Applied Mathematical Modelling 128 (2024) 137–153

Contents lists available at ScienceDirect

Applied Mathematical Modelling

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

A probabilistic peridynamic framework with an application to the study of the statistical size effect

Mark Hobbs [a,b], Hussein Rappel [a,*], Tim Dodwell [a,c]

[a] Department of Engineering, Faculty of Environment, Science and Economy, University of Exeter, United Kingdom
[b] The Alan Turing Institute, British Library, 96 Euston Road, London NW1 2DB, United Kingdom
[c] digiLab, Kings Wharf, The Quay, Exeter, EX2 4AN, United Kingdom

ARTICLE INFO

Keywords:
Peridynamics
Multilevel Monte Carlo
Uncertainty quantification
Statistical size effect
Model validation

ABSTRACT

The high computational expense of peridynamic models remains a major limitation, hindering *'outer-loop'* applications that require a large number of simulations, for example, uncertainty quantification. This contribution presents a framework that makes such computations feasible. By employing a Multilevel Monte Carlo framework, where the majority of simulations are performed using a coarse mesh, and performing relatively few simulations using a fine mesh, a significant reduction in computational cost can be realised, and statistics of structural failure can be estimated. The maximum observed speed-up factor is 16 when compared to a standard Monte Carlo estimator, thus enabling the efficient forward propagation of uncertain parameters in a computationally expensive peridynamic model. Furthermore, the multilevel method provides an estimate of both the discretisation error and sampling error, thereby improving confidence in numerical predictions. The performance of the approach is demonstrated through an examination of the statistical size effect in quasi-brittle materials.

1. Introduction

The use of finite-element based models to address probabilistic problems in solid mechanics is well established, with many examples focused on understanding the failure behaviour of quasi-brittle materials. For example, Su et al. [1] used a non-linear finite element model with cohesive elements in combination with Monte Carlo simulation to obtain statistical information of structural load-carrying capacity. More recently, Wu et al. [2] employed a phase-field model within a finite element framework in combination with Monte Carlo simulation to estimate the probability distribution of the failure load. The reader is also referred to the work of Song and Kawai [3], who provide a comprehensive review of Monte Carlo methods in structural reliability analysis, along with references to many applied examples that employ finite element-based models.

Whilst finite element-based models that solve the governing equations of classical continuum mechanics have been very successful, they are inherently limited when dealing with the complex and discrete nature of material failure. The governing partial differential equations, by definition, assume a spatially continuous and differentiable displacement field that is undefined along discontinuities, necessitating the implementation of various remedies to address these difficulties. Additionally, the classical theory assumes that all forces are contact forces that act across zero distance (local theory). However, fracture is a process in which non-

* Corresponding author at: Department of Engineering, Faculty of Environment, Science and Economy, University of Exeter, United Kingdom.
E-mail addresses: mhobbs@turing.ac.uk (M. Hobbs), h.rappel@exeter.ac.uk (H. Rappel), tim@digilab.co.uk (T. Dodwell).

https://doi.org/10.1016/j.apm.2024.01.004
Received 2 May 2023; Received in revised form 23 December 2023; Accepted 5 January 2024

Available online 11 January 2024

locality is known to be important [4], and a fundamental feature of any model capable of correctly capturing the structural size effect is the presence of some form of length scale [5].

The peridynamic theory, an integral-type non-local theory of solid mechanics introduced by Silling [6], overcomes many of the inherent limitations of the classical continuum theory by eliminating the requirement of a spatially continuous and differentiable displacement field. As a result, the theory is capable of modelling the emergence of damage localisation and fracture without any additional assumptions or techniques. However, the high computational cost of peridynamic simulations remains a major limitation, and as a consequence, the study of problems with a probabilistic/uncertain component is prohibitively expensive due to the large number of repeat simulations typically required by conventional Monte Carlo (MC) methods. Although techniques exist to reduce the number of required samples, the use of a computationally expensive fine mesh is necessary to control the discretisation error. Given that the computational cost of a peridynamic model is $\mathcal{O}(N^2)$, where N is the number of particles, it is prudent to avoid running many simulations using a fine mesh. Hence, the use of multilevel methods, where the majority of samples are taken using a coarse and computationally cheap mesh becomes imperative. This study aims to establish the feasibility of forward uncertainty quantification for costly peridynamic simulations through the utilisation of multilevel methods. To the best of the authors knowledge, there is no work within the peridynamic literature that examines the forward propagation of uncertain parameters.

In this work, we employ the multilevel Monte Carlo (MLMC) method. The aim of MLMC is attain the same solution error as MC but at a significantly reduced computational cost. The standard MC estimator is computationally expensive as all samples must be computed using a fine mesh that guarantees a small discretisation error. A significant reduction in computational cost can be realised by taking the majority of samples on a coarse mesh (low accuracy but computationally cheap), and taking relatively few samples on a fine mesh (high accuracy but computationally expensive). This is made possible by isolating the error sources in the estimator: (1) sampling error (variance) and (2) discretisation error (deterministic error). The sampling error is controlled by using a low accuracy but computationally cheap model to take a large number of samples, and the discretisation error is reduced to a defined tolerance by employing a sufficiently fine mesh.

Multilevel techniques were first introduced by Heinrich and Sindambiwe [7] and Heinrich [8], and later popularised by Giles [9] for option pricing in computational finance. The first application of multilevel methods in the field of engineering was the study of uncertainty in groundwater flows [10]. Since Cliffe et al. [10] recognised the potential of multilevel methods, there has been a wide range of applications in engineering and scientific fields (primarily for the solution of partial differential equations (PDEs)). For example, Elfverson et al. [11] explored the efficient estimation of failure probabilities using multilevel methods, Dodwell and et al. [12] computed the failure probability of composite structures, and Hamdia and Ghasemi [13] computed the failure probability of linear elastic materials. Further applications include the study of the travel time of particles through random heterogeneous porous media [14], and assessing the risk of coastal flooding [15]. For a detailed review of multilevel Monte Carlo methods, the reader is referred to the work of Giles [16].

We have chosen to demonstrate the framework through an application to the study of the structural size effect in quasi-brittle materials. We focus on quasi-brittle materials because the range of experimental data is greater than that of any other material and quasi-brittle materials exhibit a significant size effect, where large elements fail at lower stresses than small-scale elements with identical geometry. Accurately capturing and understanding this behaviour is essential for safe and robust predictions. As stated in [4], the correct modelling of the size effect on material strength should be adopted as the basic criterion of acceptability of any model.

The statistical variability in material properties is a major contributing factor to the size effect on structural strength and numerical models provide an excellent tool for studying the magnitude of this contribution. Hobbs et al. [17] previously examined the size effect in quasi-brittle materials using a deterministic bond-based model. The numerical results were generally in good agreement with experimental data but it was not possible to examine the statistical component due the high computational expense of peridynamic simulations. Through the multilevel framework, peridynamic models could yield deeper insights into the mechanisms that govern the structural size effect. It is important to clarify that this study does not aim to provide a detailed examination of the statistical size effect, and the primary focus of the case studies is to explore potential computational savings.

The aims of this paper are: (1) demonstrate that forward uncertainty quantification for expensive peridynamic simulations is feasible by employing the MLMC method, (2) quantify the computational savings, and (3) demonstrate the importance of forward UQ by selecting examples where a deeper understanding of the physical behaviour can be gained through the consideration of uncertainty. A secondary aim is to provide new insights into the convergence behaviour of peridynamic models under a probabilistic setting. Both the convergence of the discretisation error and the predicted failure load are explored.

The paper is organised as follows: Section 2 introduces the peridynamic theory and the numerical model (a bond-based peridynamic model). Section 3 details the standard and multilevel Monte Carlo methodology. Section 4 presents two carefully selected case studies. Section 5 further discusses the results before the paper is concluded in Section 6.

2. The bond-based peridynamic model

There are two main formulations of the peridynamic theory: the bond-based formulation [6] and the state-based formulation [18]. In this work, we employ a bond-based model due to its lower computational expense and proven predictive capabilities. The bond-based theory is also conceptually simpler and more accessible to a broader range of readers.

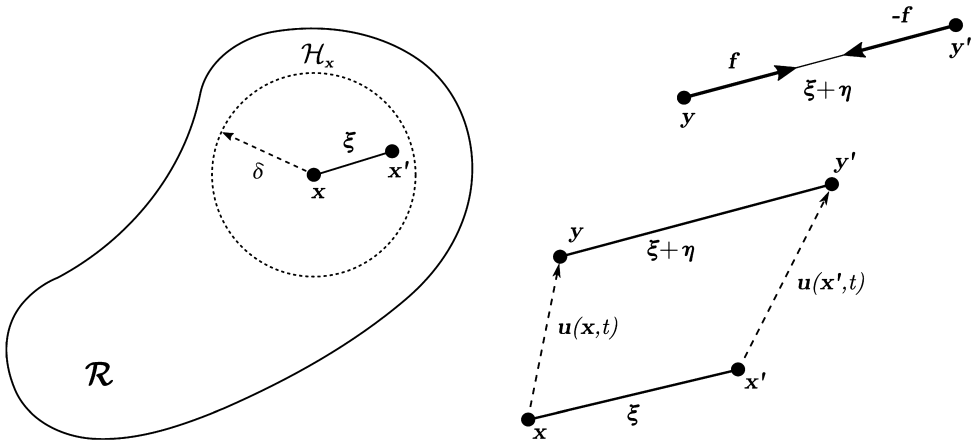

Fig. 1. Peridynamic continuum and kinematics of particle pair and bond-based pairwise force function.

2.1. Peridynamic continuum model

This section briefly introduces the bond-based peridynamic theory. However, the intention is not to provide a thorough explanation of the theory, and readers are directed to [6,19,20] for a more comprehensive understanding. Alternatively, Bobaru et al. [21] provides a less rigorous but more mechanically intuitive method for deriving the governing equations.

Let $\mathcal{R}$ denote the spatial region occupied by a material body, where $\mathbf{x}$ denotes a material point (interchangeably referred to as a particle) and $\mathbf{u}$ represents the displacement vector of a given material point $\mathbf{x}$ (see Fig. 1). Material points separated by a finite distance less than or equal to δ in the undeformed material body are assumed to interact through a pairwise force function $\mathbf{f}$. The pairwise force function for a particle pair at any given time t can be described as follows

$$\mathbf{f} = \mathbf{f}(\mathbf{x}, \mathbf{x}', \mathbf{u}(\mathbf{x}, t), \mathbf{u}(\mathbf{x}', t), t), \quad \mathbf{x}' \in \mathcal{R} : ||\mathbf{x}' - \mathbf{x}||_2 \leqslant \delta \tag{1}$$

Therefore, employing Newton's second law, the peridynamic equation of motion for a material point $\mathbf{x}$ at time t reads as

$$\rho \ddot{\mathbf{u}}(\mathbf{x}, t) = \int_{\mathcal{H}_\mathbf{x}} \mathbf{f}(\mathbf{u}(\mathbf{x}', t) - \mathbf{u}(\mathbf{x}, t), \mathbf{x}' - \mathbf{x}) dV_{\mathbf{x}'} + \mathbf{b}(\mathbf{x}, t) \tag{2}$$

where ρ denotes mass density, $\ddot{\mathbf{u}}$ denotes particle acceleration, $\mathbf{b}$ denotes body force per unit volume and $\mathcal{H}_\mathbf{x}$ (Eq. (3)) denotes the neighbourhood of material point $\mathbf{x}$. The size of the neighbourhood is defined by the horizon length δ. Note that for a 3D problem, the material point neighbourhood will be a sphere, and for a 2D problem, the neighbourhood will be circular.

$$\mathcal{H}_\mathbf{x} = \mathcal{H}_\mathbf{x}(\mathbf{x}, \delta) = \{\mathbf{x}' \in \mathcal{R} : ||\mathbf{x}' - \mathbf{x}||_2 \leqslant \delta\} \tag{3}$$

The pairwise force function $\mathbf{f}$ represents the force that particle $\mathbf{x}'$ exerts on particle $\mathbf{x}$ and contains all the constitutive information of the material under analysis. This interaction is also known as the peridynamic bond force. Particles separated by a distance greater than the horizon length δ in the undeformed (reference) state do *not* interact. The pairwise force function is defined as follows

$$\mathbf{f}(\boldsymbol{\eta}, \boldsymbol{\xi}) = f(|\boldsymbol{\xi} + \boldsymbol{\eta}|, \boldsymbol{\xi}) \frac{\boldsymbol{\xi} + \boldsymbol{\eta}}{|\boldsymbol{\xi} + \boldsymbol{\eta}|} \tag{4}$$

where $\boldsymbol{\xi} = \mathbf{x}' - \mathbf{x}$ denotes the initial relative position vector of a pair of particles, $\boldsymbol{\eta} = \mathbf{u}' - \mathbf{u}$ denotes the relative displacement vector and the current relative position vector is given by $\boldsymbol{\xi} + \boldsymbol{\eta}$. Note that in a bond-based model, the force vector $\mathbf{f}$ is parallel to the deformed bond and the scalar bond force f (vector magnitude) is proportional to the bond stretch s.

To differentiate the peridynamic theory from other non-local theories, it is important to note that many non-local theories average some measure of strain within the horizon of a material particle. The peridynamic theory does not rely on the concept of strain, which by its definition, necessitates calculating partial derivatives of displacement [19].

2.2. Non-locality

The peridynamic theory is a non-local theory in which material points interact with each other directly over finite distances. This is in contrast to the classical theory of solid mechanics, where it is assumed that all forces are contact forces that act across zero distance (local theory). Physical justification of non-locality was provided by Bažant [4], and further discussion on the origins of non-locality (with a focus on the peridynamic theory), can be found in Chapter 1 of [21] and Hobbs [22].

At the macroscale, the peridynamic horizon δ is a numerical constant with no physical meaning. This differentiates the peridynamic model from many numerical approaches, and the use of an ambiguous characteristic length parameter is avoided. For a given

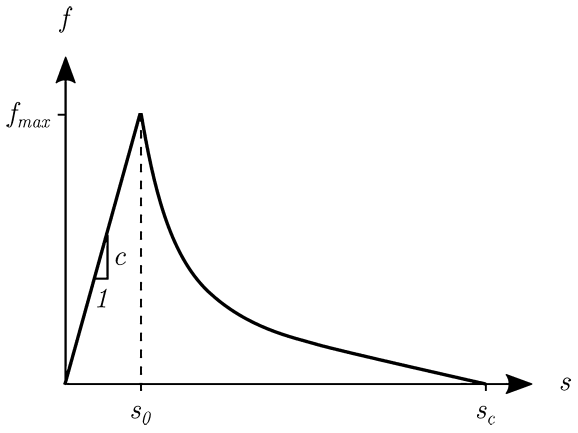

Fig. 2. Non-linear damage model (f-s). The force-stretch relationship is described by an exponentially decaying model with a linear term that forces the curve to intersect with the horizontal axis at s_c.

value of δ, the parameters in a peridynamic model can be chosen to match a given set of physically measurable material properties. Therefore, an optimum value of δ must be chosen that provides high accuracy whilst balancing computational expense. Section 2.3.2 discusses the selection of an optimum value of δ.

The reader should note the distinction between the non-local length scale in the peridynamic model (horizon δ), and the non-local length scale in a spatially correlated random field (correlation length l_c). The correlation length l_c is generally considered to be a material parameter reflecting the internal length scale of the microstructure. This will be discussed throughout the paper.

2.3. Numerical model

To illustrate the framework, we employ a two-dimensional bond-based peridynamic model, but the reader should note that the presented framework is not limited solely to bond-based models, and more complex models could be employed without any modification. The primary aim of this work is to demonstrate the multilevel framework and a detailed treatment of the numerical model is not provided. All the results presented in this paper were obtained using the explicit scheme outlined in Figure 4.14 of [22], and readers can find implementation details in the same reference. The main distinction of the model used in this work is the existence of two length scales: (1) the peridynamic horizon δ and (2) the correlation length l_c in the random field.

2.3.1. Constitutive model

Generally, it is assumed that the force-stretch (f-s) relationship of a peridynamic bond should align with the macroscopic material response. The stress-strain response of quasi-brittle materials is characterised by strain-softening behaviour in the post-peak stage, and hence we employ the non-linear softening law, illustrated in Fig. 2, first proposed in [22]. To simulate material damage, a failure mechanism is incorporated into the model by eliminating the interaction between pairs of particles if the stretch of the bond connecting them exceeds a critical value.

2.3.2. Numerical convergence

The accuracy and convergence behaviour of a peridynamic model is complicated by the presence of a length scale. To determine an optimum value of δ, an additional parameter m must be introduced. m is the ratio between the horizon radius and grid resolution ($m = \delta/\Delta x$). Bobaru and et al. [23] and Ha and Bobaru [24] define and discuss two fundamental types of convergence: (1) m-convergence: δ is fixed and $m \to \infty$. This can also be stated as δ is fixed and $\Delta x \to 0$. (2) δ-convergence: m is fixed and $\delta \to 0$. This can also be stated as m is fixed and $\Delta x \to 0$. See Fig. 3 for a graphical representation of the types of convergence. A third type of convergence can be defined: δm-convergence. This is a combination of δ- and m-convergence. See [23] for details.

In this work, we consider δ-convergence, as it is generally agreed that m should be close to 3. Madenci and Oterkus [25] investigated the choice of m for macroscale problems and it was found that values of $m = 1$ and $m = 3$ achieved the highest accuracy when compared to the classical analytical solution for the displacement of a one-dimensional bar subjected to a defined initial strain. Values of m much larger than 3 lead to excessive wave dispersion and become extremely computationally expensive. When fracture behaviour is also considered, values of m less than 3 lead to grid dependence on crack propagation [25,26]. Hu et al. [27] and Seleson [28] examined the m-convergence behaviour for two-dimensional models and [22] examined the m-convergence behaviour for three-dimensional models. Higher values of m improve the spatial integration accuracy but $m \approx 3$ provides an acceptable approximation. A value of $m = \pi$ is generally recommended for macroscale problems and is found extensively throughout the literature. The m-ratio is set to π for all problems in this paper.

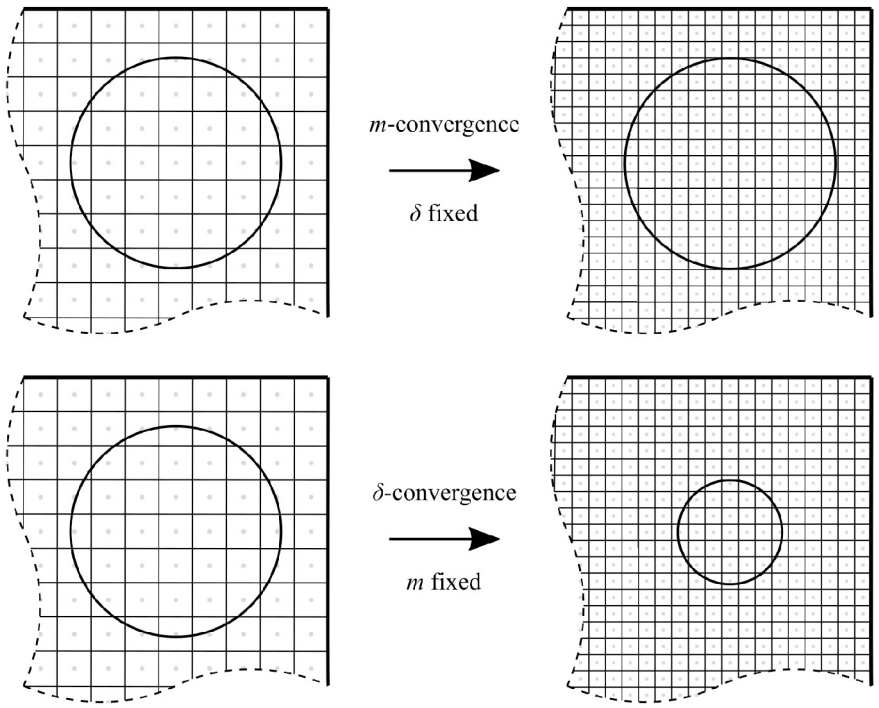

Fig. 3. Graphical representation of the two fundamental types of convergence: m-convergence and δ-convergence.

3. Multilevel Monte Carlo methodology

To explain the standard and multilevel Monte Carlo method, let us assume that we have a numerical model of a brittle/quasi-brittle structure that is subject to some uncertainty in the material properties. The accuracy and computational cost of the model is proportional to the number of degrees of freedom (M) and thus to the resolution of the mesh. Generally, we are interested in some scalar quantity of interest (QoI) $Q = Q(\mathbf{x}, t, \boldsymbol{\omega})$, for example, the failure load or the displacement at a particular point, where $\mathbf{x}$ and t are the spatial and temporal coordinates and $\boldsymbol{\omega}$ represents a vector of random variables that takes values in $\mathbb{R}^M$. $\boldsymbol{\omega}$ represents sources of uncertainty in the problem, in this case, the material properties. Note that the quantity of interest (Q) could be a function, for instance, the load-deflection response of a structure. For the presented case studies, the quantity of interest is the failure load, and the objective is to compute the expected value of Q, denoted $\mathbb{E}[Q]$, with a quantified level of uncertainty. However, for many real world applications, the probability distribution of Q is of more interest. Methods for obtaining the probability distribution of Q will be discussed.

3.1. Standard Monte Carlo simulation

In a standard Monte Carlo (MC) simulation, a large number (N) of independent random realisations (or samples) of the parameters are generated. For every sample, the solution is computed using a numerical solver (finite element model, particle model etc). The accuracy of the solution is directly proportional to the resolution of the discretisation, and it is assumed that $Q_M \to Q$ as $M \to \infty$, therefore $\mathbb{E}[Q_M] \to \mathbb{E}[Q]$ for $M \to \infty$. The required accuracy and computational budget govern the selection of M. The standard Monte Carlo estimator for the expected value $\mathbb{E}[Q_M]$ of Q_M, based on N samples, is given by Eq. (5), where $Q_M^{(j)}$ is the quantity of interest of the j^{th} sample.

$$\hat{Q}_{M,N}^{MC} = \frac{1}{N}\sum_{j=1}^{N} Q_M^{(j)} \tag{5}$$

Note that $\hat{Q}_{M,N}^{MC}$ is an unbiased estimator of $\mathbb{E}[Q_M]$, meaning that $\mathbb{E}[\hat{Q}_{M,N}^{MC}] = \mathbb{E}[Q_M]$. An estimator is *unbiased* if its expectation is the quantify of interest that we wish to estimate. The accuracy of the estimator ($\hat{Q}_{M,N}^{MC}$) can be quantified using the *root mean square error* (RMSE):

$$e(\hat{Q}_{M,N}^{MC}) = \sqrt{\mathbb{E}[(\hat{Q}_{M,N}^{MC} - \mathbb{E}[Q])^2]} \tag{6}$$

An advantage of quantifying the accuracy of the estimator in this way is that the mean square error can be expanded and two distinct sources of error can be isolated: (1) the bias error and (2) the sampling error.

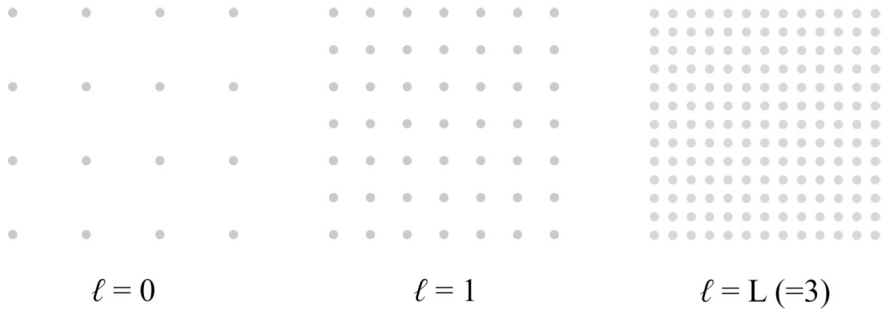

Fig. 4. Example of a hierarchy of uniformly refined meshes employed in the MLMC method. Each mesh corresponds to a level $0 \leq \ell \leq L$ in the multilevel method with $M_0 < \cdots < M_\ell < \cdots < M_L$ degrees of freedom. We restrict ourselves to the case of uniform mesh refinement where the node spacing (Δx) is halved every time.

$$e(\hat{Q}^{MC}_{M,N})^2 = \underbrace{\left(\mathbb{E}[Q_M - Q]\right)^2}_{\textit{bias error}} + \underbrace{\frac{\mathbb{V}[Q_M]}{N}}_{\textit{sampling error}} \tag{7}$$

The first term in Eq. (7) is the *bias error* (sometimes referred to as the discretisation or numerical error). This arises as we are actually interested in the expected value $\mathbb{E}[Q]$ of Q, the unobtainable random variable corresponding to the exact solution without any numerical error. If we assume that the numerical model converges to the exact solution as the mesh is refined, $\mathbb{E}[Q_M] \to \mathbb{E}[Q]$ as $M \to \infty$, then we can state the following

$$\mathbb{E}[Q_M - Q] \approx M^{-\alpha}, \quad \text{as} \quad M \to \infty \tag{8}$$

where α is the order of convergence and $\alpha > 0$.[1] The value of α is problem dependent and depends on numerous factors, such as, the chosen numerical model, the material model and the smoothness of the random field. By making M sufficiently large, the discretisation error can be reduced to any tolerance value ϵ_b.

The second term in Eq. (7) is the *sampling error* and represents the variance of the estimator and decays inversely with the number of samples N. To ensure that the sampling error is less than a defined tolerance ϵ_s, it is reasonable to determine the number of samples N using Eq. (9).

$$\epsilon_s^2 \approx \frac{\mathbb{V}[Q_M]}{N} \quad \therefore N \approx \frac{\mathbb{V}[Q]}{\epsilon_s^2} \tag{9}$$

To reduce the total error to a defined tolerance, the number of degrees of freedom M and the number of samples N must both be increased. This can be prohibitively computationally expensive when the cost to compute each sample to the required level of accuracy is high. The cost C_M to compute a single sample of Q_M is dependent on the computational complexity of the solver. The computational cost will grow as follows

$$C_M \approx M^\gamma \tag{10}$$

for some $\gamma \geq 1$. The rate at which the computational cost grows (γ) is dependent on a number of factors, such as, the dimension of the problem and the chosen solver (explicit/implicit).

Standard MC estimators are proven to be robust and accurate for many problems, however the slow convergence rate limits applications to problems where the QoI can be computed cheaply. For problems that require the solution of computationally expensive models it is not possible to achieve reasonable estimations in an acceptable time. Different strategies have been examined to accelerate MC estimators, and all are based on the idea of reducing the sampling error.

3.2. Multilevel Monte Carlo simulation

The standard MC estimator is too costly as the quantity of interest for every sample must be computed to the level of accuracy required to ensure that the discretisation error is less than a defined tolerance. The key idea of MLMC is to compute a sequence of estimates of the quantity of interest using a hierarchy of nested meshes, as illustrated in Fig. 4. A significant reduction in computational cost can be realised by taking the majority of samples on a coarse mesh (low accuracy but computationally cheap), and taking relatively few samples on a fine mesh (high accuracy but computationally expensive).

Because of the linearity of the expectation operator, the expected value of Q on the finest mesh ($\mathbb{E}[Q_{M_L}]$) can be expressed as a telescopic sum of the expectation of Q on the coarsest mesh plus a sum of correction terms that account for the difference between evaluations on consecutive mesh levels.

[1] The notation $a_n \approx n^{-\alpha}$ denotes that the sequence $\{a_1, a_2, a_3, ..., a_n\}$ decreases with a rate of $-\alpha$ as $n \to \infty$, and the lower bound of the sequence is $c_1 n^{-\alpha}$ and the upper bound is $C_1 n^{-\alpha}$ ($c_1 n^{-\alpha} \leq a_n \leq C_1 n^{-\alpha}$), where c_1 and C_1 are positive constants independent of n ($0 < c_1 \leq C_1$).

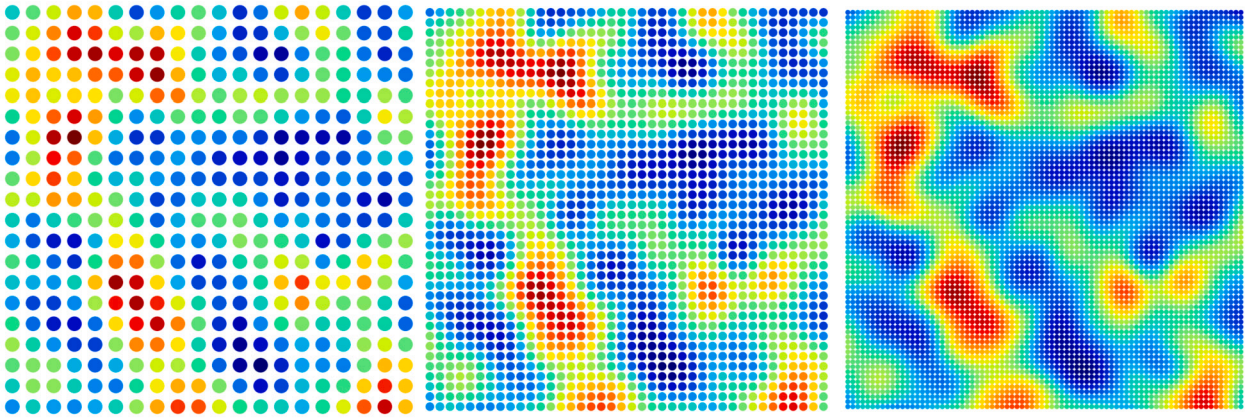

Fig. 5. This figure illustrates the same sample $\omega^{(i)}$ of a spatially correlated random field on three mesh levels. Note that the resolution of the coarsest level Δx_0 must be smaller than the correlation length l_c in the random field. Cliffe et al. [10] states that the optimal choice for the resolution of the coarsest mesh is such that Δx_0 is slightly smaller than l_c.

$$\mathbb{E}[Q_{M_L}] = \mathbb{E}[Q_{M_0}] + \sum_{\ell=1}^{L} \mathbb{E}[Y_\ell] \tag{11}$$

where Y_ℓ is the discrepancy between the QoI at successive mesh resolutions and is defined as follows

$$Y_\ell = \begin{cases} Q_{M_0} & \text{if } \ell = 0 \\ Q_{M_\ell} - Q_{M_{\ell-1}} & \text{if } 0 < \ell \leq L \end{cases} \tag{12}$$

The multilevel estimator for $\mathbb{E}[Q]$ is given by Eq. (13).

$$\hat{Q}_M^{ML} = \hat{Q}_{M_0,N_0}^{MC} + \sum_{\ell=1}^{L} \hat{Y}_{\ell,N_\ell}^{MC} \tag{13}$$

The number of samples N_ℓ on each level is determined such that the total computational cost of the estimator is minimised for a defined sampling error (see Eq. (23)). It is important to highlight that the same random sample $\omega^{(i)}$ is used to compute the quantity $Q_{M_\ell}^{(i)} - Q_{M_{\ell-1}}^{(i)}$, i.e. a coarsened version of the random sample used to compute $Q_{M_\ell}^{(i)}$ is used to compute $Q_{M_{\ell-1}}^{(i)}$ (refer to Fig. 5 for clarification).

As all the expectations $\mathbb{E}[Y_\ell]$ are estimated independently, the variance of the multilevel estimator is $\mathbb{V}[\hat{Q}_M^{ML}] = \sum_{\ell=0}^{L} N_\ell^{-1} \mathbb{V}[Y_\ell]$, where $\mathbb{V}[Y_0] = \mathbb{V}[Q_{M_0}]$. The accuracy of the estimator can be quantified by considering the mean square error.

$$e(\hat{Q}_M^{ML})^2 = \underbrace{\left(\mathbb{E}[Q_M - Q]\right)^2}_{\textit{bias error}} + \underbrace{\sum_{l=0}^{L} \frac{\mathbb{V}[Y_\ell]}{N_\ell}}_{\textit{sampling error}} \tag{14}$$

Much like the standard MC estimator, the mean square error is composed of two terms, the bias error and the sampling error. The bias error is exactly the same as in the MC estimator (see Eq. (7)), and the number of degrees of freedom on the finest level (M_L) must be sufficiently large to satisfy Eq. (8), and thus ensuring that the bias error is less than ϵ_b.

The multilevel estimator is cheaper than the standard MC estimator as the number of samples N_ℓ on every level can be chosen to ensure that the sampling error is less than ϵ_s, whilst minimising the total computational cost of the estimator. The computational cost of the multilevel Monte Carlo estimator is given by the following

$$C(\hat{Q}_M^{ML}) = \sum_{\ell=0}^{L} N_\ell C_\ell \tag{15}$$

where C_ℓ is the cost to compute a single sample of Y_ℓ on level $\ell \geq 1$ or Q_{M_0} on level 0. Note that taking a sample of Y_ℓ requires the numerical approximation of Q on two consecutive mesh levels (both $Q_{M_\ell}^{(i)}$ and $Q_{M_{\ell-1}}^{(i)}$ must be computed). The determination of the optimal sample allocation is detailed in Section 3.2.2.

To achieve a RMSE of ϵ, it can be asserted that the multilevel estimator is computationally cheaper than the standard MC estimator due to the significant reduction in variance [10]. As the MLMC estimator is unbiased, the variance of the estimator is equal to

$$\mathbb{V}[\hat{Q}_M^{ML}] = \sum_{\ell=0}^{L} \frac{\mathbb{V}[Y_\ell]}{N_\ell} \tag{16}$$

The variance of the multilevel estimator is reduced as both numerical approximations Q_{M_ℓ} and $Q_{M_{\ell-1}}$ converge to Q and consequently

$$\mathbb{V}[Y_\ell] = \mathbb{V}[Q_{M_\ell} - Q_{M_{\ell-1}}] \to 0 \quad \text{as} \quad M_\ell \to \infty \tag{17}$$

It is assumed that there exists a $\beta > 0$, where β is the order of convergence of the sampling error, such that

$$\mathbb{V}[Q_{M_\ell} - Q_{M_{\ell-1}}] \approx M_\ell^{-\beta} \tag{18}$$

3.2.1. Error estimation

The aim is to estimate $\mathbb{E}[Q]$ such that the RMSE is below a defined tolerance ϵ, whilst minimising the total computational cost of the estimator $C(\hat{Q}_M^{ML})$. The RMSE, defined by Eq. (14), is comprised of two parts: (1) the bias error and (2) the sampling error. To ensure that the RMSE is less than ϵ, it is sufficient to bound each term by $\epsilon^2/2$. To estimate the bias error, it is assumed that M_ℓ is sufficiently large so that the decay in $\left|\mathbb{E}[Q_{M_\ell} - Q]\right|$ is in the asymptotic region and satisfies the following

$$\left|\mathbb{E}[Q_{M_\ell} - Q]\right| \approx M^{-\alpha} \tag{19}$$

Following the derivation presented in [12], for uniform mesh refinement, where the number of degrees of freedom on level ℓ is given by $M_\ell \approx m^\ell M_0$, the bias error on level ℓ can be over-estimated as follows

$$\epsilon := \left|\mathbb{E}[Q_{M_\ell} - Q]\right| \leq \frac{1}{rm^\alpha - 1} \hat{Y}_{\ell,N_\ell}^{MC} \tag{20}$$

where r is set to 1. This is equivalent to the assumption that M_ℓ is sufficiently large so that the decay in $\left|\mathbb{E}[Q_{M_\ell} - Q]\right|$ is in the asymptotic region. The user may wish to select a more conservative values for r, for example 0.7 or 0.9. If the bias error is greater than the tolerance, then M_L must be increased.

To ensure that the sampling error is less than or equal to the sample tolerance ϵ_s, the following constraint is enforced

$$\sum_{\ell=0}^{L} \frac{\mathbb{V}[Y_\ell]}{N_\ell} \leq \epsilon_s^2 \tag{21}$$

As the number of samples increases, the variance of the sample mean decreases and hence precision increases. The sample variance is estimated as follows

$$s_\ell^2 = \left(\frac{1}{N_\ell} \sum_{n=1}^{N_\ell} (Y_\ell^n)^2\right) - \left(\hat{Y}_{\ell,N_\ell}^{MC}\right)^2 \approx V_\ell \tag{22}$$

3.2.2. Sample allocation

The optimal sample allocation (number of samples per level N_ℓ) is determined by solving a constrained optimisation problem that minimises $C(\hat{Q}_M^{ML})$ with respect to N_ℓ, subject to the constraint that the sampling error of the multilevel estimator is less than or equal to the defined tolerance ϵ_s.

$$N_\ell = \epsilon_s^{-2} \left(\sum_{\ell=0}^{L} \sqrt{V_\ell C_\ell}\right) \sqrt{\frac{V_\ell}{C_\ell}} \tag{23}$$

The computational cost of the MLMC estimator grows as follows

$$C(\hat{Q}_M^{ML}) = \epsilon^{-2} \left(\sum_{\ell=0}^{L} \sqrt{V_\ell C_\ell}\right)^2 \approx \epsilon^{-2-max(0,(\gamma-\beta)/\alpha)} \quad \text{as} \quad \epsilon \to 0 \tag{24}$$

The rate at which the computational cost grows with respect to the number of degrees of freedom M is given by Eq. (25), for some $\gamma \geq 1$.

$$C_\ell \approx M_\ell^\gamma \tag{25}$$

The reader is referred to Cliffe et al. [10] and [16] for a full proof of the MLMC computational complexity theorem with bounds on the RMSE.

3.2.3. MLMC implementation

Pseudo-code for the adaptive MLMC method is outlined in Algorithm 1. Optimal values of L, M_ℓ and N_ℓ are computed *on the fly* from the sample averages and the sample variances of Y_ℓ. We set the number of warm-up samples $N^\star$ to be 100. As each sample is independent and there are no shared memory requirements, Algorithm 1 can be trivially parallelised across an unlimited number of independent compute nodes.

Algorithm 1 Multilevel Monte Carlo algorithm.

Set $L = 0$, $N_\ell = N^\star$, `converged == false`
while `converged == false` **do**
  Take $N^\star$ warm-up samples on level L
  Estimate the sample variance $\mathbb{V}[Y_\ell]$ on all levels using Eq. (22)
  Estimate the optimal number of samples $\hat{N}_\ell$ on each level using Eq. (23)
  Compute $\hat{N}_\ell - N^\star$ additional samples on each level
  Estimate bias error $\hat{\epsilon}_b$ on level L using Eq. (20)
  if $\hat{\epsilon}_b < \epsilon_b$ **then**
    `converged == True`
  else
    $L = L + 1$
  end if
end while

4. Case studies

Two case studies are presented to demonstrate the power of employing the MLMC framework in combination with a peridynamic model. The case studies are selected as examples where uncertainty must be considered to gain a deeper understanding of the physical behaviour. We focus on quasi-brittle materials because the range of experimental data is greater than that of any other material and quasi-brittle materials exhibit a significant size effect. A stochastic model is required for a complete examination of the mechanisms that govern the structural size effect. The first case study examines the structural size effect in a notched beam (Type 2 problem), and the second case study examines the structural size effect in an unnotched beam (Type 1 problem).

The following subsection briefly discusses the structural size effect. For a detailed review of the structural size effect, the reader is referred to Bažant and Planas [29] and Bažant [30].

4.1. Structural size effect

According to the theory of strength-of-materials, when the maximum stress in a structure surpasses an upper limit, as determined by small-scale tests on representative material samples, structural failure is expected to occur. Basic tests such as uniaxial tension, uniaxial compression, and flexural tests are employed to establish this limit for various loading conditions. However, this approach is inadequate for quasi-brittle materials [31], as quasi-brittle materials display a size effect, meaning that larger elements fail at lower stresses than smaller-scale elements with identical geometry.

In brittle and quasi-brittle materials, there are two primary factors that contribute to the size effect on structural strength [29,30]: (1) the release of stored energy (deterministic size effect), and (2) the statistical variability in material properties, which is generally of lesser significance (statistical size effect). The deterministic size effect is governed by the size of the fracture process zone (the area where energy is dissipated during fracture), relative to the size of the structure. The statistical size effect is a consequence of the random distribution of material properties and defects. As the size of the specimen increases, the likelihood that it contains a defect that will lead to failure also increases. As there is minimal experimental data on the response of very large structures and the safety implications are often much greater, correctly predicting the influence of the statistical size effect is of utmost importance.

Two types of size effect law are defined: Type 1 applies to structures with no notches or pre-existing cracks (fracture initiates from a smooth surface), and Type 2 applies to structures with a notch (high-stress concentrations). The influence of the size effect on the mean strength of Type 1 and 2 structures is markedly different [32]. For Type 1 structures, the size effect has a significant statistical component, whereas for Type 2 structures, the statistical component is minimal.

The size effect in quasi-brittle materials has previously been examined using a deterministic bond-based peridynamic model [17]. The model did not consider the spatial variability in material properties and the magnitude of the statistical size effect remains to be established. Due to the high computational expense of peridynamic simulations, examining the statistical size effect was impracticable but the presented framework allows us to overcome the aforementioned issues. The deterministic model employed in [17] was validated against the complete set of experimental results published by Grégoire et al. [33]. This work only considers two members from the test series as the aim of this study is to demonstrate the possible computational savings that can be realised using the MLMC framework, and to demonstrate the importance of examining uncertainty. Future work will use the MLMC framework to examine the full series of tests and provide a comprehensive study of the statistical size effect.

4.2. Material strength distribution and spatial correlation

The efficient and accurate generation of random fields that capture the variability and spatial correlation of material properties is an import element of the numerical framework. In this contribution, Karhunen-Loève expansion (KL expansion) is employed because of its practical simplicity and relatively low computational cost. The correlation between material points is captured using an exponential covariance function, and the distance over which the correlation exists is determined by a length scale l_c. For quasi-brittle materials, Grassl and Bažant [34] suggested that the correlation length must, at a minimum, be as large as the fracture process zone (FPZ). For concrete, the size of the FPZ is approximately two to three times the maximum aggregate size [35].

The material strength distribution significantly influences the predicted results and the convergence of the model. In the literature, various distributions, including normal (Gaussian), log-normal, Gauss-Weibull and Weibull, have been employed for modelling

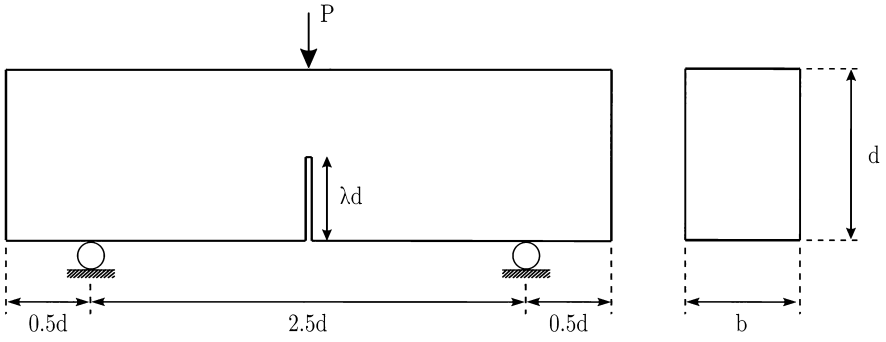

Fig. 6. Schematic of the experimental setup (adapted from [33]).

Table 1
Mesh level ℓ, corresponding mesh resolution Δx and corresponding number of nodes M. The number of nodes (degrees of freedom) on level ℓ is given by $M_\ell = m^\ell M_0$, where $m = 4$.

Level ℓ	0	1	2	3	4
Δx (mm)	10	5	2.5	1.25	0.625
no. nodes M	350	1,400	5,600	22,400	89,600

quasi-brittle materials. We examined the applicability of normal, log-normal and Weibull distributions and found that the Weibull distribution demonstrated the best agreement with experimental data while enhancing the rate of convergence for discretisation errors. These findings are discussed further in Section 5.4.

To easily generate a random field, where the probability distribution function of a material parameter at a given location is a univariate Weibull distribution, we follow the approach of Rappel et al. [36] and Rappel et al. [37]. In a Gaussian random field, the probability density function of a material parameter at a given location is a univariate Gaussian distribution. Using the copula theorem and Gaussian fields, different types of univariate marginal distributions can be produced but with the same correlation structure as Gaussian fields. Keeping the Gaussian correlation structure is advantageous as it allows us to draw samples from a Gaussian field and transform the samples into a random field with the desired distribution.

4.3. Case study 1: statistical size effect in quasi-brittle materials (Type 2)

The first problem that we consider is a notched concrete beam in three-point bending, tested experimentally in [33]. A schematic diagram of the experimental setup is illustrated in Fig. 6. The chosen beam (Specimen 3) has the following dimensions: length l = 350 mm; depth d = 100 mm; and thickness b = 50 mm. The span of the member is 250 mm and the depth of the notch is $\lambda = 0.5$ (half-notched). The properties of the concrete mixture are as follows: mean compressive strength $f_{cm,cyl}$ = 42.3 MPa, density ρ = 2346 kg/m^3 and the maximum aggregate diameter is 10 mm.

The compressive strength f_c is represented as a random field and all other properties are calculated using deterministic relations to f_c. Empirical formulas (derived from experimental data) published in *fib* Model Code 2010 [38] are used to determine the Young's modulus E, tensile strength f_t and fracture energy G_F. The correlation length l_c is set to 20 mm, and the Weibull modulus m is set to 3. The Weibull modulus m is an uncertain parameter with high sensitivity and a wide range of values can be found in the literature. According to the Weibull theory, the modulus m is a material property that is independent of the geometry and scale of the structure, however Syroka-Korol et al. [39] found that the Weibull modulus m does depend on the size of the structure and length scale l_c.

All the presented results have been obtained using a constant peridynamic horizon $\delta = 3.14\Delta x$ and regular grid spacing. Table 1 details the mesh resolution (Δx) and number of nodes (M) for every mesh level (ℓ). The non-linear softening model was calibrated to fit the experimental results for the smallest unnotched specimen (d = 50 mm). k (the rate of decay) is set to 25 and α (the position of the transition from exponential to linear decay) is set to 0.25. These parameters (k and α) are fixed for all test cases.

4.3.1. Results

We start by taking 100 samples on all levels and estimate α, β and γ. The first step is to estimate how the computational cost scales as M_ℓ increases. The time to compute each sample is recorded and it is determined that the computational costs grow linearly. The computational cost is given by Eq. (26), where $\gamma = 1$. Note that the performance of `PeriPy` scales linearly with the number of nodes $\therefore \gamma = 1$ [40].

$$C_\ell \approx M_\ell^1 \tag{26}$$

The next step is to estimate the parameters α and β for the QoI, which is taken to be the peak load. Fig. 7 illustrates the log-log plots of the estimated means and variances of Q_ℓ and $Y_\ell = Q_\ell - Q_{\ell-1}$, for $\ell = 0, \ldots, 4$, with respect to the number of degrees of freedom M_ℓ on each level. The rate of convergence of the discretisation error is given by Eq. (27), where α is approximately 0.528.

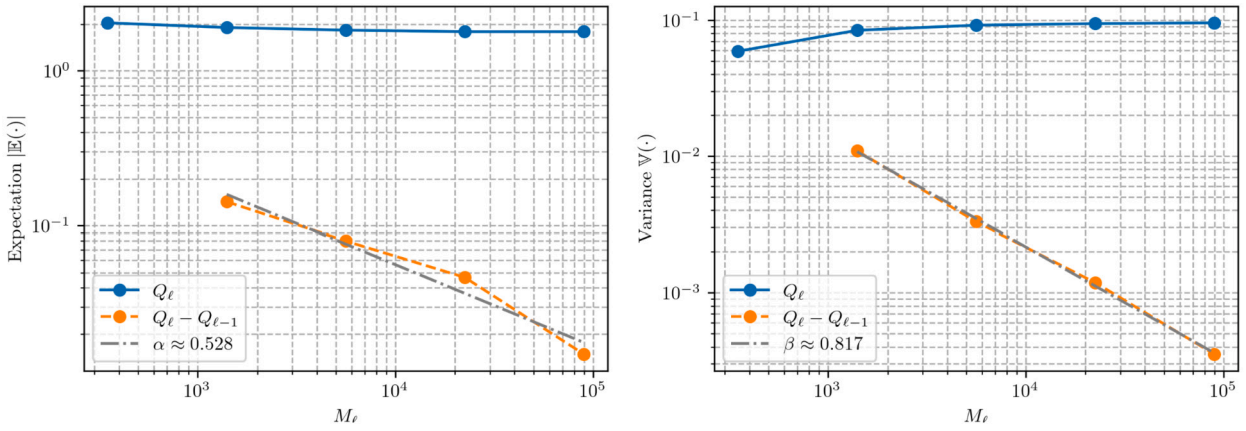

Fig. 7. Expectation (left) and variance (right) of Q_ℓ and $Y_\ell = Q_\ell - Q_{\ell-1}$ against degrees of freedom M_ℓ for problem 1. Taking 100 samples on every level, α is estimated to be 0.528 and β is estimated to be 0.817.

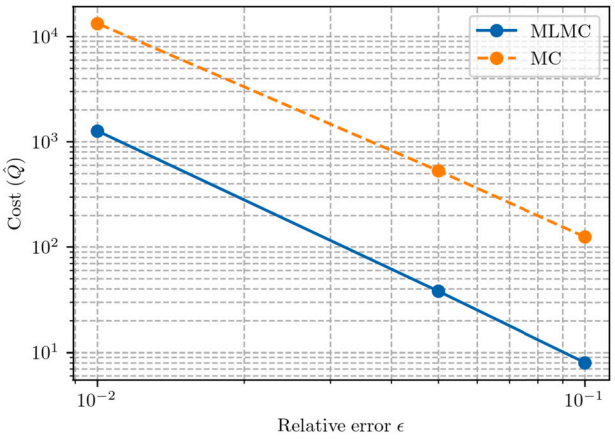

Fig. 8. Computational cost (in minutes) against ϵ for the multilevel estimator (cost $\approx \epsilon^{-2.35}$) and the standard MC estimator (cost $\approx \epsilon^{-3.89}$).

Table 2
Specimen 3 ($\lambda = 0.5$) - Sample allocation for different sampling tolerances ϵ_s. The sampling tolerance ϵ_s is given in Newtons (N). N_ℓ is computed using Eq. (23) and N is computed using Eq. (9). The bias error is approximately 0.75 N.

ϵ_s (N)	no. samples N_ℓ					N^a
	0	1	2	3	4	
100	21	5	2	0	0	9
50	86	22	8	2	0	38
10	2,152	567	204	67	18	955

[a] Number of samples required when using the standard Monte Carlo estimator. Note that all samples are computed on level 4 ($\ell = 4$).

$$\mathbb{E}[Y_\ell] \approx M_\ell^{-0.528} \tag{27}$$

The rate of convergence of the sampling error is given by Eq. (28), where β is approximately 0.817.

$$\mathbb{V}[Y_\ell] \approx M_\ell^{-0.817} \tag{28}$$

Using the estimated values of α, β and γ, Eq. (24) predicts that the cost of the MLMC simulations will grow proportionally to $\epsilon^{-2.35}$, whilst the cost of the standard MC simulations will grow proportionally to $\epsilon^{-3.89}$. Fig. 8 illustrates the computational cost (in minutes) of the multilevel estimator and the standard MC estimator against relative error ϵ. Table 2 presents the optimal number of samples N_ℓ across the 5 mesh levels for different values of sampling tolerance (ϵ_s = 10, 50 and 100 N), plus the number of samples

Table 3
Specimen 3 ($\lambda = 0.5$) - Computational cost (in minutes) for the multilevel estimator and the standard Monte Carlo estimator.

ϵ_s (N)	Cost (min)		Speed-up
	MLMC	MC	
100	8	125	16
50	38	529	14
10	1,265	13,290	10.5

required when using the standard MC estimator (N). In Table 3, the computational cost in minutes for the optimal number of sample is presented. Furthermore, using Eq. (20), the bias error is found to be relatively small at approximately 0.75 N.

Whilst the main aim of this work has been to demonstrate the computational savings of the MLMC framework, we have purposely selected examples where a deeper understanding of the physical behaviour can be gained by considering uncertainty. For Type 2 problems, the difference between the deterministic strength and the mean stochastic strength is expected to be small, as the influence of the variability in material properties is lessened due to the large stress concentrations that occur at the notch tip. This was observed numerically by Eliáš et al. [41] who found that considering spatial variability in material properties does not significantly influence the mean failure load, but does lead to an increase in the variance of the structural response. Using the finest mesh ($\ell = 4$), the deterministic model presented in [17] predicts that the specimen will fail at approximately 1800 N. Setting the sampling tolerance ϵ_s to 10 N, the mean stochastic strength is predicted to be approximately 1790 N. Note that the bias error is approximately 0.75 N. The observed results are in agreement with theory, which predicts that the difference between the deterministic strength and mean stochastic strength will be small [32,42]. Note that the experimental failure load ranged between 1580 N and 1710 N.

4.4. Case study 2: statistical size effect in quasi-brittle materials (Type 1)

The second problem that we consider is an unnotched concrete beam in three-point bending, tested experimentally in [33]. We consider Specimen 3 (illustrated in Fig. 6) again but with no notch ($\lambda = 0$). Beyond demonstrating the computational savings that can be achieved using the MLMC framework, the presented example provides insight into the following areas:

Statistical size effect - It was demonstrated in [17] that a deterministic bond-based model accurately captures the structural size effect for Type 2 (notched) problems, but fails to capture the correct response for Type 1 (unnotched) problems. This was expected as it is well known that the randomness of material properties has a significant effect on the structural strength of Type 1 problems [43,44]. In Type 1 problems, the volume of highly stressed material is much larger than that observed in Type 2 problems, and the probability that a defect is present in the stressed region is consequently higher. In Type 2 problems, the presence of a notch results in a localised region of highly stressed material, and the influence of randomness in material properties is consequently lessened. It is expected that the inclusion of statistical variability in the material properties will improve the predictive accuracy of the peridynamic model.

Convergence - It was demonstrated in [22] that a deterministic bond-based model fails to converge for Type 1 problems (the predicted strength is coupled with the mesh resolution). It was hypothesised that accounting for randomness in the material properties is required to initiate the localisation of damage and improve convergence.

4.4.1. Results

Again we start by taking 100 samples on all levels and estimate α, β and γ. As per the previous example, the computational cost grows linearly ($\gamma = 1$). Taking 100 samples on every level, α is estimated to be 0.337 and β is estimated to be 0.682 (refer to Fig. 9). The rate of convergence of the discretisation error and sampling error is slower than that observed in problem 1 (Type 2).

Using the estimated values of α, β and γ, Eq. (24) predicts that the cost of the MLMC simulations will grow proportionally to $\epsilon^{-2.94}$, whilst the cost of the standard MC simulations will grow proportionally to $\epsilon^{-4.97}$. Table 4 presents the optimal number of samples N_ℓ across the mesh levels for different values of sampling tolerance (ϵ_s = 10, 50 and 100 N), plus the number of samples required when using the standard MC estimator (N). In Table 5, the computational cost in minutes for the optimal number of sample is presented. Due to the higher variance of the estimator, the number of samples required is considerably higher than that required for the Type 2 problem. Type 1 problems are subject to a high degree of natural variability and consequently the computational cost is higher as significantly more samples are required.

By including uncertainty in the material properties, the bond-based model converges for Type 1 problems ($\alpha \approx 0.337$). This is the first time that this behaviour has been demonstrated, but the convergence behaviour is significantly worse than that observed for Type 2 problems ($\alpha \approx 0.528$) and the bias error remains large. Using Eq. (20), the bias error is estimated to be approximately 200 N.

Initial results using a normal distribution found that the model failed to converge for Type 1 problems. It was established that extreme values in the left-tail of the material strength distribution are required to initiate the localisation of damage and eliminate problems of mesh dependence. If there are no imperfections in the highly stressed region then the damage fails to localise correctly and the results exhibit a strong mesh dependency. Even when using a Weibull distribution, there will be a small number of samples where the damage fails to localise and this has a detrimental impact on the estimated convergence rate α. The convergence rate α

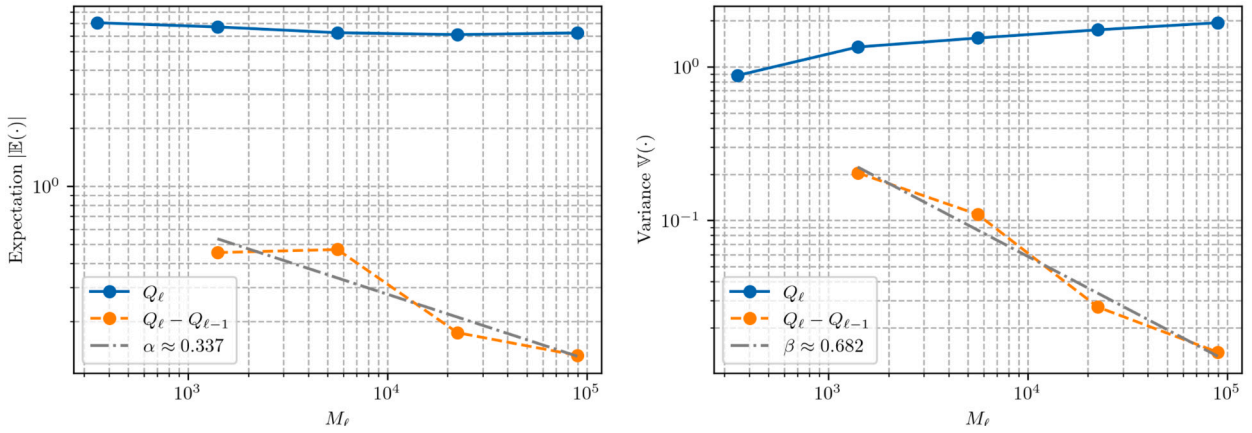

Fig. 9. Expectation (left) and variance (right) of Q_ℓ and $Y_\ell = Q_\ell - Q_{\ell-1}$ against degrees of freedom M_ℓ for problem 2. Taking 100 samples on every level, α is estimated to be 0.337 and β is estimated to be 0.682.

Table 4
Specimen 3 ($\lambda = 0$) - Sample allocation for different sampling tolerances ϵ_s. The sampling tolerance ϵ_s is given in Newtons (N). N_ℓ is computed using Eq. (23) and N is computed using Eq. (9). Due to the higher variance of the estimator, the number of samples required is considerably higher than that required for the Type 2 problem (refer to Table 2).

ϵ_s (N)	no. samples N_ℓ					N[a]
	0	1	2	3	4	
100	406	119	57	15	5	193
50	1,624	479	230	63	22	775
10	40,624	11,982	5,765	1,586	551	19,388

[a] Number of samples required when using the standard Monte Carlo estimator. Note that all samples are computed on level 4 ($\ell = 4$).

Table 5
Specimen 3 ($\lambda = 0$) - Computational cost (in minutes) for the multilevel estimator and the standard Monte Carlo estimator. The computational cost of the multilevel estimator grows proportionally to $\epsilon^{-2.94}$ and the cost of standard Monte Carlo estimator grows proportionally to $\epsilon^{-4.97}$.

ϵ_s (N)	Cost (min)		Speed-up
	MLMC	MC	
100	290	2,686	9.3
50	1,197	10,785	9
10	29,998	269,816	9

can be improved by employing a material strength distribution that is skewed towards the left (e.g. Weibull distribution with a low Weibull modulus) but this might not be physically realistic for the considered problem.

As the size of a structure increases, so does the probability that a defect will be present from which a fracture will initiate. Syroka-Korol et al. [43] determined numerically that the deterministic and mean stochastic strength start to diverge when the beam depth is greater than 50-60 mm. Specimen 3 is 100 mm deep and the magnitude of the statistical size effect is expected to be non-negligible. Setting the sampling tolerance ϵ_s to 50 N, the mean stochastic strength is estimated to be approximately 6250 N. Note that the bias error is approximately 200 N. Using the finest mesh ($\ell = 4$), the deterministic model predicts that the specimen will fail at approximately 9200 N. The experimental failure load ranged between 7620 N and 8770 N. The numerical results are consistent with the theory, i.e., the difference between the deterministic strength and mean stochastic strength is much larger than that observed for Type 2 problems. However, the deterministic model does not converge for Type 1 problems and the prediction of strength is therefore unreliable, and a rigorous comparison is not possible.

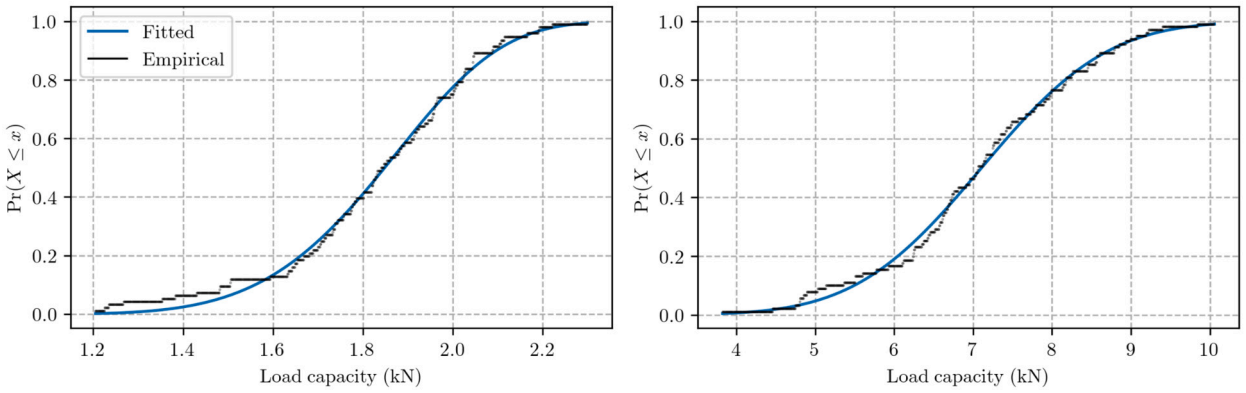

Fig. 10. The cumulative distribution function of strength (load capacity) for the Type 2 problem (left) and the Type 1 problem (right). The empirical CDF is plotted in black and the fitted Weibull CDF is plotted in blue. The y-axis represents the percentage of the population that will fail at a load less than or equal to x.

4.5. Estimating failure probabilities

The objective of the multilevel framework is to estimate the expectation of an output variable, in this case, the peak load. However, for many industrial applications, engineers are more concerned with the cumulative distribution function (CDF) of the output variable. For example, an engineer might be interested in the probability that an output variable exceeds a specific value, or as demonstrated here, the probability that a structure will fail at a load less than or equal to a specific value. Computing the CDF is complicated as the multilevel method provides relatively few values on the finest mesh. Gregory and Cotter [45] recently outlined a method that makes it possible to obtain the CDF using samples obtained on multiple mesh levels and we follow the same approach. The reader is also referred to Clare et al. [15] for further information.

The inverse transform sampling method is used to determine an approximation of the inverse CDF $F^{-1}(u)$, where $u \sim \mathcal{U}(0,1)$. If the CDF F of a random variable X is strictly increasing and absolutely continuous, then there exists a unique value $x \equiv F^{-1}(u)$ for which $F(x) = u$. By sorting the samples $\{X^i\}_{i=1,\dots,N} \sim F$ in ascending order such that $X^1 < X^2 < \cdots < X^N$, a simple consistent estimate for an evaluation to the quantile function of the distribution with CDF F is

$$\hat{F}^{-1}(u) = X^{[N\times u]}$$

This is a consistent estimate because it converges in probability to $F^{-1}(u)$ as $N \to \infty$. The inverse CDF for the multilevel approximation is then given by

$$\hat{F}_L^{-1}(u) = R(X)_0^{[N_0\times u]} + \sum_{i=1}^{L}\left(R(X)_\ell^{[N_\ell\times u]} - R(X)_{\ell-1}^{[N_{\ell-1}\times u]}\right)$$

where $R(X)_\ell^i$ is the ith ordered statistic of X_ℓ on level ℓ. The CDF of the multilevel approximation is then given by

$$\hat{F}(x) = \frac{1}{N}\sum_{i=1}^{N}\mathbb{1}_{X_i\le x}$$

where $\mathbb{1}$ is the indicator function. For a detailed description of this method we refer the reader to [45].

The empirical CDF and fitted Weibull CDF for the Type 1 and Type 2 problem is illustrated in Fig. 10. The y-axis represents the percentage of the population that will fail at a load less than or equal to x. The purpose of this figure is to demonstrate that is possible to use the multilevel framework to obtain the CDF and a discussion of the results is beyond the scope of this paper. The reader is referred to Bažant and Le [46] for further information on the computation of CDFs of structural strength and fatigue lifetime. Ideally we would include the CDF obtained using standard MC but this was not possible due to the high computational cost of obtaining samples on the finest level.

5. Further discussion

Beyond demonstrating the computational savings that can be achieved using the multilevel framework, the results presented in this paper are of interest for further reasons as discussed in this section.

5.1. Statistical size effect

A key aim of this study was to select case studies where uncertainty must be considered to gain a comprehensive understanding of the physical behaviour. We focused our studies on the structural size effect in quasi-brittle materials. Bažant [4] stated that the correct modelling of the size effect on material strength should be adopted as the basic criterion of acceptability of any model. The

results demonstrate that a bond-based peridynamic model can be used to examine both the statistical and deterministic component of the structural size effect. The intention of this study was never to provide a detailed examination of the statistical size effect, and further studies on a wider range of problems are required to improve confidence in the models predictive capabilities.

By employing the presented MLMC framework, studying the statistical component of the structural size effect using a peridynamic model becomes computationally feasible. Future work aims to employ the presented MLMC framework to study the full series of tests published in [33] and provide a detailed examination of influential factors, such as the shape of the material strength distribution and the correlation length l_c. It is stated in [34] that the ratio of the correlation length l_c to the size of the fracture process zone (FPZ) is the main parameter that influences the statistical size effect.

5.2. Convergence

This work provides new insights into the convergence behaviour of bond-based peridynamic models. It is a fundamental test of the adequacy of any numerical model that the predictions are independent of the mesh resolution. Comprehensive convergence studies of the predicted structural response are missing from the peridynamics literature, with existing investigations primarily focusing on static elastic problems [47]. Niazi et al. [48] did perform a convergence study that examined the entire structural response, but this study is limited as Type 1 problems were not considered.

To the best of the authors knowledge, [22] provides the first study on the effect of mesh refinement (δ-convergence) on the predicted peak load and load-deflection response for Type 1 and 2 problems. The deterministic bond-based peridynamic model failed to converge for Type 1 problems, and it was hypothesised that a source of randomness must be introduced to trigger the localisation of damage in Type 1 specimens, and eliminate problems of mesh dependence.

The results in this study confirm that by introducing a source of randomness, problems of mesh dependence are eliminated for Type 1 specimens. It was previously reported in [48] that the convergence behaviour is improved by randomly deleting 1% of all bonds, as first suggested by Chen et al. [49]. While this method is computationally cheap and does improve convergence behaviour, it is an oversimplified heuristic approach that lacks a robust theoretical basis, and does not consider the spatial correlation of material properties. Jones et al. [50] note that these heuristic methods are generally used to avoid problems related to symmetry, and they do not attempt to capture the true material behaviour by implementing an experimentally measured probability distribution of material properties.

5.3. Length scales

The correlation length l_c was set to be 20 mm for all considered problems. This value was selected after running a number of preliminary simulations. However, the aim of this contribution was not to identify the parameters that describe the spatial fields. It is important to note that a theoretically grounded probabilistic framework based on Bayesian inference (see [36,37,51]) is essential to identify the parameters of the spatial fields (e.g. length scale l_c) rigorously.

Furthermore, the interaction between the two length scales (peridynamic horizon δ and the correlation length l_c in the random field) requires further examination. It remains uncertain how the ratio of the two length scales influences the predictive accuracy of the model.

5.4. Material strength distribution

The material strength distribution plays an important role in the predicted results and convergence of the model. Three distributions were considered (normal, log-normal and Weibull) and it was determined that the Weibull distribution provides the best predictions of mean strength for quasi-brittle materials. This was expected and has been extensively discussed in the literature. A more novel observation is that the selected probability distribution influences the convergence rate of the bias error. Extreme values in the left-tail are required to initiate the localisation of damage and eliminate problems of mesh dependence. Note that the model failed to converge for Type 1 problems when using a normal distribution.

5.5. Model calibration

Many of the model parameters are impossible to determine exactly and are subject to significant uncertainties, for example: the length scale l_c and the Weibull modulus (shape parameter). Future work will examine the integration of the multilevel method with experimental data in a Bayesian setting to quantify modelling uncertainties as proposed by Dodwell et al. [52,53]. This will be an important step in the validation of peridynamic models, enabling the identification of model discrepancy and measurement bias, and providing better estimates of model parameters.

6. Conclusions

Peridynamic models are computationally expensive, preventing the use of standard Monte Carlo methods for the assessment of uncertainties in model outputs propagated from uncertain inputs. The aim of this study was to demonstrate the possible computational savings that can be realised using a multilevel framework. The maximum observed speed-up factor is 16 when compared to

a standard Monte Carlo estimator, thus enabling the efficient forward propagation of uncertain parameters in a computationally expensive peridynamic model. Beyond demonstrating the computational savings that can be achieved using the multilevel framework, the results presented in this paper are of interest for two further reasons:

1. Deterministic bond-based models suffer from a strong mesh dependency when simulating Type 1 problems. It has been demonstrated that by including uncertainty in the material properties, the bond-based peridynamic model converges for both Type 1 and Type 2 problems. The need to consider uncertainty is essential for robust and accurate predictions. Furthermore, the multilevel method provides an estimate of the discretisation error, thus improving the interpretability of numerical predictions.
2. A secondary aim was to select case studies where uncertainty must be considered to gain a comprehensive understanding of the physical behaviour. We examined the structural size effect in quasi-brittle materials as the random variability of material properties is known to play an important role. The correct modelling of the size effect on material strength should be adopted as the basic criterion of acceptability of any model. The results demonstrate that a bond-based peridynamic model can be used to study the statistical size effect but further studies on a wider range of problems are required to improve confidence in the models predictive capabilities.

Peridynamic models have many advantages for modelling complex failure problems, but due to their high computational cost, addressing probabilistic problems has been prohibitively expensive. This work has demonstrated a multilevel framework that makes such computations feasible and opens the door to tackling other probabilistic problems. We motivated the use of the presented framework by studying the statistical size effect in quasi-brittle materials, but forward uncertainty quantification is equally important for cases where a high degree of reliability is required, as is common in many aerospace and power generation applications.

CRediT authorship contribution statement

Mark Hobbs: Conceptualisation, Methodology, Software, Validation, Formal analysis, Writing - Original Draft, Writing - Review & Editing. **Hussein Rappel**: Conceptualisation, Formal analysis, Supervision, Writing - Original Draft, Writing - Review & Editing. **Tim Dodwell**: Conceptualisation, Methodology, Writing - Original Draft, Supervision, Funding acquisition.

Declaration of competing interest

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

Data availability

No data was used for the research described in the article.

Acknowledgements

TD and MH are funded through a Turing AI Fellowship, United Kingdom (2TAFFP\100007). The authors would like to acknowledge the use of the University of Exeter High-Performance Computing (HPC) facility in carrying out this work.

References

[1] X.T. Su, Z.J. Yang, G.H. Liu, Monte Carlo simulation of complex cohesive fracture in random heterogeneous quasi-brittle materials: a 3D study, Int. J. Solids Struct. 47 (2010) 2336–2345.
[2] J.-Y. Wu, J.-R. Yao, J.-L. Le, Phase-field modeling of stochastic fracture in heterogeneous quasi-brittle solids, Comput. Methods Appl. Mech. Eng. 416 (2023) 116332.
[3] C. Song, R. Kawai, Monte Carlo and variance reduction methods for structural reliability analysis: a comprehensive review, Probab. Eng. Mech. 73 (2023) 103479.
[4] Z.P. Bažant, Why continuum damage is nonlocal: micromechanics arguments, J. Eng. Mech. 117 (1991) 1070–1087.
[5] Z.P. Bažant, Size effect on structural strength: a review, Arch. Appl. Mech. 69 (1999) 703–725.
[6] S.A. Silling, Reformulation of elasticity theory for discontinuities and long-range forces, J. Mech. Phys. Solids 48 (2000) 175–209.
[7] S. Heinrich, E. Sindambiwe, Monte Carlo complexity of parametric integration, J. Complex. 15 (1999) 317–341.
[8] S. Heinrich, Multilevel Monte Carlo methods, in: S. Margenov, J. Waśniewski, P. Yalamov (Eds.), Large-Scale Scientific Computing, in: Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, 2001, pp. 58–67.
[9] M.B. Giles, Multilevel Monte Carlo path simulation, Oper. Res. 56 (2008) 607–617.
[10] K.A. Cliffe, M.B. Giles, R. Scheichl, A.L. Teckentrup, Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients, Comput. Vis. Sci. 14 (2011) 3.
[11] D. Elfverson, F. Hellman, A. Målqvist, A multilevel Monte Carlo method for computing failure probabilities, SIAM/ASA J. Uncertain. Quantificat. 4 (2016) 312–330.
[12] T.J. Dodwell, et al., Multilevel Monte Carlo simulations of composite structures with uncertain manufacturing defects, Probab. Eng. Mech. 63 (2021) 103116.
[13] K.M. Hamdia, H. Ghasemi, Reliability analysis of the stress intensity factor using multilevel Monte Carlo methods, Probab. Eng. Mech. 74 (2023) 103497.
[14] D. Crevillén-García, H. Power, Multilevel and quasi-Monte Carlo methods for uncertainty quantification in particle travel times through random heterogeneous porous media, R. Soc. Open Sci. 4 (2017) 170203.

[15] M.C.A. Clare, M.D. Piggott, C.J. Cotter, Assessing erosion and flood risk in the coastal zone through the application of multilevel Monte Carlo methods, Coast. Eng. 174 (2022) 104118.
[16] M.B. Giles, Multilevel Monte Carlo methods, Acta Numer. 24 (2015) 259–328.
[17] M. Hobbs, T. Dodwell, G. Hattori, J. Orr, An examination of the size effect in quasi-brittle materials using a bond-based peridynamic model, Eng. Struct. 262 (2022) 114207.
[18] S.A. Silling, M. Epton, O. Weckner, J. Xu, E. Askari, Peridynamic states and constitutive modeling, J. Elast. 88 (2007) 151–184.
[19] S.A. Silling, E. Askari, A meshfree method based on the peridynamic model of solid mechanics, Comput. Struct. 83 (2005) 1526–1535.
[20] S.A. Silling, R.B. Lehoucq, Peridynamic theory of solid mechanics, in: H. Aref, E.v.d. Giessen (Eds.), Advances in Applied Mechanics, in: Advances in Applied Mechanics, vol. 44, Elsevier, 2010, pp. 73–168.
[21] F. Bobaru, J. Foster, P. Geubelle, S. Silling, Handbook of Peridynamic Modeling, 1st ed., Chapman and Hall/CRC, New York, 2017.
[22] M.C. Hobbs, Three-dimensional peridynamic modelling of quasi-brittle structural elements, Ph.D. thesis, University of Cambridge, 2021.
[23] F. Bobaru, et al., Convergence, adaptive refinement, and scaling in 1D peridynamics, Int. J. Numer. Methods Eng. 77 (2009) 852–877.
[24] Y.D. Ha, F. Bobaru, Studies of dynamic crack propagation and crack branching with peridynamics, Int. J. Fract. 162 (2010) 229–244.
[25] E. Madenci, E. Oterkus, Peridynamic Theory and Its Applications, Springer, New York, 2014.
[26] F. Bobaru, W. Hu, The meaning, selection, and use of the peridynamic horizon and its relation to crack branching in brittle materials, Int. J. Fract. 176 (2012) 215–222.
[27] W. Hu, Y.D. Ha, F. Bobaru, Numerical integration in peridynamics, Technical Report, University of Nebraska-Lincoln, 2010.
[28] P. Seleson, Improved one-point quadrature algorithms for two-dimensional peridynamic models based on analytical calculations, Comput. Methods Appl. Mech. Eng. 282 (2014) 184–217.
[29] Z.P. Bažant, J. Planas, Fracture and Size Effect in Concrete and Other Quasibrittle Materials, Routledge, New York, 1998.
[30] Z.P. Bažant, Size effect, Int. J. Solids Struct. 37 (2000) 69–80.
[31] J.G.M.v. Mier, Fracture Processes of Concrete, CRC Press, 1996.
[32] Z.P. Bažant, Design of quasibrittle materials and structures to optimize strength and scaling at probability tail: an aperçu, Proc. R. Soc. A, Math. Phys. Eng. Sci. 475 (2019) 20180617.
[33] D. Grégoire, L.B. Rojas-Solano, G. Pijaudier-Cabot, Failure and size effect for notched and unnotched concrete beams, Int. J. Numer. Anal. Methods Geomech. 37 (2013) 1434–1452.
[34] P. Grassl, Z.P. Bažant, Random lattice-particle simulation of statistical size effect in quasi-brittle structures failing at crack initiation, J. Eng. Mech. 135 (2009) 85–92.
[35] Z.P. Bažant, G. Pijaudier-Cabot, Measurement of characteristic length of nonlocal continuum, J. Eng. Mech. 115 (1989) 755–767.
[36] H. Rappel, L. Wu, L. Noels, L.A.A. Beex, A Bayesian framework to identify random parameter fields based on the copula theorem and Gaussian fields: application to polycrystalline materials, J. Appl. Mech. 86 (2019).
[37] H. Rappel, M. Girolami, L.A. Beex, Intercorrelated random fields with bounds and the Bayesian identification of their parameters: application to linear elastic struts and fibers, Int. J. Numer. Methods Eng. 123 (2022) 3418–3463.
[38] fib, *fib* Model Code for Concrete Structures 2010, Technical Report, Fédération Internationale du Béton (*fib*), Lausanne, Switzerland, 2010.
[39] E. Syroka-Korol, J. Tejchman, Z. Mróz, FE investigations of the effect of fluctuating local tensile strength on coupled energetic–statistical size effect in concrete beams, Eng. Struct. 103 (2015) 239–259.
[40] B. Boys, T.J. Dodwell, M. Hobbs, M. Girolami, PeriPy - a high performance OpenCL peridynamics package, Comput. Methods Appl. Mech. Eng. 386 (2021) 114085.
[41] J. Eliáš, M. Vorechovský, J. Skocek, Z.P. Bažant, Stochastic discrete meso-scale simulations of concrete fracture: comparison to experimental data, Eng. Fract. Mech. 135 (2015) 1–16.
[42] Z.P. Bažant, Size effect in blunt fracture: concrete, rock, metal, J. Eng. Mech. 110 (1984) 518–535.
[43] E. Syroka-Korol, J. Tejchman, Z. Mróz, FE calculations of a deterministic and statistical size effect in concrete under bending within stochastic elasto-plasticity and non-local softening, Eng. Struct. 48 (2013) 205–219.
[44] J. Eliáš, M. Vorechovský, Fracture in random quasibrittle media: I. Discrete mesoscale simulations of load capacity and fracture process zone, Eng. Fract. Mech. 235 (2020) 107160.
[45] A. Gregory, C.J. Cotter, On the calibration of multilevel Monte Carlo ensemble forecasts, Q. J. R. Meteorol. Soc. 143 (2017) 1929–1935.
[46] Z.P. Bažant, J.-L. Le, Probabilistic Mechanics of Quasibrittle Structures: Strength, Lifetime, and Size Effect, Cambridge University Press, 2017.
[47] P. Seleson, D.J. Littlewood, Convergence studies in meshfree peridynamic simulations, Comput. Math. Appl. 71 (2016) 2432–2448.
[48] S. Niazi, Z. Chen, F. Bobaru, Crack nucleation in brittle and quasi-brittle materials: a peridynamic analysis, Theor. Appl. Fract. Mech. 112 (2021) 102855.
[49] Z. Chen, S. Niazi, F. Bobaru, A peridynamic model for brittle damage and fracture in porous materials, Int. J. Rock Mech. Min. Sci. 122 (2019) 104059.
[50] L.D. Jones, L.J. Vandeperre, T.A. Haynes, M.R. Wenman, Theory and application of Weibull distributions to 1D peridynamics for brittle solids, Comput. Methods Appl. Mech. Eng. 363 (2020) 112903.
[51] H. Rappel, L.A.A. Beex, J.S. Hale, L. Noels, S.P.A. Bordas, A tutorial on Bayesian inference to identify material parameters in solid mechanics, Arch. Comput. Methods Eng. 27 (2020) 361–385.
[52] T.J. Dodwell, C. Ketelsen, R. Scheichl, A.L. Teckentrup, A hierarchical multilevel Markov chain Monte Carlo algorithm with applications to uncertainty quantification in subsurface flow, SIAM/ASA J. Uncertain. Quantificat. 3 (2015) 1075–1108.
[53] T.J. Dodwell, C. Ketelsen, R. Scheichl, A.L. Teckentrup, Multilevel Markov chain Monte Carlo, SIAM Rev. 61 (2019) 509–545.