Aerospace Science and Technology 70 (2017) 121–136

Contents lists available at ScienceDirect

Aerospace Science and Technology

www.elsevier.com/locate/aescte

All-at-once approach to multifidelity polynomial chaos expansion surrogate modeling

Dean E. Bryson *, Markus P. Rumpfkeil

Dept. of Mechanical and Aerospace Engineering, University of Dayton, OH, 45469, USA

ARTICLE INFO

Article history:
Received 3 August 2016
Received in revised form 6 July 2017
Accepted 29 July 2017
Available online 4 August 2017

Keywords:
Multifidelity
Multi-fidelity
Variable fidelity
Polynomial chaos
Surrogate modeling

ABSTRACT

A new approach to multifidelity, gradient-enhanced surrogate modeling using polynomial chaos expansions is presented. This approach seeks complementary additive and multiplicative corrections to low-fidelity data whereas current hybrid methods in the literature attempt to balance individually calculated calibrations. An advantage of the new approach is that least squares-optimal coefficients for both corrections and the model of interest are determined simultaneously using the high-fidelity data directly in the final surrogate. The proposed technique is compared to the weighted approach for three analytic functions and the numerical simulation of a vehicle's lift coefficient using Cartesian Euler CFD and panel aerodynamics. Investigation of the individual correction terms indicates the advantage of the proposed approach is that complementary calibrations separately adjust the low-fidelity data in local regions based on agreement or disagreement between the two fidelities. In cases where polynomials are suitable approximations to the true function, the new all-at-once approach is found to reduce error in the surrogate faster than the method of weighted combinations. When the low-fidelity is a good approximation of the true function, the proposed technique out-performs monofidelity approximations as well. Sparse grid constructions alleviate the growth of the training set as root-mean-square-error is calculated for increasingly higher polynomial orders. Utilizing gradient information provides an advantage at lower training grid levels for low-dimensional spaces, but worsens numerical conditioning of the system in higher dimensions.

Published by Elsevier Masson SAS.

1. Introduction

Modeling fidelity is the degree to which a method of prediction, analysis, or experimentation accurately reproduces an effect of interest. Models of varying fidelity levels are ingrained in the aerospace design process from conceptual design through detailed design to test and evaluation. A pervasive challenge is managing the balance between highly-accurate predictions and cheaper, more expedient methods. As George Box famously stated [1], "Essentially, all models are wrong, but some are useful." Unfortunately, the boundaries of usefulness may even be unknown as aircraft designs deviate further and further from the body of historical data.

Predictions of different fidelities may be combined as an ensemble of data without preference [2,3], or more commonly used in a hierarchy in which one dataset is believed to be more trustworthy (termed high-fidelity) than another (low-fidelity). These hierarchical methods seek to predict response trends using intensely-sampled low-fidelity data and reserve high-fidelity evaluations to correct inaccuracies in the low-fidelity data. Possible examples of high-fidelity versus low-fidelity data sources for aerodynamic predictions are provided in Table 1.

Table 1
Some possible combinations of variable-fidelity methods.

High-fidelity model	Low-fidelity model
Experimental data	CFD results
Navier–Stokes	Euler
Euler	Panel methods
Finer mesh CFD results	Coarser mesh CFD results
Fully converged solutions	Partly converged solutions
Detailed geometry	Simplified geometry

A variety of examples combining predictions from multiple sources may be found in the literature, with significant cross-over between the design optimization and uncertainty quantification communities. Early works tend to focus on polynomial interpolation and regression and Taylor series expansions, while more recent research is split primarily between polynomial and krig-

* Corresponding author.
E-mail addresses: brysond1@udayton.edu (D.E. Bryson), markus.rumpfkeil@udayton.edu (M.P. Rumpfkeil).

http://dx.doi.org/10.1016/j.ast.2017.07.043
1270-9638/Published by Elsevier Masson SAS.

ing models, though other surrogate model forms (e.g., radial basis functions and support vector regression) are also in use [4]. Users of polynomial models often leverage the simplicity of its form and fast prediction of statistical moments based on the readily-available coefficients, whereas users of kriging take advantage of its nonlinear modeling capabilities and estimates of model variance and error. Corrections to the low-fidelity response may take the form of input space mappings [5], or, more commonly, output space mappings (commonly referred to as bridge, correction, or calibration functions). These bridge functions typically take the form of additive or multiplicative corrections, or a combination of the two. However, it is difficult to know *a priori* which form to use for a particular fidelity pairing and in what order to apply the corrections.

The approach taken in generating multifidelity models typically varies with application. Gradient-based optimization typically relies on local corrections of the lower-fidelity model to match the response value and gradient of the high-fidelity model, forming the basis of Trust Region Model Management [6] and Approximation Management Framework [7], which provably converge to optima of the high-fidelity model. The Global–Local Approximation technique utilizes analytic structural gradients to produce linear multiplicative scaling functions local to a design point, and was introduced by Haftka [8] for different model discretizations and extended by Chang et al. [9] to models with varying governing equations. Similarly, Lewis and Nash [10] implemented additive corrections in their multigrid optimization technique. Eldred, Giunta, and Collis [11] developed second-order additive, multiplicative, and hybrid corrections to improve the local rate of optimization convergence using Hessian information approximated by finite-differences, BFGS, and Symmetric Rank 1 updates.

To extend the region of accuracy of local corrections, nearby data may also be included. Eldred, Giunta, and Collis [11] take a multipoint approach in generating hybrid corrections, satisfying the first-order consistency condition at both the current design and the prior design point. However, they note the multipoint technique is marginally effective as it is backward looking rather than forward looking in terms of the optimization search direction. Similarly, Gano, Renaud, and Sanders [12] select the nearest data as the secondary point. An alternative multipoint approach is taken by Fischer and Grandhi [13], in which a Bayesian update is utilized to estimate the best blend of additive and multiplicative corrections. As a forward-looking approach, Bryson, Rumpfkeil, and Durscher [14,15] present a multifidelity optimization approach that leverages the concept of expected optimal point from quasi-Newton optimization, making use of the polynomial models presented in this article. Incorporating information from more than two points, Fischer and Grandhi [13] perform a Bayesian update of the weights based on all data near the current design.

Global optimization and uncertainty quantification depend on global corrections to the lower-fidelity model. There are many examples of using kriging to model the discrepancies between different models globally or over large regions, often including a low-order multiplicative scaling across the design space [12,13,16–19]. Notably, Han, Görtz, and Zimmermann [17] combined gradient enhanced kriging with a generalized hybrid bridge function to take advantage of both multiple data sources and efficiently calculated gradients. Adopting terminology from Peherstorfer, Willcox, and Gunzburger [20], these multi-step corrections are considered model adaptation as opposed to model fusion where a single high-fidelity surrogate is derived from multiple data sources, such as in co-kriging. In Kennedy and O'Hagan [21], the different models are internally linked by a constant scaling factor determined as part of the co-kriging process in addition to the discrepancy function. In Qian and Wu [22] and Han, Zimmerman, and Görtz [23], the scaling and discrepancy are both generalized to be Gaussian processes.

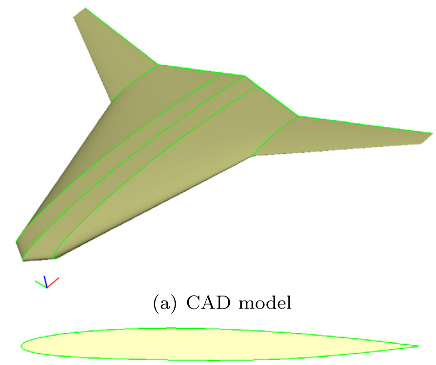

(a) CAD model

(b) Root airfoil

Fig. 1. Geometry of lambda wing vehicle.

Global multifidelity corrections to polynomial models are less common, and are frequently created in a multi-step adaptation. Ng and Eldred [24] present a method where separate polynomial expansions are developed for the low-fidelity data as well as additive and multiplicative corrections. The three surrogates are then combined in a convex combination to approximate the high-fidelity response, selecting a weighting parameter to minimize the magnitude of the correction. Alternatively, Shah et al. [25] determine the least squares-optimal linear multiplicative and constant additive correction to the low-fidelity data. Then, a single surrogate of the high-fidelity response is created using only the transformed low-fidelity data. Similarly, Palar, Tsuchiya, and Parks [26] use a regression procedure to find the least squares-optimal additive correction between fidelities, and later combine PCEs of the low-fidelity function and corrective term.

In this article, we present an all-at-once approach to generating multifidelity PCE's with hybrid additive-multiplicative bridge functions of arbitrary order. The technique seeks to generate a surrogate of the high-fidelity model by fitting both high-fidelity data and low-fidelity data corrected on a global scale. The fitting process may be considered more like a fusion process than an adaptation in that all the data participate in a single model fitting including determination of least-squares optimal additive and multiplicative corrections. The resulting corrections are complementary, eliminating the weighting parameter that must be determined in other approaches in the literature. The multifidelity approach is further augmented by the addition of gradients for both the high and low fidelities following the approaches of Li et al. [27]; Hessians may also be included following Boopathy and Rumpfkeil [28]. Demonstrations are made using analytic test functions ranging from two to one hundred dimensions, and on a three-dimensional test problem combining Cartesian Euler Computational Fluid Dynamics (CFD) with a panel aerodynamics method.

The aerodynamic application is the prediction of the lift coefficient for a tailless, lambda-wing vehicle illustrated in Fig. 1, which has been the subject of many studies at the Air Force Research Laboratory [29–32]. The lambda planform is selected to balance subsonic and supersonic performance. The outboard wing section increases subsonic performance by increasing aspect ratio, while the inboard section provides high speed performance. Flying a tailless vehicle, however, presents challenges in trim and stability, and accurate prediction of the aerodynamic forces and moments is critical. The accurate prediction of aerodynamic performance is further exacerbated by the fact that such a vehicle must both dwell in and fly through the transonic regime, for which modeling is known to be difficult. Many design decisions driving or constraining performance under these conditions are made in the conceptual- and preliminary design stages based on lower-fidelity tools. High-fidelity aerodynamic simulation at every point in the