ISEF 2011 - XV International Symposium on Electromagnetic Fields
in Mechatronics, Electrical and Electronic Engineering
Funchal, Madeira, September 1-3, 2011

EXPLORATION VERSUS EXPLOITATION USING KRIGING SURROGATE MODELLING IN ELECTROMAGNETIC DESIGN

Song Xiao, Mihai Rotaru and Jan K. Sykulski

School of Electronics and Computer Science (ECS), University of Southampton, Southampton, UK
sx1g10@ecs.soton.ac.uk, mr@ecs.soton.ac.uk, jks@soton.ac.uk

Abstract **–**This paper discusses the use of kriging surrogate modelling in multiobjective design optimisation in electromagnetics. The importance of achieving appropriate balance between exploration and exploitation is emphasised when searching for the global optimum. It is argued that this approach will yield a procedure to solve time consuming electromagnetic design problems efficiently and will also assist the decision making process to achieve a robust design of practical devices considering tolerances and uncertainties.

Introduction

Electromagnetic design almost always carries a heavy burden of high computational cost, with very few exceptions when a very simplistic analytical, empirical or equivalent circuit based model is found to be adequate for performance prediction. Most of the time throughout the design process, or at least at later stages, numerical models are required to provide necessary accuracy, typically employing 3D simulation using finite element or related technique. In the optimisation part of the design routine a single objective function evaluation may require a full field solution of the entire complicated model, often transient, or even several solutions (if averaged values are needed), which may be very 'expensive' in terms of computing times involved. Thus it is not enough to have confidence that the algorithm finds the global optimum; for practical purposes it must do so with as few objective function calls as possible. Thus within the context of searching for the optimum (usually minimum) of a particular objective function (or functions in multiobjective problems, e.g. best performance and simultaneously minimum cost), another minimum is being sought, that is looking for a strategy which finds the optimum with a minimum use of the computationally expensive performance predicting software. To complicate things further, the issue of robustness of the design comes into consideration – related to manufacturing tolerances, material variability, etc – which requires the designer not only to find the optimum design but also know more about its 'quality', in other words the 'shape' of the objective function must be estimated. In the context of stochastic optimisation this is usually expressed in terms of a compromise between exploration (searching the unexplored space) and exploitation (using information already provided) and is often supplemented and supported by various types of surrogate modelling. This paper investigates these issues and uses 'kriging' as the main technique for constructing the surrogate model.

Kriging and the Utility Functions

Kriging [1] can predict the shape of the objective function based only on limited information and estimates the accuracy of this prediction; this is helpful in assisting the main decision of the optimisation process where to put the next point for evaluation. A 'utility function' is usually constructed, based on the predicted error, which may seamlessly adjust the way of searching between the regions with confidence and uncertainty. Thus providing an efficient and robust way to achieve a balance between exploration of unknown regions with degree of uncertainty and exploitation of attractive areas with high confidence is imperative.

A brief overview of one-stage kriging methodology is first given. The method exploits the spatial correlation of data in order to build interpolation; hence the correlation function is very important. We use the standard linear regression (1) and the correlation is modelled as (2):

$$\hat{y}(x) = \sum_{k=1}^{m} \beta_k f_k(x) + \varepsilon(x) \quad (1) \qquad R(\varepsilon(x^i), \varepsilon(x^j)) = \prod_{k=1}^{n} e^{-\theta_k \left| x_k^i - x_k^j \right|^{p_k}} \quad (2)$$

where the global function $\sum_{k=1}^{m} \beta_k f_k(x)$ and an additive Gaussian noise $\varepsilon(x)$ are integrated to the predicted value $\hat{y}(x)$ of the objective function. θ_k is the correlation amongst the data in *k*-direction and p_k determines the 'smoothness' of (2). The most popular correlation function is given by the Gauss model where the value of p_k is simply taken as equal to 2.

In general, the 'expected improvement' utility function, based on the potential error predicted by a kriging model, is commonly used to select multiple design vectors for evaluation. The Expected Improvement Function [2] is defined as

$$EIF[I(x)] = \begin{cases} (f_{\min} - \hat{y}(x))\psi\left(\frac{f_{\min} - \hat{y}(x)}{s(x)}\right) + s(x)\phi\left(\frac{f_{\min} - \hat{y}(x)}{s(x)}\right) & if \quad s(x) > 0 \\ 0 & if \quad s(x) = 0 \end{cases} \quad (3)$$

where $\hat{y}(x)$ is the objective function value of *x* as predicted by the kriging model, given by equation (1), $s(x)$ is the root mean squared error in this prediction, and $\psi\left(\frac{f_{\min} - \hat{y}(x)}{s(x)}\right)$ and $\phi\left(\frac{f_{\min} - \hat{y}(x)}{s(x)}\right)$ are Gaussian density function and Gaussian distribution function, respectively.

The Expected Improvement Function (EI) may be viewed as a fixed compromise between exploration and exploitation: when the *s(x)* operator given by the Kriging method is positive, the first term of equation (3) favours searching the promising regions with high confidence, whereas the second term in the same equation favours searching the regions with high uncertainty. Through a set of practical kriging-assisted single-objective tests developed specially to assess the performance of these two terms, it has been shown that the second term representing exploration performs dramatically better in terms of finding the global optimum of the objective function, whereas the exploitation often can only find the local minimum. Since EI applies equal weights to the two terms, it may be seen as a fixed compromise between exploration and exploitation.

The balance between exploration and exploitation is a critical issue when attempting to find the global optimum of an objective function. The Weighted Expected Improvement (WEI) [3] is derived from EI by adding a tuneable parameter which can adjust the weights on exploration and exploitation, whilst the quality of the approximation of the objective function can be improved by incorporating the newly evaluated design vector at each iteration. The WEI utility function used in this work may be written as

$$WEIF[I(x)] = \begin{cases} w(f_{\min} - \hat{y}(x))\psi\left(\frac{f_{\min} - \hat{y}(x)}{s(x)}\right) + (1-w)s(x)\phi\left(\frac{f_{\min} - \hat{y}(x)}{s(x)}\right) & if \quad s(x) > 0 \\ 0 & if \quad s(x) = 0 \end{cases} \quad (4)$$

where the tuneable parameter *w* (0 < *w* < 1) controls the balance between the two terms (exploration and exploitation), therefore searching globally and locally [3]. The efficiency of the kriging with WEI has been tested with the Schwefel test function [4] as an objective function in the interval [-500 500] for different values of w. The multi-dimensional Schwefel test function is defined as

$$f(x) = \sum_{i=1}^{d} - x_i \sin(\sqrt{|x_i|}) \quad (5)$$

In the one-dimensional case used here, *d*=1. We have studied the performance of different algorithms using the Schwefel function throughout the tests. While testing using a single function may not be conclusive, the Schwefel function has in the past been found helpful when testing similar algorithms.

It has been found that the kriging model assisted by WEI when $w \in$ [0.55 1) can only find a local minimum; this is perhaps not surprising as a strong weight has been applied to the term favouring exploitation. When $w \in$(0 0.54] the kriging model is able to find the global minimum. Notable values of w are $w=1$, which puts all emphasis on exploitation, $w=0$, which focuses on exploration, and $w=0.5$, which makes the algorithm equivalent to EI. Table 1 summarizes the results of our tests.

Value of weight	Number of iterations	Value of weight	Number of iterations	Value of weight	Number of iterations
1	Fails	0.57	3 (finds LM)	0.5 (EI)	11 (finds GM)
0.9	3 (finds LM)	0.56	9 (finds LM)	0.4	7 (finds GM)
0.8	3 (finds LM)	0.55	9 (finds LM)	0.3	12 (finds GM)
0.7	3 (finds LM)	0.54	13 (finds GM)	0.2	17 (finds GM)
0.6	3 (finds LM)	0.53	14 (finds GM)	0.1	15 (finds GM)
0.59	3 (finds LM)	0.52	11 (finds GM)	0	Fails
0.58	3 (finds LM)	0.51	11 (finds GM)		

Table 1: Performance of WEI for *w* between 0 and 1 (LM – local minimum, GM – global minimum)

In order to understand better the effects of *w* more tests in the range of 0.5 to 0.6 were done. As shown in the table somewhere between 0.54 and 0.55 there is a changeover between a regime where only a local minimum is found and values of *w* which allow for the global minimum to be correctly identified. Thus too much emphasis on exploitation is a risky strategy. Equal weights (*w*=0.5 as in EI) are ‘safe’, but not optimal in a sense that there is a value of *w* around 0.4 which can provide an answer with fewer iterations (7 instead of 11). Figure 1 shows a snapshot position after the global minimum has been found after 11 iterations (using EI) and after 7 iterations (using *w*=0.4). For both cases the same six initial points were used (in practice their positions may be selected randomly) required before a particular EI or WEI strategy can be applied. The graph also shows the ‘history’ of how the points were added throughout the iterative process. Both strategies successfully find the global minimum and the quality of the final answer is comparable, but WEI with *w*=0.4 is more efficient.

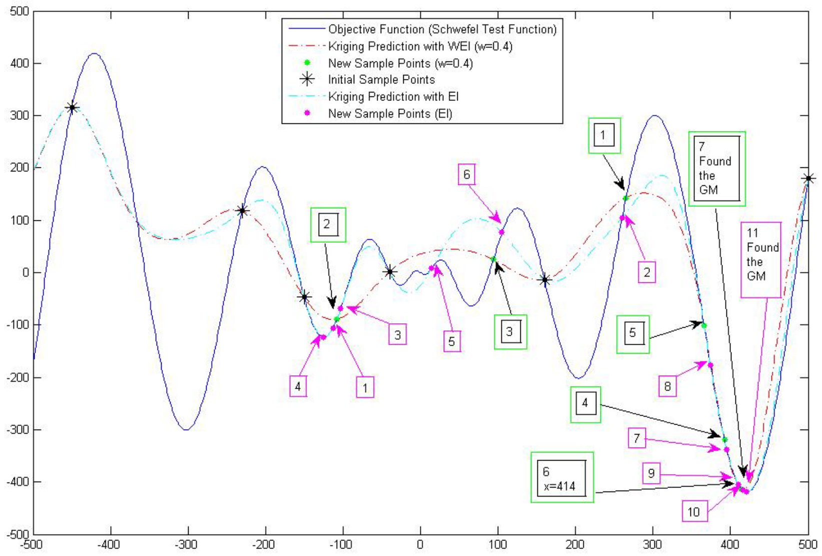

Figure 1: The performance of the kriging model with WEI ($w=0.4$) and EI
(the single square with a number indicates the iteration number of kriging with EI;
the double square with a number shows the iteration number of kriging with WEI ($w=0.4$)).

Adaptive Weighted Expected Improvement

The experiments of the previous section demonstrated the importance of the optimal choice of the weights, both in terms of the ability of the algorithm to achieve the correct answer (global minimum) and doing it efficiently (fewer iterations required); unfortunately the optimal choice of w is normally problem dependent and thus a modified strategy is required to make the method more intelligent and guide itself automatically through the process.

Reinforcement learning is a goal-directed learning approach to what to do next and how to map the situation to actions so as to maximize a numerical reward [6]. In this paper we propose to automatically tune the weighting parameter w in response to the environment feedback. In particular, the Mean Square Error (MSE) from the kriging model is used to guide the choice of the optimum weight w and the concept of an award is introduced. Thus the algorithm calculates the average value of the MSE of every predicted point and uses these values as the basis of calculating the potential rewards. Then, after comparing the rewards from different weight distributions, the weights are redistributed on the two terms which control the exploration and exploitation so that the biggest reward is achieved. The Adaptive Weighted Expected Improvement (AWEI) strategy is described as one of the possible algorithms in Fig. 2. AWEI endeavours to encourage exploration or exploitation depending of the results of the initial pre-test, one with emphasis on exploration and another on exploitation. Two rewards (Reward1 and Reward2) are calculated and compared; on the basis of this comparison w is then chosen to encourage either exploration or exploitation.

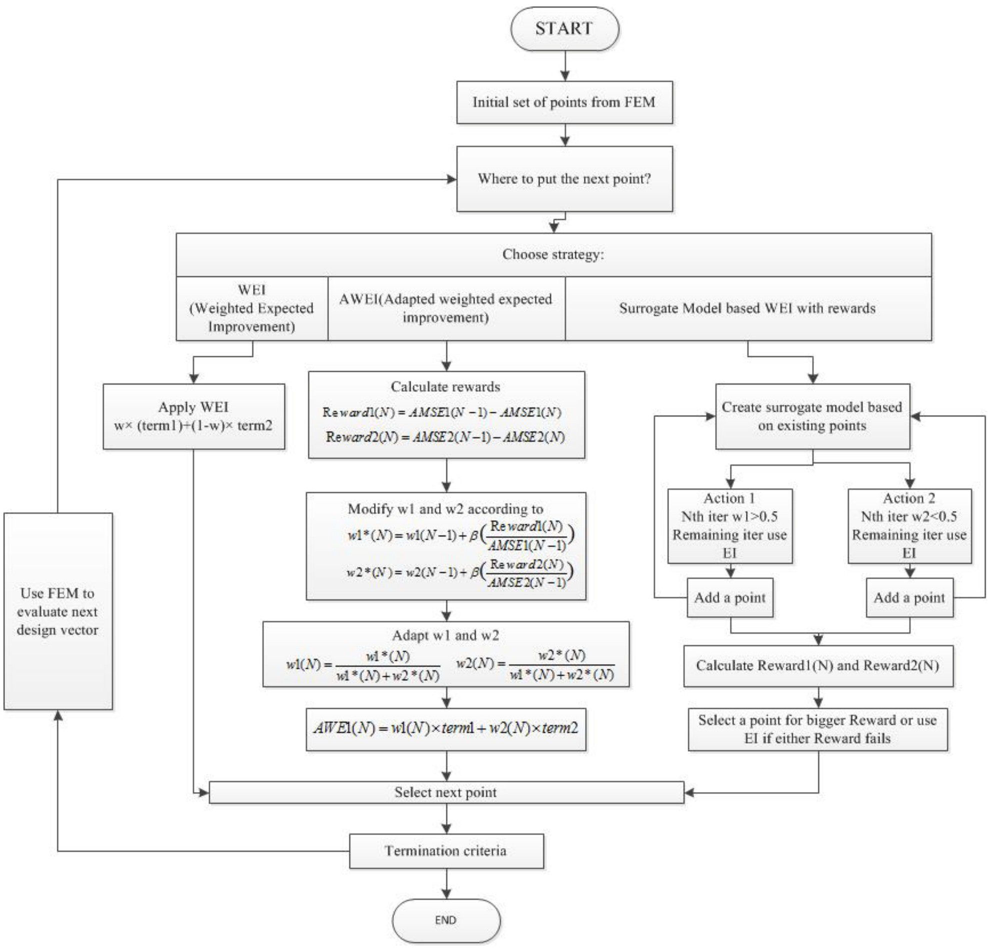

Figure 2: The decision-making chart for different strategies of balancing exploration and exploitation; term1 favours exploration while term2 favours exploitation; AMSE = average mean square error

Practical Performance of the Adaptive Weighted Expected Improvement

Several tests using the AWEI assisted kriging model for different values of β (a selectable parameter as described in Fig. 2) have been undertaken to assess its performance. However, one particular problem was identified and needed special attention. The term which encourages exploitation can sometimes cause the kriging model to stop because of choosing repeatedly the same new point for evaluation (within the specified accuracy). Should this happen (or should – for any other reason – one of the rewards not be assessed properly or fail), the algorithm is effectively reset and the Expected Improvement (EI) function is temporarily applied to select the next point for evaluation; in the next step the algorithm reverses to the AWEI. We have used the Schwefel test function again with the initial sample points imposed as $x = -450, -230, -150, -40, 160, 500$ and the tuneable parameter β varied in a controlled way. When β = 0.001, 0.005, 0.01 the model fails to find the global minimum; when β = 0.05, 0.08, 0.1 altogether 12 iterations are needed to find the global minimum; when β = 0.2 or 0.3 a better performance is observed with 8 iterations needed to find the global minimum. As demonstrated by Fig. 3, however, the best performance with only 5 iterations needed was observed when β = 0.35. Compared with the previously described WEI, the AWEI is more flexible thanks to the built-in feedback that uses the reward scheme to make decisions on how to adapt the EI function.

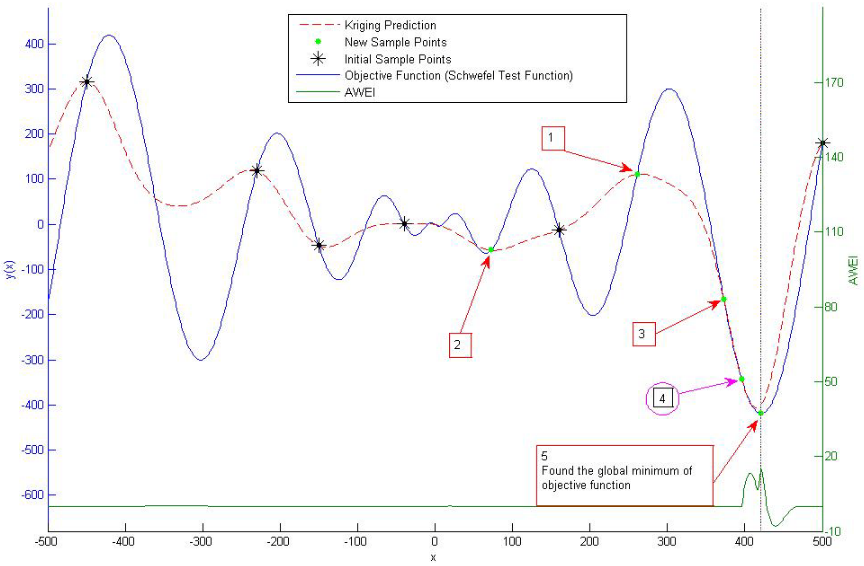

Figure 3: The performance of the kriging model with AWEI for β=0.35
(the circle with a number means more exploration at that iteration;
the square with a number means more exploitation at that iteration).

Surrogate Model based Weighted Expected Improvement approach with rewards

The AWEI is based on reinforcement learning and takes account of the feedback, which in turn uses predicted uncertainty gained from the kriging model to make a decision as a trade-off between exploitation and exploration driven by the amount of reward resulting from each action [7]. The AWEI consistently selects the action which yields the largest average reward [8] at each step of the iterative process based on the best information available, which may not necessarily be accurate or reliable. So although optimal in short term the selected action may not always be beneficial in long term. In the third strategy developed and presented in this paper an attempt is made to predict the cumulative rewards likely to occur on long terms as a consequence of a particular choice of actions. This approach follows the ideas first introduced recently in [8] in the context of games theory to a well known one-armed bandit problem. This approach requires predicting the long term awards, rather than short term at a given iteration step, which necessities some estimation of the long term consequences of the actions selected. A simple (but very inefficient in the context of electromagnetic design problems) approach would involve continuing iterations independently (in parallel) for the two initially selected (and at that point fixed) weight functions using WEI until the kriging process stops in either of the tests (because of repeating the point for evaluation) and then using the most recent calculated value of the rewards to select the more promising action. This strategy has the advantage of assessing long term benefits (rather than immediate ones) but can only be applied to problems where objective function evaluation is 'cheap' (in terms of computing times) – as was indeed the case in the original paper [8]. However, it appears that the main concept can still be useful if supplemented by another modification to the algorithm with the aid of surrogate modelling. Thus rather than using the 'expensive' model (typically the time consuming finite element field modelling software) we can create a simplified surrogate model based on existing data points and continue the parallel search for the global minimum of the surrogate model – which will be a very quick process and thus not adding noticeably to the overall computing times – before the rewards for the two alternative actions are compared and the final decision is made regarding the location of the next point for evaluation. We then use finite element (or similarly 'expensive' software) to get a new 'reliable' point, update the surrogate model and continue iterations.

Thus a particular contribution of this paper to the concept of predicting the likely long term potential rewards before the next point is evaluated is the addition of surrogate modelling so that the 'forward prediction' is cheap; inevitably such a prediction will be less reliable than using real data points (which in this application, as already stressed, would be too expensive and thus unacceptable) but may result in an overall better assessment of long term benefits of different actions than a simple one-stage algorithm developed and described earlier as the Adaptive Weighted Expected Improvement strategy.

There are of course many methods of constructing a surrogate model. We suggest, and have implemented, using the root mean square error already available in the kriging prediction, even though this is not the real error between the kriging approximation and the real objective function (which is of course unknown at this stage). A random error is added to the calculated error and thus effectively we now have two kriging surrogate models simultaneously, the original one based on the most recent 'real' data points, and a second one – used only for the purpose of 'forward prediction' of the long term effects of a particular action – which ultimately leads to an overall long term 'reward' of a particular action. As there are two possible actions and they are assessed independently we end up with two rewards; the better reward will identify the better cause of action, a new point is selected, the finite element programme executed and a new point added to the curve. This will give rise to a new surrogate model and a new 'secondary' surrogate model (or rather a pair of models as there are two parallel actions); the process will continue until some termination criteria are met. The flowchart of the decision making process can be easily followed in Fig. 2.

As before the Schwefel test function was chosen to test the Surrogate Model based Weighted Expected Improvement approach with rewards (SMWEI). The choice of the values $w1$ and $w2$ is a matter of further experiments in order to generate some guidelines about how to select the initial values. Moreover, as the two actions are independent the weights $w1$ and $w2$ do not actually have to add up to 1, although most of the testing assumed that they do. Finally, as the 'second' kriging surrogate model relies on a random distribution of error, all tests were conducted ten times with the same pair of values of $w1$ and $w2$ and performance averaged. Throughout the testing the same initial sample points were assumed to allow comparison, but in reality such points may be distributed randomly or using one of the accepted strategies such as a Latin Hypercube sampling. Some results will now be discussed. When $w1=0.6$ and $w2=0.4$ the number of iterations required to find a global minimum is between 10 and 16, when $w1=0.7$ and $w2=0.3$ it is between 9 and 16, for $w1=0.8$ and $w2=0.2$ it is 10 to 17, finally for $w1=0.9$ and $w2=0.1$ it is 11 to 15. The average number of iterations for the pairs of values above is 13, 12, 15 and 12, respectively, so it is quite steady and does not appear to be sensitive to the variation of the values. Figure 4 shows a particular set of results for $w1=0.7$ and $w2=0.3$; there is an interesting 'departure' at iteration no 10 to explore a remote region before returning to the global minimum at iteration 11. Finally, Fig. 5 demonstrates a particularly successful case when $w1=0.7$ and $w2=0.1$ where only 6 iterations were required.

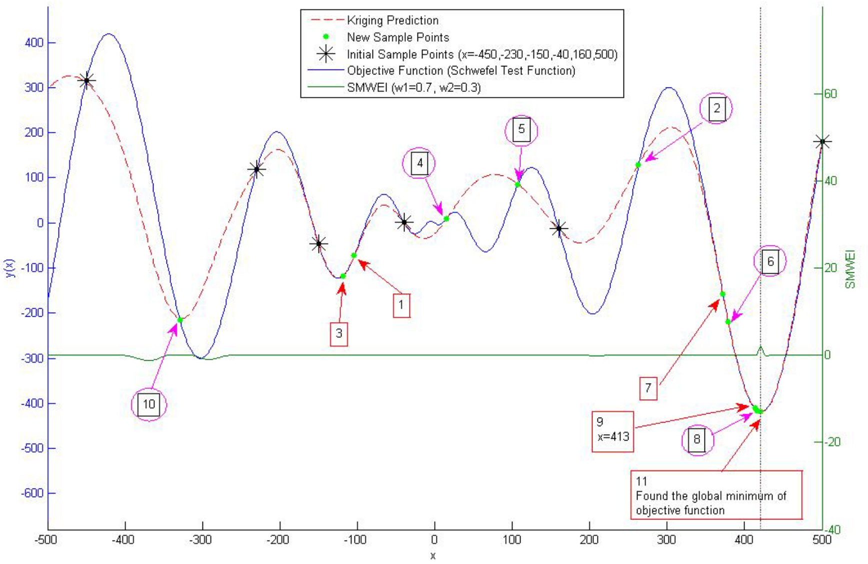

Figure 4: The performance of the kriging surrogate model with SMWEI
(the circle with a number means more exploration at that iteration;
the square with a number means more exploitation at that iteration).

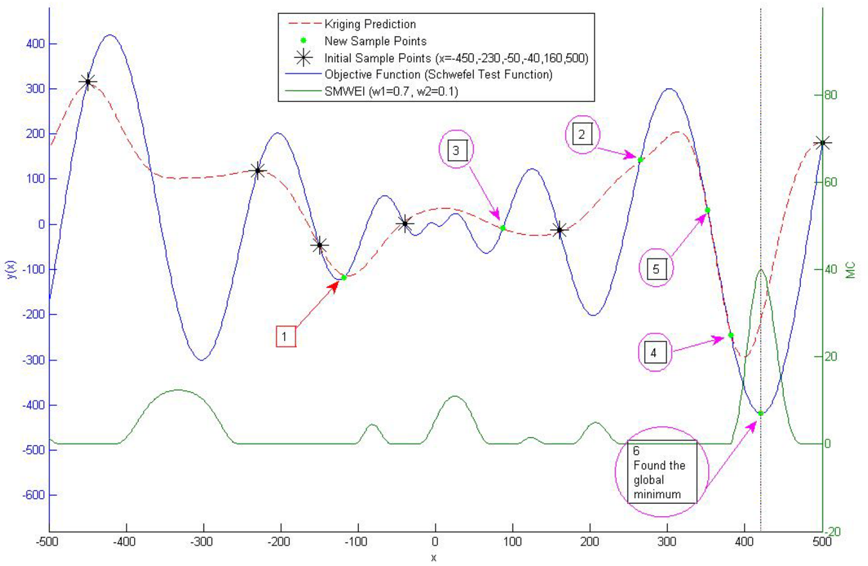

Figure 5: The 'best' performance of SMWEI with *w*1=0.7 and *w*2=0.1
(the circle with a number means more exploration at that iteration;
the square with a number means more exploitation at that iteration).

There appears to be little benefit in applying the last strategy compared with the previous AWEI algorithm, but more testing is required to draw more meaningful conclusions. In particular, it is interesting to notice that the SMWEI clearly makes a better attempt at exploring local minima – this may prove very important in the context of robust design where not only the value but also the shape of the minimum is of relevance. Thus a strategy which explores the space more thoroughly may after all be preferable even if more expensive. At this stage the three strategies are considered as alternative and all are available in the flowchart of Fig 2. An attempt will be made in the future to provide guidelines about how to select one strategy for the given problem in hand.

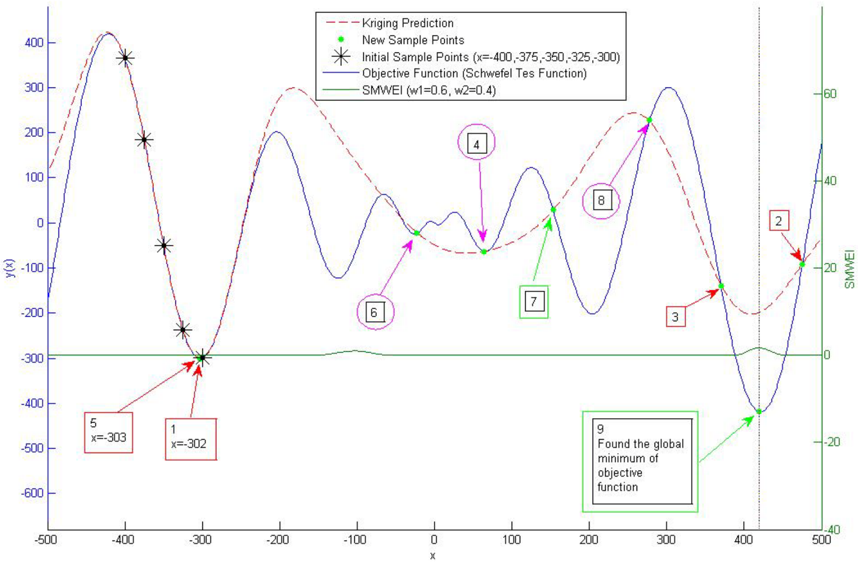

Fig. 6. The performance of SMWEI facing an ‘extreme’ case (a circle with a number means more exploration, whereas a square means more exploitation at that iteration; a double square means using EI at that iteration)

Finally, a test was conducted to see how the algorithm performs when the initial points are not distributed favourably, for example if positioned as in Fig. 6. This is clearly a challenging case as the initial points give very little clue as to the real shape of the objective function. Rather remarkably the SMWEI algorithm performs very robustly with only 9 iterations required to find the global minimum, whereas kriging with EI needs 12 iterations. Other tests of that nature were equally encouraging.

Conclusion

Two novel algorithms have been proposed, both adopting concepts of reinforcement learning, in an attempt to automatically balance exploration and exploitation in computationally expensive electromagnetic design optimisation problems. Both are based on kriging surrogate modelling and use the notion of rewards for selecting the best position of the next point for evaluation. The one-stage algorithm appears to perform very efficiently in terms of its ability to find a global minimum, whereas the strategy based on two kriging surrogate models and forward performance prediction offers more reliable information about the shape of the objective function. Both algorithms will be implemented in practical design of electromagnetic and electromechanical devices.

References

[1] L. Lebensztajn, C. A. R. Marreto, M. C. Costa and J-L. Coulomb, Kriging: a useful tool for electromagnetic device optimization, IEEE Transactions on Magnetics, Vol. 40, No. 2, pp. 1196-99, 2004.
[2] D. R. Jones, M. Schonlau and W. J. Welch, Efficient global optimization of expensive black-box functions, Journal of Global Optimization, Vol. 13, pp. 455-492, 1998.
[3] A. Sobester, S. J. Leary and A. J. Keane, On the design of optimization strategies based on global response surface approximation models, Journal of Global Optimization, Vol. 33, pp. 31-59, 2005.
[4] V. Picheny, D. Ginsbourger and Y. Richet, Noisy expected improvement and online computation time allocation for the optimization of simulators with tunable fidelity. Proc. Engineering Optimization, 2010.
[5] H. P. Schwefel. Numerical Optimization of Computer Models. John Wiley and Sons, 1981.
[6] R. Sutton and A. Barto. Reinforcement learning: An introduction. MIT press, 1998.
[7] N. Pavlidis, D. Tasoulis, N. Adams, and D. Hand. Dynamic multi-armed bandit with covariates. Proc. of 18th European Conf. on Artificial Intelligence, pages 777–779, 2008.
[8] A. Sykulski, N. Adams, and N. Jennings. On-line Adaptation of Exploration in the One-Armed Bandit with Covariate Problem. Proc. of 9th Int. Conf. on Machine Learning and Applications, 2010.