, the luminal pressure is increased to 121 mmHg, taking into account the imprinted stresses and strains acting on the imaged configuration. For the constitutive behavior of the ILT in the male71 aneurysm, the model proposed by Gasser et al. (2008) was chosen, using the parameters provided in his work. Furthermore, in all cases, the stochastic material law described above in (7) is used to model the aneurysmatic arterial wall.

To study the sensitivity with respect to the unknown correlation length d from (3), we investigated three random field models with different correlation lengths $d = 12.5$, $d = 25$ and $d = 50$ mm, respectively. The chosen range covers random fields which exhibit rather short-scale fluctuations as well as highly correlated random fields with a correlation length in the range of the diameter of a typical aneurysm. All finite element calculations were done using our in-house research code BACI (Wall and Gee 2014). In order to obtain a reference solution for comparison with the proposed approach, we performed direct MC with all high fidelity finite element models for all three correlation lengths with a sample size of $N_{\text{SAM}} = 50,000$. First, the impact of the uncertain constitutive parameter β on the relevant mechanical quantities such as Cauchy von Mises stress (σ_{vM}), Euler-

Almansi von Mises strain (e_{vM}), and strain energy Ψ is discussed on the basis of these MC simulations, and the influence of the correlation length on the distributions of these quantities is investigated. Then, in Sect. 4.2, we will compare the results of the employed Bayesian multi-fidelity MC approach to the MC reference solution and the accuracy and efficiency of the proposed multi-fidelity approach is examined.

4.1 Monte Carlo reference solution

As expected, in spite of drastic variations in the wall stiffness, the overall spatial pattern of the von Mises stress is still largely determined by the geometry of the aneurysm and rather insensitive to those variations. However, a slight dependency on the correlation length was found, where a smaller correlation length leads to greater variations in the overall stress pattern. Figure 4 shows the von Mises stress and strain resulting from five different showcase realizations of the stochastic constitutive law for the male71 aneurysm. The samples in Fig. 4 were computed with a correlation length of $d = 12.5$ mm for which the resulting stress pattern exhibits the largest variations. Nevertheless, the overall spatial pattern of the stress remains remarkably similar for all depicted realizations of the random field. Moreover, when the correlation length of the random field is increased, the variations in the spatial stress pattern become even less significant. This is in agreement with theoretical mechanical considerations, as the stress state in the aneurysm wall is dominated by in plane membrane stress, which is determined mostly by the overall geometry of the aneurysm and the traction boundary conditions.

In contrast to the spatial stress pattern, the spatial strain pattern of the samples very significantly indicate a strong dependency between local stiffness of the wall and local strain. Unlike in the stress pattern, regions with high strain are not determined by the overall shape or the geometry of the aneurysm but rather by regions with low wall stiffness as shown in Fig. 4c. The male67 aneurysm shows similar behavior and is hence not shown here.

In addition to this qualitative assessment, we examined the empirical probability distributions of the stress obtained from MC at several locations across the aneurysms, see Fig. 5. The mechanical quantities of interest were evaluated at single elements at these specific locations. Figure 6 shows the empirical probability densities of the von Mises stress, von Mises strain and strain energy at two locations on the male67 aneurysm. Location 1, shown in Fig. 5a, is located in the neck of the aneurysm above the sack. Location 2 is at the right lateral, distal end of the AAA sack, see Fig. 5a. Both locations are in regions that exhibit high wall stress. For the other aneurysm male71, Fig. 7 shows the probability densities evaluated at location 3, located at the center of the dorsal


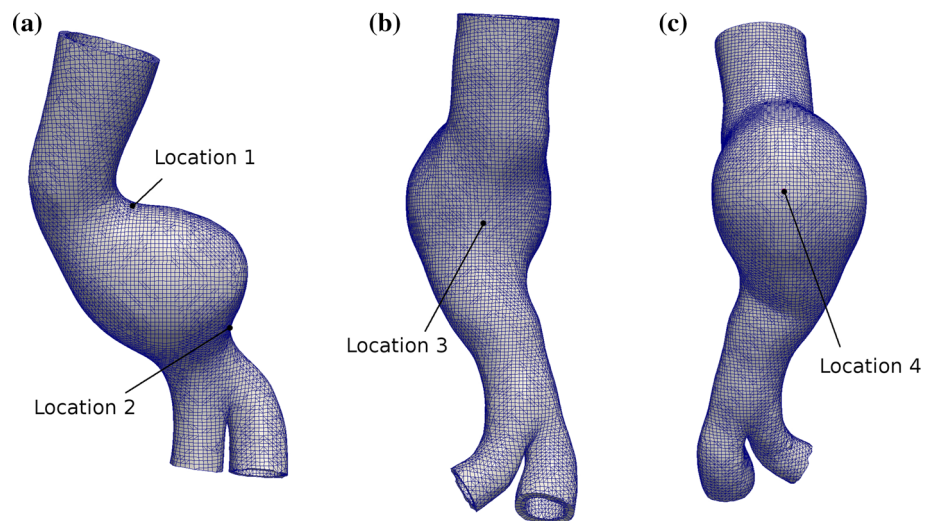


Fig. 4 Simulation results of male71 aneurysm for five realizations of the random field using a correlation length of $d = 12.5$ mm. Realizations of random field (a), and resulting von Mises stresses (MPa) (b) and von Mises strains (c)

Fig. 5 Locations in AAA wall at which mechanical quantities are evaluated. **a** Male67 model anterior view. **b** Male71 model posterior view. **c** Male71 model anterior view



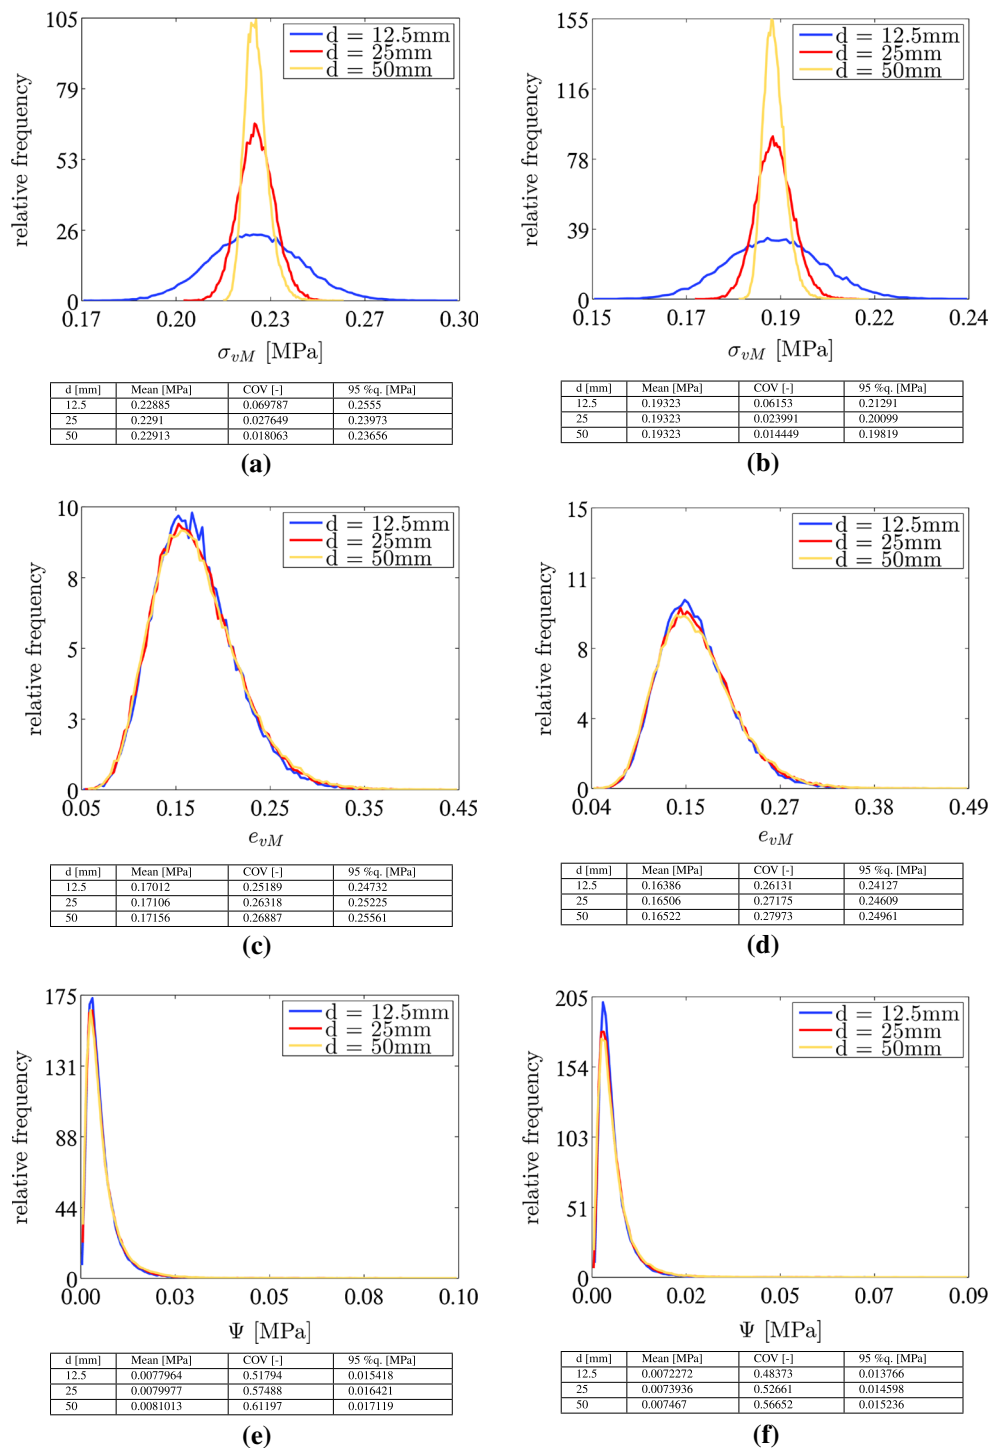


Fig. 6 Monte Carlo reference solution for male67 aneurysm. **a** Von Mises stress at location 1. **b** Von Mises stress at location 2. **c** Von Mises strain at location 1. **d** Von Mises strain at location 2. **e** Strain energy at location 1. **f** Strain energy at location 2

side of the aneurysm sack in the high stress region, and at location 4, located at the ventral side of the aneurysm sack in a low stress region with thick ILT underneath the AAA wall, see Fig. 5b, c. All empirical densities are plotted for the different correlation lengths of the random field in the sto-

chastic constitutive model. In addition to the mean and COV, the 95 % quantiles are provided for each of the distributions in Figs. 6 and 7, respectively.

The obtained distributions for the stresses confirm the qualitative assessment that the stresses are only mildly

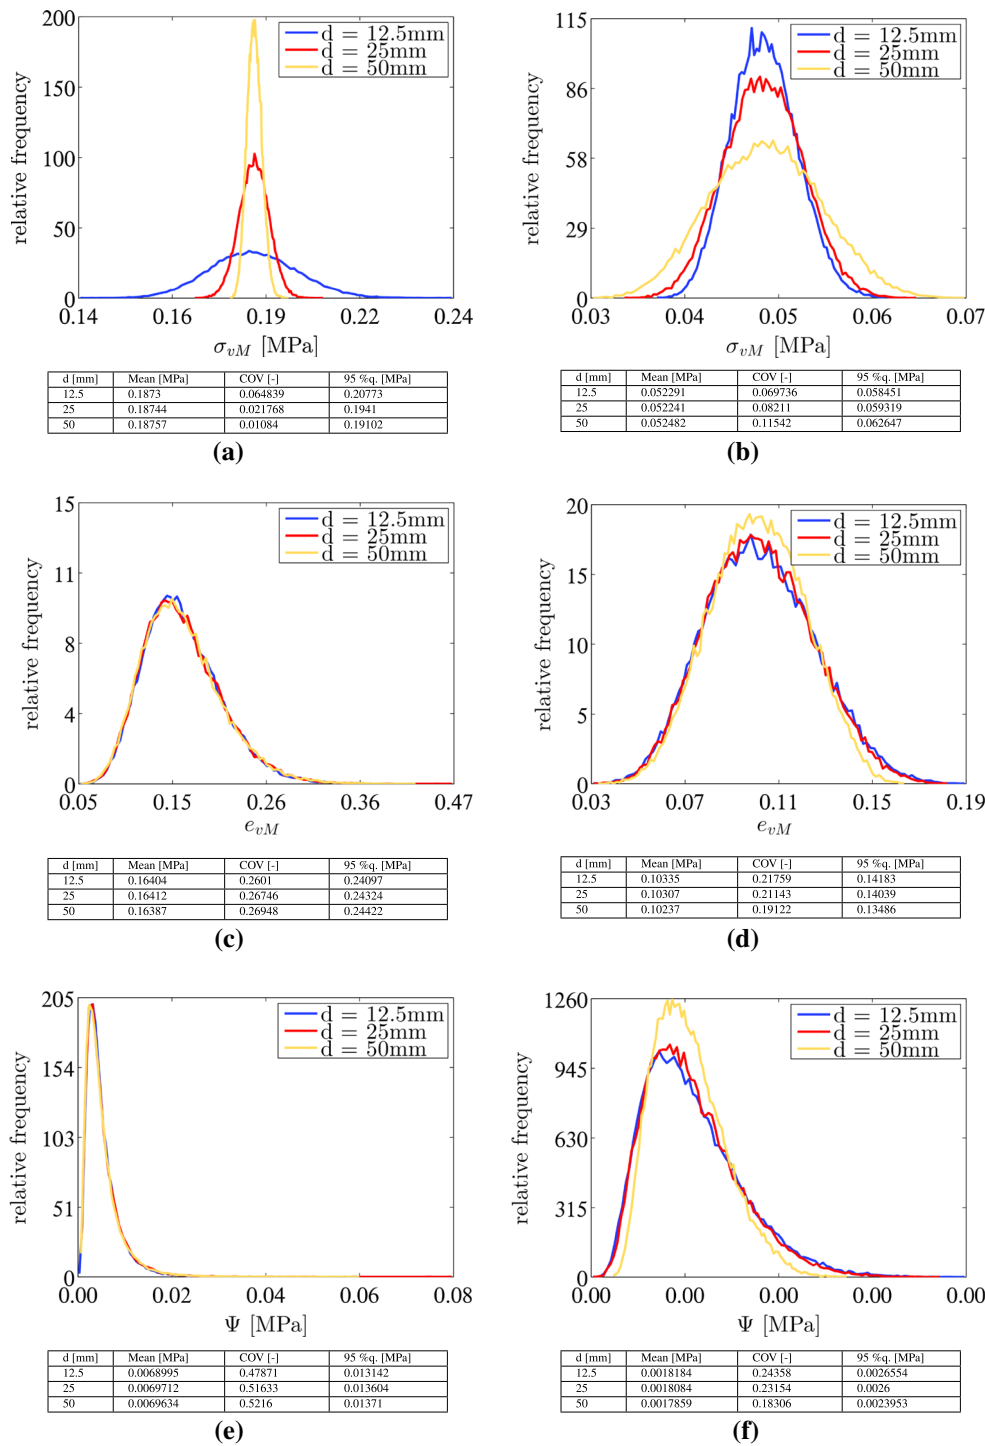


Fig. 7 Monte Carlo reference solutions for male71 aneurysm. **a** Von Mises stress at location 3. **b** Von Mises stress at location 4. **c** Von Mises strain at location 3. **d** Von Mises strain at location 4. **e** Strain energy at location 3. **f** Strain energy at location 4

affected by the local stiffness of the wall. The COVs of the stress distributions are below 0.12. If only the locations 1, 2, and 3, which all lie in high stress regions of the respective AAAs, are considered, the maximum COV of the stress distribution is 0.07. Also for these three elements, the COVs

of the stress become even smaller with increasing correlation length of the random field. In contrast to this, the COV of the stress at location 4, which lies in a region with thick ILT and low wall stress, is less influenced by the correlation length. The 95 % quantile of the stress distributions, as a

“worst-case” estimate, is approximately 2–19 % higher than the respective mean value, see Figs. 6a, b and 7a, b. This provides further evidence that uncertainties in the constitutive parameter β result in only moderate uncertainty in the computed stresses.

While this appears to be consistent with the findings from Raghavan and Vorp (2000), it should be stressed that Raghavan and Vorp only considered mild parameter variations within the 95 % confidence interval of their estimated mean value and not within the complete range of measured values. Furthermore, they did not consider spatial intra-patient variations of the constitutive parameters nor did they consider the prestressed state of the aneurysm or ILT. Findings similar to ours have also been reported for intracranial aneurysms by Ma et al. (2007) and Miller and Lu (2013). These two studies also indicate only a mild dependence of the wall stresses on the constitutive parameters; however, only spatially homogeneous parameters were considered therein as well.

Our results show that the wall stresses are only mildly sensitive to variations in the constitutive parameter β . This holds especially for the rather long-scale spatial heterogeneity obtained with the longer correlation lengths of the random field, see Fig. 6a. The sensitivity of the stress state on the material properties typically increases as the random field exhibits strong spatial gradients in the parameter field as one moves to smaller correlation lengths, indicating that only strongly localized variations in this parameter have a significant effect on the stresses. We argue that the reason for this insensitivity, in spite of strongly localized variations in β , can be ascribed to the prestressed state of the imaged geometry. The prestress in the in vivo imaged configuration of cardiovascular structures has to be accounted for in order to obtain meaningful simulation results, as has been pointed out in numerous publications (Gee et al. 2009, 2010; Speelman et al. 2009; Miller and Lu 2013). All techniques to account for this essentially approximate a stress state that equilibrates with the external load given the spatial configuration. The boundary conditions in a typical cardiovascular problem resemble those termed statically determinate. Dominated by traction boundary condition emulating luminal blood pressure, the Dirichlet boundary conditions applied far away from the region of interest have only minor influence. Thus, in prestressed cardiovascular structures which are predominantly loaded by traction boundary conditions, the overall stress state, which is largely determined by the prestress, is governed by simple equilibrium of internal stress with the external loading and therefore is relatively insensitive to variations in the constitutive parameters.

In contrast to the stresses, the strain and strain energy exhibit large COVs and are therefore drastically affected by the constitutive parameter β . The probability distributions depicted in the Figs. 6c–f and 7c–f reveal the large variations in local strain and strain energy. With COVs up to 0.28 and

0.6, respectively, the strains and strain energy exhibit significantly larger COVs than the stresses, rendering a statement about the true strain state of the aneurysm difficult when facing uncertain constitutive parameters. Whereas the distribution of the stresses are symmetric, both strains and strain energy exhibit skewed distributions. The contour plots in Fig. 4c reveal a close dependence between low β value and high local strain and vice versa, and the overall spatial pattern of the strain depends mostly on the realization of the random field and not on the geometry of the aneurysm. Since the local stress state is mostly dictated by the given spatial configuration and external load, local strains depend predominantly on the local stiffness. The softer the wall, the higher the strains need to be in order to reach a certain “predetermined” stress level.

Comparing the probability densities of the strain and strain energy at different correlation lengths in Figs. 6c–f and 7c–f, respectively, there is no noticeable difference between the three assumed correlation lengths. The variance of these probability distributions is virtually independent from the chosen correlation length for both, strains and strain energy. If the 95 % quantiles are used as an estimate for a worst-case scenario and compared to the mean values, we see that strain is potentially 60–70 % higher than the mean value and thus could be dramatically underestimated in deterministic models with population-averaged constitutive parameters. In case of the strain energy, the margin of uncertainty is even larger with a COV of typically more than 0.5, which results in a 95 % quantile that is more than twice as high as the mean value.

The explanation for these very large COVs is once again the non-stress-free in vivo imaged configuration of the aneurysm. The combination of a known deformed configuration and the load case leading to this configuration results in a computed stress state, which is only mildly sensitive to variations and uncertainties in the constitutive parameter β . In a sense not knowing the stress-free reference configuration but the deformed configuration is an advantage for the prediction of wall stress. This is not the case for the computation of the strain state. With a stress state that is to a large extent determined by the geometry of the aneurysm, the strains become very sensitive to the local stiffness of the AAA wall. Assuming that the local stress state is dictated by the geometry and hence more or less fixed, the strain state that corresponds to this stress state depends on β . Low local stiffness will result in very high local strains and vice versa. Thus, any uncertainty in β directly translates to uncertainty in the strain state, inhibiting an accurate prediction of the strains if the stress-free geometry and the constitutive properties are unknown or uncertain. The same argument holds if the strain energy is considered as quantity of interest. As a result, the probability distributions of the strains and strain energy evaluated at specific locations in the wall, exhibit very high COVs.

4.2 UQ using Bayesian multi-fidelity MC

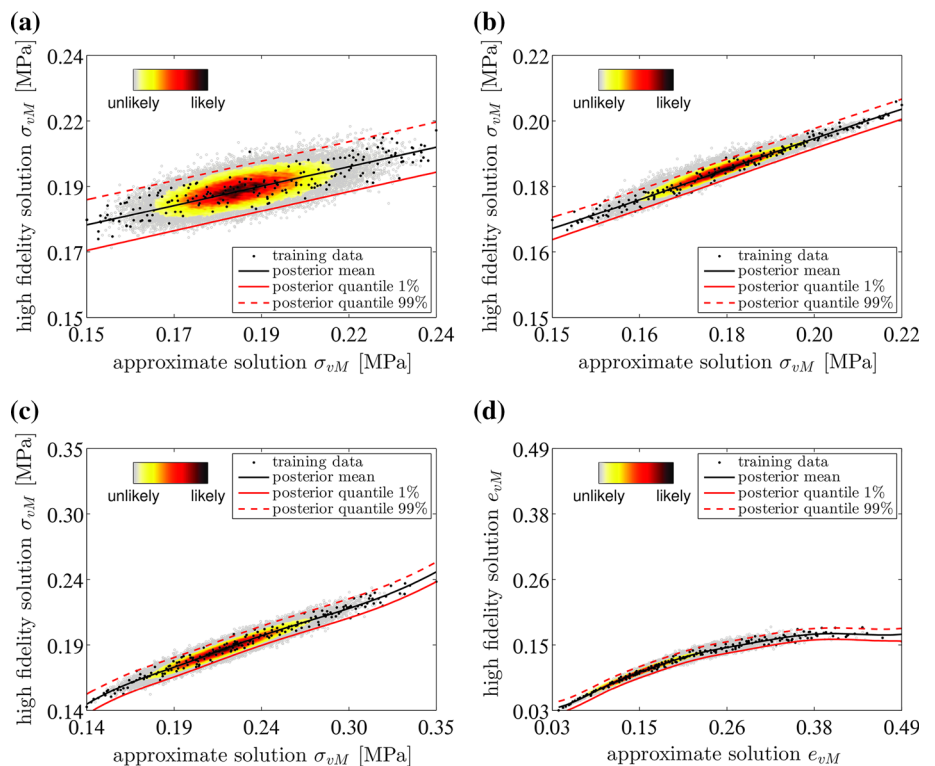
Of course, conducting direct MC for large patient cohorts in a clinical setting is impractical. Hence, we will demonstrate the capabilities of the aforementioned UQ approach, which reduces the computational cost to an acceptable level for studies with a larger patient cohort. To demonstrate the efficiency and accuracy of the proposed framework, we chose a subset of the results from the previous chapter. Four examples which are summarized in Table 3 are considered, examining von Mises stress and strain at distinct locations of the two AAA models as quantities of interest. Moreover, we take into account different correlation lengths of the random field as well as different approximation schemes for the approximate models. After computing 50,000 MC samples using the approximate models, 200 of those were selected and the corresponding high fidelity samples were taken from the MC reference solution. The 200 samples were selected such that they evenly cover the entire support of the particulate approximation of π_x . This can be achieved by sorting all samples into equally wide bins and then selecting a fixed number of

particles from each bin. Then, this set of training samples was used to determine the posterior density $\pi_{\theta, \sigma^2}(\theta, \sigma^2)$ and subsequently to compute the conditional probability distribution $p(y|x)$. Figure 8 shows a comparison between the solution of the approximate model and the high fidelity solution for all four examples. The color-coded 2D histograms, based on all 50,000 MC samples, show the interrelation between the quantity of interest computed on the approximate and the high fidelity model, respectively. The training samples that were used to infer the parameters of the regression model are shown in Fig. 8 as black dots. The obtained posterior mean as well as the 1 and 99% quantiles of $p(y|x)$ are also depicted in Fig. 8. In all four cases, the posterior mean of $p(y|x)$ captures the interrelation between the approximate and high fidelity solution very well. Furthermore, the computed quantiles readily provide confidence intervals, which contain virtually all MC samples with very few exceptions. The reason for the noisy relationship between x and y is that the coarsening of the discretization yields larger discretization errors, and therefore, spurious effects due to the coarser discretization appear. Moreover, the coarser mesh is unable

Table 3 Examples for Bayesian multi-fidelity MC approach

Example	Patient	Quantity	Location	d (mm)	Approximation scheme
1	Male67	σ_{vM}	2	25	Coarsening
2	Male71	σ_{vM}	3	25	Coarsening
3	Male71	σ_{vM}	3	12.5	Coarsening no ILT
4	Male71	e_{vM}	4	25	Coarsening no ILT

Fig. 8 Comparison between approximate solution and high fidelity solution. In addition to the posterior mean and the 1 and 99% quantiles of $p(y|x)$, the figures show the used training samples (black dots) as well as all 50,000 samples as color-coded 2d histogram. **a** Approximate versus high fidelity solution, male67 aneurysm, $d = 25$ mm, σ_{vM} at location 2. **b** Approximate versus high fidelity solution, male71 aneurysm, $d = 25$ mm, σ_{vM} at location 3. **c** Approximate versus high fidelity solution, male71 aneurysm, $d = 12.5$ mm, σ_{vM} at location 3. **d** Approximate versus high fidelity solution, male71 aneurysm, $d = 25$ mm, e_{vM} at location 4



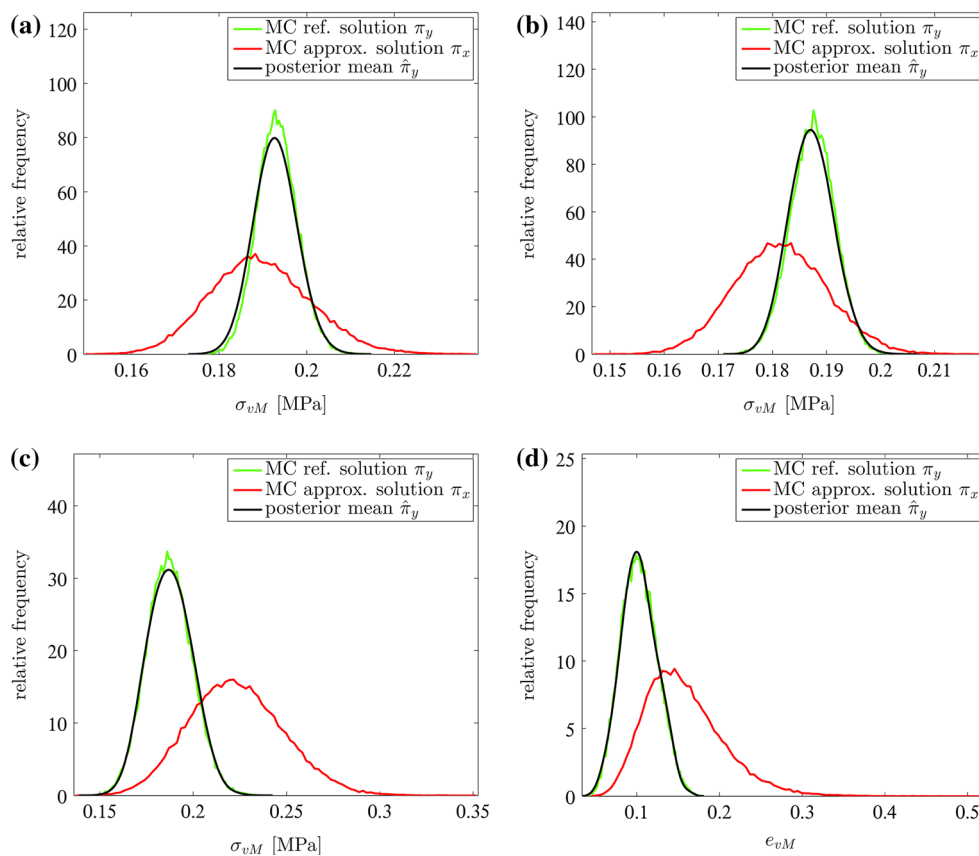


Fig. 9 Exemplary comparison of empirical densities from approximate model $\pi_x(x)$, high fidelity model $\pi_y(y)$ and posterior mean approximation $\hat{\pi}_y(y)$. **a** Example 1: male67, $d = 25$, σ_{vM} at loca-

tion 2. **b** Example 2: male71, $d = 25$, σ_{vM} at location 3. **c** Example 3: male71 $d = 12.5$, σ_{vM} at location 3. **d** Example 4: male71 $d = 25$, e_{vM} at location 4

to resolve the finer details of the random field, which results in a smeared less detailed solution. We found that these effects are more distinct if the quantity of interest is the von Mises stress as compared to von Mises strain.

While the first two examples show an almost linear relation between approximate and high fidelity solution, example three and particularly example four exhibit a distinct nonlinear dependency. The Bayesian regression model is able to capture these nonlinear interrelations as well, without any modification of the formulation. It also can be seen in Fig. 8, especially in the latter two examples, that the lack of ILT in the approximate model leads to significant differences in the overall magnitude between the stresses computed with approximate and the high fidelity model, respectively. In fact, the stress and strain are considerably higher in the approximate model, due to the lack of ILT. By comparing the two densities $\pi_x(x)$ and $\pi_y(y)$, as shown in Fig. 9, this difference becomes even more apparent.

The posterior mean approximation $\hat{\pi}_y(y)$ is computed for all examples using (23) and depicted in Fig. 9 together with $\pi_x(x)$ and the MC reference solutions of the high fidelity

model $\pi_y(y)$. Despite the differences between $\pi_x(x)$ and $\pi_y(y)$, the posterior mean approximation $\hat{\pi}_y(y)$ is in excellent agreement with the MC reference solution for all examples.

Table 4 lists the mean, COV and 95% quantile for both, the MC reference $\pi_y(y)$ and the posterior mean approximation $\hat{\pi}_y(y)$. Additionally, the relative error between the MC reference and the Bayesian approach is given, thus allowing a quantitative assessment of the accuracy. For the mean value, a relative error below 1% is achieved for all examples. The error of the estimation of the COV and the 95% quantile is slightly higher. However, considering the crude approximate models and the gain in computational efficiency, the accuracy is more than sufficient for biomedical applications.

The ability to obtain an accurate estimate for the complete PDF, rather than just the first moments is a major advantage, because we can compute the probability that the quantity of interest exceeds a certain threshold y_0 . Although the exact definition of a failure threshold can be difficult in biomechanical systems, such a failure probability is a valuable tool in many applications including rupture risk stratification of

Table 4 Comparison between MC reference solution and Bayesian multi-fidelity MC approach

	MC reference			Posterior mean			Relative error (%)		
	Mean	COV	95 % q.	Mean	COV	95 % q.	Mean	COV	95 % q.
Ex. 1	0.1932	0.0239	0.1983	0.1928	0.0259	0.2012	0.21	8.26	1.45
Ex. 2	0.1874	0.0217	0.1939	0.1872	0.0223	0.1940	0.12	2.70	0.05
Ex. 3	0.1872	0.0648	0.2073	0.1871	0.0653	0.2080	0.10	0.86	0.34
Ex. 4	0.1030	0.2114	0.1430	0.1023	0.2139	0.1395	0.71	1.20	2.44




Fig. 10 Probability of exceeding a local failure threshold y_0 for a given quantity of interest. Comparison of MC reference solution to posterior mean approximation and posterior quantiles. **a** Example 1: male67,

$d = 25$, σ_{vM} at location 2. **b** Example 2: male71, $d = 25$, σ_{vM} at location 3. **c** Example 3: male71 $d = 12.5$, σ_{vM} at location 3. **d** Example 4: male71 $d = 25$, e_{vM} at location 4

AAAs. Under the assumption that a suitable failure threshold y_0 can be defined, the probability that a threshold y_0 is exceeded is computed using (29). Figure 10 depicts the posterior mean failure probability computed for a range of failure thresholds y_0 for the chosen examples. Again an excellent agreement with the MC reference solution is achieved, at a fraction of the computational cost, see Table 5. In addition, credible intervals which are also shown in Fig. 10 can be computed using (30). The intervals are an indicator for the

accuracy of the computed failure probability estimate based on the training data. As depicted in Fig. 10, the provided bounds contain the MC reference solution.

The computational efficiency of the method depends on the ability to create cheap approximate models, which still provide “enough” relevant information. In our examples, the ratio of computational effort required to compute one sample of the accurate and the approximate model, respectively, is between 10.2 and 50. Whereas Koutsourelakis (2009) reports

Table 5 Comparison of computational costs between direct MC and Bayesian multi-fidelity MC approach

Example	Patient	Approximation scheme	Cost MC	Cost multi-fidelity MC
1	Male67	Coarsening	50,000	5,075
2	Male71	Coarsening	50,000	2,051
3	Male71	Coarsening no ILT	50,000	1,195
4	Male71	Coarsening no ILT	50,000	1,195

Costs are given as equivalent number of calls to the respective high fidelity model

ratios of over 1,000, we believe that in general biomedical engineering problems such ratios can usually not be achieved. However, a ratio of 50, accomplished by using a coarser discretization and omitting the ILT in the approximate model, results in a tremendous reduction of the computational effort, see Table 5. Compared to direct MC using the high fidelity model, the total computational effort is reduced by a factor of 42. As the male67 patient does not exhibit ILT and merely a coarser discretization was used to create the approximate model, the computational cost for the male67 aneurysm is reduced by a factor of 10 compared to direct MC. However, the vast majority of AAAs in our patient database exhibit ILT and especially for large AAA, which are computationally expensive, being able to omit the ILT in the model will result in even greater savings. However, we would like to point out that omitting the ILT in the approximate model is obviously only a viable option as long as no uncertainties in the ILT properties itself are considered in the UQ analysis. By using direct MC as a sampling strategy for the approximate model, as compared to importance sampling or SMC, we are able to access additional quantities of interest by computing few extra samples in the high fidelity model and reuse all samples of the approximate model. We would like to stress the fact that the approximate model does not have to be accurate, as in deterministic multi-level/grid schemes. As shown in Fig. 9, approximate models without ILT yield significantly higher values for strain and stress due to the missing support of ILT for the AAA wall. Additionally, no restrictions to the interrelation between the approximate and high fidelity model apply. Even highly nonlinear interrelations are detected and accurately reproduced by the Bayesian regression model, the complexity of which is determined by the information provided by the training samples. Whereas inaccurate approximate models and nonlinear interrelations do not adversely affect the accuracy of the approach per se, coarsening of the discretization yields less detailed smeared out results due to the coarser representation of the random field, a larger discretization error, and potentially aggravates spurious numerical effects such as volumetric locking. This results in a noisier relationship between the approximate and the accurate model, i.e., a higher variance of the conditional probability distribution $p(y|x)$. As the variance increases, the credible intervals will become larger too. Furthermore,

the accurate prediction of the quantity of interest on the high fidelity model becomes more difficult, especially with regard to the tails of the distribution and the assessment of small failure probabilities.

While we used a set of 200 training samples and achieved an excellent agreement with the MC reference solution, the amount of training data needed depends on a number of factors. To be specific, the number of necessary training samples and hence runs of the high fidelity model, depend on the required accuracy, the employed approximate model and on the quantity of interest as well as the statistic to be estimated. Hence, it is difficult to provide an a priori estimates of how many samples are needed to achieve a given accuracy. However, the variance of the conditional distribution $p(x|y)$ and the confidence intervals computed using (30) provide us with a good indicator whether additional data will improve the accuracy. If the variance can no longer be decreased by adding more training data, the low fidelity model has to be changed if a more accurate assessment is required. We note that more elaborate schemes to select the training samples could further reduce the number of evaluations of the high fidelity model. This could be achieved for instance by exploiting the confidence intervals; in the sense that additional samples are selected in regions of x which contribute more to the uncertainty in the estimate of $Pr[y \in \mathcal{A}]$ and where these intervals are relatively large (Koutsourelakis 2009). However, this is beyond the scope of this work and subject of ongoing research.

4.3 Implications for computational rupture risk assessment of AAAs

Considering these results with regard to AAA rupture risk assessment using deterministic computational models, it seems that stress-based failure or rupture models (Vande Geest et al. 2006a; Gasser et al. 2010; Reeps et al. 2012) are well suited for rupture risk assessment, since wall stress is a rather insensitive measure with regard to constitutive uncertainties in the AAA wall. However, while we considered uncertainty in the constitutive parameter β , it is important to keep in mind that various other sources of uncertainty exist. Most importantly, uncertainties in the wall thickness are likely to have a profound affect on the stresses in the

arterial wall. While we considered only the uncertainty in one constitutive parameter, the methodology presented here can readily be extended to include other uncertain model parameters. This is demonstrated for a simple example in the following section. The application to realistic patient-specific geometries with multiple random fields is beyond the scope of this paper, but can be done with the proposed method, and we are planning to do this in the future. We believe that rupture potential stratification employing stress-based measures can be improved by rigorous UQ and that population-averaged constitutive parameters and deterministic models are not sufficient for the accurate patient-specific assessment of AAA rupture potential.

Furthermore, while many of the conducted studies assess rupture potential by computing wall stresses based on an elastic model, it has been pointed out previously that the damage and failure process of soft tissue is more involved and cannot be captured with purely elastic models employing a von Mises stress-based failure criterion (Marini et al. 2011; Rodriguez et al. 2008; Peña et al. 2008; Volokh and Vorp 2008; Volokh 2010; Gasser and Holzapfel 2002). In addition, our results indicate that the von Mises stresses, while being only mildly sensitive to variations in β , do not contain or reveal all information about the mechanical state of the tissue. With the stresses being determined mostly by the geometry of the aneurysm and the applied transmural pressure, the strains or strain energy is almost decoupled from the stresses as shown in Fig. 4. Advanced tissue damage and failure models, e.g., based on the evolution of the isochoric stored strain energy of the material, as proposed by Simo (1987) and applied to AAAs by Marini et al. (2011) require accurate prediction of the prevailing strain and strain energy. In lack of truly patient-specific parameters, detailed statistical and probabilistic models based on empirical data and rigorous UQ are the only way to obtain meaningful and reliable results.

However, it is well known that several growth and remodeling processes constantly rebuild and transform the tissue in the arterial wall; hence, cardiovascular systems cannot be considered to exhibit purely elastic behavior. Localized areas of high strains are likely to be remodeled such that elastic strain peaks in the tissue are reduced. Therefore, we believe that 95 % quantiles of both strains and strain energy are unrealistically high, as growth and remodeling processes would diminish such high strains.

To further reduce the uncertainty in the input parameters, the use of regression techniques to create statistical models which relate, e.g., the constitutive parameters to noninvasively measurable variables could help to reduce the uncertainty in patient-specific material properties or other model parameters. Since this empirical approach has been successfully applied to estimate the spatial distribution of AAA wall strength (Reeps et al. 2012; Vande Geest et al. 2006a), we

believe that this approach could help to refine the probabilistic description of other uncertain model parameters as well.

5 Proof of concept: multiple sources of uncertainties

In this section, we want to demonstrate with a simpler proof of concept example that the proposed approach can also be used to study geometric uncertainties such as an uncertain distribution of wall thickness. Furthermore, we show that problems with more than one uncertain physical parameter, all of which are modeled as random fields, can be tackled with the Bayesian multi-fidelity approach. In addition to the uncertain constitutive parameter β , we consider an uncertain wall thickness as well. Therefore, we exemplarily consider the model of a cylinder with a radius of 15 mm and a length of 200 mm as example that can be interpreted as a model of a generic artery. As with the patient-specific AAA models, we apply simple clamped boundary conditions at all in- and outlet of the artery segment and impose a traction orthogonal to the deformed configuration, to mimic a luminal blood pressure of 160 mmHg on the inner surface of the cylinder. To reduce the computational burden, we consider only the upper half of the cylinder and apply symmetry boundary conditions on the cut surfaces. Note that due to the asymmetric spatial distribution of the uncertain properties, the deformation of the full cylinder will in general not be symmetric and computing only half the model will introduce an error. However, as this example is shown here in a mere proof of concept fashion, we are willing to make this sacrifice for the benefit of shorter computation times. The stochastic constitutive law described in Sect. 2 is used as stochastic material model for the wall of the cylinder. The wall thickness of the cylinder is also modeled by a log-normal random field, realizations of which are computed using (5) and (6). The parameters of the marginal log-normal distribution of the field are set to $\mu_t = 0.447$ and $\sigma_t = 0.333$, respectively. These parameters result in a spatially varying wall thickness that covers roughly the range between 0.5 and 4 mm. In this example, we assume that the two fields describing the constitutive parameter β and the wall thickness, respectively, are not cross-correlated. As in the patient-specific examples, we create two models of different fidelity. A high fidelity model, a realization of which is shown in the left part of Fig. 11, with 24,000 Elements and 98,532 degrees of freedom, and a low fidelity approximate model, shown in right part of Fig. 11. The low fidelity model consists of 300 elements and 2,046 degrees of freedom. To further reduce the computational cost of the low fidelity model, only a reduced luminal pressure is applied. By applying only half the load, the number of necessary load steps in the nonlinear solution scheme is cut in half as well. The result is a low fidelity model which is approximately 120 times cheaper to evaluate than the corresponding high




Fig. 11 One realization of high fidelity (*left*) and low fidelity (*right*) finite element model of the generic artery example

Table 6 Generic artery example: size of studied finite element models in degrees of freedom (dof) and solution times for one sample

Model	Solution time (CPU seconds)	Dofs
High fidelity	4,800	98,532
Low fidelity	42	2,046

fidelity model. Of course, the reduced load will result in a dramatic underestimation of, e.g., the wall stress computed with the low fidelity model. The resulting CPU times required to compute one sample along with the numbers of degrees of freedom are summarized in Table 6 for the two models. Simulations of the low fidelity model were computed on one core, whereas for the high fidelity model, four cores were used. For the sake of brevity, we restrict ourselves to one variant of this example, i.e., we consider only one correlation length of the random fields ($d = 25$ mm) and we discuss the results of the Bayesian multi-fidelity approach based on a single quantity of interest. Here, we chose the von Mises stress in the wall of the cylinder at a specific location that is indicated in Fig. 11 as location 5 and will be referred to as such from here on. The overall procedure is the same as described for the patient-specific examples and follows the steps given in Table 1. Briefly, we compute a MC solution for the low and high fidelity model using 50,000 samples. The latter is only needed for comparison with the results of the Bayesian multi-fidelity approach. Then, as described previously, we use 200 samples from both models as training data to infer the conditional probability distribution $p(y|x)$, which then allows us to compute an approximation $\hat{\pi}_y(y)$ of the desired probability density using (23).

5.1 Results

The plots in Fig. 12 provide a summary of the results. In Fig. 12a, a comparison between the solution of the approximate model and the corresponding solution of the high fidelity model is shown. In this example, the relationship is very close to linear. Furthermore, it can be seen that the stress values computed by the low fidelity model are obviously considerably lower than the ones computed with the high

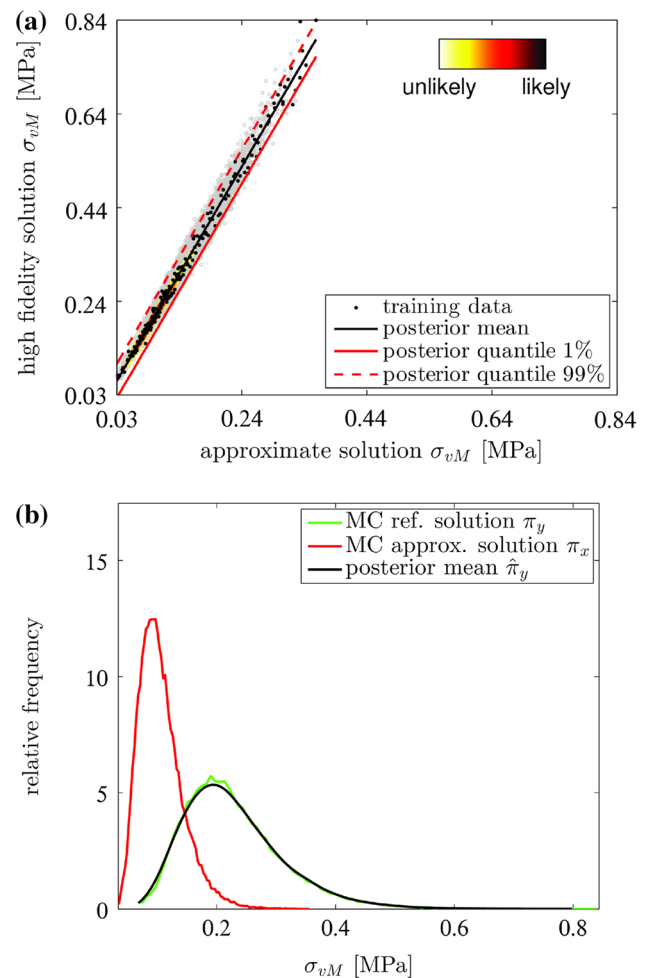


Fig. 12 Results for generic artery with uncertain constitutive parameter and wall thickness. **a** Approximate versus high fidelity solution, $d = 25$ mm, σ_{vM} at location 5. **b** Comparison of empirical densities from approximate model $\pi_x(x)$, high fidelity model $\pi_y(y)$ and posterior mean approximation $\hat{\pi}_y(y)$

fidelity model due to the reduced pressure, which is applied to the low fidelity model. This difference is also evident in the computed histograms, which are shown in Fig. 12b. The distribution of the stress computed with the coarse model is shifted to the right compared to the MC reference solution. However, the magnitude of the stress is mainly determined by the local wall thickness, and this dependency can also be captured by the coarser model with reduced load. Hence, the posterior mean approximation $\hat{\pi}_y(y)$, computed using the Bayesian multi-fidelity approach, matches the MC reference solution $\pi_y(y)$ very well. The relative errors in mean, COV, and 95% quantile are 0.059, 1.44 and 0.63%, respectively. Using the Bayesian multi-fidelity approach, the computational costs can be reduced to the equivalent of approximately 640 full evaluations of the high fidelity model.

As expected, the results show a high variability of the stress at the evaluated location (COV = 0.369) due to the

variable wall thickness. We believe that this will also be the case for patient-specific geometries of cardiovascular structures, and we plan to investigate more elaborate examples in the future using the proposed approach.

6 Conclusions

In this work, we presented a Bayesian multi-fidelity Monte Carlo (MC) framework for efficient uncertainty quantification (UQ) in large-scale nonlinear problems, with high stochastic dimension. We applied the approach to complex large-scale nonlinear biomechanical models with patient-specific geometries with an uncertain constitutive parameter. Furthermore, we demonstrated with a proof of concept example that uncertainties in wall thickness can be considered in addition to constitutive uncertainty. Thereby, the uncertain parameter are always modeled as three-dimensional non-Gaussian random field, leading to a high stochastic dimension of the problem. By combining multi-fidelity sampling with non-parametric Bayesian regression, we were able to approximate the full probability distribution function of our quantities of interest at a fraction of the computational costs as compared to sampling of the accurate high fidelity model. Furthermore, we showed that the approach achieves an excellent agreement with the MC reference solution. This renders UQ for complex biomechanical problems feasible on standard computing platforms.

In addition, we investigated the impact of the uncertain constitutive parameter β on the stress, strain, and strain energy in the arterial wall of patient-specific models of two AAAs. We showed that, as expected, the von Mises stress is rather robust against variations in the constitutive parameter β and that the stress state is to a large extent determined by the imaged geometry of the aneurysm. For the von Mises strain or strain energy, on the other hand, large coefficients of variations were found, revealing the insufficiency of deterministic models with population averaged input parameters. We conclude that careful UQ is a necessity under these circumstances, since UQ not only reveals the variability but can also provide worst-case scenarios and estimates such as, e.g., the 95 % quantile or local failure probabilities. Of course, the uncertainties in other input parameters have to be considered in such an analysis as well.

Furthermore, a parametric study was performed to examine the sensitivity of the results with respect to the correlation length d of the random field. We studied three different correlation lengths using direct MC. It was found that the probability distribution of the strains and strain energy at the tested locations were not affected by the choice of the correlation length. The stress, however, showed a greater variance with decreasing correlation length of the random field, revealing

the need for further research and experiments to narrow the range of the correlation length of the random field.

References

- Astrand H, Stalhand J, Karlsson J, Karlsson M, Sonesson B, Lanne T (2011) In vivo estimation of the contribution of elastin and collagen to the mechanical properties in the human abdominal aorta: effect of age and sex. *J Appl Physiol* 110(1):176–187
- Au S, Beck J (2001) Estimation of small failure probabilities in high dimensions by subset simulation. *Probabilist Eng Mech* 16(4):263–277
- Bocchini P, Deodatis G (2008) Critical review and latest developments of a class of simulation algorithms for strongly non-Gaussian random fields. *Probabilist Eng Mech* 23(4):393–407
- Bourgund U, Bucher CG (1986) Importance sampling procedure using design points (ISPUD)—a user's manual. Institute of Engineering Mechanics, University of Innsbruck
- Charpis DC, Schuëller GI, Pellissetti MF (2007) The need for linking micromechanics of materials with stochastic finite elements: a challenge for materials science. *Comput Mater Sci* 41(1):27–37
- Chen P, Quarteroni A, Rozza G (2013) Simulation based uncertainty quantification of human arterial network hemodynamics. *Int J Numer Method Biomed Eng* 29(6):698–721
- Cho H, Venturi D, Karniadakis GE (2013) Karhunen-Loève expansion for multi-correlated stochastic processes. *Probabilist Eng Mech* 34:157–167
- Cliffe KA, Giles MB, Scheichl R, Teckentrup AL (2011) Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients. *Comput Vis Sci* 14(1):3–15
- Del Moral P, Doucet A, Jasra A (2006) Sequential Monte Carlo samplers. *J R Stat Soc Ser B (Statistical Methodology)* 68(3):411–436
- Doucet A, de Freitas N (2001) Sequential Monte Carlo methods in practice. Springer, Berlin
- Eldred MS (2009) Recent advances in non-intrusive polynomial chaos and stochastic collocation methods for uncertainty analysis and design. In: 50th AIAA/ASME/ASCE/AHS/ASC structures, structural dynamics, and materials conference, pp 2009–2274
- Field RV Jr, Grigoriu M (2012) A method for the efficient construction and sampling of vector-valued translation random fields. *Probabilist Eng Mech* 29:79–91
- Gasser TC, Holzapfel GA (2002) A rate-independent elastoplastic constitutive model for biological fiber-reinforced composites at finite strains: continuum basis, algorithmic formulation and finite element implementation. *Comput Mech* 29(4–5):340–360
- Gasser TC, Görgülü G, Folkesson M, Swedenborg J (2008) Failure properties of intraluminal thrombus in abdominal aortic aneurysm under static and pulsating mechanical loads. *J Vasc Surg* 48(1):179–188
- Gasser TC, Auer M, Labruto F, Swedenborg J, Roy J (2010) Biomechanical rupture risk assessment of abdominal aortic aneurysms: model complexity versus predictability of finite element simulations. *Eur J Vasc Endovasc Surg* 40(2):176–185
- Gee M, Förster C, Wall W (2010) A computational strategy for prestressing patient-specific biomechanical problems under finite deformation. *Int J Numer Method Biomed Eng* 26(1):52–72
- Gee MW, Reeps C, Eckstein HH, Wall WA (2009) Prestressing in finite deformation abdominal aortic aneurysm simulation. *J Biomech* 42(11):1732–1739
- Ghanem R, Spanos P (2003) Stochastic finite elements: a spectral approach. Dover, New York
- Grigoriu M (1995) Applied non-Gaussian processes. Prentice Hall, Englewood Cliffs

- Hurtado J, Barbat A (1998) Monte Carlo techniques in computational stochastic mechanics. *Arch Comput Method E* 5(1):3–29
- Kemp F (2003) An introduction to sequential Monte Carlo methods. *J R Stat Soc Ser D (The Statistician)* 52(4):694–695
- Kennedy MC, O'Hagan A (2000) Predicting the output from a complex computer code when fast approximations are available. *Biometrika* 87(1):1–13
- Koutsourelakis P (2009) Accurate uncertainty quantification using inaccurate models. *SIAM J Sci Comput* 31(5):3274–3300
- Ma B, Lu J, Harbaugh RE, Raghavan ML (2007) Nonlinear anisotropic stress analysis of anatomically realistic cerebral aneurysms. *J Biomech Eng* 129(1):88
- Maier A, Gee MW, Reeps C, Pongratz J, Eckstein HH, Wall WA (2010) A comparison of diameter, wall stress, and rupture potential index for abdominal aortic aneurysm rupture risk prediction. *Ann Biomed Eng* 38(10):3124–3134
- Marini G, Maier A, Reeps C, Eckstein HH, Wall WA, Gee MW (2011) A continuum description of the damage process in the arterial wall of abdominal aortic aneurysms. *Int J Numer Method Biomed Eng* 28(1):87–99
- McKay MD, Beckman RJ, Conover WJ (1979) Comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* 21(2):239–245
- Metropolis N, Ulam S (1949) The Monte Carlo method. *J Am Stat Assoc* 44(247):335–341
- Miller K, Lu J (2013) On the prospect of patient-specific biomechanics without patient-specific properties of tissues. *J Mech Behav Biomed Mater* 27:154–166
- Moore JA, Steinman DA (1997) Computational blood flow modelling: errors associated with reconstructing finite element models from magnetic resonance images. *J Biomech* 31(2):179–184
- Neal RM (2001) Annealed importance sampling. *Stat Comput* 11(2):125–139
- Olsson A, Sandberg G (2002) Latin hypercube sampling for stochastic finite element analysis. *J Eng Mech* 128:121
- Panayirci HM, Schuëller GI (2011) On the capabilities of the polynomial chaos expansion method within SFE analysis—an overview. *Arch Comput Method E* 18(1):43–55
- Pellissetti M, Schuëller GI (2009) Scalable uncertainty and reliability analysis by integration of advanced Monte Carlo simulation and generic finite element solvers. *Comput Struct* 87(13–14):930–947
- Peña E, Calvo B, Martínez MA, Doblare M (2008) On finite-strain damage of viscoelastic-fibred materials. Application to soft biological tissues. *Int J Numer Meth Eng* 74(7):1198–1218
- Popescu R, Deodatis G, Prevost J (1998) Simulation of homogeneous nonGaussian stochastic vector fields. *Probabilist Eng Mech* 13(1):1–13
- Pradlwarter H, Schuëller GI, Koutsourelakis P, Charmpis D (2007) Application of line sampling simulation method to reliability benchmark problems. *Struct Saf* 29(3):208–221
- Raghavan M, Vorp D (2000) Toward a biomechanical tool to evaluate rupture potential of abdominal aortic aneurysm: identification of a finite strain constitutive model and evaluation of its applicability. *J Biomech* 33(4):475–482
- Raghavan M, Kratzberg J, Castro de Tolosa E, Hanaoka M, Walker P, da Silva E (2006) Regional distribution of wall thickness and failure properties of human abdominal aortic aneurysm. *J Biomech* 39(16):3010–3016
- Raghavan ML, Webster MW, Vorp DA (1996) Ex vivo biomechanical behavior of abdominal aortic aneurysm: assessment using a new mathematical model. *Ann Biomed Eng* 24(5):573–582
- Raghavan ML, Hanaoka MM, Kratzberg JA, Higuchi MdL, Da Silva ES (2011) Biomechanical failure properties and microstructural content of ruptured and unruptured abdominal aortic aneurysms. *J Biomech* 44(13):2501–2507
- Reeps C, Gee M, Maier A, Gurdan M, Eckstein HH, Wall WA (2010) The impact of model assumptions on results of computational mechanics in abdominal aortic aneurysm. *J Vasc Surg* 51(3):679–688
- Reeps C, Maier A, Pelisek J, Härtl F, Grabher-Meier V, Wall WA, Essler M, Eckstein HH, Gee MW (2012) Measuring and modeling patient-specific distributions of material properties in abdominal aortic aneurysm wall. *Biomech Model Mechanobiol* 12(4):1–17
- Roccabianca S, Figueroa CA, Tellides G, Humphrey JD (2014) Quantification of regional differences in aortic stiffness in the aging human. *J Mech Behav Biomed Mater* 29:618–634
- Rodriguez JF, Alastrue V, Doblare M (2008) Finite element implementation of a stochastic three dimensional finite-strain damage model for fibrous soft tissue. *Comput Method Appl M* 197(9):946–958
- Sankaran S, Marsden AL (2011) A stochastic collocation method for uncertainty quantification and propagation in cardiovascular simulations. *J Biomech Eng* 133(3):031,001
- Schulze-Bauer CA, Holzzapfel GA (2003) Determination of constitutive equations for human arteries from clinical data. *J Biomech* 36(2):165–169
- Shields MD, Deodatis G, Bocchini P (2011) A simple and efficient methodology to approximate a general non-Gaussian stationary stochastic process by a translation process. *Probabilist Eng Mech* 26(4):511–519
- Shinozuka M, Deodatis G (1996) Simulation of multi-dimensional Gaussian stochastic fields by spectral representation. *Appl Mech Rev* 49(1):29–53
- Shinozuka M, Jan C (1972) Digital simulation of random processes and its applications. *J Sound Vib* 25(1):111–128
- Simo J (1987) On a fully three-dimensional finite-strain viscoelastic damage model: formulation and computational aspects. *Comput Methods Appl Mech Eng* 60(2):153–173
- Spanos PD, Zeldin BA (1998) Monte Carlo treatment of random fields: a broad perspective. *Appl Mech Rev* 51(2):219–237
- Speelman L, Bosboom EMH, Schurink GWH, Buth J, Breeuwer M, Jacobs MJ, van de Vosse FN (2009) Initial stress and nonlinear material behavior in patient-specific AAA wall stress analysis. *J Biomech* 42(11):1713–1719
- Stefanou G (2009) The stochastic finite element method: past, present and future. *Comput Method Appl M* 198(9):1031–1051
- Stefanou G, Papadrakakis M (2007) Assessment of spectral representation and Karhunen-Loève expansion methods for the simulation of Gaussian stochastic fields. *Comput Method Appl M* 196(21):2465–2477
- Thubrikar J, Labrosse M, Robicsek F, Al-Soudi J, Fowler B (2001) Mechanical properties of abdominal aortic aneurysm wall. *J Med Eng Technol* 25(4):133–142
- Vallabhaneni SR, Gilling-Smith GL, How TV, Carter SD, Brennan JA, Harris PL (2004) Heterogeneity of tensile strength and matrix metalloproteinase activity in the wall of abdominal aortic aneurysms. *J Endovasc Ther* 11(4):494–502
- Vande Geest J, Di Martino E, Bohra A, Makaroun M, Vorp DA (2006a) A biomechanics-based rupture potential index for abdominal aortic aneurysm risk assessment: demonstrative application. *Ann NY Acad Sci* 1085(1):11–21
- Vande Geest J, Sacks M, Vorp D (2006b) The effects of aneurysm on the biaxial mechanical behavior of human abdominal aorta. *J Biomech* 39(7):1324–1334
- Vande Geest J, Wang D, Wisniewski S, Makaroun M, Vorp D (2006c) Towards a noninvasive method for determination of patient-specific wall strength distribution in abdominal aortic aneurysms. *Ann Biomed Eng* 34(7):1098–1106
- Vanmarcke E (2010) Random fields: analysis and synthesis. World Scientific, Singapore

- Vappou J, Luo J, Konofagou E (2010) Pulse wave imaging for noninvasive and quantitative measurement of arterial stiffness in vivo. *Am J Hypertens* 23(4):393–398
- Volokh K, Vorp D (2008) A model of growth and rupture of abdominal aortic aneurysm. *J Biomech* 41(5):1015–1021
- Volokh KY (2010) Comparison of biomechanical failure criteria for abdominal aortic aneurysm. *J Biomech* 43(10):2032–2034
- Vořechovský M (2008) Simulation of simply cross correlated random fields by series expansion methods. *Struct Saf* 30(4):337–363
- Wall WA, Gee MW (2014) BACI: a parallel multiphysics simulation environment. Tech. rep., Institute for Computational Mechanics, Technische Universität München
- Xiu D (2009) Fast numerical methods for stochastic computations: a review. *Commun Comput Phys* 5(2–4):242–272
- Xiu D (2010) Numerical methods for stochastic computations: a spectral method approach. Princeton University Press, Princeton
- Xiu D, Sherwin SJ (2007) Parametric uncertainty analysis of pulse wave propagation in a model of a human arterial network. *J Comput Phys* 226(2):1385–1407