9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization
4-6 September 2002, Atlanta, Georgia

AIAA 2002-5540

SCREENING AND APPROXIMATION METHODS FOR EFFICIENT STRUCTURAL OPTIMIZATION

Stephen J. Leary*
Atul Bhaskar†
Andy J. Keane‡
University of Southampton, Highfield, Southampton SO17 1BJ, U.K

Abstract

In this paper we discuss two statistical techniques for achieving computational economy during the optimization process. The first, the use of approximation methods is often applied when optimizing expensive computational models of complex engineering systems: the idea is to replace the expensive analysis code by a cheap surrogate model for the purposes of optimization. There are many approximation methods available in the literature, we focus here on kriging. The second, screening experiments, has received much attention in the statistics community. This statistical tool has been applied to the problem of structural optimization here. Indeed, one purpose of this paper is to increase awareness of these tools in the structural optimization community. In particular, a focus here is on screening multiple responses, as a structural optimization problem typically requires optimization of at least one objective subject to at least one constraint. Finally, both approaches are combined in order to provide an algorithm which appears very efficient for large dimensional (>10) structural optimization problems. A structural optimization case study of industrial interest demonstrates the approach.

*Research Fellow, Computational Engineering and Design Center, School of Engineering Sciences, S.J.Leary@soton.ac.uk.
†Senior Lecturer, Computational Engineering and Design Center, School of Engineering Sciences, A.Bhaskar@soton.ac.uk.
‡Professor, Computational Engineering and Design Center, School of Engineering Sciences, Andy.Keane@soton.ac.uk.

1 Introduction

In an industrial setting the role of a designer is far from trivial. He or she is under continual pressure to produce designs that are better than previous variants, for instance: this might mean designing a lighter and/or less costly structural component. The resultant design must also meet stringent requirements, e.g., meeting stress or displacement requirements. It may also have to satisfy multidisciplinary design criteria; an aero-engine component designed by a structural engineer may also have to satisfy certain aerodynamic criteria. Not only is a designer under pressure to produce such a design, he or she is expected to produce it in ever shorter timescales. We note that most scope for design improvements is during the preliminary design phase, where less sophisticated models are invoked for the analysis of a design. During final design only one analysis of the design may be possible so it is essential that during the preliminary design phase a good design satisfying all necessary constraints is chosen.

The idea of producing a better design leads naturally to the idea of optimization. There are however problems: even at the preliminary design stage a finite element model or a CFD analysis (depending upon the problem) would typically be used for evaluating the design. Even at this early stage the cost of running the computational model often proves to be a bottleneck to the direct use of optimization methods. An alternative strategy is required.

To ameliorate these difficulties two statistical techniques are discussed. The first of these uses approximation methods and the second is a statistical screening study.

Perhaps the most common way of tackling the problem of expensive function optimization is through the use of approximations to the expensive model. Response surface methods (see for example Myers and Montgomery) seek polynomial approximations to the function.[9] These models, once constructed, provide a cheap means of approximating the expensive function/model. Expensive function evaluation typically means an expensive computer run, e.g. of an FE or a CFD code. An elegant optimization algorithm employing kriging as the approximation method was described in Jones *et al.*[4] Other approximation methods include the use of neural networks and radial basis functions.[14,10] Of course, this list is far from exhaustive: there are a plethora of approximation methods in the literature, each with its good and bad points. A full review of this subject area is well beyond the scope of this paper and we consider one approach, kriging, in some detail. We note however that all these methods are, in effect, forms of curve fitting built using selective runs of the high fidelity model.

One potential problem with the above approach is the accuracy of the approximation given the limited information available. This is particularly the case in higher dimensions so a screening strategy to overcome this burden is demonstrated.

The idea of a screening study is to obtain as much information on how the input variables affect the response using as few runs of the high fidelity code as possible. The motivation here is that it may often be the case that a few variables are responsible for most of the variation in the response whereas most of the other variables contribute little. *Screening experiments are used to find out which variables affect the response the most.* Once again, there are a multitude of screening studies available in the literature. Amongst others Myers and Montgomery discuss screening in the context of polynomial response surface approximations.[9] Welch *et al.* suggest an algorithm using the kriging hyperparameters to suggest important variables.[13] Elster and Neumaier consider edge designs formed from conference matrices whilst Clarich *et al.* use t-tests applied to a latin hypercube design.[2,1]

Again, this is not an exhaustive list: it is by no means obvious which is the best screening strategy to use and when. The purpose here is simply to show how such stategies can lead to efficiency gains in the optimization process. In this paper we opt to use Edge designs for our screening study.[2] With this approach we know in advance how many runs of the high fidelity code are required for the screening study: usually twice the dimension of the problem. Also the algorithm is easy to implement and provides a model independent estimate of the importance of each variable.

We note further that it is natural to consider a combination of screening studies and approximation methods to further increase the efficiency of the optimization process.

The rest of this paper is laid out as follows: Section 2 introduces approximation methods (in particular, we focus on kriging). Section 3 briefly describes some screening strategies then considers one in detail, Edge designs formed from conference matrices, an algorithm we found very easy to implement. Section 4 discusses some potential optimization strategies and in section 5 an industrial case study is performed to demonstrate the approach. Finally, in section 6 some conclusions are drawn.

2 Approximation Methods

Before we build an approximation we require a systematic means of selecting the set of inputs (called a design of experiments, or DoE in short) at which to perform a computational analysis.

In k dimensions the 2^k vertices formed by the upper and lower bounds on each design variable form the design bounding box within which the experimental design is created. The idea with DoE is to in some sense fill this design space with a limited number of points. As a result the algorithms are often referred to as "space filling" designs.

Simple experimental designs include 2^k factorial designs which are created by specifying each design variable at 2 levels, the upper and lower bounds on each variable (this design considers every vertex of the design bounding box). 3^k factorial designs additionally include the midpoint of each input.

These experimental designs prove expensive (particularly for large k) so fractional factorial designs or, alternatively, D-optimal designs could be considered. Many examples of these approaches exist in the response surface literature; the interested reader should consult Myers and Montgomery for further details.[9]

One popular choice for generating an experimental design for computational experiments is the Latin Hypercube.[7] Such designs however are not guaranteed to have good "space filling" properties. As a result optimum latin hypercubes can be considered. These are latin hypercube designs that achieve optimality in some space filling sense (e.g., a maximin distance criterion).[3] In this paper our experimental designs are formed from optimal latin hypercube designs using this criterion. An algorithm for generating such designs can be found in Morris

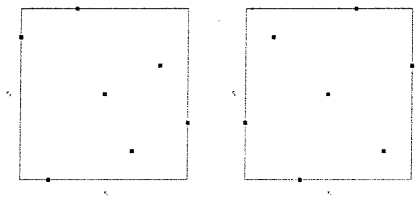

Figure 1: Latin hypercube and optimal latin hypercube design.

and Mitchell.[8] Examples of a latin hypercube design and an optimal latin hypercube design using 7 points with two independent variables (inputs) x_1 and x_2 are shown in figure 1.

Once we have chosen a suitable DoE and evaluated the high fidelity model at this set of inputs we can construct an approximation.

In a typical approximation method the relationship between observations (responses) and independent variables is expressed as

$$y = f(\mathbf{x}) \tag{1}$$

where y is the observed response, $\mathbf{x}$ is a vector of k independent variables

$$\mathbf{x} = [x_1, x_2, ..., x_k] \tag{2}$$

and $f(\mathbf{x})$ is some unknown function. An approximation to this response

$$\hat{y} = \hat{f}(\mathbf{x}), \tag{3}$$

is sought.

Perhaps the simplest approximation model is a polynomial response surface.[9] Here the response of interest is replaced by a low order polynomial approximation. This is usually implemented as a local model defined in a specific region of the design space, namely around the current best design. The model is only assumed valid in a small neighbourhood of the current best design and the optimization algorithm proceeds using a move limit or a trust region strategy. Whenever the boundary of the trust region is reached the model is updated and the process repeated.

Global approximations try to capture the behaviour of the response over the entire domain of interest. Many approximations could be considered. Amongst these are artificial neural networks, radial basis functions and kriging.[14,10,4] In this paper we focus on kriging; a brief description follows.

Kriging attempts to model some non-linear functional relationship $y = f(\mathbf{x})$. Given a set of N training data $[\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, ..., \mathbf{x}^{(N)}]$ the kriging model can be used to make a prediction $\hat{y} = \hat{f}(\mathbf{x})$ at untested points $\mathbf{x}$ in the design space.

A correlation matrix of the training data

$$\mathbf{R}(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}) = \exp[-d(\mathbf{x}^{(i)}, \mathbf{x}^{(j)})] \tag{4}$$

is first sought where d is some distance measure. For example

$$d(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}) = \sum_{h=1}^{k} \theta_h |x_h^{(i)} - x_h^{(j)}|^{p_h} \tag{5}$$

where $\theta_h \geq 0$ and $1 \leq p_h \leq 2$ are some as yet undetermined parameters. Note in the following we fix $p_h = 2$ as this has been found by experience to be a good choice..

When we wish to sample at a new point $\mathbf{x}$, we form a vector of correlations between the new points and the training data

$$\mathbf{r}(\mathbf{x}) = \mathbf{R}(\mathbf{x}, \mathbf{x}^{(i)}) = [\mathbf{R}(\mathbf{x}, \mathbf{x}^{(1)}), ..., \mathbf{R}(\mathbf{x}, \mathbf{x}^{(N)})]. \tag{6}$$

The prediction is then given by

$$\hat{y}(\mathbf{x}) = \mu + \mathbf{r}^T \mathbf{R}^{-1}(\mathbf{y} - \mathbf{1}\mu). \tag{7}$$

The mean and variance of the prediction are

$$\mu = \frac{\mathbf{1}^T \mathbf{R}^{-1} \mathbf{y}}{\mathbf{1}^T \mathbf{R}^{-1} \mathbf{1}} \tag{8}$$

$$\sigma^2 = \frac{(\mathbf{y} - \mathbf{1}\mu)^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{1}\mu)}{N}. \tag{9}$$

The parameters θ_h and p_h are determined by maximizing the likelihood

$$\frac{1}{(2\pi)^{\frac{N}{2}} (\sigma^2)^{\frac{N}{2}} |\mathbf{R}|^{\frac{1}{2}}} \exp\left[\frac{-(\mathbf{y} - \mathbf{1}\mu)^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{1}\mu)}{2\sigma^2}\right] \tag{10}$$

of the sample.

Another useful quantity here is the mean squared error of the prediction defined by

$$s^2(\mathbf{x}) = \sigma^2 \left[1 - \mathbf{r}^T \mathbf{R}^{-1} \mathbf{r} + \frac{(1 - \mathbf{1}^T \mathbf{R}^{-1} \mathbf{r})^2}{\mathbf{1}^T \mathbf{R}^{-1} \mathbf{1}}\right], \tag{11}$$

this quantity allows us to asses the accuracy of the prediction. Note that as set out here kriging strictly interpolates the training data.

This technique will be used later when considering increasing the efficiency of optimization in our industrial case study.

3 Screening Methods

The purpose of a screening study is to rank the design variables according to how they affect the response. Minimizing the number of expensive computer simulations whilst maximizing the obtained information is the ultimate goal. The more important variables can then be altered in an optimization problem of reduced dimension. We note that screening will also allow approximations to be constructed with less training data.

Myers and Montgomery lists a set of screening experiments using factorial designs.[9] As highlighted earlier, these can be computationally very expensive, requiring many runs of the high fidelity model. This is particularly the case in higher dimensions. As a result, more efficient screening designs are considered. These include fractional factorial and Plackett-Burman designs.

One approach when considering a large number of design variables is group screening.[6] Here certain variables are grouped together and a screening experiment is performed on the groups. Further screening experiments can be performed on any group which has a significant influence on the response. We note here that the choice of groupings of variables is important to the success of this approach. If there is little information available on the response of interest then it is by no means obvious how to do this.

One factor at a time studies could also be considered. These designs are computationally very efficient but only allow the influence of each variable to be determined in a small neighbourhood of the design space. This is a local sensitivity study using $k+1$ high fidelity model evaluations for a k dimensional problem.

In Elster and Neumaier an efficient strategy based on the one factor at a time idea but taking a global view of the design space is described.[2] The cost of generating the screening data is of the order of $2k$ runs. The response is evaluated at certain edges of the design bounding box and differences along these edges are used to estimate each variable's importance. These edges turn out to be optimally (in the sense of a distance criterion) placed so maximal information on the behaviour of the response should be obtained. The approach provides a model independent estimate of the importance of each variable. Due to these appealing properties, this method will shortly be described in a little more detail and used in later sections for performing the screening study on the industrial case study.

In the design and analysis of computer experiments literature, Welch *et al.* describe an approach using the model hyperparameters to screen out variables of little importance.[13] Firstly the response is evaluated at an N point DoE (latin hypercube here). The choice of N is based on a heuristic argument. In terms of our kriging model defined in section 2, the likelihood function in k dimensions depends on k hyperparameters. These are first set to a constant and the optimal value of this constant is chosen to maximize the likelihood. Then the hyperparameters are varied one at a time to see which most increases the likelihood. The process is repeated thus highlighting the most important variables.

Finally, in a recent case study Clarich *et al.* the response was again evaluated at an N point latin hypercube design.[1] The student parameter was used to determine the most important variables.[11] This information was used to reduce the dimension of the problem when optimizing an axial compressor.

In general it is not obvious which is the best screening strategy to use - ultimately it will be problem dependent. In the case of screening studies a comparison of certain techniques can be found in Trocine and Malone.[12]

We now turn our attention to the construction of edge designs. These have been built here using Elster and Neumaier and the references therein.[2] One construction when k is an odd prime is to consider

$$S_{ij} = \begin{cases} 0 & \text{if } i = j \\ 1 & \text{if } j - i \in R\text{mod}(k) \\ -1 & \text{otherwise} \end{cases} \tag{12}$$

where R is the set of quadratic residues mod(k) (see Elster and Neumaier for a table of quadratic residues for $k \leq 50$).[2] Of course, this approach is not restricted to odd prime dimensions. In general if our dimension is k and the next odd prime is k^* we simply introduce k_1 dummy variables (which have no influence on the response) such that $k + k_1 = k^*$. Then apply the above.

The design matrix (on $[-1, 1]^k$) is then defined as

$$X = \begin{bmatrix} \mathbf{1} & \mathbf{S}+\mathbf{I} \\ \mathbf{1} & \mathbf{S}-\mathbf{I} \end{bmatrix} \tag{13}$$

where $\mathbf{1}$ is a $k \times 1$ vector of 1's. The difference over the edges can then be used to determine the most influential variables. See Elster and Neumaier for a simple example.[2] In the case when $k = 3$ the edges are shown in figure 2. We now go on to discuss various possible optimization strategies.

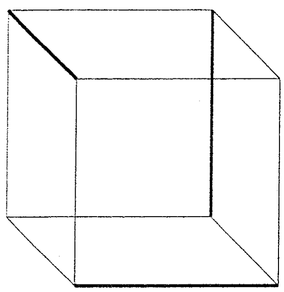

Figure 2: Edge design for $k = 3$.

4 Optimization Strategies

We first consider the use of ideas described in sections 2 and 3 to a problem involving a single objective and no constraints.

The most obvious strategy when considering approximation methods is given in figure 3. Here a de-

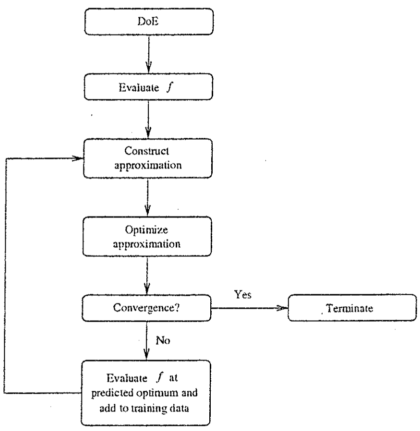

Figure 3: Algorithm using approximation methods.

sign of experiments is chosen and the expensive analysis code is run for this selection of inputs. An approximation is then constructed, this then replaces the expensive analysis code for the purposes of optimization.

Once an optimum on the surrogate surface is found, we evaluate it by making an extra run of the original analysis code. We can then check whether the design satisfies any convergence criteria and if so, terminate. If not, then we can add the predicted optimum to the training data, reconstruct the approximation and repeat this process. This can be continued until either some convergence criterion is achieved or the maximum number of allowable calls to the high fidelity model is reached.

One potential problem here is that this algorithm may locate and become trapped in a local minimum. Using a kriging model we could replace the step "optimize the approximation" with the step "optimize the expected improvement", see Jones *et al.* for details on expected improvement.[4] This makes for a sophisticated optimization strategy which balances the need for a good solution with our uncertainty in the model. This information is available to us since kriging gives us an estimate of the error in prediction via equation (11). The advantage of the expected improvement approach is that it will not become trapped in a local minimum, it's disadvantage is the extra cost in terms of calls to the high fidelity model. We opt for optimizing the approximation in the example in section 5 but mention expected improvement as an approach when the objective is multimodal.

In terms of screening strategies, the most obvious approach is to use the screening strategy to highlight important variables, then perform optimization on the reduced dimensional problem (see figure 4). The

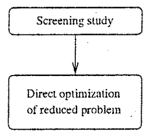

Figure 4: Algorithm using a screening strategy.

idea is that the reduced problem will require many fewer calls to the expensive model than a full direct optimization. This approach is considered in, for example, Clarich *et al.*[1]

A combination of both approaches appears natural. We propose this in the manner shown in figure 5. We could again consider expected improvement in place of direct optimization of the approximation surface. Note that the use of a screening study can significantly reduce the number of training data required to build a model of sufficient accuracy, and the use of approximation methods can significantly reduce the cost of optimization, so a dual approach is appealing. Moreover the results from the screening study may be used to help build the approximation.

When we consider multiple objectives subject to multiple constraints there are various strategies we could consider. We here consider one objective subject to one constraint only, the ideas generalize easily to more than one of both.

We first screen the objective and find the important variables (the ones that affect the objective the most) and label these as

$$v_1^{(o)}, v_2^{(o)}, \ldots, v_{\alpha_1}^{(o)} \qquad \alpha_1 << k. \qquad (14)$$

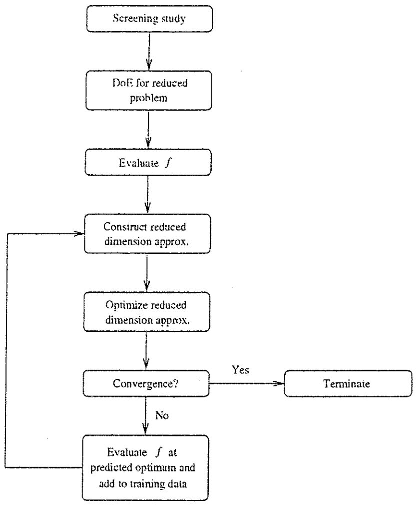

Figure 5: Algorithm using approximation methods and a screening strategy.

We then screen for the constraint and find the variables which have most influence on the constraint. We label these as

$$v_1^{(c)}, v_2^{(c)}, ..., v_{\alpha_2}^{(c)} \qquad \alpha_2 << k. \tag{15}$$

There may be some variables common to both, so the total number of variables will be

$$\alpha \leq \alpha_1 + \alpha_2 << k. \tag{16}$$

The first possibility is to consider a direct reduced optimization problem of dimension α and use the same strategy as before (figure 4). The second possibility is to include the use of approximation methods on this problem (see figure 5).

An alternative would be to consider two separate optimization problems, one of size α_1 and one of size α_2 and consider some sequential strategy. For instance we could optimize the problem of dimension α_1 and then optimize the problem of dimension α_2 to move the design toward the constraint boundary. It may be necessary to repeat the above.

The optimization could be performed using a direct strategy or by utilizing approximation methods as before. We note again, that the second approach, considering two separate problems, can reduce the number of training data required to produce an accurate approximation.

5 Industrial Case Study

To demonstrate the ideas involved we consider applying various optimization strategies to the structure shown in figure 6. This model of a spoked structure was supplied by Rolls-Royce plc. The part of

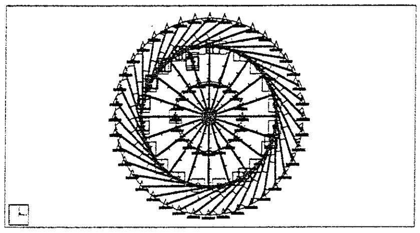

Figure 6: The structure under consideration.

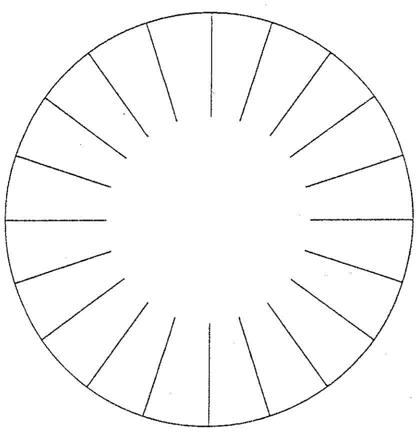

Figure 7: The part of the structure to be optimized.

the structure being optimized is shown in figure 7. The rest of the model enforces realistic boundary conditions on the structure. The structure is subjected to realistic loadings. The optimization problem is as follows: minimize

$$w(\mathbf{x}) \tag{17}$$

subject to

$$\sigma(\mathbf{x}) < 50000 \tag{18}$$

where w is the weight of the structure, σ is the maximum von Mises stress in the structure and $\mathbf{x}$ is an 11 dimensional vector defining the structural geometry. This vector contains the design variables listed in table 1. These are varied between realistic bounds, again supplied by industry. Each input is given here in terms of a percentage where 0% refers to the lower bound of each design variable and 100% refers to the upper bound of each design variable.

First, the full 11d problem is optimized, we make use of the OPTIONS design exploration system here.[5] The optimization algorithm employed is dynamic hill climbing, this is a gradient based optimization strategy using multiple restarts.[15] We

Table 1: Design variables.

Design variable	Description
x_1	spoke rotation
x_2	spoke section
x_3	ring section 1
x_4	ring section 2
x_5	ring section 3
x_6	ring section 4
x_7	ring section 5
x_8	ring section 6
x_9	ring curvature
x_{10}	ring section 7
x_{11}	ring section 8

Table 2: "Optimal" design variables using 11d DHC optimization.

Design variable	Value
x_1	0.00
x_2	0.00
x_3	100.00
x_4	100.00
x_5	86.12
x_6	94.54
x_7	99.46
x_8	100.00
x_9	98.68
x_{10}	100.00
x_{11}	100.00
Resulting objective	Resulting constraint
7.8636	49793.50

made over 2500 calls to the finite element model; the best final design is shown in table 2.

We show the optimization convergence history, plotting the number of calls to the finite element model against the best minimum weight in figure 8. We note here that subsequent restarts have failed to produce a better minimum than the first starting point (all were worse), this will not generally be the case. We further note that this direct search has not found the global minimum of the problem, a better design will be found shortly.

We next consider a screening study applied to the objective and the constraint. An 11d screening study is defined using edge designs formed from conference matrices. This is shown in figure 9. Here a + and a − refer to the upper and lower bound of each design variable respectively. The differences along edges can then be calculated and used to infer the importance of each design variable. These

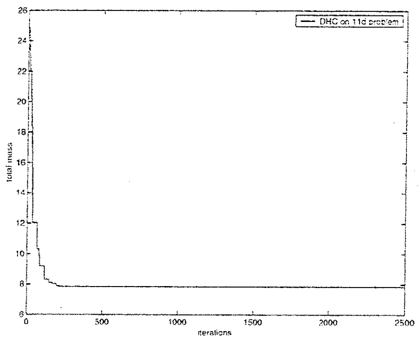

Figure 8: Convergence history of 11d DHC optimization.

	1	x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	x_9	x_{10}	x_{11}
1	+	+	+	−	+	+	+	−	−	−	+	−
2	+	−	+	+	−	+	+	+	−	−	−	+
3	+	+	−	+	+	−	+	+	+	−	−	−
4	+	−	+	−	+	+	−	+	+	+	−	−
5	+	−	−	+	−	+	+	−	+	+	+	−
6	+	−	−	−	+	−	+	+	−	+	+	+
7	+	+	−	−	−	+	−	+	+	−	+	+
8	+	+	+	−	−	−	+	−	+	+	−	+
9	+	+	+	+	−	−	−	+	−	+	+	−
10	+	−	+	+	+	−	−	−	+	−	+	+
11	+	+	−	+	+	+	−	−	−	+	−	+
12	+	−	+	−	+	+	+	−	−	−	+	−
13	+	−	−	+	−	+	+	+	−	−	−	+
14	+	+	−	−	+	−	+	+	+	−	−	−
15	+	−	+	−	−	+	−	+	+	+	−	−
16	+	−	−	+	−	−	+	−	+	+	+	−
17	+	−	−	−	+	−	−	+	−	+	+	+
18	+	+	−	−	−	+	−	−	+	−	+	+
19	+	+	+	−	−	−	+	−	−	+	−	+
20	+	+	+	+	−	−	−	+	−	−	+	−
21	+	−	+	+	+	−	−	−	+	−	−	+
22	+	+	−	+	+	+	−	−	−	+	−	−

Figure 9: Experimental design for screening study.

are plotted for weight in figure 10 and for von Mises stress in figure 11. From these figures it is clear that variables x_1, x_2, x_{10} and x_{11} appear to have most influence on the objective whereas variables x_2, x_5, x_6 and x_9 appear to have most influence on the constraint. Although variable x_2 is important for both

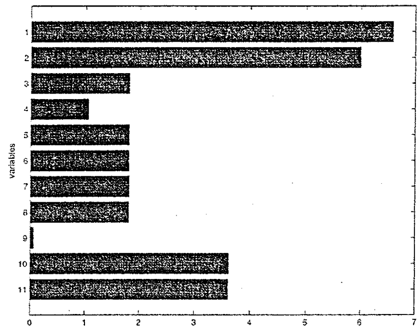

Figure 10: Screening study applied to total mass.

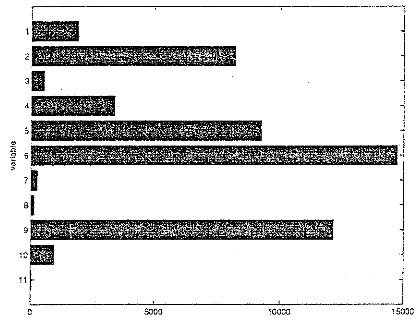

Figure 11: Screening study applied to von mises stress.

the objective and the constraint, the screening study suggests that the value of x_2 that minimizes weight (0%) also reduces stress. As a result the value of variable x_2 will be fixed at 0% and will not be considered further. We can also use the information provided by the screening strategy to suggest fixed values for the other variables which have little influence on the response (i.e. the weight and von Mises stress). These are set as x_3 = 100%, x_4 = 100%, x_7 = 100% and x_8 = 100% and will not be considered further.

The first strategy we adopt is as in figure 4, a direct 6d optimization on the remaining variables, again using dynamic hill climbing. Here 500 calls to the finite element model are taken and the results are as shown in table 3. The optimization convergence history is shown in figure 12. We note that a better solution is found whilst using fewer calls of the finite element model, even allowing for the 22 calls used in the screening study.

Table 3: "Optimal" design variables 6d DHC optimization.

Design variable	Value
x_1	0.00
x_5	87.22
x_6	95.74
x_9	95.12
x_{10}	100.00
x_{11}	99.89
Resulting objective	Resulting constraint
7.8164	49907.94

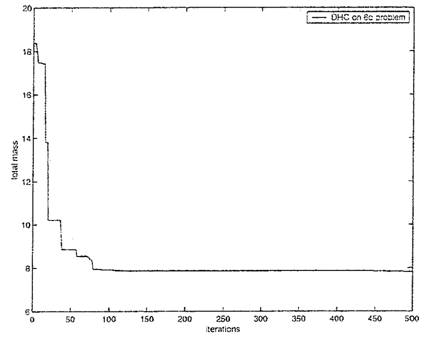

Figure 12: Convergence history for 6d DHC optimization.

The next strategy we adopt is to split the six dimensional optimization problem into two three dimensional problems again based on the screening study. We first optimize with respect to variables x_1, x_{10} and x_{11} and fix variables x_5, x_6 and x_9 at values which reduce the stress in the model. This information is also available from the screening study and suggests that we fix x_5 = 0%, x_6 = 0% and x_9 = 100%. Although the constraint is the most complex response in this model, the screening study suggests it varies little when we alter variables x_1, x_{10} and x_{11}, so we can perform optimization using a relatively small number of function evaluations, 30 here. The resulting optimum is shown in table 4 and a convergence plot is shown in figure 13.

We then perform optimization with respect to variables x_5, x_6 and x_9 using the design in table 4 as a starting point. The behaviour of the stress in now likely to be more complex so more finite element evaluations are required (we use 250 here). The final minimum is shown in table 5 and the resulting convergence history is plotted in figure 14. In total we perform 22 finite element evaluations on the screening study, 30 on the first optimization and 250 on the second optimization, a total of 302 finite element evaluations in all.

Table 4: Optimal design variables 3d DHC optimization (phase 1).

Design variable	Value
x_1	0.00
x_5	0.00
x_6	0.00
x_9	100.00
x_{10}	100.00
x_{11}	100.00
Resulting objective	Resulting constraint
11.1096	47284.34

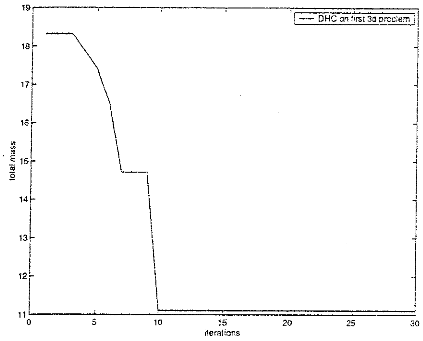

Figure 13: Convergence history for 3d DHC optimization (phase 1).

Table 5: Optimal design variables 3d DHC optimization (phase 2).

Design variable	Value
x_1	0.00
x_5	85.90
x_6	94.11
x_9	99.43
x_{10}	100.00
x_{11}	100.00
Resulting objective	Resulting constraint
7.8657	49509.96

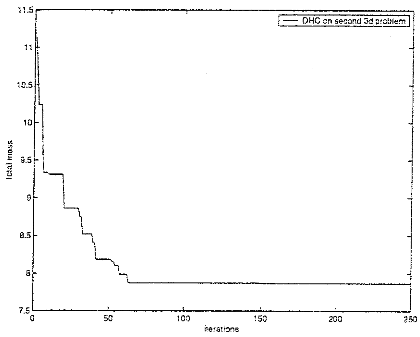

Figure 14: Convergence history for 3d DHC optimization (phase 2).

Table 6: Optimal design variables using 6d approximation.

Design variable	Value
x_1	10.34
x_5	87.36
x_6	93.92
x_9	70.22
x_{10}	100.00
x_{11}	100.00
Resulting objective	Resulting constraint
7.8794	49425.17

Finally, to further improve the efficiency of the optimization process we consider the use of approximation methods. Here we apply the kriging model discussed in section 2 to the six dimensional and the two three dimensional optimization problems.

Results of optimizing the six dimensional approximation are shown in table 6. Here an optimal latin hypercube design consisting of 30 finite element runs is used to construct the approximation. A further 27 calls to this model are then made during optimization. A total of $22 + 30 + 27 = 79$ finite element evaluations are made in total.

When considering the two three dimensional problems, as highlighted earlier, the first three dimensional problem appears to be relatively simple, so an experimental design (optimal latin hypercube) consisting of eight finite element runs is used to construct the approximation. We then optimize the approximation surface, immediately producing the design in table 4. We then add this point and reconstruct the approximation. Optimizing again gives the same point so the algorithm terminates requiring a total of nine calls to the finite element model.

We then consider the second three dimensional problem. Due to the increased complexity we start with a larger experimental design (here we chose an optimal latin hypercube design consisting of 15 runs). We then optimize the approximation and evaluate the finite element model at the predicted optimum, reconstruct a new approximation and repeat (see figure 5 for a schematic of this approach). We consider 15 further calls to the finite element model requiring 30 finite element evaluations in total. The final design obtained is shown in table 7. This requires a total of $22 + 9 + 30 = 61$ finite element evaluations.

We briefly summarize all these results it table 8. We note that in general all we require is a good design and not necessarily the true global optimum. A

Table 7: Optimal design variables using two 3d approximations (phase 2).

Design variable	Value
x_1	0.00
x_5	86.98
x_6	95.58
x_9	80.56
x_{10}	100.00
x_{11}	100.00
Resulting objective	Resulting constraint
7.8203	48194.18

Table 8: Summary of findings.

Opt. strat.	# FE evals.	Min. weight
Direct 11d DHC	2786	7.8636
6d problem DHC	522	7.8164
2 3d problems DHC	302	7.8657
6d approx.	79	7.8794
2 3d problems approx.	61	7.8203

comparison of the four approaches is shown in figure 15 (the 22 finite element evaluations required for the screening study are also included here). A combination of a screening study and an approximation method produces the best strategy in terms of efficiency allowing a very good design to be produced with relatively few calls to the finite element model.

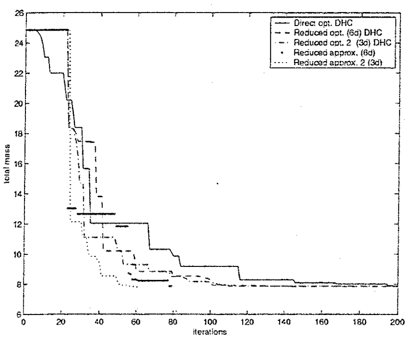

Figure 15: Comparison of approaches.

6 Conclusions

Efficient strategies for optimizing expensive models have been described. The first, approximation methods, are well known to the structural optimization community. The second, screening studies are less known and are used to reduce the dimensionality of the optimization problem. This makes optimization subject to a large number of design variable (> 10) more realistic. In particular, the combined use of both screening strategies and approximation methods is demonstrated. This allows accurate surrogate models to be constructed and optimized with very few calls to the computational model of interest.

Further, it is often the case in structural optimization that there will be multiple objectives and/or multiple constraints and several strategies for optimization are proposed, based on either grouping all the important design variables together or by using some sequential strategy. It is not obvious which will perform best in general: a sequential strategy strategy may require fewer evaluations of the expensive model but it runs the risk of potentially missing out important interaction effects between design variables.

We have demonstrated the approach with encouraging results on an industrial case study provided by Rolls-Royce plc. Despite these results much remains unanswered. Our approach involves the use of DoE, screening and approximation methods. There are many ways of doing all three but which is the best combination in general is by no means obvious. No doubt it would be problem dependent but we feel that guidelines as to which will work well in certain situations would be beneficial. This is beyond the scope of this paper and would require many different test cases encompassing a broad range of engineering problems before any real conclusions could be drawn.

In this paper we opt for approaches which in our (limited) experience work well. We are by no means suggesting that the approaches we utilize form the best overall strategy.

Further work of interest to us is the application of the above techniques to multidisciplinary optimization.

Acknowledgements

We would like to thank Rolls-Royce Plc for the financial support received as a part of the University Technology Partnership (UTP) in design as well as for providing the example used in the industrial case study. We also thank Dr. Steve Wiseall of Rolls-Royce Plc and Professor Sue Lewis of the Statistics group, Faculty of Mathematical Studies, University of Southampton for their useful comments on an earlier version of this manuscript.

References

[1] Clarich, C.; Mosetti, G.; Pediroda, V. and Poloni, C., Application of evolutive algorithms and statistical analysis in the numerical optimization of an axial compressor, The 9th International Symposium on Transport Phenomena and Dynamics of Rotating Machinery, Honolulu, Hawaii, 2002.

[2] Elster, C. and Neumaier, A., Screening by conference designs, Biometrika **82**, pp. 58-602, 1995.

[3] Johnson, M.; Moore, L. and Ylvisaker, D., Minimax and maximin distance designs, Journal of Statistical Planning and Inference **26**, pp.131–148, 1990.

[4] Jones, D. R.; Schonlau, M. and Welch, W. J., Efficient global optimization of expensive black-box functions, Journal of Global Optimization **13**, pp. 455-492, 1998.

[5] Keane, A. J., The OPTIONS design exploration system - reference manual and user guide - version B3.0, 2000.

[6] Kleijnen, J. P. C., Statistical tools for simulation practitioners, New York: Marcel Dekker, 1987.

[7] McKay, M. D.; Beckman, R. J. and Conover, W. J., A comparison of three methods for selecting values of input variables in the analysis of output from a computer code, Technometrics **21**, pp. 239-245, 1979.

[8] Morris, M. D. and Mitchell, T. J., Exploratory designs for computer experiments, Journal of Statistical Planning and Inference **43** pp. 381-402, 1995.

[9] Myers, R. H. and Montgomery, D. C., Response surface methodology: process and product optimization using designed experiments, John Wiley and Sons inc., 1995.

[10] Powell, M. J. D., Radial basis functions for multivariable interpolation: a review. In: Algorithms for Approximation, Oxford: Clarendon Press, pp. 143-167, 1987.

[11] Press, W. H.; Flannery, B. P.; Teukolsky, S. A. and Vetterling, W. T., Numerical recipes in fortran, Cambridge University Press, 1989.

[12] Trocine L. and Malone L., Finding important independent variables through screening designs: a comparison of methods, Proceedings of the 2000 Winter Simulation Conference, 2000.

[13] Welch, W. J.; Buck, R. J.; Sacks, J.; Wynn, H. P.; Mitchell, T. J. and Morris M. D., Screening, predicting, and computer experiments, Technometrics **34** pp.15-25, 1992.

[14] White, H.; Gallant, A. R.; Kornik, K.; Stinchcombe M. and Woolridge J., Artificial Neural Networks: approximation theory and learning, Blackwell Publishers, 1992.

[15] Yuret, D.; and de la Maza, M., Dynamic hill climbing: overcoming the limitations of optimization techniques, Proceedings of the 2nd Turkish Symposium on AI and ANN, pp. 254-260, 1993.