

W. Song · A.J. Keane

An efficient evolutionary optimisation framework applied to turbine blade firtree root local profiles

Received: 3 June 2004 / Revised manuscript received: 5 August 2004 / Published online: 26 October 2004
© Springer-Verlag 2005

Abstract In this paper, an efficient evolutionary optimisation of a turbine blade firtree root local profile is presented. The firtree geometry is designed using an intelligent rule-based computer-aided design system (ICAD) and analysed using an industrial-strength finite element code. A large number of geometric and mechanical constraints drawn from past experience are incorporated in the design of the model. The high computational cost associated with finding optimal designs using high-fidelity codes is addressed using a surrogate-assisted genetic algorithm. The initial surrogate model is first built based on points sampled with a design-of-experiment method. A database of designs analysed using the high-fidelity code is built and augmented while the genetic algorithm progresses. In the procedure for deciding whether the high-fidelity code should be run, a simple 3σ principle is used instead of searching for the point with maximum expected improvement. This is combined with an appropriate ranking of the design points within the database. Some benchmark test problems are first used to illustrate the effectiveness and efficiency of the framework. When applied to the problem of local shape optimisation of a turbine blade firtree root, significant improvement is achieved using a limited computational budget.

Keywords Design · Optimisation · Stress analysis

1 Introduction

The concept of using approximations is not new in the field of structural optimisation; it can be traced back at least to 1974 when Schmit and Farshi (1974) published the concept of approximation techniques for structural synthesis. Since then, various approximation methods have been developed, such as the use of intermediate variables (for example, the reciprocal of the original design variables

(Vanderplaats 1999)), force approximations for stress constraints (Vanderplaats and Salajegheh 1989), and Rayleigh quotient approximations for frequency constraints (Canfield 1990), to name a few. These techniques provide an efficient approach to obtaining responses at a reduced computational cost. For example, using member forces as intermediate variables, the optimum design of an 18-bar truss can be achieved using only eight detailed analyses (Hansen and Vanderplaats 1990). However, these methods are difficult to incorporate into existing analysis capabilities and normally require problem-specific knowledge. It is even more challenging to adopt these methods for problems in the field of multidisciplinary optimisation. Therefore more general approximation schemes need to be developed to address the increasing cost associated with running full simulation codes during optimisation.

In recent years, there has been a growing interest in methods using approximation models in optimisation, in which the complex simulation codes are treated as black-box functions (Ahn et al. 2001; Jones et al. 1998; Simpson 1998; Venter et al. 1998). Data collected via evaluations of the expensive simulation codes at predefined design sites, often chosen by experimental design methods, are used in the construction of mathematical surrogates. Various techniques for the construction of approximation models have been proposed. Perhaps the most popular techniques involve a polynomial approximation created by least-squares curve fitting to a set of data. Another line of methods is based on interpolation techniques, which are believed to be more suitable for data collected from running deterministic simulation codes. Methods of this type include radial-basis functions (RBF) developed in the field of neural networks (Bishop 1995), stochastic-based methods, developed in the field of geostatistics, commonly referred to as Kriging as used by Sacks et al. (1989). An efficient procedure for global optimisation using a Kriging model was proposed by Jones et al. (1998), in which a branch-and-bound algorithm was applied to find the point with potential maximum improvement for re-sampling. Other techniques using classifier systems and the concept of space mapping have also been proposed.

W. Song (✉) · A.J. Keane
School of Engineering Sciences, University of Southampton, SO17 1BJ, United Kingdom
E-mail: w.song@soton.ac.uk, andy.keane@soton.ac.uk

An explorative comparison was made by Guinta and Watson (1998) between polynomial and interpolation models using test problems. A comparative study between neural networks and response surface models was provided by Daberkow and Marris (1998) using a preliminary aircraft design problem. Among various methods, RBF and Kriging were identified by Jin et al. (2000) as able to produce better results than other methods under multiple modelling criteria. Moreover, the Kriging method is statistically more meaningful and allows the possibility of computing error estimates for untried data points. One of the drawbacks of the Kriging method is the relatively high computational cost in estimating the hyper-parameters in the model, especially when large numbers of sample points are involved.

Apart from the many different methods for constructing surrogates, a number of different frameworks for the management of approximation models have also been proposed. It has been argued that any effective framework for managing the surrogates in optimisation is always associated with particular optimisation algorithms. Therefore, in principle, there are at least as many possible frameworks as optimisation algorithms. In general, these frameworks can be categorised according to the algorithms with which they are used. There are two broad classes of optimisation algorithms, gradient-based methods and non-gradient-based methods in which evolutionary algorithms and direct-search methods are among the most widely used. A rigorous framework was presented by Booker et al. (1999) for the use of surrogate models with direct-search methods. Frameworks based on trust regions and gradient-based search procedures have drawn most attention in the past few years (Alexandrov et al. 1997; Alexandrov and Lewis 2003; Guinta and Eldred 2000; Sellar et al. 2003; Wujek and Renaud 1998; Alexandrov 1998). One of the features possessed by this type of rigorous framework is that they all guarantee convergence to a model local optimum. However, these work with nonlinear programming techniques or direct-search methods. The number of studies on how surrogate models can be used with evolutionary algorithms is relatively limited. In these cases, the high computational cost associated with the successive application of evolutionary algorithms to complex, high-dimensional engineering problems is a known problem, as typical evolutionary algorithms require thousands of function evaluations to converge to near-optimal solutions. Therefore, the successive use of surrogate models is believed to be crucial to the practical application of evolutionary optimisation methods to complex engineering systems.

Several attempts have been made in the last several years to tackle the problem of using surrogate modelling with evolutionary search methods, particularly genetic algorithms (GAs). Robinson and Keane employed variable-fidelity analysis models and approximation techniques to improve the efficiency of the Evolutionary Strategy (ES) (Robinson and Keane 1999). A procedure using a number of successive single-point approximation models with a genetic algorithm (GA) was proposed by Nair and Keane (1998), in which some domain knowledge was employed to construct the surrogate models and a simple generation-delay approach was

used to control the use of exact models. Ratle proposed a simple local convergence criterion to decide when the exact model should be resorted to in a procedure integrating a genetic algorithm with Kriging models (Ratle 1998). However, this does not prevent the search from converging to false optima. A revised criterion was proposed by El-Beltagy and Keane (1999) for the similar synthesis between GA and Kriging models, where a gradually reduced tolerance was used to control the switch between surrogate and exact models for each individual in the population. However, the specification of criteria values in the above methods depends largely on the users and may not be appropriate for all problems. Nevertheless, the requirement for some sort of re-sampling was identified as necessary to overcome the inadequacy of the surrogate models.

A different type of effort was attempted by Liang et al. (2000), where a hybrid search procedure was formulated with the evolutionary search working on the quadratic response surface constructed from many local optima obtained from local searches. This method essentially reduced the difficulties in building a global surrogate model without changing the overall landscape of the exact functions. However, the quality of the final global surrogate model depends very strongly on the accuracy of the local search results and the use of evolutionary search methods in finding global optima in a simplified smooth landscape with few local optima may be questionable. Jin also proposed a framework for coupling Evolutionary Strategy (ES) and neural-network-based surrogate models (Jin et al. 2000). Two types of evolution control methods were presented to decide the frequency at which the exact model should be used. The common weakness in the above methods is that neither the historical search data nor the convergence properties of the evolutionary search method are fully utilised.

In this paper, a real-coded genetic algorithm was coupled with a Kriging surrogate model in order to reduce computational cost without sacrificing the ability of the GA in finding the global optimum for complex landscapes. Instead of using a simple generation-delay criteria and user-specified tolerance control parameters, a new approach based on the posterior variance estimate is used to suggest new sample points for re-evaluation using exact models. The new sample points are then inserted into an ordered database storing all the exact solutions evaluated so far, and the surrogate model is updated when these new points fall into the section of the data set used in the construction of the surrogate model. The following sections start with a description of the Kriging method, followed by detailed discussion of the proposed framework for incorporating a Kriging model into the genetic algorithm. Test functions are first used to illustrate the effectiveness of the proposed framework, followed by the application of the framework to a local shape optimisation problem.

2 Surrogate modelling

Let $Y(\mathbf{x})$ denote the true response of the system under study and $\mathbf{x} = (x_1, \dots, x_m)^T$ denote the vector of design or con-

trol variables. Sometimes the true response of the system can be represented in explicit mathematical form, but, in most cases, knowledge of the system is incomplete or the model is too complex to represent using explicit mathematical functions. Therefore a complex computer code is often used to simulate the relationship between the responses and inputs. Whatever the case, assume that observations can be made either through physical experiments or computer simulations at some values of the design variables. In this work, let us suppose that data has been collected at n points denoted by $\mathbf{x}^{(i)} = (x_1^{(i)}, \dots, x_m^{(i)})$, ($i = 1, \dots, n$) and the associated response is denoted by $y^{(i)} = Y(\mathbf{x}^{(i)})$. Let $y(\mathbf{x})$ represent the approximation model. The relationship between the true response and approximation can be represented as follows:

$$Y(\mathbf{x}) = y(\mathbf{x}) + \Delta(\mathbf{x}) \quad (1)$$

The difference between the true response and the approximated response, the total error or residual, is due to two types of errors: one is system error (bias error) denoted by $\epsilon(\mathbf{x})$, which exists because of the incompleteness of the models employed. The second is random error denoted by $\delta(\mathbf{x})$, which exists because of a number of reasons such as the effect of uncontrollable factors in physical experiments, discretization errors typically encountered in finite element analysis and computational fluid dynamics, and round-off errors, etc. However, numerical errors can usually be controlled so that the output of deterministic simulation codes can be regarded as deterministic. Therefore the difference between evaluated response and predictions here is entirely the system error, as described in the following equation:

$$Y(\mathbf{x}) = y(\mathbf{x}) + \epsilon(\mathbf{x}) \quad (2)$$

Consequently the normal distribution assumption commonly used for $\epsilon(\mathbf{x})$ is unjustifiable in this case. This partly explains why the least-squares model does not always provide a good approximation to a deterministic computer simulation code, where multiple evaluations for the same set of inputs will give the same results. Among various techniques for interpolating such data, radial basis function neural networks (RBF) and Kriging were identified by Jin et al. (2000) as being able to produce better results than other methods under multiple modelling criteria. The choice of Kriging techniques in this work is due to the fact that this method not only provides an estimate to the function values but also an estimate of posterior variance, which is used to control the frequency of re-sampling. A brief description of the Kriging model used is provided below and detailed discussions can be found in the work by Jones et al. (1998).

The Kriging model is typically expressed as

$$y(\mathbf{x}) = \beta + Z(\mathbf{x}) \quad (3)$$

where $Z(\mathbf{x})$ is a Gaussian random process with zero mean and variance σ^2 . The covariance matrix of $Z(\mathbf{x})$ is given by

$$\text{Cov}(Z(\mathbf{x}^i), Z(\mathbf{x}^j)) = \sigma^2 R(\mathbf{x}^i, \mathbf{x}^j) \quad (4)$$

where $R(\cdot, \cdot)$ is a correlation function between \mathbf{x}^i and \mathbf{x}^j . Different types of correlation function can be employed as given by Sacks et al. (1989), Guinta and Watson (1998). A commonly used type of correlation function can be expressed as

$$R(\mathbf{x}^i, \mathbf{x}^j) = \prod_{k=1}^m \exp(-\theta_k |x_k^i - x_k^j|^{p_k}) \quad (5)$$

where $\theta_k \geq 0$ and $1 < p < 2$ are the hyperparameters. The choice of $p_k = 2$ would provide enough flexibility for modelling smooth and highly nonlinear functions for most cases. The hyperparameters θ_k are estimated by maximising the log-likelihood function given by

$$-\frac{1}{2} \left[n \ln \sigma^2 + \ln |\mathbf{R}| + \frac{1}{\sigma^2} (\mathbf{y} - \mathbf{1}\beta)^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{1}\beta) \right] \quad (6)$$

where σ^2 and β can be derived using the following equations once the θ_k are given

$$\hat{\beta} = (\mathbf{1}^T \mathbf{R}^{-1} \mathbf{1})^{-1} \mathbf{1}^T \mathbf{R}^{-1} \mathbf{y} \quad (7)$$

$$\hat{\sigma}^2 = \frac{1}{n} (\mathbf{y} - \mathbf{1}\hat{\beta})^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{1}\hat{\beta}) \quad (8)$$

A numerical optimisation procedure is required to obtain the maximum likelihood estimates (MLE) of the hyperparameters. Once the hyperparameters are obtained from the training data, the function value at a new point \mathbf{x}^* can be predicted by

$$\hat{y}(\mathbf{x}^*) = \hat{\beta} + \mathbf{r}^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{1}\hat{\beta}) \quad (9)$$

along with the posterior variance, given by

$$s^2(\mathbf{x}^*) = \sigma^2 \left[1 - \mathbf{r} \mathbf{R}^{-1} \mathbf{r} + \frac{(1 - \mathbf{1}^T \mathbf{R}^{-1} \mathbf{r})^2}{(\mathbf{1}^T \mathbf{R}^{-1} \mathbf{1})} \right] \quad (10)$$

where $\mathbf{r}(\mathbf{x}^*) = R(\mathbf{x}^*, \mathbf{x}^1), \dots, R(\mathbf{x}^*, \mathbf{x}^n)$ is the correlation vector between the new point \mathbf{x}^* and the training dataset. This quantity provides a good indication on the accuracy of the prediction at new points and will be used in our framework to decide whether exact analyses need to be used. To obtain an estimate of the accuracy of the predictions of the Kriging model, a leave-one-out cross-validation procedure can be employed, as described in Jones et al. (1998). The measure used in such a procedure is called the 'standardised cross-validated residual' (SCVR) defined as

$$\text{SCVR}_j = \frac{y(\mathbf{x}^j) - \hat{y}_{-j}(\mathbf{x}^j)}{s_{-j}(\mathbf{x}^j)} \quad (11)$$

where $\hat{y}_{-j}(\mathbf{x}^j)$ and $s_{-j}(\mathbf{x}^j)$ denote the mean and variance computed by (9) and (10) without using the j th training data. A good predictor would mean that the Gaussian process

prior is appropriate for the dataset and majority of the estimates will be scattered in the interval $[-3, 3]$. Plotting the values of $SCVR_j$ against the predicted function values would also provide clues and solutions to problems that might exist in the model, for example, if there is any linear trend in such a plot, it is sometimes possible to improve the prediction by using log transformation.

Another useful concept is the expected improvement, see Jones et al. (1998), in which the point with maximum expected improvement is found by using a branch-and-bound algorithm followed by re-sampling at that point and the reconstruction of the surrogate. In this work, a simple 3σ principle is proposed and used instead of the maximum expected improvement in determining whether or not an evaluation using the exact model is necessary. This eliminates the need for an optimiser to find the point at which maximum expected improvement can be achieved. This principle is described below

Evaluate $Y(\mathbf{x}^*)$ when

$$\hat{y}(\mathbf{x}^*) - 3s(\mathbf{x}^*) < \frac{1}{q} \sum_{j=1}^q y_j^{best} \quad (\text{minimisation}) \quad (12)$$

where $\hat{y}(\mathbf{x}^*)$ and $s(\mathbf{x}^*)$ are computed using (9) and (10), y_j^{best} represent the q best solutions in the dataset. Note that the right-hand side reduces to y_{min} when $q = 1$ for minimisation problems. The use of this principle is discussed in the following section.

It should also be mentioned that the basic Kriging model could further be extended to include derivative information as reported by Morris et al. (1993) when derivatives are available either analytically or computed using automatic differentiation tools, adjoints, etc. The availability of efficient adjoint methods for sensitivity computations makes this expansion more attractive than ever for complex simulation codes.

3 A framework for managing surrogates within genetic algorithms

One of the main challenges in using surrogate models in global optimisation is to achieve a good balance between exploiting the surrogate and improving the accuracy of the surrogate. A commonly used strategy is to re-sample the point to which the surrogate converges under certain criteria and to reconstruct the surrogate using the augmented dataset. A revised methodology would be re-sampling the point at which the expected improvement is maximised. A common feature of these strategies is the use of an incrementally augmented dataset. In this work, a general framework for managing the surrogates with genetic algorithms is proposed by efficient and effective use of a simple 3σ principle and historical data. This framework is based on one of the basic features of the genetic algorithms, that is, as the search proceeds, the population tends to be filled with more and more similar solutions close to the optima. There are two features

that make this framework distinct from previously proposed ones:

- the use of a simple 3σ principle (12) eliminates the need for a search process for the point with maximum expected improvement, and;
- ranking of the historical dataset allows more efficient and effective use of the results obtained from exact and usually computationally expensive codes.

The proposed procedure is illustrated in Fig. 1. The algorithm starts with evaluations of the exact code on data points selected using design-of-experiment methods, for example, Latin hypercube sampling. The results are then archived in a central database according to the ranking of the fitness, and the top p or all of the solutions in the database are then used to construct the initial Kriging model. What follows is a typical genetic algorithm process with the fitness being evaluated first by surrogate model, then if the 3σ principle requires it, the exact expensive code is executed and the results inserted into their proper places in the database according to their fitness. Therefore the better solutions will come before the worse solutions and possibly into the top p of all solutions used in building the surrogate and will play a role in the update of the surrogate. The surrogate model is only updated when there has been a change in the solutions used to build the surrogate. It should be noted that the hyperparameters in the Kriging model are kept constant once found in the construction of the initial Kriging model. As mentioned earlier, the number of similar solutions will increase as the GA proceeds, so the hyperparameters will have a less important role in the later stages of the GA search. This will further expedite the search as the estimation of hyperparameters itself involves the solution of an optimisation problem, which can be time consuming. The update of the Kriging model will only involve the re-computation of the mean and variance using (8) and (9) based on the updated database. However, it would be beneficial and sometimes necessary to update the hyperparameters when more points are added into the dataset, especially if the response surfaces exhibit multimodel behavior.

The 3σ principle used is based on the fact that if the average fitness of the top q designs lies outside the interval $[\hat{y}(\mathbf{x}^*) - 3s(\mathbf{x}^*), \hat{y}(\mathbf{x}^*) + 3s(\mathbf{x}^*)]$ computed at point \mathbf{x}^* , the probability of producing a better design at point \mathbf{x}^* is very small. Two control parameters p and q are used to specify the number of design points to build the surrogate and number of design points used in the right hand of the 3σ principle. It is not difficult to understand the effect of these two parameters: increasing the value of parameter p will cause more points to be used in building the surrogate model, and increasing the value of parameter q will lead to more new points falling into the p interval and therefore more exact evaluations.

The real-coded genetic algorithm used in this work is derived from the basic classes available from the C++ class library GAlib, developed by Matthew (1999). Instead of using binary strings to represent real values, as in most common implementations of genetic algorithms, real coding offers a direct way of representing solutions in numerical opti-

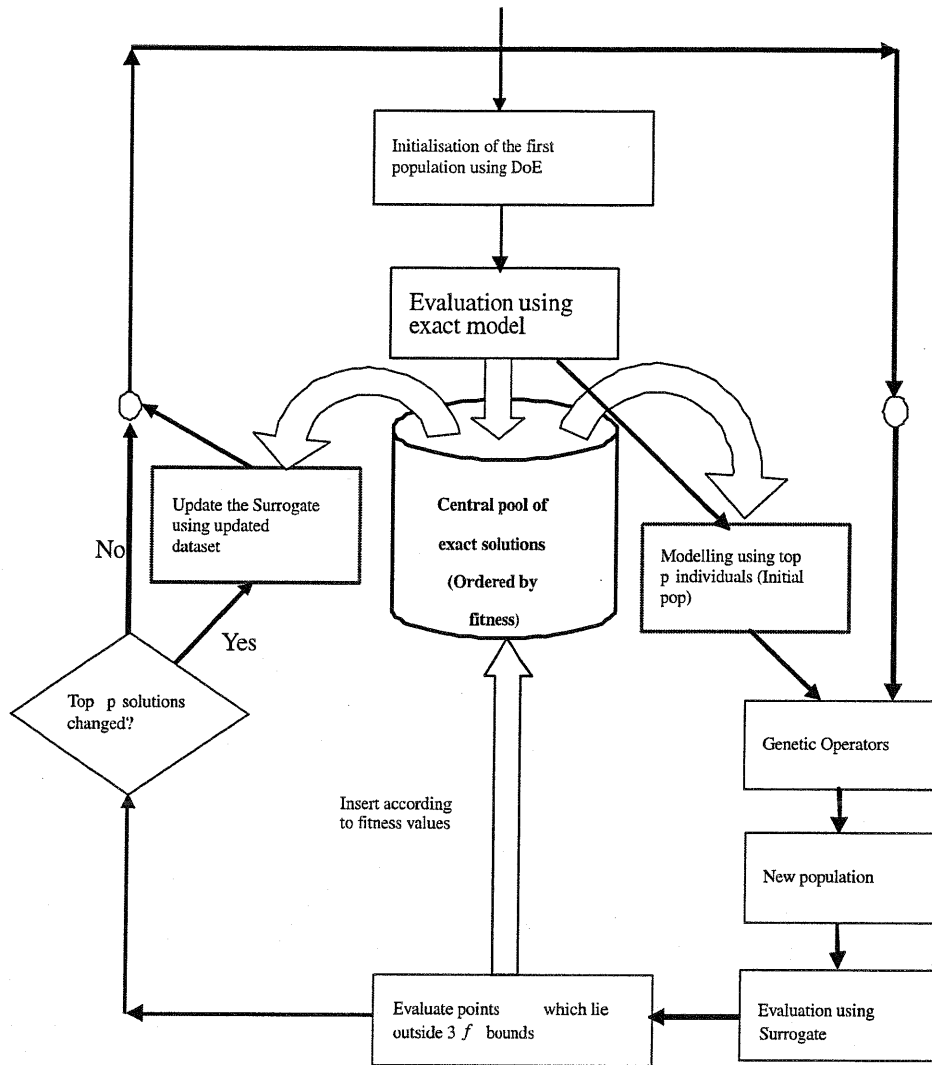


Fig. 1 Procedure for incorporating Kriging model into genetic algorithms

misation problems. Each chromosome is an array of real numbers, e.g., the design variables. The use of real coding eliminates the need for coding and decoding processes. Non-uniform crossover and random mutation operators are used in the current implementation of the GA and are briefly described below.

Let $\mathbf{x}^i = \{x_1^i, x_2^i, \dots, x_m^i\}$ and $\mathbf{x}^j = \{x_1^j, x_2^j, \dots, x_m^j\}$ represent two designs in the parent population; the child \mathbf{x}^c is generated by a BLX crossover, which is defined as

$$x_k^c = (1 - \alpha)x_k^i + \alpha x_k^j, \quad k = 1, \dots, m \quad (13)$$

Non-uniform mutation is defined as

$$x_k^m = x_k^l + \beta (x_k^u - x_k^l), \quad k = 1, \dots, m \quad (14)$$

where α and β are random numbers between 0 and 1. Finally, it should be pointed out that duplicate points are not archived to avoid problems in the computation of the inverse of the correlation matrix. Also any type of implementation of the genetic algorithms could be used in this general framework for exploiting surrogate modelling.

4 Experiments on benchmark test functions

Two commonly used benchmark test functions are used here to test the effectiveness of the proposed framework. The first is an unconstrained 20D Rastrigin function as defined by (15) and the second is the constrained 20D Bump function introduced by Keane, which is given by (16). Both of these two functions have a large number of local minima and are usually difficult for gradient-based optimisers to find global optima.

$$F_{\text{Rastrigin}} = (10n) + \sum_{i=1}^n (x_i^2 - 10 \cos(2\pi x_i)) \quad (15)$$

$$F_{\text{Bump}} = \frac{\text{abs} \left(\sum_{i=1}^n \cos^4(x_i) - 2 \prod_{i=1}^n \cos^2(x_i) \right)}{\sqrt{\sum_{i=1}^n i x_i^2}} \quad (16)$$

$$\text{subject to } \prod_{i=1}^n x_i > 0.75 \text{ and } \sum_{i=1}^n x_i < \frac{15n}{2}$$

The bounds of variables for these two problems are $-5.12 < x_i < 5.12$ and $0 < x_i < 10$, respectively. The average convergence curves against the number of generations are shown in Fig. 2 for the Rastrigin function and in Fig. 3 for the Bump function. It can be seen that the number of exact analyses has been reduced to only one third of the original number that would be required to obtain similar results using a direct optimisation approach. The standardised cross-validated residual is also shown in Fig. 4 for the initial Kriging surrogate for the Rastrigin problem. It is shown that a reasonably good approximation has been obtained. The proposed

framework is next applied to the local shape optimisation of the firtree problem.

5 Local profile optimisation of turbine blade firtree root

In this section, a turbine blade firtree root local notch profile optimisation problem (Song et al. 2002) using cubic non-uniform rational B-splines (NURBS) is chosen to illustrate the effectiveness of the proposed framework. This notch is part of the geometry at the base of a turbine blade that is used to attach the blade to the rotating disc and is a region where high contact stresses occur. Here the firtree geometry has been modelled using the rule-based design capability




Fig. 2 Genetic algorithms with surrogate modelling on 20D Rastrigin function




Fig. 3 Genetic algorithms with surrogate modelling on 20D Bump function




Fig. 4 Standard cross-validation residuals of initial kriging model for Rastrigin function




Fig. 5 NURBS representation of double-arc fillet using seven control points and its defining coordinates

provided by ICAD. The notch geometry is modelled using NURBS, rather than circular arcs, which gives more flexibility in modelling local shape variations. However, the use of NURBS introduces more design variables, thus increasing the cost of finding an optimum design. The definition of the local profile is illustrated in Fig. 5. Finite element analysis is carried out to evaluate the peak stress at the notch region and is used as the objective function in the optimisation. Results are shown in Fig. 6, in which two convergence curves are shown, one without using the surrogate model, and the other with. It can be seen that comparable results can be obtained with only one third of the number of exact function evaluations are used. Comparison between the original and optimised




Fig. 6 Genetic algorithms with surrogate modelling on local profile optimisation of fir tree root




Fig. 7 Comparison between results of GA search without and with surrogates

geometry is shown in Fig. 7. The worst principal stress at point B (where the maximum stress occurs) is reduced from 745.4 MPa to 685.6 MPa – a reduction of over 8% without worsening the stresses at point A, using around one third of the computing effort which would otherwise be required using direct GA search.

6 Conclusion

An efficient evolutionary optimisation process for improving the local profile of a turbine blade firtree root is proposed using a surrogate-assisted evolutionary framework. The geometry of the firtree root is designed using a rule-based computer-aided design system (ICAD) with the local profile defined using NURBS. The resulting model is analysed using an industrial-strength finite element code. Due to the high computational cost associated with finding the optimum using such high-fidelity codes, a surrogate-assisted genetic algorithm is applied to the problem. The initial surrogate model is first built based on points sampled with design-of-experiment methods. A database of designs evaluated using the high-fidelity code is built and augmented while the genetic algorithm progresses. In the procedure for deciding whether the high-fidelity code should be applied, a simple 3σ principle is used instead of searching for the point with maximum expected improvement. This is combined with appropriate ranking of the design points within the database. Some benchmark test problems are used to illustrate the effectiveness and efficiency of this framework. When applied to the problem of local shape optimisation of the firtree root, significant improvements are achieved under a reduced computational budget.

Acknowledgement This work was funded by EPSRC under grant GR/M53158 and the University Technology Partnership for Design, which is a collaboration between Rolls-Royce, BAE SYSTEMS and the Universities of Cambridge, Sheffield and Southampton.

References

- Ahn JA, Kim H, Lee D, Rho O (2001) Response surface method for airfoil design in transonic flow. *J Aircr* 38(2)
- Alexandrov NM (1998) On managing the use of surrogates in general nonlinear optimization and MDO. AIAA-98-4798
- Alexandrov NM, Dennis JE Jr, Lewis RM (1997) A trust region framework for managing the use of approximation models in approximation. NASA/CR-201745
- Alexandrov NM, Lewis RM (2003) First-order frameworks for managing models in engineering optimisation. 1st international workshop on surrogate modelling and space mapping for engineering optimisation, 11/16-19/2000, TDU
- Bishop C (1995) Neural networks for pattern recognition. Oxford University Press
- Booker AJ, Dennis JE, Frank PD, Serafini DB, Torczon V, Trosset MW (1999) A rigorous framework for optimization of expensive functions by surrogates. *Struct Optim* 17(1):1–13
- Canfield RA (1990) High-quality approximation of eigenvalues in structural optimisation. *AIAA J* 28(6):1116–1122
- Daberkow DD, Marris DN (1998) New approaches to conceptual and preliminary aircraft design: a comparative assessment of a neural network formulation and a response surface methodology, AIAA. World Aviation Conference, September 28–30, 1998, Anaheim, CA
- El-Beltagy MA, Keane AJ (1999) Evolutionary optimisation for computationally expensive problems using gaussian processes. Proceedings of the Genetic and Evolutionary Computation Conference (GECCO99), Morgan Kaufman, pp 196–203
- Guinta AA, Eldred MS (2000) Implementation of a trust region model management strategy in the dakota optimisation toolkit. AIAA-2000-4935
- Guinta AA, Watson LT (1998) A comparison of approximation modelling techniques: polynomial versus interpolating models. AIAA-98-4758
- Hansen SR, Vanderplaats GN (1990) An approximation method for configuration optimisation of trusses. *AIAA J* 28(1):161–172
- Jin R, Chen W, Simpson TW (2000) Comparative studies of meta-modelling techniques under multiple modelling criteria. AIAA-2000-4801
- Jin Y, Olhofer M, Sendhoff B (2000) A framework for evolutionary optimisation with approximate fitness functions. *IEEE Transactions on Evolutionary Computation*
- Jones DR, Schinlau M, Welch WJ (1998) Efficient global optimisation of expensive black-box functions. *J Glob Optim* 13:455–492

- Liang KH, Yao X, Newton C (2000) Evolutionary search of approximated n-dimensional landscapes. *Int J Knowl Based Intell Eng Systems* 4(3):172–183
- Matthew W (1999) GALib: a c++ library of genetic algorithm components. <http://www.lancet.mit.edu/ga/>
- Morris MD, Mitchell TJ, Ylvisaker D (1993) Bayesian design and analysis of computer experiments: use of derivatives in surface prediction. *Technometrics* 35:243–255
- Nair PB, Keane AJ (1998) Combining approximation concepts with genetic algorithm-based structure optimisation procedure
- Ratle A (1998) Accelerating the convergence of evolutionary algorithms by fitness landscape approximation. *Parallel problem solving from Nature V*, 87–96
- Robinson GM, Keane AJ (1999) A case for multi-level optimisation in aeronautical design *Aeronaut J* 103:481–485
- Sacks J, Welch WJ, Mitchell JJ, Wynn HP (1989) Design and analysis of computer experiments. *Stat Sci* 4(4):409–435
- Schmit LA, Farshi B (1974) Some approximation concepts for structural synthesis. *AIAA J* 12(5):692–699
- Sellar RS, Batill SM, Renaud JE (2003) Response surface based. Concurrent subspace optimisation for multidisciplinary system design
- Simpson TW (1998) Comparison of response surface and kriging models in the multidisciplinary design of an aerospike nozzle. NASA/CR-1998-206935, ICASE report No. 98-16, 1998
- Song W, Keane AJ, Rees J, Bhaskar A, Bagnall S (2002) Local shape optimisation of a fir tree root using NURBS. 9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, Atlanta, Georgia 4–6 Sep 2002
- Vanderplaats GN (1999) Structural design optimisation status and direction. *J Aircr* 36(1)
- Vanderplaats GN, Salajegheh E (1989) A new approximation method for stress constraints in structural synthesis. *AIAA J* 27(3):352–358
- Venter G, Haftka RT, Starners JH Jr (1998) Construction of response surface approximations for design optimisation. *AIAA J* 36(12)
- Wujek BA, Renaud JE (1998) New adaptive move-limit management strategy for approximate optimization, part 1. *AIAA J* 36(10):1911–1921

