ECCOMAS Congress 2016
VII European Congress on Computational Methods in Applied Sciences and Engineering
M. Papadrakakis, V. Papadopoulos, G. Stefanou, V. Plevris (eds.)
Crete Island, Greece, 5–10 June 2016

VARIABLE FIDELITY MODELLING IN MODERN AIRCRAFT DESIGN

Marian A. Zastawny[1]

[1] Aircraft Research Association Ltd.
Manton Lane, Bedford, MK41 7PF, UK
e-mail: mzastawny@ara.co.uk

Keywords: Aircraft Design, Set-Based Design, Surrogate Modelling, Variable-Fidelity Modelling, Krigging

Abstract. *The key to successful application of surrogate modelling is the ability to quickly generate high quality data over a large parameter space. This can be achieved through application of Variable Fidelity Modelling (VFM), which is a data fusion technique that allows data from different sources to be combined. VFM can be used to correct a dataset obtained by a simple and fast method with a small number of high quality samples in order to generate a high-fidelity surrogate at a marginal cost.*

This paper presents the results of a feasibility study of a VFM technique based on Hierarchical Kriging in the context of aerospace applications. A generic process for generation of a high quality surrogate using VFM and assessment of its accuracy is presented and validated for a number of realistic use case scenarios. Investigation of the impact of various parameters of the method is conducted, highlighting the importance of selection of the appropriate data correlation functions.

The potential of VFM to significantly reduce the computational cost required for generation of a high-fidelity surrogate is demonstrated for the case of a parametric study of wing cruise performance used at the early stages of the aircraft design process. In addition, an application where the VFM method is used to predict the effects of wing section optimization is shown. The VFM approach presented in this paper can be implemented in a number of different scenarios and can offer significant improvements for the surrogate generation process by quickly generating high-fidelity response surfaces.

1 INTRODUCTION

Due to the high costs involved, aircraft designers have to ensure that every new aircraft meets the required safety and performance standards long before it enters service. In practice, this is achieved during a long development process, which involves intensive calculations, tests and numerical simulations. Modern design offices make extensive use of computers, simulating almost every aspect of the aircraft system.

In practice, a significant number of important decisions, which have substantial impact on the design of a new aircraft, are taken at the preliminary design stage when the amount of available information is usually small and a large number of possible configuration variants are still considered. In a conventional approach, the so-called Point-Based Design (PBD), a baseline concept is developed at an early design stage and is subsequently modified to achieve a feasibility, and ideally, optimal design during the design process. Usually, this involves a number of iterations during which the aircraft can undergo significant modifications as more information becomes available. This means that a wrong decision taken at the beginning of the project can lead to problems downstream, once more information becomes available.

An alternative approach, using the Set-Based Design (SBD) paradigm, was proposed in the CONGA (Configuration Optimisation for Next Generation Aircraft) programme [1]. In this methodology, a number of designs are developed concurrently and, as more information becomes available, the unfeasible concepts are removed from the set. By keeping the design space open and reducing the number of additional iterations, SBD enables fast convergence on the designs which satisfy all constraints [2].

The main challenge of SBD is the fact that a large amount of data has to be developed at the early stage of the project. This can prove problematic, especially for a novel aircraft configuration, for which the amount of empirical, legacy data is limited. Even despite the continuous improvements in simulation algorithms and speed of computers, performing extensive simulations of large data sets is still not feasible in practice. A promising alternative is the use of surrogates generated from a reduced number of data points to explore the design space. This can considerably reduce the computational costs; however, a large number of simulations may still be necessary.

A further reduction of the effort required to generate an appropriate surrogate can be achieved if a low-fidelity method is used for its construction. This, however, may result in a loss of data quality. The Variable Fidelity Method (VFM) explored in this paper attempts to take advantage of speed offered by application of the low-fidelity methods, without compromising the accuracy of the final surrogate, by correcting it with a limited set of high-fidelity data points.

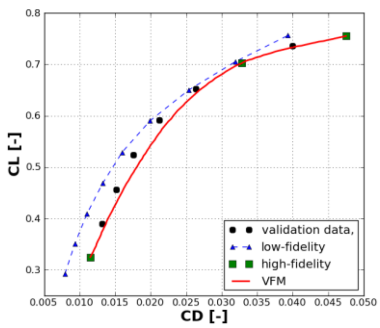

Figure 1: Illustration of the VFM concept.

The concept of VFM is illustrated in Figure 1. In principle, the VFM uses a simple fast method to construct a response surface that characterises the overall trends in the dataset. This is then corrected by a small set of high-fidelity data allowing a high-quality response surface to be obtained at a fraction of the cost required to generate it only from the high-fidelity data.

This paper evaluates the potential applications of VFM in the context of aerodynamic aircraft design. First, the VFM concept developed by DLR [3], which is used in this study, is briefly presented and discussed. This is followed by analysis of the results for four test cases aimed at the investigation of different features of the method. Next, a short discussion on possible application of the VFM approach is given. The paper finishes with short summary and main conclusions.

2 VARIABLE FIDELITY MODELLING

As described in the introduction, the VFM is a data fusion technique that allows the data with different levels of fidelity to be combined in order to generate a high-fidelity surrogate with a limited number of high-fidelity data points. The method discussed in this paper has been introduced in [3] and expanded further in [4]-[6]. This section provides an overview of the technique and outlines the main ideas used.

2.1 Method description

The VFM method discussed in this paper is based on the concept of Hierarchical Kriging (HK). Kriging is a statistical interpolation method introduced by Krige in [7], which models the interpolated variables by a Gaussian process and allows for their prediction based on covariance between the known samples. HK extends this concept to different levels of sample fidelities by adding the next step during which the low-fidelity surrogate is corrected by the high-fidelity data [4].

A low-fidelity surrogate is constructed during the first stage of the HK process. The values predicted by the surrogate can be calculated as:

$$\hat{y}_{lf}(\boldsymbol{x}) = \beta_{0,lf} + \boldsymbol{r}_{lf}^{T}(\boldsymbol{x})\boldsymbol{R}_{lf}^{-1}\left(\boldsymbol{y}_{S,lf} - \beta_{0,lf}\boldsymbol{1}\right) \tag{1}$$

where $\boldsymbol{y}_{,lf}$ are the low fidelity samples, $\boldsymbol{R}_{lf}$ is the correlation matrix representing the correlation between the available samples, $\boldsymbol{r}_{lf}$ is the correlation vector representing the correlation between the untried point, and the available samples and $\boldsymbol{1}$ is the identity matrix**.** $\beta_{0,lf}$ is an unknown constant calculated as $\beta_{0,lf} = \left(\boldsymbol{1}^{T}\boldsymbol{R}_{lf}^{-1}\boldsymbol{1}\right)^{-1}\boldsymbol{1}^{\boldsymbol{T}}\boldsymbol{R}_{lf}^{-1}\boldsymbol{y}_{s,lf}$.

The VFM response surface, which corrects the low-fidelity surrogate with high-fidelity samples, is calculated in the second stage of the process. The resulting high-fidelity predictions are evaluated as:

$$\hat{y}(\boldsymbol{x}) = \beta_0\,\hat{y}_{lf}(\boldsymbol{x}) + \boldsymbol{r}^{T}(\boldsymbol{x})\boldsymbol{R}^{-\boldsymbol{1}}(\boldsymbol{y}_S - \beta_0\boldsymbol{F}) \tag{2}$$

where $\boldsymbol{y}_s$ are the high fidelity samples, $F = \left[\hat{y}_{lf}\left(\boldsymbol{x}^{(1)}\right), \ldots, \hat{y}_{lf}\left(\boldsymbol{x}^{(n)}\right)\right]^{T}$, and $\beta_0 = (\boldsymbol{F}^{T}\boldsymbol{R}^{-1}\boldsymbol{F})^{-1}\,\boldsymbol{F}^{\boldsymbol{T}}\boldsymbol{R}^{-1}\boldsymbol{y}_s$ is the scaling factor, indicating the correlation of the low- and high-fidelity functions.

It is quite clear that the generated response surface will strongly depend on the chosen correlation model, that has an impact on both high- and low-fidelity correlation matrices $\boldsymbol{R}$ and

$\boldsymbol{R}_{lf}$ as well as on the correlation vectors $\boldsymbol{r}$ and $\boldsymbol{r}_{lf}$. The most common correlation model is the Gaussian exponential function which can be written as:

$$\boldsymbol{R}(\boldsymbol{\theta}, \boldsymbol{x}, \boldsymbol{x}') = \prod_{k=1}^{m} \exp(-\theta_k |x_k - x_k'|^{p_k}), \qquad 1 < p_k \leq 2 \tag{3}$$

where θ are unknown hyper-parameters which need to be calculated. Different correlation models, such as Gauss correlation function, cubic splines or radial-basis-functions are also available.

Although the scaling factors β can be determined analytically, they depend on the θ hyper-parameters, which are obtained during surrogate construction process by maximum likelihood estimation. This essentially requires optimisation of the hyper-parameter values in a way that leads to minimising the variances in the surrogate. More information on the process can be found in [4].

2.2 Error estimation

One of the main advantages of the applied method is that it provides an estimate for the root mean-square error (RSME), which can be used to assess the quality of the generated surrogate. The RSME estimate can be used not only to evaluate the statistical accuracy of the generated response surface (e.g. max, mean), but also to highlight the surrogate regions with highest uncertainty in the results. This can be applied for adaptive surrogate refinement, as will be shown later.

It has to be noted, however, that the RSME estimate is based on the mathematical formulation and depends on the chosen correlation model, therefore it might not be accurate. In fact, it is generally found that, although it correctly highlights the surrogate regions with least confidence, the prediction is generally too optimistic when compared to true errors [4]. Hence, an additional test is required to validate the accuracy of the obtained response surface.

Different techniques for surrogate error analysis are available. If a sufficient number of high-fidelity data points is available, a cross-validation can be performed. In such a scenario, a subset of high-fidelity data is identified and removed from the set of points used for the generation of the VFM surrogate. The surrogate prediction can be then compared to the available high-fidelity data in order to determine the values of true errors.

Most often, however, the amount of high-fidelity data is so small that it is desirable to use all available data for generation of the surrogate. In those circumstances the so-called “leave-one-out test” [8] can be adopted. The principle of this technique is to generate several variants of the response surface. For each variant a single high-fidelity point is randomly removed and the true error is calculated by comparison of the surrogate prediction with the high-fidelity data. If the true error remains small for all tested variants, the surrogate can be considered as saturated, i.e. sufficiently accurate so that addition of new high-fidelity points will not lead to significant improvement in the quality of the generated response surface.

2.3 VFM process at ARA

This section describes how the VFM procedure, suggested in [6], was implemented at ARA. Figure 2 illustrates various steps of the process. First, Design of Experiment (DoE) techniques are used to define the samples for low- and high-fidelity data used to generate the VFM surrogate. The aim is to ensure that the entire design space is covered with as few data-points as possible. Usually the number of high-fidelity points is between 5-10% of the number of low-fidelity samples.

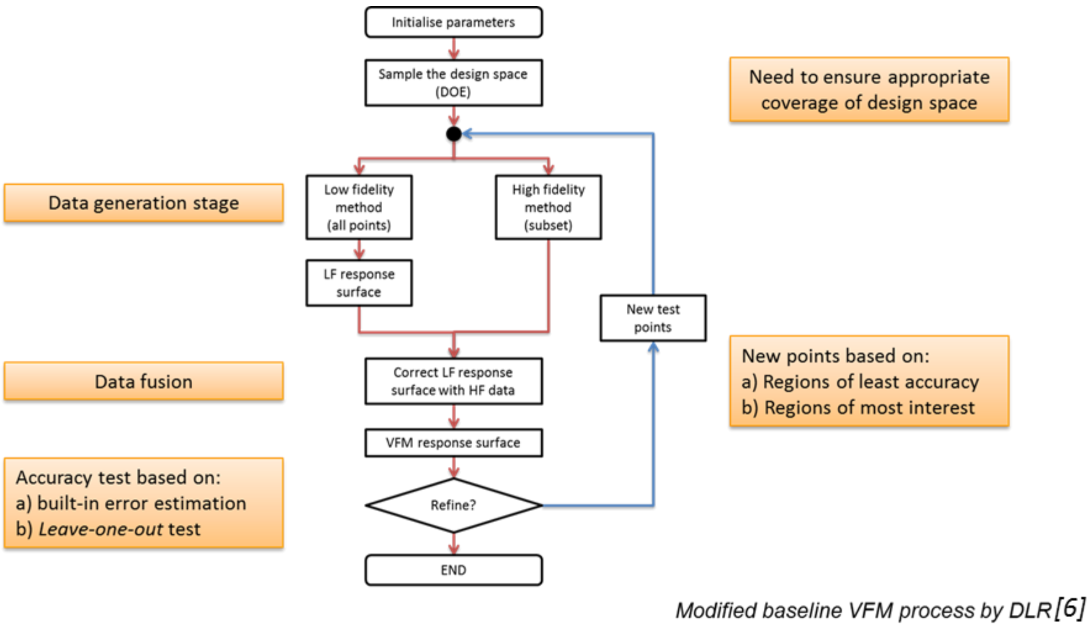

Figure 2: VFM process implemented at ARA.

Next, data for both sets of samples is generated either by simulation or through experiments. In fact, the source of data is irrelevant, as long as the fidelities of different data sets can be easily identified. Once the data are available, a low-fidelity response surface is generated and corrected with the high-fidelity data using the process described above.

The quality of the obtained surface is evaluated based on the VFM RSME prediction. If the estimated accuracy is sufficient, the surrogate is validated using either reference high-fidelity data or by means of the leave-one-out test. If this test is passed as well, the surrogate generation finishes, otherwise additional refinement is performed.

The choice of refinement points depends on the purpose of the generated surrogate, and weighting of errors in different regions of the surrogate. Either way, the new high-fidelity points are generated based on the RSME estimate, usually more than one at a time. Occasionally, if the RSME error prediction is small but the true errors are considerable, additional points can be added based on the true error values. The process concludes once a VFM model of sufficient accuracy is obtained.

3 FEASIBILITY STUDY

This section presents the results of the VFM feasibility study conducted at ARA in order to evaluate potential applications of this technique and offer a better understanding of its limitations. The tests described here focus on VFM applications in the field of aerodynamics and each of them aims at evaluation of a specific feature of the VFM approach.

3.1 Choice of the low-fidelity model

One of the obvious applications of the VFM is to combine the data obtained with methods characterized by different levels of fidelity. In general, the low-fidelity method allows for the results to be obtained quickly, but with a limited level of accuracy. The high-fidelity data, on the other hand, offers a greater accuracy, but requires more effort to be obtained. This test addresses the impact of the choice of the low-fidelity method on the resulting VFM surrogate and how the low-fidelity data affects the required number of high-fidelity correction data points.

The example presented here covers 2D simulations of the flow over a NACA 63_2-215 aerofoil resolved using the BVGK code [9] in two variants:

- low-fidelity – inviscid
- high-fidelity – inviscid with coupled boundary layer

Two scenarios are investigated. First, a lift polar at Mach number $M = 0.5$ is simulated over a linear incidence (α) range from 0 to 5 degrees. Figure 3 shows the lift coefficient as the function of the incidence angle for different levels of fidelity, along with the high-fidelity points used for generation of the VFM. It can be observed that although the low-fidelity data predicts higher values of the lift coefficient Cl, it predicts the same trend as the reference high-fidelity results. A VFM surrogate of a sufficient accuracy is be obtained for all investigated high-fidelity points configurations.

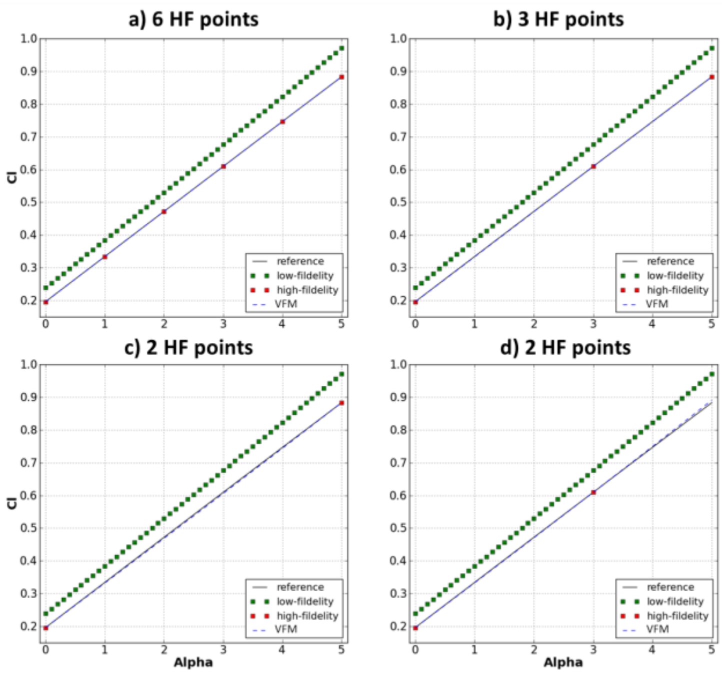

Figure 3: VFM of lift coefficient obtained with different high-fidelity points.

A comparison of error estimation obtained during the VFM generation with the ‘true’ error calculated as the difference between the lift coefficient predicted by the VFM and the reference high-fidelity data is shown in Figure 4. The figure demonstrates that reducing the number of high-fidelity points used to generate the VFM surrogate leads to higher errors. Also, apart from case (c), with 3 high-fidelity points, the error prediction has a similar behaviour to the ‘true’ error although they differ in magnitude.

The second scenario considered during this test was an investigation of a drag polar at $M = 0.6$ over a similar incidence range from 0 to 5 degrees. The resulting drag polars along with the used high-fidelity points are illustrated in Figure 5. Again, a sufficient number of high-fidelity points allows for a VFM surrogate of a reasonable quality to be obtained. On the other hand, using 3 high-fidelity points, which was sufficient in the previous scenario, now results in a VFM of a very poor accuracy. This can be explained by inspection of the low-fidelity data behaviour shown in Figure 6a. Since the low-fidelity model is inviscid, the resulting drag coefficient will be zero everywhere, apart from high incidences where wave drag effects are observed. The low-fidelity prediction of the drag coefficient has therefore a completely different behaviour to the trends observed in the high-fidelity model, thus it actually impacts the VFM surrogate generation in a negative way. In fact, for this scenario, generation of a response surface from high-fidelity points only would result in a better prediction.

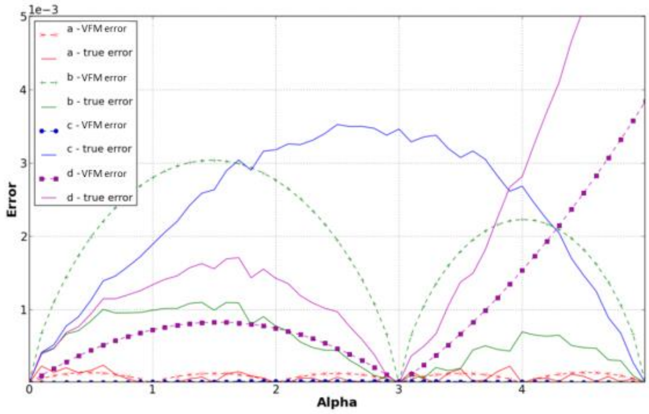

Figure 4: Comparison of errors for the lift coefficient prediction.

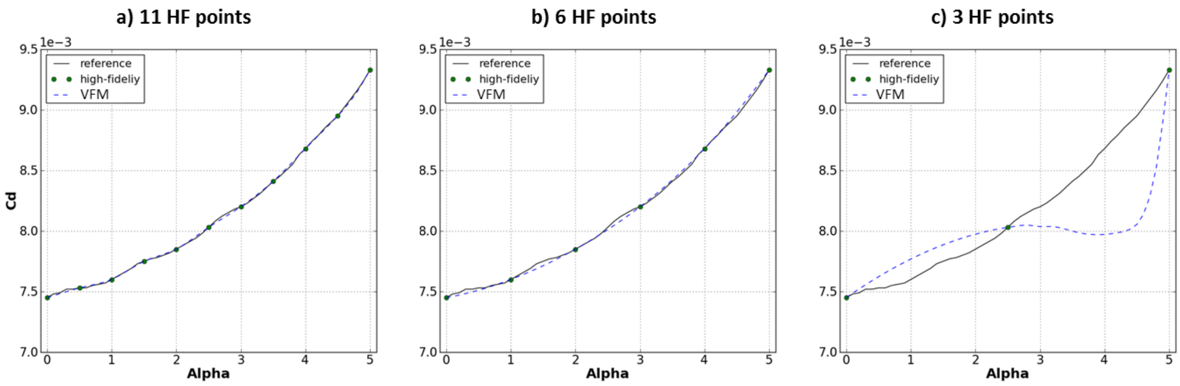

Figure 5: VFM of drag coefficient obtained with different high-fidelity points.

What is encouraging, however, is that the model errors, shown in Figure 6b, illustrate that the built-in VFM error prediction is capable of capturing the poor quality of the generated surrogate and highlighting the areas where improvement is required.

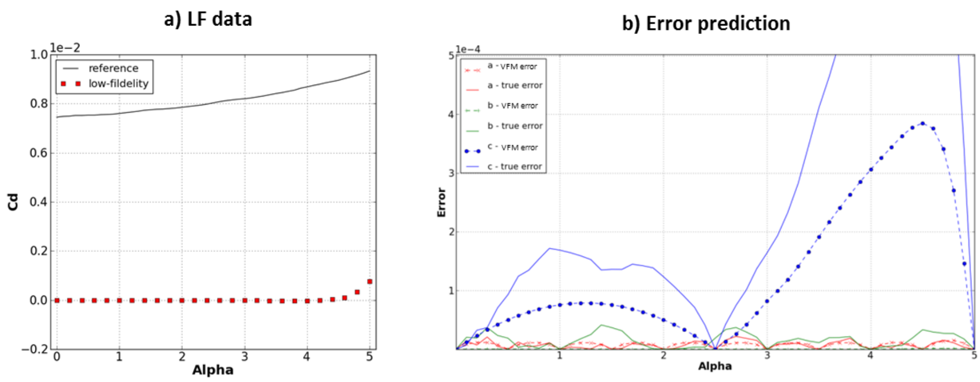

Figure 6: Comparison of errors of lift coefficient predictions.

The main conclusion from this test is that the choice of adequate low-fidelity data is critical to successful application of the VFM. The low-fidelity data should be characterised by a similar behaviour to the trends observed in the high-fidelity model. The second conclusion is that the error estimates obtained using the VFM, although not precisely correct, allow for the assessment of the quality of the surrogate and indicate the parameters for which more high-fidelity information should be provided.

3.2 VFM accuracy and surrogate refinement process

The second feasibility test focused on investigation of the accuracy improvement offered by adoption of the VFM when compared to a response surface generated only from the high-fidelity data points and on the impact of surrogate refinement process on the accuracy of the obtained data.

In this case, the drag coefficient dependency over a range of Mach numbers and incidences for a civil airliner wing was studied. The high-fidelity data was generated from the results of 3D simulations of the wing using the TAU RANS solver [10]. The low-fidelity data was obtained by spanwise integration of the results obtained for 2D simulations of the sections of the investigated wing. The concept of the process is illustrated in Figure 7.

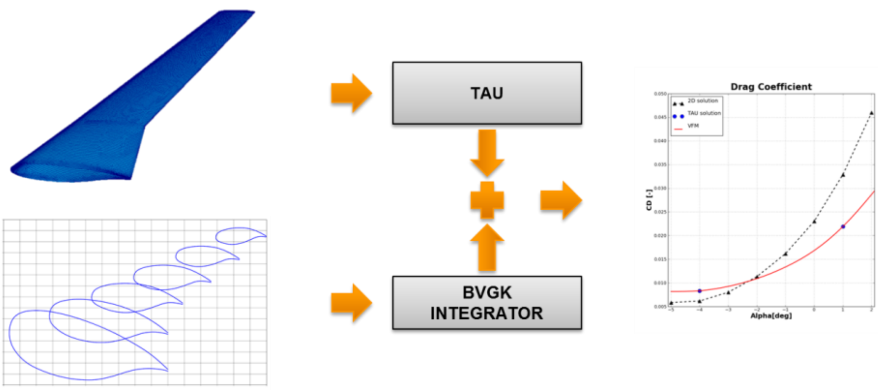

Figure 7: Illustration of the process used for surrogate refinement test.

Figure 8 shows the surrogates of the drag coefficient as a function of incidence angle at various Mach numbers obtained using VFM and by fitting a response surface (RSM) for high-fidelity points only. The initial surrogates were generated with 5 high-fidelity data points obtained by DoE (see highlighted circles in Figure 8a). The surrogates were subsequently refined by addition of new high-fidelity data points in the regions where VFM error prediction indicated the largest values of the relative error (ratio between error and drag coefficient value). The use of relative error was motivated by the fact that, in practice, the low drag coefficient values are of most interest. The surrogates were deemed as saturated, i.e. addition of new high-fidelity points resulted in a little difference in estimated errors, when 19 high-fidelity data points were used. The final surrogates are compared to each other and to the reference data (TAU) in Figure 8b.

The convergence of statistical characteristics (maximum, average, median) of the VFM error estimates is illustrated in Figure 9. A rapid, although non-monotonic, reduction in all statistical values describing the error estimates is observed at the beginning of the refinement process. Change in the relative error remains relatively small when 11 high-fidelity points are used and decreases monotonically when more than 15 points are available. Similar behaviour is observed for the absolute error values, although here, the maximum error remains relatively high.

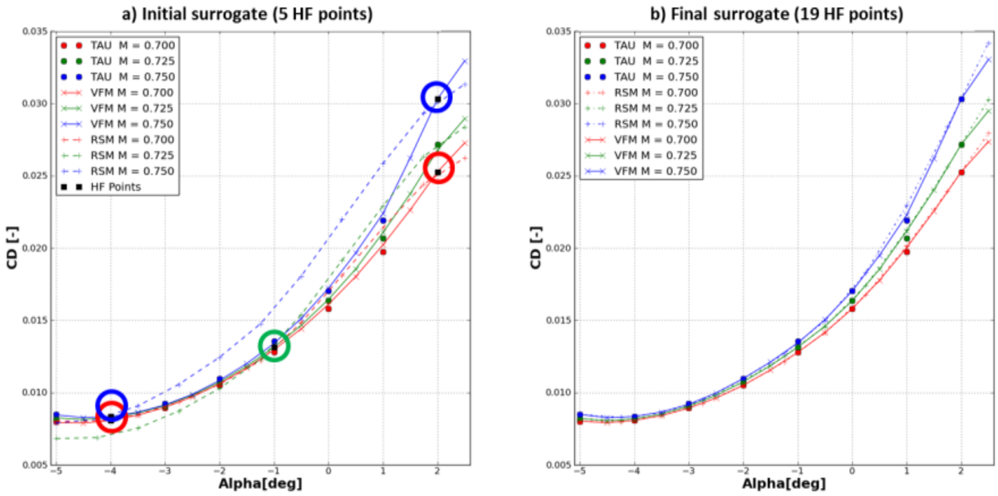

Figure 8: Surrogates of the drag coefficient polars at various Mach numbers obtained using VFM and fitting a response surface (RSM) from high-fidelity points only.

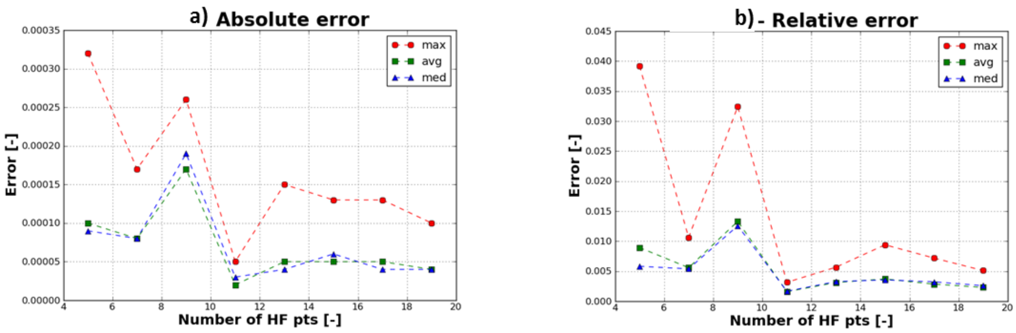

Figure 9: Convergence of the error estimates.

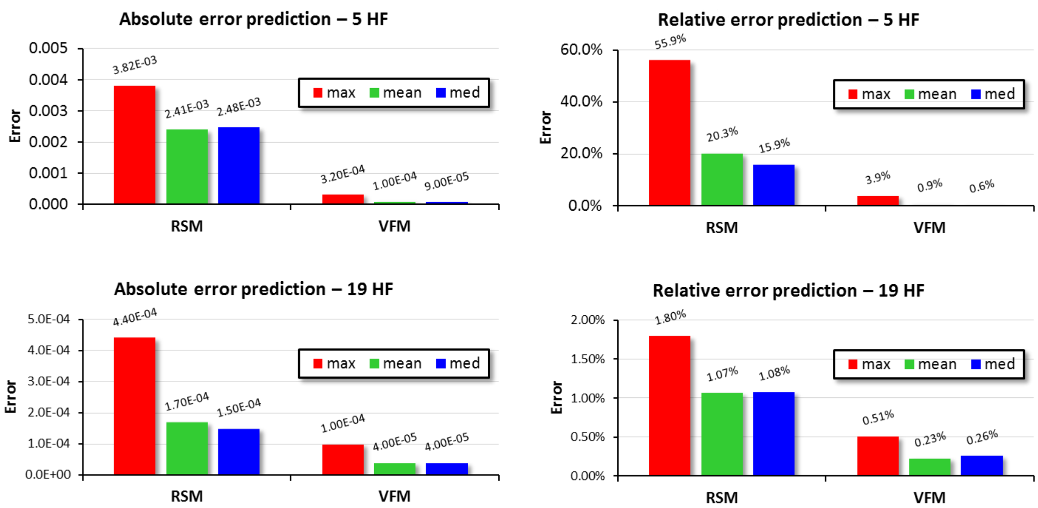

Figure 10: Comparison of absolute and relative error prediction for VFM and RSM.

A comparison of the statistics of the error predictions before and after the refinement process for the VFM and high-fidelity only surrogate (RSM) is shown in Figure 10. Although the error reduction for the RSM is considerably more pronounced, the RSM prediction remains inferior to the VFM. Moreover, the VFM surrogate obtained with 5 high-fidelity data points offers similar accuracy as the refined RSM both in terms of absolute and relative errors. The refinement process, driven by the requirement to minimise the maximum relative error, allowed both the maximum absolute error and maximum relative error to be reduced by a factor of 3. In general the improvement in relative errors is stronger than for absolute errors.

High-fidelity solutions for a number of additional points, not used for surrogate refinement, were obtained for an additional validation stage. This allowed for the real errors, i.e. the difference between VFM prediction of the drag coefficient and the actual simulated value, to be calculated. The results of this validation for the initial and saturated surrogates are demonstrated in Figure 11.

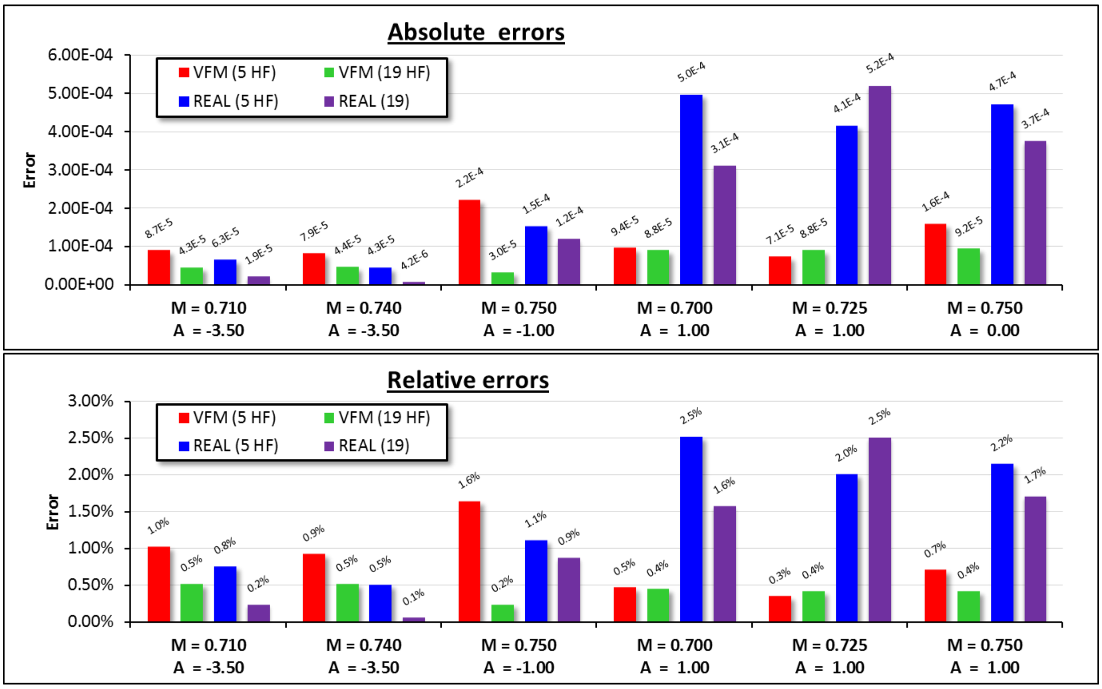

Figure 11: Comparison of the VFM error estimates and the actual prediction errors.

Both absolute and relative errors remain small for the low drag coefficient values, but the real errors are considerably larger in the high incidence range. Although in general, the surrogate refinement led to reduction of both estimated and actual errors, in some cases, i.e. $M = 0.725$ and $A = 1.0$, the errors have increased. No direct correlation between the errors predicted by the VFM and the actual errors can be identified. Nevertheless, the same trends can be observed, i.e. reduction in error estimate corresponds to the reduction of the actual inaccuracy of the prediction.

This test clearly demonstrates that given an appropriate low-fidelity model, VFM can be used to generate a surrogate of a superior quality to a response surface generated from high-fidelity data only. In practice, this means that by adopting the VFM technique the number of required high-fidelity data points can be significantly reduced or the quality of the surrogate can be improved.

The analysis of the error estimates shows how it can be used to drive the refinement process in the desired way (local vs. global, absolute vs. relative). Although the error estimates generally give a reasonable prediction of the surrogate accuracy, additional tests, using high-fidelity data, are required in order to ensure the quality of the obtained response surface.

3.3 Time benefits enabled by application of VFM

One of the main motivations for application of the VFM technique is to enable generation of a high-quality surrogate at a fraction of the cost required to generate the same surrogate based only on high-fidelity points. One of the main ARA requirements in the CONGA project was to construct a dataset with wing cruise performance (lift and drag coefficients) for a wing family with varying planform areas and spans at a range of incidences (see Figure 12 for illustration).

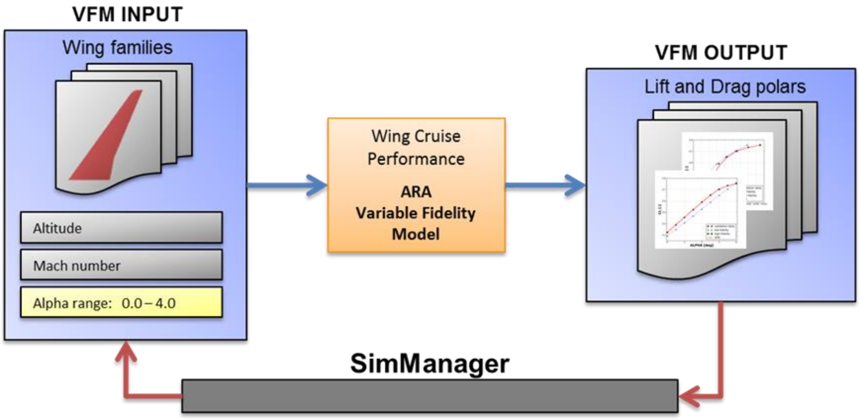

Figure 12: Illustration of the process adopted in CONGA project.

The required dataset was constructed from 108 points over a three-dimensional design space with 3 values of span, 4 values of wing surface area and 9 incidences. The low-fidelity data was computed with the Viscous Full-Potential method (VFP) [11], whereas the high-fidelity results were extracted from the results of TAU simulations. A surrogate refinement process, similar to the one described in previous section was adopted, with lift coefficient surrogate refinement driven by reduction of absolute error estimates, and drag coefficient surrogate driven by the relative error reduction. The leave-one-out test has been applied to ensure a sufficient quality of the obtained surrogates. The final VFM surrogate was constructed from 108 low-fidelity points and 20 high-fidelity results.

On average, 15 minutes were required to generate a single data point with VFP compared to 4 hours on a 20-core cluster needed for a single TAU solution. A comparison of the time needed to generate the VFM surrogate with the time required for 108 TAU simulations is shown in Figure 13. Application of VFM allowed the surrogate of appropriate quality to be obtained in 25% of the time necessary to simulate all 108 data points. In theory, it could be possible to construct a surrogate from a smaller number of high-fidelity data points, but, as already demonstrated in the previous example, it is unlikely that it would have similar quality and would still require significant computational effort.

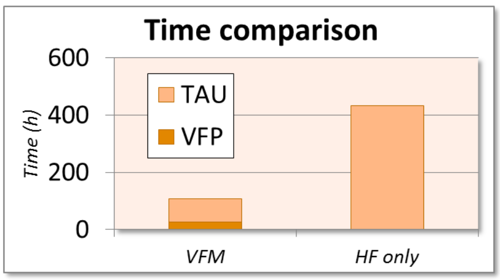

Figure 13: Comparison of the time required to generate the final data set.

This test case clearly demonstrates the time benefits of the application of the VFM technique. Generation of a surrogate has an additional advantage in that it can be used to examine the drag and lift coefficients at intermediate values of the design parameters, at virtually no cost, since the generated VFM response surface is continuous. Moreover, it is easily possible to extend the surrogate by increasing the range of parameters or by addition of new dimensions.

3.4 Impact of the correlation function

The final test presented in this paper investigates the impact of one of the most important parameters of the VFM technique, i.e. the choice of correlation function used to generate the response surface.

In order to assess the impact of the correlation function, surrogates with dependence of the aerofoil performance (L/D) on its thickness-to-chord ration (t/c) are analysed. Unlike in the previous examples, where two different methods were used to generate the solutions, here both low- and high-fidelity data are generated using the same 2D solver – BVGK. The difference in the level of data fidelity is attributed to the different means of obtaining the geometry. The high-fidelity data consists of 3 aerofoils of different thicknesses designed for optimum performance, while the low-fidelity geometries are generated by scaling the baseline section thickness while maintaining a fixed camber line (see Figure 14 for illustration).

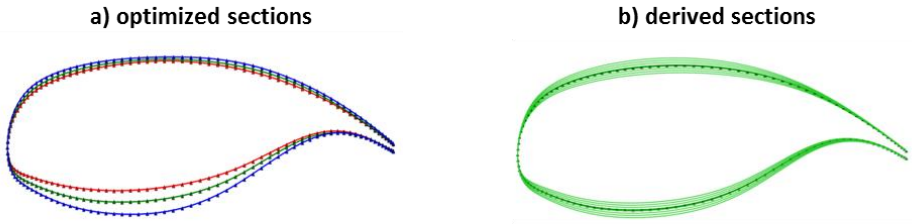

Figure 14: Sections used for generation of VFM surrogate.

This exercise is motivated by the fact that it is not normally possible, due to the time constraints, to optimise sections during the preliminary aircraft design stage. In fact, the wing sections probably do not exist at this point in time and even if some legacy data are available, this might not be suitable for the new aircraft model. Still, ability to predict the potential optimum aerodynamic performance would allow realistic performance targets to be set when sizing the aircraft, without actually performing the time and resource intensive optimisation.

The VFM surrogates generated with two different correlation functions, i.e. Gaussian Exponential function (GEXP) and Multi-Quadratic (MQ) Radial Basis Function (RBD), are shown in Figure 15. Both surrogates use the same low-fidelity data – performance (L/D) characteristics of the varying thickness sections are derived from the optimised baseline aerofoil. Two high-fidelity data points, corresponding to the baseline and thick sections, are used in both scenarios. The thin optimised aerofoil is used as a reference point for validation.

The data presented in Figure 15 highlight significant differences not only in the obtained VFM surrogates but also differences in the low-fidelity response surfaces constructed during the process. While the GEXP low-fidelity response surface closely follows the supplied data points, the MQ surface is smoother and follows the general trend rather than the actual data. Similar behaviour is observed for the generated VFM surrogates, i.e. in the extrapolation region the GEXP surrogate follows the low-fidelity data, whereas the MQ extends the original data trend. The magnitude of estimated errors is also smaller for the case of the MQ correlation function.

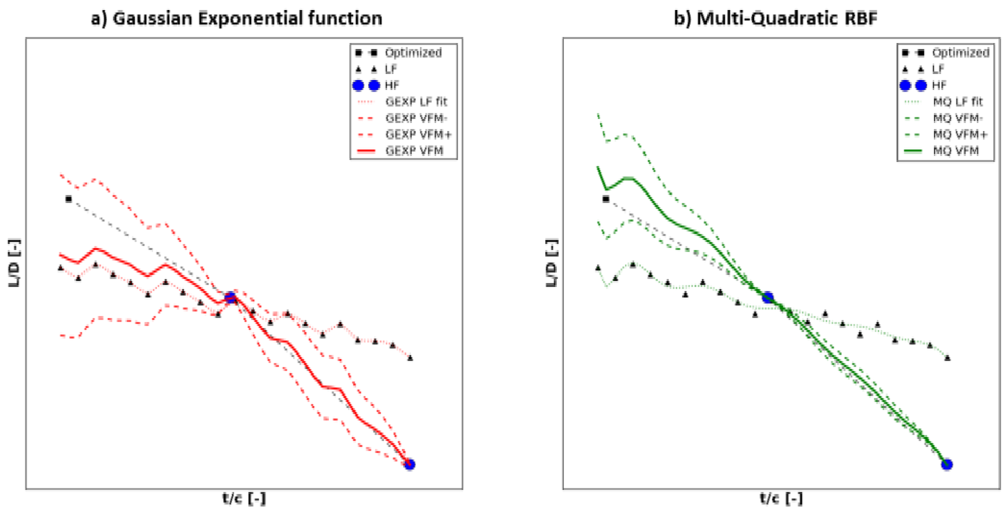

Figure 15: Comparison of the VFM surrogates obtained with different correlation functions.

The results presented for this test case show the importance of the choice of an appropriate correlation function. If the input data is characterised by a high level of noise it is beneficial to use RFB-based correlation functions, as they allow for a smooth surrogate to be generated. On the other hand, GEXP is more localised, so it is a more appropriate choice for data where the general behaviour is not necessarily smooth and local effects need to be considered. In all cases, however, care needs to be taken when extrapolating the data beyond the available high-fidelity data boundaries.

4 POSSIBLE APPLICATIONS OF VFM

The VFM technique has a range of possible applications in the context of modern aircraft design. This section attempts to summarise the various ways in which this method can be adopted. It is suggested to divide the range of possible VFM applications into two main groups, namely data fusion and local surrogate refinement.

The objective of data fusion is to combine two or more datasets with different level of fidelity into a single high-quality response surface. As demonstrated in previous sections, using the VFM not only offers significant reduction of the time needed to generate high quality surrogates from sparse datasets, but also allows for efficient assessment of the surrogate accuracy.

Below is a non-exhaustive list of possible scenarios for data fusion:

- Correcting low-fidelity code with high-fidelity code (e.g. EULER + RANS)
- Correcting partially converged CFD results with fully converged ones
- Combining results with different mesh refinement levels
- Combining wind tunnel data with CFD data
- Correcting CFD for aeroelasticity effects
- Combining results with different levels of geometry maturity

An alternative application of the VFM is to locally refine the surrogate during the optimisation process. Combined with an optimisation algorithm, VFM can be applied to find function minima efficiently (see Figure 16 for illustration). In such a case, the design space is explored with the low-fidelity response surface, while high-fidelity results are used for local refinement of the surrogate. An example process is described below:

1. Find the function minimum with the low-fidelity response surface
2. Run the high-fidelity code on the optimum point
3. Generate the VFM surrogate with the addition of high-fidelity data
4. Use the VFM as response surface in step 1

It is believed that the approach presented above could significantly reduce the time required for optimisation, as only a few high-fidelity simulations would be necessary.

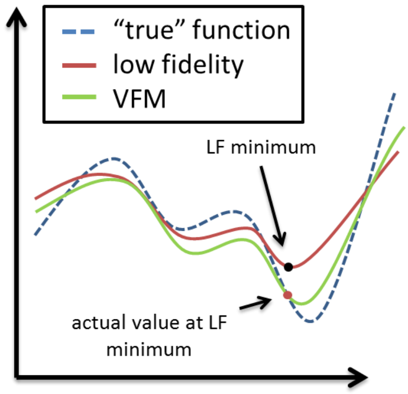

Figure 16: Illustration of VFM used for optimisation.

5 CONCLUSIONS

This paper explored the applicability of using the Variable Fidelity Modelling (VFM) in the context of aerodynamic analysis of modern aircraft designs. The VFM can be used for rapid generation of high-quality surrogate models by combining the results obtained from a low-fidelity method with a limited number of high-fidelity data points. By reducing the number of high-fidelity points required for generation of the surrogate, VFM enables the overall cost of creating the response surface to be significantly reduced.

A wide range of possible applications has been identified, which include, but are not limited to, combining results from simulations at different levels of fidelity, using different levels of input fidelities, as well as using the method to increase efficiency during an optimisation process. The studied examples have clearly demonstrated that applying VFM can offer significant time reduction without loss of surrogate quality, or can be applied to enhance the quality of an existing dataset.

Examination of the presented test cases allowed for the best practice for VFM application to be developed:

- The quality of VFM strongly depends on the trends observed in low-fidelity data
- VFM results are sensitive to the chosen correlation function. Gaussian Exponential Function is generally more robust, however Multi-Quadratic Radial Basis Function can be used to obtain smooth data and is more suitable for extrapolation

- The error estimates returned by the VFM can be used to indicate the error magnitudes and areas of poor surrogate quality, however, they do not offer accurate error values
- The best practice for surrogate construction is to drive the refinement process using the built-in error estimator and use additional techniques (e.g. leave-one-out test) for results verification

REFERENCES

[1] Al-Ashaab, A., Golob, M., Noriega, P., Torriani, F., Alvarez, P., Beltran, A., Busachi, A., Ex-Ignotis, L., Rigatti, C. and Sharma, S. (2013), "Capturing the Industrial Requirements of Set-Based Design for CONGA Framework", Proceedings of the 11th International Conference on Manufacturing Research (ICMR2013), Cranfield University, Cranfield, UK

[2] Sobek II, Durward K., Allen C. Ward, and Jeffrey K. Liker. "Toyota's principles of set-based concurrent engineering." *MIT Sloan Management Review* 40.2 (1999): 67.

[3] Han, Z-H., S. Görtz, and R. Hain. "A variable-fidelity modeling method for aero-loads prediction." *New Results in Numerical and Experimental Fluid Mechanics VII*. Springer Berlin Heidelberg, 2010. 17-25.

[4] Han, Zhong-Hua, and Stefan Görtz. "Hierarchical kriging model for variable-fidelity surrogate modeling." *AIAA journal* 50.9 (2012): 1885-1896.

[5] Han, Zhong-Hua, Ralf Zimmermann, and Stefan Görtz. "A new cokriging method for variable-fidelity surrogate modeling of aerodynamic data." *AIAA Paper* 1225 (2010): 2010.

[6] Han, Zhong-Hua, Stefan Görtz, and Ralf Zimmermann. "Improving variable-fidelity surrogate modeling via gradient-enhanced kriging and a generalized hybrid bridge function." *Aerospace Science and Technology* 25.1 (2013): 177-189.

[7] Krige, D. G. "A statistical approach to some basic mine valuation problems on the Witwatersrand." *Journal of the South African Institute of Mining and Metallurgy* 94.3 (1994): 95-112.

[8] Keane, Andy, and Prasanth Nair. Computational approaches for aerospace design: the pursuit of excellence. John Wiley & Sons, 2005.

[9] Ashill, P. R., Wood R. F., and D. J. Weeks. "An improved semi-inverse version of the Viscous Garabedian and Korn method (VGK)." *RAE TR* 87002 (1987).

[10] Schwamborn, D., Gerhold, T., and Heinrich R.. "The DLR TAU-code: Recent applications in research and industry." ECCOMAS CFD 2006: Proceedings of the European Conference on Computational Fluid Dynamics, Egmond aan Zee, The Netherlands, September 5-8, 2006. Delft University of Technology; European Community on Computational Methods in Applied Sciences (ECCOMAS), 2006.

[11] Arthur M. "A method for calculating subsonic and transonic flows over wings or wing-fuselage combinations with an allowance for viscous effects", 22nd Aerospace Sciences Meeting, Aerospace Sciences Meetings,.