



HAL
open science

Robust tuning of Robbins-Monro algorithm for quantile estimation in iterative uncertainty quantification

Bertrand Iooss

► **To cite this version:**

Bertrand Iooss. Robust tuning of Robbins-Monro algorithm for quantile estimation in iterative uncertainty quantification. 2020. hal-02918478

HAL Id: hal-02918478

<https://hal.science/hal-02918478>

Preprint submitted on 20 Aug 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Robust tuning of Robbins-Monro algorithm for quantile estimation in iterative uncertainty quantification

Bertrand Iooss, EDF R&D, France. bertrand.iooss@edf.fr

August 20, 2020

Abstract

In uncertainty quantification of a numerical simulation model output, the classical approach for quantile estimation requires the availability of the full sample of the studied variable. This approach is not suitable at exascale as large ensembles of simulation runs would need to gather a prohibitively large amount of data. This problem can be solved thanks to an on-the-fly (iterative) approach based on the Robbins-Monro algorithm. We numerically study this algorithm for estimating a discretized quantile function from samples of limited size (a few hundreds observations). As in practice, the distribution of the underlying variable is unknown, the goal is to define “robust” values of the algorithm parameters, which means that quantile estimates have to be reasonably good in most situations. This paper presents new empirically-validated iterative quantile estimators, for two different practical situations: when the final number of the model runs N is a priori fixed and when N is unknown in advance (it can then be minimized during the study in order to save cpu time cost).

keywords: Uncertainty Quantification, Quantile, Online statistics, Robbins-Monro, Averaging.

1 Introduction

In engineering studies, a numerical simulation model is often used as a tool to assess the safety of a complex industrial system [6, 13]. For such a goal, when developing and using the numerical simulation model, uncertainty quantification (UQ) and global sensitivity analysis are valuable tools [20, 3]. It is then required to run the simulation model several (or even many) times with different values of the model inputs (according to their predefined probability laws) in order to calculate statistical quantities of interest (noted QoI) on the model outputs, i.e. their mean, variance, quantiles, probability of threshold exceedence, sensitivity indices, ... [1, 27, 25]. To estimate these QoI, the usual practice is to run all the simulation and store all their results before calculating the QoI.

A major difficulty arises when the UQ ensemble runs produce massive amount of data (e.g. when state and time variables are simulated) that have to be statistically aggregated, making them extremely vulnerable to the storage issues and I/O bottleneck.

To keep a manageable amount of data, the classical approach, used in most of the studies, consists in taking a limited number of outputs (e.g. by only taking the central point of a spatial map output), see various examples in [8]. A more suitable technique would be to use one-pass statistical algorithms, also called iterative, recursive, update, online or even parallel statistics. Such algorithms only require to store the current results that can next be updated with incoming new samples: the required storage memory is therefore only the one needed for the results of one (or a limited number) simulation. For instance, if we are able to store one spatio-temporal evolution of a simulation output, we are able to compute its statistics after any number of simulations.

The iterative statistical estimation issue is a relatively classical subject in the treatment of large volumes of data, in the so-called *big data* framework [24]. One-pass variance algorithms were proposed for example in [26, 5]. Numerically stable, update formulas for arbitrary centered statistical moments and co-moments are presented in [14] and

were used, for instance, to compute large scale parallel statistics for a single simulation run in [2]. However, iterative estimation issues have been little explored in UQ of numerical models (see [18] for an overview of the different underlying update statistical issues in UQ). In global sensitivity analysis of model outputs, for the estimation of Sobol' indices [21], [9] introduced an iterative computation for the case of a scalar model output, while [22] applied the iterative covariance formulas on massive output data (a spatio-temporal model output). For estimating quantiles of model outputs, a first iterative estimation algorithm has been studied in [18].

In this paper, we focus on the estimation of quantiles, as often required in simulation-based risk assessment, but in an iterative fashion. Indeed, quantiles are essential elements for the calculation of prediction or tolerance intervals, and for the detection of outliers, in particular in safety studies. As in a lot of industrial studies, we have to deal with limited number of simulations [6], typically a few hundreds (see an example in the field of nuclear engineering in [10]). In this work, for the sake of brevity, we only consider a scalar output $Y \in \mathbb{R}$ which writes

$$Y = G(X) \quad (1)$$

where $G(\cdot)$ is the model function and $X \in \mathbb{R}^d$ is the vector of the d random input variables (defined by their joint probability density function). However, we keep in mind that iterative UQ is interesting when the dimension output is very large, which brings storage issues.

The following section presents the Robbins-Monro (RM) algorithm for quantile estimation. It allows to introduce the tuning RM parameters issue that is studied in this paper by the way of intensive numerical tests. Two different situations are then distinguished. First, the final number of iterations (i.e. number of computer model runs) N is a priori fixed, which is a classical way to deal with UQ problems. Section 3 defines well-tuned choices for the RM parameters in this context. Second, N is unknown in advance. Indeed, in practical situations, it is strongly interesting to stop the quantile estimation process when a sufficient precision has been reached in order to save cpu time cost. Section 4 shows the interest to use a stochastic adaptive step-size rule associated to the averaged RM estimator in this context. Section 5 concludes the work.

2 Quantile estimation

We look for an estimator \hat{q}_α of α -quantiles q_α (of the random variable Y) defined by:

$$q_\alpha = \inf\{y \in \mathbb{R} \mid \mathbb{P}(Y \leq y) \geq \alpha\}, \quad (2)$$

with $\alpha \in [\alpha_{\min}, \alpha_{\max}]$ where $\alpha_{\min} (\in]0, 1[)$ and $\alpha_{\max} (\in]0, 1[)$ are the minimal and maximal values of the orders of the estimated quantiles. In our study, α_{\min} (resp. α_{\max}) will be equal to 5% (resp. 95%). The empirical estimator of q_α writes

$$\hat{q}_\alpha^N = Y_{(\lfloor \alpha N \rfloor + 1)}, \quad (3)$$

where $(Y_{(1)}, \dots, Y_{(N)})$ is the ordered sample associated with the i.i.d. sample (Y_1, \dots, Y_N) . The sample (Y_1, \dots, Y_N) comes from a so-called Monte Carlo uncertainty propagation of N i.i.d. values of X through the model $G(\cdot)$ (see Eq. (1)).

Instead of this empirical estimator, the RM algorithm [19] is devoted to iterative quantile estimation. Its asymptotic properties (consistence and central limit theorem) provide essential guarantees of convergence [7]. However, these theoretical results are of little use while tuning the algorithm parameters and our restricted sample size does not reach the asymptotic regime. Indeed, as in [23], our quantile estimation problem has to be done with a small-size sample (a few hundred observations).

The RM algorithm consists in updating the current quantile estimator (noted $q_\alpha(n)$) with each new observation Y_{n+1} with $n \geq 1$ by the recurrence formula

$$q_\alpha(n+1) = q_\alpha(n) - \frac{C}{n^\gamma} (\mathbf{1}_{Y_{n+1} \leq q_\alpha(n)} - \alpha), \quad (4)$$

with $q_\alpha(1) = Y_1$ (initialization step from the first data), $C > 0$ a constant and $\gamma \in]0, 1[$ governing the rate of descent of the stochastic algorithm. The required memory storage for this estimator is two values (one realization of Y and one quantile estimate).

At finite sample size N , the RM estimator of the α -quantile of Y is therefore $\hat{q}_\alpha = q_\alpha(N)$. This estimator is consistent and asymptotically Gaussian for $\gamma \in]0.5, 1[$ [7]. The value of γ does not therefore seem to be of crucial importance but, for low N , we will see that its adjustment is important.

3 Fixed number of model evaluations

In this section, the final size N of available realizations of the model output Y is fixed. The tuning of the γ parameter is first considered; then, the tuning of the C constant is studied.

3.1 Robust tuning of γ via a linear profile

We are looking for a value of γ which gives “acceptable” results whatever the distribution of Y (unknown in practice). Our numerical test considers the cases $Y \sim \mathcal{N}(0, 1)$ and $Y \sim \mathcal{U}[0, 1]$, with $N = 1000$, $C = 1$ and three orders of quantile α (0.05, 0.5 et 0.95). For each of these cases, Figures 1 and 2 show 50 independent trajectories of the RM estimator $q_\alpha(n)$ for $n = 1, \dots, N$ by considering three different choices of γ : 0.6, 1 and a linear variation as a function of n [18] which is written

$$\gamma(n) = 0.5 + 0.5 \frac{n-1}{N-1}. \quad (5)$$

The theoretical and asymptotic properties of the RM algorithm are preserved using Eq. (5), because the γ values lie in $]0.5, 1]$.

The idea of the $\gamma(n)$ profile, given by Eq. (5), is to have strong fluctuations of the estimator at the beginning of the algorithm (to remove its dependence on the values of Y drawn first) then weak fluctuations at the end of the algorithm (to stabilize the estimator at the last iterations). This so-called *search-then-converge learning rule* [16] ensures a regular decrease of the mixing (from strong to weak) all along the RM iterations. Indeed, we can see on Figures 1 and 2 that the fluctuations with $\gamma = 1$ are too small in the Gaussian case ($\gamma = 0.6$ is satisfactory in this case) and the fluctuations with $\gamma = 0.6$ are too strong in the uniform case ($\gamma = 1$ is satisfactory in this case). The profile of a linear variation of γ achieves a compromise between these two extreme cases (and in the many other tests carried out).

3.2 Robust tuning of C

In the previous section, the constant C has been set to 1. This choice turns out to be catastrophic when the variable considered has a dispersion which is not of this order of magnitude. It should be remembered that in practice this dispersion of the

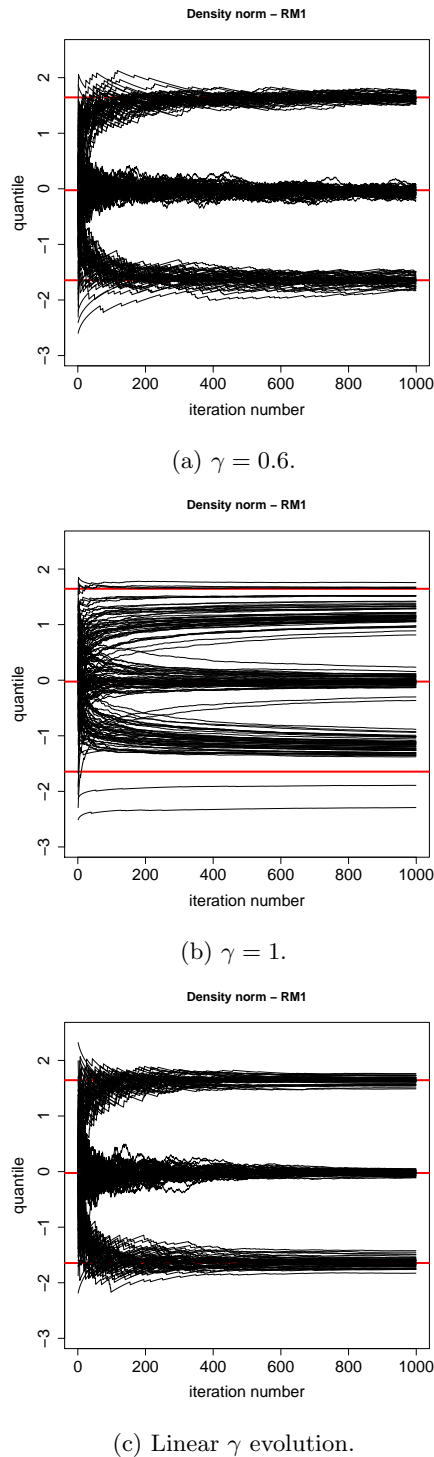


Figure 1: Simulations of trajectories of the RM algorithm ($N = 1000$, $Y \sim \mathcal{N}(0, 1)$). Red lines give exact order quantiles 0.05, 0.5 and 0.95.

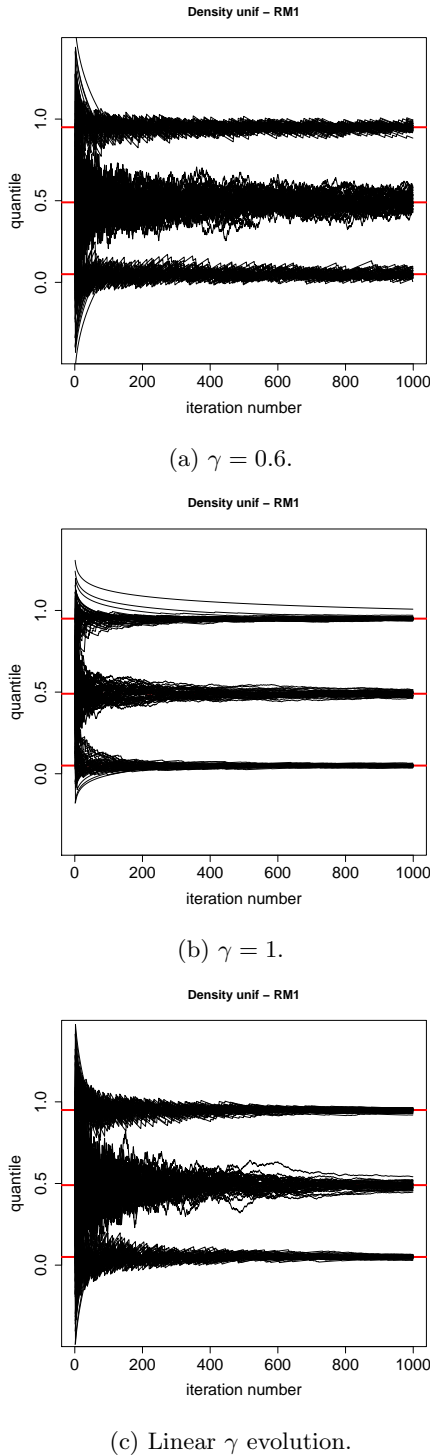


Figure 2: Simulations of trajectories of the RM algorithm ($N = 1000$, $Y \sim \mathcal{U}[0, 1]$). Red lines give exact order quantiles 0.05, 0.5 and 0.95.

studied variable Y is unknown. Then, as for γ , we look for a choice of C which gives “acceptable” results whatever the distribution of Y (unknown in practice).

Figure 3 shows 50 independent trajectories of the RM estimator $q_\alpha(n)$ for $n = 1, \dots, 1000$, Y following a lognormal law ($\log(Y) \sim \mathcal{N}(0, 1)$) and three orders of quantile α (0.05, 0.5 et 0.95). γ has a linear profile and three different settings of C are tested: 1, 10 and an adaptive tuning which is written

$$C(n) = |q_{\alpha_{\max}}(n-1) - q_{\alpha_{\min}}(n-1)|, \quad (6)$$

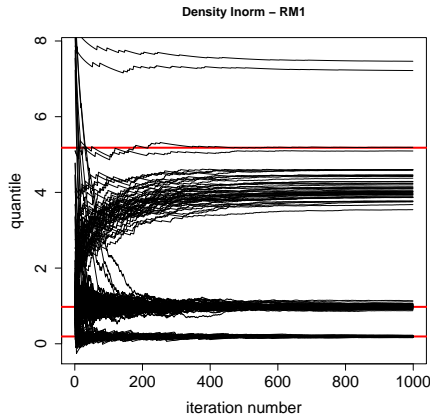
where $n \geq 2$ and $C(1) = |Y_2 - Y_1|$. In Figure 3, it is clear that, for the quantile of order 0.95, C must be large enough to obtain sufficiently large fluctuations from the beginning of the RM algorithm. The adaptive adjustment of C via Eq. (6) allows to automatically regulate these fluctuations. Many other numerical tests on distributions of different types have confirmed the correctness of this choice.

4 Non-fixed number of model evaluations

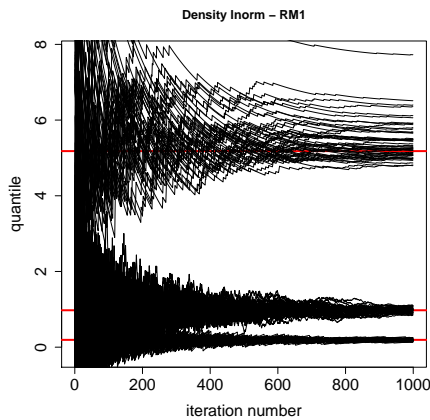
From a user point of view, fixing the number N of model evaluations at the beginning at the study can be quite impractical. Indeed, the user has often no idea of the sample size that is needed to compute the QoI with a sufficient precision. When the estimator is not iterative, a convergence control is easy to provide (for example by bootstrap) in order to stop the simulations (see, e.g., [1]). However, such solutions do not exist in iterative statistical estimation. This difficult issue is left for a future work and we start, in this section, to provide a preliminary study to adapt our RM algorithm when N is unknown. Indeed, the profile that has been chosen in Section 3.1 for γ (Eq. (5)) depends on N . To get rid of the linear profile for γ , we first introduce an averaged RM version; then, a stochastic adaptive stepsize is proposed.

4.1 Averaged version of Robbins-Monro

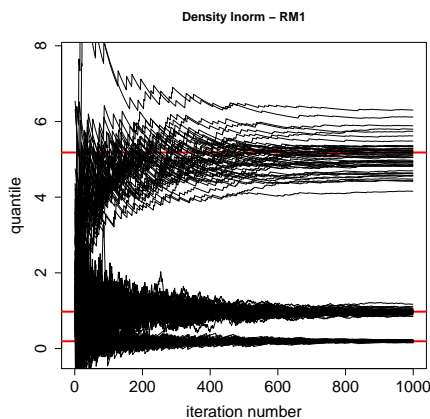
It is well known that the averaged version of RM (noted here ARM) converges faster than the classical RM algorithm of Eq. (4) [15]. The idea is to exploit the basic recursive formula for computing a



(a) $C = 1$.



(b) $C = 10$.



(c) C adaptatif.

Figure 3: Simulations of trajectories of the RM algorithm ($N = 1000$, $Y \sim \mathcal{LN}(0, 1)$). Red lines give exact order quantiles 0.05, 0.5 and 0.95.

mean, to keep a quantile mean update (after applying the RM estimator) at each iteration. By preliminary numerical tests (not shown here), we have however noted that, if an averaged quantile is introduced into (4), the fluctuations of the estimator along the iterations are not of sufficient magnitude to converge to the exact value. It is thus necessary to preserve the formulation (4) for $q_\alpha(n)$ and to store in addition, at each iteration, the averaged estimator (noted $\bar{q}_\alpha(n)$):

$$\bar{q}_\alpha(n+1) = \bar{q}_\alpha(n) + \frac{q_\alpha(n+1) - \bar{q}_\alpha(n)}{n+1}, \quad (7)$$

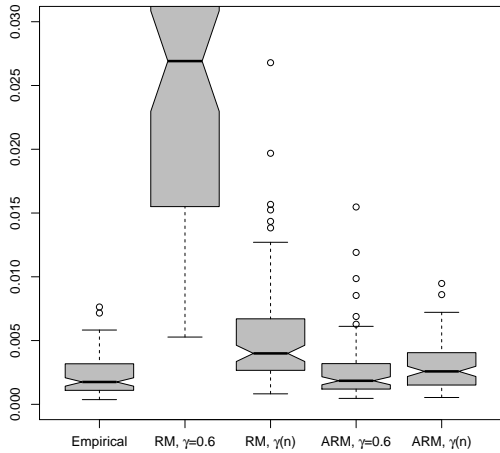
with $n \geq 1$ and $\bar{q}_\alpha(1) = Y_1$. The required memory storage for the ARM estimator is three values (one realization of Y and two quantile estimates).

Figure 4 compares the RM and ARM algorithms for the two cases $Y \sim \mathcal{N}(0, 1)$ and $Y \sim \mathcal{U}[0, 1]$, $N = 1000$ and the adaptive tuning of C (Eq. (6)). The quantiles are estimated for orders α discretized inside the interval $[0.05, 0.95]$ by step of 0.01. The metric used (on the ordinate) is the mean square error between the exact quantiles and the estimated quantiles. Estimates are repeated 100 times independently in order to capture the variability of errors due to sampling. The reference estimator is the empirical estimator (which is not iterative). In these examples, the performances of ARM with a γ -linear profile are similar and close to those of the empirical estimator, and better than those of RM. A constant and low γ (equals to 0.6) gives even better results with ARM (but not with RM) for both distribution cases (normal and uniform). In fact, the averaging in ARM (which makes the quantile estimator converges more quickly) makes it unnecessary to increase the γ towards 1 that we have with the linear profile.

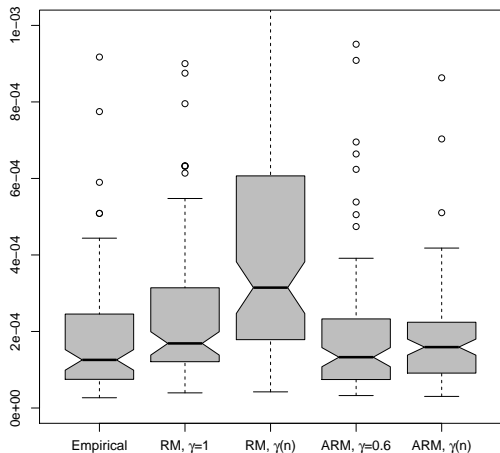
Other tests with different distributions, not shown here, present similar conclusions. They also confirm that it is necessary to keep the adaptive tuning of C .

4.2 Application of the Kesten's rule

Another way of improvement of the RM algorithm would be to apply a stochastic adaptive stepsize rule (instead of a moving γ). The old and simple Kesten's rule [11, 16] allows to do so by incrementing n (in the stepsize $1/n$) only if the two last errors have different signs. It is based on the idea that "if we are far from the optimal, the errors tend to all have the same sign; as we get close, the errors tend



(a) Normal distribution.



(b) Uniform distribution.

Figure 4: Mean square errors of discrete quantile functions for the empirical, RM and ARM estimators. $\gamma(n)$ corresponds to the γ -linear profile.

to alternate”. Then, it slows down the stepsize reduction at the beginning of the algorithm.

The RM estimator using the Kesten’s rule (noted KRM) writes

$$q_\alpha(n+1) = q_\alpha(n) - \frac{C}{k_n^\gamma} (\mathbf{1}_{Y_{n+1} \leq q_\alpha(n)} - \alpha), \quad (8)$$

with $n \geq 1$, $q_\alpha(1) = Y_1$ and

$$\begin{aligned} \text{if } n > 2, \quad k_n &= k_{n-1} + \mathbf{1}_{\delta(n)\delta(n-1) < 0}, \\ \text{if } n \leq 2, \quad k_n &= n, \end{aligned} \quad (9)$$

with $\delta(n) = q_\alpha(n) - q_\alpha(n-1)$. The required memory storage for the KRM estimator is four values (one realization of Y and three quantile estimates). If ARM (see section 4.1) is used instead of RM, the estimator is noted KARM and requires a memory storage of five values (one realization of Y and four quantile estimates).

Our numerical test considers the case $Y \sim \mathcal{N}(0, 1)$, with $N = 1000$, $\gamma = 1$, $C = 1$ and three orders of quantile α (0.05, 0.5 et 0.95). Figure 5 shows 50 independent trajectories of the RM estimator $q_\alpha(n)$ for $n = 1, \dots, N$ by applying Kesten’s rule or not. Results clearly show that the Kesten’s rule applies the *search-then-converge learning rule* [16] which led us to consider a γ -linear profile in Section 3.1. In this case, the KARM estimates provide similar result than the KRM (results not shown here).

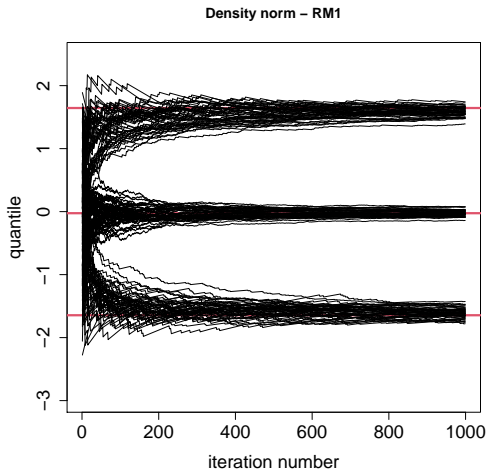
Other tests with different distributions for Y (uniform, triangular, exponential, lognormal, multimodal, etc.) have been performed. The main conclusion is that KARM (Kesten’s rule applied on the averaged RM estimator), associated to $\gamma = 1$ and the adaptive tuning of C , is the most robust approach.

5 Conclusion

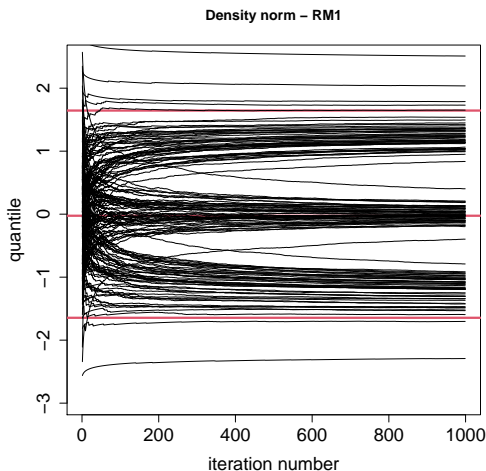
This paper has proposed the computation of quantiles by use of a parallel one-pass strategy based on new robust versions of the stochastic quantile algorithm of Robbins-Monro. This kind of iterative approach has been demonstrated in the present paper by studying a scalar output, but it is needed in case of huge-volume simulation output, as a spatio-temporal evolution of several quantities (as studied in [22, 18]). Indeed, in this situation, it allows to avoid the storage of large amount of data during the uncertainty propagation stage.

More precisely, our work made it possible to define some heuristics for the iterative estimation of quantile by the RM algorithm (eq. (4)) with a finite-size sample N (a few hundreds values). Several versions of the RM algorithm have been studied and led to some conclusions:

- The choice of an adaptive C (Eq. (6)) is beneficial in all cases;



(a) With Kesten's rule.



(b) Without Kesten's rule.

Figure 5: Simulations of trajectories of the RM algorithm ($N = 1000$, $Y \sim \mathcal{N}(0, 1)$). Red lines give exact order quantiles 0.05, 0.5 and 0.95.

- When using RM, the choice of a γ -linear profile (Eq. (5)) is robust and must be privileged;
- On the other hand, the ARM algorithm (average version of RM) gives good results while avoiding to define N at the beginning of the study. With ARM, weak γ gives better results;
- Finally, the introduction of the stochastic adaptive stepsize Kesten's rule, which led to define the KARM estimator, further improves our results. In this case, $\gamma = 1$ has to be chosen.

This paper has presented preliminary works which need to be further investigated. For example, the use of well distributed point sequences instead of i.i.d. samples (tests not shown here) allow to greatly improve the precision of the different RM estimators. This idea seems judicious and will be studied in depth in the case where the variable Y comes from a model (as in UQ) whose input dimension is large and where the choice of a good design of experiment (of the “space filling design” type) is important [17].

In the same order of idea, it will be fruitful to combine the RM algorithm and the techniques of simulation of rare events [13, 27], as already studied by [12]. Another major perspective of our work will be to control the convergence of the estimate when N is not fixed, which is an essential issue in industrial applications.

Last, in [18], the quantiles of the spatio-temporal outputs have been computed cell per cell and time-step per time-step via an iterative UQ approach, which is perfected in this paper. [18] have shown that the interpretation of this *ubiquitous* quantiles (for instance in the form of static spatial maps, temporal probes or videos) is much richer than the traditional predefined probe-based or sub-sampled approaches. However, the functional space where the spatio-temporal field lies has not been considered. Dealing with this space in a rigorous mathematical way (as in [4]), the ubiquitous quantile estimates would conserve the geometrical and temporal structure of the ensemble run study, and would then keep, more deeply, a physical significance.

Acknowledgments

Part of this work has been funded by the international ANR project INDEX (ANR-18-CE91-0007) devoted to researches on incremental design of experiments. The author thanks Luc Pronzato, Bernard Bercu, Alejandro Ribés and Clément Gauchy for fruitful discussions.

References

- [1] M. Baudin, A. Dutfoy, B. Iooss, and A-L. Popelin. Open TURNS: An industrial software for uncertainty quantification in simulation. In R. Ghanem, D. Higdon, and H. Owhadi, editors, *Springer Handbook on Uncertainty Quantification*, pages 2001–2038. Springer, 2017.

- [2] Janine C Bennett, Vaidyanathan Krishnamoorthy, Shusen Liu, Ray W Grout, Evatt R Hawkes, Jacqueline H Chen, Jason Shepherd, Valerio Pascucci, and Peer-Timo Bremer. Feature-based statistical analysis of combustion simulation data. *IEEE transactions on visualization and computer graphics*, 17(12):1822–1831, 2011.
- [3] E. Borgonovo. *Sensitivity Analysis - An Introduction for the Management Scientist*. Springer, 2017.
- [4] H. Cardot, P. Cénac, and P-A. Zitt. Efficient and fast estimation of the geometric median in hilbert spaces with an averaged stochastic gradient algorithm. *Bernoulli*, 19:18–43, 2013.
- [5] Tony F Chan, Gene Howard Golub, and Randall J LeVeque. Updating formulae and a pairwise algorithm for computing sample variances. In *COMPSTAT 1982 5th Symposium held at Toulouse 1982*, pages 30–41. Springer, 1982.
- [6] E. de Rocquigny, N. Devictor, and S. Tarrantola, editors. *Uncertainty in industrial practice*. Wiley, 2008.
- [7] M. Duflo. *Random iterative models*. Springer, Berlin, 1997.
- [8] R. Ghanem, D. Higdon, and H. Owhadi, editors. *Springer Handbook on Uncertainty Quantification*. Springer, 2017.
- [9] L. Gilquin, E. Arnaud, C. Prieur, and H. Monod. Recursive estimation procedure of Sobol’ indices based on replicated designs. *Preprint*, <http://hal.univ-grenoble-alpes.fr/hal-01291769>, 2017.
- [10] B. Iooss and A. Marrel. Advanced methodology for uncertainty propagation in computer experiments with large number of inputs. *Nuclear Technology*, 205:1588–1606, 2019.
- [11] H. Kesten. Accelerated stochastic approximation. *Ann. Math. Stat.*, 29:41–59, 1958.
- [12] M. Kohler, A. Krzyżak, and H. Walk. Nonparametric recursive quantile estimation. *Statistics and Probability Letters*, 93:102–107, 2014.
- [13] J. Morio and M. Balesdent. *Estimation of rare event probabilities in complex aerospace and other systems*. Woodhead Publishing, 2016.
- [14] Philippe Pébay. Formulas for robust, one-pass parallel computation of covariances and arbitrary-order statistical moments. *Sandia Report SAND2008-6212*, Sandia National Laboratories, 94, 2008.
- [15] B.T. Polyak and A.B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30:838–855, 1992.
- [16] W.B. Powell. *Approximate dynamic programming, Second edition*. Wiley, 2011.
- [17] L. Pronzato and W. Müller. Design of computer experiments: space filling and beyond. *Statistics and Computing*, 22:681–701, 2012.
- [18] A. Ribés, T. Terraz, Y. Fournier, B. Iooss, and B. Raffin. The many advantages of avoiding writing files for large scale uncertainty quantification in numerical simulation. In *In situ visualization for computational science*, *In press*, 2020.
- [19] H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22:400–407, 1951.
- [20] R.C. Smith. *Uncertainty quantification*. SIAM, 2014.
- [21] I.M. Sobol’. Sensitivity estimates for non linear mathematical models. *Mathematical Modelling and Computational Experiments*, 1:407–414, 1993.
- [22] T. Terraz, A. Ribes, Y. Fournier, B. Iooss, and B. Raffin. Large scale in transit global sensitivity analysis avoiding intermediate files. In *Proceedings the International Conference for High Performance Computing, Networking, Storage and Analysis (Supercomputing)*, Denver, USA, november 2017.
- [23] L. Tierney. A space-efficient recursive procedure for estimating a quantile of an unknown distribution. *SIAM Journal on Scientific and Statistical Computing*, 4:706–711, 1983.
- [24] C. Wang, M-H. Chen, E. Schifano, J. Wu, and J. Yan. Statistical methods and computing for big data. *Stat. Interface*, 9:399–414, 2016.

- [25] Pengfei Wei, Chenghu Tang, and Yuting Yang. Structural reliability and reliability sensitivity analysis of extremely rare failure events by combining sampling and surrogate model methods. *Proceedings of the Institution of Mechanical Engineers, Part O: Journal of Risk and Reliability*, *In Press*, 2019.
- [26] BP Welford. Note on a method for calculating corrected sums of squares and products. *Technometrics*, 4(3):419–420, 1962.
- [27] E. Zio and N. Pedroni. An adaptive metamodel-based subset importance sampling approach for the assessment of the functional failure probability of a thermal-hydraulic passive system. *Applied Mathematical Modelling*, 48:269–288, 2017.