

Cumulative Global Metamodels with Uncertainty: a Tool for Aerospace Integration

Patrick H. Reisenhel¹, John F. Love², Daniel J. Lesieutre³, and Robert E. Childs⁴
Nielsen Engineering & Research, Inc., Mountain View, CA 94566, USA

The integration of multidisciplinary data is key to supporting decisions during the development of aerospace products. Multidimensional metamodels can now be automatically constructed using limited experimental or numerical data, including data from heterogeneous sources. Recent progress in multidimensional response surface technology, for example, provides the ability to interpolate between sparse data points in a multidimensional parameter space. These analytical representations act as surrogates that are based on and complement higher fidelity models and/or experiments. These high-level representations or metamodels can include technical data from multiple fidelity levels and multiple disciplines, but also nontechnical data such as costs and schedules. Most importantly, these representations can be constructed on-the-fly and are cumulatively enriched as more data become available. The purpose of the present paper is to highlight applications of these Cumulative Global Metamodels (CGM), their ease of construction, and the role they can play in aerospace integration. This paper focuses on two types of applications, namely, design optimization and mutual data set enrichment via data fusion.

I. Nomenclature

C_D = drag force coefficient
 C_m = pitching moment about center of gravity
 C_R = radial force coefficient
 M = Mach number
 r_c = canopy radius
 ref = reference value
 α = angle of attack
 ϕ = roll angle
 θ_c = total canopy angle
 θ_r = canopy reference angle
 σ = standard deviation

II. Background

GLOBAL metamodels and response surface technology are increasing used in a variety of fields, including structural reliability, instrument calibration, and aerodynamic and trajectory optimization, to name a few.¹⁻¹¹ Because of their analytical nature, these models can be used for automated searches and are naturally well-suited to the acceleration of optimization tasks and rapid strategy evaluation. A central issue to constructing appropriate response surface models is the so-called curse of dimensionality, in which the number of data points required to characterize/support the surface increases exponentially with the number of independent variables. This difficulty is well-known and, in effect, precludes the use of conventional schemes, such as polynomial and piecewise-polynomial (finite-element) approximations. Neural networks, support vector machines, and multidimensional splines all have proven capable of multidimensional data generalization and (partially) circumventing this difficulty. The particular approach used in this paper is based on self-training radial basis function networks which form the basis of the NEAR-RS (response surface) technology.

¹ Chief Scientist, 605 Ellis St., Suite 200. Senior Member AIAA.

² Research Engineer, 605 Ellis St., Suite 200.

³ Senior Research Engineer, 605 Ellis St., Suite 200. Senior Member AIAA

⁴ Vice-President, 605 Ellis St., Suite 200. Senior Member AIAA

NEAR-RS is a software system consisting of two modules: a metamodel (response surface) identification module, and a metamodel evaluation/interrogation module. A graphical user interface included in this second module serves as a multidimensional viewer, which facilitates the visualization of trends in high-dimensional data.¹² A key aspect of the technology is the ability to estimate further sampling needs and model quality, based on uncertainty estimation.

Four distinct applications illustrating the use of these methods are discussed in this paper. The first three are based on ongoing projects and illustrate the data adaptivity, data fusion, and feature extraction/data mining capabilities of the method. The last application presents an example of the uncertainty modeling capabilities of CGMs.

III. Refueling Drogue Canopy Design

Rapid optimization through multidimensional surrogate modeling was used in a recent effort at Nielsen Engineering & Research involving the redesign of a refueling drogue canopy. Figure 1 depicts a standard C-130 high-speed drogue during a wind-tunnel test (flow is from left to right). The picture shows the end of the refueling hose, connected through a ball joint to the MA-3 coupling, to which the strut arms of the canopy are attached. The inflated canopy provides the appropriate drag force to both stabilize the hose-and-drogue assembly and maintain the hose at the proper angle for refueling operations.

A. Objective and Methods

For a given air speed, the design objective is to maximize the radial force on the canopy, while maintaining the drag within a narrow prescribed range. The canopy shape is modeled as a first approximation as being axisymmetric. The four geometric design parameters (see Fig. 2) are: the canopy radius, r_c , the reference angle, θ_r , the vent angle, θ_v , and the total arc, θ_c . These four parameters define a family of canopy shapes with a common downstream strut arm attachment point.

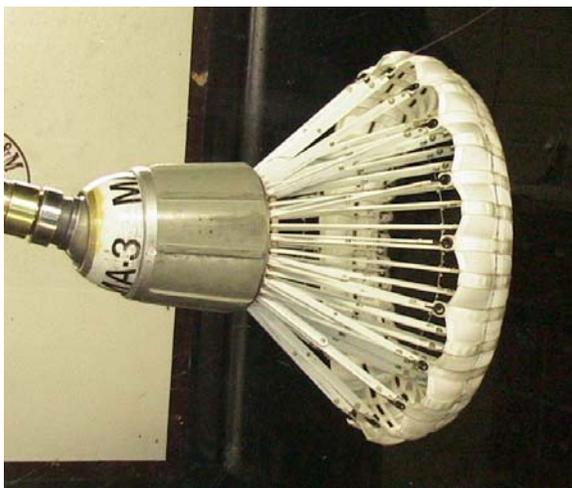


Figure 1. Standard C-130 refueling drogue.

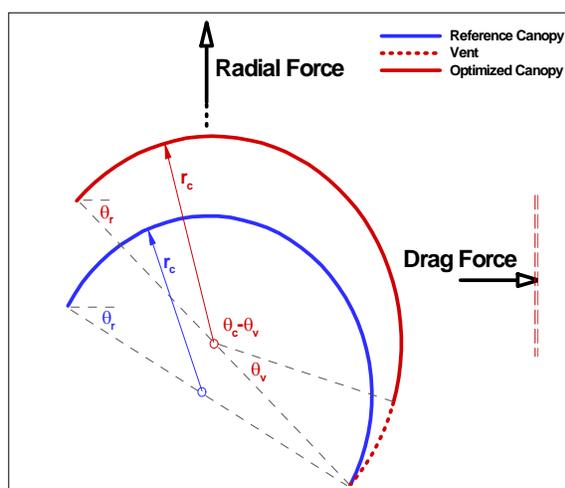


Figure 2. Canopy geometric parameter definition.

The aerodynamic performance of the inflated canopies was estimated using Sandia National Laboratory's VPARA code.¹³ VPARA predicts the axisymmetric flow field over rigid bluff bodies such as thin shells, and has been used to estimate the drag and radial force produced by canopies of a fixed shape. The solution method is based on the axisymmetric vorticity transport equation, where discrete ring vortices are generated at the body, satisfying the appropriate surface boundary conditions. These vortices are subsequently diffused and convected into the boundary layers and the wake of the body, while new vortices are generated at the body surface. VPARA was used to compute time histories of the pressure distributions and the forces on the modeled surfaces. Estimates of the mean drag and radial force coefficients were obtained by averaging the instantaneous loads over a portion of the total time history. As an example, Fig. 3 depicts the location of the shed vortices after a nondimensional convective time of 240.

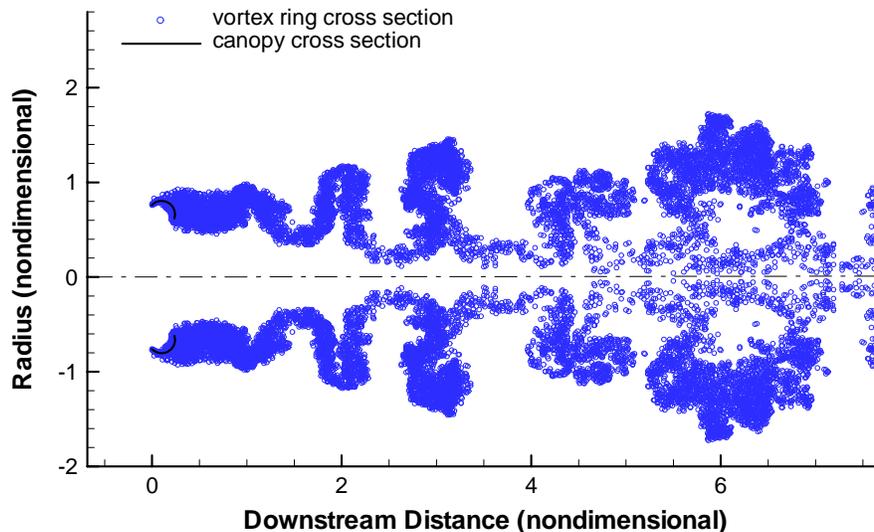


Figure 3. Canopy wake cross-section calculated using VPARA.

B. Optimization

For this design problem, there are four independent variables (the four geometric parameters defined in Fig. 2), and two dependent variables, C_D and C_R . From a known reference condition (corresponding to an existing drogue), the design space was seeded with the results of 12 additional VPARA calculations. Using these 13 initial configurations, four-dimensional response surfaces were constructed for C_D and C_R .

As previously mentioned, the objective was to maximize C_R , subject to C_D being confined to a narrow range of allowable values. Thus, the C_D and C_R response surfaces were combined into a single objective function, defining a new, highly nonlinear hypersurface. A thorough, global search of this hypersurface was then conducted through repeated interrogations of this surrogate model. The result of the search is a hierarchy of maxima which can then be investigated through further VPARA calculations.

After collecting a suitable number of additional calculations, new response surfaces are formed for C_D and C_R . In this manner, the resulting, updated hypersurface is cumulatively enriched with the addition of each new data point. An illustration of this process is shown in Fig. 4. Figure 4 shows the response surface of the radial force coefficient at various stages of the design iteration process. The surfaces are two-parameter projections of the hypersurface, i.e., r_c and θ_c are kept at fixed values (not necessarily corresponding to an optimum), for illustration purposes. The results of actual VPARA calculations for this same projection are indicated using triangular symbols. An indication of the corresponding drag coefficient response surface is also given through coloring. Figures 4(a), (b), (c) and (d) correspond to response surfaces formed after 29, 44, 95, and 136 calculations, respectively. Each figure is an instantaneous snapshot of a continuous parameter space representation, based on the available data.

An assessment of the performance of the optimization method is given in Fig. 5. Two curves are shown. Both represent the maximum radial force coefficient versus iteration number. At a given iteration, the bottom (red) curve represents the current best design based on the VPARA computations carried out to that point. By contrast, the top (green) curve at the *same* iteration number represents the new suggested optimum, based on a response surface computed using all the available data. Asymptotically, these two curves should converge. The purpose of the figure is not to show convergence, however, but to document the process actually used during the optimization, as well as to furnish insight into how the method works. For this purpose, next to these performance metric curves (maximum C_R , conditional to $C_D = 1.0$) is an indication of the independent variables for which this performance is achieved. The magnitudes of the independent (geometry definition) variables are shown in the form of radial histograms or “radiograms,” for short. Each radiogram shows four vectors at 90-degree angles, corresponding (clockwise from top) to: r_c , θ_r , θ_v , and θ_c . In these graphs, the minimum and maximum length of each vector has been rescaled from 0 to 1 for convenience. Also, a blue line connecting the ends of the vectors is added to enhance the visual perception. Using these radiograms, one can visualize on the same plot as the performance curves (a) the evolution of the actual design (red vectors), in terms of the geometric parameters, and (b) the parameters corresponding to the current suggested maximum (green vectors), based on the response surface.

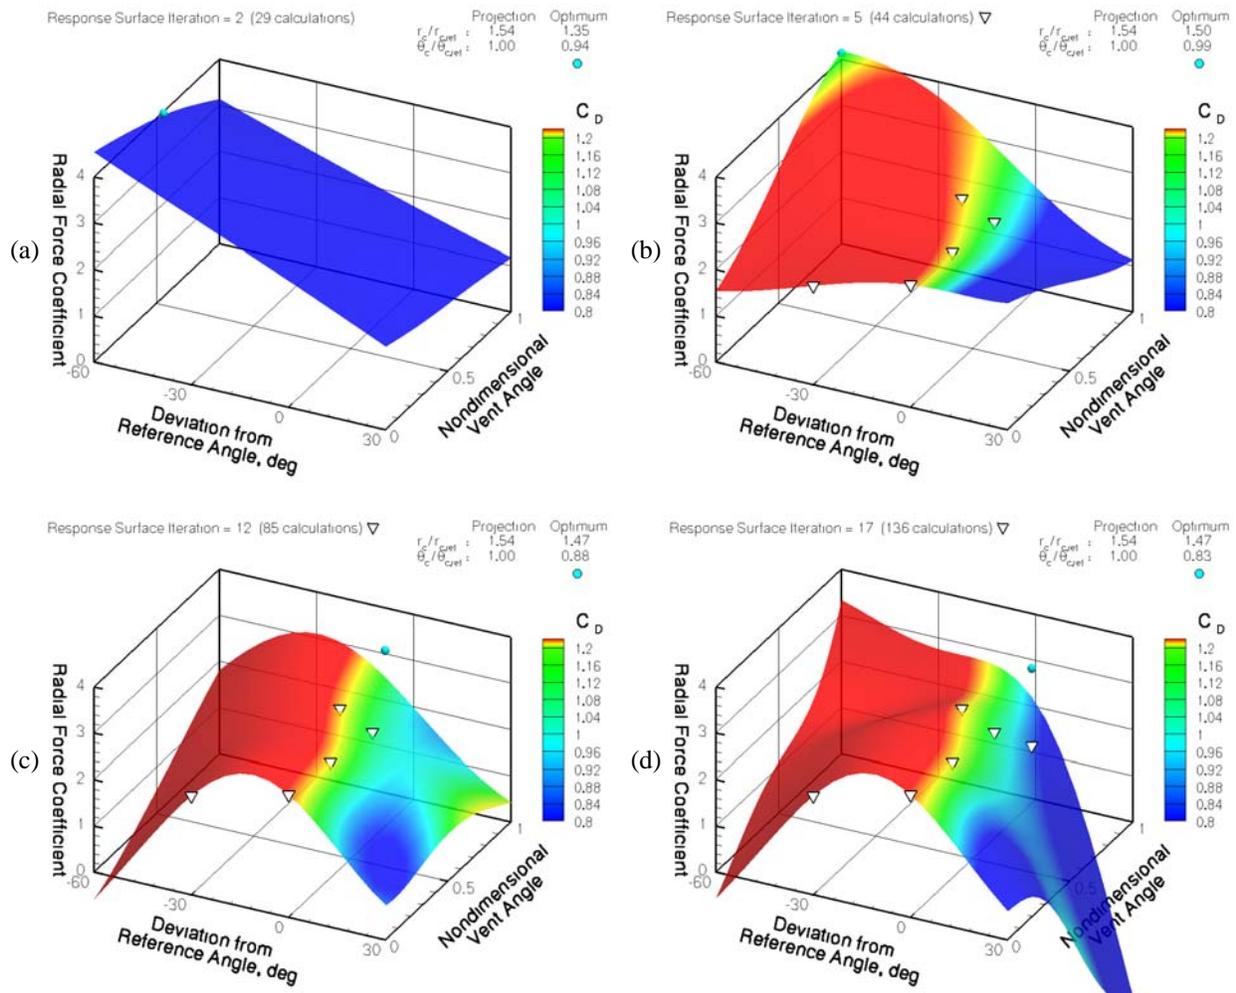


Figure 4. Radial force response surface evolution as a function of cumulative data set enrichment.

In the early stages of design, the response surface is more erratic and tends to point the design in new directions which may or may not be fruitful ultimately, as more real data are collected. The addition of these new data points enriches the response surface, which becomes more realistic and more accurate in the process. Also, rather than a systematic convergence study, the graph of Fig. 5 is a faithful documentation of the design process used. During the course of this design, the parameter space underwent several expansions, i.e., redefinitions of the parameter boundaries by the designer. Each time this happened, the current response surface pointed to new, extrapolative extrema (as high as 4.1). But, as the data set was subsequently enriched, this would be followed by a period of relaxation back to a more realistic maximum. Aside from parameter expansions, there were also changes in strategy throughout the process. For example, in the early phases, the designer chose to follow the suggestions made by the response surface optimizer (i.e., to investigate the response surface maximum). Later, the designer chose to perform calculations at some of the “corners” of the parameter space, to get a better overall representation the response surface. Later still, the designer was satisfied that the calculated optimum met the desired specifications, and was more interested in populating the design space near this optimum, in order to better understand the local shape of the response, such as the width of the allowable region, and whether there were any “dangerous” sensitivities to stay away from. In the final stages, the designer chose to investigate suboptimal maxima, essentially optimizing in subranges of the canopy radius. This was done with the goal of fabricating and wind-tunnel testing several different canopies, providing a smooth path from the existing, reference canopy to the (very different) optimized canopy. Regardless of the strategy used, the response surface's accuracy continues to be cumulatively enhanced as more data become available.

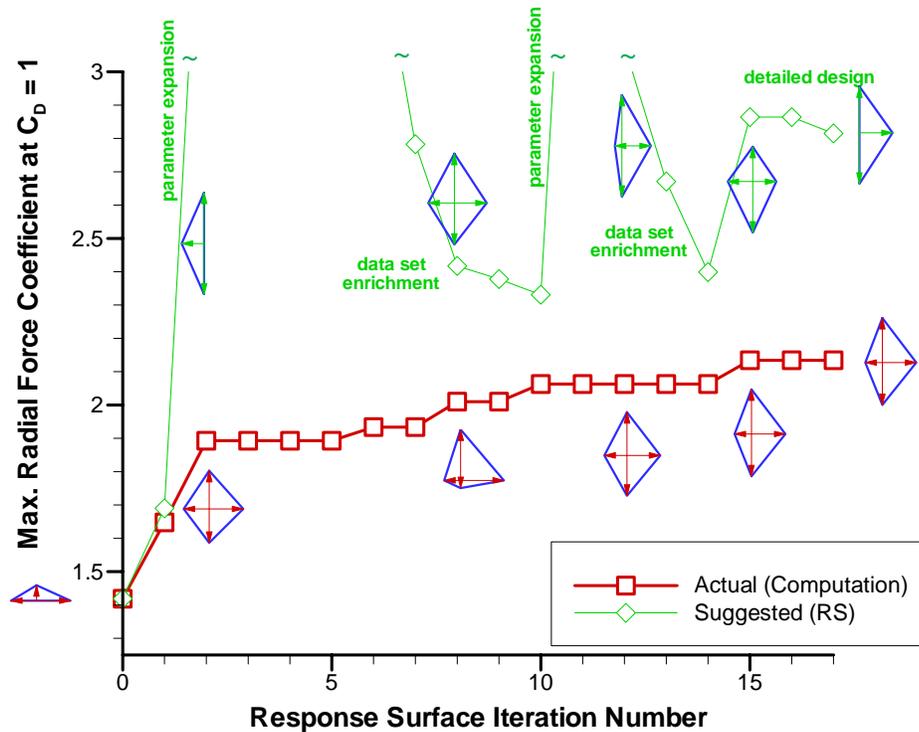


Figure 5. Design evolution as a function of response surface iteration.

Careful scrutiny of Fig. 5 also shows that this was not a mere gradient search, and that, indeed, significantly different directions and regions of the parameter space were investigated as a result of the surrogate model based optimization. Most importantly, the continuous and analytical description provided by the response surface is well-suited to rapid global search strategies, since interrogations of the surrogate model are *many* orders of magnitude faster than the individual data calculations, a fact that the authors took advantage of, giving a high degree of confidence that the parameter space was searched thoroughly.

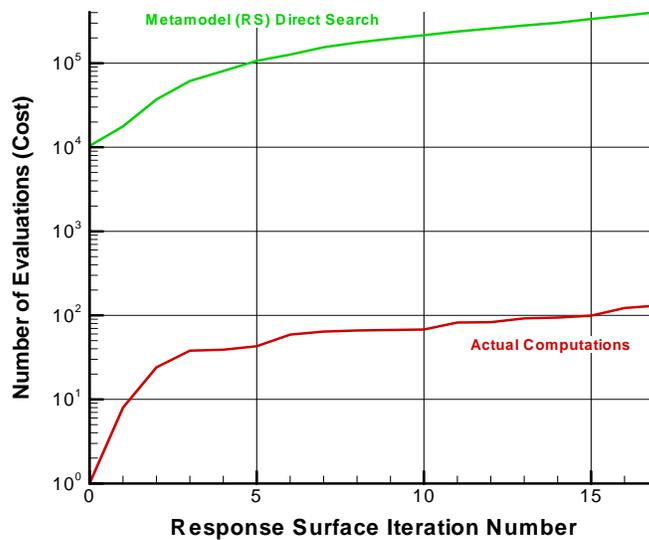


Figure 6. Efficiency of cumulative metamodel approach shown as a comparison of direct search count versus number of actual calculations.

This gain in performance is illustrated in Fig. 6, which shows the cumulative number of surrogate (metamodel) evaluations, compared to the cumulative number of VPARA calculations needed to achieve the final design, a net three-and-a-half order of magnitude gain.

IV. Large Asymmetric Launch Vehicle Payload Fairing Design

A second application involving sequential optimization concerns the aerodynamic and structural design of a payload fairing large enough to encompass a spacecraft with an optical mirror up to twice the diameter of an Evolved Expendable Launch Vehicle (EELV).¹⁴

A. Objective and Methods

The principle design objectives at the early stages of this work were related to stability and control, as well as the mass of the fairing. Optimization and computational fluid dynamics (CFD) methods were used to design an asymmetric exterior shape for the fairing that meets specific aerodynamic design goals. These aerodynamic goals include a balance of a low lateral force on the fairing and smooth variations in that force with respect to angle of attack, across a range of Mach numbers near $M = 1$.

The fairing surface was defined principally in terms of analytical functions. The cost to perform design optimization increases significantly with the number of independent design variables; therefore, it is important to minimize the number of variables needed to parameterize the surface shape. This can be accomplished with the analytical surface shape definition, rather than a surface defined by discrete points. Also, the approach sought to impose the designer's understanding of aerodynamics on the optimization process, and to focus the design on shapes that would trigger flow separation gradually to alleviate excessive aerodynamic forces on the fairing.

Because the flow around the payload fairing involves separation, a Reynolds-averaged Navier-Stokes CFD methodology (Overflow¹⁵) was used to predict the steady-state forces and moments on the launch vehicle. The payload fairing geometry was initially parameterized using nine free design variables, but three of those were found to have minimal effect on performance and were, consequently, set to fixed values. The remaining six active independent variables were used to automatically remesh the surface and recompute a new volume grid using Gridgen. Grid convergence studies were conducted to determine the optimal grid resolution needed. Using this grid resolution, the CFD calculation took between two and six hours, depending on the requested convergence level.

B. Optimization

The approach used was similar to the one described in Section III. A global optimization scheme is applied to a continuous, analytical surrogate, rather than the CFD calculations themselves. This allows a thorough search of the solution space without the cost of performing actual CFD solutions. Naturally, the approach relies on the response surface being representative of the design space, which may not be the case, particularly at the early stages of the design. Nevertheless, and as demonstrated through the previous example and this one, the procedure appears to yield solutions satisfying the problem constraints in a relatively small number of CFD calculations.

Similar to the example of Section III, multiple dependent variables were combined into a single objective function. The dependent variables were the overall pitching moment, C_m , at multiple aerodynamic conditions (M, α, ϕ). The six-dimensional independent variables parameter space was initially seeded using Latin hypercube sampling¹⁶ (LHS) design-of-experiments methods. The results of these initial calculations were used to form the first generation of response surfaces (one for each aerodynamic condition). A global search of the objective function over the response surfaces would then "suggest" the next set of CFD runs. A mixed strategy was used to enrich the data set (i.e., compute further CFD solutions). As in Section III, in the early stages of design, the strategy involved successive waves of parameter expansion, followed by periods of populating the data set. Each generation of new CFD calculations used LHS randomization for the independent variables, running 10 to 16 solutions at a time on a cluster of Linux workstations. By contrast, the intermediate and later stages of the design were characterized by a *focusing* of the parameter space on the most promising regions, at a given iteration of the design. Note that these regions of interest did not always remain the most promising one, as further CFD calculations would either confirm or invalidate the trend predicted by the response surface at the previous iteration.

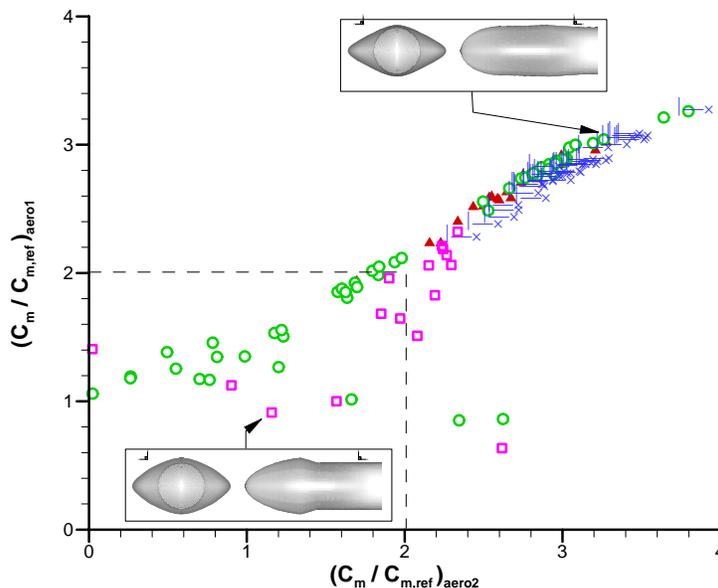


Figure 7. Rapid multiobjective optimization using cumulatively enhanced metamodel.

Figure 7 shows the result of the optimization procedure. In this example, the various symbols are associated with successive generations of CFD data points harvested in the process. Each data point generation is the result of an exhaustive “scatter and poll” search strategy on the global metamodel constructed as a result of the design’s previous iteration. Therefore, the metamodel is enhanced at each generation by the new knowledge provided by the results of the latest CFD runs. In this case, the objective was to reduce the coefficient C_m to less than twice some reference coefficient $C_{m,ref}$ for two distinct aerodynamic conditions, *aero1* and *aero2*, which are combinations of different (M, α) pairs near $M = 1$. Thus, any point such that $(C_m / C_{m,ref})_{aero1} < 2.0$ and $(C_m / C_{m,ref})_{aero2} < 2.0$ (i.e., inside the dashed line in Fig. 7) is an acceptable aerodynamic design. The main point of this figure is that, after just 10 generations (less than 230 CFD runs), a large number of acceptable designs were found, guided by the cumulatively enhanced metamodel. This rather small number of CFD runs (for a six-dimensional parameter space) is to be contrasted with the total number of metamodel interrogations used to drive the optimization, which is on the order of 3 million.

V. Mutual Enhancement of Data Sets

Global metamodels can be used to perform data fusion operations to enhance the usefulness of limited experimental data. The problem is akin to interpolating and extrapolating the data outside of the range where these data were collected, a task which, without any regularizing assumptions, constitutes a fundamentally ill-posed problem.¹⁷ Regularizing assumptions can come in various forms: physics based models, mathematical equations (such as splines), implicit smoothness assumptions, or other empiricisms. The method used here employs a particular form of regularization, in which a hypersurface going through the experimental data is “supported” by additional computational constraints. The calculation of this hypersurface uses the same technology as previously described in Section II and the applications of Sections III and IV.

A. Correction of Aerodynamic Databases Using Experimental Data

The fusion of wind-tunnel data with aerodynamic databases is considered next. The goal is to assimilate wind-tunnel data for use in flight simulations. The example shown below corresponds to a generic body-tail configuration (not shown). This application merges two data sets: an experimental (wind tunnel) data set, and a computational data set. These data are used to create a hypersurface metamodel. The data point inputs to the hypersurface calculation are referred to as the *support vectors* of the surface. The “computational” support vectors are supplied by the MISL3 code.¹⁸ This MISL3 database consists of forces and moments predictions for a wide range of angles of attack, roll angles, and Mach numbers in the subsonic, transonic, and supersonic range. The experimental support vectors are supplied by a wind tunnel test for a much smaller range of conditions consisting of three Mach numbers, four roll angles, and a subset of the angle-of-attack range. The goal was to produce an “error database,” defined as the difference between a fit to the MISL3 database and a fit to the experimental data. The fits were obtained by

constructing a low-order analytic (smoothly varying) four-dimensional response surface generated using simultaneously the MISL3 support vectors (in one “plane” of the parameter space), *and* the experimental support vectors (in a separate “plane” of the parameter space). The error database can then be used to “correct” the MISL3 database to take into account the experimental measurements. Conversely, this method is a way of constructing smart interpolation and extrapolation schemes in cases where limited quantities of data are available. From a user perspective, this process is automatic and does not require the specification of any equations. Only support vectors from alternate sources of data are needed. In this case (see Fig. 8), there is a very specific shape to the extrapolation in angle of attack and roll angle. This shape is “inherited” from the MISL3 database.

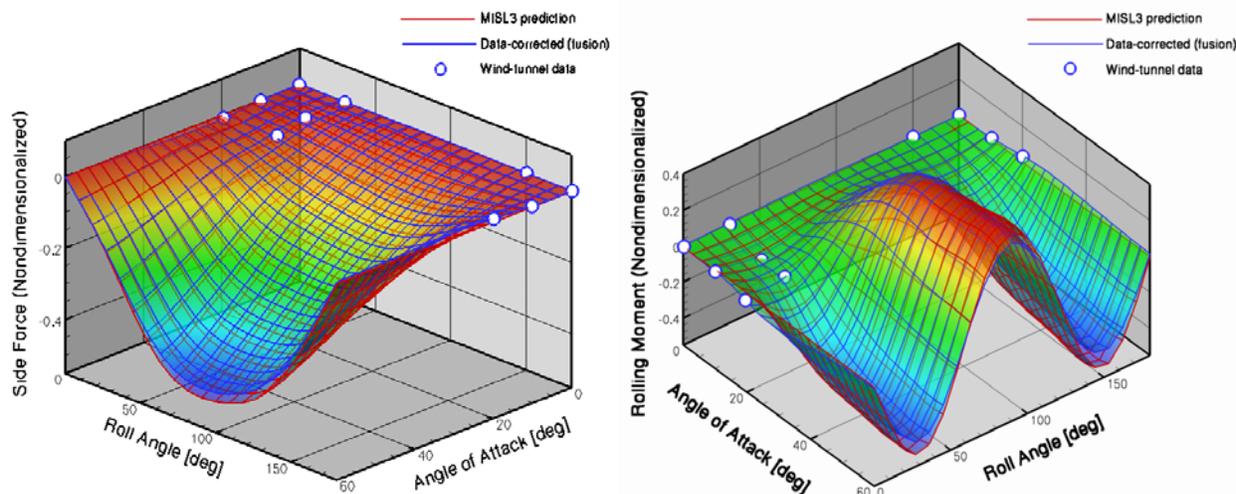


Figure 8. MISL3 and experimental side force (rolled coordinate system) and rolling moment predictions.

B. Alternate Uses

In a completely different application (not shown), NEAR has used this method to improve pressure drag predictions based on the fusion of high-resolution CFD calculations with limited pressure tap measurements in a wind-tunnel test.

It is argued that the technique outlined above is quite general and can be used, either to “fill-in” where limited quantities of experimental data are available, or to “fine tune” the results of computational analyses using the limited data available.

VI. Uncertainty Prediction

As previously mentioned, a key aspect of the technology is the ability to estimate further sampling needs and estimate the response surface quality, based on uncertainty estimation. The uncertainty estimation in NEAR-RS is based on propagating statistical descriptions of uncertainty in measurements or input data to estimates of uncertainty in the response surface coefficients themselves. This approach makes use of the covariance of the output measurements and is based on the theory of best linear unbiased estimation.

By way of illustration, an example of response surface with uncertainty estimation is given in Fig. 9, for a model of the aerodynamic performance of an X-38-type forebody configuration as a function of Mach number and pitch angle upon atmospheric reentry. The X-38 data set consists of Euler CFD computations carried out by NASA Ames, and was used as a test case in the development of CFDEExchange,¹² a system for the automated archival, search, and results retrieval of CFD solutions, using a relational database. In this system, the retrieved CFD solutions are considered as points in a multidimensional space. In the present case, the solution space is parameterized by Mach number, pitch angle, grid resolution, and solution algorithm settings.

The response surface shown in Fig. 9 is for given resolution and algorithm settings. The CFD points for these conditions are the only inputs to the response surface calculation, and are indicated in the figure in the form of triangular symbols. Although the data points are clustered only around the nominal trajectory parameters, the resulting surface is analytical and exists everywhere in space, including outside the intended X-38 reentry trajectory. The uncertainty prediction is indicated by the coloring of the response surface. For this calculation, a uniform uncertainty of 10^{-4} was assumed for each data point. Note the large amplification of the prediction uncertainty away from the data points.

It has been previously suggested¹² that, in cases where the goal is to characterize the overall response surface within an acceptable uncertainty bound, the uncertainty prediction can be used to guide the further population of the parameter space, for example, by “placing” further runs/cases in areas of highest uncertainty. In other applications, advanced optimization and design-of-experiments strategies that make efficient use of this predicted uncertainty are currently under development at NEAR.

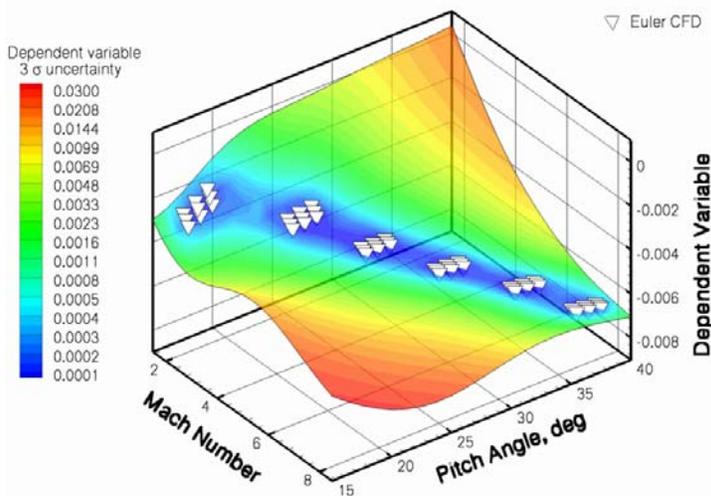


Figure 9. Uncertainty modeling for X-38 reentry.

VII. Conclusion

Cumulative global metamodels (1) allow data structure flexibility and the use of heterogeneous data sets, (2) provide a fully analytic, mathematical description that can be easily manipulated and shared between aerospace customers, and (3) have a rational basis for propagating uncertainty estimates, making these models suitable for risk assessment. Metamodel uncertainty can be used as a driver for decision making or in further populating data sets.

Acknowledgments

The authors thank the following sponsoring agencies for their financial support: (1) Air Force Research Laboratory (refueling drogue/canopy design, and launch vehicle payload fairing design), (2) Office of Naval Research, Program Officer Brian Almquist, Code 321OE (aerodynamic database enhancement), and (3) NASA Ames Research Center, Technical Monitor Terry Holst, (X-38/response surface uncertainty). The authors also thank Daniel A. Pruzan, Marnix F.E. Dillenius, Michael R. Mendenhall, and Laura C. Rodman for helping supply the application data that made possible the development of the methods presented in this paper.

References

- ¹Burman, J., Papila, N., Shyy, W., and Gebart, B. R., “Assessment of Response Surface-Based Optimization Techniques for Unsteady Flow Around Bluff Bodies,” AIAA 2002-5596.
- ²DeLoach, R. and Erickson, G. E., “Low-Order Response Surface Modeling of Wind Tunnel Data Over Truncated Inference Subspaces,” AIAA 2003-0456.
- ³Forrester, A. I. J., Bressloff, N. W. and Keane, A. J., “Response Surface Model Evolution,” AIAA 2003-14089.
- ⁴Hirokawa, N., Fujita, K., and Iwase, T., “Voronoi Diagram Based Blending of Quadratic Response Surfaces for Cumulative Global Optimization,” AIAA 2002-5460.
- ⁵Knill, D. L., Giunta, A., Baker, C. A., Grossman, B., Mason, W. H., Haftka, R. T., and Watson, L. T., “Response Surface Models Combining Linear and Euler Aerodynamics for Supersonic Transport Design,” *Journal of Aircraft*, Vol. 36, No. 1, Jan.-Feb. 1999, pp. 75-86.
- ⁶Krishnamurthy, T., “Response Surface Approximation With Augmented and Compactly Supported Radial Basis Functions,” AIAA 2003-1748.
- ⁷Papila, N., Shyy, W., Griffin, L., and Dorney, D. J., “Shape Optimization of Supersonic Turbines Using Response Surface and Neural Network Methods,” AIAA 2001-1065.

⁸Raeth, P. G., Gustafson, S. C., Little, G. R., and Puterbaugh, T. S., "Stretch and Hammer Neural Networks For N-Dimensional Data Generalization," Air Force Wright Laboratory Report WL-TR-91-1146, Jan. 1992.

⁹Poggio, T. and Girosi, F., "Network for Approximation and Learning," *Proc. IEEE*, Vol. 78, No. 9, 1990, pp. 1481-1497.

¹⁰Rais-Rohani, M., and Singh, M. N., "Efficient Response Surface Approach for Reliability Estimation of Composite Structures," AIAA 2002-5604.

¹¹Vittal, S. and Hajela, P., "Confidence Intervals for Reliability Estimated Using Response Surface Methods," AIAA 2002-5475.

¹²Rodman, L. C., Reisenhel, P. H., and Childs, R. E., "An Automated Documentation and Reporting System For CFD," AIAA 2002-0986.

¹³Strickland, J. H., "Axisymmetric Bluff-Body Flow: A Vortex Solver For Thin Shells," SAND91-2760, May 1992.

¹⁴Childs, R. E., Reisenhel, P. H., Rose, J., and Maly, J., "Large Asymmetric Launch Vehicle Payload Fairing," NEAR TR 611, Nielsen Engineering & Research, Apr. 2005.

¹⁵Buning, P.G., Jespersen, D.C., Pulliam, T.H., Klopfer, G.H., Chan, W.M., Slotnick, J.P., Krist, S.E., and Renze, K.J., "Overflow User's Manual Version 1.8r," NASA Ames Research Center, Oct. 2000.

¹⁶Giunta, A.A., Wojtkiewicz, S.F., and Eldred, M.S., "Overview of Modern Design of Experiments Methods for Computational Simulations," AIAA 2003-0649.

¹⁷Zeldin, B. A. and Meade, A. J., "Integrating Experimental Data and Mathematical Models in Simulation of Physical Systems," *AIAA Journal*, Vol. 35, No. 11, 1998, pp. 1787-1790.

¹⁸Lesieutre, D.J., Love, J. F., and Dillenius, M.F.E., "MISL3 Aerodynamic Analysis For Finned Vehicles With Axisymmetric Bodies," NEAR TR 561, Nielsen Engineering & Research, Mar. 2004.