Proper Orthogonal Decomposition & Kriging Strategies for Design

Mr. D.J.J. Toal, Dr. C. Holden, Dr. N.W. Bressloff & Prof. A.J. Keane

Introduction

The surrogate modelling approach to design optimisation aims to locate a global optimum with a significantly reduced number of function evaluations compared to, for example, a genetic algorithm. This is achieved through the construction of a model of the true design space based on an initial sampling. This model, which can be evaluated much faster than the true function, can then be searched and updated in regions of interest.

Of the many types of surrogate model, kriging is perhaps the most popular due to its ability to represent complex responses whilst providing an error estimate of the predictor. However, the expense of constructing a kriging model can be prohibitive for large data sets and at higher dimensions.

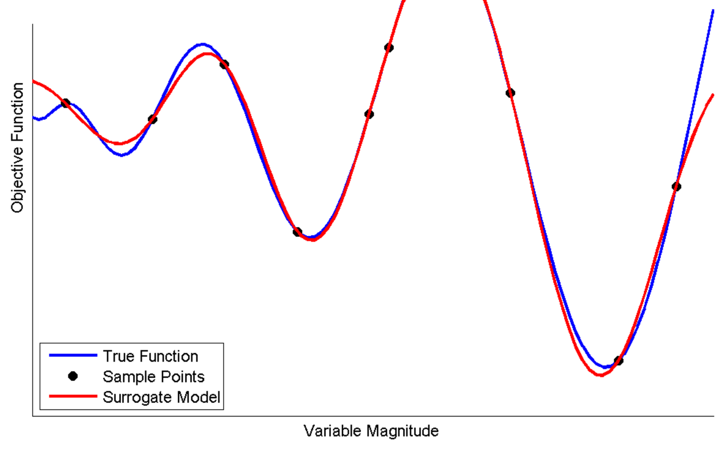

Fig 1: An example of a surrogate model

The aim of the following research is to therefore remove this obstacle to the application of kriging whilst developing a strategy for use at higher dimensions.

Hyperparameter Tuning Strategies

At its core, a kriging model consists of the following correlation function,

$$\exp\left(-\sum_{\ell=1}^{d} 10^{\theta_\ell}\|x_{i_\ell}-x_{j_\ell}\|^{p_\ell}\right).$$

The so called hyperparameters, θ and p, control the rate of correlation decrease and smoothness of the model, and must therefore be carefully selected, or tuned. Typically these hyperparameters are selected through an exhaustive global optimisation of the likelihood after every update to the model.

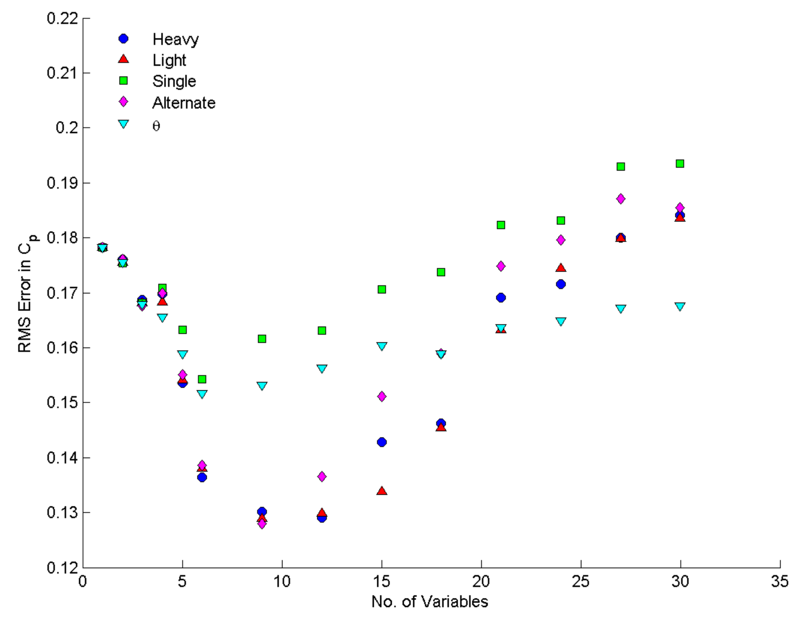

Fig 2: A comparison of five tuning strategies

The goal of an initial investigation was therefore to determine if these hyperparameters should be tuned after every update, and to what degree they should be tuned, in order for the model to remain effective. This investigation demonstrated that tuning the hyperparameters after every other set of updates reduced the overall tuning cost by 50% and had little impact on performance

Hybridised Likelihood Maximisation

Throughout the initial investigation into basic tuning strategies the method of tuning remained a constant, a genetic algorithm was followed by a hill climber to converge the solution. While this proves very effective it could be considered rather inefficient as it takes no account of the existence of an analytical gradient of the likelihood.

A novel hybrid tuning strategy was therefore developed which combined a particle swarm optimisation with an embedded local search to improve the rate of convergence. The particle swarm performs a global search while the local search, employing gradients, is used to accelerate convergence.

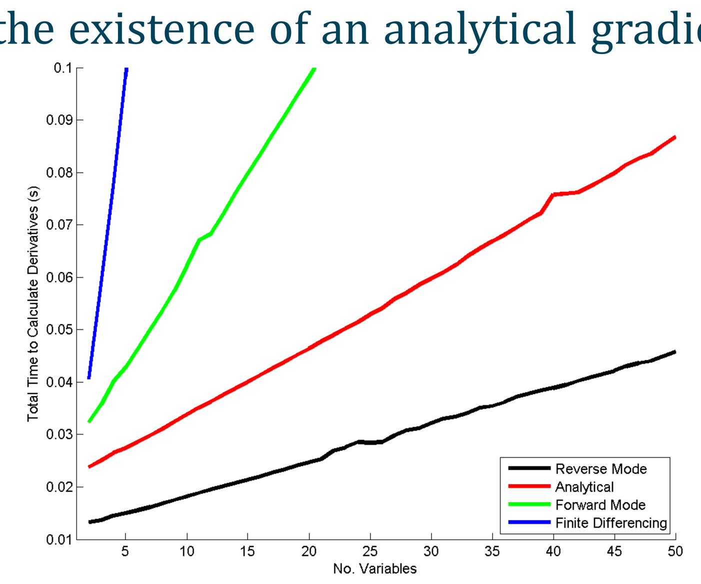

Fig 3: Comparative cost of gradient calculations

Rather than employing the analytical gradient, the local search uses a much more efficient reverse differentiation of the likelihood. This calculates the exact gradients but for a considerably reduced cost, therefore improving the overall performance of the algorithm.

Geometric Filtration Using POD

Surrogate models are more commonly used on problems with fewer than 15 variables. In an effort to improve their performance at higher dimensions given a limited evaluation budget, a new two stage optimisation strategy, termed geometric filtration, was developed. An initial kriging optimisation is performed which aims to indentify good designs. These good designs are then used to reparameterise the problem using proper orthogonal decomposition (POD). The new parameterisation has effectively filtered out the poor designs and reduced the dimensionality. A secondary optimisation then employs this new parameterisation.

This strategy was found to not only improve upon designs obtained via a basic kriging approach by 14% but also to significantly reduce the cost incurred through hyperparameter tuning by over 88%, reducing the optimisation time of an aerofoil from 22 hours to 2.5 hours.

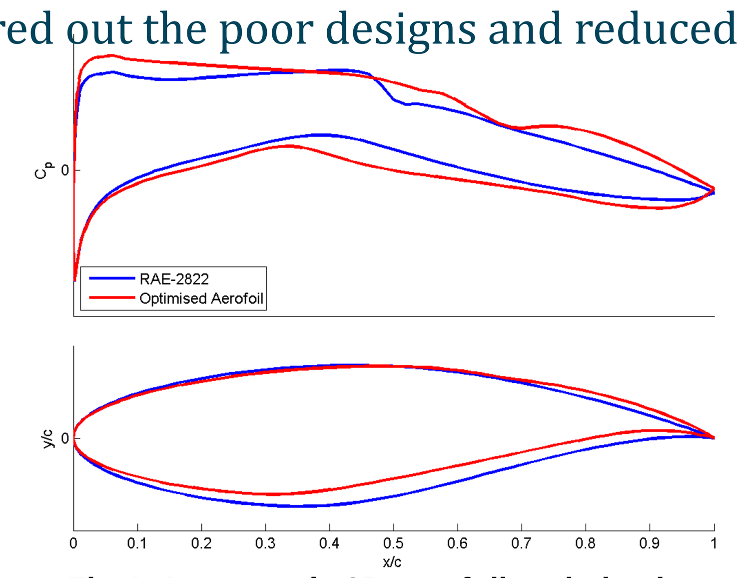

Fig 4: An example 2D aerofoil optimisation

Future Work

Although an efficient hybridised tuning strategy has been developed questions still remain as to how to effectively apply it to kriging models of varying dimensionality and with sampling plans of varying size. The objective is therefore to now provide the designer with guidance on the effective use of this hyperparameter tuning technique.

To date the geometric filtration strategy has been applied only to aerofoil design. The design optimisation of a transonic transport wing will therefore be considered. This design case will not only confirm the applicability of the strategy to 3D design problems but will also provide a demonstration of the overall impact of the other advances made over the course of the research.

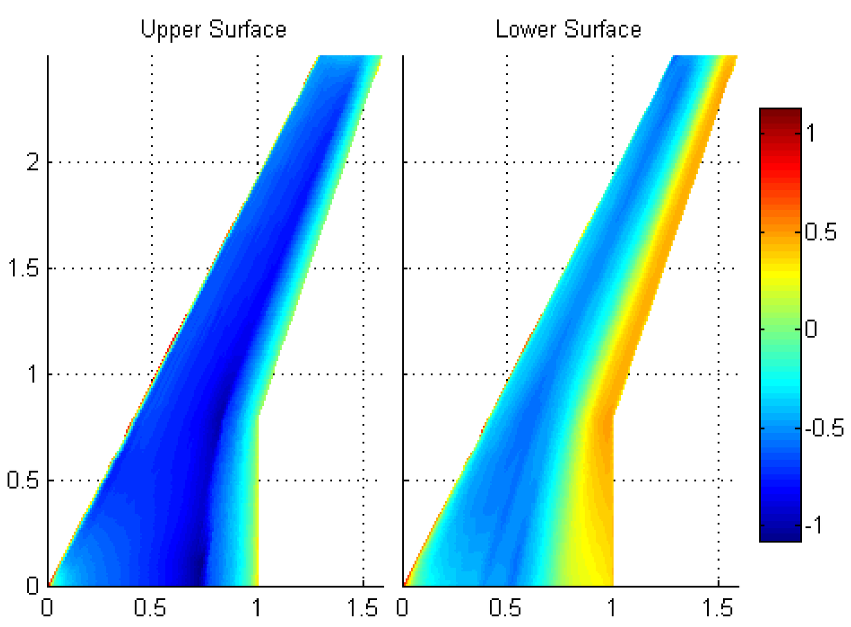

Fig 5: Pressure distribution over baseline 3D wing

Conclusions

The presented research goes some way to overcoming some of the problems of kriging based optimisations, namely the prohibitive tuning costs and the poor performance at higher dimensions.

Tuning the kriging hyperparameters after every other update to the model was found to have a negligible impact on the performance of the optimisation but reduced total tuning costs by approximately 50%.

An efficient novel hyperparameter tuning strategy has been developed employing a particle swarm hybridised with a local search which utilises an adjoint to obtain gradients of the likelihood function.

A novel two stage optimisation strategy utilising a POD based reparameterisation to reduce the dimensionality of the design space but maintain flexibility has also been presented. This strategy produced a 14% improvement in performance with a 88% reduction in tuning time over a traditional kriging optimisation.

Combining these strategies will therefore produce a marked improvement over traditional kriging allowing for more efficient optimisations within industry.

University of Southampton
School of Engineering Sciences

Email: djjt@soton.ac.uk
Computational Engineering & Design Group, University of Southampton, SO17 1BJ, UK

AIRBUS