Contents lists available at ScienceDirect



Journal of Computational Physics

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



Gradient-based optimisation of the conditional-value-at-risk using the multi-level Monte Carlo method

Sundar Ganesh, Fabio Nobile*

Institute of Mathematics, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland

ARTICLE INFO

Keywords: Multilevel Monte Carlo methods VaR CVaR Uncertainty quantification Optimisation under uncertainty Gradient descent

ABSTRACT

In this work, we tackle the problem of minimising the Conditional-Value-at-Risk (CVaR) of output quantities of complex differential models with random input data, using gradient-based approaches in combination with the Multi-Level Monte Carlo (MLMC) method. In particular, we consider the framework of multi-level Monte Carlo for parametric expectations introduced in [24] and propose modifications of the MLMC estimator, error estimation procedure, and adaptive MLMC parameter selection to ensure the estimation of the CVaR and sensitivities for a given design with a prescribed accuracy. We then propose combining the MLMC framework with an alternating inexact minimisation-gradient descent algorithm, for which we prove Q-linear convergence in the optimisation iterations under the assumptions of strong convexity and Lipschitz continuity of the gradient of the objective function. We demonstrate the performance of our approach on two numerical examples of practical relevance, which evidence the same optimal asymptotic cost-tolerance behaviour as standard MLMC methods for fixed design computations of output expectations.

1. Introduction

Optimisation algorithms play an important role across various scientific and engineering fields as valuable design tools. The key goal of optimisation is to find the best values of certain parameters (design variables) of a model, typically a differential model such as a Partial Differential Equation (PDE), used to predict the behaviour of a certain system, such that a desired output Quantity of Interest (QoI) of the model is optimised. Such differential models usually also include various other input parameters besides the design variable, which may or may not be fully characterised. There is an increasing interest in the computational science and engineering community to treat such parameters as random variables to reflect their uncertainty, either due to a lack of knowledge or to some intrinsic variability. As a result, the output QoI being optimised also becomes a random variable. Naively optimising the system for only one particular value of the inputs (e.g., the nominal value) can lead to a design that is not robust enough to the uncertainties in the system. A classical example is of civil engineering structures designed to minimise structural loads for moderate wind conditions, which are then unable to withstand local wind gusts or storms [22].

The field of PDE-constrained Optimization Under Uncertainty (OUU) seeks to characterise the randomness of the output QoI of the PDE using summary statistics such as moments, quantiles, etc., and optimise the summary statistic instead of the QoI directly. In particular, in risk-averse PDE-constrained optimisation, one aims at favouring designs with acceptable performance also in extreme

* Corresponding author.

E-mail addresses: sundar.ganesh@epfl.ch (S. Ganesh), fabio.nobile@epfl.ch (F. Nobile).

https://doi.org/10.1016/j.jcp.2023.112523

Available online 6 October 2023

Received 19 October 2022; Received in revised form 18 July 2023; Accepted 25 September 2023

^{0021-9991/© 2023} The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

conditions. In this case, the summary statistic, often called a risk-measure, should quantify the importance that is given in the design process to unfavourable scenarios. The reader is referred to [38] for a comprehensive review of several risk-measures and their main properties. An important class of risk-measures is that of coherent risk-measures [36,4], which exhibit favourable properties such as monotonicity and convexity.

In this work, we focus on the so-called Conditional-Value-at-Risk (CVaR) [35], which corresponds to the expectation of the output QoI conditional on being above a given quantile (referred to as the Value-at-Risk (VaR) in the finance literature), and is a widely used coherent risk-measure. To be more specific, let $Q(z) \in \mathbb{R}$ denote the random QoI, which depends on the design parameter $z \in \mathbb{R}^d$. We denote by $\mathbb{E}[X]$ the expected value of a random variable X, and by $c_{\tau}(z)$ the CVaR of Q(z) of significance $\tau \in [0, 1]$, i.e., $c_{\tau}(z) := \mathbb{E}[Q(z)|Q(z) \ge q_{\tau}]$, where q_{τ} is the τ -quantile of Q(z). It was demonstrated in [35] that $c_{\tau}(z)$ could be written in the following form under certain conditions on the distribution of Q(z):

$$c_{\tau}(z) = \min_{\theta \in \mathbb{R}} \left\{ \Phi(\theta; z) := \mathbb{E}\left[\phi(\theta, Q(z)) \right] \right\}, \quad \phi(\theta, Q) := \theta + \frac{(Q - \theta)^+}{1 - \tau}, \tag{1}$$

where we denote with X^+ the positive part of X; namely $X^+ := \max(0, X)$. In this work, we consider the following problem of penalised CVaR minimisation:

$$\mathcal{J}^{*} = \min_{z \in \mathbb{R}^{d}} \left\{ c_{\tau}(z) + \kappa \left\| z - z_{ref} \right\|_{l^{2}}^{2} \right\},$$

$$= \min_{\substack{z \in \mathbb{R}^{d} \\ \theta \in \mathbb{R}}} \left\{ \mathcal{J}(\theta, z) := \Phi(\theta; z) + \kappa \left\| z - z_{ref} \right\|_{l^{2}}^{2} \right\},$$
(2)
(3)

where we have added a term penalising deviation of the design z from a preferred design z_{ref} . In Eq. (3), the parameter κ controls the strength of the penalisation, and $\|\cdot\|_{l^2}$ denotes the Euclidean norm. In particular, we will consider Monte Carlo type approximations of problem (3), and since evaluating the objective function and its sensitivities at a given design z requires the solution of a costly PDE many times, we will accelerate the Monte Carlo estimation by multilevel strategies following the well established Multi-Level Monte Carlo (MLMC) paradigm [15], which has been shown to provide significant performance improvements in comparison to classical Monte Carlo methods for estimating various summary statistics of output QoI of differential models [16,15,19,25,12].

Two broad approaches can be used to solve problem (3); namely, gradient-free and gradient-based methods. Evolutionary algorithms, a type of gradient-free method, were used in combination with Monte Carlo estimators for PDE-constrained CVaR minimisation in [34,33]. A genetic algorithm was also used in combination with MLMC estimators in [32]. Multiple different risk-measures, including the CVaR, were estimated, and the framework was applied to aerodynamic shape optimisation problems. The authors of [11] proposed a multifidelity Monte Carlo estimator for the CVaR based on cross-entropy methods combined with importance sampling. This approach was used in [31] in combination with gradient-free optimisation algorithms to minimise the CVaR. However, gradient-free algorithms typically have slower rates of convergence in comparison to gradient-based methods and involve multiple expensive evaluations of the objective function. We propose instead the use of gradient based algorithms, combined with MLMC estimators, to compute sensitivities of the objective function in problem (3). In particular, the MLMC estimators developed in this work rely on the framework of parametric expectations [24] and extend the work in [5] to the computation of CVaR sensitivities, addressing the corresponding challenges as outlined hereafter. We mention that gradient-based methods have also been used for other risk-measures and sampling strategies in the context of PDE-constrained OUU [17,40,29].

The computation of the sensitivities of the CVaR $c_r(z)$ with respect to the extended design variables z and θ typically requires the estimation of expectations of the form $\mathbb{E}\left[(Q(z) - \theta)^+\right]$ and $\mathbb{E}\left[\mathbb{1}_{Q(z) \ge \theta} f(z)\right]$ for suitable design-dependent random variables f(z). Although $\mathbb{E}\left[(Q(z) - \theta)^+\right]$ and $\mathbb{E}\left[\mathbbm{1}_{Q(z) \ge \theta} f(z)\right]$ can be shown to be differentiable in θ and z [21,20] under some conditions on the distribution of Q(z), sample- or quadrature-based approximations of these expectations are typically not differentiable and may require some additional treatment. One possibility is to directly use the non-differentiable estimations in combination with non-smooth optimisation techniques that use sub-gradient information. For example, the work in [26] uses a combination of smooth and non-smooth optimisation techniques, using sub-gradients computed using Monte Carlo estimators, to minimise the CVaR. Alternatively, one could construct smoothed versions of the maximum/indicator functions, with sufficient regularity such that sample-based approximations are still differentiable. For example, a regularised version of the CVaR was constructed in [23], with second order differentiability, and optimised successfully using a trust-region method. However, although regularised or smoothed versions of the CVaR can be constructed with adequate differentiability, this property is lost in the limit of vanishing smoothing, as is required when the algorithm is close to the optimum. The method proposed in [24,5] offers an alternative to CVaR regularisation. In these works, the quantity $\mathbb{E}\left[(Q(z)-\theta)^+\right]$ is estimated directly using an MLMC estimator at a set of points in θ , all sharing the same realisations of Q(z), followed by a cubic spline interpolation over the pointwise evaluations thus obtained. Derivatives such as $\mathbb{E}\left[\mathbb{1}_{Q(z)\geq\theta}f(z)\right]$ are then approximated using derivatives of the cubic spline. We propose to follow the above path in this work. As was discussed in [5], directly using a naive MLMC estimator to estimate $\mathbb{E}\left[\mathbbm{1}_{Q(z)\geq \theta}f(z)\right]$ causes non-optimal MLMC complexity behaviour. By constructing an MLMC estimator of $\mathbb{E}\left[(Q(z) - \theta)^+ f(z)\right]$ and numerically differentiating in θ instead, the approach in [5] ameliorates this issue and preserves the same optimal complexity behaviour of the MLMC method as predicted for estimating $\mathbb{E}[Q(z)]$. Lastly, since the MLMC estimator proposed in [24,5] automatically provides an approximation $\hat{J}(\cdot, z)$ of the function $\theta \mapsto \mathcal{J}(\theta, z)$ at a given design z, we propose in this work to use an optimisation algorithm in which, at each iteration, gradient steps are taken only in the design variable z, whereas exact optimisation in θ is performed using the surrogate $\hat{J}(\cdot,z)$. Such an algorithm, introduced in [6], was applied in combination with the Monte Carlo estimation of a regularised version of the CVaR in [7].

(4)

The main contributions of this work are as follows. We propose novel expressions for the sensitivity of the objective function defined in Eq. (3) in terms of parametric expectations, thus allowing us to use and extend the framework in [5] to build cost optimal adaptive MLMC estimators for those sensitivities with error control. We then propose to use MLMC sensitivity estimators within an Alternating Minimisation-Gradient Descent (AMGD) algorithm, analogous to the one proposed in [6,7], where gradient steps are taken in the design variable z whereas exact optimisation is performed in θ using an MLMC-constructed surrogate \hat{J} of \mathcal{J} . The accuracy of the surrogate and sensitivity estimation is increased over the optimisation iterations and is set proportional to the gradient norm. Following closely the analysis in [7], we propose a convergence result for our algorithm under the assumption that the objective function $\mathcal{J}(\theta, z)$ is strongly convex with Lipschitz continuous gradients.

The structure of this paper is as follows. We present the problem formulation in Section 2, for a problem of penalised CVaR minimisation of the form in Eq. (3). The novel expression for the gradients in terms of the parametric expectations is also presented in Section 2. In Section 3, we propose the AMGD algorithm with inexact gradient and objective function estimation and demonstrate its convergence. Section 4 discusses the novel MLMC estimators, error estimation procedure, and adaptive Continuation MLMC (CMLMC)-type hierarchy selection for the gradients of $\mathcal{J}(\theta, z)$. In addition, it presents a final CMLMC-AMGD algorithm. Lastly, in Section 5, we demonstrate the above optimisation algorithm and MLMC procedure on two problems of interest. The first is a two-dimensional oscillator, typically used to model oscillatory phenomena in excitable media. The second is a more applied problem of pollutant transport modelling. We demonstrate that the procedure proposed in this work performs well and reflects the theoretical results presented in Sections 2 and 3.

2. Problem formulation

Let $(\Omega, \mathcal{F}, \mathbb{P})$ denote a complete probability space, $\omega \in \Omega$ an elementary random event and $z \in \mathbb{R}^d$ the vector of design variables. We denote by $Q(z, \omega) \in \mathbb{R}$ the random QoI, typically a functional of the solution to an underlying differential model with random input ω and design z. We are interested in minimising the CVaR $c_{\tau}(z)$ of the random variable $Q(z, \cdot)$ over the designs $z \in \mathbb{R}^d$, as indicated in Eq. (2), following the formulation presented in [35]. To this end, we first introduce the following assumptions on the random variable $Q(z, \cdot)$.

Assumption 1. For any $z \in \mathbb{R}^d$:

- (i) $Q(z, \cdot)$ is a random variable in $L^p(\Omega, \mathbb{R})$ for some $p \in [1, \infty)$.
- (ii) The measure of Q(z, ·) admits a probability density function, i.e., the measure of Q(z, ·) is free of atoms. We denote by Γ the subset of random variables in L^p(Ω, ℝ) that are free of atoms, and hence, Q(z, ·) ∈ Γ ⊂ L^p(Ω, ℝ).
- (iii) There exists a positive random variable K, possibly dependent on z, such that $\mathbb{E}[K] < \infty$ and

 $|Q(z + \Delta z, \cdot) - Q(z, \cdot)| \le K(\cdot) \|\Delta z\|_{l^2},$

for any $\Delta z \in \mathbb{R}^d$ close enough to 0 (restated here from [21,20]).

(iv) For almost every $\omega \in \Omega$, the mapping $z \mapsto Q(z, \omega)$ is differentiable in \mathbb{R}^d and the corresponding vector of partial derivatives $Q_z(z, \cdot) = [Q_{z^1}(z, \cdot), ..., Q_{z^d}(z, \cdot)]^T$ of Q with respect to the components z^k of $z, k \in \{1, ..., d\}$, is a random variable in $L^p(\Omega, \mathbb{R}^d)$.

To quantify the tails of $Q(z, \cdot)$, we first define the τ -VaR $q_{\tau}(z)$, alternatively known as the τ -quantile, of significance $\tau \in (0, 1)$ as follows:

$$q_{\tau}(z) := \min\{\theta \in \mathbb{R} \mid \mathbb{E}\left[\mathbbm{1}_{Q(z,\cdot) \le \theta}\right] \ge \tau\}.$$
(5)

The τ -CVaR $c_{\tau}(z)$ is defined as the expected value of $Q(z, \cdot)$ in the tail above and including the τ -VaR $q_{\tau}(z)$:

$$c_{\tau}(z) := \mathbb{E}\left[Q(z, \cdot)|Q(z, \cdot) \ge q_{\tau}(z)\right].$$
(6)

As was described in Section 1, [35] proposed that $c_{\tau}(z)$ could be written in the form in Eq. (1) for a random variable $Q(z, \cdot)$ satisfying Assumption 1.(ii).

In this work, we extensively use the concept of parametric expectations. In particular, let us introduce the function (parametric expectation) Φ : $\mathbb{R} \times \mathbb{R}^d \to \mathbb{R}$ as:

$$\Phi(\theta; z) := \mathbb{E}\left[\phi(\theta, Q(z, \cdot))\right], \quad \theta \in \mathbb{R}, \ z \in \mathbb{R}^d, \tag{7}$$

with $\phi : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ given by:

$$\phi(\theta, Q) := \theta + \frac{(Q - \theta)^+}{1 - \tau}, \quad \theta \in \mathbb{R}, \ Q \in \mathbb{R}.$$
(8)

The introduction of the parametric expectation Φ has the advantage that the τ -VaR $q_{\tau}(z)$ and the τ -CVaR $c_{\tau}(z)$ of any significance τ can be obtained by simple post-processing of Φ as:

$$q_{\tau}(z) = \underset{\theta \in \mathbb{R}}{\arg\min \Phi(\theta; z)}, \qquad c_{\tau}(z) = \underset{\theta \in \mathbb{R}}{\min \Phi(\theta; z)} = \Phi(q_{\tau}(z); z).$$
(9)

The framework of parametric expectations allows us to write the penalised CVaR minimisation problem in Eq. (2) as a combined minimisation over θ and *z* as in Eq. (3). The problem is restated below for reference:

$$\mathcal{J}^* = \min_{\substack{z \in \mathbb{R}^d \\ \theta \in \mathbb{R}}} \left\{ \mathcal{J}(\theta, z) := \Phi(\theta; z) + \kappa \left\| z - z_{ref} \right\|_{l^2}^2 \right\}.$$
(10)

For the remainder of this work, we address the challenge of solving problem (10). The combined objective function $\mathcal{J}(\theta, z)$ has several properties that, when combined with the properties of Q in Assumption 1, have useful implications for gradient based optimisation techniques. We first discuss the differentiability of $\mathcal{J}(\theta, z)$. Theorem 2.1 below gives a result on Fréchet differentiability of the CVaR.

Theorem 2.1. Let $\mathfrak{Q}(X, Y)$ denote the space of bounded linear operators between the normed vector spaces X and Y. We define the function $\mathcal{R} : \mathbb{R} \times L^p(\Omega; \mathbb{R}) \to \mathbb{R}$ as follows:

$$\mathcal{R}(\theta, Q) := \theta + \frac{\mathbb{E}\left[(Q - \theta)^+\right]}{1 - \tau} = \mathbb{E}\left[\phi(\theta, Q)\right].$$
(11)

Then, $\mathcal{R}(\theta, Q)$ is jointly Fréchet differentiable in $\mathbb{R} \times \Gamma$, with Fréchet derivative $D\mathcal{R}(\theta, Q) \in \mathfrak{L}(\mathbb{R} \times L^p(\Omega, \mathbb{R}), \mathbb{R})$ at the point $(\theta, Q) \in \mathbb{R} \times \Gamma$ in the direction $(\delta\theta, \delta Q) \in \mathbb{R} \times L^p(\Omega, \mathbb{R})$ given by:

$$D\mathcal{R}(\theta, Q)(\delta\theta, \delta Q) = \left(1 - \frac{\mathbb{E}\left[\mathbbm{1}_{Q \ge \theta}\right]}{1 - \tau}\right)\delta\theta + \frac{\mathbb{E}\left[\mathbbm{1}_{\{Q \ge \theta\}}\delta Q\right]}{1 - \tau}.$$
(12)

Proof. The reader is referred to Appendix A for the proof. \Box

This result, combined with Assumption 1 on Q, leads immediately to the differentiability of $\mathcal{J}(\theta, z)$.

Corollary 2.1. *The objective function* $\mathcal{J}(\theta, z)$ *is jointly Fréchet differentiable in* $\mathbb{R} \times \mathbb{R}^d$ *, with Fréchet derivative* $D\mathcal{J}(\theta, z) \in \mathfrak{L}(\mathbb{R} \times \mathbb{R}^d, \mathbb{R})$ *at the point* (θ, z) *in the direction* $(\delta\theta, \delta z) \in \mathbb{R} \times \mathbb{R}^d$ *given by:*

$$D\mathcal{J}(\theta, z)(\delta\theta, \delta z) = \left(1 - \frac{\mathbb{E}\left[\mathbbm{1}_{Q \ge \theta}\right]}{1 - \tau}\right)\delta\theta + \frac{\mathbb{E}\left[\mathbbm{1}_{\{Q \ge \theta\}}Q_z^T\delta z\right]}{1 - \tau} + 2\kappa(z - z_{ref})^T\delta z.$$
(13)

A direct implication of Corollary 2.1 is that the gradient $\nabla \mathcal{J} \in \mathbb{R}^{d+1}$ and the partial derivatives $\mathcal{J}_z(\theta, z) = [\mathcal{J}_{z^1}(\theta, z), ..., \mathcal{J}_{z^d}(\theta, z)]^T$ and $\mathcal{J}_{\theta}(\theta, z)$ exist and are given by the following expressions:

$$\mathcal{J}_{\theta}(\theta, z) = 1 - \frac{\mathbb{E}\left[\mathbbm{1}_{\mathcal{Q}(z, \cdot) \ge \theta}\right]}{1 - \tau},\tag{14}$$

$$\mathcal{J}_{z}(\theta, z) = \frac{\mathbb{E}\left[\mathbbm{1}_{Q(z, \cdot) \ge \theta} Q_{z}(z, \cdot)\right]}{1 - \tau} + 2\kappa(z - z_{ref}).$$
(15)

One of the main contributions of this work is the estimation of the sensitivities in Eqs. (14) and (15) using MLMC estimators. However, as discussed in Section 1, using MLMC to directly estimate the expectations in Eqs. (14) and (15) may result in compromised or nonoptimal MLMC performance. The reader is referred to [5,24] for a detailed discussion on the topic. To ameliorate this issue, we propose the following alternative formulation of the gradients in terms of parametric expectations:

$$\mathcal{J}_{\theta}(\theta, z) = \Phi'(\theta; z), \quad \text{with } \Phi \text{ as in Eqs. (7)-(8)}, \tag{16}$$

$$\mathcal{J}_{z}(\theta, z) = \Psi'(\theta; z) + 2\kappa(z - z_{ref}), \tag{17}$$

where
$$\Psi(\theta; z) := \mathbb{E}\left[-\frac{(Q(z, \cdot) - \theta)^+ Q_z(z, \cdot)}{1 - \tau}\right] =: \left[\mathbb{E}\left[\psi(\theta, Q, Q_{z^1})\right], ..., \mathbb{E}\left[\psi(\theta, Q, Q_{z^d})\right]\right]^T.$$
 (18)

The superscript prime of the parametric expectations in Eq. (16) and Eq. (17) denotes the derivative computed with respect to θ . In addition to $\Phi(\theta; z)$, we have introduced the parametric expectation $\Psi(\theta; z) \in \mathbb{R}^d$ and the function $\psi(\theta, Q, Q_{z^k}) \in \mathbb{R}$ where z^k and Q_{z^k} denote the k^{th} components of z and Q_z respectively, $k \in \{1, ..., d\}$. The differentiability of $\Psi(\theta; z)$ in θ follows by the same arguments of Theorem 2.1 and Corollary 2.1, under Assumption 1. It was shown in [5] that since ϕ and ψ are Lipschitz continuous in their arguments, the corresponding MLMC estimators no longer suffer from the compromised performance due to discontinuities. The idea is then to build MLMC estimators $\hat{\Phi}(\cdot, z)$ for the whole functions $\theta \mapsto \Phi(\theta; z)$ and $\theta \mapsto \Psi(\theta; z)$ respectively on a suitably chosen interval $\Theta \subset \mathbb{R}$, and then approximate \mathcal{J}_{θ} and \mathcal{J}_z as $\hat{\mathcal{J}}_{\theta}(\theta, z) = \hat{\Phi}'(\theta; z)$ and $\hat{\mathcal{J}}_z(z, z_{ref})$ respectively. As a

by-product of this approach for estimating sensitivities, we construct an approximation $\theta \in \Theta \mapsto \hat{J}(\theta, z) = \hat{\Phi}(\theta; z) + \kappa \left\| z - z_{ref} \right\|_{l^2}^2$ of the objective function itself for all $\theta \in \Theta$, at a given design $z \in \mathbb{R}^d$. This allows us to consider an optimisation problem in which exact minimisation in θ is performed at each iteration using the surrogate \hat{J} , and gradient steps are performed only in z using the approximate gradient \tilde{J}_z . Notice that the gradient approximation in z is inconsistent with the surrogate model \hat{J} , i.e., $\tilde{J}_z \neq \partial_z \hat{J}$, in contrast to \hat{J}_{θ} . We will detail this approach in the next section.

3. Gradient based optimisation algorithm

In this section, we present a gradient-based iterative procedure to find a local minimiser (θ^*, z^*) of the OUU problem in Eq. (10), should it exist. The broad goal of a gradient based algorithm is to define the iterates (θ_i, z_i), $j \in \mathbb{N}$ such that

$$\lim_{i \to \infty} (\theta_j, z_j) = (\theta^*, z^*), \tag{19}$$

where the iterates are computed using gradient information. Motivated by our interest in using MLMC estimators based on parametric expectations to estimate the objective function and its sensitivities, we consider in this section the general situation in which, at each iteration *j* of the gradient based algorithm, we build an approximation $\hat{J}^{j}(\theta, z), \theta \in \Theta$ of the objective function at the design $z \in \mathbb{R}^{d}$ on a suitably chosen interval $\Theta \subset \mathbb{R}$ (which may depend on *j*, although to ease the notation, we do not highlight such dependence), as well as approximations $\hat{J}^{j}_{\theta}(\theta, z)$ and $\tilde{J}^{j}_{z}(\theta, z), \theta \in \Theta$, where the approximation \tilde{J}^{j}_{z} may not coincide with the *z*-derivative of \hat{J}^{j} . The approximations \hat{J}^{j}_{θ} , \hat{J}^{j}_{θ} and \tilde{J}^{j}_{z} may be random, as will be the case for MLMC estimators. We then propose the following variation of the standard gradient descent algorithm, starting from an initial design z_{0} :

$$\theta_j \in \operatorname*{arg\,min}_{\theta \in \Theta} \hat{J}^j(\theta, z_j),\tag{20}$$

$$z_{j+1} = z_j - \alpha \widetilde{J}_z^j(\theta_j, z_j), \tag{21}$$

where $\alpha > 0$ denotes a step size parameter. We note that according to the procedure in [5], the interval Θ can be freely selected and, hence, we can ensure that θ_j always belongs to the interior of Θ , so that $\hat{J}_{\theta}^{j}(\theta_j, z_j) = 0 \forall j \in \mathbb{N}$.

In Theorem 3.1 in Section 3.1, we show that the iterates (θ_j, z_j) converge Q-linearly in the iteration counter *j* towards (θ^*, z^*) under additional assumptions on the objective function \mathcal{J} and its approximations $\hat{\mathcal{J}}^j$. The results of Theorem 3.1, specifically the implications of Eq. (24) introduced there, demonstrate that Q-linear convergence of the iterates z_j and θ_j in *j* can be obtained if the gradient approximation is accurate up to a tolerance that is a fraction η of the gradient magnitude at the previous iteration. Such an accuracy condition was used in [7], and has been utilised in literature prior to this work. The interested reader is referred to [10] and [8].

The step size is selected sufficiently small, and remains fixed over all optimisation iterations, although variable step sizes and line search methods could be easily added. The algorithm is terminated once the gradient magnitude drops to a specified fraction of the initial value. We introduce here the notation $w = (\theta, z)$, $\nabla \mathcal{J} = (\mathcal{J}_{\theta}, \mathcal{J}_{z})$ and $\tilde{\nabla} \hat{\mathcal{J}}^{j} = (\hat{\mathcal{J}}_{\theta}^{j}, \hat{\mathcal{J}}_{z}^{j})$ for convenience in the following.

Algorithm 1: Novel AMGD algorithm

1: Input: Initial design z_0 , iteration counter j = 0, tolerance $0 < \epsilon < 1$, step size $\alpha > 0$ and tolerance fraction $\eta > 0$. 2: Set residual $r = \epsilon + 1$. 3: while $r > \epsilon$ do if j = 0 { Compute $\hat{\mathcal{J}}^0(\cdot, z_i)$ and $\tilde{\mathcal{J}}_z^0(\cdot, z_i)$ up to a fixed tolerance.} 4: else $\left\{ \text{ Compute } \hat{\mathcal{J}}^{j}(\cdot,z_{j}) \text{ and } \widetilde{\mathcal{J}}^{j}_{z}(\cdot,z_{j}) \text{ such that } \text{MSE}\left(\tilde{\nabla}\hat{\mathcal{J}}^{j}(\cdot,z_{j})\right) \leq \eta^{2} \left\| \nabla \mathcal{J}(\theta_{j-1},z_{j}) \right\|_{L^{2}}^{2} \text{ with } \text{MSE}\left(\tilde{\nabla}\hat{\mathcal{J}}^{j}(\cdot,z_{j})\right) \text{ defined as in Eq. (24).} \right\}$ 5: Compute a minimiser $\theta_j \in \arg \min_{\theta \in \Theta} \hat{J}^j(\theta, z_j)$. 6: 7: Compute gradient step $z_{i+1} = z_i - \alpha \widetilde{\mathcal{J}}_z^j(\theta_i, z_i)$. Set residual $r = \left\| \tilde{\nabla} \hat{\mathcal{J}}^{j}(w_{j}) \right\|_{l^{2}}^{2} / \left\| \tilde{\nabla} \hat{\mathcal{J}}^{0}(w_{0}) \right\|_{l^{2}}^{2}$ 8: 9: Update $j \leftarrow j + 1$. 10: end while

3.1. Convergence analysis

For the interested reader, we present a self-contained convergence analysis of the iterates (θ_j, z_j) in Theorem 3.1, under additional assumptions on \mathcal{J} and $\hat{\mathcal{J}}^j$, based on the analysis presented in [7]. The key differences in the two analyses are related to the fact that the algorithm studied here is an AMGD algorithm instead of a pure gradient descent algorithm. We first note that the objective function $\mathcal{J}(\theta, z)$ is convex under the additional assumption that $Q(z, \cdot)$ is almost surely convex in z [35, Theorem 10]. When combined with the assumption that $\mathcal{J} \to \infty$ when $||z||_{l^2}, |\theta| \to \infty$, this ensures that a minimiser of $\mathcal{J}(\theta, z)$ exists in $\mathbb{R} \times \mathbb{R}^d$. However, we require additional assumptions on the objective function \mathcal{J} to prove Q-linear convergence of the iterates θ_j and z_j towards such a minimiser; namely Assumptions 2 and 3 below on strong convexity and on the Lipschitz continuity of the gradients, respectively. An immediate implication of Assumption 2 is that there exists a unique minimiser $(\theta^*, z^*) \in \mathbb{R} \times \mathbb{R}^d$ for the OUU problem in Eq. (10) such that $\mathcal{J}_z(\theta^*, z^*) = \mathcal{J}_\theta(\theta^*, z^*) = 0$.

In what follows, we denote by $\mathbb{E}_j[\cdot]$ the expectation conditional on all of the random variables used to define z_j (i.e., conditioned on the past up to iteration j), and by $\langle \cdot, \cdot \rangle$ the l^2 inner product. Readers interested in the implementation details of Algorithm 1 and its relation to the MLMC method can proceed directly to Section 4.

Assumption 2. The objective function \mathcal{J} is μ -strongly convex, i.e., there exists $\mu > 0$ such that, for all $w_a, w_b \in \mathbb{R} \times \mathbb{R}^d$, equivalently:

(i)
$$\mathcal{J}(w_b) \ge \mathcal{J}(w_a) + \langle w_b - w_a, \nabla \mathcal{J}(w_a) \rangle + \frac{\mu}{2} \|w_b - w_a\|_{l^2}^2$$
,
(ii) $\langle \nabla \mathcal{J}(w_b) - \nabla \mathcal{J}(w_a), w_b - w_a \rangle \ge \mu \|w_b - w_a\|_{l^2}^2$.

Assumption 3. The objective function \mathcal{J} has Lipschitz continuous gradients, i.e., there exists L > 0 such that, for all $w_a, w_b \in \mathbb{R} \times \mathbb{R}^d$:

$$\|\nabla \mathcal{J}(w_b) - \nabla \mathcal{J}(w_a)\|_{l^2} \le L \|w_b - w_a\|_{l^2}.$$
(22)

Lemma 3.1. Let \mathcal{J} satisfy Assumptions 2 and Assumptions 3. Then we have that, for $0 < \alpha \le 1/L$,

$$\frac{\mu}{2} \|w - w^*\|_{l^2}^2 + \frac{\alpha}{2} \|\nabla \mathcal{J}(w)\|_{l^2}^2 \le \langle \nabla \mathcal{J}(w), w - w^* \rangle.$$
(23)

The above result is restated here from [7, Lemma 2.1].

Theorem 3.1. Let $\Theta \subset \mathbb{R}$ be a convex set. Let $\mathcal{J} : \mathbb{R} \times \mathbb{R}^d \to \mathbb{R}$ satisfy Assumptions 2 and 3, and $\hat{\mathcal{J}}^j : \Theta \times \mathbb{R}^d \to \mathbb{R}$ satisfy the following condition:

$$MSE\left(\tilde{\nabla}\hat{\mathcal{J}}^{j}(\cdot,z_{j})\right) := \mathbb{E}_{j}\left[\left\|\hat{\mathcal{J}}_{\theta}^{j}(\cdot,z_{j}) - \mathcal{J}_{\theta}(\cdot,z_{j})\right\|_{L^{\infty}(\Theta)}^{2}\right] + \sum_{k=1}^{d} \mathbb{E}_{j}\left[\left\|\tilde{\mathcal{J}}_{z,k}^{j}(\cdot,z_{j}) - \mathcal{J}_{z,k}(\cdot,z_{j})\right\|_{L^{\infty}(\Theta)}^{2}\right] \le \eta^{2} \left\|\nabla\mathcal{J}(\theta_{j-1},z_{j})\right\|_{l^{2}}^{2},$$
(24)

for some $\eta > 0$, where (θ_{j-1}, z_j) is the jth iterate produced by Algorithm 1 with step size α satisfying $0 < \alpha \le 1/L$ and $\alpha \mu \le 1$. Then, the following result holds true:

$$\mathbb{E}\left[\left\|z_{j+1} - z^*\right\|_{l^2}^2 + C_1(\theta_j - \theta^*)^2\right] \le \xi \mathbb{E}\left[\left\|z_j - z^*\right\|_{l^2}^2 + C_1(\theta_{j-1} - \theta^*)^2\right],\tag{25}$$

for some constants $C_1 > 0$ and $0 < \xi < 1$.

Proof. From the definition of the iterate z_{i+1} in Eq. (21), we have:

$$\left\|z_{j+1} - z^*\right\|_{l^2}^2 = \left\|z_j - z^* - \alpha \widetilde{\mathcal{J}}_z^j(\theta_j, z_j)\right\|_{l^2}^2$$
(26)

$$= \left\| z_{j} - z^{*} \right\|_{l^{2}}^{2} + \alpha^{2} \left\| \widetilde{J}_{z}^{j}(\theta_{j}, z_{j}) \right\|^{2} - 2\alpha \left\langle \widetilde{J}_{z}^{j}(\theta_{j}, z_{j}), z_{j} - z^{*} \right\rangle$$

$$= \left\| z_{j} - z^{*} \right\|_{l^{2}}^{2} + \alpha^{2} \left(\left\| \widetilde{J}_{z}^{j}(\theta_{j}, z_{j}) \right\|^{2} + \left(\widehat{J}_{\theta}^{j}(\theta_{j}, z_{j}) \right)^{2} \right)$$
(27)

$$\underbrace{-2\alpha\left(\left\langle \mathcal{J}_{z}(\theta_{j},z_{j}),z_{j}-z^{*}\right\rangle + \left\langle \mathcal{J}_{\theta}(\theta_{j},z_{j}),\theta_{j}-\theta^{*}\right\rangle\right)}_{=:\hat{T}_{2}} \\ \underbrace{-2\alpha\left(\left\langle \widetilde{\mathcal{J}}_{z}^{j}(\theta_{j},z_{j})-\mathcal{J}_{z}(\theta_{j},z_{j}),z_{j}-z^{*}\right\rangle + \left\langle \widehat{\mathcal{J}}_{\theta}^{j}(\theta_{j},z_{j})-\mathcal{J}_{\theta}(\theta_{j},z_{j}),\theta_{j}-\theta^{*}\right\rangle\right)}_{=:\hat{T}_{3}}.$$

$$(28)$$

The term $\hat{T}_1 = \alpha^2 \left\| \tilde{\nabla} \hat{J}^j(w_j) \right\|_{l^2}^2$ can be bounded as follows:

$$\mathbb{E}_{j}\left[\hat{T}_{1}\right] = \alpha^{2} \mathbb{E}_{j}\left[\left\|\tilde{\nabla}\hat{\mathcal{J}}^{j}(\theta_{j}, z_{j})\right\|_{l^{2}}^{2}\right]$$
(29)

$$\leq \alpha^2 \mathbb{E}_j \left[\left\| \tilde{\nabla} \hat{\mathcal{J}}^j(\theta_j, z_j) \pm \nabla \mathcal{J}(\theta_j, z_j) \right\|_{l^2}^2 \right]$$
(30)

$$\leq \alpha^{2} \left[\mathbb{E}_{j} \left[\left\| \tilde{\nabla} \hat{\mathcal{J}}^{j}(\theta_{j}, z_{j}) - \nabla \mathcal{J}(\theta_{j}, z_{j}) \right\|_{l^{2}}^{2} \right]^{1/2} + \mathbb{E}_{j} \left[\left\| \nabla \mathcal{J}(\theta_{j}, z_{j}) \right\|_{l^{2}}^{2} \right]^{1/2} \right]^{2}$$
(31)

$$\leq \alpha^{2} \left[\eta \left\| \nabla \mathcal{J}(\theta_{j-1}, z_{j}) \right\|_{l^{2}} + \mathbb{E}_{j} \left[\left\| \nabla \mathcal{J}(\theta_{j}, z_{j}) \right\|_{l^{2}}^{2} \right]^{1/2} \right]^{2}$$
(32)

$$\leq \alpha^{2} \left[(\eta^{2} + \eta) \left\| \nabla \mathcal{J}(\theta_{j-1}, z_{j}) \right\|_{l^{2}}^{2} + (1 + \eta) \mathbb{E}_{j} \left[\left\| \nabla \mathcal{J}(\theta_{j}, z_{j}) \right\|_{l^{2}}^{2} \right] \right].$$

$$(33)$$

The term $\hat{T}_2 = -2\alpha \langle \nabla \mathcal{J}(w_j), w_j - w^* \rangle$ can be bounded as follows:

Journal of Computational Physics 495 (2023) 112523

S. Ganesh and F. Nobile

$$\mathbb{E}_{j}\left[\hat{T}_{2}\right] \leq -\alpha\mu \left(\left\|z_{j}-z^{*}\right\|_{l^{2}}^{2} + \mathbb{E}_{j}\left[\left(\theta_{j}-\theta^{*}\right)^{2}\right]\right) - \alpha^{2} \mathbb{E}_{j}\left[\left\|\nabla \mathcal{J}(\theta_{j},z_{j})\right\|_{l^{2}}^{2}\right],\tag{34}$$

where we have used Lemma 3.1. Finally, the term $\hat{T}_3 = -2\alpha \left\langle \tilde{\nabla} \hat{\mathcal{J}}^j(w_j) - \nabla \mathcal{J}(w_j), w_j - w^* \right\rangle$ can be bounded as follows:

$$\mathbb{E}_{j}\left[\hat{T}_{3}\right] \leq 2\alpha \mathbb{E}_{j}\left[\left\|\tilde{\nabla}\hat{\mathcal{J}}^{j}(\theta_{j}, z_{j}) - \nabla\mathcal{J}(\theta_{j}, z_{j})\right\|_{l^{2}} \left\|w_{j} - w^{*}\right\|_{l^{2}}\right]$$

$$(35)$$

$$\leq 2\alpha \mathbb{E}_{j} \left[\left\| \tilde{\nabla} \hat{\mathcal{J}}^{j}(\theta_{j}, z_{j}) - \nabla \mathcal{J}(\theta_{j}, z_{j}) \right\|_{l^{2}}^{2} \right]^{1/2} \mathbb{E}_{j} \left[\left\| w_{j} - w^{*} \right\|_{l^{2}}^{2} \right]^{1/2}$$

$$(36)$$

$$\leq 2\alpha\eta \left\|\nabla \mathcal{J}(\theta_{j-1}, z_j)\right\|_{l^2} \mathbb{E}_j \left[\left\|w_j - w^*\right\|_{l^2}^2\right]^{1/2}.$$
(37)

Combining the bounds for \hat{T}_1 , \hat{T}_2 and \hat{T}_3 , we have the following:

$$\mathbb{E}_{j}\left[\left\|z_{j+1}-z^{*}\right\|_{l^{2}}^{2}\right] \leq (1-\alpha\mu)\left\|z_{j}-z^{*}\right\|_{l^{2}}^{2}-\alpha\mu\mathbb{E}_{j}\left[(\theta_{j}-\theta^{*})^{2}\right] + \alpha^{2}(\eta^{2}+\eta)\left\|\nabla\mathcal{J}(\theta_{j-1},z_{j})\right\|_{l^{2}}^{2}+\alpha^{2}\eta\mathbb{E}_{j}\left[\left\|\nabla\mathcal{J}(\theta_{j},z_{j})\right\|_{l^{2}}^{2}\right] + 2\alpha\eta\left\|\nabla\mathcal{J}(\theta_{j-1},z_{j})\right\|_{l^{2}}\mathbb{E}_{j}\left[\left\|w_{j}-w^{*}\right\|_{l^{2}}^{2}\right]^{1/2}.$$
(38)

We now utilise Lemma 3.1 once again, from which we have the following result:

$$\alpha \|\nabla \mathcal{J}(w)\|_{l^{2}} \le (1 + \sqrt{1 - \alpha \mu}) \|w - w^{*}\|_{l^{2}} =: \tilde{L} \|w - w^{*}\|_{l^{2}},$$
(39)

for $0 < \alpha \le 1/L$ and $\alpha \mu \le 1$. In addition, the last term of Eq. (38) can be rewritten as follows:

$$2\alpha\eta \left\|\nabla \mathcal{J}(\theta_{j-1}, z_j)\right\|_{l^2} \mathbb{E}_{j}\left[\left\|w_{j} - w^*\right\|_{l^2}^{2}\right]^{1/2} \le \eta \left(\frac{\alpha^2 \left\|\nabla \mathcal{J}(\theta_{j-1}, z_j)\right\|_{l^2}^{2}}{\tilde{L}} + \tilde{L} \mathbb{E}_{j}\left[\left\|w_{j} - w^*\right\|_{l^2}^{2}\right]\right)$$
(40)

Applying Eqs. (39) and (40) to Eq. (38), we then have the following simplified bound:

$$\mathbb{E}_{j}\left[\left\|z_{j+1}-z^{*}\right\|_{l^{2}}^{2}\right] \leq \left(1-\alpha\mu+(\eta^{2}+2\eta)\tilde{L}^{2}+2\eta\tilde{L}\right)\left\|z_{j}-z^{*}\right\|_{l^{2}}^{2} + \left(-\alpha\mu+\eta\tilde{L}^{2}+\eta\tilde{L}\right)\mathbb{E}_{j}\left[(\theta_{j}-\theta^{*})^{2}\right] + \left((\eta^{2}+\eta)\tilde{L}^{2}+\eta\tilde{L}\right)(\theta_{j-1}-\theta^{*})^{2},$$

$$= \left(1-C_{1}+C_{2}\right)\left\|z_{j}-z^{*}\right\|_{l^{2}}^{2} - C_{1}\mathbb{E}_{j}\left[(\theta_{j}-\theta^{*})^{2}\right] + C_{2}(\theta_{j-1}-\theta^{*})^{2},$$

$$(42)$$

where we have defined the constants $C_1 = \alpha \mu - \eta \tilde{L}^2 + \eta \tilde{L}$ and $C_2 = (\eta^2 + \eta)\tilde{L}^2 + \eta \tilde{L}$. We then have the following:

$$\mathbb{E}_{j}\left[\left\|z_{j+1}-z^{*}\right\|_{l^{2}}^{2}\right]+C_{1}\mathbb{E}_{j}\left[(\theta_{j}-\theta^{*})^{2}\right] \leq (1-C_{1}+C_{2})\left\|z_{j}-z^{*}\right\|_{l^{2}}^{2}+C_{2}(\theta_{j-1}-\theta^{*})^{2}$$

$$\leq \max\left(1-C_{1}+C_{2},\frac{C_{2}}{C_{1}}\right)\left(\left\|z_{j}-z^{*}\right\|_{l^{2}}^{2}+C_{1}(\theta_{j-1}-\theta^{*})^{2}\right).$$
(43)

We note that the leading constant on the right hand side is less than 1 as long as $C_1 > C_2$, which holds true for $\eta < \sqrt{1 + \alpha \mu / \tilde{L}^2} - 1$. This in turn ensures contraction in the norm $||z||_{l^2}^2 + C_1 \theta^2$ on the space $\mathbb{R}^d \times \mathbb{R}$. This completes the proof.

Remark 1. We note that although the accuracy condition Eq. (24) is stated in the L^{∞}-norm for all θ , the proof of Theorem 3.1 uses this property only at θ_j . This condition is required since we do not know the quantile θ_j a priori, and seek to use the parametric expectation framework from [5] to do so. [5] requires that the error in the approximations \hat{J}^j be controlled at all θ , in order to estimate θ_j accurately.

Remark 2. In practical applications, it is difficult to determine whether Assumptions 2 and 3 are satisfied, since both are strongly dependent on the properties of the random QoI $Q(z, \cdot)$. These assumptions require stronger properties on $Q(z, \cdot)$ and its Probability Density Function (PDF) than those presented in Assumption 1; for example, that the PDF remains both upper bounded and lower bounded away from zero for all designs z, and that the random variable $Q(z, \cdot)$ is bounded, i.e., $Q(z, \cdot) \in L^{\infty}(\Omega, \mathbb{R})$.

Remark 3. We remark that η is a monotonically increasing function of the product $\alpha \mu$ in the interval $0 < \alpha \mu \le 1$. For an appropriate step size α , chosen such that $\alpha \mu$ is close to 1, η can be as large as 0.4.

4. Gradient estimation and error control using MLMC methods

We note that the key assumption in the proof of Theorem 3.1 is Eq. (24); namely, that the gradient approximation is accurate up to a tolerance that is proportional to the magnitude of the true gradient. As stated earlier in Section 1, we are interested in utilising the framework of MLMC estimators for parametric expectations developed in [5] for the accurate estimation of the objective function \mathcal{J} (risk-measure CVaR) and its gradient.

Expressing the gradients \mathcal{J}_z and \mathcal{J}_{θ} in terms of the first derivatives of the parametric expectations $\Phi(\theta; z)$ and $\Psi(\theta; z)$ as in Eqs. (16) and (17) and estimating the latter using MLMC estimators poses many key advantages. The first advantage was already seen earlier in Section 3; namely that $\tilde{\mathcal{J}}_z^j$ and $\hat{\mathcal{J}}_{\theta}^j$ can be estimated for all θ for a given design z in one shot. Secondly, as was demonstrated in [5], the level-wise differences for the MLMC estimator of $\Phi(\theta; z)$, analogous to the level-wise differences corresponding to the classical MLMC estimator of $\mathbb{E}[Q]$, decay at the same rate in the levels l as the differences $Q_l - Q_{l-1}$, in an appropriately selected norm over $\theta \in \mathbb{R}$. This ensures that if cost-optimal MLMC behaviour can be achieved for estimating $\mathbb{E}[Q]$, then it can be achieved also for MLMC estimators of $\Phi(\theta; z)$, using a practically computable number of samples. The last key advantage is that, using the mechanism in [5], one can select the parameters of the MLMC estimator such that a prescribed tolerance can be attained on the MLMC approximation error on Φ and Ψ . By prescribing a tolerance proportional to the gradient magnitude, one can estimate the gradient using MLMC estimators that respect the condition in Eq. (24) as required by Algorithm 1.

Although the procedure used in this work to estimate Φ accurately is identical to the one described in [5], some important modifications are required to use the same procedure for accurately estimating Ψ . We present in this section the modifications of the work developed in [5] that are required for the accurate estimation of Ψ , and consequently the gradients \mathcal{J}_{θ} and \mathcal{J}_{z} , using the MLMC method.

4.1. MLMC estimator for the gradients

We begin by recalling that the parametric expectation Ψ is defined as in Eq. (18). The proposed MLMC method relies on a sequence of approximations $\{Q_l(z)\}_{l=0}^L$ to Q(z) on a sequence of L + 1 discretisations with, for example, different mesh sizes $h_0 > h_1 > ... > h_L$, typically a geometric sequence $h_{l-1} = sh_l$ with s > 1. The MLMC estimator for the k^{th} component $\Psi_k(\cdot; z) := \mathbb{E}\left[\psi(\cdot, Q(z), Q_{z^k}(z))\right]$ of Ψ on Θ , $k \in \{1, ..., d\}$ follows the same construction as that for Φ in [5]. The first step is to estimate $\Psi_k(\theta_r, z), r \in \{1, ..., n\}$, on a set of n equidistant points $\theta = \{\theta_1, ..., \theta_n\}$ such that $\Theta = [\theta_1, \theta_n]$, by a standard MLMC estimator $\Psi_{L,k}(\theta_r; z)$, which reads:

$$\hat{\Psi}_{L,k}(\theta_r; z) := \frac{1}{N_0} \sum_{i=1}^{N_0} \psi\left(\theta_r, Q_0^{(i,0)}(z), Q_{z^k,0}^{(i,0)}(z)\right) \\
+ \sum_{l=1}^{L} \frac{1}{N_l} \sum_{i=1}^{N_l} \left[\psi\left(\theta_r, Q_l^{(i,l)}(z), Q_{z^k,l}^{(i,l)}(z)\right) - \psi\left(\theta_r, Q_{l-1}^{(i,l)}(z), Q_{z^k,l-1}^{(i,l)}(z)\right)\right],$$
(44)

where $Q_l^{(i,l)}(z) \equiv Q_l(z; \omega^{(i,l)})$ and $Q_{l-1}^{(i,l)}(z) \equiv Q_{l-1}(z; \omega^{(i,l)})$ are correlated realisations of $Q_l(z)$ and $Q_{l-1}(z)$, respectively, typically obtained by solving the underlying differential problem on meshes with discretisation parameters h_l and h_{l-1} , driven by the same realisation $\omega^{(i,l)}$ of the random parameters for the fixed design z. On the other hand, $Q_l^{(i,l)}$ and $Q_k^{(j,k)}$ are independent if $i \neq j$ or $l \neq k$. Finally, $Q_{z^k,l-1}^{(i,l)}$ and $Q_{z^k,l-1}^{(i,l)}$ are the sensitivities of the realisations $Q_l^{(i,l)}$ and $Q_{l-1}^{(i,l)}$ respectively with respect to z^k . { N_l } $_{l=0}^L$ is a decreasing sequence of sample sizes. The MLMC hierarchy is hence defined by three parameters; namely the number of interpolation points n, the number of levels L and the level-wise sample sizes N_l .

We finally construct a MLMC estimator $\hat{\Psi}_{L,k}$ of the whole function $\Psi_k(\cdot; z) : \Theta \to \mathbb{R}$ by interpolating over the pointwise estimates as below:

$$\hat{\Psi}_{L,k}(\cdot;z) = \mathcal{S}_n\left(\hat{\Psi}_{L,k}(\theta;z)\right),\tag{45}$$

where S_n denotes a uniform cubic spline interpolation operator and $\hat{\Psi}_{L,k}(\theta;z)$ denotes the set of pointwise MLMC estimates in Eq. (44), that is $\hat{\Psi}_{L,k}(\theta;z) = \{\hat{\Psi}_{L,k}(\theta_1;z), \hat{\Psi}_{L,k}(\theta_2;z), \dots, \hat{\Psi}_{L,k}(\theta_n;z)\}$. An estimate of the first derivative Ψ'_k in θ is then obtained by computing the derivative of the resultant interpolated function, for each component $\hat{\Psi}'_{L,k}(z)$.

$$\hat{\Psi}_{L,k}'(\cdot;z) := S_n'\left(\hat{\Psi}_{L,k}(\theta;z)\right) := \frac{\partial}{\partial\theta} S_n\left(\hat{\Psi}_{L,k}(\theta;z)\right). \tag{46}$$

4.2. Estimation of the Mean Squared Error (MSE) of the gradient

Since we have assumed that the gradient estimate $\tilde{\nabla} \hat{J}^j$ is a random vector in $L^p(\Omega, \mathbb{R}^{d+1})$ with $p \ge 2$, we propose to quantify the error on the gradient in an MSE sense as follows:

$$\operatorname{MSE}\left(\tilde{\nabla}\hat{\mathcal{J}}^{j}(\cdot,z_{j})\right) := \mathbb{E}\left[\left\|\hat{\mathcal{J}}_{\theta}^{j}(\cdot,z_{j}) - \mathcal{J}_{\theta}(\cdot,z_{j})\right\|_{\mathrm{L}^{\infty}(\Theta)}^{2}\right] + \sum_{k=1}^{d} \mathbb{E}\left[\left\|\tilde{\mathcal{J}}_{z,k}^{j}(\cdot,z_{j}) - \mathcal{J}_{z,k}(\cdot,z_{j})\right\|_{\mathrm{L}^{\infty}(\Theta)}^{2}\right],\tag{47}$$

where $\mathcal{J}_{z,k}$ and $\widetilde{\mathcal{J}}_{z,k}^{j}$ denote the k^{th} components of \mathcal{J}_{z} and $\widetilde{\mathcal{J}}_{z}^{j}$.

We now present a result relating MSE $(\tilde{\nabla} \hat{\mathcal{J}}^{j}(\cdot, z_{j}))$ to the MSE of the MLMC estimators $\hat{\Phi}'_{I}$ and $\hat{\Psi}'_{I}$.

Proposition 4.1. Let $\hat{\Phi}_L(\cdot; z_j)$ and $\hat{\Psi}_L(\cdot; z_j)$ denote the MLMC estimators of $\Phi(\cdot; z_j)$ and $\Psi(\cdot; z_j)$ as defined in [5] and Eq. (44) respectively. Let $\tilde{\nabla} \hat{J}^j(\cdot, z_j)$ be the approximation to the true gradient $\nabla \mathcal{J}(\cdot, z_j)$ computed using the estimates $\hat{\Phi}'_L(\cdot; z_j)$ and $\hat{\Psi}'_L(\cdot; z_j)$ at the j^{th} optimisation iteration. Let Ψ_k and $\hat{\Psi}_{L,k}$ denote the k^{th} component of Ψ and $\hat{\Psi}_L$ respectively, for $k \in \{1, ..., d\}$. Let the MSEs on $\hat{\Phi}'_L$ and $\hat{\Psi}'_{L,k}$ be defined as follows:

$$MSE\left(\hat{\Phi}'_{L}\right)(z_{j}) := \mathbb{E}\left[\left\|\hat{\Phi}'_{L}(\cdot;z_{j}) - \Phi'(\cdot;z_{j})\right\|^{2}_{L^{\infty}(\Theta)}\right],\tag{48}$$

$$MSE\left(\hat{\Psi}'_{L,k}\right)(z_j) := \mathbb{E}\left[\left\|\hat{\Psi}'_{L,k}(\cdot;z_j) - \Psi'_k(\cdot;z_j)\right\|^2_{L^{\infty}(\Theta)}\right],\tag{49}$$

for the design $z_i \in \mathbb{R}^d$. Then, we have that:

$$MSE\left(\tilde{\nabla}\hat{\mathcal{J}}^{j}(\cdot,z_{j})\right) = MSE\left(\hat{\Phi}_{L}^{\prime}\right)(z_{j}) + \sum_{k=1}^{d} MSE\left(\hat{\Psi}_{L,k}^{\prime}\right)(z_{j}).$$

$$(50)$$

Proof. We first note that:

$$\left\|\hat{\mathcal{J}}_{\theta}^{j}(\cdot,z_{j}) - \mathcal{J}_{\theta}(\cdot,z_{j})\right\|_{\mathcal{L}^{\infty}(\Theta)}^{2} = \left\|\hat{\Phi}^{\prime}(\cdot;z_{j}) - \Phi^{\prime}(\cdot;z_{j})\right\|_{\mathcal{L}^{\infty}(\Theta)}^{2}$$
(51)

$$\left\|\widetilde{\mathcal{J}}_{z,k}^{j}(\cdot, z_{j}) - \mathcal{J}_{z,k}(\cdot, z_{j})\right\|_{L^{\infty}(\Theta)}^{2} = \left\|\widehat{\Psi}_{L,k}^{\prime}(\cdot; z_{j}) - \Psi_{k}^{\prime}(\cdot; z_{j})\right\|_{L^{\infty}(\Theta)}^{2}$$
(52)

Adding together each of the contributions and taking the expectation on both sides, we have that:

$$MSE\left(\tilde{\nabla}\hat{\mathcal{J}}^{j}(\cdot,z_{j})\right) = MSE\left(\hat{\Phi}_{L}^{\prime}\right)(z_{j}) + \sum_{k=1}^{d} MSE\left(\hat{\Psi}_{L,k}^{\prime}\right)(z_{j}). \quad \Box$$
(53)

As was described earlier in this section, we seek to use the error estimation and adaptivity procedure described in [5] to accurately estimate Φ' and Ψ'_{k} , and consequently, to accurately estimate the gradient $\nabla \mathcal{J}$. From Eq. (50), it is evident that if one can control the MSE of $\hat{\Phi}'_{L}$ and $\hat{\Psi}'_{L,k}$ in an L^{∞} sense, one can control the MSE on the gradient $\nabla \mathcal{J}$ as defined in Eq. (47). Specifically, the MSE of the gradient is equal to a simple sum of the MSEs of the parametric expectations. Eq. (50) hence allows us to use the work of [5] to accurately calibrate MLMC estimators for the parametric expectations Φ' and Ψ' such that the resultant gradient estimate is accurate up to a prescribed tolerance.

4.3. Modified error estimation procedure

Since the error estimation procedure is independent of the design *z*, in the following, we drop the explicit dependence of Φ and Ψ on *z*, with the dependence being implied. We recall here that the error estimation procedure for estimating MSE $(\hat{\Phi}'_L)$ is identical to that presented in [5]. The procedure for estimating MSE $(\hat{\Psi}'_L)$ however has several modifications from the procedure for $\hat{\Phi}'_L$, that we detail in this section. We recall that MSE $(\hat{\Psi}'_L)$ was defined in Eq. (49). Proceeding similarly as in [5], we can bound MSE $(\hat{\Psi}'_L)$ as follows:

$$MSE\left(\hat{\Psi}'_{L,k}\right) \le (\hat{e}_{i}^{\Psi_{k}})^{2} + (\hat{e}_{b}^{\Psi_{k}})^{2} + (\hat{e}_{s}^{\Psi_{k}})^{2},$$
(54)

where $\hat{e}_i^{\Psi_k}$, $\hat{e}_b^{\Psi_k}$ and $\hat{e}_s^{\Psi_k}$ denote error estimators that estimate the error due to interpolation, the error due to approximation of the QoI (i.e. bias error), and the error due to finite sampling (i.e. statistical error) respectively on $\hat{\Psi}_{L,k}$. The reader is referred to [5] for a detailed discussion of each of the three errors components, as well as their corresponding estimators.

The procedure for estimating the interpolation and bias errors requires the accurate estimation of θ -derivatives of the function $\Psi_{l,k}(\theta) = \mathbb{E}\left[\psi(\theta, Q_l, Q_{z^k, l})\right]$. Although the true function $\Psi_{l,k}$ is smooth, replacing the true probability density with an empirical probability density corresponding to a Monte Carlo estimator implies that the right-hand side would be a linear combination of piecewise linear functions. The first derivative of such a function would be piecewise constant, and high order derivatives would not exist in the discontinuity points, and would be zero otherwise. A MLMC hierarchy designed based on estimates obtained in this manner would lead to non-optimal complexity behaviour. In [5, Section 3.2], a Kernel Density Estimation (KDE) based procedure was described for ameliorating this issue. Although the error estimation procedure is broadly the same for estimating Ψ as for Φ , an important distinction arises with respect to this KDE procedure, which we detail in this section.

Since the issue chiefly relates to the regularity of the empirical Monte Carlo probability density, we propose the use a KDE based smoothing technique; namely, we replace the true joint density p_l of $(Q_l, Q_{z^k,l})$ with a KDE smoothed joint probability density p_l^{kde} , which consists of a linear combination of two-dimensional kernels composed of products of two one-dimensional Gaussian kernels centred on each of the N_l fine samples $\{(Q_l^{(l,l)}, Q_{z^k,l})\}_{l=1}^{N_l}$:

S. Ganesh and F. Nobile

Journal of Computational Physics 495 (2023) 112523

$$\Psi_{l,k}(\theta) = \iint_{\mathcal{T}} \psi(\theta, q, q_{z^k}) p_l(q, q_{z^k}) dq dq_{z^k}$$
(55)

$$\approx \iint \psi(\theta, q, q_{z^k}) p_l^{kde}(q, q_{z^k}) dq dq_{z^k}$$
(56)

$$:= \frac{1}{N_l} \int_{i=1}^{N_l} \iint \psi(\theta, q, q_{z^k}) K_{\delta_l}(q, Q_l^{(i,l)}) K_{\delta_{z^k,l}}(q_{z^k}, Q_{z^{k,l}}^{(i,l)}) dq dq_{z^k}$$
(57)

$$= -\frac{1}{N_l} \sum_{i=1}^{N_l} \int q_{z^k} K_{\delta_{z^k,l}}(q_{z^k}, Q_{z^{k,l}}^{(i,l)}) dq_{z^k} \int \frac{(q-\theta)^+}{1-\tau} K_{\delta_l}(q, Q_l^{(i,l)}) dq$$
(58)

$$= -\frac{1}{N_l} \sum_{i=1}^{N_l} Q_{z^k,l}^{(i,l)} \int \frac{(q-\theta)^+}{1-\tau} K_{\delta_l}(q, Q_l^{(i,l)}) dq =: \mathbb{E}_{l,k}^{kde} \left[\psi(\theta, \cdot, \cdot) \right].$$
(59)

Here, $K_{\delta_l}(\cdot,\mu)$ denotes a Gaussian kernel with mean μ and bandwidth parameter $\delta_l > 0$, which is selected according to Scott's rule [37] for the realisations $\{Q_l^{(i,l)}\}_{i=1}^{N_l}$ and controls the "width" of the kernel. A closed form expression can be computed for the integral in Eq. (59), leading to the KDE smoothened approximation $\mathbb{E}_{l,k}^{kde} [\psi(\theta,\cdot,\cdot)]$ for Ψ_k .

According to the procedure in [5], the interpolation error requires the estimation of the quantity $\|\Psi_k^{(4)}\|$, for which we use the KDE estimator described above. To this end, we first select a level $\lceil L/2 \rceil$ from the MLMC hierarchy; this choice of level is to ensure that $\Psi_{\lceil L/2 \rceil,k}$ is sufficiently close to Ψ_k , and $N_{\lceil L/2 \rceil}$ is large enough for the KDE procedure to produce accurate estimates. We then construct the KDE approximation $\Upsilon_{\lceil L/2 \rceil,k}(\theta) := \mathbb{E}_{\lfloor L/2 \rceil,k}^{kde} [\psi(\theta,\cdot,\cdot)]$. The fourth derivative $\Upsilon_k^{(4)}$ is then constructed using a second order central finite difference scheme on a uniform grid on Θ with $n' \gg n$ points. The norm is evaluated on the same grid as follows:

$$\left\|\Psi_{k}^{(4)}\right\|_{L^{\infty}(\Theta)} \approx \max_{i \in \{1,\dots,n'\}} \left|\Upsilon_{\lfloor L/2 \rfloor,k}^{(4)}(\theta_{i})\right|.$$
(60)

For the bias error on $\hat{\Psi}_{L,k}$, we are required to estimate the quantity

$$\left\|S_{n}'\left(\mathbb{E}\left[\psi(\theta,Q_{l},Q_{z^{k},l})-\psi(\theta,Q_{l-1},Q_{z^{k},l-1})\right]\right)\right\|_{L^{\infty}(\Theta)}.$$
(61)

Replacing the expectation by a Monte Carlo estimator leads to the same regularity issue as described earlier in this section. To smooth the empirical Monte Carlo density, we propose the use of a KDE smoothed approximation $p_{l,l-1}^{kde}$ to the true density $p_{l,l-1}$ of $(Q_l, Q_{z^k,l}, Q_{l-1}, Q_{z^k,l-1})$, consisting of products of four one-dimensional Gaussian kernels:

$$\mathbb{E}\left[\psi(\theta, Q_l, Q_{z^k, l}) - \psi(\theta, Q_{l-1}, Q_{z^k, l-1})\right]$$
(62)

$$\approx \frac{1}{N_l} \sum_{i=1}^{N_l} \iiint \left[\psi(\theta, q^f, q^f_{z^k}) - \psi(\theta, q^c, q^c_{z^k}) \right]$$

$$K \leftarrow f = 2^{(i,l)} K \leftarrow f = 2^{(i,l)} K$$

$$= \frac{1}{N_l} \sum_{i=1}^{N_l} Q_{z^k,l-1}^{(i,l)} \int \frac{(q^c - \theta)^+}{1 - \tau} K_{\delta_{l-1}}(q^c, Q_{l-1}^{(i,l)}) dq^c - Q_{z^k,l}^{(i,l)} \int \frac{(q^f - \theta)^+}{1 - \tau} K_{\delta_l}(q^f, Q_l^{(i,l)}) dq^f$$
(65)

$$=: \mathbb{E}_{l,l-1,k}^{kde} \left[\psi(\theta, Q_l, Q_{z^k,l}) - \psi(\theta, Q_{l-1}, Q_{z^k,l-1}) \right].$$
(66)

The expectation in Eq. (61) can be replaced by the KDE smoothened expectation in Eq. (66), which can then be used in the bias error estimation procedure outlined in [5]. Lastly, the procedure for the statistical error follows the idea of bootstrapping developed in [5] identically without modification.

4.4. Adaptive hierarchy selection procedure and CMLMC-gradient descent algorithm

We discuss in this section how to select the parameters of the MLMC hierarchy; namely the number of interpolation points *n*, the level-wise sample sizes N_l and the number of levels *L*. The aim is to select these parameters such that a prescribed tolerance can be obtained on the gradient estimate $\tilde{\nabla} \hat{J}^j$. In what follows, we drop the dependence on *z* for notational simplicity, with the dependence being implied. We propose here a minor variation of the framework presented in [5, Section 5]. An adaptive strategy was proposed therein for the selection of the hierarchy parameters *n*, *L* and N_l for any statistic s_τ , the MSE of whose estimator \hat{s}_τ could be bounded by a linear combination of MSEs on $\hat{\Phi}_L$ and its derivatives:

$$MSE(\hat{s}_{\tau}) \le c_0 MSE(\hat{\Phi}_L) + c_1 MSE(\hat{\Phi}'_L) + c_2 MSE(\hat{\Phi}''_L), \quad c_0, c_1, c_2 > 0.$$
(67)

We first note that the same hierarchy adaptivity procedure extends trivially to any linear combination of MSEs of $\hat{\Phi}_L$, $\hat{\Psi}_{L,k}$, and their derivatives. Specifically, this includes the case of the MSE on the gradient $\tilde{\nabla} \hat{J}^j$ in Eq. (50). In addition, each of the MSEs on the

parametric expectations in Eq. (50) can be split into its three error contributions, similar to Eq. (54), leading to the following error estimator for MSE $(\bar{\nabla} \hat{J}^{j}(w))$:

$$MSE\left(\tilde{\nabla}\hat{\mathcal{J}}^{j}\right) = MSE\left(\hat{\Phi}_{L}'\right) + \sum_{k=1}^{d} MSE\left(\hat{\Psi}_{L,k}'\right)$$

$$\leq \underbrace{\left((\hat{e}_{i}^{\Phi})^{2} + \sum_{k=1}^{d}(\hat{e}_{i}^{\Psi_{k}})^{2}\right)}_{Squared interpolation \ error} + \underbrace{\left((\hat{e}_{b}^{\Phi})^{2} + \sum_{k=1}^{d}(\hat{e}_{b}^{\Psi_{k}})^{2}\right)}_{Squared \ bias \ error} + \underbrace{\left((\hat{e}_{s}^{\Phi})^{2} + \sum_{k=1}^{d}(\hat{e}_{s}^{\Psi_{k}})^{2}\right)}_{Squared \ statistical \ error}$$
(68)

Here, $\hat{\mathbf{e}}_{i}^{\Phi}$, $\hat{\mathbf{e}}_{b}^{\Phi}$ and $\hat{\mathbf{e}}_{s}^{\Phi}$ denote the interpolation, bias and statistical error estimators corresponding to MSE ($\hat{\Phi}_{L}^{\prime}$). Once in the above form, the procedure described in [5] for adapting the hierarchy parameters *n*, *L* and *N*_l for linear combinations of MSEs can be extended trivially to the current case when combined with the modifications proposed in Section 4.3.

Lastly, we comment that the above adaptive procedure is carried out within the framework of the CMLMC algorithm presented in [5]. The CMLMC algorithm works by first simulating a small "screening" hierarchy with relatively few samples and levels. The algorithm then adapts the hierarchy parameters with respect to a decreasing set of tolerances, of which the target tolerance is the final one. The optimal parameters for a given tolerance in the sequence are computed based on estimates obtained from the optimal hierarchy for the previous tolerance, or the initial "screening" hierarchy. In this way, the MLMC estimator becomes robust to large variations in the estimates produced by an initial screening hierarchy.

We now possess all the ingredients required to tailor Algorithm 1 to the specific case in which an MLMC procedure is combined with a CMLMC algorithm to estimate the gradient up to a prescribed tolerance. The algorithm is detailed below, and differs from Algorithm 1 in that the first estimate of the gradient is computed based on a screening hierarchy, and that successive gradients are computed such that the MSE on the gradient satisfies a tolerance equal to a fraction of the gradient magnitude from the previous iteration; namely, the right-hand side of Eq. (24) is estimated using $\hat{J}_w^{j-1}(w_{j-1})$. Another key difference to note is that in contrast to CMLMC algorithm described in [5], the screening hierarchy used to compute first estimates for the design z_j is the optimal hierarchy used to accurately estimate the gradient for the design z_{j-1} . In addition, the gradient at the first design point z_0 is estimated using an initial small fixed hierarchy.

Algorithm 2: CMLMC-gradient descent OUU algorithm

Input: Initial design z_0 , iterate j = 0, tolerance 0 < e < 1, step size a > 0 and $\eta > 0$. Set residual r = e + 1while r > e do if j = 0 { Simulate screening hierarchy } else $\left\{ \text{ Start CMLMC from the optimal hierarchy for } z_{j-1}; \text{ Simulate CMLMC adapting hierarchy such that MSE} (<math>\hat{\nabla} \hat{\mathcal{J}}^j(\cdot, z_j)$) $\leq \eta^2 \left\| \hat{\mathcal{J}}_w^{j-1}(w_{j-1}) \right\|_{l^2}^2 \right\}$ Compute minimiser $\theta_j \in \arg \min_{\theta \in \Theta} \hat{\mathcal{J}}^j(\theta, z_j) = \hat{\Phi}_L(\theta, z_j)$ Compute gradient $\tilde{\mathcal{J}}_2^j(\theta_j, z_j) = \hat{\Psi}'_L(\theta_j; z_j) + 2\kappa(z_j - z_{ref})$ Compute gradient step $z_{j+1} = z_j - \alpha \tilde{\mathcal{J}}_2^j(\theta_j, z_j)$ and $\hat{\nabla} \hat{\mathcal{J}}^j(w_j) = (\hat{\mathcal{J}}_{\theta}^j(\theta_j, z_j) = 0, \tilde{\mathcal{J}}_2^j(\theta_j, z_j))$ Set residual $r = \|\tilde{\nabla} \hat{\mathcal{J}}^j(w_j)\|_{l^2}^2 / \|\tilde{\nabla} \hat{\mathcal{J}}^0(w_0)\|_{l^2}^2$ Update $j \leftarrow j + 1$ end while

5. Numerical results

5.1. FitzHugh Nagumo oscillator

To demonstrate the optimisation framework, we use the FitzHugh–Nagumo system described in [14] and [30]. The FitzHugh–Nagumo model is a two dimensional simplification of the Hodgkin-Huxley model introduced by [18], which was originally proposed in the field of neuroscience to model the phenomenon of spiking neurons. The dynamical equations read as follows:

$$\begin{bmatrix} \dot{v} \\ \dot{w} \end{bmatrix} = \begin{bmatrix} v - \frac{v^3}{3} - w + I \\ \zeta \left(v + a - bw \right) \end{bmatrix}, \quad \begin{bmatrix} v(t=0) \\ w(t=0) \end{bmatrix} = \begin{bmatrix} v^0 \\ w^0 \end{bmatrix}, \quad t \in [0,T],$$
(69)

where $[v(t), w(t)]^T \in \mathbb{R}^2$ denotes the state variables and *a*, *b*, ζ and *I* denote system parameters. Fig. 1 shows a phase-space plot containing the *v* and *w*-nullclines for a nominal value of the system parameters. The oscillator enters a limit cycle for parameter values such that the intersection of the two nullclines lies in the interval $v \in [-1, 1]$, indicated by the black lines. If the intersection lies exterior to this interval, then the oscillator eventually reaches the intersection and remains at a constant value of *v* and *w*. Although initially proposed to model neuron behaviour, the FitzHugh–Nagumo model has seen widespread use in modelling wave phenomena in excitable media. Examples include blood coagulation [13,27] and cardio-electrophysiological phenomena [9], wherein the optimal control of the model plays an important role in the application. The reader is referred to [39] for an overview of existing work on the modelling applications and optimal control of the FitzHugh–Nagumo system.



Fig. 1. FitzHugh-Nagumo oscillator dynamics. (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)

In this work, we study the forced FitzHugh-Nagumo system:

$$\begin{bmatrix} \dot{v} \\ \dot{w} \end{bmatrix} = \begin{bmatrix} v - \frac{v^3}{3} - w + I + \sigma \dot{W}_1 \\ \zeta (v + a - bw) + \sigma \dot{W}_2 \end{bmatrix}, \quad \begin{bmatrix} v(t=0) \\ w(t=0) \end{bmatrix} = \begin{bmatrix} v^0 \\ w^0 \end{bmatrix}, \quad t \in [0,T],$$
(70)

where W_1 and W_2 are "formal" derivatives of standard Brownian paths and $\sigma = 0.01$ controls the noise strength. To study the behaviour of the system, we propose the following QoI:

$$Q = \frac{1}{T} \int_{0}^{T} v^{2}(t) dt.$$
(71)

We are interested in minimising an objective function of the form in Eq. (10), where we seek to minimise the CVaR with significance $\tau = 0.7$. We denote by $z = [a, b, \zeta, I]^T$ the vector of design parameters with respect to which we want to carry out the optimisation, and seek to penalise deviations from the design $z_{ref} = [0.8, 0.7, 0.08, 1.0]$.

We discretise the interval [0, T] using a hierarchy of uniform grids $t_j = j\Delta t_l, j \in \{0, 1, ..., N_{T,l}\}$, with $\Delta t_l = T/N_{T,l}$ and $N_{T,l} = N_{T,0}2^l$. We set T = 10 and $N_{T,0} = 20$, and consider an Euler-Maruyama discretisation of Eq. (70). Using the notation v_n^l to denote the approximation of $v(t_n)$ at level l, the discretised system then reads:

$$\begin{bmatrix} v_{n+1}^l \\ w_{n+1}^l \end{bmatrix} = \begin{bmatrix} v_n^l \\ w_n^l \end{bmatrix} + \Delta t_l \begin{bmatrix} v_n^l - \frac{(v_n^l)^3}{3} - w_n^l + I \\ \zeta \left(v_n^l + a - b w_n^l \right) \end{bmatrix} + \sigma \sqrt{\Delta t_l} \begin{bmatrix} \xi_{1,n}^l \\ \xi_{2,n}^l \end{bmatrix},$$
(72)

$$\begin{bmatrix} v_0 \\ w_0^{\prime} \end{bmatrix} = \begin{bmatrix} v^0 \\ w^0 \end{bmatrix}, \quad n \in \{0, ..., N_{T,l} - 1\},$$
(73)

where $\xi_{1,n}^l$ and $\xi_{2,n}^l$ are independently drawn realisations of standard normal random variables. The quantity of interest that we study is the following time average:

$$Q = \frac{1}{T} \int_{0}^{T} v^{2}(t) dt \approx \sum_{n=0}^{N_{T,l}-1} \left(\frac{(v_{n}^{l})^{2} + (v_{n+1}^{l})^{2}}{2} \right) \frac{\Delta t_{l}}{T} =: Q_{l}.$$
(74)

To compute the sensitivities $Q_{z,l}$, we utilize the method of adjoints. We consider the corresponding adjoint variables λ_n^l and v_n^l corresponding to v_n^l and w_n^l , $n \in \{1, ..., N_{T,l}\}$ respectively. The adjoint equation reads as follows:

$$\begin{bmatrix} \lambda_n^l \\ v_n^l \end{bmatrix} = \begin{bmatrix} \lambda_{n+1}^l \\ v_{n+1}^l \end{bmatrix} + \Delta t_l \left(\begin{bmatrix} (1 - (v_n^l)^2) & \zeta \\ -1 & -\zeta b \end{bmatrix} \begin{bmatrix} \lambda_{n+1}^l \\ v_{n+1}^l \end{bmatrix} + \begin{bmatrix} \frac{2v_n^l}{T} \\ 0 \end{bmatrix} \right),$$
(75)

$$\begin{bmatrix} \lambda_{N_{T,l}}^{l} \\ v_{N_{T,l}}^{l} \end{bmatrix} = \Delta t_{l} \begin{bmatrix} \frac{v_{n}^{l}}{T} \\ 0 \end{bmatrix}, \quad n \in \{1, \dots, N_{T,l} - 1\}.$$
(76)

The reader is referred to Appendix B for the details of the derivation.

Once the adjoint equation is solved backwards in time, the approximation $Q_{z,l}$ of the sensitivities Q_z at level l can then be obtained as follows:

$$Q_{a,l} = \sum_{n=0}^{N_{T,l}-1} \Delta t_l \zeta v_{n+1}^l, \qquad Q_{b,l} = -\sum_{n=0}^{N_{T,l}-1} \Delta t_l \zeta w_n^l v_{n+1}^l,$$

$$Q_{I,l} = \sum_{n=0}^{N_{T,l}-1} \Delta t_l \lambda_{n+1}^l, \qquad Q_{\zeta,l} = \sum_{n=0}^{N_{T,l}-1} \Delta t_l (v_n^l + a - bw_n^l) v_{n+1}^l.$$
(77)



Fig. 2. Error estimator performance for the CMLMC estimator of the gradient for the FitzHugh-Nagumo system.

To demonstrate the performance of Algorithm 2, we assess the performance individually of its two components; firstly, the performance of the CMLMC algorithm, the error estimation procedure and the adaptive strategy described in Section 4 for accurately estimating the gradient for a given design, and secondly, the gradient based optimisation procedure described in Algorithm 2. We first assess the performance of the CMLMC algorithm and adaptive strategy. We remark that the solution of the forward and adjoint problems, as well as the CMLMC procedure, are implemented within the XMC software library [1], which we use for the simulations presented herein.

We seek to accurately estimate the gradient $\nabla \mathcal{J}(\cdot, z_0)$ using the estimator $\hat{\mathcal{J}}_w(\cdot, z_0)$, where $z_0 = [0.7, 0.8, 0.08, 1.0]$ and we set $\tau = 0.70$ for the significance of the CVaR. The gradient and gradient error are estimated using the MLMC procedure described in Sections 4. To assess the reliability of the error bound derived in Proposition 4.1, we run a reliability study wherein we adapt the parameters of the MLMC hierarchy to attain a prescribed tolerance on MSE ($\hat{\mathcal{J}}_w(\cdot, z_0)$). We run the MLMC algorithm 20 times for each tolerance tested and compare the estimated error to the true error obtained using a reference gradient computed using a Monte Carlo estimator with 2×10^5 samples and 2×10^4 time steps. Specifically, we are interested in assessing the tightness of the inequality in Eq. (68).

The resultant plot is shown in Fig. 2a. Three errors are plotted in Fig. 2a; namely, the true error on the gradient, defined in the L^{∞} sense, corresponding to the term on the leftmost side of Eq. (68), the square root of the MSE estimate on the gradient, produced by the optimally calibrated MLMC hierarchy, corresponding to the term on the rightmost side of Eq. (68), and the true error on the gradient evaluated at the point (θ_0, z_0), where θ_0 corresponds to the 70%-VaR for the design z_0 . The true errors are computed with respect to a reference solution computed using 2×10^5 samples and 2×10^4 time steps. As can be seen from the figure, the MSE estimator provides a tight bound on the true error on the parametric expectations. However, the true error on the gradient in the L^{∞} sense is much larger than the true pointwise error. This is a natural consequence of using the L^{∞} -norm over the entire interval Θ to define the MSE, as compared to using the pointwise error. Controlling the MSE error in an L^{∞} sense, as defined in Eq. (47), is necessitated by the error accuracy condition in Eq. (24), in order to ensure Q-linear convergence of Algorithm 2.

Fig. 2b shows the complexity behaviour of the MLMC estimator calibrated using the CMLMC algorithm. We compute the cost required to obtain the final optimal hierarchy for a given tolerance ϵ^2 on MSE $(\hat{J}_w(\cdot, z_0))$. As can be seen from the figure, the cost grows as ϵ^{-2} , which is the theoretically predicated best case performance for the MLMC estimator. For comparison, we also plot the estimated cost of a comparable Monte Carlo estimator, as well as the expected cost growth rate for the case of the first order time discretisation used here. The Monte Carlo reference cost is computed as described in [5].

We now examine the performance of the gradient descent algorithm proposed in Section 3. We are interested in solving the minimisation problem given in Eq. (10), with $\tau = 0.7$. We utilize the framework of Algorithm 2, with a tolerance $\epsilon = 0.01$ on the gradient ratio. This implies that we stop the algorithm once the gradient magnitude has dropped to $1/100^{\text{th}}$ of its initial magnitude. As an initial guess, we begin with the design $z_0 = [0.7, 0.8, 0.08, 1.0]$. We also set $z_{ref} = [0.7, 0.8, 0.08, 1.0]$. We combine the above with the CMLMC algorithm detailed in [5] and detailed further in Section 4, with $\eta = 0.2$ on the relative error on the gradient.

We plot in Fig. 3a the value of the objective function for different iterations of the objective function. We observe Q-linear convergence in the number of iterations towards the final value, as predicted by Theorem 3.1, although we cannot guarantee that the hypotheses of Theorem 3.1 are satisfied for this problem. Fig. 3b shows the value of the gradient ratio r for different iterations of the optimisation algorithm. We also observe that the gradient decreases Q-linearly. Lastly, we plot in Fig. 3c the Cumulative Distribution Function (CDF) of the output QoI $Q(z_j, \cdot)$ computed at different iterations of the optimisation algorithm, as well as the predicted VaR and CVaR values. We observe that the CDF, the VaR and the CVaR all move left, reducing the mass in the right tail of the distribution. Since we are minimising the CVaR, defined as the expectation of the random variable above the VaR, this translates to moving the right tail of the distribution as much as possible to the left.

Fig. 4a shows the optimal hierarchy produced by the CMLMC algorithm at each iteration of the optimisation. We observe that since the tolerance supplied to the CMLMC algorithm is a fraction of the gradient magnitude, the optimally tuned hierarchy becomes larger for later iterations of the optimisation. In addition, Fig. 4b shows the cumulative cost required for the optimisation algorithm to reach a given gradient magnitude. The cumulative cost at a given optimisation iteration is defined as the sum of costs of all optimal hierarchies until the current optimisation iteration. Specifically, the cumulative cost is computed as $\sum_{i=0}^{i} \sum_{l=0}^{L} N_{l}^{(i)} (\text{Cost}(Q_{l}) + \text{Cost}(Q_{l-1}))$, where $\{N_{l}^{(i)}\}_{l=0}^{L}$ denote the optimal level-wise sample sizes for the *i*th optimisation iteration and Cost (Q_{l}) denotes the



Fig. 3. Performance of Algorithm 2 over different iterations for the FitzHugh–Nagumo system.



Fig. 4. Hierarchy of CMLMC estimators and complexity behaviour of Algorithm 2 for different iterations for the FitzHugh–Nagumo system.

average cost of simulating one sample of Q_l . This cost is plotted versus the gradient magnitude. We observe that after an initial pre-asymptotic regime, the cumulative cost grows as $\|\tilde{\nabla}\hat{J}^j(w_j)\|_{l^2}^{-2}$, a rate commensurate with the use of an optimally tuned MLMC hierarchy at each iteration tuned to obtain a tolerance proportional to $\|\tilde{\nabla}\hat{J}^j(w_j)\|_{l^2}$.

5.2. Pollutant transport problem

э

We now apply the methodology to a more applied problem of practical relevance. A problem of pollutant transport is studied, where the concentration of pollutant in a domain is modelled using a steady reaction-diffusion-advection equation. We consider a square domain $D = (0, 1) \times (0, 1)$, with boundary $\partial D := \Gamma_d \cup \Gamma_n$, where $\Gamma_d := \{0\} \times (0, 1)$ and $\Gamma_n := \partial D \setminus \Gamma_d$. We denote by $u : D \times \mathbb{R}^9 \times \Omega \to \mathbb{R}$ the concentration of the pollutant. The concentration satisfies the following equation:

$$-\nabla \cdot (\varepsilon \nabla u(x, z, \omega)) + \mathbb{V}(x, \omega) \cdot \nabla u(x, z, \omega) = f(x) - B(x, z), \quad x \in D,$$
(78)

subject to the following boundary conditions:

$$e \frac{\partial u}{\partial n}(x, z, \omega) = 0, \quad x \in \Gamma_n, \quad \text{for } \mathbb{P} - \text{a.e. } \omega \in \Omega$$
(79)

$$u(x, z, \omega) = 0, \quad x \in \Gamma_d, \quad \text{for } \mathbb{P} - \text{a.e. } \omega \in \Omega, \tag{80}$$

where $\epsilon > 0$ denotes a viscosity parameter. $\mathbb{V}(x, \omega)$ is a random divergence-free velocity field defined as follows:

$$\mathbb{V}(x,\omega) := \begin{bmatrix} b(\omega) - a(\omega)x_1\\ a(\omega)x_2 \end{bmatrix},\tag{81}$$

where $a \sim \mathcal{U}[4.95, 5.05]$ and $b \sim \mathcal{U}[3.95, 3.05]$ are uniformly distributed random variables, and x_1 and x_2 denote the components of x. The source f(x) is the sum of five Gaussian source terms:

$$f(x) = \sum_{i=1}^{5} s_i \exp\left(-\frac{(x-\mu_i)^T (x-\mu_i)}{2\sigma_i^2}\right),$$
(82)

where the values of s_i , μ_i and σ_i are given in Table 1. The sink term B(x, z) is defined as follows:

Table 1

Source term parameters for the ponutant transport proble
--

i	μ_i	σ_i	s _i
1	$[0.55205319, 0.65571641]^T$	0.0229487	2.3220339
2	$[0.49379544, 0.10950509]^T$	0.0205321	1.7931427
3	$[0.13032797, 0.57569277]^T$	0.0196891	2.3522452
4	$[0.33868732, 0.37971428]^T$	0.0212297	2.2850373
5	$[0.27670822, 0.15833522]^T$	0.0227373	2.3194400



Fig. 5. Error estimator performance for the pollutant transport problem.

$$B(x,z) = \sum_{k=1}^{9} z^k \exp\left(-\frac{(x-p_k)^T (x-p_k)}{2\sigma^2}\right),$$
(83)

where the locations p_k are defined as $p_k = (0.25i, 0.25j), i, j \in \{1, 2, 3\}, k = 3(i - 1) + j, \sigma = 0.05$, and z^k denotes the kth component of $z \in \mathbb{R}^9$. We are interested in studying the distribution of the random QoI Q, defined as follows:

$$Q(z,\omega) := \frac{\kappa_s}{2} \int_D u^2(x, z, \omega) dx,$$
(84)

with $\kappa_s = 10^4$.

The problem is implemented using the FEniCS finite element software [28]. The domain is discretised using a uniform triangular mesh with piecewise linear finite elements. The resultant linear system is solved using a sparse direct solver [3,2]. The number of elements per side of the square domain varies as $32 \times 2^{l/2}$, $l \in \{0, 1, ..., L\}$, leading to a mesh size h_l that varies as $h_l = h_0 \times 2^{-l}$. An in-built automatic differentiation module within the FEniCS library is used to compute the sensitivities of the QoI with respect to design parameters. Once again, the XMC software library [1] is used to implement the CMLMC procedure.

Similar to Section 5.1, we seek to examine both parts of the optimisation algorithm; namely the CMLMC and the gradient based OUU algorithm. For the CMLMC, we seek to accurately estimate $\hat{J}_{w}(\cdot, z_0)$, where $z_0 = [0.1]^9$, such that MSE $(\hat{J}_{w}(\cdot, z_0))$ satisfies a prescribed tolerance. Fig. 5 shows the results of reliability and complexity studies conducted for the above parameters, similar to the one conducted for the FitzHugh–Nagumo system in Section 5.1. For studying the reliability of the error estimators, we conduct 20 independent CMLMC simulations for a given tolerance. For each simulation, we plot three errors; namely the true L^{∞} error on the gradient, the square root of the MSE estimate produced by our error estimation procedure described in Section 4, and the true pointwise error on the gradient, computed by evaluating the parametric expectations $\hat{J}_{w}(\cdot, z_0)$ for the gradient at θ_0 , the VaR corresponding to the design z_0 . The reference value of the gradient is computed by first running 20 simulations for a tolerance that is half of the finest tested tolerance, and averaging over the gradient estimates produced by these simulations. Similar to before, we find that although our novel error estimators provide a tight bound on the true L^{∞} error of the gradient, the L^{∞} error on the gradient is significantly larger than the error on the gradient evaluated at θ_0 . Fig. 5b presents the complexity results of the CMLMC algorithm. The cost to compute the optimal hierarchy for a given tolerance ϵ^2 on MSE $(\tilde{\nabla} \hat{J}^j(\cdot, z_0))$ is plotted versus the tolerance, for each of the 20 CMLMC simulations at a given tolerance, in addition to their sample average value. In addition, the theoretical cost growth rate of a comparable Monte Carlo estimator is shown, as well as the estimated cost of the estimator for reference and comparison. The Monte Carlo reference cost is computed as described in [5]. As can be seen from the figure, the complexity follows the theoretically predicted complexity e^{-2} .

For the OUU, we wish to minimise an objective function of the form in Eq. (10), with $z_{ref} = 0$ and for significance of $\tau = 0.7$. This implies that we seek to minimise the CVaR while also minimising the amplitude of the controlled sinks. We utilise Algorithm 2, starting from a design $z_0 = [0.1]^9$, and halt the optimisation once a gradient ratio of r = 0.08 has been achieved. In Fig. 6, we show the source field f(x), the control field $B(x, z^*)$ and the solution $u(x, z^*, \omega)$ for the mean conditions $a(\omega) = 4$ and $b(\omega) = 5$ at the optimal control z^* obtained by solving problem (10).

S. Ganesh and F. Nobile

Journal of Computational Physics 495 (2023) 112523







Fig. 7. Optimization performance over different iterations for the pollutant transport problem.



Fig. 8. Hierarchy and complexity behaviour for different iterations for the pollutant transport problem.

Fig. 7a shows the decay of the objective function towards its final value. We once again observe Q-linear convergence in the optimisation counter *j*, as predicted by Theorem 3.1. In addition, we plot in 3b the gradient ratio for different iterations of the optimisation, which also decreases Q-linearly in the iteration counter *j*. Fig. 7c shows the CDF of the output QoI $Q(z_j, \cdot)$ for different iterations *j* of the optimisation algorithm, along with the estimated VaR and CVaR. The CDF, the VaR and the CVaR all move left as before in Section 5.1, which translates to moving the right tail of the distribution as much as possible to the left.

Fig. 8a shows the optimal hierarchy produced by the CMLMC algorithm at each iteration of the optimisation for a given tolerance. Similar to before, we observe that the optimally tuned hierarchy increases in size for later optimisation iterations, since the tolerance supplied to the CMLMC is a fraction of the gradient magnitude. Fig. 8b shows the cumulative cost as defined in Section 5.1 for a given gradient magnitude. We observe once again that the cumulative cost grows as $\|\tilde{\nabla} \hat{\mathcal{J}}^{j}(w_{j})\|_{l^{2}}^{-2}$ after an initial pre-asymptotic regime, as is to be expected for the use of an optimally tuned MLMC hierarchy at each iteration, tuned to obtain a tolerance proportional to $\|\tilde{\nabla} \hat{\mathcal{J}}^{j}(w_{j})\|_{l^{2}}$.

We now wish to study the performance of the AMGD algorithm for different significances τ . To this end, we compare the performance of the algorithm for $\tau = 0.7$ and $\tau = 0.9$. Since the performance of the algorithm in terms of objective function and



Fig. 9. Hierarchy and complexity behaviour for different significances for pollutant transport problem.

gradient decay in the $\tau = 0.9$ case are nearly identical to the performance observed in Fig. 8 for the $\tau = 0.7$, the corresponding results are not presented here. Fig. 9a shows the optimal hierarchy produced by the CMLMC algorithm at each optimisation iteration for the two significances tested. We observe that the level-wise sample sizes N_l decay at the same rate in the levels l for both tested significances, however with a larger constant for the $\tau = 0.9$ case. Additionally, Fig. 9b shows the cumulative cost for a given gradient magnitude, for both significances. We observe that the cumulative cost grows as $\|\tilde{\nabla}\hat{J}^j(w_j)\|_{l^2}^{-2}$ in both cases, following an initial preasymptotic regime. However, the $\tau = 0.9$ case shows a larger constant. In this case, the interval Θ is changed with each optimisation iteration such that it is centered on the quantile estimate corresponding to the previous optimisation iteration. It can be shown, for the case of a simple Monte Carlo estimator, that the constant is expected to scale in this case as $(1 - \tau)^{-1}$. The interested reader is referred to [5] for further discussion. We note that we observe a similar scaling in this case.

6. Conclusions

The aim of this work was to tackle the challenge of minimising the CVaR of a random QoI, typically the output of a differential model with random inputs, over a suitable design space, using gradient-based optimisation techniques. A main challenge in utilising gradient-based techniques was the differentiability of the CVaR in terms of the design variables. A differentiability result was presented in Section 2, which was a generalisation of the one presented in [21], showing that gradient-based algorithms could still be used to directly minimise the CVaR without requiring smoothing.

The expression for the sensitivities of the CVaR with respect to design parameters required the computation of expectations of discontinuous functions of the QoI; namely, the indicator function. Estimating this expectation naively using MLMC estimators could become impractically expensive, and possibly result in non-optimal complexity behaviour of the corresponding MLMC estimator. A similar issue was discussed and tackled in [5], and an alternative was proposed using the framework of parametric expectations. We presented a modified expression for the sensitivities of the CVaR, based on derivatives of parametric expectations, thereby allowing us to use the work in [5]. Based on this modification, we also presented a novel optimisation algorithm consisting of an alternating minimisation-gradient procedure. We demonstrated a theoretical result that, under additional assumptions on the combined objective function in Eq. (10), the novel algorithm would achieve Q-linear convergence of the design iterates towards the optimal design in the optimisation iterations.

To enable the use of the work in [5], we presented modifications of the MLMC estimator, the error estimation procedure and adaptive hierarchy selection procedure specific to computing the sensitivities of the CVaR. Namely, a relation was derived between the MSE of the sensitivities and the MSE of the parametric expectations in Section 4. In addition, a modification of the KDE smoothing procedure presented in [5] was presented, specific to CVaR minimisation. The combination of the MSE relation and KDE modification allowed us to trivially extend the error estimation and hierarchy adaptivity procedure of [5] to the current application. Lastly, a minor modification of the CMLMC procedure of [5] was presented in Algorithm 2, wherein the CMLMC was restarted from the optimal hierarchy of the previous design iterate.

The combination of gradient-based optimisation and MLMC estimation of the sensitivities of the CVaR was tested on two problems of practical relevance; namely the FitzHugh–Nagumo oscillator and a more applied problem of advection-reaction-diffusion problem used to model pollutant transport. In both cases, it was observed that the novel error estimation procedure provided tight bounds on the MSE of the gradient as defined in Eq. (47). In addition, the CMLMC algorithm was shown to produce the best-case complexity behaviour for the MLMC estimators of the sensitivities. The OUU algorithm was shown to converge Q-linearly in the optimisation iterations, while also preserving the best case MLMC cost complexity.

The numerical examples considered in this work demonstrated that the AMGD procedure performs well for the cases presented here. However, one may wish to improve on the performance of the algorithm by considering alternatives to the AMGD algorithm. Such variations could, for example, include higher-order optimisation methods such as the Newton method. It still remains to be seen whether higher-order methods can directly be used with objective functions of the type in problem (10), as well as whether the framework of parametric expectations can be combined with such an algorithm. We plan to explore such questions in future works.

CRediT authorship contribution statement

Sundar Ganesh: Conceptualization, Methodology, Software, Validation, Writing – Original draft preparation Fabio Nobile: Supervision, Writing – Reviewing and Editing

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Sundar Ganesh reports financial support was provided by European Union.

Data availability

Data will be made available on request.

Acknowledgements

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 800898 and the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research (OSR) under Award No. OSR-2019-CRG8-4033.

Appendix A. Proof of Theorem 2.1

To prove Theorem 2.1 on the Fréchet differentiability of the objective function $\mathcal{J}(\theta, z)$, we first prove an important result in Lemma A.1. We recall that $\Gamma \subset L^p(\Omega, \mathbb{R})$ is the set of L^p -integrable random variables whose measures are atom-free.

Lemma A.1. Consider random variables $Y \in \Gamma \subset L^{p}(\Omega, \mathbb{R})$ and $\delta Y \in L^{p}(\Omega, \mathbb{R})$. We then have the following:

$$\lim_{\|\delta Y\|_{L^p} \to 0} \mathbb{E}\left[\mathbb{1}_{\{0 \le Y \le -\delta Y\}}\right] = 0,$$
(85)

and
$$\lim_{\|\delta Y\|_{l_p} \to 0} \mathbb{E}\left[\mathbb{1}_{\{-\delta Y \le Y \le 0\}}\right] = 0.$$
(86)

Proof. We begin with the proof for Eq. (85), since the proof for Eq. (86) follows from identical arguments. We make use of the following result; for any $X \in L^{p}(\Omega, \mathbb{R})$, the following holds for any $\epsilon > 0$ and $p \ge 0$:

$$\mathbb{E}\left[\mathbb{1}_{\{|X|\geq\epsilon\}}\right] \leq \mathbb{E}\left[\frac{|X|^{p}}{\epsilon^{p}}\right] = \|X\|_{\mathrm{L}^{p}}^{p} \epsilon^{-p}.$$
(87)

Setting $\epsilon = ||X||_{L^p}^{\beta}$ for some $\beta \in [0, 1)$, we have that

$$\mathbb{E}\left[\mathbb{1}_{\{|X| \ge \|X\|_{L^{p}}^{\beta}\}}\right] \le \|X\|_{L^{p}}^{p-\beta p} = \|X\|_{L^{p}}^{\gamma},$$
(88)

where $\gamma := p(1 - \beta)$. We rewrite the term within the limit in Eq. (85) as follows:

$$\mathbb{E}\left[\mathbb{1}_{\{0\leq Y\leq -\delta Y\}}\right] = \mathbb{E}\left[\mathbb{1}_{\{0\leq Y\leq -\delta Y\}}\left(\mathbb{1}_{\{|\delta Y|<\|\delta Y\|_{L^{p}}^{\beta}\}} + \mathbb{1}_{\{|\delta Y|\geq \|\delta Y\|_{L^{p}}^{\beta}\}}\right)\right]$$

$$\tag{89}$$

$$= \mathbb{E}\left[\mathbbm{1}_{\{0\leq Y\leq -\delta Y\}}\mathbbm{1}_{\{|\delta Y|<\|\delta Y\|_{L^{p}}^{\beta}}\right] + \mathbb{E}\left[\mathbbm{1}_{\{0\leq Y\leq -\delta Y\}}\mathbbm{1}_{\{|\delta Y|\geq \|\delta Y\|_{L^{p}}^{\beta}\}}\right].$$
(90)

The first term can be bounded as follows:

$$\mathbb{E}\left[\mathbb{1}_{\{0\leq Y\leq-\delta Y\}}\mathbb{1}_{\{\|\delta Y\|<\|\delta Y\|_{L^{p}}^{\beta}\}}\right]\leq\mathbb{E}\left[\mathbb{1}_{\{0\leq Y\leq\|\delta Y\|_{L^{p}}^{\beta}\}}\right]$$
(91)

Due to dominated convergence, we can pass the limit into the expectation, resulting in the following:

$$\lim_{\|\delta Y\|_{L^{p}}\to 0} \mathbb{E}\left[\mathbb{1}_{\{0\leq Y\leq \|\delta Y\|_{L^{p}}^{\beta}\}}\right] = \mathbb{E}\left[\lim_{\|\delta Y\|_{L^{p}}\to 0}\mathbb{1}_{\{0\leq Y\leq \|\delta Y\|_{L^{p}}^{\beta}\}}\right] = \mathbb{E}\left[\mathbb{1}_{\{Y=0\}}\right] = 0,$$
(92)

since $Y \in \Gamma$ is atom-free. The second term can be bounded as follows, where we use a Hölder inequality:

$$\mathbb{E}\left[\mathbb{1}_{\{0\leq Y\leq-\delta Y\}}\mathbb{1}_{\{|\delta Y|\geq \|\delta Y\|_{L^{\rho}}^{\beta}\}}\right] \leq \left\|\mathbb{1}_{\{0\leq Y\leq-\delta Y\}}\right\|_{L^{\infty}}\left\|\mathbb{1}_{\{|\delta Y|\geq \|\delta Y\|_{L^{\rho}}^{\beta}\}}\right\|_{L^{1}}$$
(93)

$$\leq \mathbb{E} \left[\mathbb{1}_{\{ |\delta Y| \ge \|\delta Y\|_{1,p}^{\beta} \}} \right] \tag{94}$$

$$\leq \|\delta Y\|_{L^{p}}^{\gamma}.$$
(95)

Hence, we have that the second term in Eq. (90) goes to zero as well with the application of the limit, thus concluding the proof for Eq. (85). The proof for Eq. (86) follows from identical arguments. \Box

We now use the above result and present a proof of Theorem 2.1. We note that the function $\mathcal{R}(\theta, Q)$ is a composition of two functions. We define the functions $I_1 : \Gamma \to \mathbb{R}$ and $I_2 : \mathbb{R} \times \Gamma \to \Gamma$ as follows:

$$l_1(Y) := \mathbb{E}\left[Y^+\right] \tag{96}$$

$$l_{2}(\theta, Q) := Q - \theta,$$

$$\implies \mathcal{R}(\theta, Q) = \theta + \frac{l_{1} \circ l_{2}(\theta, Q)}{1 - \tau}.$$
(97)
(98)

Hence, to show that \mathcal{R} is Fréchet differentiable, it suffices to show that each of the functions l_1 and l_2 are Fréchet differentiable.

It is straightforward to see that l_2 is Fréchet differentiable (being linear and bounded) with Fréchet derivative $Dl_2(\theta, Q)$ in the direction $(\delta\theta, \delta Q) \in \mathbb{R} \times L^p(\Omega, \mathbb{R})$ given by:

$$Dl_2(\theta, Q)(\delta\theta, \delta Q) = \delta Q - \delta\theta.$$
⁽⁹⁹⁾

The Fréchet derivative of l_1 however, requires some consideration. We argue that the Fréchet derivative of l_1 exists at any point $Y \in \Gamma$ and is given by $Dl_1(Y)(\delta Y) = \mathbb{E}\left[\mathbb{1}_{\{Y>0\}}\delta Y\right]$. To prove this statement, we must verify the following limit:

$$\lim_{\|\delta Y\|_{L^{p}}\to 0} \frac{\left|\mathbb{E}\left[(Y+\delta Y)^{+}\right] - \mathbb{E}\left[Y^{+}\right] - \mathbb{E}\left[\mathbb{1}_{\{Y\geq 0\}}\delta Y\right]\right|}{\|\delta Y\|_{L^{p}}} = 0$$
(100)

To show the above, we begin by re-writing the numerator as follows:

$$\mathbb{E}\left[(Y+\delta Y)^{+}-Y^{+}-\mathbb{1}_{\{Y\geq0\}}\delta Y\right] = \mathbb{E}\left[\delta Y\mathbb{1}_{\{Y+\delta Y\geq0,Y\geq0\}}-\mathbb{1}_{\{Y\geq0\}}\delta Y\right]$$
$$+\mathbb{E}\left[(Y+\delta Y)\mathbb{1}_{\{Y+\delta Y\geq0,Y<0\}}\right]$$
$$-\mathbb{E}\left[Y\mathbb{1}_{\{Y+\delta Y<0,Y\geq0\}}\right].$$
(101)

Inserting Eq. (101) into Eq. (100), we have the following:

$$\frac{\left|\mathbb{E}\left[(Y+\delta Y)^{+}\right]-\mathbb{E}\left[Y^{+}\right]-\mathbb{E}\left[\mathbb{1}_{\{Y\geq 0}\delta Y\}\right]\right|}{\left\|\delta Y\right\|_{L^{p}}}\leq\frac{T_{1}+T_{2}+T_{3}}{\left\|\delta Y\right\|_{L^{p}}},$$
(102)

with the terms T_1 , T_2 and T_3 given by:

.

$$T_1 := \left| \mathbb{E} \left[\delta Y \mathbbm{1}_{\{Y + \delta Y \ge 0, Y \ge 0\}} - \delta Y \mathbbm{1}_{\{Y \ge 0\}} \right] \right|,\tag{103}$$

$$T_2 := \left| \mathbb{E} \left[(Y + \delta Y) \mathbb{1}_{\{Y + \delta Y \ge 0, Y < 0\}} \right] \right|,\tag{104}$$

$$T_{3} := \left| \mathbb{E} \left[Y \mathbb{1}_{\{Y + \delta Y < 0, Y \ge 0\}} \right] \right|.$$
(105)

We then begin with the term T_1 . We first note that T_1 can be rewritten in the following manner:

$$T_1 = \left| -\mathbb{E}\left[\delta Y \mathbb{1}_{\{0 \le Y < -\delta Y\}} \right] \right| \le \mathbb{E}\left[|\delta Y| \mathbb{1}_{\{0 \le Y < -\delta Y\}} \right]$$
(106)

$$\leq \|\delta Y\|_{L^{p}} \left\| \mathbb{1}_{\{0 \leq Y < -\delta Y\}} \right\|_{L^{q}} = \|\delta Y\|_{L^{p}} \mathbb{E} \left[\mathbb{1}_{\{0 \leq Y < -\delta Y\}} \right]^{1/q}, \tag{107}$$

$$\leq \|\delta Y\|_{L^p} \mathbb{E}\left[\mathbbm{1}_{\{0 \leq Y \leq -\delta Y\}}\right]^{1/q}.$$
(108)

The term T_2 can be bounded as follows:

$$T_{2} = \left| \mathbb{E} \left[(Y + \delta Y) \mathbb{1}_{\{Y + \delta Y \ge 0\}} \mathbb{1}_{\{Y < 0\}} \right] \right| \le \mathbb{E} \left[|\delta Y| \mathbb{1}_{\{-\delta Y \le Y < 0\}} \right],$$

$$\le \|\delta Y\|_{L^{p}} \left\| \mathbb{1}_{\{-\delta Y \le Y < 0\}} \right\|_{L^{q}} = \|\delta Y\|_{L^{p}} \mathbb{E} \left[\mathbb{1}_{\{-\delta Y \le Y \le 0\}} \right]^{1/q}.$$

$$(109)$$

Similarly, the term
$$T_3$$
 can be bounded as follows:

$$T_{3} = \left| \mathbb{E} \left[Y \mathbb{1}_{\{Y+\delta Y<0\}} \mathbb{1}_{\{Y\geq 0\}} \right] \right| \leq \mathbb{E} \left[|\delta Y| \mathbb{1}_{\{0\leq Y<-\delta Y\}} \right]$$

$$\leq \|\delta Y\|_{L^{p}} \left\| \mathbb{1}_{\{0\leq Y\leq-\delta Y\}} \right\|_{L^{q}} = \|\delta Y\|_{L^{p}} \mathbb{E} \left[\mathbb{1}_{\{0\leq Y\leq-\delta Y\}} \right]^{1/q}.$$

$$(111)$$

Inserting Eqs. (108), (110) and (112) into Eq. (102), and applying the limit using Lemma A.1, we have that:

$$\lim_{\|\delta Y\|_{L^{p}}\to 0} \frac{\left|\mathbb{E}\left[(Y+\delta Y)^{+}\right] - \mathbb{E}\left[Y^{+}\right] - \mathbb{E}\left[\mathbb{1}_{\{Y\geq 0\}}\delta Y\right]\right|}{\|\delta Y\|_{L^{p}}} = 0.$$
(113)

This concludes the proof.

Appendix B. Adjoint of first-order ODE with additive noise

We present here the derivation of the adjoints for a first-order Ordinary Differential Equation (ODE) with white noise forcing for an objective function containing the CVaR of a time-averaged quantity of the trajectory. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space, $\omega \in \Omega$ denote an elementary random event, and $z \in \mathbb{R}^d$ the set of design variables. Let $u(t, z, \omega) \in U \subset \mathbb{R}^{N_u}$ be the state vector at time $t \in [0, T]$ for a given random input ω and design z. The state vector u is governed by the following ODE with additive noise.

$$\dot{u}(t,z,\omega) = g(u,z) + \tau \dot{W}(t,\omega) \quad \text{over} \ (0,T], \tag{114}$$

$$u(0,z,\omega) = u^0,$$
 (115)

where $g: U \times \mathbb{R}^d \to \mathbb{R}^{N_u}$, and $W: [0,T] \times \Omega \to \mathbb{R}^{N_u}$ is a N_u -dimensional standard Wiener process.

We discretise the problem on a uniform temporal grid \mathbb{T} where the interval [0,T] is divided into $N \in \mathbb{N}$ segments of step size $\Delta t = T/N$, $\mathbb{T} := \{t_n := n\Delta t : n \in [0, N_1]\}$. The ODE is discretised using the Euler–Maruyama scheme, which reads as follows:

$$u^{n+1} = u^n + \Delta t g(u^n, z) + \tau \sqrt{\Delta t} \xi^n,$$

$$u^0 = u_0,$$

where u^n denotes the approximation to $u(t_n, z, \omega)$, $\xi^n \in \mathbb{R}^{N_u}$ are N_u -dimensional random vectors whose components are independent identically distributed standard normal variables. We are interested in computing the statistics of time-averages of functions of the trajectory.

$$Q = \langle f(u) \rangle_T. \tag{116}$$

We approximate the time integral using the trapezoid rule on the aforementioned temporal grid, leading to

$$Q(z,\omega) \approx Q_h(z,\omega) := \sum_{n=0}^{N-1} \left(\frac{f(u^n) + f(u^{n+1})}{2} \right) \frac{\Delta t}{T}.$$
(117)

We are interested in minimising the CVaR of this quantity over the parameters z but use the combined formulation in Eq. (10). The corresponding Lagrangian for the problem reads

$$\mathcal{L}(\theta, z, \{u^n\}, \{\lambda^n\}) = \theta + \frac{\mathbb{E}\left[(Q(z, \cdot) - \theta)^+\right]}{1 - \tau} + \mathbb{E}\left[\sum_{n=0}^{N-1} \lambda^{n+1} \left(u^n + \Delta t g^n + \tau \sqrt{\Delta t} \xi^n - u^{n+1}\right) - \lambda^0 (u^0 - u_0)\right],\tag{118}$$

where we use $g^n := g(u^n, z)$, and $\lambda^n \in \mathbb{R}^{N_u}$, $n \in [0, N]$ denote the Lagrange multipliers for the initial condition and the steps of the discretised equations.

Differentiating with respect to z gives

$$\frac{d\mathcal{L}}{dz} = \mathbb{E}\left[\frac{\mathbb{1}_{Q_h \ge \theta}}{(1 - \tau)T} \sum_{n=0}^{N-1} \left(\frac{f_u^n u_z^n + f_u^{n+1} u_z^{n+1}}{2}\right) \Delta t\right] \\
+ \mathbb{E}\left[\sum_{n=0}^{N-1} \lambda^{n+1} \left(u_z^n + \Delta t (g_u^n u_z^n + g_z^n) - u_z^{n+1}\right)\right] \\
=: \mathbb{E}\left[\hat{\mathcal{L}}\right].$$
(119)

Re-arranging the terms leads to

$$\hat{\mathcal{L}} = u_z^0 \left[\lambda^1 (1 + \Delta t g_u^0) + \frac{\mathbb{1}_{Q_h \ge \theta}}{(1 - \tau)T} \frac{f_u^0 \Delta t}{2} \right] + \Delta t \lambda^1 g_z^0 + u_z^N \left[\frac{\mathbb{1}_{Q_h \ge \theta}}{(1 - \tau)T} \frac{f_u^N \Delta t}{2} - \lambda^N \right] \\ + \sum_{n=1}^{N-1} u_z^n \left[\lambda^{n+1} (1 + \Delta t g_u^n) - \lambda^n + \frac{\mathbb{1}_{Q_h \ge \theta}}{(1 - \tau)T} \Delta t f_u^n \right] + \Delta t \lambda^{n+1} g_z^n,$$
(121)

where we have used the subscript notation for partial derivatives.

We have in our case that $u_z^0 = 0$. To remove terms dependent on u_z^n , we set

$$\lambda^{n} = \lambda^{n+1} (1 + \Delta t g_{u}^{n}) + \frac{\mathbb{1}_{Q_{h} \ge \theta}}{(1 - \tau)T} \Delta t f_{u}^{n}, \quad n = 1, ..., N - 1$$
(122)

$$\lambda^{N} = \frac{\mathbb{I}_{Q_{h} \ge \theta}}{(1 - \tau)T} \frac{f_{u}^{N} \Delta t}{2}.$$
(123)

This gives us the adjoint equations which are solved backwards in time. It is noteworthy to mention that since Eq. (122) is linear, that it can be solved for $\{\lambda^n\}$ without the factor $\frac{\mathbb{I}_{Q_h \ge \theta}}{(1-\tau)T}$, and equivalently, the sensitivities can be computed as:

$$\frac{\mathrm{d}\mathcal{L}}{\mathrm{d}z} = \frac{\mathbbm{1}_{Q_l \geq \theta}}{(1-\tau)T} \mathbb{E}\left[\sum_{n=0}^{N-1} \Delta t \,\lambda^{n+1} g_z^n\right].$$

That is, setting

$$\mathcal{J}(\theta, z) = \theta + \frac{\mathbb{E}\left[(\mathcal{Q}_{h}(z, \cdot) - \theta)^{+}\right]}{1 - \tau},$$
(125)

we have that $\mathcal{J}_{z}(\theta, z) = \mathbb{E}\left[\sum_{n=0}^{N-1} \Delta t \lambda^{n+1} g_{z}^{n}\right].$

References

- Ramon Amela, Quentin Ayoul-Guilmard, Rosa M. Badia, Sundar Ganesh, Fabio Nobile, Riccardo Rossi, Riccardo Tosi, ExaQUte XMC, Comp. software. ExaQUte Consortium, Oct. 2020, https://dx.doi.org/10.5281/zenodo.4265429.
- [2] P.R. Amestoy, A. Buttari, J.-Y. L'Excellent, T. Mary, Performance and scalability of the block low-rank multifrontal factorization on multicore architectures, ACM Trans. Math. Softw. 45 (1 2019) 2:1–2:26.
- [3] P.R. Amestoy, I.S. Duff, J. Koster, J.-Y. L'Excellent, A fully asynchronous multifrontal solver using distributed dynamic scheduling, SIAM J. Matrix Anal. Appl. 23 (1) (2001) 15–41.
- [4] Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, David Heath, Coherent measures of risk, Math. Finance 9 (3) (1999) 203-228.
- [5] Quentin Ayoul-Guilmard, Sundar Ganesh, Sebastian Krumscheid, Fabio Nobile, Quantifying uncertain system outputs via the multilevel Monte Carlo method distribution and robustness measures, Int. J. Uncertain. Quantificat. 13 (5) (2023) 61–98, https://doi.org/10.1615/Int.J.UncertaintyQuantification.2023045259.
- [6] Quentin Ayoul-Guilmard, Sundar Ganesh, Fabio Nobile, Report on stochastic optimisation for simple problems, Deliverable 6.3. Version 1.3, ExaQUte Consortium, July 27, 2020, https://dx.doi.org/10.23967/exaqute. 2021.2.001.
- [7] Florian Beiser, Brendan Keith, Simon Urbainczyk, Barbara Wohlmuth, Adaptive sampling strategies for risk-averse stochastic optimization with constraints, IMA J. Numer. Anal. (2023), https://doi.org/10.1093/imanum/drac083.
- [8] Raghu Bollapragada, Richard Byrd, Jorge Nocedal, Adaptive sampling strategies for stochastic optimization, SIAM J. Optim. 28 (4) (2018) 3312–3343.
- [9] Tobias Breiten, Karl Kunisch, Riccati-based feedback control of the monodomain equations with the Fitzhugh–Nagumo model, SIAM J. Control Optim. 52 (6) (2014) 4057–4081.
- [10] Richard H. Byrd, Gillian M. Chin, Jorge Nocedal, Yuchen Wu, Sample size selection in optimization methods for machine learning, Math. Program. 134 (1) (2012) 127–155.
- [11] Anirban Chaudhuri, Benjamin Peherstorfer, Karen Willcox, Multifidelity cross-entropy estimation of conditional value-at-risk for risk-averse design optimization, in: AIAA Scitech 2020 Forum, 2020, p. 2129.
- [12] Nathan Collier, Abdul-Lateef Haji-Ali, Fabio Nobile, Erik von Schwerin, Raúl Tempone, A continuation multilevel Monte Carlo algorithm, BIT Numer. Math. (2014) 1–34.
- [13] Elena A. Ermakova, Mikhail A. Panteleev, Emmanuil E. Shnol, Blood coagulation and propagation of autowaves in flow, Pathophysiol. Haemost. Thromb. 34 (2–3) (2005) 135–142.
- [14] Richard FitzHugh, Impulses and physiological states in theoretical models of nerve membrane, Biophys. J. 1 (6) (1961) 445-466.
- [15] Michael B. Giles, Multilevel Monte Carlo methods, Acta Numer. 24 (2015) 259–328.
- [16] Michael B. Giles, Multilevel Monte Carlo path simulation, Oper. Res. 56 (3) (2008) 607–617.
- [17] Philipp A. Guth, Vesa Kaarnioja, Frances Y. Kuo, Claudia Schillings, Ian H. Sloan, Parabolic PDE-constrained optimal control under uncertainty with entropic risk measure using quasi-Monte Carlo integration, arXiv preprint, arXiv:2208.02767, 2022.
- [18] Alan L. Hodgkin, Andrew F. Huxley, A quantitative description of membrane current and its application to conduction and excitation in nerve, J. Physiol. 117 (4) (1952) 500–544.
- [19] Håkon Hoel, Erik von Schwerin, Anders Szepessy, Raúl Tempone, Adaptive multilevel Monte Carlo simulation, in: Numerical Analysis of Multiscale Computations, Springer, 2012, pp. 217–234.
- [20] L. Jeff Hong, Guangwu Liu, Monte Carlo estimation of value-at-risk, conditional value-at-risk and their sensitivities, in: Proceedings of the 2011 Winter Simulation Conference (WSC), IEEE, 2011, pp. 95–107.
- [21] L. Jeff Hong, Guangwu Liu, Simulating sensitivities of conditional value at risk, Manag. Sci. 55 (2) (2009) 281-293.
- [22] Anoop Kodakkal, Brendan Keith, Ustim Khristenko, Andreas Apostolatos, Kai-Uwe Bletzinger, Barbara Wohlmuth, Roland Wüchner, Risk-averse design of tall buildings for uncertain wind conditions, Comput. Methods Appl. Mech. Eng. 402 (2022) 115371.
- [23] Drew P. Kouri, Thomas M. Surowiec, Risk-averse PDE-constrained optimization using the conditional value-at-risk, SIAM J. Optim. 26 (1) (2016) 365–396.
- [24] Sebastian Krumscheid, Fabio Nobile, Multilevel Monte Carlo approximation of functions, SIAM/ASA J. Uncertain. Quantificat. 6 (3) (2018) 1256–1293.
- [25] Sebastian Krumscheid, Fabio Nobile, M. Pisaroni, Quantifying uncertain system outputs via the multilevel Monte Carlo method—Part I: central moment estimation, J. Comput. Phys. 414 (2020) 109466.
- [26] Churlzu Lim, Hanif D. Sherali, Stan Uryasev, Portfolio optimization by minimizing conditional value-at-risk via nondifferentiable optimization, Comput. Optim. Appl. 46 (3) (2010) 391–415.
- [27] A.I. Lobanov, T.K. Starozhilova, The effect of convective flows on blood coagulation processes, Pathophysiol. Haemost. Thromb. 34 (2-3) (2005) 121-134.
- [28] Anders Logg, Kent-Andre Mardal, Garth Wells, Automated Solution of Differential Equations by the Finite Element Method: The FEniCS Book, vol. 84, Springer Science & Business Media, 2012.
- [29] Matthieu Martin, Fabio Nobile, Panagiotis Tsilifis, A multilevel stochastic gradient method for PDE-constrained optimal control problems with uncertain parameters, arXiv preprint, arXiv:1912.11900, 2019.
- [30] Jinichi Nagumo, Suguru Arimoto, Shuji Yoshizawa, An active pulse transmission line simulating nerve axon, Proc. IRE 50 (10) (1962) 2061–2070.
- [31] Benjamin Peherstorfer, Boris Kramer, Karen Willcox, Multifidelity preconditioning of the cross-entropy method for rare event simulation and failure probability estimation, SIAM/ASA J. Uncertain. Quantificat. 6 (2) (2018) 737–761.
- [32] Michele Pisaroni, Fabio Nobile, Penelope Leyland, Continuation multilevel Monte Carlo evolutionary algorithm for robust aerodynamic shape design, J. Aircr. 56 (2) (2019) 771–786.
- [33] Domenico Quagliarella, Value-at-risk and conditional value-at-risk in optimization under uncertainty, in: Uncertainty Management for Robust Industrial Design in Aeronautics, Springer, 2019, pp. 541–565.
- [34] Domenico Quagliarella, Elisa Morales Tirado, Andrea Bornaccioni, Risk measures applied to robust aerodynamic shape design optimization, in: Flexible Engineering Toward Green Aircraft, Springer, 2020, pp. 153–168.

21

- [35] R. Tyrrell Rockafellar, Stanislav Uryasev, Conditional value-at-risk for general loss distributions, J. Bank. Finance 26 (7) (2002) 1443–1471.
- [36] Andrzej Ruszczyński, Alexander Shapiro, Optimization of convex risk functions, Math. Oper. Res. 31 (3) (2006) 433–452.
- [37] David W. Scott, On optimal and data-based histograms, Biometrika 66 (3) (1979) 605-610.

[38] Alexander Shapiro, Darinka Dentcheva, Andrzej Ruszczyński, Lectures on Stochastic Programming: Modeling and Theory, SIAM, 2009.

- [39] Murat Uzunca, Tugba Küçükseyhan, Hamdullah Yücel, Bülent Karasözen, Optimal control of convective FitzHugh–Nagumo equation, Comput. Math. Appl. 73 (9) (2017) 2151–2169.
- [40] Andreas Van Barel, Stefan Vandewalle, Robust optimization of PDEs with random coefficients using a multilevel Monte Carlo method, SIAM/ASA J. Uncertain. Quantificat. 7 (1) (2019) 174–202.