Variable-Fidelity and Reduced-Order Models for Aero Data for Loads Predictions

Stefan Görtz[1], Ralf Zimmermann[1,2], and Zhong-Hua Han[1,3]

[1] German Aerospace Center (DLR), Lilienthalplatz 7, 38108 Braunschweig, Germany
stefan.goertz@dlr.de
[2] TU Braunschweig, Institute Computational Mathematics,
38106 Braunschweig, Germany
ralf.zimmermann@tu-bs.de
[3] Northwestern Polytechnical University, 710072 Xi'an, People's Republic of China
hanzh@nwpu.edu.cn

Abstract. This paper summarizes recent progress in developing metamodels for efficiently predicting the aerodynamic loads acting on industrial aircraft configurations. We introduce a physics-based approach to reduced-order modeling based on proper orthogonal decomposition of snapshots of the full-order CFD model, and a mathematical approach to variable-fidelity modeling that aims at combining many low-fidelity CFD results with as few high-fidelity CFD results as possible using bridge functions and variants of Kriging and Cokriging. In both cases, the goal is to arrive at a model that can be used as an efficient surrogate to the original high-fidelity or full-order CFD model but with significantly less evaluation time and storage requirements. Both approaches are demonstrated on industrial aircraft configurations at subsonic and transonic flow conditions.

Keywords: Variable fidelity modeling (VFM), reduced order modeling (ROM), proper orthogonal decomposition (POD), Kriging, loads, aerodynamics, computational fluid dynamics (CFD).

1 Introduction

Nowadays, computational fluid dynamics (CFD) is intensively being used in aerodynamic design [1] and is becoming more interdisciplinary, helping provide closer ties between aerodynamics, structures, propulsion, and flight controls. Nevertheless, routine use of CFD based on the Navier-Stokes equations to generate the large data bases needed for loads and stability and control applications is not feasible at present [2] and requires orders of magnitude improvement in throughput (productivity) through improvements in geometry and grid generation, improved turbulence models, and improvement in algorithm and hardware performance [3]. However, improvements in CFD alone will not be enough and efficient tools for making rapid but accurate predictions of the aerodynamic loads acting on an aircraft throughout its flight envelope and beyond are sought after [1,4].

N. Kroll et al. (eds.), *Computational Flight Testing*, Notes on Numerical Fluid Mechanics and Multidisciplinary Design 123,
DOI: 10.1007/978-3-642-38877-4_8,

Here, we focus on two approaches that promise to provide a full map of the required aerodynamic data at a minimum number of high-fidelity CFD simulations: reduced-order modeling and variable-fidelity modeling. Reduced-order modeling refers to approximating the governing equations of fluid dynamics using, instead of the millions of degrees of freedom typical of a full-order CFD simulation, only a handful of degrees of freedom, while variable-fidelity modeling refers to numerical methods for combining the results of many low-fidelity CFD simulations with a few high-fidelity CFD results. Both approaches require several solutions of the full-order discrete system upfront to build the model but offer the possibility to compute many approximate solutions, each at a fraction of the cost of a single simulation of the high-fidelity or full-order model. Thus, large increases in computational efficiency are obtained, thereby justifying the initial computational expense of constructing these models and motivating their use for predicting aerodynamic data for loads or for multi-disciplinary analysis and optimization.

The article begins with a brief introduction to the variable-fidelity and reduced-order modeling strategies we have developed, followed by their application to predicting the aerodynamic loads of several industrial aircraft configurations.

2 Variable-Fidelity Modeling

Variable-fidelity modeling (VFM) offers an efficient way to generate aerodynamic data based on a set of CFD methods with varying degree of fidelity and computational expense (potential theory, Euler, and Reynolds-Averaged Navier-Stokes (RANS) equations) or based on a single physical model evaluated on meshes of varying refinement. The unknown aerodynamic data is approximated as a function of the input parameters such as Mach number, angle of attack, etc. The goal is to reduce the number of expensive high-fidelity computations used to build the model. A low-fidelity CFD method is used to automatically compute hundreds or thousands of solutions at points in the parameter space selected with a Design of Experiments (DoE) tool. The remainder of the parameter space is "filled in" using interpolation procedures. A few points in the parameter space are computed using a high-fidelity CFD method. These points are also selected using a DoE tool. Then, data fusion is used to combine the data stemming from the different methods, with low-fidelity data indicating trends and a small number of high-fidelity data correcting the absolute values. Obviously, one of the key challenges is how to manage the different models of varying fidelity or how to correct the low-fidelity model to approximate the high-fidelity data.

The most popular method currently used for VFM is a correction-based method. The correction is called "bridge function", "scaling function" or "calibration". The correction can be multiplicative, additive or hybrid multiplicative/additive. Multiplicative bridge functions are used to "locally" scale the low-fidelity function to approximate the high-fidelity function [5]. The additive bridge function was then developed as a "global" correction and has become the most popular method for variable-fidelity optimization or for data fusion [6]. It

was shown to be more accurate and robust than a multiplicative bridge function. To represent more complicated correlation between low- and high-fidelity functions, a hybrid multiplicative/additive bridge function method was proposed in [7] for variable-fidelity optimization.

Our VFM approach [8,9] is based on two novel ingredients, the generalized hybrid bridge function and hierarchical Kriging. In the generalized hybrid bridge function a low-order polynomial is used to describe the appropriate multiplicative scaling between the high- and low-fidelity models. In a second step, a Kriging model is used to describe the difference between the high-fidelity model and the scaled low-fidelity model. This new bridge function was developed to fully exploit the trend underlying the low-fidelity data, which is used to assist the prediction of the high-fidelity function. Details of our implementation can be found in [10]. Hierarchical Kriging refers to a surrogate model of a high-fidelity function that uses a Kriging model of a sampled lower-fidelity function as a model trend. As a consequence, the variation in the lower-fidelity data is mapped to the high-fidelity data implicitly, and a more accurate surrogate model for the high-fidelity function is obtained. As such, this method does not use an explicit bridge function to correct the low-fidelity data. We have implemented this approach for two and three levels of fidelity [11]. Gradient information, e.g., from an adjoint solver, can also be included, if available [10]. The unique feature of our hierarchical Kriging modeling method compared to other approaches such as Cokriging [12] is that it does not require modeling the "cross correlation" between the low- and high-fidelity functions, resulting in reduced model complexity and smaller size of the correlation matrix. In addition, the root-mean squared error (RMSE) predicted by our method is more consistent with the trend of the corresponding true error. This can be explained by the fact that the behavior of the low-fidelity function is explicitly included in the RMSE formulation. If the low-fidelity function captures the trend of the high-fidelity function, a better prediction of the RMSE can be obtained. This improved RSME estimation can be used, e.g., to implement an effective adaptive sampling strategy [11].

3 Reduced-Order Modeling

A good review of popular reduced-order modeling (ROM) methods used in the fields of unsteady aerodynamics and aeroelasticity can be found in [13]. A ROM approach is essentially any method that produces a model that is of a much smaller size than the full-order CFD model and thus much cheaper to evaluate. This is typically achieved by choosing a reduced basis, which usually requires several up-front solutions of the high-dimensional, full-order model. In order to be accurate and useful, however, it must still contain the important linear and non-linear characteristics of the full size system. An approximate solution is then computed as a linear combination of the reduced basis functions. The coefficients in the linear combination can be determined in different ways, e.g., by projecting the governing equations onto the reduced basis (Galerkin projection) [14], or simply by interpolating between the coefficients of the reduced basis [15]. The

cost of such a computation would be very small if the dimension of the reduced basis is very small and if one ignores the cost of the off-line determination of the reduced basis. Thus, as a consequence, ROMs are mainly of interest when the cost of generating the ROM is outweighed by the saving achieved when using it for predictions. This is the case when a large number of predictions will be made with the ROM as, for example, in multi-disciplinary optimization and analysis.

Here, we focus on proper orthogonal decomposition (POD) [16] to determine the reduced basis, which requires the generation of a snapshot set and removing "redundant" information from the snapshot set. The generation of the snapshot set requires several solutions of the high-dimensional full-order CFD model. POD is then used to compute an orthonormal basis ordered by information content spanning the same space as the snapshot set. This requires performing an eigenvalue or singular value decomposition of the snapshot matrix. Subsequently, we restrict all computations of interest to the POD subspace. In particular, we determine the coefficients of the linear combination of POD modes such that the defect of the approximate solution to the governing equations is minimized. In other words, the solution closest to fulfilling the governing equations of fluid dynamics in the low-dimensional POD subspace is sought after. In this way, the full-order CFD flow problem is replaced by a low-order non-linear least-squares optimization problem. The least squares method solves the optimization problem by minimizing the sum of the squared residuals over the computational domain with a gradient-based Levenberg-Marquardt algorithm with Broyden rank-one updates [17]. The approach is referred to as least-squares ROM or LSQ-ROM in the following and has originally been introduced in [18] in the context of inviscid airfoil analysis and (inverse) design optimization. Details of our implementation for steady, turbulent flow problems can be found in [19,20,21,22].

4 Results

In this section we present results for both the variable-fidelity and the reduced-order modeling approach. The focus is on demonstrating the application of both approaches to industrial aircraft configurations. As this article is concerned with meta-modeling and not with CFD code development, we consider the full-order CFD model validated and do not present comparisons with wind-tunnel or flight-test data, but compare the ROM and VFM results with reference CFD solutions instead.

4.1 Variable-Fidelity Modeling of DLR F12 Full Aircraft Configuration

Here, we apply the hierarchical Kriging method sketched above to the efficient prediction of the aerodynamic raw data of an industrial configuration, the DLR F12, with three independent variables: angle of attack, $\alpha \in [-4°,+6°]$, Mach number, $Ma \in [0.6,0.9]$, and the deflection angle of the horizontal tail, $\delta \in [-6°,+6°]$. Note that only the integrated aerodynamic data such as C_L, C_D, and C_M are considered here.

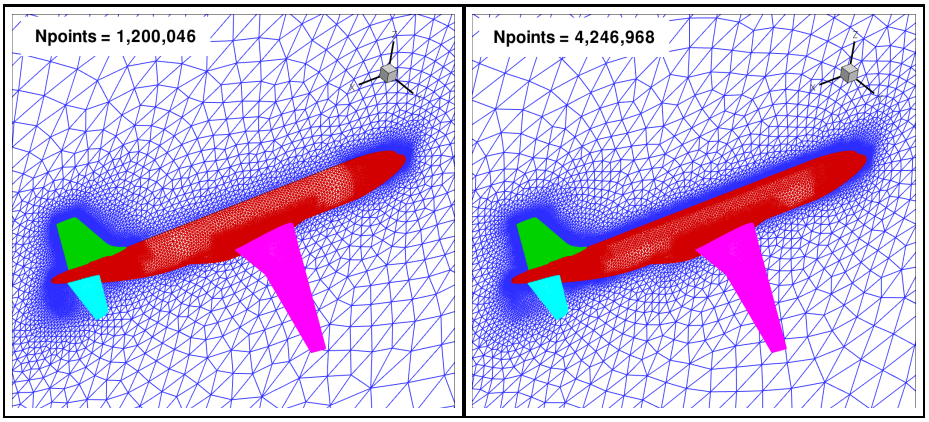

Fig. 1. Computational grids for DLR F12 (left: Euler grid; right: RANS grid) [11]

The computational grids used for the Euler and RANS computations with the DLR TAU code [23] are shown in Fig. 1. The Euler and RANS computations were performed in parallel on 24 domains.

An Euler computation on the grid with 1 million grid points roughly takes about 24 min, whereas a single RANS computation takes about 3 h of wall clock time on the grid with around 4 million grid points. Note that the Spalart–Allmaras one-equation turbulence model is used for the RANS computations, assuming a Reynolds number of 6.5 million.

The sample points were chosen using the quasi-Monte Carlo method. A total of 94 sample points were chosen for the Euler calculations, and 80 sample points are chosen for the RANS computations. All the Euler samples were used to build the low-fidelity Kriging that aims to capture the global trend of the unknown high-fidelity aerodynamic function, while 25 RANS samples were used to build a hierarchical Kriging model for the high-fidelity aerodynamic function with the Kriging model of the low-fidelity function being used as the trend model. The remaining RANS samples were used for validation purposes.

The results are examined by making cuts through the resulting response surfaces for the drag and pitching moment coefficients. The results for a fixed Mach number of Ma=0.8 and a deflection angle of the horizontal tail of δ=0° are shown as a function of the angle of attack in 4.2. The Cokriging predictions are denoted by the blue dashed lines, while the hierarchical Kriging results are represented by a black solid line. The high-fidelity RANS validation data is given by black square symbols. The red dash-dotted line, which shows the Kriging model of the Euler samples, is included to show the good correlation between the low- and high-fidelity data. The blue dash-dot-dotted line denotes the predictions of a Kriging model through the 25 high-fidelity samples.

It is observed that both the Cokriging and hierarchical Kriging models give accurate results if compared with the RANS validation data. Hierarchical Kriging is also observed to make better predictions than Cokriging, although both models

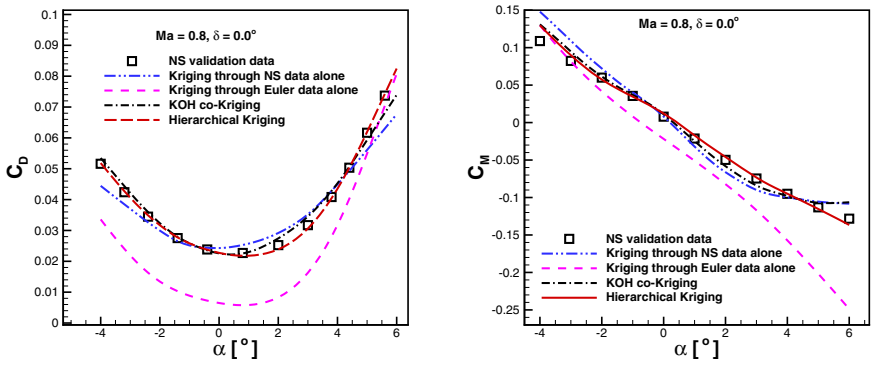

Fig. 2. Drag (left) and pitching moment (right) coefficients for DLR F12. Comparison of Cokriging, Hierarchical Kriging and Kriging results with TAU validation data [11]. KOH Cokrig. refers to Cokriging according to Kennedy and O'Hagen [24]

rely on the same set of samples. On the other hand, it is not surprising that hierarchical Kriging outperforms Kriging of the RANS samples alone as the low-fidelity samples are considered in addition to the high-fidelity samples. More detailed results and an estimate of the efficiency of the chosen VFM approach can be found in [11].

4.2 Reduced-Order Models for Flow around Industrial Aircraft Configurations

ROM for Subsonic Flow around Full Aircraft Configuration. The basic POD-ROM method has previously been used by the authors to model the inviscid transonic flow around the NACA 0012 airfoil [19] and the turbulent transonic flow around the NACA 64A010 airfoil [20,21] and the NASA Common Research Model [20]. Here, we present a reduced-order model of the compressible RANS equations for the turbulent flow around a complete passenger aircraft in clean configuration. The model relies on four snapshots computed with the DLR flow solver TAU [23] at a constant subsonic Mach number and at angles of attack of $\alpha = [1.0°, 0.0°, 1.0°, 2.0°]$ using the Menter-SST turbulence model. The hybrid unstructured computational grid used to compute the snapshots consists of 8,898,749 points.

A POD of the snapshot data gives an orthogonal representation, where the first three modes capture the complete information contained in the snapshots. All three modes and the mode of the mean flow were kept for POD-based flow approximation. Note that the mean flow vector is a constant vector that is not subject to any estimation/optimization procedure. Here, the initial flow problem of order 71,189,992 (number of gird points times number of unknowns per grid point) is reduced to an order-3 optimization problem.

An approximate flow solution was computed for α=7°, i.e., for a flow condition beyond the range of angles of attack of the snapshots, keeping the remaining

free-stream conditions constant. Note that a prediction at an angle of attack that is five degrees beyond the range covered by the snapshots is comparably large and is chosen here only to show the capability of the method. In general one cannot expect the ROM to behave well at conditions for which no comparable snapshot information has been provided and therefore, in realistic applications, such extreme cases of extrapolation should be avoided.

The LSQ-ROM optimization procedure was initialized with an approximate solution derived from thin-plate spline (TPS) interpolation of the POD coefficients associated with the POD of the given snapshots. In order to obtain adequate problem scaling, that is, adequate orders of magnitude in the residual vectors and the Jacobian, the point-wise residuals as outputted by the CFD solver were weighted by one over the square-root of the grid-cell volumes. The weighting by the (reciprocal) cell volumes puts emphasis on residuals in small grid cells, while it dampens the influence of residuals in large ones.

Table 1 summarizes the number of full-order CFD residual evaluations that were due to solve the least-squares optimization problem, the computational effort in terms of CPU time, and the relative error in the integrated lift coefficient with respect to the CFD reference solution. The corresponding values for the TPS-interpolated solution used to initialize the LSQ-ROM, a Kriging-interpolated solution and a fully-converged reference TAU solution at 7° are shown as well.

With regard to the number of residual evaluations, acceleration by a factor of 470 is obtained with the LSQ-ROM when compared to a full-order CFD simulation. The absolute computational time is reduced by a factor of 335 by the LSQ-ROM approach and by a factor of 2,753 by the interpolation-based approaches. Note, however, that the error in C_L for the LSQ-ROM approximation is one order of magnitude lower than that for the interpolation-based approximations.

The corresponding surface pressure distributions are shown in 4.2. The reference TAU solution is shown in the upper left corner and the ROM solution in the upper right corner. The TPS and Kriging-interpolated solutions are shown in the lower left and right corner, respectively. The contour values in the legend have been omitted for proprietary reasons. The ROM solution is seen to be very similar to the reference TAU solution, while the interpolated solutions differ notably from the TAU reference solution over the front of the fuselage. The

Table 1. Results for industrial aircraft configuration at extrapolatory flow conditions; subsonic Mach number, typical wind-tunnel Reynolds number

Procedure	No. residual eval.	CPU time	Error C_L
Kriging	-	4.5 min	-1.99%
TPS	-	4.5 min	+1.33%
ROM (TPS initialization)	17	37.0 min	+0.28%
TAU CFD reference	8,000	12,390.0 min	0.00%

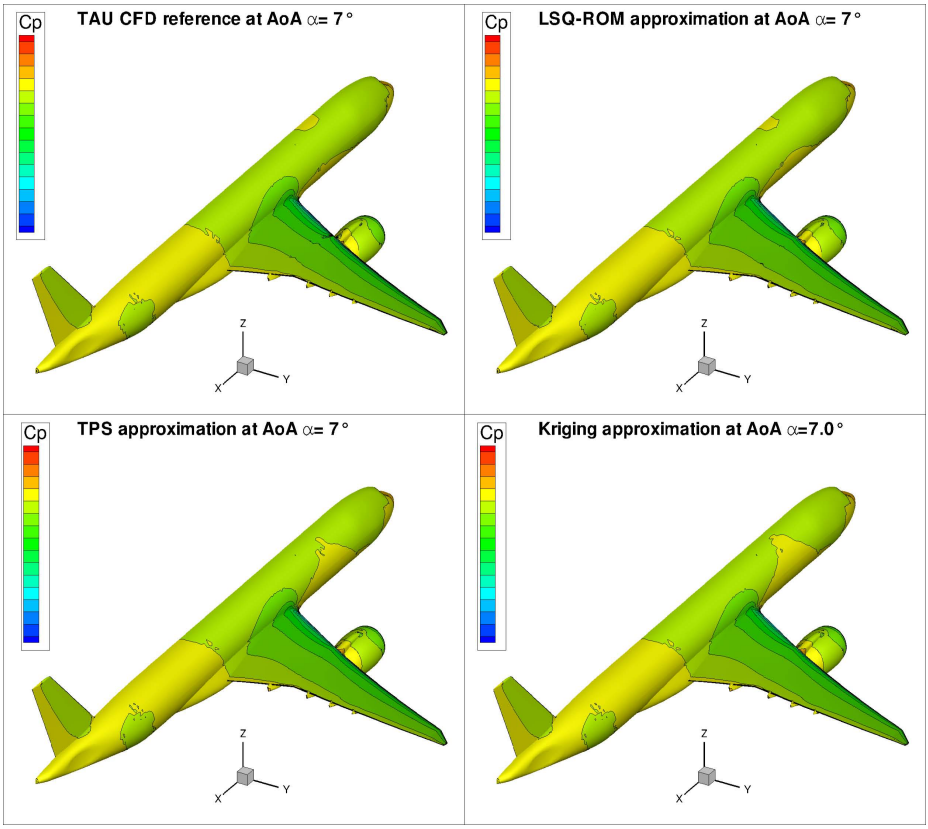

Fig. 3. Comparison of surface pressure contours; α=7°, subsonic Mach number, typical wind-tunnel Reynolds number

footprint of the nacelle vortex over the wing is still visible in the ROM solution, but is smeared out in the interpolated solutions.

ROM for Transonic Flow around Wing-Body Aircraft Configuration. Next, we apply the ROM approach to predict the fully turbulent flow around an industrial wing-body aircraft configuration at transonic flow conditions. Nine snapshots at angles of attack between 0.0° and 8.0° in 1° intervals were computed with TAU using the turbulence model due to Spalart and Allmaras for a Mach number of 0.85 and a Reynolds number of 3.3 million. The hybrid unstructured grid used for all computations featured 5.8 million grid points.

Here, we used a volume-based scalar product in the POD of the nine snapshots and the point-wise residuals as outputted by the CFD solver were weighted by one over the square-root of the grid-cell volumes to put emphasis on residuals associated with small grid cells. Our best practice for transonic flow predictions derived from a 2D test case [21] is to use the residuals associated with the 20%

smallest grid cells only for residual optimization, i.e., about 1.17 million grid points. This improves the accuracy of the predicted surface pressure solution, as the smallest grid cells are typically located close to the surface. Also note that only the primitive mean-flow variables were considered for residual optimization.

First, we made a prediction for an angle of attack of α=3.5°, that is, for an angle of attack within the interval of snapshots. A total of 32 flow solver residual evaluations were required to find a set of POD coefficients corresponding to a ROM solution with the minimum flow solver residual. The total wall-clock time of the ROM computation, including reading and processing of the snapshots, was less than 5 min on 12 cores, while the computation of the converged reference CFD solution took 4,000 iterations or 5.5 h on 12 cores. Note that the convergence of the reference solution was accelerated by initializing the TAU flow solver with the converged solution at 3.0°. In summary, a speed-up by a factor of 125 was achieved with respect to the number of flow solver iterations and by a factor of more than 66 in terms of wall-clock time. Note, however, that one flow-solver residual evaluation is less expensive than a complete iteration, which also included updating the boundary conditions, etc., so the number of iterations required to compute the reference CFD solutions cannot directly be compared to the number of residual evaluations due to evaluate the ROM. Also, if more than one prediction were to be performed with the ROM, the snapshots need to be read and processed only once, making a subsequent prediction cheaper, which translated into a higher speed-up in terms of wall-clock time.

The surface pressure distribution predicted by the ROM is compared to the reference TAU solution in Fig. 4. Please note that the contour legend has been omitted for proprietary reasons. Excellent agreement between the ROM solution (left) and the TAU reference solution (right) is observed over the fuselage. Both solutions feature a lambda-shock over the wing. Good agreement is seen over most of the wing, apart from the region between the wing tip and the middle of the wing, which is dominated by a strong shock.

A more quantitative comparison is presented in Fig. 5. The figure shows the surface pressure coefficient as a function of the dimensionless chord at four different span-wise section cuts. Please note that in this and all subsequent plots, the y-axis values have been omitted for proprietary reasons. The filled circle symbols represent the TAU reference solution while the solid blue line corresponds to the predicted ROM solution. Note that the agreement between the two solutions is very good at all sections but the section corresponding to η=0.84, which lies in the outboard wing region that is dominated by a strong shock. At this section, the ROM solution is observed to feature a somewhat upstream shock location and an unphysical pressure level on the upper surface just upstream of the shock. This may be due to the sought-after solution not being contained in the POD subspace, or due to a sub-optimal set of POD coefficients from the residual optimization procedure. It is also possible that there is a negative influence of snapshots at significantly lower or higher angles of attack with different flow physics on the ROM solution, as the POD couples all snapshots through the

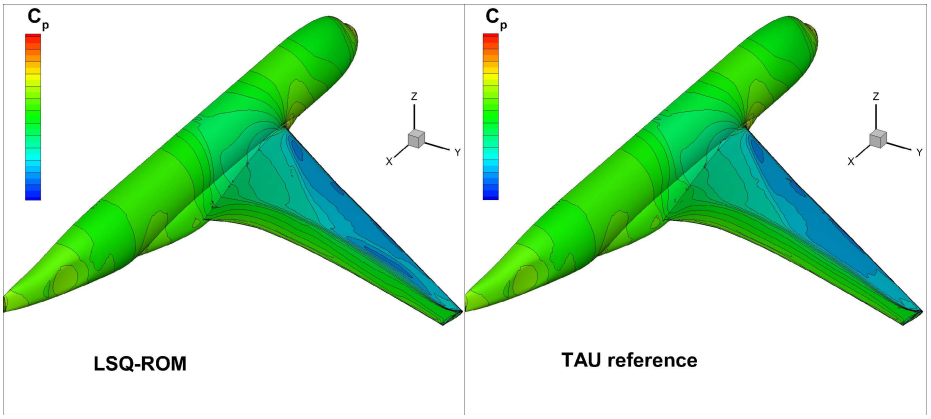

Fig. 4. Comparison of surface pressure distribution, M=0.85, α=3.5°, Re=3.3×10^6

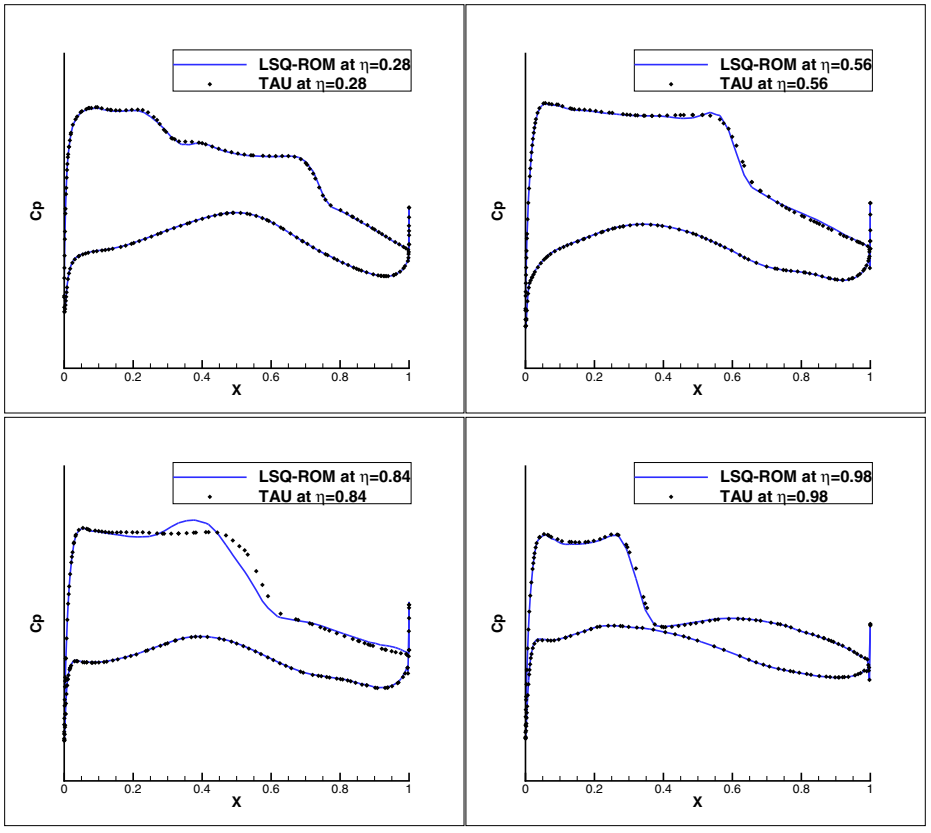

Fig. 5. Comparison of pressure coefficient at different wing section cuts, M=0.85, α=3.5°, Re=3.3×10^6

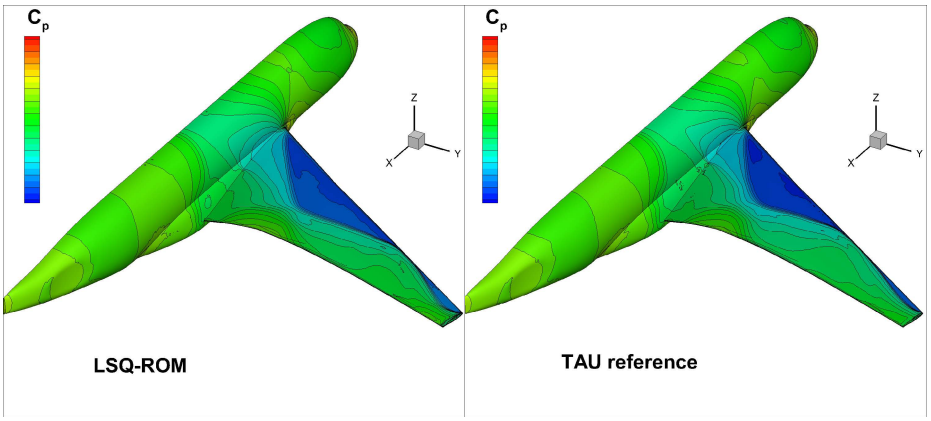

Fig. 6. Comparison of surface pressure distribution, M=0.85, α=8.0°, Re=3.3×10^6

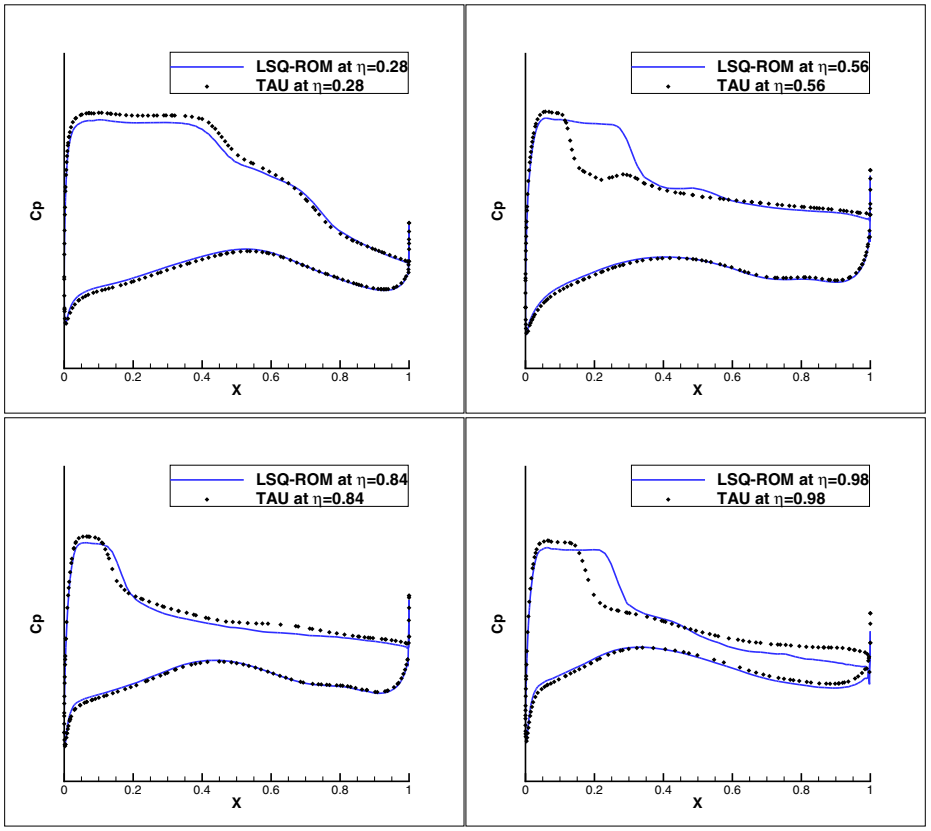

Fig. 7. Comparison of pressure coefficient at different wing section cuts, M=0.85, α=8.0°, Re=3.3×10^6

modes. The situation may be improved, for example, by reducing or properly selecting the number of modes, or by weighting or clustering the snapshots.

Next, we considered eight snapshots at angles of attack between 0.0° and 7.0° in 1° intervals to make a prediction for an angle of attack of α=8.0°, that is, for an angle of attack outside the interval of snapshots. This is considered a more challenging scenario for any ROM approach, because we are now dealing with extrapolatory flow condition and thus, potentially, with flow physics that are not contained in the snapshots used to construct the POD subspace. This time, a total of only 26 flow solver residual evaluations were required to find a set of POD coefficients corresponding to a ROM solution with the minimum flow solver residual. Note that the number of residual evaluations is linked to the number of POD coefficients that need to be tuned, which in turn depends on the number of retained modes and thus on the number of snapshots. The total wall-clock time of the ROM computation, including reading and processing of the snapshots, was less than 4 min on 12 cores, while the computation of the converged reference CFD solution took again about 4,000 iterations or 5.5 h on 12 cores.

The surface pressure distribution predicted by the ROM for 8° is compared to the reference TAU solution in Fig. 6. Small differences between the ROM solution (left) and the TAU reference solution (right) are observed over the forward fuselage this time. Larger discrepancies are seen over the wing, although the overall agreement in terms of the rather complex flow topology is very good.

The more quantitative comparison in Fig. 7 reveals that the agreement between the two solutions is good for the inboard section at η=0.28 and the more outboard section at η=0.84, both in terms of the pressure levels as well as the shock locations. Fair agreement is observed for the mid-wing section at η=0.56 and the outboard section at η=0.98. For both sections, the ROM predicted a more downstream shock location. As mentioned above, the disagreement is due to the fact that we extrapolate beyond the range of the snapshots and that the sought-after solution is not contained in the POD subspace.

5 Conclusions

We have developed Variable-Fidelity Modeling (VFM) and Reduced-Order modeling (ROM) methods and applied them to predict the aerodynamic loads of three different aircraft configurations at subsonic and transonic Mach numbers. Both methods require a set of samples or snapshots to be computed offline with the expensive full-order model to build the surrogate or reduced-order models before a prediction can be made online. Once a model has been built, it can be evaluated repeatedly at low computational cost, making their use in multi-disciplinary analysis, design and optimization attractive.

References

1. Abbas-Bayoumi, A., Becker, K.: An industrial view on numerical simulation for aircraft aerodynamic design. Journal of Mathematics in Industry 1(10) (2011), doi:10.1186/2190-5983-1-10
2. Salas, M.D.: Digital Flight: The Last CFD Aeronautical Grand Challenge. Journal of Scientific Computing 28(2/3), 479–505 (2006), doi:10.1007/s10915-006-9087-7
3. Tinoco, E.N., Bogue, D.R., Kao, T.J., Yu, N.J., Li, P., Ball, D.N.: Progress toward CFD for full flight envelope. The Aeronautical Journal 109(1100), 451–460 (2005)
4. Klenner, J., Becker, K., Cross, M., Kroll, N.: Future Simulation Concept. In: Proceedings 1st CEAS Conference, Paper D07027256, Berlin, Germany (2007)
5. Chang, K.J., Haftka, R.T., Giles, G.L., Kao, P.-J.: Sensitivity-based scaling for approximating structural response. Journal of Aircraft 30(2), 283–288 (1993)
6. Choi, S., Alonso, J.J., Kim, S., Kroo, I., Wintzer, M.: Two-level multi-fidelity design optimization studies for supersonic jets. AIAA Paper 2005-0531 (2005)
7. Gano, S.E., Renaud, J.E., Martin, J.D., Simpson, T.W.: Update strategies for kriging models for using in variable fidelity optimization, AIAA Paper 2005–2057 (April 2005)
8. Han, Z.-H., Görtz, S., Hain, R.: A Variable-Fidelity Modeling Method for Aero-Loads Prediction. In: Dillmann, A., Heller, G., Klaas, M., Kreplin, H.-P., Nitsche, W., Schröder, W. (eds.) New Results in Numerical and Experimental Fluid Mechanics VII. NNFM, vol. 112, pp. 17–25. Springer, Heidelberg (2008)
9. Han, Z.-H., Görtz, S., Zimmermann, R.: On Improving Efficiency and Accuracy of Variable-Fidelity Modeling in Aero-data for Loads Context. In: Proceedings of CEAS 2009 European Air and Space Conference, Manchester, UK, October 26-29 (2009)
10. Han, Z.H., Görtz, S., Zimmermann, R.: Improving variable-fidelity surrogate modeling via gradient-enhanced kriging and a generalized hybrid bridge function. Journal of Aerospace Science and Technology (2012) (in press), doi:10.1016/j.ast.2012.01.006
11. Han, Z.H., Görtz, S.: Hierarchical Kriging Model for Variable-Fidelity Surrogate Modeling. AIAA Journal 50(9), 1885–1896 (2012)
12. Han, Z.H., Zimmermann, R., Görtz, S.: Alternative Co-Kriging Model for Variable-Fidelity Surrogate Modeling. AIAA Journal 50(5) (May 2012)
13. Lucia, D., Beran, P., Silva, W.: Reduced-order modeling: new approaches for computational physics. Progress in Aerospace Sciences 40(1-2), 51–117 (2004)
14. Rowley, W., Colonius, T., Murray, R.M.: Model reduction for compressible flows using POD and Galerkin projection. Physica D. Nonlinear Phenomena 189(1-2), 115–129 (2004)
15. Bui-Thanh, T., Damodaran, M., Willcox, K.: Proper orthogonal decomposition extensions for parametric applications in transonic aerodynamics. AIAA Journal 42(8), 1505–1516 (2004)
16. Holmes, P., Lumley, J., Berkooz, G.: Turbulence, Coherent Structures, Dynamical Systems and Symmetry. Cambridge University Press, Cambridge (1996)
17. Ruszczynski, A.: Nonlinear Optimization. Princeton University Press, Princeton (2006)
18. LeGresley, P.A., Alonson, J.J.: Investigation of Non-Linear Projection for POD Based Reduced Order Models for Aerodynamics. AIAA Paper 2001-0926 (2001)

19. Zimmermann, R., Görtz, S.: Non Linear Reduced Order Models For Steady Aerodynamics. In: Procedia Computer Science, Proceedings of the International Conference on Computational Science, ICCS 2010, vol. 1(1), pp. 165–174 (May 2010)
20. Zimmermann, R., Görtz, S.: Non-Linear Reduced Order Models for Steady Turbulent Aerodynamics. The Aeronautical Journal 116(1184), 1079–1100 (2012)
21. Zimmermann, R.: Towards Best-Practice Guidelines for POD-based Reduced Order Mod-eling of Transonic Flows. In: Proceedings International Conference on Evolutionary and Deterministic Methods for Design, Optimization and Control with Applications to Industrial and Societal Problems (Eurogen 2011), Capua, Italy, September 14-16 (2011)
22. Mifsud, M., Zimmermann, R., Görtz, S.: A POD-based reduced order modeling approach for the efficient computation of high-lift aerodynamics. In: Proceedings International Conference on Evolutionary and Deterministic Methods for Design, Optimization and Control with Applications to Industrial and Societal Problems (Eurogen 2011), Capua, Italy, September 14-16 (2011)
23. Gerhold, T., Friedrich, O., Evans, J., Galle, M.: Calculation of complex three-dimensional configurations employing the DLR-TAU-code. AIAA Paper 1997-0167 (1997)
24. Kennedy, M.C., O'Hagan, A.: Predicting the Output from a Complex Computer Code When Fast Approximations Are Available. Biometrika 87(1), 1–13 (2000)