FOREWORD: SPECIAL ISSUE ON MULTILEVEL-MULTIFIDELITY APPROACHES FOR UNCERTAINTY QUANTIFICATION

M.S. Eldred,^{1,} G. Geraci,¹ & G. Iaccarino²*

¹*Sandia National Laboratories, Albuquerque, New Mexico, USA, 87185*

²*Stanford University, Stanford, California, USA, 94305*

**Address all correspondence to: M.S. Eldred, Sandia National Laboratories, P.O. Box 5800, Albuquerque, New Mexico, USA, 87185, E-mail: mseldre@sandia.gov*

Original Manuscript Submitted: 11/2/2020; Final Draft Received: 11/18/2020

This three-part special issue is focused on recent developments in multilevel/multifidelity algorithms and their effective demonstration in real-world applications, spanning both forward and inverse UQ and related optimization contexts.

KEY WORDS: *uncertainty quantification, multifidelity, multilevel*

1. INTRODUCTION

Computational simulation continues to advance in its predictive capability through the development of high-fidelity multi-physics/multi-scale simulation models, with unprecedented resolution enabled by the latest high-performance computers. Uncertainty quantification (UQ) methodologies are challenged in this environment, both by the prohibitive cost of computing high-fidelity ensembles and by the increasing number of uncertainty sources that is often induced by this model complexity. To address these challenges, researchers are effectively harnessing the utility that exists within hierarchies of model forms (multifidelity) and resolution levels (multilevel and multi-index) in order to control multiple sources of error while efficiently allocating simulation resources. By relaxing the need for exclusive reliance on the most expensive models, high-fidelity UQ studies become tractable. This special issue is focused on recent developments in multilevel/multifidelity (ML/MF) algorithms and their effective demonstration in real-world applications, spanning both forward and inverse UQ and related optimization contexts.

We first overview the contributions within this three-part special issue, including work in ML/MF algorithms that spans forward UQ (Section 2.1), inverse UQ (Section 2.2), and optimization (Section 2.3), and including realistic deployments of ML/MF methods that demonstrate their utility in a variety of areas of computational science (Section 3). Each paper is annotated as part I, II, or III, corresponding to the printed issues in which they appear. We conclude with a set of summary observations in Section 4.

2. ML/MF ALGORITHMS

2.1 Forward UQ

Forward UQ is the process of propagating input uncertainties to estimate statistics on output quantities of interest (QoIs). This typically involves an ensemble of simulation evaluations drawn from the random parameter space, and in the case of ML/MF approaches, this ensemble is expanded to fuse simulation data from multiple model forms and/or resolutions.

- In Adcock et al. (Part II), the standard multilevel Monte Carlo (MLMC) estimator is extended to modern, heterogeneous computer architectures. The optimal sample allocation strategy for this novel MLMC estimator is derived by solving a convex optimization problem and the samples are distributed on both levels and processing units (GPUs and CPUs). A numerical test is also provided for a heated channel flow with stochastic forcing

for which the authors demonstrate a cost reduction of more than 80% with respect to a corresponding MLMC where samples are distributed on GPUs only.

- In Badia et al. (Part III), the authors derive a MLMC estimator for problems with random geometry and topology. The proposed approach streamlines the generation of the mesh for each realization using embedded methods. This simplifies creation of a mesh hierarchy and avoids robustness issues, relative to traditional body-fitted unstructured mesh generation. The authors consider their recently introduced aggregated finite element method as a starting point for the definition of a novel embedded MLMC estimator and demonstrate that their method performs consistently, in terms of error, variance decay, and complexity, relative to standard MLMC using deterministic geometries and body-fitted meshes. The authors consider three test problems ranging from the solution of a Poisson equation in a simple domain to a heat equation problem with uncertain topology and geometry.
- In Barth (Part III), a new MLMC estimator is derived for the computation of expectation. This new estimator is designed to remove the approximation error bias of the standard MLMC and it can be applied in all situations in which a sequence of model approximations have a rate-dependent and decreasing approximation error that can be extrapolated to the zero error limit. The new MLMC estimator is observed to exhibit a factor of 2 cost reduction with respect to the standard MLMC. Moreover, the reuse of samples across levels is also proposed for both the standard and newly derived MLMC versions.
- In De et al. (Part II), multifidelity concepts are applied in the rapidly expanding area of machine learning via deep neural networks. Adopting a bifidelity transfer learning approach, the authors demonstrate improved predictive accuracy using a teacher-student formulation, incorporating a high-fidelity Gaussian process (GP) model as the teacher, relative to a number of other candidate architectures and reference benchmarks. Multifidelity performance is demonstrated for cantilever beam, lithium-ion battery, and NACA airfoil applications.
- In Gorodetsky et al. (Part II), the authors introduce a general Bayesian multifidelity UQ framework (MFNets) that provides flexibility in handling different data sources (e.g., nonhierarchical models) and in combining models with disparate parameterizations, while also retaining robustness in managing corrupted data. Numerical examples demonstrate the accuracy and efficiency of MFNets compared to existing strategies such as multilevel and control variate (CV) estimators. Further, the MFNets theoretical framework enables analysis of different approaches, and the authors demonstrate that the Monte Carlo (MC) CV estimator can be derived entirely from the use of Bayes rule and linear-Gaussian models.
- In Jornet et al. (Part I), the authors develop a strategy to estimate the full probability density function of solutions of linear random differential equations using MLMC. After interpreting the density as an expectation, their goal is to accelerate the convergence of the expectation estimators while reducing the effect of noise—i.e., numerical error—in the process. The authors compare and contrast different approaches for variance reduction, including a combination of MC sampling and quadrature rules and a MLMC estimator applied to the series expansions of the solutions. The results are consistent with theoretical expectations and several demonstrations show considerable speedup over classical multilevel sampling.

2.2 Inverse UQ

Inverse UQ is the process of inferring a consistent characterization of the random input uncertainties based on the available reference data for observable QoI. When taking a Bayesian perspective, for example, this typically involves computing a Markov chain to form observations of the posterior distribution of the parameters, given the available prior information and the QoI data (embedded within a likelihood function). In an ML/MF setting, we may leverage different model fidelities to accelerate the Markov chain Monte Carlo (MCMC) evolution and compute the posterior distribution at reduced cost.

- In Bruder et al. (Part I), a measure-theoretic formulation for stochastic inversion is described, for which a central need is the efficient construction of a forward mapping from random input variables to high-fidelity

response QoIs. This motivates the use of a multifidelity approach that employs a GP model to approximate the conditional densities between model fidelities, enabling a computationally efficient MC-based approximation of the forward propagation of an initial density through the model hierarchy. The methodology is demonstrated using examples from fluid flow in porous media and solid mechanics.

- In Catanach et al. (Part II), the authors consider an alternative approach to traditional Bayesian inference for problems involving expensive evaluations of the chemical master equation. They develop a multifidelity extension of the sequential tempered MCMC sampler which replaces the overall inference problem with a sequence of subproblems that gradually increase model fidelity and the influence of the observed data. These problems can be solved efficiently and are designed using an information-theoretic criterion that seeks to maximize the impact of each model in the hierarchy on the final Bayesian posterior. Computational savings are demonstrated for applications involving models of gene expression and other biologically motivated problems.
- In Narayan et al. (Part III), the authors investigate the use of low-rank multifidelity approaches by considering nonlinear operators when identifying the low-rank structure of the solutions. Specifically, they consider general kernel structures and select the one that minimizes the resulting solution rank; the initial selection is based on a library of kernel functions that is further expanded using a convex mixture of computing models or data-driven optimization. Applications spanning diverse disciplines demonstrate the efficacy of the approach compared to methods based on linear low-rank decompositions.

2.3 Optimization

An important consumer of both forward and inverse UQ is optimization, whether for design optimization under uncertainty (repeated usage of forward UQ for computing design metrics) or for optimal experimental design (repeated usage of inverse UQ for updating posterior distributions given new data). This additional level of iteration provides a rich set of options for ML/MF acceleration.

- In Sacher et al. (Part III), a multifidelity global optimization approach is developed, centered around adaptive enrichment of GP models within the context of the efficient global optimization (EGO) algorithm. Here, a recursive/autoregressive approach is employed for the multifidelity GP modeling, and the enrichment strategy is constructed to select the best individual fidelity samples using a prediction variance metric for the multifidelity surrogate. The methodology relaxes common sample nesting requirements and is demonstrated using an application in nonlinear fluid-structure interaction.

3. DEPLOYMENTS AT SCALE

ML/MF algorithms are now enjoying broad deployment and community adoption within a number of areas of computational science. These experiences are providing a critical proving ground for new methods and are providing essential feedback on R&D needs that could enhance impact in realistic settings. The following papers, while not restricted to such, have a strong application deployment component.

- In Geraci et al. (Part III), the authors consider the novel application of multifidelity strategies to applications in computer network security. Specifically the authors study different target models and construct low-fidelity counterparts to demonstrate consistent acceleration compared to single-fidelity MC. The authors further explore the possibility of forming near-optimal LF models by minimizing the overall cost of the multifidelity estimator over the space of LF model parameterizations. This is one of the first rigorous applications of UQ in this emerging field.
- In Jofre et al. (Part II), the authors consider a large-scale application in irradiated particle-laden turbulent flow and investigate the performance of a collection of sampling-based approaches, including multilevel, multifidelity, and multilevel-multifidelity Monte Carlo (MLMC, MFMC, and MLMF MC). They target a key challenge with the deployment of ML/MF strategies to complex multiphysics applications, that of the selection

of an effective low-fidelity (LF) model ensemble in the case where a much larger set of possibilities exists, e.g., arising from combining fidelity definitions for each sub-physics component. Based on a pilot study over 31 LF models, the authors first optimize the sequence of models within MLMC, demonstrating approximately two orders of magnitude improvement with respect to MC. Next, the CV component within MLMF MC is optimized through LF model selection, demonstrating an additional order of magnitude speedup relative to its corresponding MLMC.

- In Mukhopadhyaya et al. (Part I), multifidelity concepts are applied to generate probabilistic aerodynamic databases for a full configuration aircraft, the NASA Common Research Model. An autoregressive multifidelity Gaussian process approach is used to recursively emulate performance data across multiple model forms while incorporating observation noise. This noise comprises the model form uncertainty associated with the different fidelities and is generated either using input from subject-matter experts or using a Reynolds-averaged Navier-Stokes eigenspace perturbation methodology. These multifidelity probabilistic databases are intended to quantify the risk of not meeting a certification requirement during the development of new aircraft.
- In Seo et al. (Part I), the authors investigate the performance of multifidelity control variate estimators for coronary circulation models. Uncertain random inputs, statistically characterized by clinical data collection, are propagated through ensembles of 3D, 1D, and 0D computational models to obtain statistics of clinically relevant model outputs. A three-dimensional anatomic model of the left coronary artery, with geometry obtained directly from clinical tomography images, is considered for a realistic numerical demonstration. The authors demonstrate that use of the multifidelity estimator, via pilot sample estimation and performance extrapolation, results in a cost reduction of one or two orders of magnitude compared to a single-fidelity estimator.

4. SUMMARY AND OUTLOOK

In aggregate, this special issue demonstrates that multilevel/multifidelity research, development, and application is an extremely active area, with emerging success stories in a variety of areas of computational science, as well as an evolving set of R&D needs for enhancing the effectiveness of these approaches in expanding contexts.

While an extensive survey of emerging research areas is outside the scope of this foreword, the guest editors would highlight the following as important directions for ongoing and future research, many of which are evident within the collection of work in this special issue:

- *Generalizations for complex model ensembles:* This direction encompasses multiple interrelated issues:
 - Expand algorithm support *beyond classical hierarchical assumptions*. While recursive pairing is very effective for highly correlated cases (e.g., discretization of elliptic PDEs), nonidealized cases benefit from more general graph relationships.
 - Support *expansion in the composition of the model ensemble*, e.g., the progression from multilevel to multi-index approaches and the incorporation of multiphysics/multiscale relationships.
 - Diverse ensembles may require identification of shared input-output manifolds among models to *bridge dissimilar input parameterizations* and enable *correlated mappings for the underlying physical processes* that define key response QoI.
 - Given a rich set of modeling alternatives, potentially including hyperparameter resolution controls, integrate *model tuning* to maximize predictive value per unit cost for a single model as well as *ensemble selection* to improve the aggregate ML/MF performance, for example, by identifying well-correlated model pairs.
- *Goal orientation:* ML/MF strategies are being expanded beyond expected value metrics to target more general statistical analysis goals, such as global sensitivity (Sobol' indices), robustness (variance), and reliability (failure probability) metrics. These extensions can be challenging (e.g., dissimilarity in failure regions), but are essential for enabling their use in general contexts such as design optimization under uncertainty.

- *Comprehensive error management:* An essential concept in ML/MF methods is effectively balancing multiple sources of error, most commonly including deterministic resolution bias and stochastic estimator variance. With the generalization of ML/MF strategies into new areas, there is an emerging need to encompass surrogate modeling errors (e.g., prediction error estimates from GPs, reduced-order models, Bayesian neural networks, et al.), online simulation error estimation (e.g., variational multiscale), and stochastic simulation processes (e.g., direct simulation Monte Carlo).
- *Embrace heterogeneity:* Given expanding requirements for diverse model ensembles with widely varying computing needs, there is an emerging need for tools and processes that support high levels of heterogeneity and an opportunity to marry the adaptive management of diverse simulation ensembles with techniques that can extract concurrency from modern peta and exascale computer architectures.

ACKNOWLEDGMENTS

The authors extend their thanks to Begell House and the IJUQ editorial leadership for sponsoring this collection, and to all of the authors for their diligent attention to the work herein. Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.