Chinese Journal of Aeronautics, 2013,26(4): 841–849

Chinese Society of Aeronautics and Astronautics & Beihang University
Chinese Journal of Aeronautics
cja@buaa.edu.cn
www.sciencedirect.com

Variable-fidelity optimization with design space reduction

Mohammad Kashif Zahir, Gao Zhenghong *

School of Aeronautics, Northwestern Polytechnical University, Xi'an 710072, China

Received 10 May 2012; revised 19 June 2012; accepted 19 October 2012
Available online 2 July 2013

KEYWORDS
Airfoil optimization;
Curse of dimensionality;
Design space reduction;
Genetic algorithms;
Kriging;
Surrogate models;
Surrogate update strategies;
Variable fidelity

Abstract Advanced engineering systems, like aircraft, are defined by tens or even hundreds of design variables. Building an accurate surrogate model for use in such high-dimensional optimization problems is a difficult task owing to the curse of dimensionality. This paper presents a new algorithm to reduce the size of a design space to a smaller region of interest allowing a more accurate surrogate model to be generated. The framework requires a set of models of different physical or numerical fidelities. The low-fidelity (LF) model provides physics-based approximation of the high-fidelity (HF) model at a fraction of the computational cost. It is also instrumental in identifying the small region of interest in the design space that encloses the high-fidelity optimum. A surrogate model is then constructed to match the low-fidelity model to the high-fidelity model in the identified region of interest. The optimization process is managed by an update strategy to prevent convergence to false optima. The algorithm is applied on mathematical problems and a two-dimensional aerodynamic shape optimization problem in a variable-fidelity context. Results obtained are in excellent agreement with high-fidelity results, even with lower-fidelity flow solvers, while showing up to 39% time savings.

1. Introduction

With advances in computational fluid dynamics (CFD) and computer hardware, CFD has now become an integral part of the aircraft design process. The high-fidelity (HF) aerodynamic data it provides has contributed to cutting aerodynamic design cost and time scales by reducing the number of required wind tunnel tests.[1] However, major benefits can be achieved if CFD is included in the conceptual design phase, where as many as tens of thousands of analyses must be performed, and global optimization can play a key role. Given the high cost of CFD and optimization, a prominent area of research today is to find ways to reduce the computational time while retaining the high fidelity of the analysis. In the area of aerodynamic optimization, the variable-fidelity (VF) (also called multi-fidelity) method has quickly grown in popularity.[2–22]

Variable-fidelity and other model management methods have been developed to solve optimization problems that involve simulations with high computational expense.[9,11] In many engineering design problems, differing levels of fidelity can model the system of interest. Higher-fidelity models typically incorporate more detailed physics and are computationally expensive to evaluate than lower-fidelity (LF) models. Lower-fidelity models are typically much cheaper to evaluate,

* Corresponding author. Tel.: +86 29 88495971.
E-mail addresses: mkz2a770@gmail.com (M.K. Zahir), zgao@nwpu.edu.cn (Z. Gao).
Peer review under responsibility of Editorial Committee of CJA.

Production and hosting by Elsevier

1000-9361 © 2013 Production and hosting by Elsevier Ltd. on behalf of CSAA & BUAA. Open access under CC BY-NC-ND license.
http://dx.doi.org/10.1016/j.cja.2013.06.002

but designs produced by using these models neglect important physical effects included in more expensive higher-fidelity models. In aircraft design, the Navier–Stokes and Euler equations are examples of two computational models with different fidelities, where the latter is obtained by removing the viscosity terms from the Navier–Stokes equations.

Variable-fidelity optimization (VFO) has emerged as an attractive method of performing both high-speed and high-fidelity optimization and the past three decades have seen rapid increase in its development and usage.[2–23] These algorithms attempt to leverage information from computationally inexpensive low-fidelity models to reduce the time required to converge to the optimum of high-fidelity functions. This is usually accomplished by using a low-fidelity solver plus a correction term – the difference between high-fidelity and low-fidelity solvers – modeled by a surrogate model calibrated at selected sample points in the design space.[16] A variety of methods have been used for generating these surrogate models including Kriging, [4,7–9,11,14,15,18,24–30] radial basis functions (RBFs),[28,31] neural networks,[4,9,15,31,32] and support vector regression (SVR).[33,34] Insightful reviews of surrogate models and variable-fidelity methods have appeared in the literature.[11,19,23,24,35]

Advanced engineering systems, like aircraft, are defined by tens or even hundreds of design variables. Building an accurate surrogate model for use in such high-dimensional optimization problems is a difficult task. In essence, a surrogate model is a data-fit and is only accurate in the region where it is adequately trained. Intelligent techniques of generating sampling plans (also called design of experiments or DoE) – a sparse set of points where the surrogate will be trained – exist to achieve uniform coverage of the design space. However, if the problem being dealt with has many dimensions, the number of training points required for reasonable uniform coverage of the design space rises exponentially – the so-called curse of dimensionality.[23,36] A surrogate-based optimizer may converge to a local optimum, or worse a false optimum, due to inaccuracies of the surrogate model.[25,37]

While it is easy to control the range of validity of the surrogate model in gradient-based optimization algorithms by using ad hoc move limits or a trust-region framework, it is not straightforward in global optimization schemes like genetic algorithms (GAs).[38] This issue has been addressed by other researchers in the past. Ratle[39] uses a heuristic convergence criterion to determine when the approximate model must be updated. The basic idea is that the convergence of the search process should be stable and therefore, the change of the best solution should not be larger than a user-defined value. This, however, relies on the assumption that the first sets of data points are weakly correlated with the global optimum of the original problem, which is not necessarily true for high-dimensional systems. Others perform on-line learning of the approximate model based on a prescribed generation delay.[2,40] Another concept of evolution control is applied by Jin et al.[37] where the surrogate model and the original fitness function are both used in tandem during the evolutionary process based on a fixed[32] or adaptive[17,37,41] criterion.

One way to solve this problem is to limit the range of design variables so that the shape being modeled is sufficiently simple to be approximated from very sparse data.[36] This begs the question: what should these limits be? This paper presents a technique of intelligently narrowing down a search space to a smaller region of interest using low-fidelity methods. The surrogate model developed in this small region is found to be very accurate. It is then combined with several update strategies and used in a variable-fidelity optimization context to predict the global optimum in a larger design space. The method is demonstrated on a two-dimensional (2D) aerodynamic design optimization problem with good accuracy.

The remainder of the paper is organized as follows. Section 2 describes the design space reduction (DSR) technique. Update strategies for the surrogate model are discussed in Section 3. Optimization of mathematical functions is performed in Section 4, and a more complex 2D aerodynamic optimization problem is introduced in Section 5 along with the analytical methods and tools used. Airfoil optimization results are presented in Section 6 followed by conclusions in Section 7.

2. Design space reduction

For a surrogate model to be useful in an optimization context, it is important that the surrogate model is accurate at the sequence of iterates generated by the search algorithm as it converges towards the true optimum. How the model performs at other points in the parameter space is of no concern in this specific context.[38] This observation provides the pretext for development of the design space reduction technique. In previous research, the authors have pointed out the importance of selecting a low-fidelity solver capable of predicting the aerodynamic behavior that is consistent with the high-fidelity solver.[2,42] It is reasoned that such a low-fidelity solver will converge towards the region of the high-fidelity optimum – the desired region of interest. This region can be determined by examining the search trajectory of the low-fidelity solver. An accurate surrogate can be created in this small region of interest and thereafter be used for variable-fidelity optimization.

The design space reduction algorithm proceeds as follows: (1) the optimization is initially performed using a genetic algorithm (GA) coupled to a low-fidelity solver in a large design space; (2) the search trajectory is analyzed to identify the reduced design space; (3) this is followed by another optimization using the low-fidelity solver and the surrogate model in a variable-fidelity context. The flowchart of the complete variable-fidelity framework with design space reduction (DSR–VFO) is shown in Fig. 1.

Fig. 2 shows a sample design space from a low-fidelity optimization run on a 10-variable problem. Since a GA progressively converges towards the optimum, only the population members for the last five generations are analyzed. Three methods are considered for selecting a small region for generating the surrogate model:

(1) The extreme minimum and maximum values of each design variable.
(2) A normal distribution fit to the design variables with 95% confidence levels as the bounds.
(3) A 5% region around the low-fidelity optimum point. The upper and lower bounds are defined by a region that forms a 5% locus around this point.

Fig. 2 shows the design space produced by the above methods. The location of the high-fidelity optimum is also shown for reference. All methods produce regions smaller than the