

OPTIMISATION OF PROCESS EQUIPMENT USING GLOBAL SURROGATE MODELS

D.W. Stephens^{1*} and P.D. Fawell²

¹Applied CCM Pty Ltd, Dandenong North, Victoria, AUSTRALIA

²Parker Centre (CSIRO Process Science and Engineering), Perth, Western Australia, AUSTRALIA.

*d.stephens@appliedccm.com.au

ABSTRACT

Computational Fluid Dynamics (CFD) modelling offers powerful simulation capability, but only provides information for a given design and set of process conditions. While it allows comparison between different outcomes, CFD results alone do not provide an optimised outcome.

The computational cost associated with the use of high-fidelity CFD models poses a serious impediment to the successful application of optimisation algorithms in engineering design. Advances in hardware and parallel processing have reduced costs by orders of magnitude over the last few decades, but the fidelity with which engineers desire to model systems has also increased considerably. Evaluation of such models may take significant computational time for complex geometries.

In many design problems, thousands of function evaluation may be required to undertake an optimisation study. As a result, CFD models are often impractical for design optimisation. In contrast, surrogate models are compact and cheap to evaluate (order of seconds or less) and can, therefore, be easily used for such tasks.

This paper applies surrogate modelling techniques to a CFD model of sediment transport in a raked industrial thickener. The global surrogate produced using radial basis functions is used to demonstrate single and multi-objective optimisation for the case study. For multi-objective problem, use from a practical design perspective of the information contained in the set of optimal solutions is illustrated through four examples.

NOMENCLATURE

f	function
\mathbf{f}	function data
\hat{f}	model of f
k	number of folds
N	integer number
s	real number
x	input sample points
y	true model output
\bar{y}	mean true model output
\tilde{y}	predicted model output
β	RBF model parameters
ϵ	model error

Subscripts

i index

INTRODUCTION

For many industrial fluid dynamics problems, it is impractical to perform experiments on the physical world directly. Instead, complex, physics-based simulation codes are used to run experiments on computer hardware. Accurate, high-fidelity Computational Fluid Dynamics (CFD) models are typically time consuming and computationally expensive, a serious impediment to the successful application of formal sensitivity analysis in engineering design. While advances in High Performance Computing and multi-core architectures have helped, routine tasks such as visualisation, design space exploration, sensitivity analysis and optimisation quickly become impractical (Simpson et al., 2008; Forrester et al., 2008). As a result, researchers have turned to methods that mimic the behaviour of the simulation model as closely as possible, while being computationally cheaper to evaluate (e.g. Morgans et al., 2007 and Stephens et al., 2011). This work concentrates on the use of data-driven, global approximations using compact surrogate models in the context of computer experiments. The objective is to construct a surrogate model that is as accurate as possible over the complete design space of interest using as few simulation points as possible. Once constructed, the global surrogate model is reused in other stages of the computational engineering pipeline, such as optimisation.

This paper describes the application of surrogate modelling using Radial Basis Functions to a case study of a CFD model of sediment transport in an industrial thickener. Multi-objective optimisation is then conducted on the produced surrogate model, clearly demonstrating the advantages over optimisation by simple comparison of CFD outputs.

SURROGATE MODELLING

Two different approaches have been developed to deal with the problem of costly simulation codes: model-driven and data-driven. Model-driven approximation is commonly known as Model Order Reduction (MOR). Taking a top-down approach, MOR starts with the original model's equations and derives approximations using rigorous mathematical techniques. In contrast, data-driven approximation, also known as Reduced Order Modelling (ROM), takes a bottom-up approach. The exact, inner working of the simulation code is not assumed to be known; only the input-output behaviour is important. A model is constructed based on modelling the simulation code response to intelligently chosen input configurations. The application of surrogates in this paper concentrates on the use of the data-driven approach.

Surrogate modelling can be seen as a non-linear inverse problem for which the aim is to determine a continuous function (f) of a set of design variables from a limited amount of available data (\mathbf{f}), as illustrated in Figure 1. While deterministic in nature, \mathbf{f} can represent exact evaluations of the function f or noisy observations and in general cannot carry sufficient information to uniquely identify f , which may allow multiple surrogates to be consistent with the available data (Figure 2). Thus, surrogate modelling deals with the twin problems of: (a) constructing a model (\hat{f}) from the available data \mathbf{f} , and (b) assessing the errors (ε) attached to it.

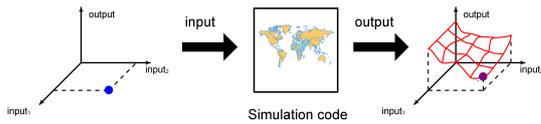


Figure 1: Surrogate model input and output relationship.

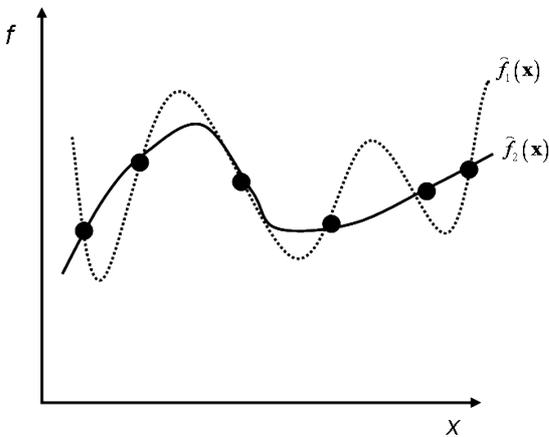


Figure 2: Multiple surrogates may be consistent with the data.

For a given problem there may be a preference for one or the other, based upon expert knowledge of the function f or the simulation code.

Hence, using the surrogate modelling approach the prediction of the simulation-based model output is formulated as $f_p(\mathbf{x}) = \hat{f}(\mathbf{x}) + \varepsilon(\mathbf{x})$.

Model-driven modelling has the advantage of staying true to the “real” simulation model since physical laws are conserved. In contrast, the advantage of data-driven modelling is its generality; it can be applied to any problem where the process can be described as a data generating “black box”. This is useful when the available simulation code is such that it cannot be altered in a domain-specific way. In addition, often it is the global input-output behaviour that is important. In this case, all the approximation intricacies of the different subsystems in the global system no longer play a role and a full MOR may not be worthwhile.

APPLICATION OF SURROGATE MODELS

An important distinction must be made between two different application areas of surrogate models. The first is by far the most popular and involves constructing *local surrogates* that are accurate in local regions of the parameter. Alternatively, a *global surrogate* describes the global behaviour of the system. Here the surrogate is

tuned to mimic the underlying models as closely as needed over the complete parameter space. Such surrogates are a useful, cheap way to gain insight into the global behaviour of the system. The global approach has been utilised in this paper.

Local surrogates

The most common use for local surrogates is to guide the search towards a global optimum, an approach often referred to as surrogate-assisted optimisation (SAO) or surrogate-based optimisation (SBO). Once the optimum is found the surrogates are discarded as they only contain accurate data in the vicinity of local optima of the system and serve no further use. Another drawback to local surrogates is that the output function is generally the objective function of the optimisation problem, which can be a complicated function of the simulation outputs and hence the actual simulation outputs are lost in the objective function. This can be overcome by storing the simulation outputs during the surrogate construction phase.

Global surrogates

Global surrogate models are particularly useful for design space exploration, sensitivity analysis, visualisation, and *what-if* analysis. If optimisation is the goal, it can be argued that a global model is less useful, since significant time savings could be achieved if more effort were directed at finding the optimum rather than accurately modelling regions of poor designs. This paper will demonstrate where global surrogates are very useful for multi-objective optimisation.

It could be argued that to obtain an accurate global surrogate, numerous simulations still need to be performed, with the same problem of computational cost. However, this is not the case since: (1) constructing a global surrogate is a one-time, up-front investment, (2) adaptive modelling and sampling (Gorissen et al., 2010) can drastically decrease the required number of data points to produce an accurate model and (3) the global surrogate can often be reused for various analyses such as sensitivity analysis, optimisation and design tools.

CONSTRUCTION OF THE SURROGATE MODEL

Accurate surrogate modelling of an expensive simulation code involves a number of design choices that need to be addressed.

Sampling strategy

In general, theory covering the topic of sampling strategy is referred to as Design and Analysis of Computer Experiments (DACE). When dealing with computer experiments, the data is typically deterministic and noise-free. Thus, in computer experiments there is no need to do replicates (as opposed to physical experiments) and the classical theory of optimal design does not apply. Instead, a major concern is to create an experimental design which can sample the complete design space in a representative way with a minimum number of samples.

Mathematically, this means that in order to construct a surrogate model ($\hat{f}(\mathbf{x})$), a dataset needs to be populated, for which two approaches exist. In the first, the size and distribution of the dataset is chosen up-front. Common sampling methods that fit in this category are uniform

grid, factorial design and Latin hypercube (Santner et al., 2003). However, since the simulation code used for evaluating the dataset is expensive to compute, it becomes important to avoid unnecessary simulations. At the same time, the complexity of the response surface is not known up-front and defining a data distribution *a priori* is difficult. The second approach involves selecting the data points iteratively, at locations where the information gain will be the greatest. This process is called Adaptive sampling or Active Learning (Gorissen et al., 2010). A large variety of Adaptive sampling methods have been developed, e.g. the Expected Improvement sampling strategy (Jones et al., 1998 and Schonlau, 1997) and the LOLA-Voronoi sequential design strategy (Crombecq et al., 2009).

Modelling strategy

Popular surrogate model types include Response Surface (RS) models, Radial Basis Function (RBF) models, Artificial Neural Networks (ANN), Support Vector Machines (SVM) and Kriging models. Different model types are preferred for different application domains. However, in general there is no theory that can be used *a priori* for selection of a model type. This work has exclusively used Radial Basis functions with a Gaussian basis function, as they were found to be the best suited to the case studied. Radial basis function (RBF) models use linear combinations of radially symmetric functions to interpolate samples data points. A RBF with a Gaussian basis function models is

$$\hat{f}(x) = \sum_{i=1}^N w_i e^{-\gamma_i (\|x-x_i\|)} \quad (1)$$

Where $\|\cdot\|$ is the Euclidean distance, N is the total number of sample points, x_i is the i^{th} sample point, and the γ 's are model parameters (often set to a constant) and the w 's are weighting factors found solving a linear system of N equations. RBF models are shown to produce a good fit for arbitrary contours (Powell, 1987).

Model accuracy assessment and cross-validation

For global modelling it is essential to assess the accuracy of a surrogate for prediction before it is used, usually by comparing some response data produced by the analysis code (CFD output) with corresponding response data predicted by the model. A metric is calculated based on these two sets of values (such as root-mean-squared error) to quantify the degree of accuracy.

After a surrogate model is built, a straightforward way to assess its accuracy is to run the analysis code on a large additional set of sample points and compare the output with those predicted by the model (Wang and Lowther, 2006). However, this comes at a high cost, as CFD simulations may take significant computational time. The cross-validation method can estimate the model's accuracy without requiring any additional sample points. In general, the data is divided into k subsets (k -fold cross-validation) of approximately equal size. A surrogate model is constructed k times, each time leaving out one of the subsets from training, and using the omitted subset to compute the error measure of interest. The average of the error measure is calculated from the result in each of the k iterations.

Cross-validation works by estimating a model accuracy measure with only a limited number of samples points. The computational cost of cross-validation (i.e. generating k surrogates) is justified by the fact that model fitting and assessment usually takes only seconds to minutes, while a single run of the CFD code could last hours to days, if not longer.

While cross-validation almost always performs very well (Bengio and Chapados, 2003), Gorissen et al. (2009a) showed that it is not always efficient at preventing unwanted ripples or bumps in the final model response. A new measure called Linear Reference Model (LRM) has been developed by Gorissen (2010) that helps reduce this problem. Any visible behavioural complexity (bumps, ripples, etc) should be explainable by data points at (or near) those locations. Thus, intuitively, if there is no such evidence nearby, the bump should be regarded as a model artefact. The LRM metric is based on the core assumption that if nothing else is known, the model behaviour between two neighbouring points should be linear. This is achieved by penalising a model proportional to how much it deviates from a linear fit.

CASE STUDY

Problem description

CFD modelling of process unit operations is a tool that is being used increasingly within the minerals processing industry to reduce operating and capital costs and increase throughputs. One such unit operation where CFD has been applied is gravity thickening (Johnston et al., 1998; Fawell et al., 2009; Kahane et al., 2002; Rudman et al., 2008; Owen et al., 2009; Rudman et al., 2010). The principle is simple – adding flocculant to aggregate the fine particles and allowing them to settle to produce clear overflow liquor and concentrated underflow slurry.

Understanding material transport generated by raking in a thickener is important to overall thickener performance. Predictions of velocity vectors, shear rate and residence time throughout the bed region may be obtained from detailed knowledge of the rake geometry and the physical properties of the bed material. Rudman et al. (2008) presented details of a single-phase CFD model developed specifically for the analysis of the material transport aspects of raking within industrial thickeners, including torque and residence time distribution prediction. Such a model can be used to investigate the effect of rake design on material transport within industrial thickeners.

The aim of this study was to build a surrogate utilising the outputs from the CFD model detailed in Rudman et al. (2008). Once built, the surrogate can be used to investigate optimum rake blade angle and speed.

In this case study, the system was a 2 m diameter pilot thickener, 2 m high with a 14° floor angle. The rake was comprised of two arms, each with 5 equi-spaced blades, on a 75 mm diameter rake shaft. The blade angle could be varied from 0-65°, as shown in Figure 3. Each blade was 290 × 75 mm and they were positioned with a gap of 23 mm between the blade and the thickener floor. Slurry is fed into the thickener through a ring manifold at a rate of 3 m³ h⁻¹. The experimental rheogram for the slurry is shown in Figure 4.

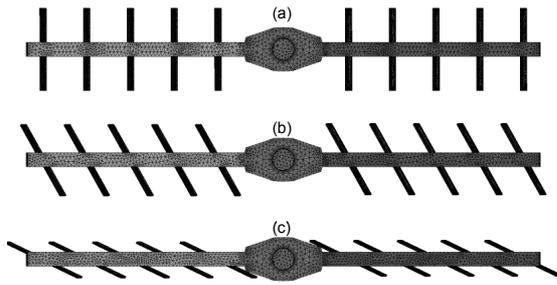


Figure 3: Example rake blade configurations, (a) 0°, (b) 30°, (c) 65°.

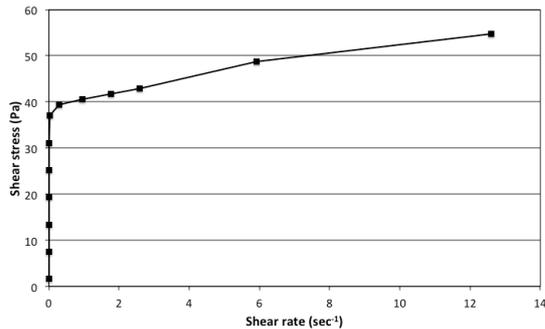


Figure 4: Rheogram of shear stress versus strain rate for the slurry used in the simulations.

The parameters relating to rake design and operation evaluated for their influence on material transport were:

- Rake blade angle (0 - 65°), and
- Rake speed (0.2 - 2 rpm).

CFD computations were performed with ANSYS-CFX (release 12.1) using the sediment transport model described in Rudman et al. (2008). Simulation meshes contained 0.7 to 1.2 million nodes depending on the blade angle. The ANSYS-CFX “high resolution” differencing scheme was used for the convection terms. The simulation is performed in a rotating coordinate frame attached to the rake. In this frame, the geometry is stationary (although the tanks walls rotate), making the simulation simpler to undertake. The bed material is assumed homogeneous and no dewatering is included in the model. The measured rheogram (Figure 4) is used in the simulation with a 1-D interpolation being used for its input. Each simulation took approximately 12 hours using 1 CPU core. The information produced by the rake model is presented in terms of CFD images and calculated parameters. Figure 5 is an example of the graphical output, where a vertical slice plane through the middle of the thickener is presented displaying the residence time.

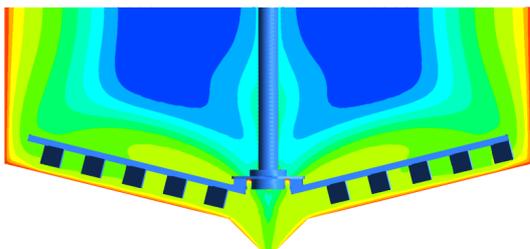


Figure 5: Residence time contours on a plane through the rake (rake speed 1.25 rpm, blade angle 30°). Ten linearly distributed colour bands of blue to red from 0 to 4000 s.

The residence time is calculated using a convection only transport equation. The output quantities from CFD used for the surrogate building were the components of rake torque (pressure and viscous), total torque and a measure of plug flow. For plug flow, the residence time for bed material reaching the outlet should be constant. This means that the standard deviation of the residence time needs to be small. This standard deviation has been divided by the mean residence time (for plug flow) to develop a plug flow metric. A value of zero represents plug flow and values greater than one indicate very poor performance.

Sampling strategy

For this case study, 60 sample points were used to build a surrogate model (Figure 6). Latin hypercube sampling (Santner et al., 2003) with corner inclusion was used for the sample selection. This strategy samples all regions of the design space equally, making it very suitable for modelling with computer analysis code.

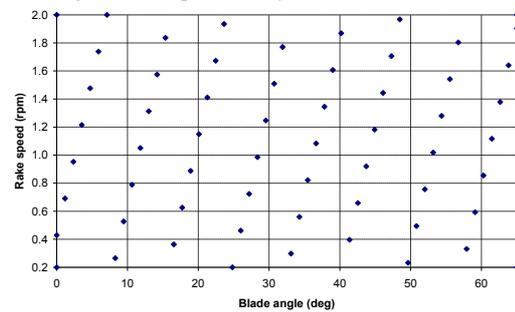


Figure 6: Samples of the input parameters used to build the global surrogate models.

RESULTS

Construction of the global RBF model

All of the RBF surrogate models were built using the SUrogate MOdeling (SUMO) MATLAB toolbox (Gorissen et al., 2010), a plug-in based adaptive tool that automatically generates a surrogate model within the predefined accuracy and time limits set by the user. Different plug-ins are supported: model types, model parameter optimization algorithms, sample selection, and sample evaluation methods.

For the work described here, the toolbox was used with a set control flow as outlined below. A set of samples were chosen as described in the Sampling strategy section. The CFD model was run for each of the chosen sample points to generate the model output for each of these input points. Based upon this set, one or more surrogate models of the chosen type (e.g. RBF) are constructed and their parameters optimized using a Genetic Algorithm (GA). Models are assigned a score based on an equal weighting of two measures (e.g. cross-validation and LRM). The score for each model is used as the fitness to drive the GA to a good optimum in the model parameter optimization landscape. The optimization continues until one of the following three conditions is satisfied: (1) no further improvement is possible, (2) the maximum allowed time has been reached, or (3) the user required accuracy has been met. The maximum time allowed for the GA was 1 hour. In all cases this was longer than necessary, since the best solution was found after a few generations.

The GA as implemented in the MATLAB GADS toolbox was used to search the parameter space of possible models. The RBF model has a set of hyper-parameters; the parameter space was searched for the optimal basis function combination and the parameters of each selected basis function.

The population size of the GA is set to 10 as is the maximum number of generations. The GA cross over fraction was set to 0.7.

The error function used to measure the fitness of each model is defined as:

$$Error(y, \tilde{y}) = \frac{\sum_{i=1}^N |y_i - \tilde{y}_i|}{\sum_{i=1}^N |y_i - \bar{y}|} \quad (2)$$

where y_i , \tilde{y}_i , \bar{y} are the true, predicted and mean true response values, respectively. The metric used to drive the hyper-parameter optimisation was an equally weighted combination of the error function in (2) on a 5-fold cross-validation and the LRM measure.

Once the RBF model has been built, a 20-fold cross-validation was used to measure the accuracy of the final model. Results from this cross-validation are presented in Figure 7 for each of the CFD model outputs (a separate model was built). It can be seen that the surrogates for each of the torque outputs have similar level of accuracy. The plug flow metric shows some deviation between the predicted and measured values at the higher end of the range. This is caused by the scarcity of data points with this magnitude of the metric. The predictions in this region could be improved by the addition of further samples into the building of the surrogate.

Once the surrogate models have been created they can be used for many purposes, including visualisation of the relationship of model inputs to outputs, design tools (what-if analysis), sensitivity analysis and optimisation.

Optimisation

The case study discussed in this paper involves a multi-objective optimisation problem (MOOP). There are two competing objectives – minimise rake torque and maximise the plug flow behaviour (minimise the plug flow metric).

Single composite objective function

In the composite case, a single objective function using combinations of the individual objective functions is constructed (minimise the rake torque and minimise the plug flow metric). To illustrate this approach, three composite objective functions have been selected - (i) more preference is given to (80:20), (ii) equal preference to minimise torque and plug flow metric (50:50) and (iii) more preference given to minimising the plug flow metric (20:80). In each of these cases the torque and plug flow metric were scaled by the maximum value found in the entire design space, before forming the weighted composite objective function. Contour plots of the three objective functions are shown in Figure 8, with the minimum marked with a white star in each of the cases.

The optimum values shown in this figure illustrate the subjective nature of the preference-based approach. The next section illustrates a novel technique to finding the Pareto-optimum set using the surrogates.

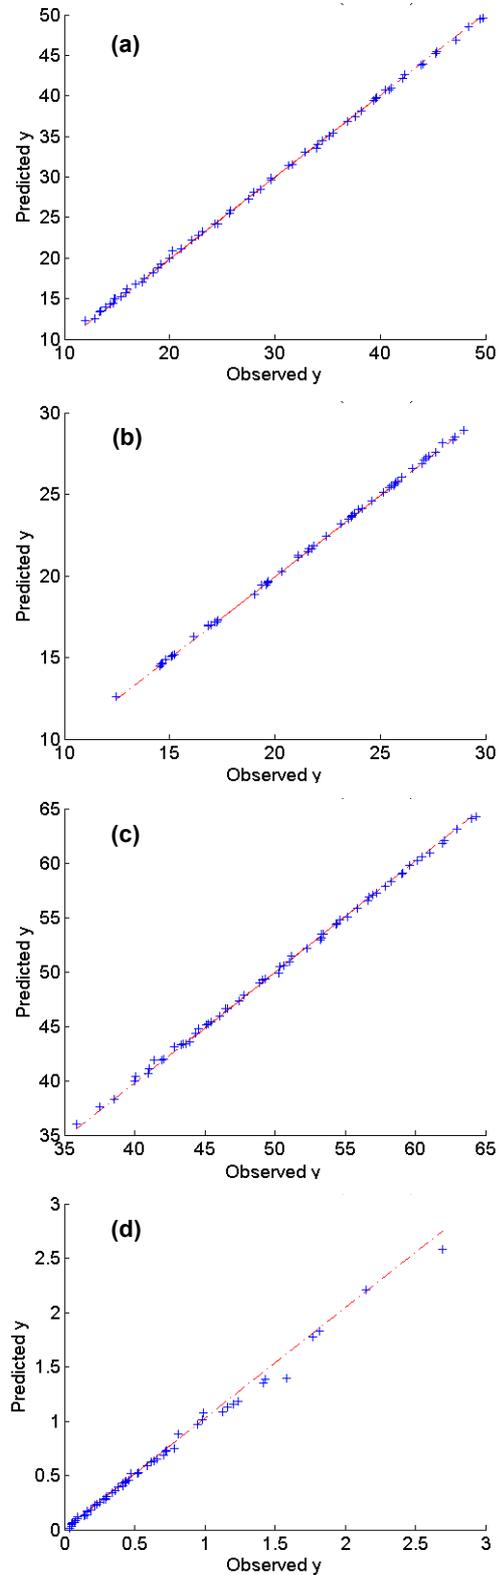


Figure 7: Cross-validated predictions versus actual values for (a) pressure component of torque, (b) viscous component of torque (c) total torque and (d) plug flow metric.

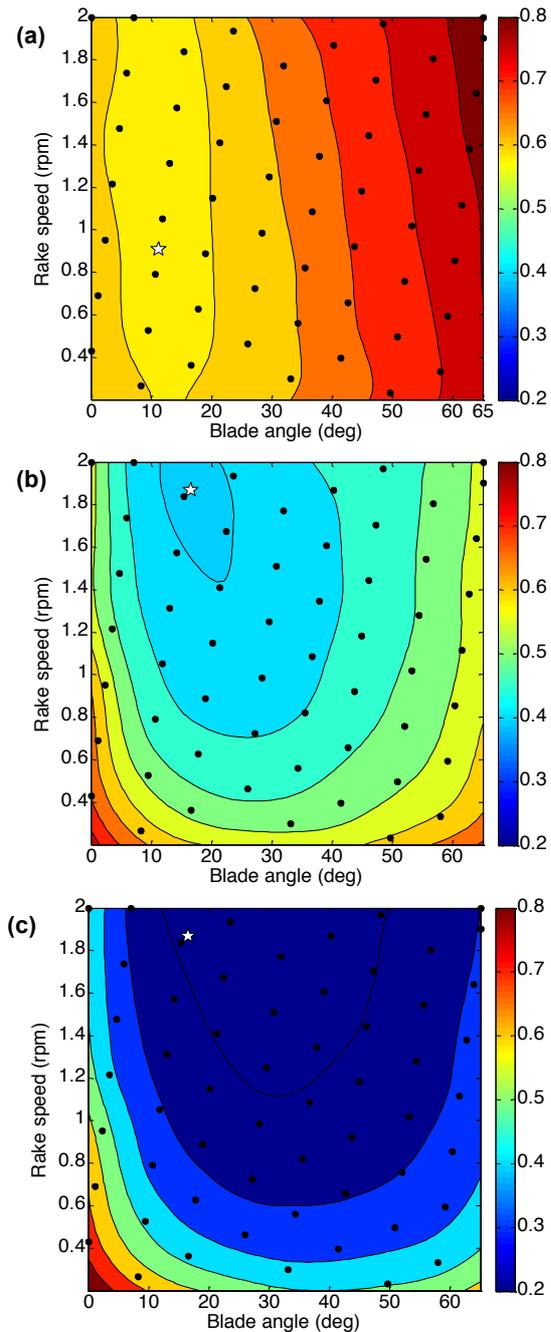


Figure 8: Contour plots showing the three different objective functions and the optimum value (white star) for (a) 80:20, (b) 50:50 and (c) 20:80 preference to torque and plug flow metric, respectively.

Multiple objective functions

A multi-objective optimisation problem deals with more than one objective function. In most practical decision-making problems, multiple objectives are evident. Because of a lack of suitable solution methodologies, a MOOP has been mostly cast and solved as a single-objective optimisation problem in the past (Deb, 2005), as shown in the previous section.

Consider the decision-making involved in designing a thickener rake. Start by taking two extreme hypothetical designs, i.e. one that requires large torque and produces

near plug flow behaviour (A) and one that requires low torque and produces significant non-plug flow (B). If low torque is the only objective of this decision-making process, the optimal choice is solution B. However, if plug flow were the only objective then the optimal choice would be solution (A). This so-called two-objective optimisation problem need not be considered as the two independent optimisation problems, the results of which are the two extreme solutions discussed above. Between these two extreme solutions, there exist many other solutions, where a trade-off between torque and plug flow exists. Thus, between any two such solutions, one is better in terms of one objective, but this betterment comes only from a sacrifice on the other objective. In this sense, all such solutions are optimal solutions to a multi-objective optimisation problem. Often, such trade-off solutions provide a clear front on an objective space plotted with the objective values. This front is called the Pareto-optimal front and all trade-off solutions are called Pareto-optimal solutions. Not all rake designs will lie on the Pareto-optimal front; in fact there are often more sub-optimal configurations than optimal.

In the case of multi-objective optimisation, the design engineer now has a dilemma. Which of these optimal solutions should they choose? If a set of trade-off solutions are already worked out or available, one can evaluate the pros and cons of each of these solutions based on all such non-technical and qualitative (yet still important) considerations and compare them to make a choice. Thus, in a multi-objective optimisation, ideally the effort must be made in finding the set of trade-off optimal solutions (Pareto-optimal-front) by considering all objectives to be important. After a set of such trade-off solutions are found, a user can then apply higher-level qualitative considerations to make a choice.

Rather than using multi-objective optimisation algorithms as discussed by Deb (2005), a novel approach is taken here for finding the Pareto-optimal front. One of the properties of Pareto-optimal fronts is that they always form part of the boundary of the feasible region in objective space. All points in the design variable space can be mapped to the objective space, as shown in Figure 9. Using this knowledge, if one could plot enough points in the objective space, then the Pareto-optimal front could be identified.

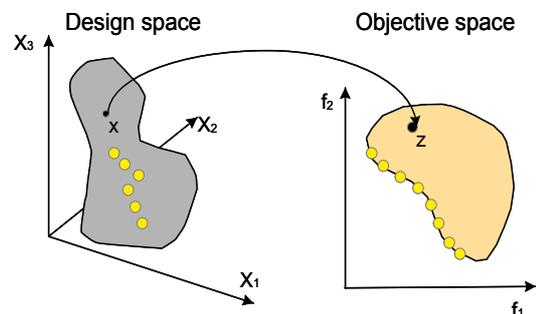


Figure 9: Representation of the design variable space and the corresponding objective space.

Surrogates allow rapid evaluation of model output for given inputs. Therefore, the surrogate model could be evaluated at many different design variable combinations to allow generation of the objective space feasible region and hence locate the Pareto-optimal front. A rectangular

grid comprising of 2600 input point combinations was evaluated for each of the objective functions, with the time taken for this evaluation being less than 1 minute. The objective function values for each of these input points can be plotted in objective space, as shown in Figure 10, with the Pareto-optimal front marked by the red line. The coloured stars on this line indicate the locations of the three optimal solutions found using the different preference composite objective functions. Each blue diamond in this figure represents one design variable combination of blade angle and rake speed.

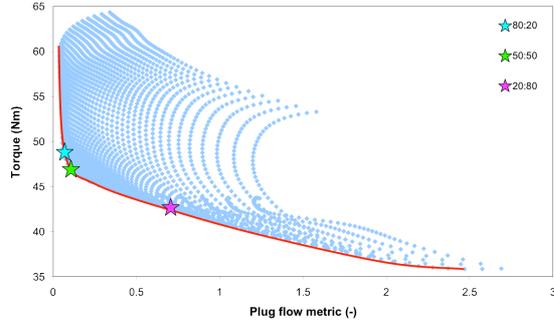


Figure 10: Responses for the rake transport case study plotted in objective space.

Apart from determining optimal solutions, Figure 10 can also be used to determine how to move from a sub-optimal solution to one of the Pareto-optimal solutions. As an example, assume a rake design with a blade angle of 40° and a rake speed of 0.25 rpm. These conditions are placed at the position marked with the red star in Figure 11. Consider the following scenarios:

- (i) Find a design that has the same torque but improves the plug flow behaviour. This is shown by the horizontal arrow in Figure 11. The point of the Pareto front (end of the arrow) can then be mapped back from objective space into design space allowing the blade angle and rake speed to be determined. The design point in the Pareto-optimal set for this torque is a blade angle of 18.2° and rake speed of 2 rpm.
- (ii) The plug flow metric is acceptable, but would like to reduce the torque. This is shown by the vertical arrow in Figure 11. The design point in the Pareto-optimal set for this plug flow metric is a blade angle of 10.2° and a rake speed of 0.6 rpm.
- (iii) The blade angle cannot be changed. Are there any possible improvements by changing the rake speed? To investigate this, information is extracted from the global surrogates for rake torque and the plug flow metric at a fixed blade angle of 40° ; this is plotted in objective space, shown by the green line in Figure 12. Following the green line to the left of the star is for increasing rake speed, but the torque increases as well. If rake torque is not an issue then the optimal speed for this blade angle is 1.75 rpm.
- (iv) The rake speed cannot be changed. Are there any possible improvements that can be made by changing the blade angle? Again, information is extracted at the fixed rake speed of 0.25 rpm from both global surrogates of the rake torque and rake speed and plotted in objective space (magenta line in Figure 12). In this case we can see that the

current blade angle is giving approximately the best achievable plug flow metric for this speed. Moving upwards from the red star is for increasing blade angles and downwards for decreasing angles. The only possibility in this case is to decrease the rake torque whilst at the same time accepting deterioration in the plug flow behaviour.

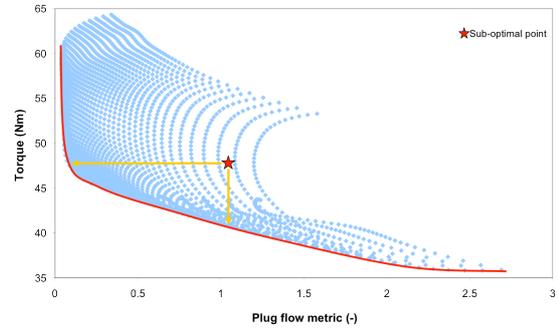


Figure 11: Objective space illustrating design scenarios (i) and (ii).

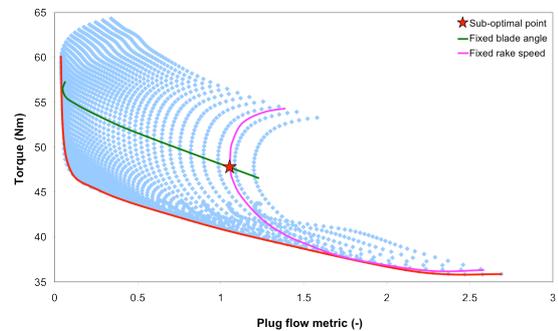


Figure 12: Objective space illustrating design scenarios (iii) and (iv).

CONCLUSION

A Radial Basis Function (RBF) surrogate model has been described, along with the approach to optimising its parameters to make a globally relevant model. The output from a CFD model of sediment transport in a raked thickener was used to produce surrogate models representing rake torque and plug flow behaviour. Examples were provided demonstrating how these surrogate models can be used to optimise the rake design.

The most important component of this case study was the demonstration of how a global surrogate could be utilised for multi-objective optimisation. Solution of multi-objective problems is the ultimate goal in design optimisation, as very few real world problems are single objective problems. For the rake transport problem, the full set of trade-off optimal solutions could be found. Utilisation of the information contained in this set of optimal solutions from a practical design perspective was illustrated through four examples.

It has been demonstrated through the case study how surrogate models can enable the realisation of the full potential of CFD models, and the effort required does not impede the practical use for engineering design tasks such as optimisation. It must be stressed that surrogate models are only as good as the underlying CFD model from which they are created. If the CFD model is not accurate and validated, neither will the surrogate model.

ACKNOWLEDGEMENTS

This work was conducted via the AMIRA project P266F: "Improving Thickener Technology". The authors acknowledge the financial support given to this research from the project sponsors.

REFERENCES

- BENGIO, Y. and CHAPADOS, N. (2003). "Extensions to metric based model selection", *Journal of Machine Learning Research*, **3**, 1209-1227.
- CROMBECQ, K., GORISSEN, D., TOMMASI, L. and DHAENE, T. (2009). "A novel sequential design strategy for global surrogate modelling", Proceedings of the 41st Winter Simulation Conference.
- DEB, K. (2005). "Multi-objective Optimisation", In BURKE, E. K. and KENDALL, G. (Eds.), *Search methodologies: Introductory tutorials in optimization and decision support techniques*. Berlin: Springer-Verlag, pages 273-316.
- FAWELL, P.D., FARROW, J.B., HEATH, A.R., NGUYEN, T.V., OWEN, A.T., PATERSON, D., RUDMAN, M., SCALES, P.J., SIMIC, K., STEPHENS, D.W., SWIFT, J.D. and USHER, S.P. (2009). "20 years of AMIRA P266 'Improving Thickener Technology' - how has it changed the understanding of thickener performance?", Proceedings of the 12th International Seminar on Paste and Thickened Tailings, 21-24 April 2009, Vina del Mar, Chile. JEWEL, R., FOURIE, A., BARRERA, S. and WIERTZ, J. (Editors), Australian Centre for Geomechanics: Nedlands, Australia, pages 59-68.
- FORRESTER, A., SOBESTER, A. and KEANE, A. (2008). *Engineering Design Via Surrogate Modelling: A Practical Guide*. Wiley.
- GORISSEN, D. (2010). "Grid-Enabled Adaptive Surrogate Modelling for Computer Aided Engineering", PhD thesis, Ghent University.
- GORISSEN, D., TOMMASI, L.D., CROMBECQ, K. and DHAENE, T. (2009a). "Sequential modelling of a low noise amplifier with neural networks and active learning," *Neural Computing and Applications*, **18**, 485-494.
- GORISSEN, D., CROMBECQ, K., COUCKUYT, I., DHAENE, T. and DEMEESTER, P. (2010). "A surrogate modelling and adaptive sampling toolbox for computer based design", *Journal of Machine Learning Research*, **11**, 2051-2055.
- JOHNSTON, R.R.M., SCHWARTZ, P., SIMIC, K., NEWMAN, M., FARROW, J.B. and SWIFT, J.D. (1998). "Improving thickener performance at the Pasminco Hobart smelter", in: DUTRIZAC, J.E., GONZALEZ, J.A., BOLTON, G.I. and HANCOCK, P. (Editors), *Zinc and Lead Processing*, Met. Soc. CIM, 697-713.
- JONES, D.R., SCHONLAU, M. and WELCH, W.J. (1998). "Efficient global optimization of expensive black-box functions", *Journal of Global Optimization*, **13**, 455-492.
- KAHANE, R., NGUYEN, T. and SCHWARZ, M.P. (2002). "CFD modelling of thickeners at Worsley Alumina Pty Ltd", *Applied Mathematical Modelling*, **26**, 281-296.
- MORGANS, R.C., DOOLAN, C.J. and STEPHENS, D.W. (2007). "Derivative Free Global Optimisation of CFD Simulations", Proceedings of 16th Australasian Fluid Mechanics Conference, 2-7 December, Gold Coast, Australia.
- OWEN, A.T., NGUYEN, T.V. and FAWELL, P.D. (2009). "The effect of flocculant solution transport and addition conditions on feedwell performance in gravity thickeners", *International Journal of Mineral Processing*, **93**, 115-127.
- POWELL, M.J.D. (1987). "Radial basis functions for multivariable interpolation: A review, in Algorithms for Approximation". Oxford, U.K.: Oxford Univ. Press, 105-210.
- RUDMAN, M., SIMIC, K., PATERSON, D.A., STRODE, P., BRENT, A. and ŠUTALO, I.D. (2008). "Raking in gravity thickeners". *Int. J. Miner. Process.*, **86**, 114-130.
- RUDMAN, M., PATERSON, D.A. and SIMIC, K. (2010). "Efficiency of raking in gravity thickeners". *Int. J. Miner. Process.*, **95**, 30-39.
- SANTNER, T. J., WILLIAMS, B. J. and NOTZ, W.I. (2003). *The Design and Analysis of Computer Experiments*, New York: Springer-Verlag.
- SCHONLAU, M. (1997). "Computer experiments and global optimization". Ph.D. thesis, University of Waterloo.
- SIMPSON, T.W., TOROPOV, V., BALABANOV, V. and VIANA, F.A.C. (2008). "Design and analysis of computer experiments in multidisciplinary design optimization: a review of how far we have come or not", Proceedings of the 12th AIAA/ISSMO Multidisciplinary Analysis and Optimisation Conference, 10-12 September, Victoria, British Columbia, Canada.
- STEPHENS, D.W., GORISSEN, D., CROMBECQ, K. and DHAENE, T. (2011). "Surrogate Based Sensitivity Analysis of Process Equipment", *Applied Mathematical Modelling*, **35**, 1676-1687.
- WANG, L. and LOWTHER, D.A. (2006). "Selection of Approximation Models for Electromagnetic Device Optimization", *IEEE Transactions on Magnetics*, **42**, 1227-1230.

