

# Geometric Filtration Using POD for Aerodynamic Design Optimization

David J. J. Toal,<sup>\*</sup> Neil W. Bressloff<sup>†</sup> and Andy J. Keane<sup>‡</sup>

*University of Southampton, Southampton, SO17 1BJ, United Kingdom*

When carrying out design searches, traditional variable screening techniques can find it extremely difficult to distinguish between important and unimportant variables. This is particularly true when only a small number of simulations is combined with a parameterization which results in a large number of variables of seemingly equal importance. Here the authors present a variable reduction technique which employs proper orthogonal decomposition to filter out undesirable or badly performing geometries from an optimization process. Unlike traditional screening techniques, the presented method operates at the geometric level instead of the variable level. The filtering process uses the designs which result from a geometry parameterization instead of the variables which control the parameterization. The method is shown to perform well in the optimization of a two dimensional airfoil for the minimization of drag to lift ratio, producing designs better than those resulting from traditional kriging based surrogate model optimization and with a significant reduction in surrogate tuning cost.

## I. Introduction

THE optimization of complex geometries is prevalent throughout the field of aerodynamic design, from the optimization of two dimensional airfoils, to the optimization of complete wings and aircraft. Designers are, however, restricted somewhat in the parameterization of such shapes. A typical airfoil parameterization, for example, may have anything from a few to more than 30 variables.<sup>1-5</sup> Parameterize an aircraft wing using a series of such airfoils and the number of variables can quickly be in the order of hundreds. Such large numbers of variables naturally restrict the ability of an optimization algorithm to achieve an optimal design given a limited budget of simulations.

To enable the efficient optimization of such geometries it is therefore normal for a designer to employ some method of screening to identify those variables which influence the objective function most. With these variables identified the optimization can proceed using the reduced variable set. Painchaud et al.,<sup>3</sup> for example, reduced a 34 variable NURBS parameterization of an airfoil to a total of 11 variables, while Song et al.<sup>6</sup> reduced a 33 variable parameterization of an engine nacelle to seven important variables. Such large reductions in the complexity of an optimization are especially beneficial when the objective function is evaluated using an expensive high fidelity computational simulation. The expense of such simulations, even with the growth of parallel computations, prohibits the exhaustive search of design spaces with large numbers of variables. In such problems the reduction of the design space offered by variable screening is invaluable, allowing the designer to use a specified simulation budget more effectively.

Variable screening does however suffer from a number of significant problems. The cost of screening a complex optimization problem which requires evaluations using high fidelity simulations may rival or exceed that of the optimization of the reduced variable set. Given that the screening procedure may not in fact return a definite set of variables which contribute most significantly to the change in objective function, the screening budget might be better used in the optimization of the original complete variable set. Variable screening also results in a reduction in the flexibility of the geometry parameterization, reducing the ability

---

<sup>\*</sup>Graduate Research Student, School of Engineering Sciences, Student Member AIAA

<sup>†</sup>Lecturer, School of Engineering Sciences.

<sup>‡</sup>Professor of Computational Engineering, School of Engineering Sciences.

of the optimization to achieve certain high performance designs which may only be found with the original, complete, variable set.

Optimization of complex geometries with large numbers of variables and a limited simulation budget therefore equates to a simple trade-off. The designer can attempt to optimize the original parameterization and potentially achieve a high performance design, though with a limited budget this optimal design will be difficult to locate. Alternatively designers can spend a proportion of their simulation budget to screen out unimportant or insignificant variables and optimize the reduced design space using the remainder of the budget. The optimal design within the reduced space may be easier to locate, but the actual design may be significantly worse than that which could be obtained using the complete variable set. In the worst case scenario the variable screening procedure may indicate no significant difference in the importance of each variable. A sub-optimal reduced set of variables may therefore be chosen, significantly hampering the optimizer's ability to achieve a good design.

An optimization strategy which combines the advantages of design space reduction associated with variable screening, whilst retaining the majority of the flexibility of the original optimization would therefore be desirable. To this end the authors have devised a strategy which combines surrogate modeling techniques and proper orthogonal decomposition (POD) in an attempt to optimize aerodynamic problems consisting of large numbers of variables but given only a limited simulation budget. Within the framework of this optimization strategy proper orthogonal decomposition is used to reparameterize the problem in an attempt to filter out badly performing geometries.

Aerodynamicists already employ a form of geometric filtration, perhaps without even realizing it, for some aerodynamic optimization problems. Consider, for instance, the design of an airfoil for optimal performance at transonic speeds. Here a designer utilizes their knowledge of existing airfoil shapes and bases the optimization on an existing supercritical airfoil, such as the RAE-2822. An optimization for a particular transonic flight condition therefore results in a series of perturbations to this baseline geometry, be it through the manipulation of NURBS<sup>1</sup> control points or through the addition of analytical functions.<sup>4</sup> The choice of this baseline airfoil has immediately filtered out badly performing designs and in doing so has reduced the design space before even commencing the parameterization.

Robinson et al.<sup>7</sup> used this concept to construct an airfoil parameterization based on a series of orthogonal bases derived from an ensemble of supercritical airfoils. This parameterization technique operated on the principle that airfoils which perform well at a particular flight condition, the transonic regime in this case, had a number of common geometric features which could be extracted through the orthogonalization process and used for optimization. This is perfectly feasible in the case of 2D airfoil design, where large databases of the performance of different airfoils exist, Abbott and Doenhoff<sup>8</sup> being a popular example. Problems arise when an optimization problem is encountered for which there is no such literature. There is, for example, no similar database of the performance of different wing body fairings, the best designs of which could be used to construct a similar set of orthogonal bases. While the generation of such basis functions is undoubtedly a useful tool, this approach cannot be applied to every aerodynamics problem.

The following paper begins with an overview of the proposed geometric filtration procedure which is then applied to the optimization of a transonic airfoil for minimum drag to lift ratio. This test optimization problem is investigated using three different optimization techniques. An extensive optimization of the airfoil is first performed directly using a genetic algorithm. The problem is then investigated using a traditional kriging based approach, which is considered a direct competitor to the geometric filtration strategy. Finally the performance of the geometric filtration strategy is considered. The application of the geometric filtration strategy to this design problem sees an improvement in the overall quality of the final design with the additional benefit of a substantial reduction in the cost of the tuning of the kriging hyperparameters.

## II. Optimization Via Geometric Filtration

### A. Overview of the Complete Methodology

At its heart the proposed optimization methodology consists of a basic surrogate modeling optimization strategy. This approach to optimization is particularly popular when objective function evaluations are extremely expensive. The surrogate model attempts to model the response of the objective function to changes in the design variables. A stochastic optimizer, for example, a genetic algorithm, can then draw on the surrogate model instead of the expensive computer simulation in an attempt to locate regions of optimal design. Typically the expensive computer simulation is subsequently used to evaluate the true objective

function at these positions. The true objective function can then be used to update the surrogate model, which can be searched again. This cycle is then repeated until some stopping criterion is reached.

The proposed optimization strategy, shown in figure 1, also begins with an initial surrogate based optimization, in this case employing a krig to construct the surrogate. This initial optimization utilizes the designer’s initial geometry parameterization and a proportion of the total simulation budget. This is followed by a reparameterization procedure using proper orthogonal decomposition which attempts to both reduce the number of design variables and filter out badly performing designs. A secondary surrogate model optimization is then performed utilizing this new geometry parameterization and the remainder of the total simulation budget.

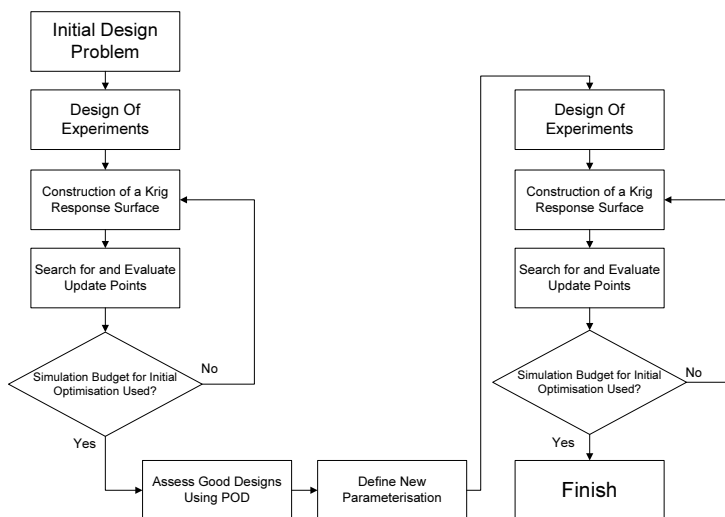


Figure 1. An overview of the geometric filtration optimization methodology

In the context of the proposed optimization strategy, the initial surrogate model optimization followed by the POD, performs a job similar to that of traditional variable screening, except that the screening is performed at the geometric instead of the variable level. This prevents the reduction in variables being restricted to the those of the original parameterization. It is hoped that the initial optimization provides a number of good designs which are then decomposed using POD into a series of orthogonal bases which capture the common features of these good designs. A similar process has been used by Kamali et al.<sup>9</sup> to reduce a design space but using points from a Latin hypercube sampling of the design space and not a subset of points from an initial optimization.

As shown throughout the literature<sup>10,11</sup> POD cannot accurately represent data outside of the initial snapshot ensemble. Here, this particular feature of POD, combined with an initial search offers an advantage by considering only the current best geometries in the construction of the orthogonal bases: the bad geometries cannot be recreated through a combination of POD bases, hence they have been filtered out. The second surrogate model optimization therefore has the benefit of a reduced number of variables with a minimal reduction in geometric flexibility.

## B. Kriging

The implementation of the geometric filtration strategy considered within this paper utilizes a kriging based optimization in both the initial and secondary optimizations. Although a different surrogate modelling method, such as a simple polynomial or a radial basis function could conceivably be used within this framework, kriging is used due to its ability to more accurately represent complicated responses whilst providing an error estimate of the predictor. First used by geologists in the estimation of mineral concentrations, it has since been popularised by Sacks et al.<sup>12</sup> in the creation of surrogate models of deterministic computational experiments and has been used successfully in the optimization of a number of different design problems.<sup>13–16</sup>

To demonstrate the basic process of kriging we consider the optimization of an objective function,  $y$ , which is dependant on the vector of variables,  $\mathbf{x}$ , which in this case are the variables defining the initial geometry

parameterization. The initial bound constrained optimization problem in  $d$  dimensions is therefore,

$$\begin{aligned} & \text{Minimise} && y(\mathbf{x}) \\ & \text{subject to} && \\ & && \mathbf{l} < \mathbf{x} < \mathbf{u}. \end{aligned} \tag{1}$$

where  $\mathbf{l}$  and  $\mathbf{u}$  refer to the lower and upper bounds of the variables,  $(\mathbf{x})$  respectively. In general the objective function values  $y(\mathbf{x}_i)$  and  $y(\mathbf{x}_j)$ , which depend on  $d$  variables will be similar if the distance between  $\mathbf{x}_i$  and  $\mathbf{x}_j$  is small. This can be modelled statistically by considering the correlation between two points as,

$$\exp\left(-\sum_{l=1}^d \theta_l \|\mathbf{x}_{i_l} - \mathbf{x}_{j_l}\|^{p_l}\right) \tag{2}$$

where  $\theta_l$  and  $p_l$  are known as the hyperparameters and determine the rate of correlation decrease and the degree of smoothness in the  $l^{\text{th}}$  direction respectively. These hyperparameters, including a regression constant if required,<sup>17</sup> are chosen to maximise the likelihood on the observed dataset  $\mathbf{y}$ , where  $\mathbf{y}$  is a vector of  $n$  objective function values found by sampling the problem space.

The concentrated likelihood function,<sup>18</sup>

$$-\frac{n}{2} \ln(\hat{\sigma}^2) - \frac{1}{2} \ln(|\mathbf{R}|), \tag{3}$$

is evaluated by first calculating the mean,

$$\hat{\mu} = \frac{\mathbf{1}^T \mathbf{R}^{-1} \mathbf{y}}{\mathbf{1}^T \mathbf{R}^{-1} \mathbf{1}} \tag{4}$$

and then the variance,

$$\hat{\sigma}^2 = \frac{1}{n} (\mathbf{y} - \mathbf{1}\hat{\mu})^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{1}\hat{\mu}), \tag{5}$$

where  $\mathbf{1}$  is an  $n \times 1$  vector of ones and  $\mathbf{R}$  is the correlation matrix, the  $i, j^{\text{th}}$  elements of which are calculated using Eq. (2).

The concentrated likelihood is dependent only on the symmetric matrix  $\mathbf{R}$  and hence on the hyperparameters which are then optimized to maximise the likelihood. Any optimization strategy could conceivably be used for this optimization problem although a global optimizer is generally preferred due to the multimodal nature of the likelihood.<sup>19</sup> With the hyperparameters defined, the surrogate model can be used to predict regions which either minimise the models prediction of the objective function or maximise its expected improvement.<sup>18</sup>

### C. Proper Orthogonal Decomposition

Proper orthogonal decomposition (POD), otherwise known as Principal Component Analysis, or Karhunen-Loève expansion, has been extensively used throughout engineering with regard to computational fluid dynamics (CFD). It has been used in the derivation of reduced order models for the purposes of control,<sup>20-23</sup> optimization<sup>24-26</sup> and has also been used as a flow analysis tool.<sup>27</sup> POD decomposes a series, or ensemble, of snapshots<sup>28</sup> of data into a set of optimal orthogonal basis functions of decreasing importance. The bases functions are optimal in the sense that no other bases functions will capture as much information in as few dimensions.<sup>29</sup> Applying POD, therefore, to an ensemble of airfoils results in a series of orthogonal bases similar to those of Robinson et al.<sup>7</sup>

The decomposition process, as per Sirovich's method of snapshots,<sup>28</sup> begins with the definition of an ensemble of snapshot vectors,  $\mathbf{S}$ . In the case of the geometric optimizations considered within this paper, the snapshot ensemble is constructed from a series of  $M$  vectors consisting of the  $x$  and  $y$  coordinates of a number of designs selected from the initial optimization,

$$\mathbf{S} = [\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_M]. \tag{6}$$

The matrix of snapshots is then decomposed into a mean,  $\bar{\mathbf{s}}$  and a matrix of the fluctuations of each snapshot from this mean,

$$\mathbf{S} = \bar{\mathbf{s}} + \mathbf{F} \tag{7}$$

$$\text{where } \bar{\mathbf{s}} = \frac{1}{M} \sum_{i=1}^M \mathbf{s}_i. \quad (8)$$

The orthogonal basis functions are calculated by considering the solution to the eigenvalue problem,

$$\mathbf{C}\mathbf{V} = \mathbf{\Lambda}\mathbf{V}, \quad (9)$$

where  $\mathbf{\Lambda}$  is a vector of eigenvalues and the square symmetric correlation matrix  $\mathbf{C}$  is given by,

$$\mathbf{C} = \mathbf{F}^T \mathbf{F}. \quad (10)$$

The matrix of eigenvectors,  $\mathbf{V}$ , can then be used to calculate the matrix of eigenfunctions,

$$\mathbf{\Phi} = \mathbf{F}\mathbf{V}, \quad (11)$$

where  $\mathbf{\Phi}$  is a matrix of  $M$  eigenfunctions. These eigenfunctions, along with the corresponding vector of modal coefficients,  $\boldsymbol{\alpha}$ , or matrix of modal coefficients,  $\mathbf{A}$ , allow the fluctuations and hence the original snapshots to be reconstructed,

$$\begin{aligned} \mathbf{s}_i &= \bar{\mathbf{s}} + \mathbf{\Phi}\boldsymbol{\alpha}_i \\ &\text{or} \\ \mathbf{S} &= \bar{\mathbf{s}} + \mathbf{\Phi}\mathbf{A}. \end{aligned} \quad (12)$$

The advantage of POD is that not all of the  $M$  POD bases (eigenfunctions) are necessary to recreate the original snapshot ensemble to a required degree of accuracy. The cumulative percentage variation<sup>30</sup> can be used to define a reduced number of bases functions with which the original ensemble can be approximately recreated. The importance of each POD basis function is related to the relative magnitude of the corresponding eigenvalue, a large eigenvalue therefore indicates an important basis function. The cumulative percentage variation,

$$\frac{\sum_{i=1}^N \lambda_i}{\sum_{i=1}^M \lambda_i} \times 100, \quad (13)$$

is therefore a measure of the combined importance of the first  $N$  bases. Using this simple calculation a reduced number of bases can be selected in order to meet a minimum required percentage variation. Using this reduced number of bases, the original snapshot vectors can be approximated by,

$$\mathbf{s}_i \approx \bar{\mathbf{s}} + \sum_{i=1}^N \boldsymbol{\alpha}_i \boldsymbol{\phi}_i, \quad (14)$$

where  $\boldsymbol{\phi}_i$  and  $\boldsymbol{\alpha}_i$  are the  $N$  most important basis vectors and corresponding modal coefficients. The modal coefficients corresponding to each of the original snapshot vectors can be calculated using the orthonormality property of the basis vectors.

The geometric filtration optimization methodology therefore moves from a surrogate modelling optimization based on the magnitude of the original design variables to one which considers the magnitude of the POD modal coefficients. The bounds of the secondary optimization are defined by the minimum and maximum modal coefficients of the original snapshot ensemble. The optimization problem therefore moves from the formulation of Eq. (1) in  $d$  dimensions to one in  $N$  dimensions,

$$\begin{aligned} &\text{Minimise } y(\boldsymbol{\alpha}) \\ &\text{subject to} \\ &\min(\mathbf{A}) < \boldsymbol{\alpha} < \max(\mathbf{A}). \end{aligned} \quad (15)$$

### III. Optimization of a Transonic Airfoil

#### A. Description of the Test Problem

To demonstrate the effectiveness of the proposed strategy the optimization of a 2D airfoil is used as a test case. The RAE-2822 airfoil is parameterized using two NURBS curves, one each for the upper and lower

surface with the positions and weights of the controls points defining both NURBS curves optimized using a local BFGS optimization as per Lépine et al.<sup>1</sup> The resulting parameterization of the RAE-2822, shown in figure 2, consists of a total of 20 variables. The control points at the leading and trailing edges are fixed, the two control points on the line,  $x = 0$ , are only permitted to move vertically to maintain curvature continuity at the leading edge, while the remaining control points can move in both axes. Bounds are placed on the movement of each of the control points to help prevent unrealistic designs. Although the weight of each control point remains fixed for the purposes of the following optimizations they could be permitted to vary, increasing the dimensionality of the problem to 31.

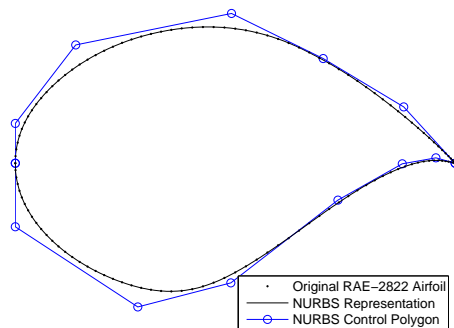


Figure 2. NURBS parameterization of the RAE-2822 airfoil

The airfoil is optimized to minimize the drag to lift ratio at Mach 0.725, Reynolds number of  $6 \times 10^6$  and a fixed angle of attack of  $2^\circ$ , using the full potential solver VGK.<sup>31</sup> At these flow conditions the RAE-2822 has a drag to lift ratio of 0.0148 equating to a lift to drag ratio of 67.8, and exhibits a noticeable upper surface shockwave just before the mid-chord point, figure 3(a).

The speed of the VGK solver, approximately one second per simulation, offers a major advantage when analyzing the performance of the geometric filtration strategy. The speed of the solver allows extensive averaging to be carried out thus giving a more accurate picture of the performance of both the traditional kriging and geometric filtration approaches to design optimization.

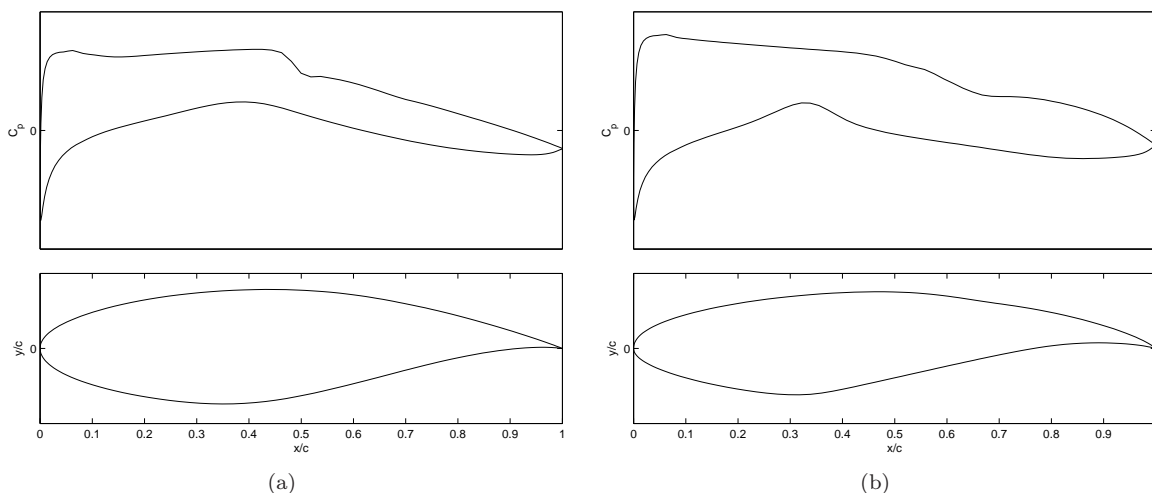


Figure 3. Pressure distribution and geometry of the original RAE-2822 airfoil with  $\frac{C_D}{C_L} = 1.48 \times 10^{-2}$  or  $\frac{C_L}{C_D} = 67.8$  (a), and for an example airfoil resulting from an optimization utilising a genetic algorithm followed by dynamic hill climber with  $\frac{C_D}{C_L} = 1.08 \times 10^{-2}$  or  $\frac{C_L}{C_D} = 93.0$  (b)

## B. Exhaustive Optimization Using a Genetic Algorithm

The speed of the objective function evaluations also allows a number of exhaustive searches of the design space using a genetic algorithm, (GA). Stochastic methods, such as the genetic algorithm, provide a reliable way of locating the region of the global optimum within a design space, however, as they typically require a large number of function evaluations, it is not generally feasible to use such methods when expensive high fidelity simulations are required. Applying a genetic algorithm to the above airfoil optimization problem provides a useful indication as to the true optimum which both the geometric filtration strategy and the traditional kriging strategy are attempting to attain.

The airfoil design problem was optimized using a GA followed by a dynamic hill climber (DHC), both implemented using the OptionsMatlab<sup>32</sup> design exploration system. A budget of 10,000 function evaluations was used in the optimization, with the budget split evenly between the GA and DHC. The 5,000 available function evaluations for the genetic algorithm equated to 100 generations of 50 points each. Although genetic algorithms are typically very good at locating the general region of the global optimum, they can be very slow to converge to a precise answer. The DHC is therefore used to converge the optimization towards a more accurate solution.

This extensive optimization was carried out a total of 10 times, producing an average  $\frac{C_D}{C_L}$  of  $1.06 \times 10^{-2}$  with a standard deviation of  $4.59 \times 10^{-4}$  which equates to a  $\frac{C_L}{C_D}$  of 94.7, an improvement of some 28.4% over the original RAE-2822. Figure 3(b) provides an indication of the designs this extensive optimization produced. Upon comparison to the original airfoil, Figure 3(a), one can clearly see a complete removal of the upper surface shockwave, a reduction in the upper surface pressure and an increase in the lower surface pressure, which results in an overall decrease in the drag to lift ratio. The optimization process has resulted in a reduction to the leading edge curvature and an overall reduction in thickness. The trailing edge camber of the airfoil has also been increased and the thickest region of the airfoil has moved forward compared to that of the original RAE-2822 airfoil.

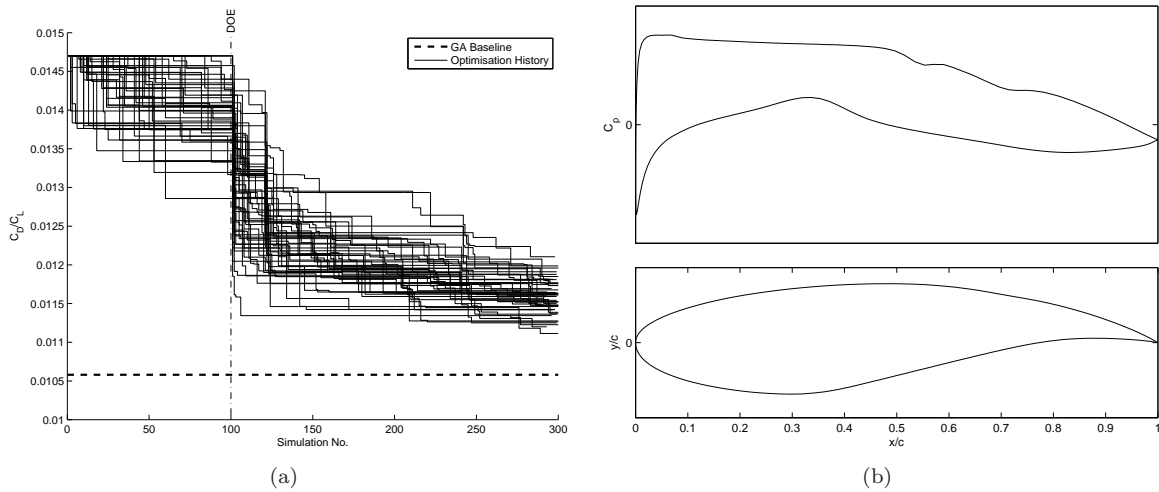
## C. Standard Kriging Based Optimization

As previously discussed, surrogate modelling techniques, such as kriging, attempt to model the response of the objective function to changes in the design variables. This model is then extensively searched and updated resulting in fewer evaluations of the objective function compared to that of a direct optimizer. The geometric filtration methodology considered within this paper comprises of two such kriging based optimizations linked by a reparameterization of the design problem. It is, therefore, necessary to determine the performance of a standard kriging optimization with respect to the current airfoil design problem in order to provide a meaningful measure of the performance gains offered by the geometric filtration strategy.

Consider now a basic kriging optimization consisting of a total budget of 300 objective function evaluations, a value much more indicative of what would be available in a typical design optimization than the 10,000 evaluations used in the previous exhaustive search. Of this total simulation budget, one third are used in the initial design of experiments (DOE), in accordance with the work of Sóbester et al.,<sup>33</sup> with the remaining budget reserved for updates to the surrogate model which are evaluated in batches of ten. All updates to the kriging models used in each of the following investigations are based on the model's prediction of the objective function. A genetic algorithm is used to minimise the objective function predicted by the model with the cluster centroids of the final population selected as the update points.

The kriging model has its hyperparameters tuned after every other set of updates. This strategy was found by Toal et al.<sup>34</sup> to offer a significant reduction in tuning cost whilst having little impact on the performance of the optimization. Each of the hyperparameters,  $(\theta_l$  and  $p_l)$ , are permitted to vary and regression is included in the kriging model as per the work of Forrester et al.<sup>17</sup> The 41 hyperparameters controlling the kriging model are optimized using a genetic algorithm with a population size of 50 for 100 generations followed by a dynamic hill climber.<sup>34</sup>





**Figure 4.** optimization histories for each of the 50 traditional kriging airfoil optimizations (a), pressure distribution and geometry for an example airfoil resulting from a traditional kriging optimization process with  $\frac{C_D}{C_L} = 1.16 \times 10^{-2}$  or  $\frac{C_L}{C_D} = 86.2$  (b)

This traditional kriging based optimization strategy is applied to the airfoil design problem a total of 50 times, with the random number seed used to generate the DOE changing each time. The optimization histories for each of these 50 optimizations are presented in Figure 4(a), with a typical airfoil resulting from one such optimization presented in Figure 4(b).

The 50 optimizations result in an average  $\frac{C_D}{C_L}$  of  $1.16 \times 10^{-2}$  with a standard deviation of  $2.17 \times 10^{-4}$  which equates to a  $\frac{C_L}{C_D}$  of 86.4 and an improvement in the objective function of 21.7%. The improvement over the initial airfoil design is therefore not as significant as that obtained using the direct GA. However, one must note that the direct search used over 33 times the number of function evaluations. The kriging strategy should therefore be commended for attaining 76.3% of the improvement obtained by the extensive optimization with 3% of the simulation budget. Given enough time and updates one would expect the traditional kriging strategy to approach the results obtained by the GA, denoted by the dotted line in figure 4(a).

Similar to the airfoils resulting from the exhaustive search, (figure 3(b)), the kriging based optimization results in a reduction in the upper surface pressure and an increase in lower surface pressure over the whole airfoil. The shockwave appears to have been removed but there are a number of noticeable oscillations in the upper surface pressure distribution compared to that resulting from the direct genetic algorithm. This results in an overall decrease in drag to lift ratio, but one which is not quite as significant as that obtained using the genetic algorithm.

When one considers the average time taken for the traditional kriging optimization, approximately 22.0 hours, one can observe a serious drawback associated with the use of this particular optimization strategy: the cost of hyperparameter tuning. Given the low cost of each simulation in this particular optimization the cost of the complete simulation budget is therefore a small fraction of the total optimization time (approximately 0.38%). The majority of the time is due to the hyperparameter tuning process, and more specifically the  $O(n^3)$  decomposition of the correlation matrix necessary in each evaluation of the likelihood. The size of this correlation matrix and hence the cost of every evaluation of the likelihood increases as more update points are added to the surrogate model. The cost of tuning the hyperparameters will therefore only increase as an optimization progresses and become a considerable bottleneck towards the end of the process.

A designer must therefore carefully consider the cost incurred by hyperparameter tuning relative to that of the objective function evaluations to determine if kriging is an appropriate optimization strategy for the problem. Taking the optimization of the RAE-2822 airfoil as an example, the cost of constructing the kriging model vastly outweighs the cost of each objective function evaluation and a designer could use a simpler strategy but with a larger number of evaluations, such as the direct GA. However, if the design problem is much more complex, for example, involving the optimization of a 3D wing using Reynolds Averaged Navier-Stokes (RANS) simulations, then the cost of each objective function evaluation may be in the region of tens of hours. For such a case the extra time incurred through the hyperparameter tuning is more acceptable.



In terms of the total optimization time incurred, the implementation of a standard kriging strategy is therefore more practical given expensive objective function evaluations but less so as the cost of the objective function decreases. Even so, the reduction in overall tuning cost, even for expensive simulations, is a worthwhile goal allowing the designer to complete an optimization in a shorter time, or allowing additional objective function evaluations within a similar time frame.

## IV. Geometric Filtration

The optimization of the RAE-2822 airfoil, utilizing the geometric filtration strategy described in Section II, is considered next. In this initial investigation a total budget of 300 simulations, (identical to that used in the traditional kriging strategy), is employed, with half of this budget used in the initial optimization and the remainder used in the secondary optimization.

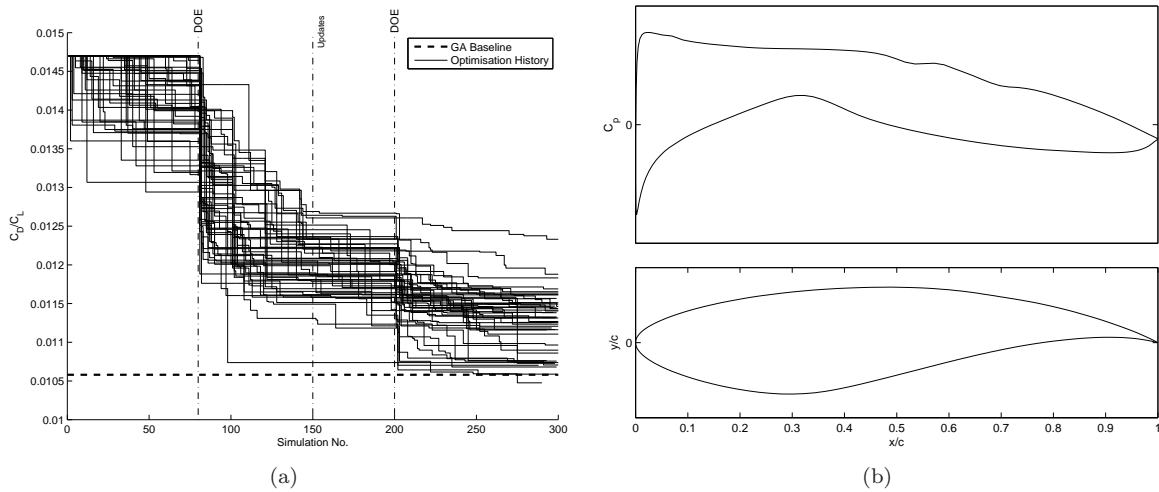
In the initial kriging based optimization the available budget is split evenly between the initial DOE and the updates to the kriging model instead of one third for the DOE as the initial response surface resulting from the use of a very small number of sample points is deemed inadequate. The secondary optimization however uses one third of the available simulations in its DOE. A smaller ratio of DOE to updates is more reasonable for this optimization as there is a smaller number of variables after the reparameterization, hence fewer points are required to construct an adequate response surface.

In this particular implementation of the geometric filtration strategy a total of 30 airfoil geometries are selected from those generated during the initial kriging based optimization. The airfoils are selected using a KMEANS<sup>35</sup> clustering algorithm in an attempt to maintain some measure of diversity between the selected airfoils. Upon decomposing the snapshot ensemble into the POD bases, the first  $N$  bases which captured a cumulative total percentage variation of greater than 99.99% are selected as the POD modes for the secondary optimization. Such a large cumulative percentage variation ensures that the original airfoils contained within the snapshot ensemble can be recreated to a high degree of accuracy. This high accuracy means that the original 30 airfoils exist within the design space resulting from the reparameterization. As the objective function values are already known for these airfoils, they allow, what is essentially, 30 additional “free” design points to be added to the design of experiments of the secondary optimization.

In summary the optimization begins with a DOE of 80 points, followed by 70 updates to the kriging model in batches of 10. From this initial optimization 30 geometries are selected and reparameterized using POD. The secondary optimization then optimizes the coefficients of the POD bases, beginning with a DOE of 50 points, (plus the 30 airfoils from the snapshot ensemble), and followed by 100 updates to the model, again in batches of ten. As with the traditional kriging optimization discussed previously, the hyperparameters are tuned after alternate updates in accordance with Toal et al.<sup>34</sup>

**Table 1. Comparison of the three optimization strategies with respect to the optimization of the RAE-2822 for minimum drag to lift ratio**

| Strategy             | No. of Evaluations | Mean $\frac{C_D}{C_L}$ | Std $\frac{C_D}{C_L}$ | Mean $\frac{C_L}{C_D}$ | Std $\frac{C_L}{C_D}$ | Time (hrs) |
|----------------------|--------------------|------------------------|-----------------------|------------------------|-----------------------|------------|
| Genetic Algorithm    | 10,000             | $1.06 \times 10^{-2}$  | $4.59 \times 10^{-4}$ | 94.7                   | 4.60                  | 2.78       |
| Traditional Krig     | 300                | $1.16 \times 10^{-2}$  | $2.17 \times 10^{-4}$ | 86.4                   | 1.62                  | 22.0       |
| Geometric Filtration | 300                | $1.13 \times 10^{-2}$  | $3.76 \times 10^{-4}$ | 89.0                   | 2.98                  | 4.15       |



**Figure 5.** optimization histories for each of the 50 optimizations utilising the geometric filtration strategy with an equal simulation budget for both the initial and secondary optimizations (a), pressure distribution and geometry for an example airfoil resulting from the geometric filtration optimization process with  $\frac{C_D}{C_L} = 1.13 \times 10^{-2}$  or  $\frac{C_L}{C_D} = 88.9$  (b)

Once again a total of 50 different optimizations are carried out, with the random number seed used in the Latin Hypercube changing each time. The histories for each of these optimizations are presented in figure 5(a). The 50 optimizations result in an average drag to lift ratio of  $1.13 \times 10^{-2}$  with a standard deviation of  $3.76 \times 10^{-4}$  equating to an average lift to drag ratio of 89.0. Applying the geometric filtration optimization strategy to this design problem has therefore resulted in an average of 84.1% of the improvement obtained with the genetic algorithm. From figure 5(a) it is also observed that a number of the final designs actually achieve a drag to lift ratio similar to that of the direct GA. The average results of these optimizations along with those of the traditional krig and the GA are presented together in Table 1.

Figure 5(b) shows a typical airfoil resulting from an optimization using this implementation of the geometric filtration strategy. Upon comparison to the airfoil in figure 4(b), resulting from the traditional kriging optimization, one can observe the removal of the shockwave but with a reduction to the oscillations in upper surface pressure. The reduction in the upper surface pressure and increase in the lower surface pressure has been maintained. The optimization has also resulted in an airfoil with a more significant reduction to the leading edge curvature, something which was observed with the direct GA, (figure 3(b)), but which is not as pronounced in figure 4(b).

Upon comparison of the optimization histories of figures 4(a) and 5(a) a disadvantage of the current implementation of the geometric filtration strategy can be observed. The application of geometric filtration appears to have resulted in an increase in the variance of the final designs. The standard deviation of the drag to lift ratio for the traditional kriging strategy was  $2.17 \times 10^{-4}$  while geometric filtration resulted in an increased standard deviation of  $3.76 \times 10^{-4}$ . Therefore, although geometric filtration results in a better design on average, there is slightly less consistency in the quality of the final designs.

Consideration of the total optimization time leads to the observation of another advantage of geometric filtration over the traditional kriging based optimization. Instead of an average optimization time of 22.0 hours, the utilization of geometric filtration has resulted in an average optimization time of 4.15 hours. As the krig tuning process is again responsible for the majority of the optimization time, geometric filtration has resulted in a substantial reduction in tuning cost. This reduction in cost can be explained when one considers the actual tuning process in each optimization, particularly the size of the correlation matrix,  $\mathbf{R}$ , used to calculate the likelihood, Eq. (3).

In the traditional kriging based optimization the correlation matrix begins at  $100 \times 100$  in size and steadily grows as more update points are evaluated until, when the hyperparameters are to be tuned for the final time, the matrix is  $280 \times 280$  in size. The above implementation of the geometric filtration strategy however begins with a correlation matrix of  $80 \times 80$  which grows to  $140 \times 140$  by the final tune of the initial optimization. Following the reparameterization process, 110 of these design points are filtered out and 50 new points added from the design of experiments of the secondary optimization. The correlation matrix

therefore grows steadily from  $80 \times 80$  to  $160 \times 160$  when the hyperparameters of the secondary krig are tuned for the final time. As the correlation matrix remains consistently smaller when geometric filtration is used, the  $O(n^3)$  cost of the factorisation in each likelihood evaluation is reduced hence reducing the overall cost of the hyperparameter tuning considerably.

In conclusion, the optimization of the RAE-2822 using the proposed geometric filtration strategy has resulted in better designs than the traditional kriging strategy for a considerably reduced tuning overhead but with a slight reduction in the consistency of these designs.

## V. Conclusions

An optimization strategy involving two kriging based response surface optimizations and a POD based reparameterization has been introduced and applied to the optimization of an airfoil for minimum drag to lift ratio. This strategy, termed geometric filtration, was found to outperform a traditional kriging based optimization, producing better designs for a considerable reduction in overall optimization cost.

The geometric filtration strategy applies an initial kriging response surface model optimization to the original problem. From the results of this optimization a number of good design points are selected to form a snapshot ensemble for the purposes of proper orthogonal decomposition (POD). The POD basis functions then act as a reparameterization of the original problem, filtering out badly performing designs and reducing the number of variables. A secondary kriging response surface based optimization is then carried out in which the modal coefficients of the POD bases are optimized.

The optimization of a transonic airfoil for minimum drag to lift ratio was used as a test case to compare the geometric filtration strategy to a traditional kriging based optimization and an extensive direct optimization using a genetic algorithm. The genetic algorithm used a total of 10,000 evaluations of the objective function and achieved a 28.4% improvement in the drag to lift ratio. Such a large number of function evaluations is not typically possible when a high fidelity simulation is used, and is only included here to provide an indication of the true optimum to which the traditional krig and geometric filtration are attempting to attain. The traditional kriging strategy achieved 76.3% of the improvement obtained by the genetic algorithm but with only 300 objective function evaluations. However, applying geometric filtration to the same problem, again using 300 objective function evaluations, produced designs achieving 84.1% of the improvement obtained with the genetic algorithm, a substantial improvement over the traditional kriging strategy.

Due to the reduction in the size of the correlation matrix used in the calculation of the concentrated likelihood, the application of the geometric filtration strategy realises a drastic reduction in the cost of surrogate model construction reducing from a total of 22.0 hours when using the traditional kriging strategy to a much more respectable 4.2 hours.

Although the implementation of geometric filtration resulted in a slight increase in the variance of the final designs obtained, the authors are confident that through careful management of the setup of the strategy this can be overcome. The size of the initial optimization, the number of airfoils selected to form the POD snapshot ensemble and the number of bases used in the secondary optimization will all require careful consideration in order to obtain the most from this strategy.

As the geometric filtration strategy is based upon a series of surrogate model optimizations it stands to reason that any techniques that can be applied to a normal surrogate model can also be applied to geometric filtration. The initial or secondary optimizations could make use of gradient or Hessian enhanced surrogate models as well as more "exotic" updating formulations or could even incorporate variable fidelity simulations through co-kriging. The basic geometric filtration strategy presented within this paper could therefore be improved on a number of fronts.

In conclusion geometric filtration offers an improvement in both the quality of the final designs and the overall cost of an optimization compared to a traditional kriging strategy but further investigations are required to reach this strategy's full potential.

## Acknowledgments

The presented work was undertaken as part of an Airbus funded activity. The authors would therefore like to thank Dr. C. Holden of Airbus for her much appreciated help and support. The authors would also like to thank Dr. A. Forrester and Dr. A. Sóbester of the University of Southampton for their advice and input.

## References

- <sup>1</sup>Lépine, J., Guibault, F., and Trépanier, J., “Optimized Nonuniform Rational B-Spline Geometrical Representation for Aerodynamic Design of Wings,” *AIAA Journal*, Vol. 39, No. 11, 2001, pp. 2011–2041.
- <sup>2</sup>Song, W. and Keane, A., “A Study of Shape Parameterisation Methods for Airfoil Optimization,” *10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 2004, pp. 2031–2038.
- <sup>3</sup>Painchaud-Ouellet, S., Tribes, C., Trépanier, J., and Pelletier, D., “Airfoil Shaped Optimization Using a Nonuniform Rational B-Spline Parameterization Under Thickness Constraint,” *AIAA Journal*, Vol. 44, No. 10, 2006, pp. 2170–2178.
- <sup>4</sup>Hicks, R. and Henne, P., “Wing Design by Numerical Optimisation,” *Journal of Aircraft*, Vol. 15, No. 7, 1978, pp. 407–412.
- <sup>5</sup>Sobieczky, H., “Parametric Airfoils and Wings,” *Notes on Numerical Fluid Mechanics*, Vol. 68, 1998, pp. 71–88.
- <sup>6</sup>Song, W. and Keane, A., “Surrogate-Based Aerodynamic Shape Optimization of a Civil Aircraft Engine Nacelle,” *AIAA Journal*, Vol. 45, No. 10, 2007, pp. 2562–2574.
- <sup>7</sup>Robinson, G. and Keane, A., “Concise Orthogonal Representation of Supercritical Aerofoils,” *Journal of Aircraft*, Vol. 38, No. 3, 2001, pp. 580–583.
- <sup>8</sup>Abbott, I. and Von Doenhoff, A., *Theory of Wing Sections*, Dover Publications Inc., 1960.
- <sup>9</sup>Kamali, M., Ponnambalam, K., and Soulis, E., “Integration of Surrogate Optimization and PCA for Calibration of Hydrologic Models, A WATCLASS Case Study,” *IEEE International Conference on Systems, Man and Cybernetics*, 2007, pp. 2733–2737.
- <sup>10</sup>LeGresley, P. and Alonso, J., “Dynamic Domain Decomposition and Error Correction for Reduced Order Models,” *41st AIAA Aerospace Sciences Meeting & Exhibit*, 2003.
- <sup>11</sup>Bui-Thanh, T. and Willcox, K., “Aerodynamic Data Reconstruction and Inverse Design Using Proper Orthogonal Decomposition,” *AIAA Journal*, Vol. 42, No. 8, 2004, pp. 1505–1516.
- <sup>12</sup>Sacks, J., Welch, W., Mitchell, T., and Wynn, H., “Design and Analysis of Computer Experiments,” *Statistical Science*, Vol. 4, No. 4, 1989, pp. 409–435.
- <sup>13</sup>Hoyle, N., Bressloff, N., and Keane, A., “Design Optimization of a Two-Dimensional Subsonic Engine Air Intake,” *AIAA Journal*, Vol. 44, No. 11, 2006, pp. 2672–2681.
- <sup>14</sup>Keane, A., “Statistical Improvement Criteria for Use in Multiojective Design Optimization,” *AIAA Journal*, Vol. 44, No. 4, 2006, pp. 879–891.
- <sup>15</sup>Huang, D., Allen, T., Notz, W., and Miller, R., “Sequential Kriging Optimization Using Multiple-Fidelity Evaluations,” *Structural and Multidisciplinary Optimization*, Vol. 32, No. 5, 2006, pp. 369–382.
- <sup>16</sup>Sakata, S., Ashida, F., and Zako, M., “Structural Optimization Using Kriging Approximation,” *Computer Methods in Applied Mechanics and Engineering*, Vol. 192, No. 7-9, 2003, pp. 923–939.
- <sup>17</sup>Forrester, A., Keane, A., and Bressloff, N., “Design and Analysis of “Noisy” Computer Experiments,” *AIAA Journal*, Vol. 44, No. 10, 2006, pp. 2331–2339.
- <sup>18</sup>Jones, D., “A Taxonomy of Global Optimization Methods Based on Response Surfaces,” *Journal of Global Optimization*, Vol. 21, No. 4, 2001, pp. 345–383.
- <sup>19</sup>Hollingsworth, P. and Mavris, D., “Gaussian Process Meta-Modelling: Comparison of Gaussian Process Training Methods,” *AIAA 3<sup>rd</sup> Annual Aviation Technology, Integration and Operations (ATIO)*, 2003.
- <sup>20</sup>Lewin, G. and Haj-Hariri, H., “Reduced-Order Modelling of a Heaving Airfoil,” *AIAA Journal*, Vol. 43, No. 2, 2005, pp. 270–283.
- <sup>21</sup>Siegal, S., Cohen, K., Seidel, J., and McLaughlin, T., “Proper Orthogonal Decomposition Snapshot Selection for State Estimation of Feedback Controlled Flows,” *44th AIAA Aerospace Sciences Meeting and Exhibit*, 2006.
- <sup>22</sup>Hung, V. and Hien, T., “Modeling and Control of Physical Processes Using Proper Orthogonal Decomposition,” *Mathematical and Computer Modelling*, Vol. 33, No. 3, 2001, pp. 223–236.
- <sup>23</sup>Gross, A. and Fasel, H., “Control-Oriented Proper Orthogonal Decomposition Models for Unsteady Flows,” *AIAA Journal*, Vol. 45, No. 814-827, 2007, pp. 4.
- <sup>24</sup>Li, G., Li, M., Azarm, S., Rambo, J., and Joshi, Y., “Optimizing Thermal Design of Data Centre Cabinets With a New Multi-Objective Genetic Algorithm,” *Distributed and Parallel Databases*, Vol. 21, No. 2-3, 2007, pp. 167–192.
- <sup>25</sup>My-Ha, D., Lim, K., Khoo, B., and Willcox, K., “Real-time Optimisation Using Proper Orthogonal Decomposition: Free Surface Prediction Due to Underwater Bubble Dynamics,” *Computers & Fluids*, Vol. 36, No. 3, 2007, pp. 499–512.
- <sup>26</sup>LeGresley, P. and Alonso, J., “Airfoil Design Optimization Using Reduced Order Models Based on Proper Orthogonal Decomposition,” *Fluids 2000 Conference and Exhibit*, 2000.
- <sup>27</sup>Kim, Y., Rockwell, D., and Liakopoulos, A., “Vortex Buffeting of Aircraft Tail: Interpretation via Proper Orthogonal Decomposition,” *AIAA Journal*, Vol. 43, No. 3, 2005, pp. 550–559.
- <sup>28</sup>Sirovich, L., “Turbulence and Dynamics of Coherent Structures Part 1: Coherent Structures,” *Quarterly of Applied Mathematics*, Vol. 45, No. 3, 1987, pp. 561–571.
- <sup>29</sup>Lucia, D. and Beran, S., “Reduced-Order Model Development Using Proper Orthogonal Decomposition and Volterra Theory,” *AIAA Journal*, Vol. 42, No. 6, 2004, pp. 1181–1190.
- <sup>30</sup>Jolliffe, I., *Principle Component Analysis*, Springer, 2002.
- <sup>31</sup>Engineering Sciences Data Unit, *VGK Method for Two-Dimensional Airfoil Sections*, ESDU-96028, London, 1996.
- <sup>32</sup>Pound, G. and Price, A., *The Geodise OptionsMatlab Toolbox - A User's Guide*, <http://www.geodise.org/documentation/OptionsMatlab/html/index.htm>.
- <sup>33</sup>Sóbester, A., Leary, S., and Keane, A., “On the Design of Optimization Strategies Based on Global Response Surface Approximation Models,” *Journal of Global Optimization*, Vol. 33, No. 1, 2005, pp. 31–59.

<sup>34</sup>Toal, D., Bressloff, N., and Keane, A., “Kriging Hyperparameter Tuning Strategies,” *AIAA Journal*, Vol. 46, No. 5, 2008, pp. 1240–1252.

<sup>35</sup>Anderberg, M., *Cluster Analysis for Applications*, Academic Press, 1975.