Multi-Fidelity Optimization for Aerospace Vehicle Design

†Ramana V. Grandhi[1] and Christopher C. Fischer[2]

[1]Distinguished Research Professor, Department of Mechanical and Materials Engineering, Wright State University Dayton, OH, USA.
[2]Graduate Research Assistant, Wright State University, Dayton, OH

*Presenting author: Ramana.grandhi@wright.edu
†Corresponding author: Ramana.grandhi@wright.edu

Introduction

Computational optimization methodologies are growing in popularity for use in aircraft design as physical interactions are understood and system requirements increase in complexity and demand. These methodologies offer the potential of increased system performance through enabling numerous design options to be explored systematically. While optimization methods have traditionally been predominant in the latter stages of a design process (preliminary and detailed), there is a growing interest and need for the utilization of higher-fidelity physics in the earlier stages of design (conceptual). However, simulation models based on these higher-fidelity physics tend to have higher computational cost in comparison to their lower-fidelity counterparts. Therefore, optimization, which typically requires a large number of model evaluations, can be prohibitively expensive given the higher computational cost of these physics-based models. The traditional method for reducing the computational cost inherently present in optimization is to employ surrogate-based optimization (SBO) techniques. This paper presents an alternative/modified technique that leverages high-fidelity response data to correct low fidelity models for use in an SBO environment.

The aforementioned need for high-fidelity physics to be brought into the design cycle at an earlier stage of the design process is evident in the design of next generation military aircraft, both manned and unmanned. These aircraft demand increased capabilities in speed, range, survivability, mission versatility, and reliability. To satisfy these demands, one must achieve synergy between aircraft constituent sub-systems including, among others, propulsion, structures, flight controls, and materials. This necessary synergy, and resulting maximum platform performance, is only attainable through the use of a truly integrated design process. Such a process, along with the desire to identify and exploit beneficial coupling within the physics of the design domain, inherently requires leveraging higher fidelity computational simulations among various disciplines early in the design process. This stands contrary to conventional conceptual design practices that utilize the use of handbooks, spreadsheets, and legacy information. However, it is currently unclear as to when it is appropriate and/or necessary to bring in higher fidelity simulation models, or even experimental data, that will provide the best benefit to the design process.

The conceptual and preliminary design stages of aircraft design have traditionally been two separate, time intensive design phases. However, the concept of dialable/multi-fidelity design is one in which the gap between these two phases is bridged by introducing physics into the design process at an earlier stage than traditionally employed. This introduction of physics at an earlier stage eliminates the need for separate conceptual and preliminary design phases and consequently reduces total design time.

The concept of multi-fidelity design does however pose certain obstacles such as determining how and when to “dial” or switch between different fidelity models. This work explores the ability of applying an adjustment factor to the response of a low-fidelity model so as to predict the true system response taken to be the response obtained from a high-fidelity simulation throughout an optimization routine. A surrogate model is constructed for the purpose of determining an adjustment factor given any design point (thus a function of design variables) using information of previous high and low fidelity simulations from previous optimization iterations. In doing so, sensitivity information of the high-fidelity simulation model can be estimated through a combination of sensitivity information for the adjustment factor surrogate model and low fidelity models. It is shown that optimization on the high-fidelity as well as adjusted low-fidelity models converge to the same local optimum whereas, optimization on adjusted low-fidelity model does so in an order of magnitude fewer high-fidelity function evaluations.

In this research, an adaptation to traditional Trust Region Model Management schemes has been developed and employed within an optimization routine designed to facilitate the decision making process of determining “when” to utilize higher-fidelity response information. This is implemented in parallel with surrogate modeling techniques, such as Kriging, for the purpose of constructing a model of adjustment factors. This surrogate adjustment factor model is then used to correct the low-fidelity model.

Multi-Fidelity Optimization (MFO)

Numerous heuristic techniques have been used to optimize a high-fidelity function using lower-fidelity information. We consider heuristic multi-fidelity optimization (MFO) approaches to be approaches that generally converge in practice to an optimum of the high-fidelity function, but in which there is no formal mathematical proof or guarantee. These methods vary from problem specific necessities to rigorous methods that compute a probability of finding an improved high-fidelity function value. Examples of problem specific multi-fidelity approaches include adding global response surface corrections to low-fidelity models[**? ?**], using the low-fidelity function gradient as the optimization direction, but performing the line search with the high-fidelity function value, creating a response surface using both high- and low-fidelity analysis results, and running higher-fidelity models when two or more lower-fidelity models disagree. In contrast, we consider a non-heuristic method to be one in which given a set of requirements for the initial design(s) and behavior of the high and low-fidelity functions, there is a mathematical guarantee that with enough time the multi-fidelity method will find a high-fidelity optimum. Non-heuristic multi-fidelity methods may converge slower than single-fidelity methods, but nonetheless they are guaranteed to work eventually. Our discussion of MFO methods focuses on both heuristic and non-heuristic methods that are broadly applicable and likely to find a high-fidelity optimum for general problems.

In MFO, global and local approaches define two sides of a coin in which research focuses. Global methods search the entire feasible domain for the best design, whereas local methods attempt to find the nearest design that has better performance than all other designs in that neighborhood. Some approaches combine either an augmented Lagrangian, exact penalty method, constraint filtering, or barrier with either a pattern-search or a Simplex method. Other methods use linear interpolation of both the objective function and the constraint. Global methods have the benefit that they typically do not require estimates of the high-fidelity functions’ gradients. This is important because frequently a high-fidelity functions’ gradient is unavailable and cannot be estimated accurately. However, given the extreme advancements in sensitivity calculation approaches, this work makes the assumption that gradient information is not beyond the scope of attainment. A detriment to using global optimization methods is that they typically require considerably more high-fidelity evaluations than local methods. So, there is clearly a need for both types. This work focuses on local methods which may utilize global approach ideology and methodology such as enhanced surrogate modeling techniques.

Open Issues in MFO

This work only considers local optimization and does not attempt to find the globally optimal design due to the challenges addressed. For local multi-fidelity optimization four major challenge themes have emerged, (i) how to rigorously and effectively combine multiple low-fidelity models to best predict the high-fidelity function behavior, (ii) when is it appropriate to utilize high-fidelity model, an when is it safe for lower fidelity models be used to drive the design process, (iii) how to find a high-fidelity optimum in the presence of potentially computationally expensive function evaluations, and (iv) how best to use gradient information such that surrogate models use as much information as is known about the high-fidelity function.

Summary Remarks

In this work, an adjustment factor technique (surrogate-based Bayesian influenced hybrid bridge function) combined with a TRMM optimization scheme was presented for use in multi-fidelity design processes. The novel approach developed is a weighted average of additive and multiplicative adjustment factors where the weighting coefficients are calculated using a Bayesian Updating technique. This techniques is the basis of Bayesian statistics and is used in uncertainty quantification and inference-based statistics. The key to the presented surrogate-based adjustment factor technique is the use of a Gradient Enhanced Kriging surrogate model constructed over a localized trust region. This localized trust region is adaptive in the sense that its relative size and location are determined by the accuracy of the corrected low-fidelity model through the presented TRMM methodology. Therefore, a surrogate hybrid bridge function model is constructed utilizing all data available within the trust region (from previous optimization iterations) and thus adjustment factors obtained at said points utilized to correct low-fidelity simulation response; thus, more accurately predicting a high-fidelity response.

This multi-fidelity optimization methodology is demonstrated using three different optimization problems. The first problem is an unconstrained minimization in which there are two fidelities that define the objective. This problem proved the capability of arriving at the high-fidelity optimum through the use of a Bayesian inspired bridge function implemented in a TRMM optimization environment. The Second problem extends the first by adding multi-fidelity constraints to the optimization problem. Such a case is new to literature in handling multiple fidelities in the objective and constraints simultaneously without the use of Lagrange multipliers to simplify the problem via combining constraints into the objective. This problem proved the ability to maintain high-fidelity accuracy while reducing computational cost associated with the design process. Finally, third demonstration illustrates the application and benefits associated with implementing this methods on a real world engineering problem of wing body design in structural and fluid mechanics disciplines.